Random Self-Reducibility Properties of Learning Problems over Burnside Groups of Exponent 3

Nelly Fazio* Kevin Iga† Antonio Nicolosi‡ Ludovic Perret§
William E. Skeith III*

July 26, 2011

Abstract

In this work we investigate the hardness of a computational problem introduced in the recent work of Baumslag *et al.* in [1, 2]. In particular, we study the B_n-LHN problem, which is a generalized version of the learning with errors (LWE) problem, instantiated with a particular family of non-abelian groups (free Burnside groups of exponent 3). In our main result, we demonstrate a random self-reducibility property for B_n-LHN. Along the way, we also prove a sequence of lemmas regarding homomorphisms of free Burnside groups of exponent 3 that may be of independent interest.

Keywords. Random self-reducibility. Learning with errors. Post-quantum cryptography. Non-commutative cryptography. Burnside groups.

1 Introduction

Motivation & Background. In the recent work of Baumslag *et al.* [1, 2], the authors propose a generalization of the *learning parity with noise* (LPN) and *learning with errors* (LWE) problems to an abstract class of group-theoretic learning problems, termed *learning homomorphisms with noise* (LHN). As shown in [1], this class of problems contains LPN and LWE as special cases, but also allows instantiations based on non-abelian groups. Specifically, the work of [1] describes a combinatorial instantiation of LHN using a class of relatively free groups known as *Burnside groups*, which are in some sense the "most general" groups for which every element has a finite order dividing some constant n. The Burnside group instantiation of the problem was termed learning Burnside homomorphisms with noise (B_n-LHN). While many aspects of the B_n-LHN problem were addressed in [1], the question of complexity reductions (*e.g.,* worst-case to average-case reductions) remained open. We take steps toward resolving these questions by showing a certain random self-reducibility property for B_n-LHN.

Random self-reducibility. Since any practical implementation of a cryptographic scheme must include an algorithm which generates hard problem instances, it is desirable that such instances

*The City College of CUNY, {fazio,wes}@cs.ccny.cuny.edu
†Pepperdine University, kiga@pepperdine.edu
‡Stevens Institute of Technology, nicolosi@cs.stevens.edu
§UPMC Univ Paris 06 & INRIA Paris-Rocquencourt Center, SALSA project, ludovic.perret@lip6.fr

do not take much effort to find. One notion which in some sense captures this idea, is that of *random self-reducibility.* Roughly speaking, a random self-reducibility property makes an assertion about the average-case hardness of a computational problem. In particular, it says that solving the problem on a *random* instance is not any easier than solving the problem on an *arbitrary* instance. Hence, if a computational problem satisfies random self-reducibility, it is a trivial matter to sample "good" instances: a random instance will suffice. Indeed, random self-reducibility is one of the hallmarks of intractability assumptions that have withstood the test of time. Notable examples include the RSA problem [10]; the discrete logarithm problem and the Diffie-Hellman problem [3]; the quadratic residuosity assumption [4]; the composite residuosity assumption [8]; and the learning with errors (LWE) problem [9].

Our Contributions. In this paper, we make progress towards understanding the computational hardness of learning Burnside homomorphisms with noise. In particular, we establish a random self-reducibility property for B_n-LHN, by showing that learning under uniform surjective secret homomorphisms is no easier than learning under an arbitrary one. We remark that the original formulation of B_n-LHN did not require that the homomorphism be a surjection. However, this limitation seems rather inconsequential for the cryptographic application of the assumption. First, as the security parameter grows, the probability of sampling a non-surjective secret diminishes exponentially. Hence the distributions of instances coming from the two variations on the assumption are in fact statistically close (*cf.* Section 4). Moreover, as shown in Section 5, there is an efficiently computable test for surjectivity, so that the distribution of instances for the modified assumption remains efficiently sampleable (via rejection sampling).

Techniques. Most of the technical lemmas regarding homomorphisms of free Burnside groups of exponent 3 (denoted B_n) involve relating the groups and their morphisms to their abelianized counterparts ($B_n/[B_n, B_n]$), as well as finding certain useful facts that are preserved under this relation. In a study with such a focus on homomorphisms, a number of elementary ideas from homological algebra apply naturally. In particular, we make frequent use of commutative diagrams, exact sequences, and occasionally, the five lemma. These techniques are briefly reviewed in Section 2.

Organization. Section 2 briefly reviews basic notions from group-theory and homological algebra, while Section 3 provides some background on free Burnside groups of exponent 3. The learning Burnside homomorphisms with noise (B_n-LHN) problem is defined in Section 4. Section 5 presents the random self-reducibility result for B_n-LHN.

2 Background: Group Theory and Homological Algebra

Free groups. If X is a subset of a group G, let $X^{-1} = \{x^{-1} \mid x \in X\}$. An expression w of the form $a_1 \dots a_n$ ($n \geq 0,\ a_i \in X \cup X^{-1}$) is termed a ***word*** or an ***X-word***. Such an X-word is said to be ***reduced*** if $n > 0$ and no subword $a_i a_{i+1}$ takes either of the forms xx^{-1} or $x^{-1}x$. If F is a group and X is a subset of F such that X generates F and every reduced X-word is different from 1_F, then one says that F is a ***free group***, freely generated by the set X, and refers to X as a ***free set*** of generators of F, and writes F as $F(X)$. A key property of a free group F freely generated by a set X is that for every group H, every mapping θ from X into H can be extended uniquely to a homomorphism θ_* from F into H. If θ_* is a surjection, and if K is the kernel of θ_*, then the quotient group F/K is isomorphic to H. If R is a subset of F, then in the event that K is generated by all of the conjugates of the elements of R, we express this by writing $H = \langle X; R \rangle$

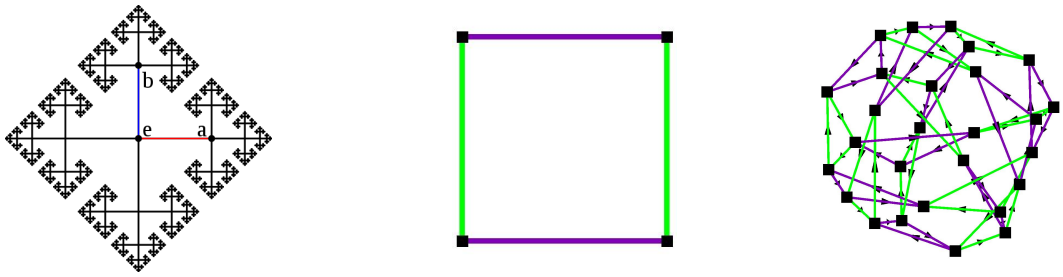

(a) Free group, $F(a, b)$ (b) $B(2, 2) \cong (\mathbb{F}_2^2, +)$ (c) $B(2, 3) = B_2$

Figure 1: Cayley graphs for various groups

and term the pair $\langle X; R\rangle$ a ***presentation*** of H (notice that the mapping θ is usually implicit).

Relatively Free Groups. If F is a free group and K a normal subgroup of F, then the factor group F/K is called ***relatively free*** if K is ***fully invariant***, *i.e.,* if $\alpha(K) \leq K$ for any endomorphism α of F. If $x_1, \ldots, x_n$ are free generators of F, then $x_1K, \ldots, x_nK$ are called relatively free generators of F/K, and typically denoted simply by $x_1, \ldots, x_n$ when there is no risk of confusion. Let E_n denote a relatively free group of rank n, *i.e.,* $F_n = F(x_1, \ldots, x_n)$ and $E_n = F_n/K$ for some fully invariant K. One key property of such a group is that any set map on its generators into E_n can be extended to an endomorphism of E_n. Hence, one is immediately equipped with an exponential number of homomorphisms, provided that the image is non-trivial.

Cayley distance. Finitely generated groups can also be viewed as geometric objects via the notion of the ***Cayley graph***. The Cayley graph of a group G relative to a particular set of generators has the group elements as vertexes, and an edge between two vertexes if and only if multiplication by a generator (or its inverse) translates one to the other. Figure 1 depicts Cayley graphs for several groups, including the 27-element *Burnside* group $B(2, 3)$ of exponent 3 with 2 generators. (Burnside groups are discussed in Section 3.) The ***Cayley distance*** between two group elements is defined as the length of the shortest path between the corresponding nodes in the Cayley graph. The maximum Cayley distance between any two elements in the graph is the ***diameter*** of the Cayley graph. The ***Cayley norm*** of an element x, denoted $\|x\|$, is its distance from the identity element in the Cayley graph. We remark that $\max_{x\in G}(\|x\|)$ corresponds precisely to the diameter.

Commutators. In non-abelian groups, the ***commutator*** of two group elements a, b, denoted $[a, b]$, is the group element satisfying the identity $ab = ba[a, b]$, that is, $[a, b] = a^{-1}b^{-1}ab$. Starting with the generators $x_1, \ldots, x_n$ of the group as the recursive basis, one obtains an ordered sequence of ***formal commutators*** by combining two formal commutators a, b into the formal commutator $[a, b]$. The ***weight*** of a formal commutator is defined by assigning weight 1 to the generators, and defining the weight of $[a, b]$ as the sum of the weights of a and b. The weight imposes a partial order on formal commutators, which is typically made total by assuming an arbitrary ordering among formal commutators of any given weight greater than 1, and by adopting the lexicographical order among the generators.

Exact Sequences and the Five Lemma. Consider a sequence of homomorphisms of groups

$$A_1 \xrightarrow{f_1} A_2 \xrightarrow{f_2} A_3 \xrightarrow{f_3} \cdots \xrightarrow{f_{k-1}} A_k.$$

Such a sequence is said to be ***exact*** if for every $j \in \{1, \ldots, k-2\}$ we have that $\mathrm{Im}(f_j) = \ker(f_{j+1})$. One common example is $0 \longrightarrow H \xrightarrow{i} G \xrightarrow{p} G/H \longrightarrow 0$ where $H \lhd G$, i is the inclusion of H into G, and p is the canonical epimorphism of G onto the quotient, sending $g \mapsto gH$. Consider the following commutative diagram, where the rows are exact.

$$\begin{array}{ccccccccc} A & \longrightarrow & B & \longrightarrow & C & \longrightarrow & D & \longrightarrow & E \\ \downarrow{\scriptstyle e} & & \downarrow{\scriptstyle f} & & \downarrow{\scriptstyle g} & & \downarrow{\scriptstyle h} & & \downarrow{\scriptstyle i} \\ A' & \longrightarrow & B' & \longrightarrow & C' & \longrightarrow & D' & \longrightarrow & E' \end{array}$$

The ***five lemma*** states that if e is surjective and i is injective, then if f and h are isomorphisms, so is g. Furthermore, if i is injective and f and h are surjective, then g is also surjective.[1]

3 Brief Background on Burnside Groups

For a positive integer k, consider the class of groups for which all elements x satisfy $x^k = 1$. Such a group is said to be of *exponent* k. We will be interested in a certain family of such groups called the *free Burnside groups of exponent* k, which are in some sense the "largest." The free Burnside groups are uniquely determined by two parameters: the number of generators n, and the exponent k. We will denote these groups by $B(n, k)$:

Definition 3.1 (Free Burnside group) *For any $n, k \geq 0$, the* Burnside group of exponent k with n generators *is defined as*

$$B(n,k) = \langle \{x_1, \ldots, x_n\}; \{w^k \mid \textit{ for all words } w \textit{ over } x_1, \ldots, x_n\} \rangle.$$

The question of whether $B(n, k)$ is finite or not is known as the *bounded Burnside problem*. For sufficiently large k, $B(n, k)$ is generally infinite [7]. For small exponents, it is known that $k \in \{2, 3, 4, 6\}$ yields finite groups for all n. (We remark that with the exception of $k = 2$, these are non-trivial results.) For other small values of k (most notably, $k = 5$), the question remains open.

For the purposes of this paper, we will be interested primarily in groups of exponent 3; hence in what follows we will denote $B(n, 3)$ simply by B_n for brevity. Next, we review some important facts about B_n (see [6, 5] for a fuller account).

B_n IS FREE. In the category of groups of exponent 3, B_n is a free object on the set of generators $\{x_1, \ldots, x_n\}$. That is, if G is any group such that $g^3 = 1$ for all $g \in G$, then for any set map $f : \{x_1, \ldots, x_n\} \longrightarrow G$, there exists a unique homomorphism $\overline{f} : B_n \longrightarrow G$ such that $\overline{f}(x_i) = f(x_i)$ for every $i \in [n]$.

NORMAL FORM OF B_n. Although B_n is non-abelian, an interesting consequence of the order law $w^3 = 1$ for $w \in B_n$ is that B_n has a simple normal form: Each B_n-element can be written uniquely as an ordered sequence of (a subset of) generators (or their inverses[2]), appearing in lexicographical order, followed by (a subset of) the commutators of weight 2 (or their inverses), and finally by (a

[1] Dually, if e is surjective and f, h injective, then g is also injective.

subset of) the commutators of weight 3 (or their inverses):

$$x_1^{\alpha_1} \cdots x_i^{\alpha_i} \cdots x_n^{\alpha_n} [x_1, x_2]^{\beta_{1,2}} \cdots [x_i, x_j]^{\beta_{i,j}} \cdots [x_{n-1}, x_n]^{\beta_{n-1,n}} [x_1, x_2, x_3]^{\gamma_{1,2,3}}$$
$$\cdots [x_i, x_j, x_k]^{\gamma_{i,j,k}} \cdots [x_{n-2}, x_{n-1}, x_n]^{\gamma_{n-2,n-1,n}} = \prod_{i=1}^{n} x_i^{\alpha_i} \prod_{i<j} [x_i, x_j]^{\beta_{i,j}} \prod_{i<j<k} [x_i, x_j, x_k]^{\gamma_{i,j,k}}$$

where all $\alpha_i, \beta_{i,j}, \gamma_{i,j,k} \in \{0, 1, -1\}$ for all $1 \leq i < j < k \leq n$, and $[x_i, x_j, x_k] = [[x_i, x_j], x_k]$.

ORDER OF B_n. From the above normal form, it follows that B_n has exactly $3^{n+\binom{n}{2}+\binom{n}{3}}$ elements.

CENTER OF B_n. The center, $Z(B_n) = \{g \in B_n \mid [g, h] = 1 \, \forall h \in B_n\}$ is the subgroup $[[B_n, B_n], B_n]$ generated by all commutators of weight 3. This follows in part from the fact that all commutators of weight 4 are the identity in B_n.

HOMOMORPHISMS FROM B_n TO B_r. There are $3^{n\left(r+\binom{r}{2}+\binom{r}{3}\right)}$ homomorphisms from $B_n \longrightarrow B_r$. This follows immediately from the order of B_r and from the fact that B_n is a free object in the category of groups of exponent 3 with generating set of size n.

4 Learning Burnside Homomorphisms with Noise

In this section, we review (a variant of) a group-theoretic learning problem introduced in [1], under the name of learning Burnside homomorphisms with noise (B_n-LHN). Our formulation of the problem samples only surjective homomorphisms for problem instances, in contrast to [1] which samples uniformly over all homomorphisms. As we show in Section 4.2, this modification is of essentially no consequence from a computational perspective. In what follows, for groups G, H we'll denote the set of epimorphisms (*i.e.,* surjective homomorphisms) from G to H by $\mathrm{Epi}(G, H)$.

4.1 The B_n-LHN Problem

For a security parameter $n > 0$, the B_n-LHN setting consists of the groups $G_n \doteq B_n$ and $P_n \doteq B_r$, where $2 \leq r \leq 4$. Let Φ_n be the uniform distribution over the set of *surjective* homomorphisms from B_n to B_r: $\Phi_n \doteq \mathbf{U}(\mathrm{Epi}(B_n, B_r))$. At a high level, the B_n-LHN problem is to distinguish random $G_n \times P_n$ pairs from random (preimage, "noisy" image) pairs under a hidden homomorphism $\varphi \xleftarrow{\$} \Phi_n$. The "noise" in the pairs is determined by an error distribution Ψ_n on B_r, which amounts to taking a randomly ordered product of a random subset of the generators and their inverses. More precisely, the probability mass function of Ψ_n is defined as:

$$\forall e \in B_r, \qquad \Pr_{E \xleftarrow{\$} \Psi_n}[E = e] = \Pr_{\mathbf{v} \xleftarrow{\$} \mathbb{F}_3^r, \sigma \xleftarrow{\$} S_r} \left[e = \prod_{i=1}^{r} x_{\sigma(i)}^{v_i} \right] \tag{1}$$

where the x_i's are the generators of B_r, the v_i's are the components of $\mathbf{v}$, and S_r denotes the symmetric group on r letters. Since $x^2 = x^{-1}$ in B_r, the norm $\|e\|$ of a Ψ_n-sample e is at most r.

Definition 4.1 (LHN-Decision Problem) *For a (B_n, B_r)-homomorphism φ, define the distribution $\mathbf{A}_\varphi^{\Psi_n}$ on $B_n \times B_r$ whose samples are preimage / distorted image pairs (a, b) where $a \xleftarrow{\$} \mathbf{U}(B_n)$ and $b = \varphi(a)e$ for $e \xleftarrow{\$} \Psi_n$. The* LHN-decision problem *is to distinguish the uniform distribution $\mathbf{U}(B_n \times B_r)$ from $\mathbf{A}_\varphi^{\Psi_n}$, for $\varphi \xleftarrow{\$} \Phi_n$.*

[2]Note that $x^{-1} = x^2$ in B_n, as B_n has exponent 3.

Since B_n is a relatively free group, any mapping of its n generators uniquely extends to a homomorphism, and hence $\mathbf{U}(\hom(B_n, B_r))$ is efficiently sampleable. Furthermore, we argue that surjective homomorphisms account for an overwhelming fraction of $\hom(B_n, B_r)$ (*cf.* Section 4.2) and are efficiently recognizable (*cf.* Section 5). It follows that Φ_n is efficiently sampleable via rejection sampling.

4.2 Surjective Homomorphisms of B_n

The original phrasing of the B_n-LHN problem sampled secrets uniformly from all of $\hom(B_n, B_r)$, in contrast to the above definition, in which $\Phi_n = \mathbf{U}(\mathrm{Epi}(B_n, B_r))$. We argue below that this modification has minimal impact on the computational aspects of the problem.

First, we note that as the security parameter n grows, the probability of sampling an instance φ that is not surjective is negligible in n. As we will show in Lemma 5.2, a homomorphism $\varphi \in \hom(B_n, B_r)$ is surjective if and only if its corresponding "abelianized" map $\overline{\varphi} \in \hom(\mathbb{F}_3^n, \mathbb{F}_3^r)$ is surjective. Furthermore, any two $\overline{\varphi}, \overline{\varphi}' \in \hom(\mathbb{F}_3^n, \mathbb{F}_3^r)$ are associated (via abelianization) to the same number of (B_n, B_r)-homomorphisms. Hence, to compute the fraction of $\hom(B_n, B_r)$ which is not surjective, we need only compute the fraction of $\hom(\mathbb{F}_3^n, \mathbb{F}_3^r)$ which is not surjective. A crude upper bound that suffices for our purposes can be obtained via the union bound: We estimate the probability that a randomly selected $\overline{\varphi} \in \hom(\mathbb{F}_3^n, \mathbb{F}_3^r)$ is not surjective by bounding the probability that its image is contained in some $r-1$ dimensional subspace V of $\mathbb{F}_3^r$. Specifically, for $\overline{\varphi} \xleftarrow{\$} \hom(\mathbb{F}_3^n, \mathbb{F}_3^r)$ and any subspace $V \subset \mathbb{F}_3^r$ such that $\dim(V) = r-1$, denote with E_V the event that $\mathrm{Im}(\overline{\varphi}) \subset V$. Then we have

$$\Pr\left[\overline{\varphi} \notin \mathrm{Epi}(\mathbb{F}_3^n, \mathbb{F}_3^r)\right] = \Pr\left[\bigcup_{V \subset \mathbb{F}_3^r} E_V\right]$$

where the probability is over $\overline{\varphi} \xleftarrow{\$} \hom(\mathbb{F}_3^n, \mathbb{F}_3^r)$, and the union on the right hand side is over all subspaces of dimension $r-1$. Since each $(r-1)$-dimensional subspace corresponds uniquely (up to sign) to a non-trivial linear equation over the basis of $\mathbb{F}_3^r$, we see that there are $\frac{3^r-1}{2}$ such subspaces. Furthermore, the image of each of the n generators of $\mathbb{F}_3^n$ satisfies the linear equation of the subspace wth a $1/3$ chance, and hence $\Pr[E_V] = \frac{1}{3^n}$ for each V. By the union bound we then have

$$\Pr\left[\bigcup_{V \subset \mathbb{F}_3^r} E_V\right] \leq \frac{3^r - 1}{2 \cdot 3^n} < 3^{r-n}$$

which is negligible in n as long as the gap between r and n is superlogarithmic (*e.g.,* if r is constant fraction of n). We remark that in fact r is bounded by a small constant both in the formulation of the B_n-LHN assumption in this paper (*cf.* Definition 4.1) and in that of [1].

As a consequence, our distribution of instances Φ_n is statistically close to the uniform distribution $\mathbf{U}(\hom(B_n, B_r))$. Indeed, for any $X_n \subset S_n$, we have

$$\Delta(\mathbf{U}(X_n), \mathbf{U}(S_n)) = \frac{|S_n \setminus X_n|}{|S_n|}$$

where $\mathbf{U}(X_n)$ is considered a distribution on S_n by assigning probability 0 to all elements in $S_n \setminus X_n$, and where Δ denotes statistical distance (total variation distance). Hence, whenever

$\nu(n) = |S_n \setminus X_n| \,/\, |S_n|$ is negligible in n (as in our case), then the ensemble of distributions $\mathbf{U}(X_n)$ is statistically close to $\mathbf{U}(S_n)$.

To summarize, for typical choices of the parameters r and n, the computational hardness of the B_n-LHN problem under uniform unconstrained homomorphisms (the original assumption from [1]) *vs.* uniform epimorphisms (Definition 4.1) are information-theoretically equivalent. In other words, constraining the sampling of instances to $\mathrm{Epi}(B_n, B_r)$ does not alter the computational characteristics of the B_n-LHN assumption.

5 Random Self-Reducibility of B_n-LHN

In this section, we establish a random self-reducibility property of the learning Burnside homomorphisms with noise problem: Learning under uniform surjective secret homomorphisms is no easier than learning under an arbitrary one (Theorem 5.6). We start with a general observation regarding the LHN problem over arbitrary groups G_n, P_n (Lemma 5.1). We then characterize (B_n, B_r)-epimorphisms as precisely those maps whose abelianization is an $(\mathbb{F}_3^n, \mathbb{F}_3^r)$-epimorphism (Lemma 5.2). Next, we use Lemma 5.2 to show that any two epimorphisms from B_n to B_r can be converted into each other via automorphisms of B_n (Lemma 5.4). Together with Lemma 5.4, the result of Lemma 5.1 essentially establishes the random self-reducibility of B_n-LHN (Theorem 5.6).

Lemma 5.1 *Let $\big(a, b = \varphi(a) \cdot e\big) \in G_n \times P_n$ be an instance of* LHN *sampled according to $\mathbf{A}_{\varphi}^{\Psi_n}$, and α be a permutation on G_n. It holds that $(a', b) = \big(\alpha(a), b\big) \in G_n \times P_n$ is sampled according to $\mathbf{A}_{\varphi\circ\alpha^{-1}}^{\Psi_n}$.*

Proof: We have $(a' = \alpha(a), b) = \big(\alpha(a), \varphi(a) \cdot e\big) = \big(\alpha(a), \varphi \circ \alpha^{-1}(\alpha(a)) \cdot e\big) = \big(a', \varphi \circ \alpha^{-1}(a') \cdot e\big)$. ■

Let $\rho : G \longrightarrow G/[G,G]$ denote the projection onto the commutator-factor, and consider the following diagram:

$$\begin{array}{ccc} B_n & \xrightarrow{\ \rho\ } & \mathbb{F}_3^n \\ \Big\downarrow{\scriptstyle \varphi} & & \Big\downarrow{\scriptstyle \overline{\varphi}} \\ B_r & \xrightarrow{\ \rho\ } & \mathbb{F}_3^r \end{array}$$

Lemma 5.2 *Let $\varphi \in \hom(B_n, B_r)$, and let $\overline{\varphi} \in \hom(\mathbb{F}_3^n, \mathbb{F}_3^r)$ be the corresponding map on the abelianization. Then φ is surjective $\iff$ $\overline{\varphi}$ is surjective.*

Proof: Consider the following diagram:

$$\begin{array}{ccccccccc} 0 & \longrightarrow & [B_n, B_n] & \xrightarrow{\ i\ } & B_n & \xrightarrow{\ \rho\ } & \mathbb{F}_3^n & \longrightarrow & 0 \\ & & \Big\downarrow{\scriptstyle \hat{\varphi}} & & \Big\downarrow{\scriptstyle \varphi} & & \Big\downarrow{\scriptstyle \overline{\varphi}} & & \\ 0 & \longrightarrow & [B_r, B_r] & \xrightarrow{\ i\ } & B_r & \xrightarrow{\ \rho\ } & \mathbb{F}_3^r & \longrightarrow & 0 \end{array} \tag{2}$$

The short exact sequences are the result of abelianization of B_n and B_r to $\mathbb{F}_3^n$ and $\mathbb{F}_3^r$, respectively. The central vertical map is the given homomorphism φ, and $\hat{\varphi}$ is φ restricted to the commutator subgroup $[B_n, B_n]$. Since homomorphisms map commutators to commutators, $\hat{\varphi}$ maps into $[B_r, B_r]$. The map $\overline{\varphi}$ is obtained by considering the map $\rho \circ \varphi : B_n \longrightarrow \mathbb{F}_3^r$, which is a map to an abelian group and therefore factors through the abelianization of B_n as $\overline{\varphi} \circ \rho$. Thus, this diagram is commutative.

($\Longrightarrow$) If $\varphi \in \mathrm{Epi}(B_n, B_r)$, then a diagram chase around the right-hand square in (2) shows that $\overline{\varphi}$ is also surjective.

($\Longleftarrow$) Now suppose that $\overline{\varphi}$ is surjective. If we can show that $\hat{\varphi}$ is surjective, then by the Five Lemma, φ will be surjective.

Let $\{x_1, \ldots, x_r\}$ be the generators for B_r. By the surjectivity of $\overline{\varphi}$ and the commutativity of the right square of (2), there exist elements $\mathcal{B} = \{e_1, \ldots, e_r\} \subset B_n$ so that

$$\varphi(e_i) = x_i \gamma_i$$

where $\gamma_i \in [B_r, B_r]$. We now consider the following commutative diagram:

$$\begin{array}{ccccccccc} 0 & \longrightarrow & [[B_n, B_n], B_n] & \xrightarrow{j} & [B_n, B_n] & \xrightarrow{\pi} & [B_n, B_n]/[[B_n, B_n], B_n] & \longrightarrow & 0 \\ & & \downarrow{\scriptstyle \alpha} & & \downarrow{\scriptstyle \hat{\varphi}} & & \downarrow{\scriptstyle \beta} & & \\ 0 & \longrightarrow & [[B_r, B_r], B_r] & \xrightarrow{j} & [B_r, B_r] & \xrightarrow{\pi} & [B_r, B_r]/[[B_r, B_r], B_r] & \longrightarrow & 0 \end{array} \tag{3}$$

The map α is the restriction of $\hat{\varphi}$ to $[[B_n, B_n], B_n]$, and as before, it maps into $[[B_r, B_r], B_r]$. The existence of the homomorphism β and the commutativity of the diagram is a standard diagram chase. We now show that α and β are surjective, and then by the Five Lemma, $\hat{\varphi}$ will then be surjective.

Because the lower central series for B_r terminates at $[[B_r, B_r], B_r]$, we know that $[B_r, B_r]$ is generated by $[x_i, x_j]$ and $[[x_i, x_j], x_k]$. So to show that β is surjective, it suffices to show that for every i and j, $\pi([x_i, x_j])$ is in the image of β. Equivalently, for each i and j, there is a $z_{ij} \in [[B_r, B_r], B_r]$ so that $[x_i, x_j] z_{ij}$ is in the image of $\hat{\varphi}$.

We recall that there exist e_i so that $\varphi(e_i) = x_i \gamma_i$. We now compute:

$$\begin{aligned} \hat{\varphi}([e_i, e_j]) &= \varphi([e_i, e_j]) \\ &= [x_i \gamma_i, x_j \gamma_j] \\ &= \gamma_i^{-1} x_i^{-1} \gamma_j^{-1} x_j^{-1} x_i \gamma_i x_j \gamma_j \\ &= x_i^{-1} \gamma_i^{-1} [\gamma_i^{-1}, x_i^{-1}] \gamma_j^{-1} x_j^{-1} x_i \gamma_i x_j \gamma_j \\ &= x_i^{-1} \gamma_i^{-1} \gamma_j^{-1} x_j^{-1} x_i \gamma_i x_j \gamma_j^{-1} [\gamma_i^{-1}, x_i^{-1}] \\ &\vdots \\ &= x_i^{-1} x_j^{-1} x_i x_j \prod z_k \\ &= [x_i, x_j] z_{ij} \end{aligned}$$

where all $z_{ij} \in [[B_r, B_r], B_r]$. Therefore, β is surjective.

Now we wish to show α is surjective. Again we compute

$$\begin{aligned}
\alpha([[e_i,e_j],e_k]) =& \varphi([[e_i,e_j],e_k]) \\
=& [[x_i\gamma_i, x_j\gamma_j], x_k\gamma_k] \\
=& [[x_i,x_j]z_{ij}, x_k\gamma_k] \\
=& [[x_i,x_j],x_k\gamma_k] \\
=& [x_i,x_j]^{-1}\gamma_k^{-1}x_k^{-1}[x_i,x_j]x_k\gamma_k \\
=& [x_i,x_j]^{-1}\gamma_k^{-1}x_k^{-1}x_k[x_i,x_j][[x_i,x_j],x_k]\gamma_k \\
=& [x_i,x_j]^{-1}\gamma_k^{-1}[x_i,x_j]\gamma_k[[x_i,x_j],x_k] \\
=& [x_i,x_j]^{-1}\gamma_k^{-1}\gamma_k[x_i,x_j][[x_i,x_j],\gamma_k][[x_i,x_j],x_k] \\
=& [[x_i,x_j],\gamma_k][[x_i,x_j],x_k].
\end{aligned}$$

At this point we note that the first factor, $[[x_i,x_j],\gamma_k]$, is in $[[B_r,B_r],[B_r,B_r]]$, which is trivial. What remains is $[[x_i,x_j],x_k]$, which we have now shown to be in the image of α. Therefore α is an epimorphism. The Five Lemma on (3) then proves that $\hat{\varphi}$ is an epimorphism, and the Five Lemma on (2) proves that φ is an epimorphism. ■

We also make note of a useful special case of the lemma.

Corollary 5.3 *Let $\{t_1,\ldots,t_m\} \subset \mathbb{F}_3^n$ be a spanning set, and let $\rho : B_n \longrightarrow \mathbb{F}_3^n$ denote the projection onto the abelianization. Then for all $\{a_1,\ldots,a_m\} \subset B_n$ such that $\rho(a_i) = t_i$, the set $\{a_1,\ldots,a_m\}$ generates B_n.*

Proof: Let $\{t_1,\ldots,t_m\} \subset \mathbb{F}_3^n$ and $\{a_1,\ldots,a_m\} \subset B_n$ be subsets satisfying the hypothesis. Consider the map $\varphi : B_m \longrightarrow B_n$ defined by $x_i \mapsto a_i$ for each $i \in [m]$. By hypothesis, the corresponding map $\overline{\varphi} : \mathbb{F}_3^m \longrightarrow \mathbb{F}_3^n$ on the abelianization is surjective, and thus by Lemma 5.2 so is φ. Therefore $\langle a_1,\ldots,a_m\rangle = \mathrm{Im}(\varphi) = B_n$. ■

Remarks. We use lemma 5.2 below in our proof that the randomization from lemma 5.1 is in fact a complete random self reduction, but it also has computational significance: given a description of $\varphi \in \hom(B_n,B_r)$ as mappings of the generators, we now have an easy test for surjectivity: simply compute the rank of the corresponding map of linear spaces in the abelianization.

Next, we show that $\mathrm{Aut}(B_n)$ acts transitively on $\mathrm{Epi}(B_n,B_r)$ by composition on the right, so that the construction from Lemma 5.1 provides a random self-reduction.

Lemma 5.4 $\mathrm{Aut}(B_n)$ *acts transitively on* $\mathrm{Epi}(B_n,B_r)$ *by composition on the right. That is, for any $\varphi,\varphi^* \in \mathrm{Epi}(B_n,B_r)$, there exists $\alpha \in \mathrm{Aut}(B_n)$ such that $\varphi^* = \varphi \circ \alpha$.*

Proof: Let $\varphi^* \in \mathrm{Epi}(B_n,B_r)$ denote the "target" surjection, and let $\varphi \in \mathrm{Epi}(B_n,B_r)$ be an arbitrary surjection. We would like to find $\alpha \in \mathrm{Aut}(B_n)$ such that $\varphi^* = \varphi \circ \alpha$. It will be useful

to refer to the following diagram, where the maps ρ denote the corresponding projections onto the abelianization.

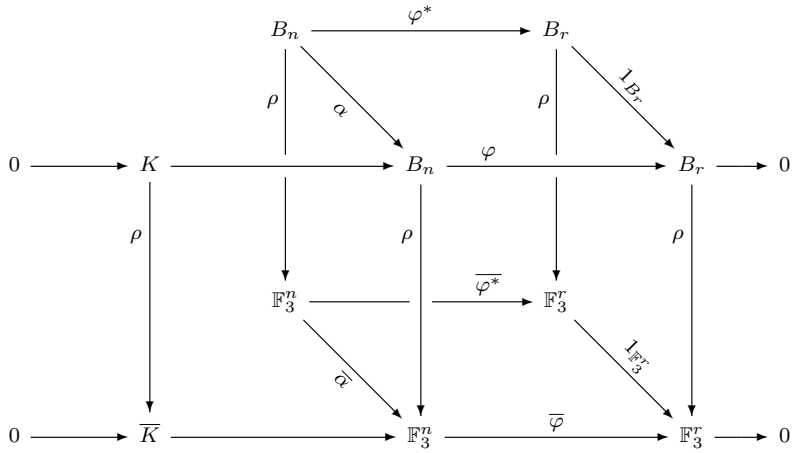

(4)

Let $x_1, \ldots, x_n$ be free generators of B_n and let $t_i = \varphi^*(x_i)$ for $i \in [n]$. Since φ^* is surjective, the t_i must generate B_r. Let $\mathcal{T} \subset \{t_1, \ldots, t_n\}$ be a minimal generating set, and let $S \subset [n]$ denote the corresponding set of indexes (so that $\mathcal{T} = \{t_i\}_{i \in S}$). Note that $|\mathcal{T}| = r$. First, observe that r elements suffice to generate $\rho(B_r)$, and by Corollary 5.3, we know that if $\{\rho(t_i)\}_{i \in S}$ generate $\mathbb{F}_3^r$, then indeed $\{t_i\}_{i \in S}$ generate B_r. Hence $|\mathcal{T}| \leq r$. However, it cannot be the case that $|\mathcal{T}| \lneq r$, since $\rho(\{x_i\}_{i \in S})$ must span a space of dimension at least r by commutativity of (4). Therefore $|\mathcal{T}| = r$.

Since φ is surjective, $\exists a_i \in B_n$ such that $\varphi(a_i) = t_i$ for every $i \in [n]$. Let $\mathcal{A} = \{a_i\}_{i \in S}$. We now modify the list of "leftover" elements $\{a_i\}_{i \in [n] \setminus S}$, so that each may be expressed in terms of $\mathcal{A}$. For each $i \in [n] \setminus S$, if $t_i \neq 1 \in B_r$, then t_i may be expressed as a word w_i in $\mathcal{T}$. For each $i \in [n] \setminus S$, we redefine a_i correspondingly, using the rule

$$a_i = \begin{cases} w_i(\mathcal{A}) & \text{if } t_i \neq 1; \\ 1 & \text{otherwise.} \end{cases}$$

Note that the modified set $\{a_i\}_{i \in [n]}$ still satisfies the property that $\varphi(a_i) = t_i$, yet $\{a_i\}_{i \in [n]} \subset \langle \mathcal{A} \rangle$.

Next, we prove that $\rho : K \longrightarrow \overline{K}$ is an epimorphism, *i.e.,* that $\rho(K) = \overline{K}$. Toward this end, let $y \in \overline{K} = \ker(\overline{\varphi})$. By the surjectivity of $\rho : B_n \longrightarrow \mathbb{F}_3^n$, there exists a $b \in B_n$ so that $\rho(b) = y$. Since $0 = \overline{\varphi}(y) = \overline{\varphi}(\rho(b)) = \rho(\varphi(b))$, we have that $\varphi(b) \in [B_r, B_r]$. By the proof of Lemma 5.2, we know that φ restricted to $[B_n, B_n]$ maps surjectively to $[B_r, B_r]$, so there exists a $\gamma \in [B_n, B_n]$ so that $\varphi(\gamma) = \varphi(b)$. Then $\varphi(b\gamma^{-1}) = 1$ so $b\gamma^{-1} \in K$. Now $\rho(b\gamma^{-1}) = \rho(b) + \rho(\gamma^{-1}) = y - 0 = y$. Therefore, $\rho(K) = \overline{K}$. Hence, if $\overline{\mathcal{K}}$ is a basis for $\ker(\overline{\varphi})$, for each $\overline{k_i} \in \overline{\mathcal{K}}$ we can find pre-images $k_i \in K = \ker(\varphi)$. We denote this set by $\mathcal{K}$. It will be convenient to use the $n - r$ elements of $[n] \setminus S$ as indices so that $\mathcal{K} = \{k_i\}_{i \in [n] \setminus S}$. We essentially construct α by sending $x_i \mapsto a_i$, but in order to ensure $\alpha \in \mathrm{Aut}(B_n)$ we tack on additional kernel elements when necessary. More precisely, we set

$$\alpha(x_i) = \begin{cases} a_i, & \text{if } i \in S; \\ a_i k_i & \text{if } i \in [n] \setminus S. \end{cases}$$

Since B_n is free on the x_i, this assignment defines a unique homomorphism α on B_n. Since $\varphi^*(x_i) = \varphi \circ \alpha(x_i)$ for all x_i, we have that $\varphi^* = \varphi \circ \alpha$. We now study the abelianization $\overline{\alpha}$ of

α. Note that $\rho(\mathcal{A})$ must have dimension r, and furthermore, $\langle \rho(\mathcal{A}) \rangle \cap \overline{K} = \{0\}$ by commutativity of (4).[3] Therefore, $\rho(\mathcal{A}) \cup \overline{\mathcal{K}}$ is a basis of $\mathbb{F}_3^n$. Furthermore, letting $u_i = \rho(x_i)$, we see that $\overline{\alpha}(u_i) = \rho(\alpha(x_i)) = \rho(a_i k_j) = \rho(a_i) + \overline{k}_j$.

By construction, we see that $\rho(\mathcal{A}) \cup \overline{\mathcal{K}} \subset \mathrm{Span}(\overline{\alpha}(x_1), \ldots, \overline{\alpha}(x_n))$, and hence the $\overline{\alpha}(x_i)$ span all of $\mathbb{F}_3^n$. Therefore $\overline{\alpha}$ is an automorphism of $\mathbb{F}_3^n$. By Lemma 5.2, it follows that α is also an automorphism of B_n. ■

Lemma 5.5 *Let G be a finite group, and S a set on which G acts transitively. Let $s \in S$ be an arbitrary element, and consider the distribution A_s on S whose samples are $g \cdot s$ where $g \xleftarrow{\$} \mathbf{U}(G)$. Then $A_s = \mathbf{U}(S)$.*

Proof: Let $t \in S$ be an arbitrary element. We wish to compute $\Pr[A_s = t]$; that is, the probability that $g \cdot s = t$, over the uniform choice of $g \xleftarrow{\$} G$. Recall that the ***stabilizer*** of $s \in S$ (denoted by $\mathrm{stab}(s)$) is the subgroup of G defined by $\mathrm{stab}(s) = \{g \in G \mid g \cdot s = s\}$. Note that $\#\{g \mid g \cdot s = t\}$ is given by $|\mathrm{stab}(s)|$ since $g \cdot s = g' \cdot s \iff g'^{-1} g \in \mathrm{stab}(s)$, which states that g, g' are in the same coset modulo $\mathrm{stab}(s)$.[4] Recall that

$$[G : \mathrm{stab}(t)] = |G \cdot t| = |S|$$

with the last equality following from the transitivity of the action. Hence, $|\mathrm{stab}(t)| = |G|/|S|$, and

$$\Pr[A_s = t] = \frac{|G|}{|S|} \cdot \frac{1}{|G|} = \frac{1}{|S|}$$

which completes the proof. ■

Theorem 5.6 (B_n-LHN Random Self-Reducibility) *With notation as in Definition 4.1, any instance of the B_n-LHN-decision problem in which φ is an* arbitrary *surjection from B_n onto B_r can be reduced to a B_n-LHN-instance in which $\varphi \xleftarrow{\$} \mathrm{Epi}(B_n, B_r)$.*

Proof: This is a straightforward consequence of Lemma 5.1, Lemma 5.4, and Lemma 5.5. Let $\varphi \in \mathrm{Epi}(B_n, B_r)$ be an arbitrary surjection and suppose we are given a distribution $\mathbf{R}$ which is either $\mathbf{U}(B_n \times B_r)$ or $\mathbf{A}_{\varphi}^{\Psi_n}$. Let $\alpha \xleftarrow{\$} \mathrm{Aut}(B_n)$. (Note that by Lemma 5.2, we can efficiently, sample such an α.) We then construct a new distribution $\mathbf{R}'$ whose samples are $(\alpha(a), b)$ where $(a, b) \xleftarrow{\$} \mathbf{R}$. If $\mathbf{R} = \mathbf{U}$, then $\mathbf{R}' = \mathbf{U}$ as well, since α is a bijection. Otherwise, by Lemma 5.1, $\mathbf{R}' = \mathbf{A}_{\varphi'}^{\Psi_n}$, where $\varphi' = \varphi \circ \alpha^{-1}$. Moreover, by Lemma 5.4 and Lemma 5.5, we see that φ' is distributed according to $\mathbf{U}(\mathrm{Epi}(B_n, B_r))$. It follows that an algorithm to solve the B_n-LHN-decision problem on a random epimorphism can be used to solve the B_n-LHN-decision problem on an arbitrary one. ■

6 Conclusions and Future Work

In this work, we take steps towards understanding the computational hardness of the B_n-LHN problem put forth in [1]. With a minor modification to the problem formulation (which results in an instance distribution statistically close to the original), we demonstrate a strong random self-reducibility property, giving evidence that the B_n-LHN problem is difficult in the average case.

[3] This can also be seen via the rank theorem, which states that $\mathbb{F}_3^n = \langle \rho(\mathcal{A}) \rangle \oplus \overline{K}$, with $\rho(\mathcal{A}) \cup \overline{\mathcal{K}}$ serving as a basis.

[4] This argument requires the existence of at least one g such that $g \cdot s = t$; we are given such a g by transitivity.

Future work includes continued efforts to assess the hardness of the B_n-LHN problem—either via explicit algorithms that demonstrate upper bounds on its complexity, or via further reductions to other computational problems.

References

[1] G. Baumslag, N. Fazio, A. R. Nicolosi, V. Shpilrain, and W.E. Skeith III. Generalized learning problems and applications to non-commutative cryptography. In *International Conference on Provable Security—ProvSec '11 (to appear)*. Springer, 2011. LNCS.

[2] Gilbert Baumslag, Nelly Fazio, Antonio R. Nicolosi, Vladimir Shpilrain, and William E. Skeith III. Generalized learning problems and applications to non-commutative cryptography. Cryptology ePrint Archive, Report 2011/357, 2011. Full version of [1]. Available at `http://eprint.iacr.org/2011/357`.

[3] W. Diffie and M.E. Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, IT-22(6):644–654, 1976.

[4] S. Goldwasser and S. Micali. Probabilistic encryption. *JCSS*, 28(2):270–299, 1984.

[5] N. Gupta. On groups in which every element has finite order. *Amer. Math. Month.*, 96:297–308, 1989.

[6] M. Hall. *The Theory of Groups*. Macmillan Company, New York, 1959.

[7] Sergei V. Ivanov. The free Burnside groups of sufficiently large exponents. *Internat. J. Algebra Comput.*, 4(1-2):ii+308, 1994.

[8] Pascal Paillier. Public-key cryptosystems based on composite degree residuosity classes. In *EUROCRYPT*, pages 223–238, 1999.

[9] O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In *STOC*, pages 84–93. ACM Press, 2005.

[10] R. Rivest, A. Shamir, and L. Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Communications of the ACM*, 21:120–126, 1978.