

Hardness of Learning Problems over Burnside Groups of Exponent 3

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

Abstract

In this work we investigate the hardness of a computational problem introduced in the recent work of Baumslag *et al.* in [3, 4]. 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.

*The City College of CUNY, {fazio,wes}@cs.cuny.cuny.edu. Supported in part by NSF grants CNS 1117675.

†Pepperdine University, kiga@pepperdine.edu

‡Stevens Institute of Technology, nicolosi@cs.stevens.edu. Supported in part by NSF grants CNS 1117679.

§UPMC Univ Paris 06 & INRIA Paris-Rocquencourt Center, SALSA project, ludovic.perret@lip6.fr

1 Introduction

MOTIVATION & BACKGROUND. In the recent work of Baumslag *et al.* [3, 4], the authors derive a number of basic cryptographic primitives (*e.g.*, symmetric encryption) from a generalization of the *learning parity with noise* (LPN [1, 11, 5]) and *learning with errors* (LWE [15, 14, 12, 2]) problems to an abstract class of group-theoretic learning problems, termed *learning homomorphisms with noise* (LHN). As shown in [3], this class of problems contains LPN and LWE as special cases, but also allows instantiations based on non-abelian groups. Specifically, the work of [3] describes a combinatorial instantiation of LHN using a class of finite groups known as *free 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 a number of cryptographic aspects of the B_n -LHN problem were addressed in [3] (*cf.* Appendix A for a discussion on the computational aspects of Burnside groups, and their relevance for cryptographic applications), several important matters were left open; perhaps the most prominent being the question of complexity reductions (*e.g.*, worst-case to average-case reductions). 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 do not take much effort to find. One notion that 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 [16]; the discrete logarithm problem and the Diffie-Hellman problem [6]; the quadratic residuosity assumption [7]; the composite residuosity assumption [13]; and the learning with errors (LWE) problem [15]. As it turns out, however, random self-reducibility properties come in several shapes. For example, the type of random self-reducibility enjoyed by the LWE is, in a sense, the strongest, in that *the secret key itself* can be randomized: given instances relative to a secret \mathbf{s} , new instances relative to a uniformly random secret \mathbf{s}' can be constructed in a way that solutions to the latter yield solutions to the former. This is a more complete form of random self-reducibility than what is known for many number-theoretic assumptions, like RSA, where it is possible to randomize *individual instances* based on a given private key, but for which there is no apparent way to re-randomize the *key itself*. More concretely, given an instance $c = m^e \bmod n$, one can compute a new instance $c' = cr^e \bmod n = (mr)^e \bmod n$, whose solution (together with knowledge of r) yields a solution for c , yet there is no apparent way to find a connected instance relative to a different modulus $n' \neq n$. We stress that the reduction shown in our work is of the LWE type: the worst-case to average-case reduction applies to the secret keys.

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.* Appendix B). 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). Finally, in Section 6, we present a limited form of a search-to-decision reduction for B_n -LHN.

TECHNIQUES. Most of the technical lemmas regarding homomorphisms of free Burnside groups of expo-

ment 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 Appendix 2.

ORGANIZATION. 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. Section 6 looks into the relationship between the search and decision versions of 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 $H = \langle X; R \rangle$ 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, \dots, x_n are free generators of F , then $x_1 K, \dots, x_n K$ are called relatively free generators of F/K , and typically denoted simply by x_1, \dots, x_n when there is no risk of confusion. Let E_n denote a relatively free group of rank n , *i.e.*, $F_n = F(x_1, \dots, x_n)$ 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, \dots, 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.

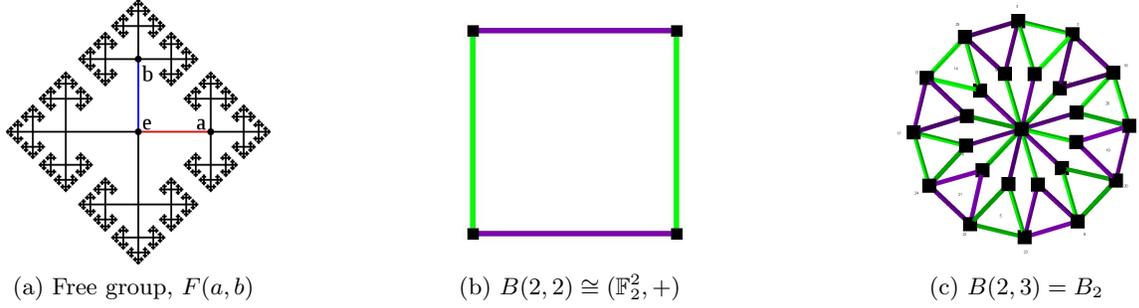


Figure 1: Cayley graphs for various groups

COMMUTATOR SUBGROUPS. If G is a group, then the **commutator subgroup** $[G, G]$ is the subgroup of G generated by elements of the form $[a, b]$ where $a, b \in G$. The commutator subgroup of G is a normal subgroup of G . More generally, if A and B are normal subgroups of G , then $[A, B]$ is the subgroup of G generated by elements of the form $[a, b]$ where $a \in A$ and $b \in B$. Thus, we can define a sequence of subgroups $G_1 = G$, $G_2 = [G, G]$, $G_3 = [[G, G], G]$, and recursively, $G_{n+1} = [G_n, G]$. This sequence of groups is called the **lower central series** for G .

ABELIANIZATION. If G is a group, then $G/[G, G]$ is called the **abelianization of G** : It is an abelian group obtained from G by creating new relations to ensure that all elements of the group commute. The canonical epimorphism $\rho_G : G \rightarrow G/[G, G]$ that takes $g \in G$ into its $g[G, G]$ in $G/[G, G]$ is usually referred to as the **projection onto the abelianization** of G . When there is no risk of confusion, we will drop the subscript and denote the projection onto the abelianization simply by ρ .

It is a basic fact from category theory that ρ is *universal* for homomorphisms from G into abelian groups: that is, if A is any abelian group, then any homomorphism $\theta : G \rightarrow A$ splits as $\theta = \theta' \circ \rho$, for a unique homomorphism $\theta' : G/[G, G] \rightarrow A$. If G and H are groups, and $\phi : G \rightarrow H$ is a homomorphism, then using the above construction on $\theta = \rho_H \circ \phi$ yields a homomorphism $\bar{\phi} : G/[G, G] \rightarrow H/[H, H]$ so that the following diagram commutes:

$$\begin{array}{ccc}
 G & \xrightarrow{\phi} & H \\
 \rho_G \downarrow & & \downarrow \rho_H \\
 G/[G, G] & \xrightarrow{\bar{\phi}} & H/[H, H]
 \end{array}$$

We refer to $\bar{\phi}$ as the **abelianization of the homomorphism ϕ** .

EXACT SEQUENCES AND THE FIVE LEMMA. Consider a sequence of homomorphisms of groups

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

Such a sequence is said to be **exact** if for every $j \in \{1, \dots, k-2\}$ we have that $\text{Im}(f_j) = \ker(f_{j+1})$. One common example is $0 \rightarrow H \xrightarrow{j} G \xrightarrow{p} G/H \rightarrow 0$ where $H \triangleleft G$, j 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.

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

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*

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

Since we are interested in average-case hardness, it is important that $B(n, k)$ be finite, else even basic issues regarding the probability distribution become unclear. 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 [10]. 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 $k = 2$, for which $B(n, k) = \mathbb{F}_2^n$ is abelian, these are non-trivial results.) For other small values of k (most notably, $k = 5$), the question remains open.

To ensure finiteness, our current knowledge of Burnside groups would require k to be in the set $\{2, 3, 4, 6\}$; however, following the work of [3], we will focus on $k = 3$. The main reasons are as follows: $k = 2$ would give the more familiar (and already studied) case $B(n, k) = \mathbb{F}_2^n$; it is convenient for k to be prime (hence eliminating $k = 4$ and $k = 6$); and perhaps most importantly, the structure of $B(n, 3)$ is much better understood in comparison to than that of $k = 4, 6$. Hence, in what follows we will deal only with $B(n, 3)$ and denote it simply by B_n for brevity.

Next, we review some important facts about B_n (see also Appendix B, or [9, 8] 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 $\{x_1, \dots, x_n\}$. That is, if G is any group such that $g^3 = 1$ for all $g \in G$, then for any set map $f : \{x_1, \dots, x_n\} \rightarrow G$, there exists a unique homomorphism $\bar{f} : B_n \rightarrow G$ such that $\bar{f}(x_i) = f(x_i)$ for every $i \in [n]$. In other words, to define a homomorphism from B_n to G we need only define the function on $\{x_1, \dots, x_n\}$. Any such assignment will extend uniquely to a group homomorphism.

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 inverses²), appearing in lexicographical order, followed by (a subset of) the commutators of weight 2 (or their inverses), and finally by (a subset of) the commutators of weight 3 (or their inverses):

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

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.

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

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

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 $3^{n(r+\binom{r}{2}+\binom{r}{3})}$ 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 [3], 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 [3] which samples uniformly over all homomorphisms. As we show in Appendix B, this modification is of essentially no consequence from a computational perspective. In what follows, for groups G, H we will denote the set of epimorphisms (*i.e.*, surjective homomorphisms) from G to H by $\text{Epi}(G, H)$.

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 $2 \leq r$.³ Let Φ_n be the uniform distribution over the set of *surjective* homomorphisms from B_n to B_r : $\Phi_n \doteq \mathbf{U}(\text{Epi}(B_n, B_r))$. 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:

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

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 $\varphi \xleftarrow{\$} \Phi_n$.

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

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), which immediately yields a partial key-randomization property for LHN in general. We then show that this

³For the cryptographic applications described in [3], it was required that $r \leq 4$ so that a certain computational problem in B_r remained feasible. We do not need any such restrictions here.

randomization is in fact complete for the specific case of B_n -LHN if we restrict the B_n -LHN secret key to be surjective. This essentially follows from proving that any two epimorphisms from B_n to B_r can be converted into each other via automorphisms of B_n (Lemma 5.4). In turn, this “transitivity” property hinges upon a technical lemma that characterizes (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). Together with Lemma 5.1, Lemma 5.4 essentially establishes the random self-reducibility of B_n -LHN under surjective homomorphisms (Theorem 5.6). (In light of the argument of Appendix B about the prominence of epimorphisms among (B_n, B_r) -homomorphisms, Theorem 5.6 additionally yields the random self-reducibility of the original B_n -LHN assumption from [3].)

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

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

Let $\rho : G \longrightarrow G/[G, G]$ denote the projection onto the abelianization. Consider the following diagram:

$$\begin{array}{ccc} B_n & \xrightarrow{\rho} & \mathbb{F}_3^n \\ \varphi \downarrow & & \downarrow \bar{\varphi} \\ B_r & \xrightarrow{\rho} & \mathbb{F}_3^r \end{array} \quad (2)$$

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

Proof: Consider the following diagram:

$$\begin{array}{ccccccccc} 0 & \longrightarrow & [B_n, B_n] & \xrightarrow{i} & B_n & \xrightarrow{\rho} & \mathbb{F}_3^n & \longrightarrow & 0 \\ & & \hat{\varphi} \downarrow & & \varphi \downarrow & & \bar{\varphi} \downarrow & & \\ 0 & \longrightarrow & [B_r, B_r] & \xrightarrow{i} & B_r & \xrightarrow{\rho} & \mathbb{F}_3^r & \longrightarrow & 0 \end{array} \quad (3)$$

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 $\bar{\varphi}$ 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 $\bar{\varphi} \circ \rho$. Thus, this diagram is commutative.

(\implies) If $\varphi \in \text{Epi}(B_n, B_r)$, then a diagram chase around (2) shows that $\bar{\varphi}$ is also surjective.

(\impliedby) Now suppose that $\bar{\varphi}$ is surjective. Let $\{x_1, \dots, x_r\}$ be the generators for B_r . Ideally, we would like to argue that x_i is in the image of φ for all i , which would yield the desired result. Surjectivity of $\bar{\varphi}$ does not immediately imply that φ hits all generators x_i of B_r ; nevertheless, along with commutativity of the right square of (3), it guarantees that φ hits a collection of “quasi-generators”: there exist elements $e_1, \dots, e_r \in B_n$ and $\gamma_1, \dots, \gamma_r \in [B_r, B_r]$ so that

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

Next, we leverage the existence of the e_i ’s and the properties of their commutators to argue surjectivity of $\hat{\varphi}$; surjectivity of φ will follow by the Five Lemma.

To show that $\hat{\varphi}$ is surjective, in turn we invoke the Five Lemma on the commutative diagram below:

$$\begin{array}{ccccccc}
0 & \longrightarrow & [[B_n, B_n], B_n] & \xrightarrow{j} & [B_n, B_n] & \xrightarrow{\pi} & [B_n, B_n]/[[B_n, B_n], B_n] \longrightarrow 0 \\
& & \downarrow \alpha & & \downarrow \hat{\varphi} & & \downarrow \beta \\
0 & \longrightarrow & [[B_r, B_r], B_r] & \xrightarrow{j} & [B_r, B_r] & \xrightarrow{\pi} & [B_r, B_r]/[[B_r, B_r], B_r] \longrightarrow 0
\end{array} \tag{4}$$

The map α is the restriction of $\hat{\varphi}$ (and hence of φ) to $[[B_n, B_n], B_n]$, and since commutators map to commutators, 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. All that remains to prove is then surjectivity of α and β .

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 $\{[x_i, x_j]\}_{i,j}$ and $\{[[x_i, x_j], x_k]\}_{i,j,k}$. Therefore, $[B_r, B_r]/[[B_r, B_r], B_r]$ is generated by $\{\pi[x_i, x_j]\}_{i,j}$, and 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 .

Recall that there exist e_i so that $\varphi(e_i) = x_i\gamma_i$. We show below that $\pi([e_i, e_j])$ is a pre-image of $\pi([x_i, x_j])$ under β . We compute:

$$\hat{\varphi}([e_i, e_j]) = \varphi([e_i, e_j]) = [\varphi(e_i), \varphi(e_j)] = [x_i\gamma_i, x_j\gamma_j] = \gamma_i^{-1}x_i^{-1}\gamma_j^{-1}x_j^{-1}x_i\gamma_i x_j\gamma_j \tag{5}$$

We then introduce a commutator to reverse the order of the first two elements:

$$= x_i^{-1}\gamma_i^{-1}[\gamma_i^{-1}, x_i^{-1}]\gamma_j^{-1}x_j^{-1}x_i\gamma_i x_j\gamma_j$$

and then use the fact that $[\gamma_i^{-1}, x_i^{-1}] \in [[B_n, B_n], B_n]$ is in the center, to move this commutator past the other terms, to the far right:

$$= x_i^{-1}\gamma_i^{-1}\gamma_j^{-1}x_j^{-1}x_i\gamma_i x_j\gamma_j[\gamma_i^{-1}, x_i^{-1}]. \tag{6}$$

Comparing this result with (5), we note that this sequence of manipulations allowed us to move the first term γ_i^{-1} one element to the right, at the expense of generating one element $[\gamma_i^{-1}, x_i^{-1}] \in [[B_r, B_r], B_r]$ at the far right. Using this same technique, we move γ_i^{-1} to the right again and again, until it is adjacent to γ_i , at which point it cancels γ_i . As before, each such move produces another $[[B_r, B_r], B_r]$ -factor on the right. Next, we deal with γ_j^{-1} analogously: We move it to the right (each time producing additional $[[B_r, B_r], B_r]$ -elements) until it cancels γ_j . At the end of this procedure, continuing from (6), we have

$$\dots = x_i^{-1}x_j^{-1}x_i x_j \prod y_k$$

where each y_k is some element of $[[B_r, B_r], B_r]$. If we call this product $z_{i,j} = \prod y_k$, we then have

$$\hat{\varphi}([e_i, e_j]) = [x_i, x_j]z_{i,j}$$

where $z_{i,j} \in [[B_r, B_r], B_r]$. We then see that

$$\beta(\pi([e_i, e_j])) = \pi(\hat{\varphi}([e_i, e_j])) = \pi([x_i, x_j]z_{i,j}) = \pi([x_i, x_j]).$$

Therefore, β is surjective.

Now we wish to show α is surjective. Since $[[B_r, B_r], B_r]$ is generated by $\{[[x_i, x_j], x_k]\}_{i,j,k}$, it suffices to show that these are all in the image of α . In particular, we show below that, for all i, j, k , $\alpha([e_i, e_j], e_k) = [[x_i, x_j], x_k]$. Again we compute

$$\alpha([e_i, e_j], e_k) = \varphi([e_i, e_j], e_k) = [[x_i\gamma_i, x_j\gamma_j], x_k\gamma_k].$$

Recalling from the previous computation that $[x_i\gamma_i, x_j\gamma_j] = [x_i, x_j]z_{i,j}$, we have

$$= [[x_i, x_j]z_{i,j}, x_k\gamma_k]$$

Since $z_{i,j} \in [[B_r, B_r], B_r]$ is central, we can commute it and its inverse together in this expression and cancel them:

$$= [[x_i, x_j], x_k\gamma_k].$$

We then expand this commutator

$$= [x_i, x_j]^{-1}\gamma_k^{-1}x_k^{-1}[x_i, x_j]x_k\gamma_k$$

and introduce $[[x_i, x_j], x_k]$ to move x_k to the left:

$$= [x_i, x_j]^{-1}\gamma_k^{-1}x_k^{-1}x_k[x_i, x_j][[x_i, x_j], x_k]\gamma_k$$

We then use the fact that $[[x_i, x_j], x_k] \in [[B_r, B_r], B_r]$ is central to move this to the right.

$$= [x_i, x_j]^{-1}\gamma_k^{-1}[x_i, x_j]\gamma_k[[x_i, x_j], x_k]$$

and move γ_k to the left using another commutator:

$$= [x_i, x_j]^{-1}\gamma_k^{-1}\gamma_k[x_i, x_j][[x_i, x_j], \gamma_k][[x_i, x_j], x_k]$$

at which point we can first cancel γ_k^{-1} with γ_k , and then cancel $[x_i, x_j]^{-1}$ with $[x_i, x_j]$, to get

$$= [[x_i, x_j], \gamma_k][[x_i, x_j], x_k].$$

Next, 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 (4) then proves that $\hat{\varphi}$ is an epimorphism, and the Five Lemma on (3) proves that φ is an epimorphism. ■

We also make note of the following simple but useful consequence of this Lemma.

Corollary 5.3 *Let $\rho : B_r \longrightarrow \mathbb{F}_3^r$ denote the projection onto the abelianization. Let $\{t_1, \dots, t_n\} \subset B_r$. Then $\{t_1, \dots, t_n\}$ generates B_r if and only if $\{\rho(t_1), \dots, \rho(t_n)\}$ generates \mathbb{F}_3^r .*

Proof: Consider the map $\varphi : B_n \longrightarrow B_r$ defined by $x_i \mapsto t_i$ for each $i \in [n]$, and let $\bar{\varphi} : \mathbb{F}_3^n \longrightarrow \mathbb{F}_3^r$ be the abelianization. Then $\{t_1, \dots, t_n\}$ generates B_r if and only if φ is surjective, and $\{\rho(t_1), \dots, \rho(t_n)\}$ generates \mathbb{F}_3^r if and only if $\bar{\varphi}$ is surjective. By Lemma 5.2, these two conditions are equivalent. ■

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 $\varphi \in \text{hom}(B_n, B_r)$ 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 $\text{Aut}(B_n)$ acts transitively on $\text{Epi}(B_n, B_r)$ by composition on the right, and thus for the case of B_n -LHN, the construction from Lemma 5.1 provides a random self-reduction.

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

Proof: Let $\varphi^* \in \text{Epi}(B_n, B_r)$ denote the “target” surjection, and let $\varphi \in \text{Epi}(B_n, B_r)$ be an arbitrary surjection. We would like to find $\alpha \in \text{Aut}(B_n)$ such that $\varphi^* = \varphi \circ \alpha$. In other words, we wish to define a bijective map α so that the following diagram commutes:

$$\begin{array}{ccc}
 B_n & \xrightarrow{\varphi^*} & B_r \\
 \alpha \downarrow & & \downarrow 1_{B_r} \\
 B_n & \xrightarrow{\varphi} & B_r
 \end{array} \tag{7}$$

Let x_1, \dots, x_n be free generators of B_n . To define α , it suffices to define $\alpha(x_i)$ for each $i \in [n]$. To derive suitable $\alpha(x_i)$ values such that α as a whole is bijective, it is convenient to study the abelianization of all the groups and maps in (7), which results in the following diagram:

$$\begin{array}{ccccccc}
 & & B_n & \xrightarrow{\varphi^*} & B_r & & \\
 & & \rho \downarrow & \searrow \alpha & \downarrow \rho & \searrow 1_{B_r} & \\
 0 & \longrightarrow & K & \xrightarrow{\quad} & B_n & \xrightarrow{\varphi} & B_r \longrightarrow 0 \\
 & & \tau \downarrow & & \rho \downarrow & & \downarrow \rho \\
 & & \mathbb{F}_3^n & \xrightarrow{\quad} & \mathbb{F}_3^n & \xrightarrow{\varphi^*} & \mathbb{F}_3^r \\
 & & \downarrow \alpha & & \downarrow \rho & \searrow 1_{\mathbb{F}_3^r} & \\
 0 & \longrightarrow & \overline{K} & \xrightarrow{\quad} & \mathbb{F}_3^n & \xrightarrow{\overline{\varphi}} & \mathbb{F}_3^r \longrightarrow 0
 \end{array} \tag{8}$$

In this diagram, the vertical maps ρ denote the projections onto the abelianization. K is the kernel of φ , \overline{K} is the kernel of $\overline{\varphi}$, and τ is an epimorphism from K to \overline{K} (essentially just the restriction of $\rho : B_n \longrightarrow \mathbb{F}_3^n$ to $K \subset B_n$) that is defined in Step 3 below.

Step 1: Finding a minimal subset \mathcal{T} of the $\{\varphi^*(x_i)\}_{i \in [n]}$ that generates B_r .

Let $t_i = \varphi^*(x_i)$ for $i \in [n]$. Since φ^* is surjective, the t_i must generate B_r . Let $\mathcal{T} \subset \{t_1, \dots, t_n\}$ be a minimal generating set for B_r , and let $S \subset [n]$ denote the corresponding set of indexes (so that $\mathcal{T} = \{t_i\}_{i \in S}$). By Corollary 5.3, $\rho(\mathcal{T})$ is also a minimal generating set for \mathbb{F}_3^r . Since \mathbb{F}_3^r is a vector space of dimension r , we know that \mathcal{T} , and thus S , has r elements.

Step 2: Finding a set $\{a_i\}_{i \in [n]}$ of φ -preimages of the $\{t_i\}_{i \in [n]}$ of minimal “rank”.

We would like to define elements $\{a_i\}_{i \in [n]} \subset B_n$ such that $\varphi(a_i) = t_i$ for every $i \in [n]$, yet the subgroup of B_n generated by the $\{a_i\}_{i \in [n]}$ admits a generating set $\mathcal{A} = \{a_i\}_{i \in S}$ with only r elements. To this aim, for $i \in S$ we invoke surjectivity of φ on t_i to find $a_i \in B_n$ such that $\varphi(a_i) = t_i$. For $i \in [n] \setminus S$, instead, we leverage the fact that \mathcal{T} is a generating set for B_r and define a_i as a word over \mathcal{A} . In detail, let w_i be a word that expresses t_i in terms of \mathcal{T} , *i.e.*, such that $t_i = w_i(\mathcal{T})$. We then define

$$a_i = w_i(\mathcal{A}).$$

Note that $\varphi(a_i) = \varphi(w_i(\mathcal{A})) = w_i(\varphi(\mathcal{A})) = w_i(\mathcal{T}) = t_i$; moreover, \mathcal{A} generates $\langle a_1, \dots, a_n \rangle$ by construction.

Step 3: Defining $\tau : K \longrightarrow \overline{K}$.

We define a morphism $\tau : K \longrightarrow \overline{K}$ to make the leftmost portion of diagram (8) above commute. To do this, we note that the maps from K to B_n and from \overline{K} to \mathbb{F}_3^n are inclusions. Now if $k \in K$,

then $\rho(\varphi(k)) = \rho(1) = 0$. By commutativity of the right square in (8), $\overline{\varphi}(\rho(k)) = \rho(\varphi(k)) = 0$, *i.e.*, $\rho(k) \in \ker(\overline{\varphi}) = \overline{K}$. Defining $\tau : k \in K \mapsto \rho(k) \in \overline{K}$ maps K to \overline{K} and makes the diagram commute.

Step 4: Proving $\tau : K \longrightarrow \overline{K}$ is a surjection.

Next, we prove that $\tau : K \longrightarrow \overline{K}$ is an epimorphism, *i.e.*, that $\tau(K) = \overline{K}$. 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 $\tau(b\gamma^{-1}) = \rho(b\gamma^{-1}) = \rho(b) + \rho(\gamma^{-1}) = y - 0 = y$, and thus $\tau(K) = \overline{K}$.

Step 5: Defining a minimal generating set $\mathcal{K} \subset K$.

Consider the map $\overline{\varphi} : \mathbb{F}_3^n \longrightarrow \mathbb{F}_3^r$ in (8). This is a surjective linear map of vector spaces, and so its kernel is a linear subspace of dimension $n - r$. Choose a basis $\overline{\mathcal{K}}$ for $\ker(\overline{\varphi})$. Since τ is surjective, for each $\overline{k}_i \in \overline{\mathcal{K}}$, we find pre-images $k_i \in K$ with $\tau(k_i) = \overline{k}_i$. 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 $\mathcal{K} = \{k_i\}_{i \in [n] \setminus S}$.

Step 6: Defining the homomorphism α .

We are now ready to define α . Since x_1, \dots, x_n are free generators, to define α it will suffice to define $\alpha(x_i)$ for each $i \in [n]$. The natural choice might seem to be defining $\alpha(x_i) = a_i$, but in order to ensure α is bijective, it will be necessary to tack on elements of K when necessary. In particular, define

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

Since B_n is free on the x_i , this assignment defines a unique homomorphism α on B_n .

Step 7: Proving that $\varphi^* = \varphi \circ \alpha$.

By freeness of B_n , it suffices to show that φ^* and $\varphi \circ \alpha$ agree on x_1, \dots, x_n . For all $i \in S$, we have $\varphi \circ \alpha(x_i) = \varphi(a_i) = \varphi^*(x_i)$. For all $i \notin S$, we have $\varphi \circ \alpha(x_i) = \varphi(a_i k_i) = \varphi(a_i) \varphi(k_i) = \varphi(a_i) = \varphi^*(x_i)$.

Step 8: Proving that the abelianization $\overline{\alpha}$ of α is an epimorphism.

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 (8).⁴ 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 $\overline{\alpha}(u_i) = \rho(\alpha(x_i))$. This, in turn, depends on whether $i \in S$ or $i \notin S$:

$$\overline{\alpha}(u_i) = \begin{cases} \rho(a_i), & i \in S \\ \rho(a_i) + \overline{k}_i, & i \notin S \end{cases} \quad (9)$$

Now for $i \in S$, $\overline{\alpha}(u_i) = \rho(a_i)$, so $\rho(\mathcal{A}) \subset \text{Span}(\overline{\alpha}(u_1), \dots, \overline{\alpha}(u_n))$. For $i \notin S$, $\overline{\alpha}(u_i) = \rho(w_i(\mathcal{A})) + \overline{k}_i$, and subtracting off $\rho(w_i(\mathcal{A})) = w_i(\rho(\mathcal{A}))$ (which is clearly in $\text{Span}(\overline{\alpha}(u_1), \dots, \overline{\alpha}(u_n))$), we see that also $\overline{k}_i \in \text{Span}(\overline{\alpha}(u_1), \dots, \overline{\alpha}(u_n))$. Therefore, we see that $\rho(\mathcal{A}) \cup \overline{\mathcal{K}} \subset \text{Span}(\overline{\alpha}(u_1), \dots, \overline{\alpha}(u_n))$, and hence the $\overline{\alpha}(u_i)$ span all of \mathbb{F}_3^n . Thus, $\overline{\alpha}$ is an epimorphism.

Step 9: Proving α is an isomorphism.

By Lemma 5.2, it follows that α is also an epimorphism of B_n . Since α is a surjective map from a finite set into itself, it is also bijective. Therefore α is an automorphism of \mathbb{F}_3^n , which completes the proof. ■

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 $g \stackrel{\$}{\leftarrow} \mathbf{U}(G)$. 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 $g \stackrel{\$}{\leftarrow} G$. Recall that the *stabilizer* of $s \in S$ (denoted by $\text{stab}(s)$) is the subgroup of G defined by $\text{stab}(s) = \{g \in G \mid g \cdot s = s\}$. Note that $\#\{g \mid g \cdot s = t\}$ is given by $|\text{stab}(s)|$ since $g \cdot s = g' \cdot s \iff g'^{-1}g \in \text{stab}(s)$, which states that g, g' are in the same coset modulo $\text{stab}(s)$.⁵ Recall that

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

⁴This 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.

⁵This argument requires the existence of at least one g such that $g \cdot s = t$; we are given such a g by transitivity.

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

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

which completes the proof. \blacksquare

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 $\varphi \stackrel{\$}{\leftarrow} \text{Epi}(B_n, B_r)$.*

Proof: This is a straightforward consequence of Lemma 5.1, Lemma 5.4, and Lemma 5.5. Let $\varphi \in \text{Epi}(B_n, B_r)$ 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 $\alpha \stackrel{\$}{\leftarrow} \text{Aut}(B_n)$. (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 $(a, b) \stackrel{\$}{\leftarrow} \mathbf{R}$. 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 $\mathbf{U}(\text{Epi}(B_n, B_r))$. 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. \blacksquare

6 A Weak Decision-to-Search Equivalence

In this part, we investigate the relation between B_r -LHN and LHN-Decision. We first observe the following.

Proposition 6.1 *Let q be the cardinality⁶ of B_r . Let also \mathcal{A} be an algorithm distinguishing $\mathbf{A}_\varphi^{\Psi_n}$ from $\mathbf{U}(B_n \times B_r)$ in time t with advantage at least ϵ , i.e. \mathcal{A} is an algorithm solving LHN-Decision. Finally, let $x \stackrel{\$}{\leftarrow} B_n$. There exists an algorithm \mathcal{B} working in $\text{poly}(q, t)$ such that:*

$$\Pr_{(a,b) \stackrel{\$}{\leftarrow} \mathbf{A}_\varphi^{\Psi_n}} \left[\mathcal{B}(a, b, x) = \varphi(x) \right] \geq \frac{\epsilon + 1}{q},$$

Proof: We can w.l.o.g. assume that the success probability of \mathcal{A} is greater on $\mathbf{A}_\varphi^{\Psi_n}$ than on $\mathbf{U}(B_n \times B_r)$, i.e. it holds that:

$$\Pr_{(a,b) \stackrel{\$}{\leftarrow} \mathbf{A}_\varphi^{\Psi_n}} \left[\mathcal{A}(a, b) = 1 \right] - \Pr_{(a,b) \stackrel{\$}{\leftarrow} \mathbf{U}(B_n \times B_r)} \left[\mathcal{A}(a, b) = 1 \right] \geq \epsilon.$$

Let $\text{Pr}_1 = \Pr_{(a,b) \stackrel{\$}{\leftarrow} \mathbf{A}_\varphi^{\Psi_n}} \left[\mathcal{A}(a, b) = 1 \right]$, and $\text{Pr}_2 = \Pr_{(a,b) \stackrel{\$}{\leftarrow} \mathbf{U}(B_n \times B_r)} \left[\mathcal{A}(a, b) = 1 \right]$.

Algorithm \mathcal{B} works as follows. It samples $(a, b) \stackrel{\$}{\leftarrow} \mathbf{A}_\varphi^{\Psi_n}$, a guess $h \stackrel{\$}{\leftarrow} \mathbf{U}(B_r)$ for the value of $\varphi(x)$, and invokes algorithm \mathcal{A} on $(x \cdot a, h \cdot b) = (a', b')$. Finally, \mathcal{B} returns h if \mathcal{A} returns 1 and any value of $B_r \setminus \{h\}$ otherwise.

Suppose $h = \varphi(x)$, then $h \cdot \varphi(a) = \varphi(x \cdot a)$. It follows that $(a', b') = (x \cdot a, h \cdot \varphi(a) \cdot e)$ is sampled according to $\mathbf{A}_\varphi^{\Psi_n}$. Hence:

$$\Pr_{(a,b) \stackrel{\$}{\leftarrow} \mathbf{A}_\varphi^{\Psi_n}} \left[\mathcal{B}(a, b, x) = \varphi(x) \right] = \Pr_{(a',b') \stackrel{\$}{\leftarrow} \mathbf{A}_\varphi^{\Psi_n}} \left[\mathcal{A}(a', b') = 1 \right] = \text{Pr}_1.$$

Assume now that $h \neq \varphi(x)$. In this case, $(a', b') = (x \cdot a, h \cdot \varphi(a) \cdot e)$ is distributed according to $\mathbf{U}(B_n \times B_r)$. Thus:

$$\Pr_{(a,b) \stackrel{\$}{\leftarrow} \mathbf{A}_\varphi^{\Psi_n}} \left[\mathcal{B}(a, b, x) = \varphi(x) \right] = \frac{\Pr_{(a',b') \stackrel{\$}{\leftarrow} \mathbf{U}(B_n \times B_r)} \left[\mathcal{B}(a', b') = 0 \right]}{q - 1} = \frac{1 - \text{Pr}_2}{(q - 1)}.$$

⁶We recall that B_r is independent of the security parameter. Thus, we enumerate all elements of B_r efficiently, i.e. $O(1)$.

As a consequence:

$$\Pr_{(\mathbf{p}, \mathbf{b}) \xleftarrow{\$} \mathbf{A}_\varphi^{\Psi_n}} \left[B(\mathbf{p}, \mathbf{b}, x) = \varphi(x) \right] = \frac{1}{q} \left(\Pr_1 + (q-1) \frac{(1 - \Pr_2)}{q-1} \right) = \frac{\Pr_1 - \Pr_2 + 1}{q} \geq \frac{\epsilon + 1}{q}.$$

■

Thus, Proposition 6.1 proves that an oracle for LHN-Decision allows to predict the results of $\varphi(x)$, for arbitrary $x \xleftarrow{\$} B_n$, slightly better than guessing randomly.

In the high advantage case (i.e. the distinguisher is perfect), we get:

Corollary 6.2 (“Weak” Decision-to-Search) *Let q be the cardinality of B_r . If $\mathbf{A}_\varphi^{\Psi_n}$ and $\mathbf{U}(B_n \times B_r)$ are perfectly distinguishable, i.e. there exists a distinguisher \mathcal{A} working in time t accepting with probability exponentially close to 1 elements from $\mathbf{A}_\varphi^{\Psi_n}$ and rejecting with probability exponentially close to 1 elements from $\mathbf{U}(B_n \times B_r)$ then there is an algorithm \mathcal{C} working in $\text{poly}(t, q)$ such that:*

$$\Pr_{(a,b) \xleftarrow{\$} \mathbf{A}_\varphi^{\Psi_n}, y \xleftarrow{\$} \mathcal{C}(a,b)} \left[\varphi(y) = b \right] \text{ with probability exponentially close to 1.}$$

Proof: The algorithm \mathcal{C} to consider is exactly the algorithm described in the proof of Prop. 6.1. ■

We emphasize that the general case (arbitrary advantage) remains an open problem. Interestingly enough, this reduces to generalize the famous Goldreich-Levin Theorem to non-abelian groups. The obstacle on the proofs seems to be the impossibility – due to non-commutativity – to sufficiently amplify the success probability of Proposition 6.1.

7 Conclusions and Future Work

In this work, we take steps towards understanding the computational hardness of the B_n -LHN problem put forth in [3]. 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.

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. In particular, one interesting open problem is to fully reduce the search version of B_n -LHN to the corresponding decision version.

References

- [1] D. Angluin and P. Laird. Learning from noisy examples. *Machine Learning*, 2(4):343–370, 1988.
- [2] Sanjeev Arora and Rong Ge. New algorithms for learning in presence of errors. Manuscript, 2011.
- [3] 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.
- [4] 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 [3]. Available at <http://eprint.iacr.org/2011/357>.
- [5] A. Blum, A. Kalai, and H. Wasserman. Noise-tolerant learning, the parity problem, and the statistical query model. *J. ACM*, 50:2003, 2003.
- [6] W. Diffie and M.E. Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, IT-22(6):644–654, 1976.
- [7] S. Goldwasser and S. Micali. Probabilistic encryption. *JCSS*, 28(2):270–299, 1984.
- [8] N. Gupta. On groups in which every element has finite order. *Amer. Math. Month.*, 96:297–308, 1989.
- [9] M. Hall. *The Theory of Groups*. Macmillan Company, New York, 1959.
- [10] Sergei V. Ivanov. The free Burnside groups of sufficiently large exponents. *Internat. J. Algebra Comput.*, 4(1-2):ii+308, 1994.
- [11] M. Kearns. Efficient noise-tolerant learning from statistical queries. In *Journal of the ACM*, pages 392–401. ACM Press, 1993.
- [12] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In *EUROCRYPT*, pages 1–23, 2010.
- [13] Pascal Paillier. Public-key cryptosystems based on composite degree residuosity classes. In *EUROCRYPT*, pages 223–238, 1999.
- [14] Chris Peikert. Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In *STOC*, pages 333–342, 2009.
- [15] O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In *STOC*, pages 84–93. ACM Press, 2005.
- [16] 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.

A Computational Aspects of Burnside Groups

In order for the Burnside groups to be of use in cryptography, at a minimum, they must have a concise representation, and the group operation must be efficiently computable. We demonstrate here that both criteria are met. First, we note that as described above, each element of B_n has a unique normal form as a product of the generators and certain commutators. Hence by storing an array of the exponents (each of which is in the set $\{0, 1, -1\}$) we can uniquely represent an element. The size of the array is cubic in n .

As for the group operation, this can be computed simply by concatenating two normal forms, and then reducing the resulting word back into normal form. This process, referred to as the *collection process*, takes cubic time (see [9], chap. 11) in the length of the input (which is itself cubic in n). However, all commutators of weight 3 are in the center $Z(B_n)$ of B_n , and hence there is no need to expand them and apply the collection process—one can simply add the corresponding exponents modulo 3. Furthermore, since all commutators of weight 4 are trivial (see [9], chap. 18), we know that $[B_n, B_n]$ is commutative. Hence, we can again avoid the collection process when moving the weight-2 commutators amongst themselves, and in cubic time, we can reduce the expression to a “nearly” normal form consisting of a product of at most $2n$ generators (or their inverses) followed by commutators in normal form. Therefore we need only to apply the collection process on linear input, and so the overall running time of computing the product is indeed $\mathcal{O}(n^3)$. Inverses can also be computed over B_n in at most cubic time by a similar (yet somewhat simpler) collecting process.

The last and most challenging computational aspect of B_n relates to its *geodesics*—the computation of distances in the Cayley graph. For the applications introduced in [3], it suffices to compute the *norm* (i.e., the distance to the identity of the group), in the codomain group P_n , which is generally small, and does not necessarily grow with the security parameter (although it may grow with a correctness parameter). For the case of the free Burnside group B_r , one possible solution is to perform a breadth-first search of the Cayley graph, storing the norm of every element in a table. This process will begin to become infeasible around $r = 5$. However, even with this small number of generators, the diameter is large enough to properly decode for many interesting error distributions Ψ_n .

B Surjective Homomorphisms of B_n

The original phrasing of the B_n -LHN problem sampled secrets uniformly from all of $\text{hom}(B_n, B_r)$, in contrast to the above definition, in which $\Phi_n = \mathbf{U}(\text{Epi}(B_n, B_r))$. 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 $\varphi \in \text{hom}(B_n, B_r)$ is surjective if and only if its corresponding “abelianized” map $\bar{\varphi} \in \text{hom}(\mathbb{F}_3^n, \mathbb{F}_3^r)$ is surjective. Furthermore, any two $\bar{\varphi}, \bar{\varphi}' \in \text{hom}(\mathbb{F}_3^n, \mathbb{F}_3^r)$ are associated (via abelianization) to the same number of (B_n, B_r) -homomorphisms. Hence, to compute the fraction of $\text{hom}(B_n, B_r)$ which is not surjective, we need only compute the fraction of $\text{hom}(\mathbb{F}_3^n, \mathbb{F}_3^r)$ 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 $\bar{\varphi} \in \text{hom}(\mathbb{F}_3^n, \mathbb{F}_3^r)$ 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 $\bar{\varphi} \stackrel{\$}{\leftarrow} \text{hom}(\mathbb{F}_3^n, \mathbb{F}_3^r)$ and any subspace $V \subset \mathbb{F}_3^r$ such that $\dim(V) = r - 1$, denote with E_V the event that $\text{Im}(\bar{\varphi}) \subset V$. Then we have

$$\Pr [\bar{\varphi} \notin \text{Epi}(\mathbb{F}_3^n, \mathbb{F}_3^r)] = \Pr \left[\bigcup_{\substack{V \subset \mathbb{F}_3^r \\ \dim(V) = r - 1}} E_V \right]$$

where the probability is over $\bar{\varphi} \stackrel{\$}{\leftarrow} \text{hom}(\mathbb{F}_3^n, \mathbb{F}_3^r)$, 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 with a $1/3$ chance, and hence $\Pr[E_V] = \frac{1}{3^n}$ for each V . By the union bound we then have

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

which is negligible in n as long as the gap between r and n is superlogarithmic (*e.g.*, if r is a 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 [3].

As a consequence, our distribution of instances Φ_n is statistically close to the uniform distribution $\mathbf{U}(\text{hom}(B_n, B_r))$. 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 [3]) *vs.* uniform epimorphisms (Definition 4.1) are information-theoretically equivalent. In other words, constraining the sampling of instances to $\text{Epi}(B_n, B_r)$ does not alter the computational characteristics of the B_n -LHN assumption.