Random Self-Reducibility of Learning Problems over Burnside Groups

William Skeith

CCNY and Graduate Center
CAISS

Joint work with Nelly Fazio, Kevin Iga, Antonio Nicolosi and Ludovic Perret

Outline

1. **Motivation & Background**
   - Why Group-Theoretic Cryptography?
   - Random self-reducibility
2. **Learning Problems Over Burnside Groups**
   - Background: LWE
   - LHN Problem
   - Burnside Groups and B_n-LHN
3. **The Reduction, in 3 Easy Steps**
   - Step 1: An Observation
   - Step 2: Completeness for Surjections
   - Step 3: Irrelevance of the Restriction

Outline

Motivation

- Interesting mathematical problem on its own . . .
- Tackling crypto challenges of post-quantum era [Sh'94]
  - Shor's algorithm: Efficient *quantum* procedure to compute the order of any element in a cyclic group
  - Hardness of order-finding at the heart of most popular public-key cryptosystems (RSA, DH, ECDH)
  - If quantum computing becomes practical, we'll need alternative crypto platforms
- Quantum computing aside, **diversifying assumptions** still seems prudent

Motivation

- Interesting mathematical problem on its own . . .
- Tackling crypto challenges of post-quantum era [Sh'94]
  - Shor's algorithm: Efficient *quantum* procedure to compute the order of any element in a cyclic group
  - Hardness of order-finding at the heart of most popular public-key cryptosystems (RSA, DH, ECDH)
  - $\therefore$ If quantum computing becomes practical, we'll need alternative crypto platforms
- Quantum computing aside, **diversifying assumptions** still seems prudent

Motivation

- Interesting mathematical problem on its own . . .
- Tackling crypto challenges of post-quantum era [Sh'94]
  - Shor's algorithm: Efficient *quantum* procedure to compute the order of any element in a cyclic group
  - Hardness of order-finding at the heart of most popular public-key cryptosystems (RSA, DH, ECDH)
  - $\therefore$ If quantum computing becomes practical, we'll need alternative crypto platforms
- Quantum computing aside, **diversifying assumptions** still seems prudent

Motivation

- Interesting mathematical problem on its own . . .
- Tackling crypto challenges of post-quantum era [Sh'94]
  - Shor's algorithm: Efficient *quantum* procedure to compute the order of any element in a cyclic group
  - Hardness of order-finding at the heart of most popular public-key cryptosystems (RSA, DH, ECDH)
  - ∴ If quantum computing becomes practical, we'll need alternative crypto platforms
- Quantum computing aside, **diversifying assumptions** still seems prudent

Motivation

- Interesting mathematical problem on its own . . .
- Tackling crypto challenges of post-quantum era [Sh'94]
  - Shor's algorithm: Efficient *quantum* procedure to compute the order of any element in a cyclic group
  - Hardness of order-finding at the heart of most popular public-key cryptosystems (RSA, DH, ECDH)
  - $\therefore$ If quantum computing becomes practical, we'll need alternative crypto platforms
- Quantum computing aside, **diversifying assumptions** still seems prudent

Motivation

- Interesting mathematical problem on its own ...
- Tackling crypto challenges of post-quantum era [Sh'94]
  - Shor's algorithm: Efficient *quantum* procedure to compute the order of any element in a cyclic group
  - Hardness of order-finding at the heart of most popular public-key cryptosystems (RSA, DH, ECDH)
  - $\therefore$ If quantum computing becomes practical, we'll need alternative crypto platforms
- Quantum computing aside, **diversifying assumptions** still seems prudent

Prior Work in Non-Commutative Cryptography

Challenging computational problems abound in group theory, however...

- Many hard problems are based on infinite groups
- This makes probabilistic modeling difficult
- Average-case hardness for many problems seems to be not well-understood

Prior Work in Non-Commutative Cryptography

Challenging computational problems abound in group theory, however...

- Many hard problems are based on infinite groups
- This makes probabilistic modeling difficult
- **Average-case hardness for many problems seems to be not well-understood**

In this Work

Main Results

- In this work we demonstrate a random self-reducibility property for a new group-theoretic problem put forth in the work of Baumslag *et al.* [BFNSS11].
- In particular, we show a worst-case to average-case reduction for the B_n-LHN problem (more on that later...)

In this Work

Main Results

- In this work we demonstrate a **random self-reducibility** property for a new group-theoretic problem put forth in the work of Baumslag *et al.* [BFNSS11].
- In particular, we show a **worst-case to average-case** reduction for the B_n-LHN problem (more on that later...)

In this Work

Main Results

- In this work we demonstrate a **random self-reducibility** property for a new group-theoretic problem put forth in the work of Baumslag *et al.* [BFNSS11].
- In particular, we show a **worst-case to average-case** reduction for the B_n-LHN problem (more on that later...)

Outline

1. **Motivation & Background**
   - Why Group-Theoretic Cryptography?
   - Random self-reducibility
2. **Learning Problems Over Burnside Groups**
   - Background: LWE
   - LHN Problem
   - Burnside Groups and B_n-LHN
3. **The Reduction, in 3 Easy Steps**
   - Step 1: An Observation
   - Step 2: Completeness for Surjections
   - Step 3: Irrelevance of the Restriction

Random self-reducibility

Random self-reducibility makes a statement about the average case complexity of a problem. In particular, it says that

Random self-reducibility

Solving a **random** instance is not any easier than solving an **arbitrary** instance.

Random self-reducibility

Random self-reducibility makes a statement about the average case complexity of a problem. In particular, it says that

Random self-reducibility

Solving a **random** instance is not any easier than solving an **arbitrary** instance.

Random Self-Reducibility

- Random self-reducibility has been a hallmark of every successful cryptographic assumption to-date.
- This is not so surprising:
  - Any cryptosystem implementation must include an algorithm which samples hard instances of a computational problem.
  - RSR ensures that hard instances are not difficult to find: a random instance will suffice.

Random Self-Reducibility

- Random self-reducibility has been a hallmark of every successful cryptographic assumption to-date.
- This is not so surprising:
  - Any cryptosystem implementation must include an algorithm which samples **hard instances** of a computational problem.
  - RSR ensures that hard instances are not difficult to find: a random instance will suffice.

- Random self-reducibility has been a hallmark of every successful cryptographic assumption to-date.
- This is not so surprising:
  - Any cryptosystem implementation must include an algorithm which samples **hard instances** of a computational problem.
  - RSR ensures that hard instances are not difficult to find: a random instance will suffice.

Random Self-Reducibility

- Random self-reducibility has been a hallmark of every successful cryptographic assumption to-date.
- This is not so surprising:
  - Any cryptosystem implementation must include an algorithm which samples **hard instances** of a computational problem.
  - RSR ensures that hard instances are not difficult to find: a random instance will suffice.

The B_n-LHN Problem

B_n-LHN

- The problem is a generalization of LWE, moving from vector spaces and inner products to the setting of groups and homomorphisms.
- As shown in [BFNSS11], this assumption suffices for some basic cryptographic tasks, *e.g.,* symmetric encryption.
- We'll start with a quick review of LWE.

The B_n-LHN Problem

B_n-LHN

- The problem is a generalization of LWE, moving from vector spaces and inner products to the setting of groups and homomorphisms.
- As shown in [BFNSS11], this assumption suffices for some basic cryptographic tasks, *e.g.*, symmetric encryption.
- We'll start with a quick review of LWE.

The B_n-LHN Problem

B_n-LHN

- The problem is a generalization of LWE, moving from vector spaces and inner products to the setting of groups and homomorphisms.
- As shown in [BFNSS11], this assumption suffices for some basic cryptographic tasks, *e.g.,* symmetric encryption.
- We'll start with a quick review of LWE.

The B_n-LHN Problem

B_n-LHN

- The problem is a generalization of LWE, moving from vector spaces and inner products to the setting of groups and homomorphisms.
- As shown in [BFNSS11], this assumption suffices for some basic cryptographic tasks, *e.g.,* symmetric encryption.
- We'll start with a quick review of LWE.

Outline

1. **Motivation & Background**
   - Why Group-Theoretic Cryptography?
   - Random self-reducibility
2. **Learning Problems Over Burnside Groups**
   - Background: LWE
   - LHN Problem
   - Burnside Groups and B_n-LHN
3. **The Reduction, in 3 Easy Steps**
   - Step 1: An Observation
   - Step 2: Completeness for Surjections
   - Step 3: Irrelevance of the Restriction

Let $\mathbf{s} \in \mathbb{F}_p^n$. The picture is as follows:

$$\begin{tikzcd} \mathbb{F}_p^n \arrow[d, "\mathbf{s}\cdot\_"'] & \ni & \mathbf{a} \arrow[d, mapsto, "\approx \mathbf{s}\cdot\mathbf{a}"] & & \\ \mathbb{F}_p & \ni & b & = & \mathbf{s}\cdot\mathbf{a} + e \end{tikzcd}$$

LWE, Informally

Roughly, the **Learning With Errors** problem is to recover $\mathbf{s}$ by sampling preimage-image pairs in the presence of some small "noise"

Learning With Errors [Reg05]

Let $\mathbf{s} \in \mathbb{F}_p^n$. The picture is as follows:

$$\begin{tikzcd}
\mathbb{F}_p^n \arrow[d, "\mathbf{s} \cdot _"'] & \ni & \mathbf{a} \arrow[d, mapsto, "\approx \mathbf{s} \cdot \mathbf{a}"] & & \\
\mathbb{F}_p & \ni & b & = & \mathbf{s} \cdot \mathbf{a} + e
\end{tikzcd}$$

LWE, Informally

Roughly, the **Learning With Errors** problem is to recover $\mathbf{s}$ by sampling preimage-image pairs in the presence of some small "noise"

Learning With Errors [Reg05]

Let $\mathbf{s} \in \mathbb{F}_p^n$. The picture is as follows:

$$\begin{tikzcd}
\mathbb{F}_p^n \arrow[d, "\mathbf{s}\cdot_"'] & \ni & \mathbf{a} \arrow[d, mapsto, "\approx \mathbf{s}\cdot\mathbf{a}"] & & \\
\mathbb{F}_p & \ni & b & = & \mathbf{s}\cdot\mathbf{a} + e
\end{tikzcd}$$

LWE, Informally

Roughly, the **Learning With Errors** problem is to recover $\mathbf{s}$ by sampling preimage-image pairs in the presence of some small "noise"

More precisely, let

- $\mathbf{s} \in \mathbb{F}_p^n$
- Ψ be a discrete gaussian distribution over $\mathbb{F}_p$ centered at 0
- Define a distribution $\mathbf{A}_{\mathbf{s},\Psi}$ on $\mathbb{F}_p^n \times \mathbb{F}_p$ whose samples are pairs $(\mathbf{a}, b)$ where $\mathbf{a} \xleftarrow{\$} \mathbb{F}_p^n$, $b = \mathbf{s} \cdot \mathbf{a} + e$, $e \xleftarrow{\$} \Psi$

Definition (LWE Search)

The Learning With Errors problem is to recover $\mathbf{s}$ by sampling the distribution $\mathbf{A}_{\mathbf{s},\Psi}$.

Definition (LWE Decision)

Distinguish the distribution $\mathbf{A}_{\mathbf{s},\Psi}$ from the uniform distribution $\mathbf{U}(\mathbb{F}_p^n \times \mathbb{F}_p)$.

More precisely, let

- $\mathbf{s} \in \mathbb{F}_p^n$
- Ψ be a discrete gaussian distribution over $\mathbb{F}_p$ centered at 0
- Define a distribution $\mathbf{A}_{\mathbf{s},\Psi}$ on $\mathbb{F}_p^n \times \mathbb{F}_p$ whose samples are pairs $(\mathbf{a}, b)$ where $\mathbf{a} \xleftarrow{\$} \mathbb{F}_p^n$, $b = \mathbf{s} \cdot \mathbf{a} + e$, $e \xleftarrow{\$} \Psi$

Definition (LWE Search)

The Learning With Errors problem is to recover $\mathbf{s}$ by sampling the distribution $\mathbf{A}_{\mathbf{s},\Psi}$.

Definition (LWE Decision)

Distinguish the distribution $\mathbf{A}_{\mathbf{s},\Psi}$ from the uniform distribution $\mathbf{U}(\mathbb{F}_p^n \times \mathbb{F}_p)$.

Learning With Errors [Reg05]

More precisely, let

- $\mathbf{s} \in \mathbb{F}_p^n$
- Ψ be a discrete gaussian distribution over $\mathbb{F}_p$ centered at 0
- Define a distribution $\mathbf{A}_{\mathbf{s},\Psi}$ on $\mathbb{F}_p^n \times \mathbb{F}_p$ whose samples are pairs $(\mathbf{a}, b)$ where $\mathbf{a} \xleftarrow{\$} \mathbb{F}_p^n$, $b = \mathbf{s} \cdot \mathbf{a} + e$, $e \xleftarrow{\$} \Psi$

Definition (LWE Search)

The Learning With Errors problem is to recover $\mathbf{s}$ by sampling the distribution $\mathbf{A}_{\mathbf{s},\Psi}$.

Definition (LWE Decision)

Distinguish the distribution $\mathbf{A}_{\mathbf{s},\Psi}$ from the uniform distribution $\mathbf{U}(\mathbb{F}_p^n \times \mathbb{F}_p)$.

Learning With Errors [Reg05]

More precisely, let

- $\mathbf{s} \in \mathbb{F}_p^n$
- Ψ be a discrete gaussian distribution over $\mathbb{F}_p$ centered at 0
- Define a distribution $\mathbf{A}_{\mathbf{s},\Psi}$ on $\mathbb{F}_p^n \times \mathbb{F}_p$ whose samples are pairs $(\mathbf{a}, b)$ where $\mathbf{a} \xleftarrow{\$} \mathbb{F}_p^n$, $b = \mathbf{s} \cdot \mathbf{a} + e$, $e \xleftarrow{\$} \Psi$

Definition (LWE Search)

The Learning With Errors problem is to recover $\mathbf{s}$ by sampling the distribution $\mathbf{A}_{\mathbf{s},\Psi}$.

Definition (LWE Decision)

Distinguish the distribution $\mathbf{A}_{\mathbf{s},\Psi}$ from the uniform distribution $\mathbf{U}(\mathbb{F}_p^n \times \mathbb{F}_p)$.

Learning With Errors [Reg05]

More precisely, let

- $\mathbf{s} \in \mathbb{F}_p^n$
- Ψ be a discrete gaussian distribution over $\mathbb{F}_p$ centered at 0
- Define a distribution $\mathbf{A}_{\mathbf{s},\Psi}$ on $\mathbb{F}_p^n \times \mathbb{F}_p$ whose samples are pairs $(\mathbf{a}, b)$ where $\mathbf{a} \xleftarrow{\$} \mathbb{F}_p^n$, $b = \mathbf{s} \cdot \mathbf{a} + e$, $e \xleftarrow{\$} \Psi$

Definition (LWE Search)

The **Learning With Errors** problem is to recover $\mathbf{s}$ by sampling the distribution $\mathbf{A}_{\mathbf{s},\Psi}$.

Definition (LWE Decision)

Distinguish the distribution $\mathbf{A}_{\mathbf{s},\Psi}$ from the uniform distribution $\mathbf{U}(\mathbb{F}_p^n \times \mathbb{F}_p)$.

Learning With Errors [Reg05]

More precisely, let

- $\mathbf{s} \in \mathbb{F}_p^n$
- Ψ be a discrete gaussian distribution over $\mathbb{F}_p$ centered at 0
- Define a distribution $\mathbf{A}_{\mathbf{s},\Psi}$ on $\mathbb{F}_p^n \times \mathbb{F}_p$ whose samples are pairs $(\mathbf{a}, b)$ where $\mathbf{a} \xleftarrow{\$} \mathbb{F}_p^n$, $b = \mathbf{s} \cdot \mathbf{a} + e$, $e \xleftarrow{\$} \Psi$

Definition (LWE Search)

The **Learning With Errors** problem is to recover $\mathbf{s}$ by sampling the distribution $\mathbf{A}_{\mathbf{s},\Psi}$.

Definition (LWE Decision)

Distinguish the distribution $\mathbf{A}_{\mathbf{s},\Psi}$ from the uniform distribution $\mathbf{U}(\mathbb{F}_p^n \times \mathbb{F}_p)$.

Outline

Learning Homomorphisms With Errors

Observation

LWE's formulation was mainly algebraic:

- Expressed in terms of homomorphisms
- Complexity reductions (worst case to average case, search to decision) also algebraic

This motivates the following

Question

Can similar learning problems yield viable intractability assumptions based on group theory?

Learning Homomorphisms With Errors

Observation

LWE's formulation was mainly algebraic:

- Expressed in terms of homomorphisms
- Complexity reductions (worst case to average case, search to decision) also algebraic

This motivates the following

Question

Can similar learning problems yield viable intractability assumptions based on group theory?

Learning Homomorphisms With Errors

Observation

LWE's formulation was mainly algebraic:

- Expressed in terms of homomorphisms
- Complexity reductions (worst case to average case, search to decision) also algebraic

This motivates the following

Question

Can similar learning problems yield viable intractability assumptions based on group theory?

Learning Homomorphisms With Errors

Observation

LWE's formulation was mainly algebraic:

- Expressed in terms of homomorphisms
- Complexity reductions (worst case to average case, search to decision) also algebraic

This motivates the following

Question

Can similar learning problems yield viable intractability assumptions based on group theory?

Learning Homomorphisms With Errors

Observation

LWE's formulation was mainly algebraic:

- Expressed in terms of homomorphisms
- Complexity reductions (worst case to average case, search to decision) also algebraic

This motivates the following

Question

Can similar learning problems yield viable intractability assumptions based on group theory?

LWE Over Groups [BFNSS11]

Vector Spaces

$$\begin{array}{ccc} \mathbb{F}_p^n & \ni & \mathbf{a} \\ \Big\downarrow {\scriptstyle \mathbf{s}\cdot\_} & & \Big\downarrow {\scriptstyle \approx\, \mathbf{s}\cdot\mathbf{a}} \\ \mathbb{F}_p & \ni & b \\ & & \| \\ & & \mathbf{s}\cdot\mathbf{a}+e \end{array}$$

Groups

$$\begin{array}{ccc} G_n & \ni & a \\ \Big\downarrow {\scriptstyle \varphi} & & \Big\downarrow {\scriptstyle \approx\, \varphi(a)} \\ P_n & \ni & b \\ & & \| \\ & & \varphi(a)e \end{array}$$

Learning Homomorphisms from Images with Errors

Setup

- Let G_n and P_n be groups
- Set Γ_n, Ψ_n, distributions on G_n, P_n, resp.
- Let Φ_n be a distribution on the set of all homomorphisms, $\hom(G_n, P_n)$

The Distribution $\mathbf{A}_{\varphi,\Psi_n}$

For $\varphi \xleftarrow{\$} \Phi_n$, define the analogous distribution $\mathbf{A}_{\varphi,\Psi_n}$ on $G_n \times P_n$ whose samples are (a, b) where

- $a \xleftarrow{\$} \Gamma_n$;
- $e \xleftarrow{\$} \Psi_n$;
- $b = \varphi(a)e$

Learning Homomorphisms from Images with Errors

Setup

- Let G_n and P_n be groups
- Set Γ_n, Ψ_n, distributions on G_n, P_n, resp.
- Let Φ_n be a distribution on the set of all homomorphisms, $\hom(G_n, P_n)$

The Distribution $\mathbf{A}_{\varphi,\Psi_n}$

For $\varphi \xleftarrow{\$} \Phi_n$, define the analogous distribution $\mathbf{A}_{\varphi,\Psi_n}$ on $G_n \times P_n$ whose samples are (a, b) where

$$a \xleftarrow{\$} \Gamma_n;$$
$$e \xleftarrow{\$} \Psi_n;$$
$$b = \varphi(a)e$$

Learning Homomorphisms from Images with Errors

Setup

- Let G_n and P_n be groups
- Set Γ_n, Ψ_n, distributions on G_n, P_n, resp.
- Let Φ_n be a distribution on the set of all homomorphisms, $\hom(G_n, P_n)$

The Distribution $\mathbf{A}_{\varphi,\Psi_n}$

For $\varphi \xleftarrow{\$} \Phi_n$, define the analogous distribution $\mathbf{A}_{\varphi,\Psi_n}$ on $G_n \times P_n$ whose samples are (a, b) where

Learning Homomorphisms from Images with Errors

Setup

- Let G_n and P_n be groups
- Set Γ_n, Ψ_n, distributions on G_n, P_n, resp.
- Let Φ_n be a distribution on the set of all homomorphisms, $\hom(G_n, P_n)$

The Distribution $\mathbf{A}_{\varphi,\Psi_n}$

For $\varphi \xleftarrow{\$} \Phi_n$, define the analogous distribution $\mathbf{A}_{\varphi,\Psi_n}$ on $G_n \times P_n$ whose samples are (a, b) where

- $a \xleftarrow{\$} \Gamma_n$;
- $e \xleftarrow{\$} \Psi_n$;
- $b \leftarrow \varphi(a)e$

Learning Homomorphisms from Images with Errors

Setup

- Let G_n and P_n be groups
- Set Γ_n, Ψ_n, distributions on G_n, P_n, resp.
- Let Φ_n be a distribution on the set of all homomorphisms, $\hom(G_n, P_n)$

The Distribution $\mathbf{A}_{\varphi,\Psi_n}$

For $\varphi \stackrel{\$}{\leftarrow} \Phi_n$, define the analogous distribution $\mathbf{A}_{\varphi,\Psi_n}$ on $G_n \times P_n$ whose samples are (a, b) where

- $a \stackrel{\$}{\leftarrow} \Gamma_n$;
- $e \stackrel{\$}{\leftarrow} \Psi_n$;
- $b = \varphi(a)e$

Learning Homomorphisms from Images with Errors

Setup

- Let G_n and P_n be groups
- Set Γ_n, Ψ_n, distributions on G_n, P_n, resp.
- Let Φ_n be a distribution on the set of all homomorphisms, $\hom(G_n, P_n)$

The Distribution $\mathbf{A}_{\varphi,\Psi_n}$

For $\varphi \xleftarrow{\$} \Phi_n$, define the analogous distribution $\mathbf{A}_{\varphi,\Psi_n}$ on $G_n \times P_n$ whose samples are (a, b) where

- $a \xleftarrow{\$} \Gamma_n$;
- $e \xleftarrow{\$} \Psi_n$;
- $b = \varphi(a)e$

Learning Homomorphisms from Images with Errors

Setup

- Let G_n and P_n be groups
- Set Γ_n, Ψ_n, distributions on G_n, P_n, resp.
- Let Φ_n be a distribution on the set of all homomorphisms, $\hom(G_n, P_n)$

The Distribution $\mathbf{A}_{\varphi,\Psi_n}$

For $\varphi \xleftarrow{\$} \Phi_n$, define the analogous distribution $\mathbf{A}_{\varphi,\Psi_n}$ on $G_n \times P_n$ whose samples are (a, b) where

- $a \xleftarrow{\$} \Gamma_n$;
- $e \xleftarrow{\$} \Psi_n$;
- $b = \varphi(a)e$

Learning Homomorphisms from Images with Errors

Setup

- Let G_n and P_n be groups
- Set Γ_n, Ψ_n, distributions on G_n, P_n, resp.
- Let Φ_n be a distribution on the set of all homomorphisms, $\hom(G_n, P_n)$

The Distribution $\mathbf{A}_{\varphi,\Psi_n}$

For $\varphi \xleftarrow{\$} \Phi_n$, define the analogous distribution $\mathbf{A}_{\varphi,\Psi_n}$ on $G_n \times P_n$ whose samples are (a, b) where

- $a \xleftarrow{\$} \Gamma_n$;
- $e \xleftarrow{\$} \Psi_n$;
- $b = \varphi(a)e$

Learning Homomorphisms from Images with Errors

Search Problem

Given $\mathbf{A}_{\varphi,\Psi_n}$, recover φ.

Decision Problem

Given samples from an unknown distribution $\mathbf{R} \in \{\mathbf{A}_{\varphi,\Psi_n}, \mathbf{U}(G_n \times P_n)\}$, determine $\mathbf{R}$.

Hardness Assumption (Decision Version)

$$\mathbf{A}_{\varphi,\Psi_n} \underset{\text{PPT}}{\approx} \mathbf{U}(G_n \times P_n)$$

Learning Homomorphisms from Images with Errors

Search Problem

Given $\mathbf{A}_{\varphi,\Psi_n}$, recover φ.

Decision Problem

Given samples from an unknown distribution $\mathbf{R} \in \{\mathbf{A}_{\varphi,\Psi_n}, \mathbf{U}(G_n \times P_n)\}$, determine $\mathbf{R}$.

Hardness Assumption (Decision Version)

$$\mathbf{A}_{\varphi,\Psi_n} \underset{\mathrm{PPT}}{\approx} \mathbf{U}(G_n \times P_n)$$

Learning Homomorphisms from Images with Errors

Search Problem

Given $\mathbf{A}_{\varphi,\Psi_n}$, recover φ.

Decision Problem

Given samples from an unknown distribution $\mathbf{R} \in \{\mathbf{A}_{\varphi,\Psi_n}, \mathbf{U}(G_n \times P_n)\}$, determine $\mathbf{R}$.

Hardness Assumption (Decision Version)

$$\mathbf{A}_{\varphi,\Psi_n} \underset{\text{PPT}}{\approx} \mathbf{U}(G_n \times P_n)$$

Outline

Instantiation: Free Burnside Groups

Question

For which groups (if any) does the abstract problem make sense?

- The authors of [BFNS11] suggested the use of free Burnside groups.
- We'll review some of the intuition for this choice, as well as some of the key facts about these groups below.

Instantiation: Free Burnside Groups

Question

For which groups (if any) does the abstract problem make sense?

- The authors of [BFNS11] suggested the use of free Burnside groups.
- We'll review some of the intuition for this choice, as well as some of the key facts about these groups below.

Instantiation:
Free Burnside Groups

Question

For which groups (if any) does the abstract problem make sense?

- The authors of [BFNS11] suggested the use of free Burnside groups.
- We'll review some of the intuition for this choice, as well as some of the key facts about these groups below.

Varieties of Groups

The free Burnside groups can be thought of as living in a certain variety of groups.

Variety of Groups (Informal)

Roughly speaking, a **variety** is the class of all groups whose elements satisfy a certain set of equations.

Example

Abelian groups can be seen as the variety corresponding to the equation

$$XY = YX.$$

The Burnside groups live in the variety defined by the equation $X^m = 1$.

Varieties of Groups

The free Burnside groups can be thought of as living in a certain variety of groups.

Variety of Groups (Informal)

Roughly speaking, a **variety** is the class of all groups whose elements satisfy a certain set of equations.

Example

Abelian groups can be seen as the variety corresponding to the equation

$$XY = YX.$$

The Burnside groups live in the variety defined by the equation $X^m = 1$.

Varieties of Groups

The free Burnside groups can be thought of as living in a certain variety of groups.

Variety of Groups (Informal)

Roughly speaking, a **variety** is the class of all groups whose elements satisfy a certain set of equations.

Example

Abelian groups can be seen as the variety corresponding to the equation

$$XY = YX.$$

The Burnside groups live in the variety defined by the equation $X^m = 1$.

Varieties of Groups

The free Burnside groups can be thought of as living in a certain variety of groups.

Variety of Groups (Informal)

Roughly speaking, a **variety** is the class of all groups whose elements satisfy a certain set of equations.

Example

Abelian groups can be seen as the variety corresponding to the equation

$$XY = YX.$$

The Burnside groups live in the variety defined by the equation $X^m = 1$.

Varieties of Groups

Via the usual “abstract nonsense”, it is easy to see that varieties of groups contain free objects—just take a free group and factor out the normal subgroup resulting from all the “equations”...

$$\begin{tikzcd} S \arrow[r, "i"] \arrow[dr, "f"'] & UF(S) \arrow[d, dashed, "Uf'"] \\ & UG \end{tikzcd} \qquad \begin{tikzcd} F(S) \arrow[r, "p"] \arrow[d, dashed, "f'"] & F(S)/E \arrow[dl, dashed, "f''"] \\ G & \end{tikzcd}$$

Sets **Groups**

Varieties of Groups

Via the usual “abstract nonsense”, it is easy to see that varieties of groups contain free objects—just take a free group and factor out the normal subgroup resulting from all the “equations”...

$$\begin{tikzcd} S \arrow[rr, "i"] \arrow[drr, "f"'] & & UF(S) \arrow[d, dashed, "Uf'"] \\ & & UG \end{tikzcd} \qquad \begin{tikzcd} F(S) \arrow[rr, "p"] \arrow[d, dashed, "f'"] & & F(S)/E \arrow[dll, dashed, "f''"] \\ G & & \end{tikzcd}$$

Sets **Groups**

Varieties of Groups

Question

Which varieties of groups contain **finite free objects**???

If the equations are say,

$$\begin{aligned}[X, Y] &= 1 \\ X^p &= 1\end{aligned}$$

then the free objects are exactly $\mathbb{Z}_p^n$, which are the objects of study in LWE (if p is prime).

Question

What happens if the $[X, Y] = 1$ equation is removed?[a] In general, the answer is not so simple...

[a] Note: $[X, Y] = X^{-1}Y^{-1}XY$.

Varieties of Groups

Question

Which varieties of groups contain **finite free objects**???

If the equations are say,

$$\begin{aligned} [X, Y] &= 1 \\ X^p &= 1 \end{aligned}$$

then the free objects are exactly $\mathbb{Z}_p^n$, which are the objects of study in LWE (if p is prime).

Question

What happens if the $[X, Y] = 1$ equation is removed?[a] In general, the answer is not so simple...

[a] Note: $[X, Y] = X^{-1}Y^{-1}XY$.

Varieties of Groups

Question

Which varieties of groups contain **finite free objects**???

If the equations are say,

$$\begin{aligned} [X, Y] &= 1 \\ X^p &= 1 \end{aligned}$$

then the free objects are exactly $\mathbb{Z}_p^n$, which are the objects of study in LWE (if p is prime).

Question

What happens if the $[X, Y] = 1$ equation is removed?[a] In general, the answer is not so simple...

[a]Note: $[X, Y] = X^{-1}Y^{-1}XY$.

Burnside Groups

Notation

For the variety of groups defined by the equation $X^m = 1$, denote the free group on n generators in this variety by $B(n, m)$.

Determining the finiteness of $B(n, m)$ is known as the **Bounded Burnside Problem**.

Burnside Groups

Notation

For the variety of groups defined by the equation $X^m = 1$, denote the free group on n generators in this variety by $B(n, m)$.

Determining the finiteness of $B(n, m)$ is known as the **Bounded Burnside Problem**.

Bounded Burnside Problem

For $n > 1$ and for sufficiently large m, it is known that $|B(n,m)| = \infty$, yet for small m, our understanding is far from complete:

$B(n,2)$	Finite (also abelian)
$B(n,3)$	Finite
$B(n,4)$	Finite
$B(n,5)$	**Unknown**
$B(n,6)$	Finite
$B(n,7)$	**Unknown**
$\vdots$	$\vdots$

B_n-LHN **Problem**

- The authors of [BFNSS11] chose to use $B(n, 3)$ to instantiate the abstract LHN problem.
  - It's finite
  - It's the smallest non-abelian case
  - The structure of $B(3, n)$ is fairly well understood
- From here out, we'll denote $B(3, n)$ by B_n for brevity.

B_n-LHN **Problem**

- The authors of [BFNSS11] chose to use $B(n, 3)$ to instantiate the abstract LHN problem.
  - It's finite
  - It's the smallest non-abelian case
  - The structure of $B(3, n)$ is fairly well understood
- From here out, we'll denote $B(3, n)$ by B_n for brevity.

B_n-LHN **Problem**

- The authors of [BFNSS11] chose to use $B(n, 3)$ to instantiate the abstract LHN problem.
  - It's finite
  - It's the smallest non-abelian case
  - The structure of $B(3, n)$ is fairly well understood
- From here out, we'll denote $B(3, n)$ by B_n for brevity.

- The authors of [BFNSS11] chose to use $B(n, 3)$ to instantiate the abstract LHN problem.
  - It's finite
  - It's the smallest non-abelian case
  - The structure of $B(3, n)$ is fairly well understood
- From here out, we'll denote $B(3, n)$ by B_n for brevity.

B_n-LHN **Problem**

- The authors of [BFNSS11] chose to use $B(n, 3)$ to instantiate the abstract LHN problem.
  - It's finite
  - It's the smallest non-abelian case
  - The structure of $B(3, n)$ is fairly well understood
- From here out, we'll denote $B(3, n)$ by B_n for brevity.

The B_n-LHN **Problem**

This is simply the LHN problem, instantiated with free Burnside groups.

- The homomorphisms are sampled uniformly from $\hom(B_n, B_r)$.
- We'll ignore the error distribution for the moment, since those details are not important to the reduction.

The B_n-LHN Problem

This is simply the LHN problem, instantiated with free Burnside groups.

- The homomorphisms are sampled uniformly from $\hom(B_n, B_r)$.
- We'll ignore the error distribution for the moment, since those details are not important to the reduction.

B_n-LHN **Problem**

The B_n-LHN **Problem**

This is simply the LHN problem, instantiated with free Burnside groups.

- The homomorphisms are sampled uniformly from $\hom(B_n, B_r)$.
- We'll ignore the error distribution for the moment, since those details are not important to the reduction.

Outline

High Level / Intuition

We can break the argument into 3 easy steps:

1. Start with a simple observation for a partial randomizaiton.
2. Show this randomization is complete for a restricted version of the problem.
3. Show that the restricted version is statistically equivalent to the original problem.

   Hence the reduction applies to the original problem as well. Any efficient algorithm that solves the modified problem would solve the original- no efficient procedure can do anything substantially different on one versus the other.

High Level / Intuition

We can break the argument into 3 easy steps:

1. Start with a simple observation for a partial randomizaiton.
2. Show this randomization is complete for a restricted version of the problem.
3. Show that the restricted version is statistically equivalent to the original problem.
   - Hence the reduction applies to the original problem as well.
   - Any efficient algorithm that solves the modified problem would solve the original- no efficient procedure can do anything substantially different on one versus the other.

High Level / Intuition

We can break the argument into 3 easy steps:

1. Start with a simple observation for a partial randomizaiton.
2. Show this randomization is complete for a restricted version of the problem.
3. Show that the restricted version is statistically equivalent to the original problem.
   - Hence the reduction applies to the original problem as well
   - Any efficient algorithm that solves the modified problem would solve the original- no efficient procedure can do anything substantially different on one versus the other.

High Level / Intuition

We can break the argument into 3 easy steps:

1. Start with a simple observation for a partial randomizaiton.
2. Show this randomization is complete for a restricted version of the problem.
3. Show that the restricted version is statistically equivalent to the original problem.
    - Hence the reduction applies to the original problem as well
    - Any efficient algorithm that solves the modified problem would solve the original- no efficient procedure can do anything substantially different on one versus the other.

High Level / Intuition

We can break the argument into 3 easy steps:

1. Start with a simple observation for a partial randomizaiton.
2. Show this randomization is complete for a restricted version of the problem.
3. Show that the restricted version is statistically equivalent to the original problem.
   - Hence the reduction applies to the original problem as well
   - Any efficient algorithm that solves the modified problem would solve the original- no efficient procedure can do anything substantially different on one versus the other.

High Level / Intuition

We can break the argument into 3 easy steps:

1. Start with a simple observation for a partial randomizaiton.
2. Show this randomization is complete for a restricted version of the problem.
3. Show that the restricted version is statistically equivalent to the original problem.
    - Hence the reduction applies to the original problem as well
    - Any efficient algorithm that solves the modified problem would solve the original- no efficient procedure can do anything substantially different on one versus the other.

Outline

An Observation

Lemma

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 of* 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.

Observe that

$$\begin{aligned}(a' = \alpha(a), b) &= (\alpha(a), \varphi(a) \cdot e)\\ &= (\alpha(a), \varphi \circ \alpha^{-1}(\alpha(a)) \cdot e)\\ &= (a', \varphi \circ \alpha^{-1}(a') \cdot e).\end{aligned}$$

$\square$

An Observation

Lemma

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 of* 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.

Observe that

$$\begin{aligned}(a' = \alpha(a), b) &= (\alpha(a), \varphi(a) \cdot e) \\ &= (\alpha(a), \varphi \circ \alpha^{-1}(\alpha(a)) \cdot e) \\ &= (a', \varphi \circ \alpha^{-1}(a') \cdot e).\end{aligned}$$

∎

An Observation

- So, we can take instances from any $\mathbf{A}_{\varphi}^{\Psi_n}$ and transform them to instances from $\mathbf{A}_{\varphi\circ\alpha}^{\Psi_n}$ for some bijection α, giving at least a partial randomization.
- Next, we show that this randomization is complete for a subset of homomorphisms...

An Observation

- So, we can take instances from any $\mathbf{A}_{\varphi}^{\Psi_n}$ and transform them to instances from $\mathbf{A}_{\varphi\circ\alpha}^{\Psi_n}$ for some bijection α, giving at least a partial randomization.
- Next, we show that this randomization is complete for a subset of homomorphisms...

Outline

Completeness of the Randomization

Observation

Right-composition by an automorphism will not change the image of φ.

- Okay, so the technique from the lemma will not suffice to randomize all instances, but what about **surjective homomorphisms???**
- The following would be ideal:

Lemma

The action of $\mathrm{Aut}(B_n)$ *on* $\mathrm{Epi}(B_n, B_r)$ *is transitive.*

- This is true, but requires some work...
- Wait- what's this about "work", you say? I know... but still, $\frac{2}{3}$ easy steps isn't so bad :)

Completeness of the Randomization

Observation

Right-composition by an automorphism will not change the image of φ.

- Okay, so the technique from the lemma will not suffice to randomize all instances, but what about **surjective homomorphisms???**
- The following would be ideal:

Lemma

The action of $\mathrm{Aut}(B_n)$ *on* $\mathrm{Epi}(B_n, B_r)$ *is transitive.*

- This is true, but requires some work...
- Wait- what's this about "work", you say? I know... but still, $\frac{2}{3}$ easy steps isn't so bad :)

Completeness of the Randomization

Observation

Right-composition by an automorphism will not change the image of φ.

- Okay, so the technique from the lemma will not suffice to randomize all instances, but what about **surjective homomorphisms???**
- The following would be ideal:

Lemma

The action of $\mathrm{Aut}(B_n)$ *on* $\mathrm{Epi}(B_n, B_r)$ *is transitive.*

- This is true, but requires some work...
- Wait- what's this about "work", you say? I know... but still, $\frac{2}{3}$ easy steps isn't so bad :)

Completeness of the Randomization

Observation

Right-composition by an automorphism will not change the image of φ.

- Okay, so the technique from the lemma will not suffice to randomize all instances, but what about **surjective homomorphisms???**
- The following would be ideal:

Lemma

The action of $\mathrm{Aut}(B_n)$ *on* $\mathrm{Epi}(B_n, B_r)$ *is transitive.*

- This is true, but requires some work...
- Wait- what's this about "work", you say? I know... but still, $\frac{2}{3}$ easy steps isn't so bad :)

Completeness of the Randomization

Observation

Right-composition by an automorphism will not change the image of φ.

- Okay, so the technique from the lemma will not suffice to randomize all instances, but what about **surjective homomorphisms???**
- The following would be ideal:

Lemma

The action of $\mathrm{Aut}(B_n)$ *on* $\mathrm{Epi}(B_n, B_r)$ *is transitive.*

- This is true, but requires some work...
- Wait- what's this about "work", you say? I know... but still, $\frac{2}{3}$ easy steps isn't so bad :)

Completeness of the Randomization

Observation

Right-composition by an automorphism will not change the image of φ.

- Okay, so the technique from the lemma will not suffice to randomize all instances, but what about **surjective homomorphisms???**
- The following would be ideal:

Lemma

The action of $\mathrm{Aut}(B_n)$ *on* $\mathrm{Epi}(B_n, B_r)$ *is transitive.*

- This is true, but requires some work...
- Wait- what's this about “work”, you say? I know... but still, $\frac{2}{3}$ easy steps isn't so bad :)

Proving Transitivity

Consider the following commutative diagram, where ρ is the projection on to the commutator factor, taking $B_n \longrightarrow B_n/[B_n, B_n] \cong (\mathbb{F}_3^n, +)$:

$$\begin{tikzcd}
B_n \arrow[r, "\rho"] \arrow[d, "\varphi"'] & \mathbb{F}_3^n \arrow[d, "\overline{\varphi}"] \\
B_r \arrow[r, "\rho"] & \mathbb{F}_3^r
\end{tikzcd}$$

The main technical lemma used to prove transitivity is the following:

Lemma

Surjections from $B_n \longrightarrow B_r$ are precisely the maps whose abelianization is also surjective.

Proving Transitivity

Consider the following commutative diagram, where ρ is the projection on to the commutator factor, taking $B_n \longrightarrow B_n/[B_n, B_n] \cong (\mathbb{F}_3^n, +)$:

$$\begin{tikzcd} B_n \arrow[r, "\rho"] \arrow[d, "\varphi"'] & \mathbb{F}_3^n \arrow[d, "\overline{\varphi}"] \\ B_r \arrow[r, "\rho"] & \mathbb{F}_3^r \end{tikzcd}$$

The main technical lemma used to prove transitivity is the following:

Lemma

Surjections from $B_n \longrightarrow B_r$ are precisely the maps whose abelianization is also surjective.

Proving Transitivity

Consider the following commutative diagram, where ρ is the projection on to the commutator factor, taking $B_n \longrightarrow B_n/[B_n, B_n] \cong (\mathbb{F}_3^n, +)$:

$$\begin{tikzcd} B_n \arrow[r, "\rho"] \arrow[d, "\varphi"'] & \mathbb{F}_3^n \arrow[d, "\overline{\varphi}"] \\ B_r \arrow[r, "\rho"] & \mathbb{F}_3^r \end{tikzcd}$$

The main technical lemma used to prove transitivity is the following:

Lemma

Surjections from $B_n \longrightarrow B_r$ are precisely the maps whose abelianization is also surjective.

Proving the Lemma

- The proof is somewhat involved, and makes use of some specific details of the structure of free Burnside groups.
- However, some of the details can be abstracted away by a few invocations of the **Five Lemma**.

The Five Lemma

Consider the following commutative diagram, where the rows are exact.

$$
\begin{tikzcd}
A \arrow[r] \arrow[d, "e"'] & B \arrow[r] \arrow[d, "f"'] & C \arrow[r] \arrow[d, "g"'] & D \arrow[r] \arrow[d, "h"'] & E \arrow[d, "i"'] \\
A' \arrow[r] & B' \arrow[r] & C' \arrow[r] & D' \arrow[r] & E'
\end{tikzcd}
$$

Lemma (Five Lemma)

*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.*[a]

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

The Five Lemma

Consider the following commutative diagram, where the rows are exact.

$$
\begin{tikzcd}
A \arrow[r] \arrow[d, "e"'] & B \arrow[r] \arrow[d, "f"'] & C \arrow[r] \arrow[d, "g"'] & D \arrow[r] \arrow[d, "h"'] & E \arrow[d, "i"'] \\
A' \arrow[r] & B' \arrow[r] & C' \arrow[r] & D' \arrow[r] & E'
\end{tikzcd}
$$

Lemma (Five Lemma)

*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.*[a]

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

Proving the Lemma

We'll apply the lemma to the following diagram:

$$
\begin{tikzcd}
0 \arrow[r] & {[B_n, B_n]} \arrow[r, "i"] \arrow[d, "\hat{\varphi}"'] & B_n \arrow[r, "\rho"] \arrow[d, "\varphi"'] & \mathbb{F}_3^n \arrow[r] \arrow[d, "\overline{\varphi}"'] & 0 \\
0 \arrow[r] & {[B_r, B_r]} \arrow[r, "i"] & B_r \arrow[r, "\rho"] & \mathbb{F}_3^r \arrow[r] & 0
\end{tikzcd} \tag{1}
$$

- By the Five Lemma, proving $\hat{\varphi}$ is onto would suffice to prove our lemma, since then φ would be onto as well.
- Intuitively, dealing with the restriction to $[B_n, B_n]$ should be easier than the original map φ.[1]

[1] We actually invoke the five lemma yet again to show that $\hat{\varphi}$ is surjective...

Proving the Lemma

We'll apply the lemma to the following diagram:

$$\begin{tikzcd}
0 \arrow[r] & {[B_n, B_n]} \arrow[r, "i"] \arrow[d, "\hat{\varphi}"'] & B_n \arrow[r, "\rho"] \arrow[d, "\varphi"'] & \mathbb{F}_3^n \arrow[r] \arrow[d, "\overline{\varphi}"'] & 0 \\
0 \arrow[r] & {[B_r, B_r]} \arrow[r, "i"] & B_r \arrow[r, "\rho"] & \mathbb{F}_3^r \arrow[r] & 0
\end{tikzcd} \tag{1}$$

- By the Five Lemma, proving $\hat{\varphi}$ is onto would suffice to prove our lemma, since then φ would be onto as well.
- Intuitively, dealing with the restriction to $[B_n, B_n]$ should be easier than the original map φ.[1]

[1] We actually invoke the five lemma yet again to show that $\hat{\varphi}$ is surjective...

Proving the Lemma

We'll apply the lemma to the following diagram:

$$\begin{tikzcd}
0 \arrow[r] & {[B_n, B_n]} \arrow[r, "i"] \arrow[d, "\hat{\varphi}"'] & B_n \arrow[r, "\rho"] \arrow[d, "\varphi"'] & \mathbb{F}_3^n \arrow[r] \arrow[d, "\overline{\varphi}"'] & 0 \\
0 \arrow[r] & {[B_r, B_r]} \arrow[r, "i"] & B_r \arrow[r, "\rho"] & \mathbb{F}_3^r \arrow[r] & 0
\end{tikzcd} \tag{1}$$

- By the Five Lemma, proving $\hat{\varphi}$ is onto would suffice to prove our lemma, since then φ would be onto as well.
- Intuitively, dealing with the restriction to $[B_n, B_n]$ should be easier than the original map φ.[1]

[1] We actually invoke the five lemma yet again to show that $\hat{\varphi}$ is surjective...

Now Back to Transitivity...

We proceed in a straightforward manner:

Goal

Given an arbitrary epimorphism φ and a target epimorphism φ^* we want to find an automorphism α such that

$$\varphi^* = \varphi \circ \alpha.$$

Now Back to Transitivity...

We proceed in a straightforward manner:

Goal

Given an arbitrary epimorphism φ and a target epimorphism φ^* we want to find an automorphism α such that

$$\varphi^* = \varphi \circ \alpha.$$

Proving Transitivity

We'd like to find an automorphism α so that the following diagram commutes:

$$\begin{tikzcd}
 & & B_n \arrow[rr, "\varphi^*"] \arrow[dr, "\alpha"] \arrow[dd, "\rho"'] & & B_r \arrow[dr, "1_{B_r}"] \arrow[dd, "\rho"'] & & \\
0 \arrow[r] & K \arrow[rr] \arrow[dd, "\rho"'] & & B_n \arrow[rr, "\varphi"] \arrow[dd, "\rho"'] & & B_r \arrow[r] \arrow[dd, "\rho"'] & 0 \\
 & & \mathbb{F}_3^n \arrow[rr, "\overline{\varphi^*}"] \arrow[dr, "\overline{\alpha}"] & & \mathbb{F}_3^r \arrow[dr, "1_{\mathbb{F}_3^r}"] & & \\
0 \arrow[r] & \overline{K} \arrow[rr] & & \mathbb{F}_3^n \arrow[rr, "\overline{\varphi}"] & & \mathbb{F}_3^r \arrow[r] & 0
\end{tikzcd} \tag{2}$$

The Idea

Idea

- The idea is simple—after all, B_n is free!
- This allows us to define α to explicitly send basis elements where they need to go to make the composition work.

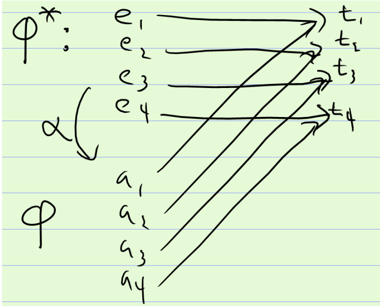

Proving Transitivity

- From the fact that B_n is free, we know that such an α exists.
- With the help of the previous lemma, we can show there is always a way to choose α to be bijective. ■

Proving Transitivity

- From the fact that B_n is free, we know that such an α exists.
- With the help of the previous lemma, we can show there is always a way to choose α to be bijective. $\blacksquare$

One More Lemma...

All that remains to show RSR for our restricted problem is to show the following

Lemma

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.

A simple counting argument (say, using the orbit-stabilizer theorem) suffices to show that each element $t \in S$ has the same number of preimages under the map from $G \longrightarrow S$ defined by $g \mapsto g \cdot s$. ∎

One More Lemma...

All that remains to show RSR for our restricted problem is to show the following

Lemma

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.

A simple counting argument (say, using the orbit-stabilizer theorem) suffices to show that each element $t \in S$ has the same number of preimages under the map from $G \longrightarrow S$ defined by $g \mapsto g \cdot s$. ∎

One More Lemma...

All that remains to show RSR for our restricted problem is to show the following

Lemma

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.

A simple counting argument (say, using the orbit-stabilizer theorem) suffices to show that each element $t \in S$ has the same number of preimages under the map from $G \longrightarrow S$ defined by $g \mapsto g \cdot s$. ∎

Outline

How Many Surjective Maps?

- Most homomorphisms $\varphi : B_n \longrightarrow B_r$ are surjective.
- In fact, if there is just a superlogarithmic gap between r and n then non-surjective maps comprise only a negligible fraction of the set of all homomorphisms.
- Even a crude estimate gives a 3^{r-n} fraction of all homomorphisms being non-surjective.

How Many Surjective Maps?

- Most homomorphisms $\varphi : B_n \longrightarrow B_r$ are surjective.
- In fact, if there is just a superlogarithmic gap between r and n then non-surjective maps comprise only a negligible fraction of the set of all homomorphisms.
- Even a crude estimate gives a 3^{r-n} fraction of all homomorphisms being non-surjective.

How Many Surjective Maps?

- Most homomorphisms $\varphi : B_n \longrightarrow B_r$ are surjective.
- In fact, if there is just a superlogarithmic gap between r and n then non-surjective maps comprise only a negligible fraction of the set of all homomorphisms.
- Even a crude estimate gives a 3^{r-n} fraction of all homomorphisms being non-surjective.

Observation

As a result, the altered distribution of instances (coming from sampling uniform surjective maps) is statistically close to the uniform distribution $\mathbf{U}(\text{hom}(B_n, B_r))$. In general,

Observation

For any $X_n \subset S_n$,

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

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 $\mathbf{U}(S_n)$.

Observation

As a result, the altered distribution of instances (coming from sampling uniform surjective maps) is statistically close to the uniform distribution $\mathbf{U}(\hom(B_n, B_r))$. In general,

Observation

For any $X_n \subset S_n$,

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

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 $\mathbf{U}(S_n)$.

Irrelevance of the Restriction

- The modified problem is no different than the original from a computational perspective
- Any efficient algorithm breaking the modified scheme could be used to break the original scheme (and vice versa).
- This proves the random self reducibility of the B_n-LHN problem.

Irrelevance of the Restriction

- The modified problem is no different than the original from a computational perspective
- Any efficient algorithm breaking the modified scheme could be used to break the original scheme (and vice versa).
- This proves the random self reducibility of the B_n-LHN problem.

Irrelevance of the Restriction

- The modified problem is no different than the original from a computational perspective
- Any efficient algorithm breaking the modified scheme could be used to break the original scheme (and vice versa).
- This proves the random self reducibility of the B_n-LHN problem.

Work in Progress / Open Questions

- Upper bounds on complexity of B_n-LHN?
- More complexity reductions: Search to decision?

Work in Progress / Open Questions

- Upper bounds on complexity of B_n-LHN?
- More complexity reductions: Search to decision?

Fin.

Questions?