

The product replacement algorithm is polynomial

Igor Pak*

Department of Mathematics
Yale University
paki@math.yale.edu

Abstract

The product replacement algorithm is a heuristic designed to generate random group elements. The idea is to run a random walk on generating k -tuples of the group, and then output a random component. The algorithm was designed by Leedham-Green and Soicher ([31]), and further investigated in [12]. It was found to have an outstanding performance, much better than the the previously known algorithms (see [12, 22, 26]). The algorithm is now included in two major group algebra packages GAP [42] and MAGMA [10].

In spite of the many serious attempts and partial results, (see [6, 14, 15, 21, 22, 32, 39, 40]), the analysis of the algorithm remains difficult at best. For small values of k even graph connectivity becomes a serious obstacle (see [19, 37, 39, 40]). The most general results are due to Diaconis and Saloff-Coste [22], who used a state of the art analytic technique to obtain polynomial bounds in special cases, and (sub)-exponential bounds in general case. The main result of this paper is a polynomial upper bound for the cost of the algorithm, provided k is large enough.

1. Introduction

The product replacement algorithm is defined as follows ([12]). Given a finite group G , let $\mathcal{N}_k(G)$ be a 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 as follows. Choose uniformly a pair (i, j) , such that $1 \leq i \neq j \leq k$, and apply one of the following four operations with equal probability:

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

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

*Supported by the NSF Postdoctoral Research Fellowship.

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* $\mathcal{W} = \mathcal{W}_k(G)$.

For a technical reason which will be clear later, it is useful to define a *lazy product replacement random walk*. Flip a fair coin at every walk step and if heads, do as above, and if tails stay put. This can slow down the walk by a factor of at most 2, but helps avoid parity problems with the graph being bipartite or nearly bipartite (see [13]).

About presentation of groups. We assume that 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 [6]). We will not use the latter operation. The group is then defined by a set of generators (g_1, \dots, g_r) . Now, in the algorithm one should take $k \geq r$ and set $g_{r+1} = \dots = g_k = \text{id}$ (see [12]).

A few words about the parameters k and t . In the original paper [12] the authors showed that when $k \geq 2 \log_2 |G|$ the product replacement graph $\Gamma_k(G)$ of moves on generating k -tuples, is (strongly) connected. It was further noted in [6] that when $k \geq 2 \log_2 |G|$ the diameter of $\Gamma_k(G)$ is $O(k \log |G|)$.

Let \mathbf{P} be a distribution on set X , and let \mathbf{U} be a uniform distribution. Define the *total variation distance* as follows:

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

It is easy to see that $0 \leq \|\mathbf{P} - \mathbf{U}\|_{\text{tv}} \leq 1$.

Main Theorem Let G be a finite group, and let $k = \Omega(\log |G| \log \log |G|)$. Denote by Q^t the probability distribution of the t -th step of the lazy product replacement random walk. Then

$$\|Q^t - U\|_{\text{TV}} < \varepsilon, \text{ given}$$

$t > C (\log^4 |G| \log^2 k + k \log k \log |G|)^2 k \log |G| \log(1/\varepsilon)$, where C is a universal constant.

In other words, when $k = \theta(\log |G| \log \log |G|)$, the mixing time of the walk is $O(\log^9 |G| (\log \log |G|)^5)$. This is a significant improvement over the general results in [22], which involve a nasty parameter $\Delta(G, k)$ defined as the largest diameter of the Cayley graph of G on k generators. On the other hand, the bounds we obtain are quite worse when compared to the bound in [32] in special cases of abelian and nilpotent groups. We elaborate further on previous results in the next section.

Before we conclude, let us emphasize however that the Main Theorem partly closes the gap between theoretical and practical results.

2. Previous results and Applications

An extensive review of the previous results and related subjects can be found in a review article [40]. Thus we shall restrict ourselves to a very brief sketch.

In [14, 15] the asymptotic for the mixing time when $k \rightarrow \infty$ is obtained. While somewhat better than ours (we get $O(k^3 \log^2 k)$ versus $O(k^2 \log k)$ known bound), the constants implied by $O(\cdot)$ notation in their case were roughly $|G|^{d(G)}$, and thus inferior to our poly-log constants.

In an analytically elaborate papers [21, 22], the authors obtained general bounds, which seem subexponential for certain abelian and nilpotent groups, and conjectured polynomial for simple groups. In [22] various specific examples were considered, but even for all abelian groups the authors do not obtain polynomial bounds. This was established by the author in [39], and then improved in [32] by an algebraic technique. In a different direction, it was shown in [32] that when k is fixed and $|G| \rightarrow \infty$, the mixing time of the walk is $O(\log |G|)$, when G is nilpotent of bounded class. This suggests that perhaps our polynomial but admittedly weak bounds might be improved in the future.

A few words about the applications. Starting with the first algorithms of Sims [44], a large number of group algorithms has been developed. Recently, various randomized algorithms has been introduced in a generality of matrix and black box groups (see e.g. [11, 29, 33]), which assume an access to (nearly) uniform group elements (see [6, 28]). This assumption is justified by a pioneer work of Babai [4], who gave the first polynomial time algorithm for generation of random group elements. His algorithms runs

in $O(\log^5 |G|)$ time, which is quite superior when compared to our $O^*(\log^9 |G|)$ bound.

In practice, however, only the product replacement algorithm is used as it exhibited a remarkable performance ([12, 26, 31]). It is natural to conjecture that the algorithm remains polynomial even for relatively small k . This was partly justified in [32] by reducing the problem to a long standing open algebraic problem. It is nevertheless clear that one can always add trivial elements to fill the rest of the k -tuple even if a given generating set is small¹. This shows that the algorithm has a rigorously polynomial modification.

3. Proof outline

In a nutshell, the main underlying idea is to “emulate” the analysis of Babai algorithm in the case of product replacement. While we do not wish to give a general method for such “emulation”, it is available indeed, and will be presented in the future publication. The idea is somewhat technical, so for the sake of brevity we do not spell it out in this case, but rather present an independent proof, which we outline below.

The proof is based on the use of a well known multi-commodity flow technique, basic results of which we recall below in section 4. Roughly, one wishes to send one unit of commodity between every pair of vertices in a graph so as to minimize congestion along the edges. The congestion achieved gives a bound on the mixing time (see e.g. [45]).

Now, the construction of the multicommodity flow is based on a modified version of the Babai’s algorithm. Speaking loosely, one considers the first vertex of the graph as an input in Babai’s algorithm, and lets the commodity move along the edges according to the algorithm operations. At the end, the distribution of commodities is somewhat nonuniform, so a standard fill up argument is used (cf. [1]).

Unfortunately a good portion of the analysis in [4] can’t be translated to our somewhat more general setting. The problem is that our starting vertex can be arbitrary and we have a version of reversibility problem, not unlike that in quantum computing (cf. [30]). An analytic approach (cf. [23]), combined with various probabilistic bounds, resolves the problem. As a byproduct, we obtain a better bound $O(\log^4 |G|)$ for the (modified) Babai algorithm. This will be included in the full version of the paper.

The rest of the paper is constructed as follows. In sections 4, 5 we present some background material on multi-commodity flows and random walks on groups. While most definitions are standard, some important notation are introduced and a few important technical results are recalled.

Section 6 is the key section where we introduce a multicommodity flow used in the proof main Theorem. In sec-

¹This is also done in practice (see [12, 31])

tions 7, 8 the congestion of the flow is computed based on a key technical Theorem 8.1, which is roughly analogous to the main result of Babai [4] (and stronger in effect). Finally, Theorem 8.1 is proved in sections 9, 10.

4. Random walks on graphs

By $\Gamma = (V, E)$ we denote oriented graphs, possibly with loops, on a set of vertices V with a set of edges E . For convenience, we write $e \in \Gamma$ in place of $e \in E$. A nearest neighbor random walk in graph Γ starting at $v \in V$, denotes $\mathcal{W} = \mathcal{W}_v(\Gamma)$, is defined as a walk which starts at v and at each step moves along a uniformly chosen edge leaving v .

We say that Γ is *symmetric* if the number of edges $(v, w) \in \Gamma$ is equal to the number of edges $(w, v) \in \Gamma$. Graph Γ is called *D-regular* if every vertex $w \in V$ has the same in-degree and out-degree $\deg(w)$. From now on we always assume that our graph Γ is symmetric and *D-regular*.

Let X be a finite set, and let \mathbf{P} be a probability distribution on X . For the rest of the paper \mathbf{U} will always denote a uniform distribution. There are several ways to measure a distance between \mathbf{Q} and \mathbf{U} (cf. [2]). First, define the *total variation distance* :

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

where $\mathbf{P}(B) = \sum_{x \in B} \mathbf{P}(x)$.

Similarly, define the *separation distance* :

$$s(\mathbf{P}) = |X| \cdot \max_{x \in X} \left(\frac{1}{|X|} - \mathbf{P}(x) \right).$$

One can think of $s(\mathbf{P})$ as of one-sided ℓ_∞ -distance. Note that $0 \leq \|\mathbf{P} - \mathbf{U}\|_{\text{TV}} \leq s(\mathbf{P}) \leq 1$.

Denote by Q_v^t the probability distribution of the t -th step of the walk $\mathcal{W}_v(\Gamma)$. If Γ is symmetric, connected, and is not bipartite, it is well known that $Q_v^t(w) \rightarrow 1/N$, where $N = |V|$. By $A = (a_{v,w})$ denote the transition matrix of the walk: $a_{v,w} = \#\{(v, w) \in \Gamma\} / \deg(v)$. By $1 = \lambda_0 > \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_{N-1} > -1$ denote the eigenvalues of A , where $N = |V|$. Let $\mu = \max\{|\lambda_1|, |\lambda_{N-1}|\}$ be the eigenvalue gap. It is easy to see that

$$\left| Q_v^t(w) - \frac{1}{N} \right| \leq \mu^t \cdot N, \quad \|Q_v^t - \mathbf{U}\|_{\text{TV}} \leq \mu^t \cdot N,$$

for all $v \in V$ (see e.g. [2]). When the walk is *lazy* (see introduction), we have $\lambda_{N-1} \geq 0$, $\mu = \lambda_1$, and

$$\left| Q_v^t(w) - \frac{1}{N} \right| \leq \lambda_1^t \cdot N, \quad \|Q_v^t - \mathbf{U}\|_{\text{TV}} \leq \lambda_1^t \cdot N,$$

for all $v \in V$.

Let $\Gamma = (V, E)$ be an oriented graph, $u, v \in V$. A *u - v flow* $f^{u,v}$ is a function on the edges of Γ , such that $f^{u,v}(e) \geq 0$ for all $e \in \Gamma$ and

$$\sum_j f^{u,v}(j, i) = \sum_j f^{u,v}(i, j) \quad \text{for all } i \in V, i \neq u, v.$$

The *value* of a flow is defined as

$$\text{val}(f) = \sum_j f^{u,v}(j, v) = \sum_j f^{u,v}(u, j).$$

A *multicommodity flow* is a collection $f = (f^{u,v})$ for all pairs of vertices $u \neq v$. We say that multicommodity flow f is *uniform* if

$$\sum_v \text{val}(f^{u,v}) = \sum_v \text{val}(f^{v,u}) = \frac{1}{N^2} \quad \text{for all } u \in V.$$

One can think of a flow $f^{u,v}$ as of a way to send $\text{val}(f^{u,v})$ units of (u, v) commodity from u to v through edges of Γ . Similarly, a uniform flow f is a way to send $1/N$ units of each of the N^2 commodities through edges of Γ .

One way to construct a uniform multicommodity flow is to use paths in Γ . Let $X^{u,v}$ be a set of simple paths between u and v . Take

$$f^{u,v}(e) = \frac{1}{N^2 \cdot |X^{u,v}|} \cdot |\{\gamma \in X^{u,v} \mid \gamma \ni e\}|$$

for any $e \in \Gamma$. Now, if a collection of paths $X = (X^{u,v})$ is given, this construction defines a multicommodity flow. We refer to [45] for the inverse procedure and further details. By $\ell(X)$ denote the maximum length of paths in X .

Let $e \in E$ be an edge in a *D-regular* graph Γ . Define *flow through e* as

$$f(e) = \sum_{u,v} f^{u,v}(e).$$

Define the *cost* of a flow as

$$\text{cost}(f) = \sum_{e \in \Gamma} f(e).$$

Observe now that for a uniform flow f which corresponds to set of path X the cost is given by

$$\begin{aligned} \text{cost}(f) &= \sum_{e \in \Gamma} f(e) \leq \sum_{u,v \in V} \ell(X) \cdot \text{val}(u, v) \\ &\leq N^2 \ell(X) \frac{1}{N^2} = \ell(X). \end{aligned}$$

If all paths in X have the same length, the above inequality becomes an equality.

Define the *congestion* $\rho(f)$ to be the scaled maximum flow through edges in Γ :

$$\rho(f) = N \cdot D \cdot \max_{e \in \Gamma} f(e).$$

A *conductance* $\Phi = \Phi(\Gamma)$ of a graph Γ is defined by

$$\Phi = \min_{X \subset V, 1 \leq |X| \leq N/2} \frac{|E(X, V \setminus X)|}{D \cdot |X|},$$

where $N = |V|$, $E(X, Y) = \{e = (v, w) \in \Gamma \mid v \in X, w \in Y\}$. A well known bound of Jerrum and Sinclair [27] (see also [2, 23, 45]) states that

$$\frac{\Phi^2}{2} \leq 1 - \lambda_1 \leq 2\Phi.$$

Let $f = (f^{u,v})$ be a uniform multicommodity flow with maximal congestion $\rho(f)$, in a D -regular graph Γ . Then

$$\Phi \geq \frac{1}{2\rho(f)}$$

Now, if f is a uniform multicommodity flow with small congestion, one can obtain bounds on the eigenvalue gap via bounds on conductance:

$$1 - \lambda \geq \frac{1}{8\rho^2(f)}$$

In general case, a shortcut was discovered by Diaconis and Stroock [23] (see also [2, 20, 45]).

Theorem 4.1 ([23]) Let f be a uniform multicommodity flow, given by a set of paths X . Then

$$1 - \lambda_1 \geq \frac{1}{\rho(f) \cdot \ell(X)}$$

5. Probability on groups

Let G be a finite group. By $d(G)$ denote the minimum number of generators of G . A generating set is called *redundant* if one generator can be omitted so that the remaining generators still generate G . By $\bar{d}(G)$ denote the maximum size of the nonredundant generating set. Note that

$$1 \leq d(G) \leq \bar{d}(G) \leq s(G) \leq \log_2 |G|,$$

where $s(G)$ is the length of the longest subgroup chain in G .

We denote generating k -tuples by $(g) = (g_1, \dots, g_k)$, where $\langle g_1, \dots, g_k \rangle = G$, $k \geq d(G)$. The set of generating k -tuples is denoted by $\mathcal{N}_k(G)$. Also, let $N_k(G) = |\mathcal{N}_k(G)|$ denote the number of generating k -tuples of G .

Theorem 5.1 ([37]) For any finite group G , $1 > \varepsilon > 0$, we have $N_k(G) > (1 - \varepsilon) \cdot |G|^k$, given $k > \log_2 |G| + 2 + \log_2 1/\varepsilon$.

A weaker version of the result in Theorem 5.1, with $k > 2 \log_2 |G| + \Omega(1)$ follows easily from $\mathbf{P}(g_i \notin \langle g_1, \dots, g_{i-1} \rangle) \geq 1/2$. We refer to [37] for references and generalizations.

By $\Gamma_k(G)$ denote the *product replacement graph* with edges defined by the moves $R_{i,j}^\pm$, $L_{i,j}^\pm$, $I_{i,j}^\pm$ and $J_{i,j}^\pm$, where $R_{i,j}^\pm$ and $L_{i,j}^\pm$ are as in the introduction, $1 \leq i \neq j < n$, and $I_{i,j}^\pm$, $J_{i,j}^\pm$ are defined as loops in $\Gamma_k(G)$. We will say that the edges corresponding to the same move have the same *label*.

Observe now that the lazy product replacement random walk, denotes $\mathcal{W}_k(G)$, can be defined now as the nearest neighbor random walk on $\Gamma_k(G)$.

Let G be a finite group, $(g) = (g_1, \dots, g_k) \in G^k$ be any k -tuple. Define *random subproducts* of (g) as follows:

$$h = g_1^{\varepsilon_1} \cdot \dots \cdot g_k^{\varepsilon_k},$$

where $\varepsilon_i \in \{0, 1\}$, $1 \leq i \leq k$, are chosen by independent flips of a fair coin (see [25]). Note that distribution of h may depend on the order of elements in a k -tuple.

Theorem 5.2 ([25]) Let (g) be a random k -tuple, chosen uniformly from G^k . By $\mathbf{P}_{(g)}$ denote be the probability distribution on G of the random subproducts of (g) . Then for all $\varepsilon, \delta > 0$ we have

$$\mathbf{P} \left(\max_{h \in G} \left| \mathbf{P}_{(g)}(h) - \frac{1}{|G|} \right| \leq \frac{\varepsilon}{|G|} \right) > 1 - \delta,$$

where the probability $\mathbf{P}(\cdot)$ is over all (g) , and $k \geq 2 \log_2 |G| + 2 \log_2(1/\varepsilon) + \log_2(1/\delta)$.

Theorem 5.3 ([8, 16, 17]) Let h_1, \dots, h_r be a collection of independently chosen random subproducts of generators. Then $\langle h_1, \dots, h_r \rangle = G$ with probability $> 1 - e^{-\varepsilon^2 r/8}$, given $r \geq 2 \log_2 |G|/(1 - \varepsilon)$, $1 > \varepsilon > 0$.

We will need to use a slightly modified version of the Theorem. Recall that the proof is based on the following well known Lemma.

Lemma 5.4 ([8, 16, 17]) Let $H < G$ be a proper subgroup of $G = \langle g_1, \dots, g_k \rangle$. Then a random subproduct $h = g_1^{\varepsilon_1} \cdot \dots \cdot g_k^{\varepsilon_k} \notin H$ with probability $\geq 1/2$.

Let $S = \{g_1, \dots, g_k\}$ be a generating set of G . Define a *symmetric random walk* on G of length t to be a random product

$$h = g_{i_1}^{\pm 1} \cdot \dots \cdot g_{i_t}^{\pm 1},$$

and a *lazy symmetric random walk* on G of length t to be a random product

$$h' = (g_{i_1}^{\pm 1})^{\varepsilon_1} \cdot \dots \cdot (g_{i_t}^{\pm 1})^{\varepsilon_t},$$

where $i_j \in \{1, \dots, k\}$, $\epsilon_j \in \{0, 1\}$, $1 \leq j \leq t$, and signs are chosen uniformly and independently. One can think of this random walk as of lazy nearest neighbor random walk on the Cayley graph $\Gamma = \text{Cayley}(G, R)$ of G with a generating set $R = S \cup S^{-1}$.

Denote by Q^t the probability distribution of h as above. From the above, it is easy to see that $Q^t(h) \rightarrow 1/|G|$ as $t \rightarrow \infty$, for all $h \in G$. A general bound on the rate of convergence can be obtained as follows.

Construct a uniform multicommodity flow f as follows. Let S be the generating set as above. For every elements $w \in G$ fix a shortest decomposition $w = s_1 \cdot s_2 \dots s_l$, where $s_i \in S$. Now send $1/|G|^2$ of a commodity $g - h$ from g to h via a path γ , which corresponds to the decomposition of $w = g^{-1}h$:

$$\gamma(g, h) = (g \rightarrow g s_1 \rightarrow g s_1 s_2 \rightarrow \dots \rightarrow g s_1 s_2 \dots s_l = h).$$

For the conductance, the Jerrum-Sinclair bound gives:

$$\Phi \geq \frac{1}{2\eta|R|},$$

where η is the maximal number of times generator s appears in the decomposition $\gamma(w)$, maximum taken over all $s \in S$, $w \in G$. Note that $\eta \leq \Delta$, where $\Delta = \Delta(G, R)$ is the diameter of $\Gamma = \text{Cayley}(G, R)$. Further, the Diaconis-Stroock bound (Theorem 4.1) implies the following result.

Theorem 5.5 ([23, 45]) *In notation above,*

$$1 - \lambda_1 \geq \frac{1}{\eta|R|\Delta}. \quad \square$$

Let $\Gamma = \text{Cayley}(G, R)$ be a Cayley graph of group G with a generating set $R = S \cup S^{-1}$, and $d = |R|$. For any $g \in G$, let $\ell(g)$ denote a distance between id and g in Γ . Define a ball of radius r , $\Gamma^r = \{g \in G \mid \ell(g) \leq r\}$. As before, by X_t denote the t -th step of the random walk on Γ starting at id .

The following result is a local version of Diaconis-Stroock theorem.

Theorem 5.6 (Babai) *In notation above, let $|\Gamma^{4r}| \leq |\Gamma|/2$, and let t be chosen uniformly from $\{k+1, \dots, l\}$, where*

$$l \geq 712 r^2 d \ln |\Gamma^{4r}|.$$

Then $\mathbf{P}(X_t \notin \Gamma^r) \geq 1/16$.

We say that a set of generators $S = \{s_1, \dots, s_n\}$ of group G is c -covering if each element of g can be written as a product of s_i , where each generator is used at most c times.

Clearly, the diameter Δ of the corresponding Cayley graph is at most cn .

We say that S is c -covering a subset $B \subset G$ if for every $g \in B$ there is a shortest decomposition $g = s_{i_1} \dots s_{i_l}$ such that the path $\gamma(\text{id}, g) \in \text{Cayley}(G, S)$, and each generator g_j is used at most c times. Note that the c -covering generating set S is the same as a set S which is c -covering the whole group G .

Theorem 5.7 *In notation above, let S be c -covering $B \subset G$, and let $|B^2| \leq |\Gamma|/2$. Let t be chosen uniformly from $\{1, \dots, l\}$, where*

$$l \geq 120 c d^2 \ln |B^2|.$$

Then $\mathbf{P}(X_t \notin B) \geq 1/24$.

The proof is a combination of the proof technique of Diaconis-Stroock [23] and Babai approach in [4]. The complete proof of Theorem 5.7 will be given elsewhere.

6. Construction of the multicommodity flow

Let $\Gamma = \Gamma_k(G)$ be the product replacement graph on generating k -tuples. We will construct a multicommodity flow in Γ with special properties.

The construction will depend on the integers r, t, m, T, M . These will be chosen appropriately later on. By N everywhere in this section we denote $N = N_k(G) = |\mathcal{N}_k(G)|$.

Let (g) be a generating k -tuple. We will use $[p, q]$ to denote a subset $\{g_p, \dots, g_q\}$. We say “multiply (g) in place j by a random walk of length T on generators $[p, q]$ ”, to denote the following procedure. We send a $1/(q-p+1)$ fraction of commodity from every $(g) = (g_1, \dots, g_k) \in \mathcal{N}_k(G)$, where it is currently contained to all k -tuples of the form $(g_1, \dots, g_j \cdot g_i^{\pm \epsilon}, \dots, g_k)$, where $i \in \{p, \dots, q\}$, $\epsilon \in \{0, 1\}$, $j \notin \{p, \dots, q\}$. Repeat this for T steps. At the end, $1/(q-p+1)^T$ fraction of the commodity will be contained in $(g_1, \dots, g_j \cdot h, \dots, g_k)$, where h is given by the lazy symmetric random walk on G with $[p, q]$ as a generating set.

Similarly, we say “multiply (g) in place j by a random subproduct on generators $[p, q]$ ”, to denote the procedure when a $1/(q-p+1)^{q-p+1}$ fraction of commodity is sent along the path $(g_1, \dots, g_k) \rightarrow (g_1, \dots, g_j \cdot g_p^{\epsilon_p}, \dots, g_k) \rightarrow (g_1, \dots, g_j \cdot g_p^{\epsilon_p} \cdot g_{p+1}^{\epsilon_{p+1}}, \dots, g_k) \rightarrow \dots \rightarrow (g_1, \dots, g_j \cdot h, \dots, g_k)$, where $\epsilon_j \in \{0, 1\}$, and h is given by the random subproducts with $[p, q]$ as a generating set.

Start with a generating k -tuple $(g) \in \mathcal{N}_k(G)$. Assume that $\langle g_1, \dots, g_r \rangle = G$. Use the following steps to define a multicommodity flow.

0) Set $1/N$ units of commodity in (g) .

1) For every $j = 1, \dots, r$, in this order, with probability $2/3$, multiply (g) in place j by a random subproduct on generators $[1, j-1] \cup [j+1, k]$.

2) Multiply (g) in place $r+1$ by a random walk on generators $[1, r]$, and of length l_1 which is chosen at random in $\{1, \dots, T\}$. Then multiply (g) in place $r+2$ by a random walk on generators $[1, r+1]$, and of length l_2 , which is chosen at random in $\{1, \dots, T\}$. Proceed in this manner until multiplication of (g) in place $r+t$ by a random walk on generators $[1, r+t-1]$, and of length l_t , which is chosen at random in $\{1, \dots, T\}$.

3) Multiply (g) in place j by a random walk of length M on generators $[1, r+t]$, for every $j = r+t+1, \dots, r+t+m$, in this order.

4) Multiply (g) in place j by a random subproduct on generators $[r+t+1, r+t+m]$, for every $j = 1, \dots, r+t, r+t+m+1, \dots, k$, in this order.

Denote by $F = (F^{(g),(h)})$ the flow defined in 1) – 4). We think of F as of $1/N$ -commodity flow which sends $1/N$ unit of commodity (g) to various generating k -tuples (h) , so each $(h) \in \mathcal{N}_k(G)$ gets some (possibly zero) fraction of the commodity.

Now consider a natural action of the symmetric group S_k by permutation of generators in k -tuples. This induces an action of S_k on edges in $\Gamma_k(G)$:

$$\sigma : R_{i,j}^{\pm} \rightarrow R_{\sigma(i),\sigma(j)}^{\pm}, L_{i,j}^{\pm} \rightarrow L_{\sigma(i),\sigma(j)}^{\pm},$$

$$\sigma : I_{i,j}^{\pm} \rightarrow I_{\sigma(i),\sigma(j)}^{\pm}, J_{i,j}^{\pm} \rightarrow J_{\sigma(i),\sigma(j)}^{\pm},$$

where $\sigma \in S_k$. This action can be extended to all paths and flows in $\Gamma_k(G)$.

Define a multicommodity flow $X = (X^{(g),(h)})$ as a flow which sends $1/N$ units of commodity from every (g) to some k -tuples (h) according to the flow

$$X^{(g),(h)} = \frac{1}{|S_k|} \cdot \bigcup_{\omega \in S_k} \omega^{-1} \cdot F^{\omega(g),\omega(h)},$$

One can think of performing X as of randomly choosing the order on elements in $\{1, \dots, k\}$ before doing 1) – 4).

Consider an involution π which acts on edges in $\Gamma_k(G)$ by reversing the order of multiplication:

$$\pi : R_{i,j}^{\pm} \leftrightarrow L_{i,j}^{\pm}, I_{i,j}^{\pm} \leftrightarrow J_{i,j}^{\pm}.$$

This action can be extended to paths and flows in Γ . Define a flow

$$Y = \frac{1}{2} (X + \pi \cdot X).$$

This corresponds to flipping a fair coin in advance and then performing 1) – 4), with the order of multiplication depending on the outcome.

Finally, define an involution ι which acts on edges in Γ by inverting their orientation. When $\Gamma = \Gamma_k(G)$ this action is defined as follows:

$$\iota : R_{i,j}^{\pm} \leftrightarrow R_{i,j}^{\mp}, L_{i,j}^{\pm} \leftrightarrow L_{i,j}^{\mp}, I_{i,j}^{\pm} \leftrightarrow I_{i,j}^{\mp}, J_{i,j}^{\pm} \leftrightarrow J_{i,j}^{\mp}.$$

Clearly, this action can be also extend to paths and flows in Γ . Now define a flow $Z = (Z^{(g),(h)})$ as follows:

$$Z = \frac{1}{2} (Y + \iota \cdot Y).$$

We shall think of Z as of multicommodity flow which sends $\text{val}(Z^{(g),(h)})$ units of commodity $(g) - (h)$ from (g) to (h) .

This completes the construction of the flow $Z = Z(r, t, m, T, M)$. In the following sections we shall prove that Z is nearly uniform (see below) and has a small congestion. This in turn gives bounds on conductance and the eigenvalue gap for the product replacement graph $\Gamma_k(G)$.

7. Congestion

This is normally the hardest quantity to calculate. It is somewhat easier in our case.

Observe that the flow Z corresponds to a certain large set of paths \mathbf{Z} in $\Gamma_k(G)$. By construction, all path in \mathbf{Z} have length at most

$$L = \ell(\mathbf{Z}) = r \cdot (k-1) + t \cdot T + m \cdot M + (k-m) \cdot m,$$

where the summands corresponds to the steps 1) – 4) of the construction.

Lemma 7.1 $\rho(Z) \leq L$.

Proof. Consider the labels of edges in a path $\gamma \in \mathbf{Z}$ (see section 2). Clearly, each path $\gamma \in \mathbf{Z}$ is uniquely defined by the labels and the starting point. Similarly, if an i -th point of the path is given along with all the labels, one can reconstruct the whole path as well.

Now, let $e = ((g), (g'))$ be an edge in $\Gamma_k(G)$. Suppose we are also given a sequence of labels of edges of the path, and we know that edge e is the i -th edge of the path. Then the starting point (h) of the path is uniquely determined, and is in one-to-one correspondence with (g) . Therefore in $\Gamma_k(G)$ the flow Z through any edge with the same label as e has must be the same.

Now, by the symmetry in the definition of \mathbf{Z} , this implies that the flow through all edges of $\Gamma_k(G)$ is the same. Indeed, the flow through $R_{i,j}^{\pm}, L_{i,j}^{\pm}, I_{i,j}^{\pm}, J_{i,j}^{\pm}$ is the same for all $1 \leq i, j \leq k$ due to the averaging over all permutations $\omega \in S_k$. The flow through $R_{i,j}^{\pm}, L_{i,j}^{\pm}$ is the same as through

$I_{i,j}^\pm, J_{i,j}^\pm$ by the symmetry in the definition of the subproducts and lazy random walks. The flow through $R_{i,j}^\pm, I_{i,j}^\pm$, is the same as through $L_{i,j}^\pm, J_{i,j}^\pm$ due to the symmetry in the change of the order of multiplication, given by involution π . Finally, the symmetry between $R_{i,j}^\pm$ and $R_{i,j}^\mp$, etc., comes from the involution ι .

Denote by $D = \deg(\Gamma_k(G))$ the total number of labels. Then for congestion we have:

$$\begin{aligned} \rho(Z) &= ND \max_{e \in \Gamma_k(G)} Z(e) = ND \frac{\sum_{e \in \Gamma_k(G)} Z(e)}{|\Gamma_k(G)|} \\ &= \text{cost}(Z) \leq L, \end{aligned}$$

where the last inequality follows from

$$\text{cost}(Z) \leq \sum_{(g),(h) \in \mathcal{N}_k(G)} \text{val}(Z^{(g),(h)}) \leq \ell(\mathbf{Z}) \leq L.$$

This completes the proof. \square

8. Uniformity

Recall that by construction flow Z is not uniform. The proof of the Main Theorem relies on the fact that Z is nearly uniform in the following precise sense.

Let $f = f(f^{u,v})$ be a multicommodity flow in graph $\Gamma = (V, E)$. We say that a flow f is ε -uniform if for all $u, v \in V$ we have

$$\frac{1-\varepsilon}{N} < \text{val}(f^{u,v}) < \frac{1+\varepsilon}{N},$$

where $N = |V|$.

Theorem 8.1 Let $k = \Omega(\log |G| \log \log |G|)$, $\varepsilon = 1/4$. Then Z is ε -uniform for certain parameters r, t, m, T, M , such that $L = O(\log^4 |G| + k \log k \log |G|)$, where $L = \ell(\mathbf{Z})$.

The proof of Theorem 8.1 is based on several intermediate results, each of them dealing with either random subproducts or with random walks on finite groups. The parameters $r \dots M$ will be explicit in the proof (see below).

Let us deduce the main result of the paper from Theorems 4.1, 8.1. We the following elementary observation.

Proposition 8.2 Let f be a ε -uniform multicommodity flow in graph $\Gamma = (V, E)$ given by a set of paths \mathbf{X} . Then

$$1 - \lambda_1 \geq \frac{1 - \varepsilon}{\rho(f)\ell(\mathbf{X})}.$$

Proof. Consider a flow \hat{f} obtained by decreasing flow along paths, so that $\text{val}(\hat{f}^{u,v}) = (1-\varepsilon)/N$, where $N = |V|$. Then $1/(1-\varepsilon) \cdot \hat{f}$ is a uniform flow with congestion

$$\rho\left(\frac{1}{1-\varepsilon} \cdot \hat{f}\right) = \frac{1}{1-\varepsilon} \rho(\hat{f}) \leq \frac{1}{1-\varepsilon} \rho(f).$$

Now apply Theorem 4.1 to get the result. \square

Proof of Main Theorem. From Proposition 8.1 and Lemma 7.1, we conclude that $1 - \lambda_1 \geq (1-\varepsilon)/L^2$, where L is as in section 7. Observe now that the Theorem follows immediately from here and Theorem 8.1, given the choice of parameters as in the proof of Theorem 8.1. We omit straightforward calculation. \square

9. Proof of Theorem 8.1

From now on we will think of uniform flows in graphs in terms of probability distributions of getting from a given vertex to other vertices. The proof of Theorem 8.1 follows from the following lemmas.

Lemma 9.1 Let $(g) = (g_1, \dots, g_k)$ be any generating k -tuple. Consider r random subproducts defined as

$$h_i = g_1^{\varepsilon_{i,1}} \cdot \dots \cdot g_{i-1}^{\varepsilon_{i,i-1}} \cdot g_{i+1}^{\varepsilon_{i,i+1}} \cdot \dots \cdot g_k^{\varepsilon_{i,k}},$$

where $1 \leq i \leq r$, and $\varepsilon_{i,j} \in \{0, 1\}$ are determined by independent coin flips. Let $g'_i = g_i \cdot (h_i)^{\nu_i}$, where $\nu_i \in \{0, 1\}$, $\mathbf{P}(\nu_i = 1) = \frac{2}{3}$, are determined by independent Bernoulli trials. Then with probability $> 1-\varepsilon$ we have $\langle g'_1, \dots, g'_r \rangle = G$, given $r > 6 \log |G| \log(1/\varepsilon)$.

Lemma 9.2 Let $k \geq r + t$, and let $(g) = (g_1, \dots, g_k)$ be a generating k -tuple such that $\langle g_1, \dots, g_r \rangle = G$. As in step 2), consider a succession of random walks of length T starting at g_j , $r+1 \leq j \leq r+t$, on a generating set $[1, j-1]$. Then, with probability $> 1-\varepsilon$, the obtained generating set $[1, r+t]$ is 4-covering, given $t > 48/(1-\varepsilon) \log |G|$, $T > 480(r+t)^2 \log |G|$.

Lemma 9.3 Let $S = \{g_1, \dots, g_n\}$ be a 4-covering generating set. Denote by h , the results of a lazy symmetric random walks of length M , starting at any $g'_1, \dots, g'_m \in G$. Denote by Q the distribution of h . Then

$$Q(h) \geq \frac{1-\varepsilon}{|G|}, \text{ for all } h \in G,$$

given $M > 16n^2(\log |G| + \ln(1/\varepsilon))$.

Lemma 9.4 Let (h_1, \dots, h_m) be m elements independently chosen from distribution Q on G , such that $Q(h) \geq (1 - \varepsilon)/|G|$ for all $h \in G$. By \mathbf{P} denote the probability distribution of the random subproducts in h_i . Then

$$\mathbf{P}(h) \geq \frac{1 - \delta}{|G|}, \text{ for all } h \in G,$$

given $m > 2 \log |G| + 3 \log(1/\delta) + \log(1/\varepsilon)$.

Proof of Theorem 8.1 Denote by $(g^{(i)}) = (g_1^{(i)}, \dots, g_k^{(i)})$ a generating k -tuple obtained after application of steps 1) – i , $i = 1..4$. Then, by Lemma 9.1, with probability $\beta_1 > 1 - \varepsilon_1$ we have $\langle g_1^{(1)}, \dots, g_r^{(1)} \rangle = G$. By lemma 9.2, after step 2), with probability $\beta_2 > \beta_1 \cdot (1 - \varepsilon_2)$, we have the set $S = \{g_1^{(2)}, \dots, g_{r+t}^{(2)}\}$ is 4-covering. Further, by Lemma 9.3, the probability distribution \mathbf{Q} of g_j , $r + t < j \leq r + t + m$ satisfies $s(\mathbf{Q}) < 1 - \beta_3$, where $\beta_3 = \beta_2(1 - \varepsilon_3)$. Finally, by Lemma 9.4, we have $s(\mathbf{P}) < 1 - \beta_4$.

Now take $\varepsilon = 1/4$, $\beta_4 = \varepsilon/k$. It suffices to have $\varepsilon_i = \varepsilon/6k$. Take $r = C_1 \log |G| \log k$, $t = C_2 \log |G|$, $T = C_3 \log^3 |G| \log^2 k$, $m = C_4 (\log |G| + \log k)$, $M = C_5 \log^2 |G| (\log |G| + \log k)$, where C_i are universal constants chosen to satisfy the lemmas. Then Z is indeed $(1/4)$ -uniform and L is as desired. \square

10. Proof of Lemmas

It is instructive to start with a quick proof of Theorem 5.3 and Lemma 5.4, as the proof of Lemma 9.1 is completely analogous. In our presentation we follow [16].

Proof of Lemma 5.4 Consider the smallest i such that $g_i \notin H$. Let $u = g_1^{\varepsilon_1} \cdot \dots \cdot g_{i-1}^{\varepsilon_{i-1}}$, $w = g_{i+1}^{\varepsilon_{i+1}} \cdot \dots \cdot g_k^{\varepsilon_k}$. Now $h = u h_i^{\varepsilon_i} w$, and $u \in H$. When $w \in H$, with probability $1/2$ we have $\varepsilon_i = 1$, and $h \notin H$. When $w \notin H$, with probability $1/2$ we have $\varepsilon_i = 0$, and $h \notin H$. Thus in either case $h \notin H$ with probability $\geq 1/2$. \square

Proof of Theorem 5.3 Consider a sequence of subgroups $H_0 = \{\text{id}\}$, $H_{i+1} = \langle H_i, h_{i+1} \rangle$. By the lemma, for every i we have $\mathbf{P}(h_{i+1} \notin H_i) \geq 1/2$, provided $H_i \neq G$. Thus we need to estimate the probability that for r flips of a fair coin there are $\leq (1 - \varepsilon)r/2$ heads. Now apply Chernoff bound. We omit the easy details. \square

Proof of Lemma 9.1 Let $H \supsetneq G$ be a proper subgroup of $G = \langle g_1, \dots, g_k \rangle$. Consider a random subproduct $g'_i = g_i \cdot (h_i)^{\nu_i}$, where h_i is as in the Lemma. By analogy with the proof of Lemma 5.4, there are two cases. If $g_i \notin H$,

then with probability $1/3$ we have $\nu_i = 0$, and $g'_i = g_i \notin H$. Assume $g_i \in H$. The random subproduct $h_i \notin H$ with probability $1/2$, and with probability $2/3$ we have $\nu_i = 1$. Therefore with probability $\geq 1/2 \cdot 2/3 = 1/3$, we have $g'_i = g_i \cdot h_i \notin H$ in this case.

Now, we showed that from every $H \supsetneq G$ we have $g'_i \notin H$ with probability $\geq 1/3$. The rest of the proof follows verbatim the proof of Theorem 5.3. \square

Proof of Lemma 9.2 We follow the proof of Babai [4], with several changes which will be indicated below.

Let $B_j = \{h = g_1^{\varepsilon_1} \cdot \dots \cdot g_j^{\varepsilon_j}, \varepsilon_i \in \{0, 1\}\}$, $C_j = B_j^{-1} B_j$. Now if $g_{j+1} \notin C_j$, then $h_1 \cdot g_{j+1} \neq h_2$ for any $h_1, h_2 \in B_j$, and $|B_{j+1}| = 2|B_j|$. Following Babai-Szemerédi [9] (see also [4]), we call this *cube-doubling*.

We claim that after a random walk $j + 1$ -th place, $g_{j+1} \notin C_j$ with probability $> 1/24$. Here we cannot explicitly use Babai Theorem 5.6 since the starting point of the walk can be any group element.

Assume that $\mathbf{P}(\mathbf{X}_t \notin B_j) > \varepsilon$, where t is random in $\{1, \dots, T\}$, and \mathbf{X}_t is a random walk on Γ starting at id . We claim that $\mathbf{P}(a\mathbf{X}_t \notin B_j) > \varepsilon/2$, for any $a \in G$, and t random in $\{1, \dots, 2T\}$.

Indeed, if $a \in B$, then with probability $> \varepsilon$ we have $a\mathbf{X}_t \notin a \cdot C_j \supset B_j$, and the claim follows. On the other hand, let $a \notin B_j$. Consider the smallest i such that $a\mathbf{X}_i \in B_j$. If $i \leq T$, from the previous observation, with probability $> \varepsilon$ we have $a\mathbf{X}_t \notin B_j$, where t is random in $\{i+1, \dots, i+T\}$. Otherwise $a\mathbf{X}_t \notin B_j$, where t is random in $\{1, \dots, T\}$. We conclude that in both cases with probability $> \varepsilon/2$ we have $a\mathbf{X}_t \notin B_j$, where t is random in $\{1, \dots, i+T\}$. This completes proof of the claim.

From Theorem 5.6 and the above observation, we conclude that given $T \leq Cn^2 \log |G|$, we have $g_j \notin C_j$ with probability $> 1/48$, where $n = (r + t)$ and C is a universal constant. Proceed as in Babai [4] to finish the proof.

Proof of Lemma 9.3 This follows easily from Theorem 5.5. Indeed, consider a Cayley graph $\Gamma = \text{Cayley}(G, R)$, where $R = S \cup S^{-1}$. Then $|R| \leq 2|S| = 2n$. Recall that S is 4-covering. For every $h \in G$ consider a path between h and id in Γ , defined as in section 2, obtained from a decomposition of h implied by 4-covering. Then the diameter $\Delta(\Gamma) \leq 4n$. Also, in notation of section 2, $\eta \leq 4$. Therefore by Theorem 5.5

$$1 - \lambda_1 \leq \frac{1}{\eta |R| \Delta} = \frac{1}{4 \cdot 2n \cdot 4n} = \frac{1}{32n^2}.$$

Finally,

$$\frac{1}{|G|} - Q(h) \leq \lambda_1^M \cdot |G| \leq \left(1 - \frac{1}{32n^2}\right)^M |G| \leq \frac{\varepsilon}{|G|},$$

for all $h \in G$, given $M > 32n^2(2 \log |G| + \log(1/\varepsilon))$. This implies the result. \square

Proof of Lemma 9.4 The proof follows from the following general observation (see[36]). Let Q be a distribution of random subproducts

$$h = g_1^{\varepsilon_1} \cdots g_k^{\varepsilon_k}$$

where g_i is chosen uniformly and independently in G . Fix any $x \in G$, and consider a distribution Q' of random subproducts

$$h' = g_1^{\varepsilon_1} \cdots g_i^{\varepsilon_i} \cdot x \cdot g_{i+1}^{\varepsilon_{i+1}} \cdots g_k^{\varepsilon_k}$$

Then $Q'(h) = Q(hx^{-1})$. Indeed, we have

$$\begin{aligned} h'x^{-1} &= g_1^{\varepsilon_1} \cdots g_i^{\varepsilon_i} \cdot (xg_{i+1}^{\varepsilon_{i+1}}x^{-1}) \cdots (xg_k^{\varepsilon_k}x^{-1}) \\ &= g_1^{\varepsilon_1} \cdots g_i^{\varepsilon_i} \cdot (xg_{i+1}x^{-1})^{\varepsilon_{i+1}} \cdots (xg_kx^{-1})^{\varepsilon_k}, \end{aligned}$$

which implies the claim. In separation distance notation, this implies that $s(Q') = s(Q)$.

Now let Q be a distribution on G such that $s(Q) \leq \varepsilon$. This implies that $Q = (1 - \varepsilon) \cdot U + \varepsilon \cdot P$ for some distribution Q_1 . Every subproduct h of elements chosen from Q can be written as a product of g_i, x_j , where g_i are chosen uniformly and independently, and x_j are chosen independently from Q_1 . Take the length of subproducts to be $m > k(1 + \beta)/(1 - \varepsilon)$. Then, by Chernoff bound, there are at least k elements g_i in h with probability $p > 1 - e^{-\beta^2(1 - \varepsilon)m/2}$.

Now assume that h indeed contains $\geq k$ elements g_i . We can use the “pulling through” method as above for all elements x_j . Denote by R the distribution of random subproducts of length k , sampled from uniform distribution, and by P the distribution of subproducts of length m sampled from Q , $s(Q) \leq \varepsilon$. We obtain $s(P) \geq p \cdot s(R)$. Now apply Theorem 5.2 to get the result. \square

Acknowledgments

We would like to thank László Babai, Gene Cooperman, Persi Diaconis, Charles Leedham-Green and Eamonn O’Brien for helpful conversations. We also thank the anonymous referees for pointing out errors in the preliminary version of this paper.

References

- [1] D. J. Aldous, Some inequalities for reversible Markov chains, *J. London Math. Soc.* **25**, 1982, 564–576.
- [2] D. Aldous, J. Fill, *Reversible Markov Chains and Random Walks on Graphs*, monograph in preparation, 1996

- [3] N. Alon, J.H. Spencer, P. Erdős, *The Probabilistic Method*, Wiley, New York, 1992
- [4] L. Babai, Local expansion of vertex-transitive graphs and random generation in finite groups, Proc. STOC’91, 164–174
- [5] L. Babai, Automorphism groups, isomorphism, reconstruction, in *Handbook of Combinatorics*, Elsevier, 1996
- [6] L. Babai, Randomization in group algorithms: Conceptual questions, in *Groups and Computation II*, DIMACS Series, vol 28, AMS, Providence, 1997
- [7] L. Babai, I. Pak, Strong bias of group generators: an obstacle to the “product replacement algorithm”, Proc. SODA’00, 627–635
- [8] L. Babai, G. Cooperman, L. Finkelstein, E. Luks and A. Seress, Fast Monte Carlo algorithms for permutation groups, Proc. STOC’91
- [9] E. Babai, E. Szemerédi, On complexity of matrix groups problems I, Proc. FOCS’84, 229–240
- [10] 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
- [11] 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.*, **29**, 2000, 33–57
- [12] F. Celler, C.R. Leedham-Green, S. Murray, A. Niemeyer, E.A. O’Brien, Generating random elements of a finite group, *Comm. Algebra*, **23**, 1995, 4931–4948
- [13] F.R.K. Chung, *Spectral Graph Theory* (CBMS Regional Conference Series in Mathematics, No. 92), AMS, Providence, RI, 1994
- [14] F.R.K. Chung, R.L. Graham, Random walks on generating sets for finite groups, *The Electronic J. of Comb.*, **4**, 1997, #R7
- [15] F.R.K. Chung, R.L. Graham, Stratified random walk on the n -cube, *Random Struct. Algor.*, **1**, 1997, 199–222
- [16] G. Cooperman, L. Finkelstein, Random algorithms for permutation groups, *CWI Quarterly*, **5**, 1992, no. 2, 107–125

- [17] G. Cooperman, L. Finkelstein, Combinatorial tools for computational group theory, in *Groups and Computation II*, DIMACS Series, vol 28, AMS, Providence, 1997
- [18] P. Diaconis, *Group Representations in Probability and Statistics*, IMS, Hayward, California, 1988
- [19] P. Diaconis, R. Graham, The graph of generating sets of an abelian group, *Colloq. Math.*, **80**, 1999, 31–38
- [20] P. Diaconis, L. Saloff-Coste, Comparison techniques for random walk on finite groups, *Ann. Prob.*, **21**, 1993, 2131–2156
- [21] P. Diaconis, L. Saloff-Coste, Walks on generating sets of abelian groups, *Prob. Th. Rel. Fields*, **105**, 1996, 393–421
- [22] P. Diaconis, L. Saloff-Coste, Walks on generating sets of groups, *Invent. Math.*, **134**, 1998, 251–199
- [23] P. Diaconis, D. Stroock, Geometric bounds for eigenvalues of Markov chains, *Ann. Appl. Prob.*, **1**, 1991, 36–61
- [24] M.J. Dunwoody, Nielsen Transformations, in *Computational Problems in Abstract Algebra*, 45–46, Pergamon, Oxford, 1970
- [25] P. Erdős, A. Rényi, Probabilistic methods in group theory, *Jour. Analyse Mathématique*, **14**, 1965, 127–138
- [26] 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
- [27] M. Jerrum and A. Sinclair, Conductance and the rapid mixing property for Markov chains: the approximation of the permanent resolved, *Proc. STOC'88*, 235–243.
- [28] W.M. Kantor, Simple groups in computational group theory, *Proc. ICM Berlin*, Vol. II, 1998, 77–86
- [29] W. Kantor, A. Seress, Black box classical groups, *Memoirs AMS*, to appear
- [30] A. Yu. Kitaev, A.Kh. Shen, M.N. Vyalyı̄, *Classical and Quantum Algorithms* (in Russian), Indep. University, Moscow, 1999.
- [31] C.R. Leedham-Green, personal communication
- [32] A. Lubotzky, I. Pak, The product replacement algorithm and Kazhdan's property (T), to appear in *Journal of AMS*, 2000
- [33] P. Neumann, C. Praeger, A recognition algorithms for special linear groups, *Proc. London Math. Soc.*, **65**, 1992, 555 – 603
- [34] I. Pak, *Random walks on groups: strong uniform time approach*, Ph.D. Thesis, Harvard U, 1997
- [35] I. Pak, When and how n choose k , in *AMS DIMACS Series*, vol. 43, 191–238, AMS, Providence, 1998
- [36] I. Pak, Random walks on finite groups with few random generators, *Electr. J. Prob.*, **4**, 1999, 1–11
- [37] I. Pak, On probability of generating a finite group, preprint, 1999
- [38] I. Pak, On the graph of generating sets of a simple group, preprint, 1999
- [39] I. Pak, Generating random elements in solvable groups, in preparation, 1999
- [40] I. Pak, What do we know about the product replacement algorithm, to appear in *Groups and Computation III*, DeGruyter, 2000
- [41] I. Pak, S. Bratus, On sampling generating sets of finite groups and the product replacement algorithm, *Proc. ISSAC'99*, 91–96
- [42] M. Schönert et al., *GAP – Groups, Algorithms, and Programming*, Lehrstuhl D für Mathematik, RWTH Aachen, Germany, 1995
- [43] A. Shalev, Probabilistic group theory, *St. Andrews Lectures*, Bath, 1997
- [44] C. C. Sims, Group-theoretic algorithms, a survey, *Proc. ICM Helsinki*, 1978, 979–985
- [45] A. Sinclair, Improved bounds for mixing rates of Markov chains and multicommodity flow, *Combin. Probab. Comput.*, **1**, 1992, 351–370