Statistical Zap Arguments
from Quasi-Polynomial LWE

Abhishek Jain
Johns Hopkins University

Zhengzhong Jin
Johns Hopkins University

Abstract

We give the first construction of *statistical* Zaps. Our construction satisfies computational soundness and relies on the quasi-polynomial hardness of learning with errors assumption.

1 Introduction

The notion of zero-knowledge (ZK) proofs [GMR85] is fundamental to modern cryptography. Put simply, ZK proofs allow one to prove the validity of a statement while maintaining full privacy of the witness.

The feasibility of ZK proofs was first established in the interactive setting, where the prover and the verifier execute a two party protocol. Subsequently, a rich line of work aimed at minimizing interaction in ZK proofs was initiated. Goldreich and Oren [GO94] established the first barrier in this direction, proving the impossibility of ZK proofs in two rounds. This led to the study of relaxed notions of privacy for proof systems. One such notable notion is *witness indistinguishability* (WI), which guarantees that the proof does not reveal which one of multiple witnesses for the statement was used in the computation.

Dwork and Naor [DN00] proved that unlike ZK, WI can, in fact, be achieved in two rounds, without relying on a trusted setup. Specifically, they constructed two round WI with a public-coin verifier message, which they termed *Zaps*, from non-interactive zero-knowledge (NIZK) proofs in the common random string model [DMP88, FLS90]. By relying on known constructions of such NIZKs, this method can be used to obtain Zaps from quadratic residuosity [DMP88], trapdoor permutations [FLS90] and the decisional linear assumption over bilinear groups [GOS06]. More recently, Zaps were also constructed based on indistinguishability obfuscation [BP15]. Over the years, Zaps have found numerous applications in cryptography.

Statistical Zaps. All of the above constructions of Zaps only achieve computational WI property. In this work, we study the stronger notion of *statistical* Zaps that achieve the WI property against computationally unbounded verifiers. In other words, statistical Zaps achieve everlasting security.

Despite two decades of research, no constructions of statistical Zaps are currently known. This is in contrast to NIZK, which is indeed known with statistical privacy [CCH$^+$19, PS19] or even perfect privacy [GOS06]. One notable reason for this disparity is that the method of [DN00] for constructing Zaps is not applicable in the statistical case.

The recent work of Kalai, Khurana and Sahai [KKS18] comes close to achieving this goal. They constructed two round statistical WI with *private-coin* verifier message based on two round statistical sender-private oblivious transfer (OT) [NP01, Kal05, HK12, BD18]. The use of a private-coin verifier message is, in fact, instrumental to their approach (which builds on [JKKR17, BGI$^+$17]). As such, a different approach is required for constructing statistical Zaps with a *public-coin* verifier.

The public-coin verifier property of Zaps is crucial to many of its applications in cryptography. Indeed, public-coin property immediately implies *public verifiability*, a property which is often used in the design of round-efficient secure multiparty computation protocols (see, e.g., [HHPV18]). Moreover, it also allows for the verifier message to be *reusable* across multiple proofs, a property which is often used, for example, in the design of resettably-secure protocols (see, e.g., [DGS09]). We remark that neither public-verifiability nor reusability, even in isolation, was previously known for statistical Zaps. Moreover, even in the computational setting, these properties were only known to be achievable from a small set of assumptions.

1.1 Our Results

We give the first construction of statistical Zaps with computational soundness, a.k.a. *statistical Zap arguments*. Our construction requires two key ingredients: a statistical sender-private OT scheme where the receiver message is quasi-polynomially pseudorandom, and collision-intractable hash functions with quasi-polynomial pseduorandom fake key property [PS19, CCH+19]. Both of these primitives can be realized from the learning with errors (LWE) assumption with quasi-polynomial hardness.

Theorem 1.1. *Assuming quasi-polynomial LWE, there exists a statistical Zap argument system.*

In order to obtain our result, we depart significantly from previous approaches for constructing Zaps. Specifically, our approach combines the recent statistical NIZK arguments of Peikert and Shiehian [PS19] in a non-black-box manner with a two round public-coin statistically hiding extractable commitment scheme (Section 3.2). Previously, such a commitment scheme in the private-coin setting was constructed and used by [KKS18].

While we focus on the statistical setting, we note that our construction also yields the first computational Zap arguments based on quasi-polynomial LWE. Previously, computational Zaps were known based on quadratic residuosity, trapdoor permutations, decisional linear assumption over bilinear groups, and indistinguishability obfuscation.

2 Preliminaries

For any two (discrete) probability distributions P and Q, let $\mathsf{SD}(P,Q)$ denote *statistical distance* between P,Q. Let $\mathbb{Z}$ denote the set containing all integers. For any positive integer q, let $\mathbb{Z}_q$ denote the set $\mathbb{Z}/q\mathbb{Z}$. Let S be a discrete set, and let $\mathcal{U}(S)$ denote the uniform distribution over S. Throughout the paper, unless specified otherwise, we use λ to denote the security parameter.

2.1 Learing with Errors

We first recall the learing with errors (LWE) distribution.

Definition 2.1 (LWE distribution)**.** *For positive integer n and modulus q, and an error distribution χ over $\mathbb{Z}$, the LWE distribution $A_{\mathbf{s},\chi}$ is the following distribution. First sample a uniform random vector $\mathbf{a} \leftarrow \mathbb{Z}_q^n$, and an error $e \leftarrow \chi$, then output $(\mathbf{a}, \langle \mathbf{a}, \mathbf{s}\rangle + e) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$.*

Standard instantiations of LWE distribution usually choose χ to be discrete Gaussian distribution over $\mathbb{Z}$.

Definition 2.2 (Quasi-polynomial LWE Assumption)**.** *There exists a polynomial $n = n(\lambda)$ and a small real constant $c \in (0, 1/2)$ such that for any non-uniform probabilistic oracle adversary $\mathcal{D}^{(\cdot)}(\cdot)$ that runs in time $2^{O(\log^4 \lambda)}$, we have*

$$\mathsf{Adv}_\lambda(\mathcal{D}) = \left|\Pr\left[\mathcal{D}^{\mathcal{U}(\mathbb{Z}_q^n \times \mathbb{Z}_q)}(1^\lambda) = 1\right] - \Pr\left[\mathbf{s} \leftarrow \mathbb{Z}_q^n : \mathcal{D}^{A_{\mathbf{s},\chi}}(1^\lambda) = 1\right]\right| < c$$

Where the adversary is given oracle access to the uniform distribution $\mathcal{U}(\mathbb{Z}_q^n \times \mathbb{Z}_q)$ or the LWE distribution $A_{\mathbf{s},\chi}$.

In the following Lemma 2.3, we show that quasi-polynomial LWE assumption implies that any adversary running in a slower quasi-polynomial time can only have inverse quasi-polynomial advantage.

Lemma 2.3. *Assuming quasi-polynomial hardness of LWE, for any non-uniform probabilistic adversary $\mathcal{D}$ that runs in time $2^{O(\log^2 \lambda)}$, we have*

$$\mathsf{Adv}_\lambda(\mathcal{D}) = \left| \Pr\left[\mathcal{D}^{\mathcal{U}(\mathbb{Z}_q^n \times \mathbb{Z}_q)}(1^\lambda) = 1 \right] - \Pr\left[\mathbf{s} \leftarrow \mathbb{Z}_q^n : \mathcal{D}^{A_{\mathbf{s},\chi}}(1^\lambda) = 1 \right] \right| < 2^{-\Omega(\log^4 \lambda)}$$

Proof. We prove by contradiction. Suppose there exists an adversary $\mathcal{D}$ such that $\mathsf{Adv}_\lambda(\mathcal{D}) \geq 2^{-\log^4 \lambda}$ for infinitely many λ. Let $\epsilon = \mathsf{Adv}_\lambda(\mathcal{D})$. Then we construct the following adversary $\mathcal{D}'^{(\cdot)}(\cdot)$. The adversary $\mathcal{D}'$ is given access to an oracle $\mathcal{O}$, and is required to output a bit to tell if $\mathcal{O} = A_{\mathbf{s},\chi}$ or $\mathcal{O} = \mathcal{U}(\mathbb{Z}_q^n \times \mathbb{Z}_q)$. The strategy of $\mathcal{D}'$ is described as follows.

Let $N_\lambda = 2^{100 \log^4 \lambda}$.

1. Execute $\mathcal{D}$ for N_λ times. In i-th execution, $i \in [N_\lambda]$, sample an $\mathbf{s}_i \leftarrow \mathbb{Z}_q^n$. Execute $\mathcal{D}^{\mathcal{O}}(1^\lambda)$ with fresh randomness. For each oracle query made by $\mathcal{D}$, forward the query to oracle $\mathcal{O}$, and then obtain a response $(\mathbf{a}, b)$. Let $b' = b + \langle \mathbf{a}, \mathbf{s}_i \rangle \in \mathbb{Z}_q$.[1] Forward $(\mathbf{a}, b')$ to $\mathcal{D}$. Let $S_{\mathcal{O}}$ be the number of executions where $\mathcal{D}$ outputs 1.
2. Execute $\mathcal{D}^{\mathcal{U}(\mathbb{Z}_q^n \times \mathbb{Z}_q)}(1^\lambda)$ for N_λ times with fresh randomness for every execution. For each oracle query made by $\mathcal{D}$, sample an element uniform at random from $\mathbb{Z}_q^n \times \mathbb{Z}_q$, and forward the sample to $\mathcal{D}$. Let $S_{\mathcal{U}}$ be the number of executions where $\mathcal{D}$ outputs 1.
3. If $S_{\mathcal{O}} > S_{\mathcal{U}}$, output 1. If $S_{\mathcal{O}} < S_{\mathcal{U}}$, output 0. If $S_{\mathcal{O}} = S_{\mathcal{U}}$, output a random bit.

In the following, we assume $\Pr[\mathbf{s} \leftarrow \mathbb{Z}_q^n : \mathcal{D}^{A_{\mathbf{s},\chi}}(1^\lambda) = 1] = \Pr[\mathcal{D}^{\mathcal{U}(\mathbb{Z}_q^n \times \mathbb{Z}_q)}(1^\lambda) = 1] + \epsilon$. The proof for the other case follows in the same manner, and is omitted.

When $\mathcal{O} = \mathcal{U}(\mathbb{Z}_q^n \times \mathbb{Z}_q)$, $S_{\mathcal{O}}$ and $S_{\mathcal{U}}$ are subjected to two independent and identical distributions. Thus, $\mathcal{D}'^{\mathcal{U}(\mathbb{Z}_q^n \times \mathbb{Z}_q)}(1^\lambda)$ outputs a random bit. We have that $\Pr[\mathcal{D}'^{\mathcal{U}(\mathbb{Z}_q^n \times \mathbb{Z}_q)}(1^\lambda) = 1] = 1/2$.

When $\mathcal{O} = A_{\mathbf{s},\chi}$, denote $\mu_O = E[S_{\mathcal{O}}]$, $\mu_U = E[S_{\mathcal{U}}]$. Now we lower bound the probability

$$\begin{aligned} \Pr[\mathcal{D}'^{A_{\mathbf{s},\chi}}(1^\lambda) = 1] &= 1 - \Pr[\mathcal{D}'^{A_{\mathbf{s},\chi}}(1^\lambda) = 0] \geq 1 - \Pr[S_{\mathcal{O}} \leq S_{\mathcal{U}}] \\ &\geq 1 - \left(\Pr\left[S_{\mathcal{O}} \leq \frac{\mu_O + \mu_U}{2} \right] + \Pr\left[S_{\mathcal{U}} \geq \frac{\mu_O + \mu_U}{2} \right] \right) \end{aligned}$$

The first line comes from the fact that $\mathcal{D}'$ outputs 0 only when $S_{\mathcal{O}} < S_{\mathcal{U}}$ or $S_{\mathcal{O}} = S_{\mathcal{U}}$. The second line follows from a union bound, since $S_{\mathcal{O}} \leq S_{\mathcal{U}}$ implies $S_{\mathcal{O}} \leq \frac{\mu_P + \mu_U}{2}$ or $S_{\mathcal{U}} \geq \frac{\mu_P + \mu_U}{2}$.

From Chernoff bound, we have

$$\Pr\left[S_{\mathcal{O}} \leq \frac{\mu_O + \mu_U}{2} \right] \leq \exp\left(-\frac{1}{2} \left(\frac{\mu_O - \mu_U}{2\mu_O} \right)^2 \mu_O \right) \leq \exp\left(-\frac{1}{8} \epsilon^2 N \right)$$

$$\Pr\left[S_{\mathcal{U}} \geq \frac{\mu_O + \mu_U}{2} \right] \leq \exp\left(-\frac{\left(\frac{\mu_O - \mu_U}{2\mu_U} \right)^2}{2 + \frac{\mu_O - \mu_U}{2\mu_U}} \mu_U \right) \leq \exp\left(-\left(\frac{1}{2} \epsilon^2 + O(\epsilon^3) \right) N \right)$$

Hence, $\Pr[\mathcal{D}'^{A_{\mathbf{s},\chi}}(1^\lambda) = 1] \geq 1 - \exp(-\Omega(\epsilon^2 N))$. Thus, we have $\mathsf{Adv}_\lambda(\mathcal{D}') \geq 1/2 - \exp(-\Omega(\epsilon^2 N)) = 1/2 - \mathsf{neg}(\lambda)$. Note that $\mathcal{D}'$ runs in time $2^{O(\log^4 \lambda)}$. We reach a contradiction with quasi-polynomial LWE assumption. □

[1] Here, we use the worst-case to average-case reduction for LWE [Reg05].

2.2 Statistical Zap Arguments

Zaps [DN00] are two-round witness indistinguishable proof systems with a public-coin verifier message. Below, we define statistical Zap arguments, i.e., Zaps that achieve statistical WI property and computational soundness.

Let $\mathcal{P}$ denote the prover and $\mathcal{V}$ denote the verifier. We use $\mathsf{Trans}(\mathcal{P}(1^\lambda, x, \omega) \leftrightarrow \mathcal{V}(1^\lambda, x))$ to denote the transcript of an execution between $\mathcal{P}$ and $\mathcal{V}$, where $\mathcal{P}$ and $\mathcal{V}$ both have input a statement x and P also has a witness ω for x.

Definition 2.4. *Let L be a language in* NP. *We say that a two round protocol $\langle \mathcal{P}, \mathcal{V} \rangle$ with a public-coin verifier message is a statistical Zap argument for L if it satisfies the following properties:*

Completeness *For every $x \in L$, and witness ω for x, we have that*

$$\Pr\left[\mathsf{Trans}(\mathcal{P}(1^\lambda, x, \omega) \leftrightarrow \mathcal{V}(1^\lambda, x)) \textit{ is accepted by } \mathcal{V}\right] = 1$$

Computational Soundness *For any non-uniform probabilistic polynomial time (cheating) prover $\mathcal{P}^*$, there exists a negligible function $\nu(\cdot)$ such that for any $x \notin L$, we have that*

$$\Pr\left[\mathsf{Trans}(\mathcal{P}^*(1^\lambda, x) \leftrightarrow \mathcal{V}(1^\lambda, x)) \textit{ is accepted by } \mathcal{V}\right] < \nu(\lambda)$$

Statistical Witness Indistinguishability *For any (unbounded cheating) verifier $\mathcal{V}^*$, there exists a negligible function $\nu(\lambda)$ such that for every $x \in L$, and witnesses ω_1, ω_2 for x, we have that*

$$\mathsf{SD}\left(\mathsf{Trans}(\mathcal{P}(1^\lambda, x, \omega_1) \leftrightarrow \mathcal{V}^*(1^\lambda, x)), \mathsf{Trans}(\mathcal{P}(1^\lambda, x, \omega_2) \leftrightarrow \mathcal{V}^*(1^\lambda, x))\right) < \nu(\lambda)$$

3 Building Blocks

We need the following building blocks.

3.1 Statistical Sender Private Oblivious Transfer

Definition 3.1. *A statistical sender private oblivious transfer (OT) is a tuple of algorithms* $(\mathsf{OT}_1, \mathsf{OT}_2, \mathsf{OT}_3)$*:*

$\mathsf{OT}_1(1^\lambda, b)$**:** *On input security parameter λ, a bit $b \in \{0, 1\}$,* OT_1 *outputs the first round message* ot_1 *and a state* st.

$\mathsf{OT}_2(1^\lambda, \mathsf{ot}_1, m_0, m_1)$**:** *On input security parameter λ, a first round message* ot_1*, two bits $m_0, m_1 \in \{0, 1\}$,* OT_2 *outputs the second round message* ot_2.

$\mathsf{OT}_3(1^\lambda, \mathsf{ot}_2, \mathsf{st})$**:** *On input security parameter λ, the second round message* ot_2*, and the state generated by* OT_1, OT_3 *outputs a message m.*

We require the following properties:

Correctness *For any $b, m_0, m_1 \in \{0, 1\}$,*

$$\Pr\left[(\mathsf{ot}_1, \mathsf{st}) \leftarrow \mathsf{OT}_1(1^\lambda, b), \mathsf{ot}_2 \leftarrow \mathsf{OT}_2(1^\lambda, \mathsf{ot}_1, m_0, m_1), m \leftarrow \mathsf{OT}_3(1^\lambda, \mathsf{ot}_2, \mathsf{st}) : m = m_b\right] = 1$$

Statistical Sender Privacy *There exists a negligible function $\nu(\lambda)$ and an deterministic exponential time extractor* OTExt *such that for any (potential maliciously generated)* ot_1, $\mathsf{OTExt}(1^\lambda, \mathsf{ot}_1)$ *outputs a bit $b \in \{0,1\}$. Then for any $m_0, m_1 \in \{0,1\}$, we have*

$$\mathsf{SD}\left(\mathsf{OT}_2(1^\lambda, \mathsf{ot}_1, m_0, m_1), \mathsf{OT}_2(1^\lambda, \mathsf{ot}_1, m_b, m_b)\right) < \nu(\lambda)$$

Quasi-polynomial Pseudorandom Receiver's Message *For any $b \in \{0,1\}$, let* ot_1 *be the first round message generated by* $\mathsf{OT}_1(1^\lambda, b)$*. For any non-uniform probabilistic adversary $\mathcal{D}$ that runs in time $2^{O(\log^2 \lambda)}$, we have*

$$\mathsf{Adv}_\lambda(\mathcal{D}) = \left|\Pr\left[\mathcal{D}(1^\lambda, \mathsf{ot}_1) = 1\right] - \Pr\left[u \leftarrow \{0,1\}^{|\mathsf{ot}_1|} : \mathcal{D}(1^\lambda, u) = 1\right]\right| < 2^{-\Omega(\log^4 \lambda)}$$

Lemma 3.2. *Assuming quasi-polynomial hardness of LWE, there exists a statistical sender private oblivious transfer scheme.*

Proof. A statistical sender-private OT scheme was recently constructed by [BD18]. Their construction satisfies correctness and statistical sender privacy properties. Further, the receiver's message in their scheme is pseudorandom, assuming LWE. We observe that assuming quasi-polynomial LWE and using Lemma 2.3, their scheme also satisfies quasi-polynomially pseudorandom receiver's message property. □

3.2 Public Coin Statistical-Hiding Extractable Commitment Scheme

We now define a statistical-hiding extractable commitment scheme. The notion and its construction are adapted from [KKS18], with some slight modifications to fit in our application. The main difference between our definition and that of [KKS18] is that we require the receiver's first round message to be public coin as opposed to private-coin.

Our syntax departs from the classical definition of commitment schemes. We consider a tuple of four algorithms $(\mathsf{Com}_1, \mathsf{FakeCom}_1, \mathsf{Com}_2, \mathsf{Dec})$, where Com_1 corresponds to the honest receiver's algorithm that simply outputs a uniformly random string. Com_2 corresponds to the committer's algorithm that takes as input a message m as well as a random string $\mathbf{b}'$ of length μ and outputs a commitment string. We require two additional algorithms: (1) $\mathsf{FakeCom}_1$ that takes a binary string $\mathbf{b}$ of length μ as input and produces a first round message that "hides" the string $\mathbf{b}$, and (2) Dec that takes as input a transcript generated using $\mathsf{FakeCom}_1$ and Com_2 and outputs the committed message if the strings $\mathbf{b}$ and $\mathbf{b}'$ used for computing the transcript are equal.

Let $\mathcal{C}$, $\mathcal{R}$ denote the committer and the receiver, respectively. We now proceed to give a formal definition.

Definition 3.3. *A public coin statistical hiding extractable commitment is a tuple* $(\mathsf{Com}_1, \mathsf{FakeCom}_1, \mathsf{Com}_2, \mathsf{Dec})$. *The commit phase and open phase are defined as follows.*

Commitment Phase

Round 1 *On input parameters $(1^\lambda, 1^\mu)$, $\mathcal{R}$ executes* Com_1 *to sample a uniform random string* com_1*. $\mathcal{R}$ sends* com_1 *to $\mathcal{C}$.*

Round 2 *On input $(1^\lambda, m)$, $\mathcal{C}$ chooses $\mathbf{b}' \leftarrow \{0,1\}^\mu$ uniformly at random and computes* $\mathsf{com}_2 \leftarrow \mathsf{Com}_2(1^\lambda, 1^\mu, \mathsf{com}_1, \mathbf{b}', m; r)$ *with randomness r. $\mathcal{C}$ sends $(\mathbf{b}', \mathsf{com}_2)$ to $\mathcal{R}$.*

Opening Phase

$\mathcal{C}$ sends the message and the randomness (m, r) to $\mathcal{R}$. $\mathcal{R}$ checks if $\mathsf{com}_2 = \mathsf{Com}_2(1^\lambda, 1^\mu, \mathsf{com}_1, \mathbf{b}', m; r)$.

We require the following properties of the commitment scheme.

Statistical Hiding *There exists a negligible function $\nu(\cdot)$, a deterministic exponential time algorithm* ComExt, *and a randomized simulator* Sim, *such that for any fixed (potentially maliciously generated)* com_1, $\mathsf{ComExt}(1^\lambda, 1^\mu, \mathsf{com}_1)$ *outputs* $\mathbf{b} \in \{0,1\}^\mu$, *and for any* $\mathbf{b}' \neq \mathbf{b}$, *and* $m \in \{0,1\}$, *we have*

$$\mathsf{SD}\left(\mathsf{Com}_2(1^\lambda, 1^\mu, \mathsf{com}_1, \mathbf{b}', m), \mathsf{Sim}(1^\lambda, 1^\mu, \mathsf{com}_1)\right) < \mu \cdot \nu(\lambda) \tag{1}$$

Quasi-polynomial Pseudorandom Receiver's Message *For any* $\mathbf{b} \in \{0,1\}^\mu$, $\mathsf{FakeCom}_1(1^\lambda, 1^\mu, \mathbf{b})$ *and a uniform random string outputted by* $\mathsf{Com}(1^\lambda, 1^\mu)$ *are quasi-polynomially indistinguishable. Specifically, for any non-uniform adversary* $\mathcal{D}$ *that runs in time* $2^{O(\log^2 \lambda)}$, *we have*

$$\left|\Pr[\mathcal{D}(1^\lambda, 1^\mu, \mathsf{Com}_1(1^\lambda, 1^\mu)) = 1] - \Pr[\mathcal{D}(1^\lambda, 1^\mu, \mathsf{FakeCom}_1(1^\lambda, 1^\mu, \mathbf{b})) = 1]\right| \leq \mu \cdot 2^{-\Omega(\log^4 \lambda)}$$

Extractable $\mathsf{FakeCom}_1$ *and* Dec *satisfy the following property. For any* $\mathbf{b} \in \{0,1\}^\mu$, *we have*

$$\Pr\left[\substack{(\mathsf{com}_1, \mathsf{st}) \leftarrow \mathsf{FakeCom}_1(1^\lambda, 1^\mu, \mathbf{b}), \\ \mathsf{com}_2 \leftarrow \mathsf{Com}_2(1^\lambda, 1^\mu, \mathsf{com}_1, \mathbf{b}, m)} : \mathsf{Dec}(1^\lambda, 1^\mu, \mathsf{st}, \mathsf{com}_2) = m\right] = 1$$

Lemma 3.4. *Assuming quasi-polynomial hardness of LWE, there exists a public coin statistical-hiding extractable commitment scheme.*

We construct a public coin statistical hiding extractable commitment by slightly modifying the commitment scheme of [KKS18]. Their construction already satisfies extractability and statistical hiding properties. However, their construction, as originally described, is private coin.

We note that the receiver's message in their scheme simply consists of multiple receiver messages of a statistical sender-private OT scheme. Then, by instantiating their construction with an OT scheme that satisfies quasi-polynomial pseudorandom receiver's message property (see Section 3.1), their scheme can be easily adapted to obtain a *public coin* statistical hiding extractable commitment. Specifically, in the modified construction, the honest receiver's algorithm $\mathsf{Com}(1^\lambda, 1^\mu)$ simply computes a uniform random string, while $\mathsf{FakeCom}_1$ corresponds to the receiver algorithm in the construction of [KKS18].

Construction. For completeness, here we describe the full construction adapted from [KKS18].

$\mathsf{Com}_1(1^\lambda, 1^\mu)$ **:** Output a uniform random string $\mathsf{com}_1 \leftarrow \{0,1\}^{|\mathsf{com}_1|}$.

$\mathsf{FakeCom}_1(1^\lambda, 1^\mu, \mathbf{b})$ **:** Parse $\mathbf{b} = (b_1, b_2, \ldots, b_\mu)$. For each $i \in [\mu]$, execute $(\mathsf{ot}_{1,i}, \mathsf{st}_i) \leftarrow \mathsf{OT}_1(1^\lambda, b_i)$. Output $\mathsf{com}_1 = (\mathsf{ot}_{1,i})_{i \in [\mu]}$ and $\mathsf{st} = (\mathsf{st}_i)_{i \in [\mu]}$.

$\mathsf{Com}_2(1^\lambda, 1^\mu, \mathsf{com}_1, \mathbf{b}', m)$ **:** Parse $\mathbf{b}' = (b'_1, b'_2, \ldots, b'_\mu)$, and $\mathsf{com}_1 = (\mathsf{ot}_{1,i})_{i \in [\mu]}$. Sample uniform random $m_1, m_2, \ldots, m_\mu \in \{0,1\}$ such that $\bigoplus_{i \in [\mu]} m_i = m$. For each $i \in [\mu]$, let $m_{b'_i, i} = m_i$, and sample $m_{1-b'_i, i} \leftarrow \{0,1\}$. Execute $\mathsf{ot}_{2,i} \leftarrow \mathsf{OT}_2(1^\lambda, 1^\mu, m_{0,i}, m_{1,i})$. Output $\mathsf{com}_2 := (\mathsf{ot}_{2,i})_{i \in [\mu]}$.

$\mathsf{Dec}(1^\lambda, 1^\mu, \mathsf{st}, \mathsf{com}_2)$ **:** Parse $\mathsf{st} = (\mathsf{st}_i)_{i \in [\mu]}$, and $\mathsf{com}_2 = (\mathsf{ot}_{2,i})_{i \in [\mu]}$. For each $i \in [\mu]$, execute $m'_i \leftarrow \mathsf{OT}_3(1^\lambda, \mathsf{ot}_{2,i}, \mathsf{st}_i)$. Let $m' = \bigoplus_{i \in [\mu]} m'_i$. Output m'.

This completes the description of the scheme. We now argue each of the required properties.

Statistical Hiding We construct the following extracting algorithm $\mathsf{ComExt}(1^\lambda, 1^\mu, \mathsf{com}_1 = (\mathsf{ot}_{1,i})_{i \in [\mu]})$. For each $i \in [\mu]$, execute $b_i = \mathsf{OTExt}(1^\lambda, \mathsf{ot}_{1,i})$. Output $\mathbf{b} = (b_i)_{i \in [\mu]}$.

Let $\mathbf{b} = \mathsf{ComExt}(1^\lambda, 1^\mu, \mathsf{com}_1)$, then for any $\mathbf{b}' \neq \mathbf{b}$, consider the following hybrids.

$\mathsf{Com}_2(1^\lambda, 1^\mu, \mathsf{com}_1, \mathbf{b}', m)$ **:** Sample $(m_i)_{i\in[\mu]}$ uniformly at random such that $\bigoplus_{i\in[\mu]} m_i = m$. For each $i \in [\mu]$, set $m_{i,b'_i} = m_i$, and $m_{i,1-b'_i} \leftarrow \{0,1\}$. Output $(\mathsf{OT}_2(1^\lambda, \mathsf{ot}_{1,i}, m_{i,0}, m_{i,1}))_{i\in[\mu]}$.

$\mathsf{Hybrid}(1^\lambda, 1^\mu, \mathsf{com}_1, \mathbf{b}', m)$ **:** Sample $(m_i)_{i\in[\mu]}$ uniformly at random such that $\bigoplus_{i\in[\mu]} m_i = m$. For each $i \in [\mu]$, set $m_{i,b'_i} = m_i$, and $m_{i,1-b'_i} \leftarrow \{0,1\}$. Output $(\mathsf{OT}_2(1^\lambda, \mathsf{ot}_{1,i}, \underline{m_{i,b_i}, m_{i,b_i}}))_{i\in[\mu]}$.

$\mathsf{Sim}(1^\lambda, 1^\mu, \mathsf{com}_1)$ **:** $\underline{\text{Sample } m_1, m_2, \ldots, m_\mu \leftarrow \{0,1\}}$. Output $(\mathsf{OT}_2(1^\lambda, \mathsf{ot}_{1,i}, m_i, m_i))_{i\in[\mu]}$.

From the statistical-hiding property of underlying OT scheme, it follows that Com_2 and Hybrid are statistically close. Specifically, there exists a negligible function $\nu(\cdot)$ such that:

$$\mathsf{SD}\left(\mathsf{Com}_2(1^\lambda, 1^\mu, \mathsf{com}_1, \mathbf{b}', m), \mathsf{Hybrid}(1^\lambda, 1^\mu, \mathsf{com}_1, \mathbf{b}', m)\right) < \mu \cdot \nu(\lambda)$$

Next, we prove that Hybrid and Sim are identical distributions. Denote $\mathcal{I} = \{i^* \in [\mu] | b_{i^*} \neq b'_{i^*}\}$. Since $\mathbf{b} \neq \mathbf{b}'$, we have $\mathcal{I} \neq \phi$. Hence, the joint distribution $(m_{i,b_i})_{i\in[\mu]\setminus\mathcal{I}}$ is uniformly random. Since $b_{i^*} \neq b'_{i^*}$ for all $i^* \in \mathcal{I}$, $(m_{i^*,b_{i^*}})$ is sampled uniformly at random for all $i^* \in \mathcal{I}$. Hence, $(m_{i,b_i})_{i\in[\mu]}$ is uniformly random. Hence, $\mathsf{Hybrid}(1^\lambda, 1^\mu, \mathsf{com}_1, \mathbf{b}', m)$ and $\mathsf{Sim}(1^\lambda, 1^\mu, \mathsf{com}_1)$ are identical distributions.

The statistical hiding property of the construction now follows by combining the above claims.

Quasi-polynomial Pseudorandom Receiver's Message This property directly follows from the quasi-polynomial pseudorandom receiver message property the OT scheme.

Extractable This property directly follows from the correctness of the OT scheme.

3.3 Correlation Intractable Hash Function

We start by defining searchable relations. The following definition is taken verbatim from [PS19].

Definition 3.5 (Searchable Relation [PS19])**.** *We say that a relation $R \subseteq \mathcal{X} \times \mathcal{Y}$ is searchable in* size S *if there exists a function $f : \mathcal{X} \to \mathcal{Y}$ that is implementable as a* Boolean *circuit of size S, such that if $(x, y) \in R$ then $y = f(x)$.*

Correlation intractable hash function is a family of keyed hash functions satisfying following property: for any searchable relation R, it is hard for a computationally unbounded adversary to find an element x such that $(x, f(x)) \in R$.

Definition 3.6 (Correlation Intractable Hash Function, slightly modified from [PS19])**.** Correlation Intractable Hash Function *(CIH) is a triple of algorithms* $(\mathsf{KGen}, \mathsf{FakeGen}, \mathsf{H}_{(\cdot)}(\cdot))$*, with the following properties:*

Let $s = s(\lambda), \ell = \ell(\lambda), d = d(\lambda)$ be $\mathsf{poly}(\lambda)$*-bounded functions. Let $\{\mathcal{R}_{\lambda,s,\ell,d}\}_\lambda$ be a family of searchable relations, where each relation $R \in \mathcal{R}_{\lambda,s,\ell,d}$ is searchable by a circuit of size $s(\lambda)$, output length $\ell(\lambda)$ and depth $d(\lambda)$.*

Statistical Correlation Intractable *There exists a negligible function $\nu(\cdot)$ such that, for any relation $R \in \mathcal{R}_{\lambda,s,\ell,d}$, and circuit C_λ that searches for a witness for R, we have*

$$\Pr\left[k \leftarrow \mathsf{FakeGen}(1^\lambda, 1^{|C_\lambda|}, C_\lambda) : \exists x \text{ s.t. } (x, \mathsf{H}_k(x)) \in R\right] < \nu(\lambda)$$

Quasi-polynomial Pseudorandom Fake Key *For any circuit C_λ with size s, output length ℓ, and depth d,* $\mathsf{KGen}(1^\lambda, 1^{|C_\lambda|})$ *outputs an uniform random string. Furthermore, for any non-uniform adversary $\mathcal{D}$ that runs in time $2^{O(\log^2 \lambda)}$, we have*

$$\left|\Pr\left[\mathcal{D}(1^\lambda, 1^{|C_\lambda|}, \mathsf{KGen}(1^\lambda, 1^{|C_\lambda|})) = 1\right] - \Pr\left[\mathcal{D}(1^\lambda, 1^{|C_\lambda|}, \mathsf{FakeGen}(1^\lambda, 1^{|C_\lambda|}, C_\lambda)) = 1\right]\right| \leq 2^{-\Omega(\log^4 \lambda)}$$

Theorem 3.7. *Assuming quasi-polynomial hardness of LWE, there exists a construction of correlation intractable hash function with quasi-polynomial pseudorandom fake key.*

Proof. The construction of such a function is given in [PS19, CCH+19]. Specifically, we use the construction of [PS19], which satisfies *statistical correlation intractability*. Moreover, the $\mathsf{FakeGen}$ algorithm in their construction simply consists of some ciphertexts that are pseudorandom assuming LWE. Thus, if we assume quasi-polynomial hardness of LWE, their construction satisfies quasi-polynomial pseudorandom fake key property. □

For our application, We require a slightly stronger property than statistical correlation intractability as defined above. Specifically, we require that the distinguishing probability in statistical correlation intractability is $2^{-\lambda}$ for a special class of relations.

We show in Corollary 3.8 that by using parallel repetition, we can construct a CIH with the above property from any CIH.

Corollary 3.8 (Amplification of Statistical Correlation Intractability). [2] *There exists a correlation intractable hash function* $(\mathsf{KGen}, \mathsf{FakeGen}, \mathsf{H}_{(\cdot)}(\cdot))$ *such that the following additional property holds.*

$2^{-\lambda}$-Statistical Correlation Intractability *Let $\{C_\lambda\}_\lambda$ be a family of* Boolean *circuits, where C_λ has polynomial size $s(\lambda)$, polynomial depth $d(\lambda)$, and outputs a single bit. There exists a polynomial $\ell = \ell(\lambda)$ such that the following holds. Let $\overrightarrow{C_{\lambda,\ell}}$ be the circuit $\overrightarrow{C_\lambda}(c_1, c_2, \ldots, c_\ell) = (C_\lambda(c_1), C_\lambda(c_2), \ldots, C_\lambda(c_\ell))$, then for large enough λ,*

$$\Pr\left[k \leftarrow \mathsf{FakeGen}\left(1^\lambda, 1^{|\overrightarrow{C_{\lambda,\ell}}|}, \overrightarrow{C_{\lambda,\ell}}\right) : \exists x \text{ s.t. } \mathsf{H}_k(x) = \overrightarrow{C_{\lambda,\ell}}(x)\right] < 2^{-\lambda}$$

Proof. Let C_{in} be the length of input to C_λ. We prove this corollary from any CIH $(\mathsf{KGen}', \mathsf{FakeGen}', \mathsf{H}'_{(\cdot)}(\cdot))$, where H' is a hash function family $\{0,1\}^{C_{in}\cdot\ell'} \to \{0,1\}^{\ell'}$. Denote $R_{\overrightarrow{C_{\lambda,\ell'}}} = \left\{(x, \overrightarrow{C_{\lambda,\ell'}}(x))\right\}$. We construct the following new CIH.

Parameters Set $\ell(\lambda) = \ell'(\lambda) \cdot \lambda$.

$\mathsf{KGen}(1^\lambda, 1^{|\overrightarrow{C_{\lambda,\ell}}|}, \overrightarrow{C_{\lambda,\ell}})$: For each $i \in [\lambda]$, execute $k_i \leftarrow \mathsf{KGen}'(1^\lambda, 1^{|\overrightarrow{C_{\lambda,\ell'}}|}, \overrightarrow{C_{\lambda,\ell'}})$ with fresh randomness. Output $k = (k_i)_{i\in[\lambda]}$.

$\mathsf{FakeGen}(1^\lambda, 1^{|\overrightarrow{C_{\lambda,\ell}}|}, \overrightarrow{C_{\lambda,\ell}})$: For each $i \in [\lambda]$, execute $k_i \leftarrow \mathsf{FakeGen}'(1^\lambda, 1^{|\overrightarrow{C_{\lambda,\ell'}}|}, \overrightarrow{C_{\lambda,\ell'}})$ with fresh randomness. Output $k = (k_i)_{i\in[\lambda]}$.

$\mathsf{H}_k(c_1, c_2, \ldots, c_\ell)$: For each $i \in [\lambda]$, execute $\mathbf{b}_i = \mathsf{H}_k(c_{\ell'(i-1)+1}, c_{\ell'(i-1)+2}, \ldots, c_{\ell' i})$, output $\mathbf{b} = (\mathbf{b}_i)_{i\in[\lambda]}$.

[2] In fact, the CIH construction in [PS19] already satisfies this additional property. Here we give a generic transformation from any CIH

We now prove that the above construction satisfies $2^{-\lambda}$-statistical correlation intractability. For large enough λ, $\nu(\lambda) < 1/2$. Hence we have

$$\begin{aligned} &\Pr\left[k \leftarrow \mathsf{FakeGen}(1^\lambda, 1^{|\overrightarrow{C_{\lambda,\ell}}|}, \overrightarrow{C_{\lambda,\ell}}) : \exists x = (x_i)_{i\in[\ell]} \text{ s.t. } \mathsf{H}_k(x) = \overrightarrow{C_{\lambda,\ell}}(x)\right] \\ =&\Pr\left[\forall i \in [\lambda], k_i \leftarrow \mathsf{FakeGen}(1^\lambda, 1^{|\overrightarrow{C_{\lambda,\ell'}}|}, \overrightarrow{C_{\lambda,\ell'}}) : \exists x_i, x_i \in R_{\overrightarrow{C_{\lambda,\ell'}}}\right] \leq (\nu(\lambda))^\lambda < 2^{-\lambda} \end{aligned}$$

The second line follows from the fact that k_i are generated independently. □

4 Our Construction

In this section, we describe our construction of a statistical Zap argument system for Graph Hamiltonicity, which is an NP-Complete problem.

Notation. We describe some notation that is used in our construction. Let L_{HAM} denote the Graph Hamiltonicity language over graphs $G = (V, E)$ of n vertices, where V denotes the set of vertices and E denotes the set of edges in G. We slightly abuse notation and use G to denote its adjacency matrix $G = (G_i[s,t])_{s,t\in[n]}$.

Let $(\mathsf{Com}_1, \mathsf{FakeCom}_1, \mathsf{Com}_2, \mathsf{Dec})$ be a public coin statistically hiding extractable commitment scheme (Definition 3.3). We set the parameter μ of the commitment scheme as $\Theta(\log^2 \lambda)$. Let $(\mathsf{KGen}, \mathsf{FakeGen}, \mathsf{H}_{(\cdot)}(\cdot))$ be a family of CIH (Definition 3.6). We choose the polynomial $\ell = \ell(\lambda)$ in Corollary 3.8 such that the CIH is $2^{-\lambda}$-statistical correlation intractable.

Circuit C_{st}. Let C_{st} denote the following Boolean circuit.
Input: a $n \times n$ matrix $c = (c_{s,t})_{s,t\in[n]}$.
Output: a boolean value.

1. For any $s, t \in [n]$, execute $G[s,t] = \mathsf{Dec}(1^\lambda, 1^\mu, \mathsf{st}, c_{s,t})$.
2. If $G = (G_i[s,t])_{s,t\in[n]}$ is a cycle graph, then output 0. Otherwise output 1.

For ease of exposition, we extend the notation C_{st} to a series of matrices $(c_1, c_2, \ldots, c_\ell)$. Specifically, $C_{\mathsf{st}}(c_1, c_2, \ldots, c_\ell)$ is defined as $(C_{\mathsf{st}}(c_1), C_{\mathsf{st}}(c_2), \ldots, C_{\mathsf{st}}(c_\ell))$.

Construction. The verifier $\mathcal{V}$ and prover $\mathcal{P}$ are both given input the security parameter λ and a graph $G = (V, E)$ of n vertices. The prover is additionally given as input a witness ω for G.

Round 1 Verifier $\mathcal{V}$ computes and sends uniform random strings $(\mathsf{com}_1 \leftarrow \mathsf{Com}_1(1^\lambda, 1^\mu), k \leftarrow \mathsf{KGen}(1^\lambda, 1^{|C_{\mathsf{st}}|})$, where C_{st} takes ℓ separate $n \times n$ matrices as input, and outputs ℓ bits.

Round 2 Prover $\mathcal{P}$ does the following:

1. Choose a random $\mathbf{b}' \leftarrow \{0,1\}^\mu$.
2. Compute ℓ first round messages of Blum's sigma protocol for Graph Hamiltonicity. Specifically, for every $i \in [\ell]$, first sample a random cycle graph $G_i = (G_i[s,t])_{s,t\in[n]}$. Next, for each $s, t \in [n]$, compute $\mathsf{c}_i[s,t] \leftarrow \mathsf{Com}_2(1^\lambda, 1^\mu, \mathsf{com}_1, \mathbf{b}', G_i[s,t]; r_i^{(s,t)})$ using randomness $r_i^{(s,t)}$. Finally let $\mathsf{c}_i = (\mathsf{c}_i[s,t])_{s,t\in[n]}$.
3. Compute $(b_1, b_2, \ldots, b_\ell) = \mathsf{H}_k(\mathsf{c}_1, \ldots, \mathsf{c}_\ell)$.

4. For every $i \in [\ell]$, compute the answer to challenge b_i in Blum's sigma protocol. Specifically, if $b_i = 0$, then set $z_i = (G_i, (r_i^{(s,t)})_{s,t\in[n]})$. Else, if $b_i = 1$, then compute a one-to-one map $\phi : G \to G_i$ such that $\phi(w)$ is the cycle G_i, and set $z_i = (\phi, (r_i^{(s,t)})_{(s,t)=\phi(e),e\notin E})$.
5. Send $\Pi = (\mathbf{b}', (\mathsf{c}_i)_{i\in[\ell]}, (z_i)_{i\in[\ell]})$ to the verifier.

Verification Upon receiving the proof $\Pi = (\mathbf{b}', (\mathsf{c}_i)_{i\in[\ell]}, (z_i)_{i\in[\ell]})$, the verifier first computes $(b_1, b_2, \cdots, b_\ell) = \mathsf{H}_k(\mathsf{c}_1, \mathsf{c}_2, \ldots, \mathsf{c}_\ell)$, and then verifies each copy (c_i, b_i, z_i) of the proof as in Blum's protocol. Specifically, if $b_i = 0$, then parse $z_i = (G_i, (r_i^{(s,t)})_{s,t\in[n]})$ and check if $\mathsf{c}_i = (\mathsf{Com}_2(1^\lambda, 1^\mu, \mathsf{com}_1, \mathbf{b}', G_i[s,t]; r_i^{(s,t)})_{s,t\in[n]}$ and G_i is a cycle graph. Otherwise if $b_i = 1$, then parse $z_i = (\phi, (r_i^{(s,t)})_{(s,t)=\phi(e),e\notin E})$ and check if ϕ is a one-to-one map, and for each $e \notin E$, and $(s,t) = \phi(e)$, check if $\mathsf{c}_i[s,t] = \mathsf{Com}_2(1^\lambda, 1^\mu, \mathsf{com}_1, \mathbf{b}', 0; r_i^{(s,t)})$. If all of the checks succeed, then accept the proof, otherwise reject.

This completes the description of our construction. Below, we prove that our construction satisfies completeness.

Completeness. In our construction, both the prover and the verifier compute the challenges as $(b_1, b_2, \ldots, b_\ell) = \mathsf{H}_k(\mathsf{c}_1, \mathsf{c}_2, \ldots, \mathsf{c}_\ell)$. Hence, to prove that the verification succeeds, it suffices to prove that for each $i \in [\ell]$, z_i is a valid answer to c_i for the challenge b_i. In a nutshell, this follows from the completeness of Blum's sigma protocol.

More specifically, if $b_i = 0$, then in step 2, $\mathcal{P}$ computes $\mathsf{c}_i = (\mathsf{Com}_2(1^\lambda, 1^\mu, \mathsf{com}_1, \mathbf{b}', G_i[s,t]; r_i^{(s,t)}))_{s,t\in[n]}$ honestly with a random cycle graph G_i. Therefore, the verification in this case succeeds. Otherwise if $b_i = 1$, we need to show that $\mathsf{c}_i[s,t] = \mathsf{Com}_2(1^\lambda, 1^\mu, \mathsf{com}_1, \mathbf{b}', 0; r_i^{(s,t)})$ for every $e \notin E$ and $(s,t) = \phi(e)$. It suffices to show that $G_i[s,t] = 0$ for such (s,t). Note that if $e \notin E$, then $\phi(e) \notin \phi(G)$, since ϕ is a one-to-one map. Hence, if $(s,t) = \phi(e)$, then $G_i[s,t] = 0$. This completes the proof.

In the next section, we prove that our construction satisfies statistical witness indistinguishability and computational soundness.

5 Proofs of Security

5.1 Statistical Witness Indistinguishability

Theorem 5.1. *The construction in Section 4 satisfies statistical witness indistinguishability. Specifically, there exists a negligible function $\nu(\lambda)$ such that for every $G \in L_{\mathsf{HAM}}$ every two witness ω_1 and ω_2 for G, every (potentially maliciously computed) fixed first round message (com_1, k), the second round prover messages Π_1 Π_2 computed using ω_1 and ω_2 respectively, satisfy*

$$\mathsf{SD}(\Pi_1, \Pi_2) < 2^{-\Omega(\mu)} + 2n^2(\ell+1) \cdot \nu(\lambda)$$

Proof. We prove the theorem via a hybrid argument. For any fixed (com_1, k), let $\mathbf{b} = \mathsf{ComExt}(1^\lambda, 1^\mu, \mathsf{com}_1)$. Since $\mathbf{b}'$ is sampled uniformly at random by the prover, $\Pr[\mathbf{b} = \mathbf{b}'] = 2^{-\mu}$. Hence, with probability $1 - 2^{-\mu}$, $\mathbf{b} \neq \mathbf{b}'$. We now build a series of hybrids.

Hybrid H_0 : $(\mathbf{b}', \mathsf{c}_1, \mathsf{c}_2, \ldots, \mathsf{c}_\ell, z_1, z_2, \ldots, z_\ell) = \Pi_1$, where each z_j is computed honestly using ω_1.

Hybrid H_j : Same as above except that $z_1, z_2, \ldots, z_{j-1}$ are computed using witness ω_2, and $z_j, z_{j+1}, \ldots, z_\ell$ are computed using ω_1.

Hybrid $\mathsf{H}_j^1, (j = 0, 1, \ldots, \ell)$:

1. Sample $\mathbf{b}' \leftarrow \{0,1\}^{\mu}$. Generate $(\mathsf{c}_i)_{i\in[\ell]\setminus\{j\}}$ honestly in the same way as in the construction.
2. Compute $b'_j \leftarrow \{0,1\}$. Compute c_j honestly in the same way as in the construction.
3. Let $(b_1, b_2, \ldots, b_j, \ldots, b_\ell) \leftarrow \mathsf{H}_k(\mathsf{c}_1, \mathsf{c}_2, \ldots, \mathsf{c}_j, \ldots, \mathsf{c}_\ell)$. If $b'_j \neq b_j$, then goto 1, otherwise goto 4.
4. For $i \in [1, j-1]$, compute z_i honestly for challenge b_i using ω_2. For $i \in [j, \ell]$, compute z_i honestly for challenge b_i using ω_1. Output $(\mathbf{b}', \mathsf{c}_1, \mathsf{c}_2, \ldots, \mathsf{c}_\ell, z_1, z_2, \ldots, z_\ell)$.

Hybrid $\mathsf{H}^2_j, (j = 0, 1, \ldots, \ell)$:

1. Sample $\mathbf{b}' \leftarrow \{0,1\}^{\mu}$. Generate $(\mathsf{c}_i)_{i\in[\ell]\setminus\{j\}}$ honestly in the same way as in the construction.
2. Sample $b'_j \leftarrow \{0,1\}$. If $b'_j = 0$, then compute c_j honestly in the same way as in the construction, and generate z_j honestly. If $b'_j = 1$, then sample a uniformly random one-to-one map ϕ. For each $e \in \omega$, $(s,t) = \phi(e)$, set $G_j[s,t] = 1$. For other edges, set $G_j[s,t] = 0$. For each $s, t \in [n]$, sample uniformly random $r_j^{(s,t)}$, and compute $\mathsf{c}_j[s,t] := \mathsf{Com}_2(1^\lambda, 1^\mu, \mathsf{com}_1, \mathbf{b}', G_j[s,t]; r_j^{(s,t)})$. Set $z_j = (\phi, (r_j^{(s,t)})_{e\notin G,(s,t)=\phi(e)})$.
3. Let $(b_1, b_2, \ldots, b_j, \ldots, b_\ell) \leftarrow \mathsf{H}_k(\mathsf{c}_1, \mathsf{c}_2, \ldots, \mathsf{c}_j, \ldots, \mathsf{c}_\ell)$. If $b'_j \neq b_j$, then goto 1, otherwise goto 4.
4. For $i \in [1, j-1]$, generate z_i according to the challenge b_i honestly using ω_2. For $i \in [j+1, \ell]$, generate z_i according to the challenge b_i honestly using ω_1. Output $(\mathbf{b}', \mathsf{c}_1, \mathsf{c}_2, \ldots, \mathsf{c}_\ell, z_1, z_2, \ldots, z_\ell)$.

Hybrid $\mathsf{H}^3_j, (j = 0, 1, \ldots, \ell)$:

1. Sample $\mathbf{b}' \leftarrow \{0,1\}^{\mu}$. Generate $(\mathsf{c}_i)_{i\in[\ell]\setminus\{j\}}$ honestly in the same way as in the construction 4.
2. Sample $b'_j \leftarrow \{0,1\}$. If $b'_j = 0$, then compute c_j honestly in the same way as in the construction, and generate z_j honestly. If $b'_j = 1$, then sample a uniformly random one-to-one map ϕ. For each $e \notin E$, $(s,t) = \phi(e)$, sample a uniformly random $r_j^{(s,t)}$, and compute $\mathsf{c}_j[s,t] := \mathsf{Com}_2(1^\lambda, 1^\mu, \mathsf{com}_1, \mathbf{b}', G_j[s,t]; r_j^{(s,t)})$. Further, for each $e \in E$, $(s,t) = \phi(e)$, compute $\mathsf{c}_j[s,t] \leftarrow \mathsf{Sim}(1^\lambda, 1^\mu, \mathsf{com}_1)$, where Sim is the simulator for the public-coin statistical-hiding extractable commitment scheme.
3. Let $(b_1, b_2, \ldots, b_j, \ldots, b_\ell) \leftarrow \mathsf{H}_k(\mathsf{c}_1, \mathsf{c}_2, \ldots, \mathsf{c}_j, \ldots, \mathsf{c}_\ell)$. If $b'_j \neq b_j$, then goto 1, otherwise goto 4.
4. For $i \in [1, j-1]$, generate z_i according to the challenge b_i honestly using ω_2. For $i \in [j+1, \ell]$, generate z_i according to the challenge b_i honestly using ω_1. Output $(\mathbf{b}', \mathsf{c}_1, \mathsf{c}_2, \ldots, \mathsf{c}_\ell, z_1, z_2, \ldots, z_\ell)$.

Hybrid $\mathsf{H}_{\ell+1}$: $(\mathbf{b}', \mathsf{c}_1, \mathsf{c}_2, \ldots, \mathsf{c}_\ell, z_1, z_2, \ldots, z_\ell) = \Pi_2$, where each z_j are generated using ω_2.

This completes the description of the hybrids. We now prove a series of lemmas to bound the statistical distance between different adjacent hybrids. The proof then follows by combining the claims of the lemmas.

Lemma 5.2. $\mathsf{SD}(\mathsf{H}_j, \mathsf{H}^1_j) = 0$

Proof. The difference between H_j and H^1_j is that H^1_j has a rejection sampling process on b_j. Hence, we have

$$\begin{aligned}\Pr[\mathbf{b}', \mathsf{c}_1, \mathsf{c}_2, \ldots, \mathsf{c}_\ell, z_1, z_2, \ldots, z_\ell | \mathsf{H}^1_j] &= \Pr[\mathbf{b}', \mathsf{c}_1, \mathsf{c}_2, \ldots, \mathsf{c}_\ell, z_1, z_2, \ldots, z_\ell | \mathsf{H}_j, b_j = b'_j] \\ &= \Pr[\mathbf{b}', \mathsf{c}_1, \mathsf{c}_2, \ldots, \mathsf{c}_\ell, z_1, z_2, \ldots, z_\ell | \mathsf{H}_j]\end{aligned}$$

The second equality comes from the fact that b'_j is chosen uniformly at random. □

Lemma 5.3. $\mathsf{SD}(\mathsf{H}_j^1, \mathsf{H}_j^2) = 0$

Proof. The only difference between H_j^1 and H_j^2 is that in H_j^1, we sample a cycle graph G_j uniformly at random whereas in H_j^2, we first sample the one-to-one map ϕ uniformly at random and then generate $G_i = \phi(w)$. Hence, the distributions over (ϕ, G_i) in H_j^1 and H_j^2 are identical. □

Lemma 5.4. $\mathsf{SD}(\mathsf{H}_j^2, \mathsf{H}_j^3) < n^2 \cdot \nu(\lambda) + 2^{-\Omega(\mu)}$

Proof. The difference between H_j^2 and H_j^3 is that in H_j^3, we use the simulator of the public-coin statistical-hiding commitment scheme for computing $\mathsf{c}_i[s,t]$ for each $e \in E$, $(s,t) = \phi(e)$. However, since the randomness $r_j^{(s,t)}$ for each such $\mathsf{c}_i[s,t]$ is never opened, the claim follows from the statistical hiding property of the commitment scheme. □

Lemma 5.5. $\mathsf{SD}(\mathsf{H}_j^3, \mathsf{H}_{j+1}) < n^2 \cdot \nu(\lambda) + 2^{-\Omega(\mu)}$

Proof. Note that proving $\mathsf{SD}(\mathsf{H}_j^3, \mathsf{H}_{j+1}) < n^2 \cdot \nu(\lambda) + 2^{-\Omega(\mu)}$ is symmetric to proving that $\mathsf{SD}(\mathsf{H}_j, \mathsf{H}_j^3) < n^2 \cdot \nu(\lambda) + 2^{-\Omega(\mu)}$. The latter follows by combining Lemmas 5.2, 5.3, 5.4. The proof follows the same strategy as previous lemmas. □

□

5.2 Computational Soundness

Theorem 5.6. *The construction in Section 4 satisfies computational soundness.*

Proof. Suppose $G \notin L_{\mathsf{HAM}}$ and there exists a cheating prover $\mathcal{P}^*$ such that $\Pr[\mathcal{P}^* \text{ succeeds}] \geq 1/\lambda^c$ for infinite many λ. Then for each such λ, there must exist a $\mathbf{b}'_0$ such that $\Pr[\mathcal{P}^* \text{ succeeds} \wedge \mathbf{b}' = \mathbf{b}'_0] \geq \lambda^{-c}2^{-\mu}$, where $\mathbf{b}'$ is outputted by the cheating prover $\mathcal{P}^*$ in the second round.

$\mathbf{b}'_0$-Extractor Ext. We first describe an algorithm Ext that extracts a $\mathbf{b}'_0$ from any cheating prover $\mathcal{P}^*$, such that $\Pr[\mathcal{P}^* \text{ succeeds } \wedge \mathbf{b}' = \mathbf{b}'_0] \geq \lambda^{-c}2^{-\mu-1}$. Ext receives oracle access to $\mathcal{P}^*$.

1. Initialize an empty multiset $S = \{\}$.
2. For $j \in [2^{1.5\mu}]$, set fresh random tape for $\mathcal{P}^*$. Compute and send uniformly random first round message $(\mathsf{Com}_1(1^\lambda, 1^\mu), k \leftarrow \mathsf{KGen}(1^\lambda, 1^{|C_{\mathsf{st}}|}))$ to $\mathcal{P}^*$. Let $(\mathbf{b}'^{(j)}, (\mathsf{c}_i^{(j)})_{i\in[\ell]}, (z_i^{(j)})_{i\in[\ell]})$ be the response of $\mathcal{P}^*$. Execute the verifier algorithm; if verification suceeds, then append multiset $S = S \cup \{\mathbf{b}'^{(j)}\}$.
3. Output $\mathbf{b}'_0$ that appears for the maximum number of times in the multiset S.

In the sequel, we denote $p_\lambda = \Pr[\mathcal{P}^* \text{ succeeds}]$.

Claim 5.7. *The algorithm* Ext *runs in time* $O(2^{1.5\mu}) = 2^{O(\log^2 \lambda)}$. *Furthermore, with probability* $1 - \exp(-\Omega(2^{0.5\mu}p_\lambda))$, *it outputs a* $\mathbf{b}'_0$ *such that* $\Pr[\mathcal{P}^* \text{ succeeds } \wedge \mathbf{b}' = \mathbf{b}'_0] \geq p_\lambda/2^{-\mu-1}$.

We defer the proof of the Claim 5.7 to the end of the proof of Theorem 5.6.

Now we use the extractor Ext to build the following hybrids.

Hybrid H_0 **:** Compute $\mathbf{b}'_0 \leftarrow \mathsf{Ext}(\mathcal{P}^*)$. Generate uniformly random string $(\mathsf{com}_1 \leftarrow \mathsf{Com}_1(1^\lambda, 1^\mu), k \leftarrow \mathsf{KGen}(1^\lambda, 1^{|C_{\mathsf{st}}|}))$. Send (com_1, k) to $\mathcal{P}^*$. Let $(\mathbf{b}', (\mathsf{c}_i)_{i\in[\ell]}, (z_i)_{i\in[\ell]})$ be the output of $\mathcal{P}^*$.

If $\mathbf{b}' = \mathbf{b}'_0$ and $(\mathbf{b}', (\mathsf{c}_i)_{i\in[\ell]}, (z_i)_{i\in[\ell]})$ passes the verification, then the hybrid outputs 1, otherwise outputs 0.

Hybrid H_1 **:** Compute $\mathbf{b}'_0 \leftarrow \mathsf{Ext}(\mathcal{P}^*)$. $\underline{\text{Generate } (\mathsf{com}_1, \mathsf{st}) \leftarrow \mathsf{FakeCom}(1^\lambda, 1^\mu, \mathbf{b}'_0)}$, $k \leftarrow \mathsf{KGen}(1^\lambda, 1^{|C_{\mathsf{st}}|})$. Send (com_1, k) to $\mathcal{P}^*$. Let $(\mathbf{b}', (\mathsf{c}_i)_{i\in[\ell]}, (z_i)_{i\in[\ell]})$ be the output of $\mathcal{P}^*$.

If $\mathbf{b}' = \mathbf{b}'_0$ and $(\mathbf{b}', (\mathsf{c}_i)_{i\in[\ell]}, (z_i)_{i\in[\ell]})$ passes the verification, then the hybrid outputs 1, otherwise output 0.

Hybrid H_2 **:** Compute $\mathbf{b}'_0 \leftarrow \mathsf{Ext}(\mathcal{P}^*)$. Generate $(\mathsf{com}_1, \mathsf{st}) \leftarrow \mathsf{FakeCom}(1^\lambda, 1^\mu, \mathbf{b}'_0)$, $\underline{k \leftarrow \mathsf{FakeGen}(1^\lambda, 1^{|C_{\mathsf{st}}|}, C_{\mathsf{st}})}$. Send (com_1, k) to $\mathcal{P}^*$. Let $(\mathbf{b}', (\mathsf{c}_i)_{i\in[\ell]}, (z_i)_{i\in[\ell]})$ be the output of $\mathcal{P}^*$.

If $\mathbf{b}' = \mathbf{b}'_0$ and $(\mathbf{b}', (\mathsf{c}_i)_{i\in[\ell]}, (z_i)_{i\in[\ell]})$ passes the verification, then the hybrid outputs 1, otherwise outputs 0.

This completes the description of the hybrids. We now prove Lemmas 5.8 and 5.9 to establish the indistinguishability of the hybrids.

Lemma 5.8. $|\Pr[\mathsf{H}_0 = 1] - \Pr[\mathsf{H}_1 = 1]| < 2^{-\Omega(\log^4 \lambda)}$.

Proof. We prove this Lemma by relying on *quasi-polynomial pseudorandom receiver's message* property of the commitment scheme (Definition 3.3). We build the following adversary $\mathcal{D}$ trying to distinguish the receiver's message of commitment scheme from random string.

$\mathcal{D}$ takes as input $(1^\lambda, 1^\mu, \mathsf{com}_1)$. Firstly, $\mathcal{D}$ computes $\mathbf{b}'_0 \leftarrow \mathsf{Ext}(\mathcal{P}^*)$. Then, it generates $k \leftarrow \mathsf{KGen}(1^\lambda, 1^{|C_{\mathsf{st}}|})$ and sends (com_1, k) to $\mathcal{P}^*$. Let $(\mathbf{b}', (\mathsf{c}_i)_{i\in[\ell]}, (z_i)_{i\in[\ell]})$ be the response of $\mathcal{P}^*$. If $\mathbf{b}' = \mathbf{b}'_0$ and $(\mathbf{b}, (\mathsf{c}_i)_{i\in[\ell]}, (z_i)_{i\in[\ell]})$ passes the verification, then output 1. Otherwise output 0.

Now $\mathcal{D}(1^\lambda, 1^\mu, \mathsf{Com}_1(1^\lambda, 1^\mu))$ simulates the environment of H_0 for $\mathcal{P}^*$. Hence,

$$\Pr\left[\mathcal{D}(1^\lambda, 1^\mu, \mathsf{Com}_1(1^\lambda, 1^\mu)) = 1\right] = \Pr[\mathsf{H}_0 = 1]$$

Also, $\mathcal{D}(1^\lambda, 1^\mu, \mathsf{FakeCom}(1^\lambda, 1^\mu, \mathbf{b}'_0))$ simulates the environment of H_1. Hence,

$$\Pr\left[\mathcal{D}(1^\lambda, 1^\mu, \mathsf{FakeCom}_1(1^\lambda, 1^\mu, \mathbf{b}'_0)) = 1\right] = \Pr[\mathsf{H}_1 = 1]$$

From Claim 5.7, $\mathcal{D}$ runs in time $2^{O(\log^2 \lambda)}$. Since the distributions $\mathsf{Com}(1^\lambda, 1^\mu)$ and $\mathsf{FakeCom}(1^\lambda, 1^\mu, \mathbf{b}'_0)$ are quasi-polynomially indistinguishable,

$$\left|\Pr\left[\mathcal{D}(1^\lambda, 1^\mu, \mathsf{Com}_1(1^\lambda, 1^\mu)) = 1\right] - \Pr\left[\mathcal{D}(1^\lambda, 1^\mu, \mathsf{FakeCom}_1(1^\lambda, 1^\mu, \mathbf{b}'_0) = 1\right]\right| < 2^{-\Omega(\log^4 \lambda)}$$

Thus, we derive that $|\Pr[\mathsf{H}_0 = 1] - \Pr[\mathsf{H}_1 = 1]| \leq 2^{-\Omega(\log^4 \lambda)}$. □

Lemma 5.9. $|\Pr[\mathsf{H}_1 = 1] - \Pr[\mathsf{H}_2 = 1]| < 2^{-\Omega(\log^4 \lambda)}$.

Proof. We prove this lemma by relying on *quasi-polynomial pseduorandom fake key* property of CIH. We build adversary $\mathcal{D}$ trying to distinguish the fake CIH key from uniform random string.

$\mathcal{D}$ takes as input $(1^\lambda, 1^\mu, k)$. It first computes $\mathbf{b}'_0 \leftarrow \mathsf{Ext}(\mathcal{P}^*)$. Next, it generates $\mathsf{com}_1 \leftarrow \mathsf{FakeCom}_1(1^\lambda, 1^\mu, \mathbf{b}'_0)$ and sends (com_1, k) to $\mathcal{P}^*$. Let $(\mathbf{b}', (\mathsf{c}_i)_{i\in[\ell]}, (z_i)_{i\in[\ell]})$ be the response of $\mathcal{P}^*$. If $\mathbf{b}' = \mathbf{b}'_0$ and $(\mathbf{b}, (\mathsf{c}_i)_{i\in[\ell]}, (z_i)_{i\in[\ell]})$ passes the verification, then output 1. Otherwise output 0.

Now $\mathcal{D}(1^\lambda, 1^{|C_{\mathsf{st}}|}, k \leftarrow \mathsf{KGen}(1^\lambda, 1^{|C_{\mathsf{st}}|}))$ simulates the environment of H_1 for $\mathcal{P}^*$. Hence,

$$\Pr[\mathcal{D}(1^\lambda, 1^{|C_{\mathsf{st}}|}, k \leftarrow \mathsf{KGen}(1^\lambda, 1^{|C_{\mathsf{st}}|})) = 1] = \Pr[\mathsf{H}_1 = 1]$$

Also, $\mathcal{D}(1^\lambda, 1^{|C_{\mathsf{st}}|}, k \leftarrow \mathsf{FakeGen}(1^\lambda, 1^{|C_{\mathsf{st}}|}, C_{\mathsf{st}}))$ simulates the environment of H_2. Hence,

$$\Pr[\mathcal{D}(1^\lambda, 1^{|C_{\mathsf{st}}|}, k \leftarrow \mathsf{FakeGen}(1^\lambda, 1^{|C_{\mathsf{st}}|}, C_{\mathsf{st}})) = 1] = \Pr[\mathsf{H}_2 = 1]$$

From Claim 5.7, $\mathcal{D}$ runs in time $2^{O(\log^2 \lambda)}$. Since the distributions $\mathsf{KGen}(1^\lambda, 1^{|C_{\mathsf{st}}|})$ and $\mathsf{FakeGen}(1^\lambda, 1^{|C_{\mathsf{st}}|}, C_{\mathsf{st}})$ are quasi-polynomially indistinguishable, we have

$$|\Pr[\mathcal{D}(1^\lambda, 1^{|C_{\mathsf{st}}|}, k \leftarrow \mathsf{KGen}(1^\lambda, 1^{|C_{\mathsf{st}}|})) = 1] - \Pr[\mathcal{D}(1^\lambda, 1^{|C_{\mathsf{st}}|}, k \leftarrow \mathsf{FakeGen}(1^\lambda, 1^{|C_{\mathsf{st}}|}, C_{\mathsf{st}})) = 1]| < 2^{-\Omega(\log^4 \lambda)}$$

Thus, we derive $|\Pr[\mathsf{H}_1 = 1] - \Pr[\mathsf{H}_2 = 1]| \leq 2^{-\Omega(\log^4 \lambda)}$. □

We now prove the following lemma to lower bound the probability that the output of H_2 is 1.

Lemma 5.10. $\Pr[\mathsf{H}_2 = 1] \geq \lambda^{-c} 2^{-\mu-2} - 2 \cdot 2^{-\Omega(\log^4 \lambda)}$

Proof. From Claim 5.7, we have

$$\begin{aligned}
\Pr[\mathsf{H}_0 = 1] &= \Pr[\mathbf{b}_0' \leftarrow \mathsf{Ext}(\mathcal{P}^*) : \mathcal{P}^* \text{ succeeds } \wedge \mathbf{b}' = \mathbf{b}_0'] \\
&\geq \Pr\left[\mathbf{b}_0' \leftarrow \mathsf{Ext}(\mathcal{P}^*) : \mathcal{P}^* \text{ succeeds } \wedge \mathbf{b}' = \mathbf{b}_0' \wedge \Pr[\mathcal{P}^* \text{ succeeds } \wedge \mathbf{b}' = \mathbf{b}_0'] > p_\lambda 2^{-\mu-1}\right] \\
&= \Pr[\mathcal{P}^* \text{ succeeds } \wedge \mathbf{b}' = \mathbf{b}_0' | \Pr[\mathcal{P}^* \text{ succeeds } \wedge \mathbf{b}' = \mathbf{b}_0'] > p_\lambda 2^{-\mu-1}] \\
&\cdot \Pr[\mathbf{b}_0' \leftarrow \mathsf{Ext}(\mathcal{P}^*) : \Pr[\mathcal{P}^* \text{ succeeds } \wedge \mathbf{b}' = \mathbf{b}_0'] > p_\lambda 2^{-\mu-1}] \\
&> \lambda^{-c} 2^{-\mu-1} \cdot \left(1 - \exp\left(-\Omega(2^{0.5\mu} p_\lambda)\right)\right) \geq \lambda^{-c} 2^{-\mu-2}
\end{aligned}$$

Combining the above with the Lemma 5.8 and Lemma 5.9, we have $\Pr[\mathsf{H}_2 = 1] \geq \lambda^{-c} 2^{-\mu-2} - 2 \cdot 2^{-\Omega(\log^4 \lambda)}$. □

In the remainder of the proof, we use the $2^{-\lambda}$-correlation intractability property of the CIH to reach a contradiction. Towards this, we first show in the following lemma that $\mathsf{H}_2 = 1$ implies that there exists a 'collision' for CIH and C_{st}. Specifically, we show that any accepting proof in hybrid H_2 such that $\mathbf{b}' = \mathbf{b}_0'$, we can find a 'collision' for CIH and C_{st}.

Lemma 5.11. *If hybrid* H_2 *outputs* 1*, denote* $\mathsf{COM} = (\mathsf{c}_1, \mathsf{c}_2, \ldots, \mathsf{c}_\ell)$ *in the accepting proof. Then* $\mathsf{H}_k(\mathsf{COM}) = C_{\mathsf{st}}(\mathsf{COM})$.

Proof. We will prove by contradiction. Denote $(b_1, b_2, \ldots, b_\ell) = \mathsf{H}_k(\mathsf{COM})$. Suppose there is an $i \in [\ell]$ such that $b_i \neq C_{\mathsf{st}}(\mathsf{c}_i)$. Now we consider two cases: (1). $b_i = 0, C_{\mathsf{st}}(\mathsf{c}_i) = 1$, (2). $b_i = 1, C_{\mathsf{st}}(\mathsf{c}_i) = 0$.

For case (1), since $b_i = 0$, z_i must be of the form $(G_i, (r_i^{(s,t)})_{s,t \in [n]})$, where G_i is a cycle graph, and $\mathsf{c}_i[s,t] = \mathsf{Com}_2(1^\lambda, 1^\mu, \mathsf{com}_1, \mathbf{b}', G_i[s,t]; r_i^{(s,t)})$ for each $s, t \in [n]$. From the extractability property of the commitment scheme and $\mathbf{b}' = \mathbf{b}_0'$, we have $\mathsf{Dec}(1^\lambda, 1^\mu, \mathsf{st}, \mathsf{c}_i[s,t]) = G_i[s,t]$. Since G_i is a cycle graph, $C_{\mathsf{st}}(\mathsf{c}_i) = 0$. Therefore, we reach a contradiction.

For case (2), since $b_i = 1$, z_i must be the form $(\phi, (r_i^{(s,t)})_{e \notin E, (s,t) = \phi(e)})$, where ϕ is a one-to-one map, and $\mathsf{c}_i[s,t] = \mathsf{Com}_2(1^\lambda, 1^\mu, \mathsf{com}_1, \mathbf{b}', 0; r_i^{(s,t)})$ for each $e \notin E, (s,t) = \phi(e)$. Let $G_i[s,t] = \mathsf{Dec}(1^\lambda, 1^\mu, \mathsf{st}, \mathsf{c}_i[s,t])$ for each $s, t \in [n]$. Since $C_{\mathsf{st}}(\mathsf{c}_i) = 0$, G_i is a cycle graph. For each edge $e' = (s', t')$ of the cycle graph, $G_i[s', t'] = 1$. Now we will show that $(\phi^{-1}(s'), \phi^{-1}(t')) \in E$. We show this by contradiction. Suppose $(\phi^{-1}(s'), \phi^{-1}(t')) \notin E$, then $\mathsf{c}_i[s', t'] = \mathsf{Com}_2(1^\lambda, 1^\mu, \mathsf{com}_1, \mathbf{b}', 0; r_i^{(s',t')})$. From extractable property of commitment scheme, $\mathsf{Dec}(1^\lambda, 1^\mu, \mathsf{st}, \mathsf{c}_i[s', t']) = 0$, which implies $G_i[s', t'] = 0$. Thus, we find a contradiction. Hence, for each edge e in cycle graph G_i, $\phi^{-1}(e)$ is an edge in G. Now we have found a Hamiltonian cycle $\phi^{-1}(G_i) \subseteq G$, which is a contradiction to $G \notin L_{\mathsf{HAM}}$. □

Combining Lemmas 5.10 and Lemma 5.11, we derive that

$$\Pr\left[k \leftarrow \mathsf{FakeGen}(1^\lambda, 1^{|C_{\mathsf{st}}|}, C_{\mathsf{st}}) : \exists \mathsf{COM}, \mathsf{H}_k(\mathsf{COM}) = C_{\mathsf{st}}(\mathsf{COM})\right] \geq \lambda^{-c} 2^{-\mu-2} - 2 \cdot 2^{-\Omega(\log^4 \lambda)}$$

However, the above contradicts the $2^{-\lambda}$-statistical correlation intractability of CIH.

We now finish the proof by proving Claim 5.7.

Proof of Claim 5.7. Extractor Ext clearly runs in time $O(2^{1.5\mu})$. To lower bound the probability $\Pr[\mathcal{P}^* \text{ succeeds} \wedge \mathbf{b}' = \mathbf{b}'_0]$, we first give a lower bound on the size of multiset S. Note that in Step 2 of description of Ext, a new element is added to S with probability p_λ. From Chernoff bound,

$$\begin{aligned}\Pr[|S| > 2^{1.5\mu}p_\lambda/2] &= 1 - \Pr[|S| \leq 2^{1.5\mu}p_\lambda/2]\\ &\geq 1 - \exp(-2^{1.5\mu}p_\lambda/8)\end{aligned}$$

From pigeonhole principle, $\mathbf{b}'_0$ outputted by Ext must appear at least $|S|/2^\mu$ times in S. Now we have

$$\begin{aligned}&\Pr\left[\mathbf{b}'_0 \leftarrow \mathsf{Ext}(\mathcal{P}^*) : \Pr[\mathcal{P}^* \text{ succeeds} \wedge \mathbf{b}' = \mathbf{b}'_0] < \frac{2^{-\mu}p_\lambda}{2}\right]\\ =&\Pr\left[\mathbf{b}'_0 \leftarrow \mathsf{Ext}(\mathcal{P}^*) : \Pr[\mathbf{b}' = \mathbf{b}'_0|\mathcal{P}^* \text{ succeeds}]p_\lambda < \frac{2^{-\mu}p_\lambda}{2}\right]\\ =&\Pr\left[\mathbf{b}'_0 \leftarrow \mathsf{Ext}(\mathcal{P}^*) : \Pr[\mathbf{b}' = \mathbf{b}'_0|\mathcal{P}^* \text{ succeeds}] < 2^{-\mu}/2\right]\\ \leq&\Pr\left[\mathbf{b}'_0 \text{ appears at least } |S|/2^\mu \text{ times in } S \wedge \Pr[\mathbf{b}' = \mathbf{b}'_0|\mathcal{P}^* \text{ succeeds}] < 2^{-\mu}/2\right]\\ \leq&\exp\left(-\frac{1}{6}|S|2^{-\mu}\right)\end{aligned}$$

The last inequality follows from Chernoff bound. When $|S| \geq 2^{1.5\mu}p_\lambda/2$, this probability is upper bounded by $\exp\left(-\frac{1}{12}2^{0.5\mu}p_\lambda\right)$. By the union bound, we have

$$\Pr\left[\mathbf{b}'_0 \leftarrow \mathsf{Ext}(\mathcal{P}^*) : \Pr[\mathcal{P}^* \text{ succeeds} \wedge \mathbf{b}' = \mathbf{b}'_0] < \frac{2^{-\mu}p_\lambda}{2}\right] \leq \exp\left(-\frac{1}{8}2^{1.5\mu}p_\lambda\right) + \exp\left(-\frac{1}{12}2^{0.5\mu}p_\lambda\right)$$

□

□

References

[BD18] Zvika Brakerski and Nico Döttling. Two-message statistically sender-private OT from LWE. In Amos Beimel and Stefan Dziembowski, editors, *TCC 2018, Part II*, volume 11240 of *LNCS*, pages 370–390. Springer, Heidelberg, November 2018.

[BGI$^+$17] Saikrishna Badrinarayanan, Sanjam Garg, Yuval Ishai, Amit Sahai, and Akshay Wadia. Two-message witness indistinguishability and secure computation in the plain model from new assumptions. In Tsuyoshi Takagi and Thomas Peyrin, editors, *ASIACRYPT 2017, Part III*, volume 10626 of *LNCS*, pages 275–303. Springer, Heidelberg, December 2017.

[BP15] Nir Bitansky and Omer Paneth. ZAPs and non-interactive witness indistinguishability from indistinguishability obfuscation. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *TCC 2015, Part II*, volume 9015 of *LNCS*, pages 401–427. Springer, Heidelberg, March 2015.

[CCH$^+$19] Ran Canetti, Yilei Chen, Justin Holmgren, Alex Lombardi, Guy N. Rothblum, Ron D. Rothblum, and Daniel Wichs. Fiat-Shamir: from practice to theory. In Moses Charikar and Edith Cohen, editors, *51st ACM STOC*, pages 1082–1090. ACM Press, June 2019.

[DGS09] Yi Deng, Vipul Goyal, and Amit Sahai. Resolving the simultaneous resettability conjecture and a new non-black-box simulation strategy. In *50th FOCS*, pages 251–260. IEEE Computer Society Press, October 2009.

[DMP88] Alfredo De Santis, Silvio Micali, and Giuseppe Persiano. Non-interactive zero-knowledge proof systems. In Carl Pomerance, editor, *CRYPTO'87*, volume 293 of *LNCS*, pages 52–72. Springer, Heidelberg, August 1988.

[DN00] Cynthia Dwork and Moni Naor. Zaps and their applications. In *41st FOCS*, pages 283–293. IEEE Computer Society Press, November 2000.

[FLS90] Uriel Feige, Dror Lapidot, and Adi Shamir. Multiple non-interactive zero knowledge proofs based on a single random string (extended abstract). In *31st FOCS*, pages 308–317. IEEE Computer Society Press, October 1990.

[GMR85] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof-systems (extended abstract). In *17th ACM STOC*, pages 291–304. ACM Press, May 1985.

[GO94] Oded Goldreich and Yair Oren. Definitions and properties of zero-knowledge proof systems. *Journal of Cryptology*, 7(1):1–32, December 1994.

[GOS06] Jens Groth, Rafail Ostrovsky, and Amit Sahai. Perfect non-interactive zero knowledge for NP. In Serge Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 339–358. Springer, Heidelberg, May / June 2006.

[HHPV18] Shai Halevi, Carmit Hazay, Antigoni Polychroniadou, and Muthuramakrishnan Venkitasubramaniam. Round-optimal secure multi-party computation. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part II*, volume 10992 of *LNCS*, pages 488–520. Springer, Heidelberg, August 2018.

[HK12] Shai Halevi and Yael Tauman Kalai. Smooth projective hashing and two-message oblivious transfer. *Journal of Cryptology*, 25(1):158–193, January 2012.

[JKKR17] Abhishek Jain, Yael Tauman Kalai, Dakshita Khurana, and Ron Rothblum. Distinguisher-dependent simulation in two rounds and its applications. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part II*, volume 10402 of *LNCS*, pages 158–189. Springer, Heidelberg, August 2017.

[Kal05] Yael Tauman Kalai. Smooth projective hashing and two-message oblivious transfer. In Ronald Cramer, editor, *EUROCRYPT 2005*, volume 3494 of *LNCS*, pages 78–95. Springer, Heidelberg, May 2005.

[KKS18] Yael Tauman Kalai, Dakshita Khurana, and Amit Sahai. Statistical witness indistinguishability (and more) in two messages. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part III*, volume 10822 of *LNCS*, pages 34–65. Springer, Heidelberg, April / May 2018.

[NP01] Moni Naor and Benny Pinkas. Efficient oblivious transfer protocols. In S. Rao Kosaraju, editor, *12th SODA*, pages 448–457. ACM-SIAM, January 2001.

[PS19] Chris Peikert and Sina Shiehian. Noninteractive zero knowledge for NP from (plain) learning with errors. *IACR Cryptology ePrint Archive*, 2019:158, 2019.

[Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th ACM STOC*, pages 84–93. ACM Press, May 2005.