

Statistical Zaps and New Oblivious Transfer Protocols

Vipul Goyal¹, Abhishek Jain², Zhengzhong Jin², and Giulio Malavolta³

¹ Carnegie Mellon University, Pittsburgh, PA
vipul@cmu.edu

² Johns Hopkins University, Baltimore, MD
abhishek@cs.jhu.edu
zzjin@cs.jhu.edu

³ UC Berkeley & Carnegie Mellon University
giulio.malavolta@hotmail.it

Abstract. We study the problem of achieving *statistical privacy* in interactive proof systems and oblivious transfer – two of the most well studied two-party protocols – when limited rounds of interaction are available.

- **Statistical Zaps:** We give the first construction of statistical Zaps, namely, two-round statistical witness-indistinguishable (WI) protocols with a *public-coin* verifier. Our construction achieves computational soundness based on the quasi-polynomial hardness of learning with errors assumption.
- **Three-Round Statistical Receiver-Private Oblivious Transfer:** We give the first construction of a three-round oblivious transfer (OT) protocol – in the plain model – that achieves statistical privacy for receivers and computational privacy for senders against malicious adversaries, based on *polynomial-time* assumptions. The round-complexity of our protocol is optimal.

We obtain our first result by devising a public-coin approach to compress sigma protocols, without relying on trusted setup. To obtain our second result, we devise a general framework via a new notion of *statistical hash commitments* that may be of independent interest.

1 Introduction

We study the problem of achieving statistical privacy in two-party cryptographic protocols. Statistical privacy is very appealing in cryptography since it guarantees *everlasting security* – even if the adversary is computationally unbounded during the protocol execution and later post-processes the protocol transcript for as long as it wants, it cannot violate the privacy guarantee. For this reason, perhaps unsurprisingly, statistical privacy is typically much harder to achieve than computational privacy. For example, achieving statistical privacy for *both* participants in two-party protocols is impossible in general.

Nevertheless, in many scenarios, “one-sided” statistical privacy is possible to achieve. In other words, it is typically possible to design protocols that guarantee statistical privacy for one participant and computational privacy for the other. In this work, we investigate the possibility of achieving such asymmetric guarantees when *limited* rounds of interaction are available. We narrow the focus of our study on interactive proof systems [24,2] and oblivious transfer [39,17], two of the most well-studied two-party protocols in the cryptography literature.

Statistical Zaps. The notion of witness-indistinguishable (WI) proofs [19] allows a prover to convince a verifier about the validity of a statement (say) x in a manner such that the proof does not reveal which one of possibly multiple witnesses that attest to the validity of x was used in the computation. More specifically, if w_1, w_2 are both witnesses for x , then the verifier should not be able to distinguish between an honest prover using w_1 from an honest prover using w_2 . Despite offering a weaker privacy guarantee than zero-knowledge (ZK) proofs [24], WI has found wide applications in cryptography. One reason for its appeal is that most known round-complexity lower bounds for ZK do not apply to WI.

The seminal work of Dwork and Naor [15] proved that unlike ZK [23], WI can be achieved in two rounds, without relying on a trusted setup. They constructed two-round WI protocols with a *public-coin* verifier message, which they termed *Zaps*, from non-interactive zero-knowledge (NIZK) proofs in the common random string model [12,18]. By relying on known constructions of such NIZKs, their methodology can be used to obtain Zaps from quadratic residuosity [12], trapdoor permutations [18] and the decisional linear assumption over bilinear groups [26]. More recently, Zaps were also constructed based on indistinguishability obfuscation [6].

Over the years, Zaps have found numerous applications in cryptography. Part of their appeal is due to the public-coin verifier property which is crucial to many applications. In particular, it implies *public verifiability*, a property which is often used in the design of round-efficient secure multiparty computation protocols (see, e.g., [27]). 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 resettable-secure protocols (see, e.g., [13]).

Remarkably, all known constructions of Zaps (as well as non-interactive WI [5,25,6]) only achieve *computational* WI property. Despite several years of research, the following fundamental question has remained open:

Do there exist statistical Zaps?

In fact, even two-round statistical WI that only satisfy public-verifiability or reusability, in isolation, are not known currently. This is in contrast to NIZKs, which are indeed known with statistical privacy [8,38] or even perfect privacy [26]. One reason for this disparity is that the methodology of [15] for constructing Zaps is not applicable in the statistical case.

The recent work of Kalai, Khurana and Sahai [31] 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) [36,1,30,28,7]. The use of a private-coin verifier message is, in fact, instrumental to their approach (which builds on [29,4]). As such, a different approach is required for constructing statistical Zaps with a public-coin verifier.

Statistical Receiver-Private Oblivious Transfer. An oblivious transfer (OT) [39,17] protocol allows a “sender” to transfer one of its two inputs to a “receiver” without learning which of the inputs was obtained by the receiver. OT is of special importance to the theory and practice of secure computation [41,22] since OT is both necessary and complete [33].

Nearly two decades ago, the influential works of Naor and Pinkas [36] and Aiello et. al. [1] constructed two-round OT protocols that achieve game-based security against malicious adversaries in the plain model. An important property of these protocols is that they guarantee *statistical privacy for senders* (and computational privacy for receivers). Subsequent to these works, new constructions of such protocols were proposed based on a variety of assumptions (see, e.g., [30,28,7]). Over the years, such OT protocols have found many applications such as constructions of two-round (statistical) WI [29,4,31], non-malleable commitments [32], and more.

A natural question is whether it is possible to construct such OT protocols with a “reverse” guarantee, namely, *statistical privacy for receivers* (and computational privacy for senders). As observed in [31], two rounds are insufficient for this task: statistical receiver privacy implies that there exists different randomness tapes for receiver that explains a fixed receiver message for both input bits 0 and 1. Thus, a non-uniform malicious PPT receiver could simply start a two-round protocol with non-uniform advice that consists of such a message and randomness tapes, and then use both random tapes to learn *both* inputs of the sender, thereby violating sender privacy.

In the same work, [31] also proved that three rounds are sufficient for this task. Namely, they constructed three round statistical receiver-private OT with game-based security against malicious adversaries, in the plain model. However, they achieve this result by relying upon *super-polynomial-time* hardness assumptions. In contrast, two-round statistical sender-private OT protocols are known from polynomial-time assumptions. This leaves open the following important question:

Does there exist three-round statistical receiver-private OT in the plain model based on polynomial-time assumptions?

1.1 Our Results

In this work, we resolve both of the aforementioned questions in the affirmative.

I. Statistical Zap Arguments. We give the first construction of statistical Zaps with computational soundness, a.k.a. *statistical Zap arguments*. The soundness of our protocol is based on the quasi-polynomial hardness of the learning with errors (LWE) assumption. While we focus on achieving statistical privacy, we note that our construction, in fact, also yields the first computational Zap argument system based on (quasi-polynomial) LWE.

Theorem 1 (Informal). *Assuming quasi-polynomial LWE, there exists a statistical Zap argument system.*

In order to obtain our result, we depart significantly from prior approaches for constructing Zaps. Specifically, our approach combines the recent statistical NIZK arguments of Peikert and Shiehian [38] in a non-black-box manner with a two-round *public-coin* statistically-hiding extractable commitment scheme (see Section 4.1). Previously, such a commitment scheme in the private-coin setting was constructed by [31].

Roughly speaking, while the work of [38] (following [8]) instantiates the Fiat-Shamir methodology [19] for compressing sigma protocols [10] into a NIZK using

collision-intractable hash (CIH) functions [9], our approach can be seen as a way to compress sigma protocols into statistical Zaps using CIH and two-round public-coin statistically-hiding extractable commitments, without using a trusted setup. Importantly, while prior approaches for compressing sigma protocols into two-round WI [29,4,31] lose the public-coin property of the sigma protocol, our approach retains it. We refer the reader to Section 2.1 for more details on our technical approach.

Related work. In a concurrent and independent work, Badrinarayanan et al. [3] also construct statistical Zap arguments from quasi-polynomial LWE. In another concurrent and independent work, Lombardi et al. [34] construct computational Zap arguments from quasi-polynomial LWE. In a follow up work, Lombardi et al. [35] construct statistical Zaps with private verifier randomness from quasi-polynomial decisional linear assumption over groups with bilinear maps.

II. Three-Round Statistical Receiver-Private Oblivious Transfer. We devise a general framework for constructing three-round statistical receiver-private OT via a new notion of *statistical hash commitments* (SHC). This notion is inspired by hash proof systems [11] that were previously used to design two-round statistical sender-private OT [30,28]. Roughly speaking, an SHC scheme is a two-round statistically hiding commitment scheme where the opening verification simply involves an equality check with a hash output (computed w.r.t. a hashing algorithm associated with the scheme).

We devise a generic transformation from any SHC scheme with statistical hiding property to three-round statistical receiver-private OT. The resulting OT scheme achieves game-based security against malicious adversaries in the plain model. For the case of senders, we in fact, achieve a stronger notion of distinguisher-dependent simulation security [16,29]. Next, we provide two instantiations of an SHC scheme:

- A direct construction based on a search assumption, specifically, the computational Diffie-Hellman (CDH) problem. This construction, in fact, achieves *perfect* hiding property.
- We provide another construction of SHC based on any two-round statistical sender-private OT. Such schemes are known based on a variety of assumptions, including DDH, Quadratic (or N^{th}) Residuosity, and LWE. This yields a new approach for *OT reversal* [40] in the context of game-based security.

Putting these together, we obtain the following result:

Theorem 2 (Informal). *Assuming the existence of any two-round statistical sender-private OT (resp., polynomial hardness of CDH), there exists a three-round statistical (resp., perfect) receiver-private OT in the plain model.*

2 Technical Overview

2.1 Statistical Zap Arguments

We now prove a high-level overview of the main ideas underlying our construction of statistical Zaps. Roughly speaking, we devise a strategy to compress sigma protocols into statistical Zaps. While the idea of compressing sigma protocols to two-round WI

arguments has been considered before [29,4,31], the resulting protocol in these works were inherently private coin as they use oblivious transfer to “hide” the verifier message in the underlying sigma protocol. To obtain a public-coin protocol, we take a different approach.

Our starting point is the recent construction of NIZKs from LWE [38,8] that compresses any “trapdoor” sigma protocol into a NIZK by instantiating the Fiat-Shamir transformation [19] in the CRS model. We start by briefly recalling these constructions.

Recent Constructions of NIZKs from LWE. The main tool underlying the constructions of NIZK in [38,8] is the notion of Correlation Intractable Hash (CIH) functions. Roughly speaking, correlation intractability means that for any multi-bit-output circuit f , if we sample a hash function $H_k(\cdot)$ from the CIH function family, it is hard to find an input x such that $H_k(x)$ coincides with $f(x)$.

The work of [38] construct a NIZK for the Graph Hamiltonian Language⁴ starting from a sigma protocol for the same language. Recall that the first round prover message in the sigma protocol consists of commitments to some random cycle graphs. Let α denote the cycle graphs. The compression strategy works as follows: first, the prover prepares commitments to α by using a public-key encryption scheme, where the public-key is a part of the CRS setup. Next, the prover computes the verifier’s challenge in the sigma protocol by evaluating the CIH function over the first round message, where the CIH key is also fixed by the CRS setup. Given this challenge, the prover finally computes the third round message of the sigma protocol. The NIZK proof simply consists of this transcript.

Roughly speaking, the zero knowledge property of this construction relies on the semantic security of the public key encryption scheme (used to commit α) as well as the programmability of the CIH. Moreover, when the public key is *lossy*, then the NIZK in fact achieves *statistical zero knowledge* property.

The soundness property crucially relies upon the ability to *extract* the values α from the commitments by using the secret key corresponding to the public-key fixed by the CRS, as well as the correlation intractability of the CIH. Specifically, for any instance that is not in the language, given the secret key of the public key encryption, one can extract α from the commitment by decrypting it using the secret key, and then check if α corresponds to cycle graphs or not. Note that this checking procedure can be viewed as a function f . Then, if the malicious prover can find an accepting proof for the false statement, it implies that the output of the function f (with the secret key hardwired) evaluated over first round prover message coincides with the verifier’s challenge bits, which are outputted by the CIH function. However, from the correlation intractability of CIH, such a prover shouldn’t exist.

Starting Observations. Towards constructing statistical Zaps in the plain model, a naive first idea would be to simply let the verifier generate and send the CRS of the (statistical) NIZK in the first round, and then require the prover to compute and send the NIZK proof based on this CRS in the second round. This attempt, however, fails immediately since the verifier may use the trapdoor corresponding to the CRS (specifi-

⁴ Their construction, in fact, works for any trapdoor sigma protocol.

cally, the secret key corresponding to the public-key encryption) to extract the prover's witness.

One natural idea to address this issue is to replace the public-key encryption scheme with a two-round statistically-hiding commitment scheme. However, while this seems to address witness privacy concerns, it is no longer clear how to argue soundness since the proof of soundness (as discussed above) crucially requires the ability to extract the α values.

Achieving Weak Privacy. In order to devise a solution to the above problems, let us first consider a significantly weaker goal of constructing a two-round protocol that achieves computational soundness but only a very weak form of privacy guarantee, namely, that the verifier can learn the prover's witness with probability at most one-half. Moreover, we do not require the protocol to be public-coin, but only satisfy the weaker property of public verifiability.

To obtain such a protocol, we rely on a 2-round statistical sender-private oblivious transfer protocol in plain model [36,30,28,7]. In such an OT scheme, even if the receiver is malicious, at least one of the sender's messages remains statistically hidden from the receiver. Given such an OT scheme, we construct the desired two-round protocol as follows:

- In the first round, the verifier acts as the OT receiver, and sends a first round OT message with a random input bit b .
- In the second round, the prover prepares a transcript of the sigma protocol in the same manner as in the NIZK construction earlier, with the following key difference: it flips a coin b' and instead of computing the first round prover message as encryptions of α values, it computes OT sender messages where in each message, he uses inputs m_0, m_1 , where $m_{b'} = \alpha$ and $m_{1-b'} = \perp$.

With probability one-half, the random bit b of the verifier and the random coin b' of the prover are *different*. In this case, the statistical sender-privacy of the OT ensures that the α values remain hidden from the verifier. As such, the construction satisfies weak privacy, as required.

For computational soundness, consider any instance that is not in the language. Suppose we have an efficient cheating prover that can generate an accepting proof with non-negligible probability. In this case, we can run the cheating prover multiple times to estimate the distribution of the random coin b' . Note that at least one side of the random coin appears with probability no less than half. Without loss of generality, let assume such side is 0. Now we can switch the verifier's random hidden bit b in the first round message of OT to 0. Since the first round message of OT computationally hides b , the efficient cheating prover should not notice the switch, and hence the two random bits coincide with constant probability. However, when the two bits coincide, we can extract α by using the receiver's trapdoor of the OT. This allows us to contradict the correlation intractability of CIH, in the same manner as before.

Finally, note that the verifier does not need to use the randomness of the OT receiver to verify the proof; as such the above construction is publicly verifiable.

Amplifying Privacy. In order to amplify the privacy guarantee of the above scheme, we consider a modified approach where we replace the random bits b and b' – which

collide with probability one-half – with random strings of length ℓ that collide with $\frac{1}{2^\ell}$ probability. Specifically, consider a two-round protocol where the receiver’s input is a random string \mathbf{b} of length ℓ , while the sender also chooses a random string \mathbf{b}' of length ℓ and “encrypts” some message m . Suppose that the protocol satisfies the following “extractability” property, namely, if \mathbf{b} and \mathbf{b}' are equal, then the receiver can extract the encrypted message; otherwise, m remains statistically hidden.

Now consider a modified version of our weakly-private two-round argument system where we replace the two-round OT with the above “string” variant. Note that with probability $1 - 2^{-\ell}$, \mathbf{b} and \mathbf{b}' chosen by the prover and the verifier would be different, in which case, the α values would remain statistically hidden. This observation can, in fact, be turned into a formal proof for statistical witness indistinguishability.

The proof of computational soundness, however, now requires more work. Specifically, we now run the cheating prover for $\approx 2^\ell$ times, and estimate a \mathbf{b}'_0 that the cheating prover is most likely to output (with probability $\geq 1/2^\ell$). We then switch \mathbf{b} to \mathbf{b}'_0 . If the first round message of the receiver is secure against 2^ℓ -time adversaries, then the cheating prover would not notice the switch. We can now extract α values and derive a contradiction in a similar manner as before.

Two Round Public-Coin Statistical-Hiding Extractable Commitments. A two-round protocol that achieves statistical hiding property for the sender as well as extractability property of the aforementioned form was first formalized as a *statistical-hiding extractable commitment scheme* in the work of [31]. Their construction, however, is private coin for the receiver. Below, we briefly recall their construction, and then discuss how it can be adapted to the public-coin setting.

- In the first round, the receiver samples a uniformly random string \mathbf{b} of length ℓ . For each bit of the \mathbf{b} , the receiver sends a first round 1-out-of-2 OT message with the input bit specified by \mathbf{b} .
- The committer first samples a uniformly random string \mathbf{b}' of length ℓ . To commit to a message m , the committer firstly uses the xor secret sharing to share m to ℓ shares. It then generates ℓ second round OT messages: for the i -th second round OT message, if the i -th bit of \mathbf{b}' is 0, then the committer puts the share in the first input slot, and puts a random value in the second slot. Otherwise, the committer puts the share in the second slot, and put a random value in the first slot.

From statistical sender-privacy of the underlying OT, the above construction achieves statistically hiding with probability $1 - 2^{-\ell}$, even if the first round messages are maliciously generated.

Let us now explain the extractability property. For any committer, there exists a string \mathbf{b}_0 of length ℓ , such that the second string coincides with \mathbf{b}_0 with probability no less than $2^{-\ell}$. Therefore, we can switch the first round message of the commitment to hide \mathbf{b}_0 . If we set ℓ to be sub-linear, and assume the first round message is secure against sub-exponential-time adversaries, then the committer would not notice the switching. Hence, when the two strings coincide, we can extract the committed message.

The aforementioned statistical-hiding extractable commitment scheme is a private coin scheme. To obtain a public-coin scheme, we rely on the fact that in many known

statistical sender-private OT schemes, the first round message is pseudorandom. For example, in the recent construction of two-round statistical sender-private OT from LWE [7], the first round message is either statistical close to uniformly random, or is an LWE instance, which is computationally indistinguishable from the uniform distribution.

Putting it all together. Our final construction combines the above ideas to obtain a statistical Zap argument system:

- In the first round, the receiver simply sends the first round message of a two-round public-coin statistical-hiding extractable commitment scheme.
- Next, the prover samples a random string \mathbf{b}' and computes a transcript of the sigma protocol in the same manner as before, except that it commits to α values within the second round messages of the public-coin statistical-hiding extractable commitment scheme.

We argue the statistical WI property by relying on the statistical-hiding property of the commitment scheme. The proof of soundness relies on the ideas discussed above. In order to base security on quasi-polynomial hardness assumptions, we set the parameter ℓ for the commitment scheme to be super-logarithmic rather than sub-linear. Given any cheating prover with inverse polynomial advantage, we run the cheating prover several times to estimate a string \mathbf{b}_0 of length ℓ such that the string chosen by the prover coincides with \mathbf{b}_0 with some inverse quasi-polynomial probability. This estimation takes quasi-polynomial time. Next, we switch the first round verifier message to one that is computed using \mathbf{b}_0 . This switch is not noticeable to the prover since the first round message hides \mathbf{b}_0 even from adversaries that run in time 2^ℓ . This allows us to extract the α values and then invoke the correlation intractability of the CIH function as before. Note that we can construct the function f for CIH explicitly by using the receiver randomness for the first round message.

2.2 Three Round Statistical Receiver-Private OT

In this section, we describe our main ideas for constructing statistical receiver-private OT in three rounds in the plain model.

Prior work based on super-polynomial time assumptions. We start by briefly recalling the recent work of [31] who investigated the problem of statistical receiver-private OT in three rounds. Since security w.r.t. black-box polynomial-time simulation is known to be impossible to achieve in three rounds [20], [31] settled for the weaker goal of achieving security w.r.t. super-polynomial time simulation [37]. To achieve their goal, [31] implemented an OT reversal approach, starting from a two-round statistical sender-private OT to obtain a three-round statistical receiver-private OT based on super-polynomial-time hardness assumptions. In fact, the use of super-polynomial-time hardness assumptions seems somewhat inherent to their approach.

Motivated by our goal of basing security on standard polynomial-time hardness assumptions, we take a different approach, both in our security definition as well as techniques. On the definitional side, we consider distinguisher-dependent simulation security [16,29] for senders. On the technical side, we develop a general framework for

three round statistical receiver-private OT via a new notion of *statistical hash commitment*. We elaborate on both of these aspects below.

Defining Security. In the setting of interactive proof systems, a well-studied security notion is weak zero-knowledge [16] which relaxes the standard notion of zero knowledge by reversing the order of quantifiers, namely, by allowing the simulator to depend upon the distinguisher. A recent work of [29] dubbed this idea as *distinguisher-dependent simulation* and studied it for proof systems and some other two-party functionalities. Following their approach, in this work, we formalize security for senders in three round OT via distinguisher-dependent simulation. Roughly speaking, this notion requires that for every malicious PPT receiver and PPT distinguisher, there must exist a PPT simulator that can simulate an indistinguishable view of the receiver.

Towards achieving distinguisher-dependent simulation security for senders, we first consider (computational) game-based security definition for senders. Interestingly, it is not immediately clear how to define game-based security for senders when we also require statistical receiver privacy. This is because in any protocol that achieves statistical receiver privacy, the protocol transcript does not fix the receiver message in an information-theoretic sense. As such, unlike the case of two-round computational receiver-private OT (where the receiver’s input is information-theoretically fixed by the transcript), we cannot simply require indistinguishability of views generated using (say) sender inputs (m_b, m_{1-b}) and (m_b, m'_{1-b}) , where b is presumably the input bit of the receiver.

We resolve this conundrum by using an observation from [29]. In order to build proof systems with distinguisher-dependent simulation security, the work of [29] used the following natural property of two-round OT with computational privacy for senders and receivers – the distribution over receiver views generated using (say) sender inputs (m_0, m_1) must be indistinguishable from at least one of the following:

- Distribution over receiver views generated using sender inputs (m_0, m_0) .
- Distribution over receiver views generated using sender inputs (m_1, m_1)

Intuitively, the first case corresponds to receiver input bit 0, while the second case corresponds to receiver input bit 1.

It is not difficult to see that the above stated property is, in fact, meaningful even when the receiver’s input is only fixed in a computational sense by the protocol transcript, which is the case in our setting. A recent work of [14] formulated a game-based security definition for senders that captures the above intuition, and we adopt it in this work. We also show that for our three round setting, game-based security for senders can be used to achieve distinguisher-dependent simulation security for senders.

So far, we have focused on formalizing security for senders. Formalizing security for receivers is easier; we consider game-based security that requires statistical/perfect indistinguishability of views generated with receiver inputs 0 and 1, against unbounded-time malicious senders.

In the remainder of this section, we describe our main ideas for constructing three-round OT with game-based security for senders and receivers.

A General Framework via Statistical Hash Commitment. We introduce a new notion of an statistical hash commitment (SHC) scheme – a two-round statistically hiding

commitment scheme where the decommitment verification simply involves an equality check with a hash output (computed w.r.t. a hashing algorithm associated with the scheme). We start by informally defining this notion and then discuss how it can be used to construct three-round OT with our desired security properties.

An SHC scheme is a two-round commitment scheme between a committer \mathcal{C} and a receiver \mathcal{R} , that comes equipped with three additional algorithms – a key generation algorithm KGen , a commitment algorithm Com , and a hash algorithm H .

- In the first round, the Receiver \mathcal{R} samples a key pair $(pk, k) \leftarrow \text{KGen}$ and sends pk to the committer \mathcal{C} .
- In the second round, to commit a bit $b \in \{0, 1\}$, the committer \mathcal{C} executes $(c, \rho) \leftarrow \text{Com}(pk, b)$, and sends c to the receiver \mathcal{R} .
- In the opening phase, the committer \mathcal{C} sends (b, ρ) to the receiver \mathcal{R} .
- The verification algorithm only involves an equality check: \mathcal{R} computes the hash algorithm H using the private key k on input (c, b) and then matches the resulting value against ρ . If the check succeeds, then \mathcal{R} accepts the opening, else it rejects.
- **Computational Binding** This property requires that no PPT malicious committer \mathcal{C} can successfully compute a commitment c , and a opening ρ_0 and ρ_1 for *both* bits $b = 0$ and $b = 1$. Put differently, for an instance x and a second round message α , a PPT malicious committer cannot compute $\text{H}(k, c, b)$ for both $b = 0$ and $b = 1$.
- **Statistical (Perfect) Hiding** This property requires that, every (possibly maliciously computed) public key pk , the commitment of 0 and 1 are statistically close.

Looking ahead, we use computational binding property of SHC to achieve computational game-based security for senders in our construction of three-round OT. The statistical (resp., perfect) hiding property, on the other hand, is used to achieve statistical (resp., perfect) game-based security for receivers.

From SHC to Three-Round OT. We next describe a generic transformation from an SHC scheme statistical/perfect receiver-private OT. In our protocol design, the OT sender plays the role of the receiver in SHC, while the OT receiver plays the role of the committer for SHC. In the discussion below, let b denote the input bit of the OT receiver and let (m_0, m_1) denote the input bits of the OT sender.

- In the first round, the sender samples a key pair (pk, k) using the key generation algorithm KGen for SHC, and sends pk to the sender.
- In the second round, it runs the commitment algorithm Com for SHC on input (pk, b) to compute a second round message c and an opening ρ , and sends c to the sender.
- In the last round, the sender samples two random strings (r_0, r_1) and then computes two “mask” bits z_0 and z_1 , one each for its inputs m_0, m_1 . The mask z_i (for $i \in \{0, 1\}$) is computed as $hc(\text{H}(k, c, i), r_i)$, where $hc(\cdot, \cdot)$ is the Goldreich-Levin universal hardcore predicate [21].

To argue computational game-based security for senders, we crucially rely upon the strong soundness of SHC. In particular, the strong soundness of SHC, coupled with the

security of the hardcore predicate ensures that at least one of the two mask bits z_i must be hidden from a malicious PPT receiver when the instance x is sampled from a hard distribution. Statistical (resp., perfect) security for receivers, on the other hand, follows from the statistical (resp., perfect) hiding property of the commitment.

We next discuss two different constructions of SHC.

Instantiating SHC from CDH. We first describe a construction of SHC that achieves *perfect* hiding property, based on CDH.

Let $\mathbf{M} = \begin{pmatrix} 1 & 0 \\ y & 1 \end{pmatrix}$, which must be full rank. Note that $g^{\mathbf{M}}$ can be computed using g^y .

- In the first round, the receiver \mathcal{R} samples a random 2-by-1 column vector \mathbf{k} as the secret key of the hash function, and sets the public key pk to be $\text{pk} = (g^y, g^{\mathbf{M}\cdot\mathbf{k}})$. It then sends pk to the committer \mathcal{C} .
- The committer \mathcal{C} (with input bit $b \in \{0, 1\}$) samples a random 2-by-1 matrix α , and uses pk to compute $c = g^{\alpha^T \cdot \mathbf{M}} \cdot g^{[0,b]}$. The committer sends c to the verifier, and then compute $\rho = g^{\alpha^T \mathbf{M}\cdot\mathbf{k}}$
- The receiver \mathcal{R} parse $c = g^z$, and computes $H(\mathbf{k}, c, b) = g^{(z - [0,b])\cdot\mathbf{k}}$. If $H(\mathbf{k}, c, b) = \rho$, then accept, otherwise reject.

We next informally argue the security of the above construction. Let us first consider computational binding property. Intuitively, for any prover who wants to compute two accepting last round messages ρ_0, ρ_1 for both $b = 0$ and $b = 1$, it must compute the inverse of \mathbf{M} , which requires that the prover knows the witness y . More formally, to prove the computational binding property, we build a PPT extractor that extracts y to derive a contradiction. Specifically, for any cheating committer who can output two accepting ρ_0, ρ_1 for $b = 0$ and $b = 1$, we can divide them to derive $g^{[0,1]\cdot\mathbf{k}}$. If we parse \mathbf{k} as $\mathbf{k} = (s, t)$, then this implies that given $(g^y, g^{\mathbf{M}\mathbf{k}}) = (g^y, g^{sy}, g^{sy+t})$, an efficient algorithm can compute $g^{[0,1]\cdot\mathbf{k}} = g^t$. We can then divide it from g^{sy+t} and derive g^{sy} . This gives us an efficient adversary for CDH.

To prove statistical hiding property, for any (potentially maliciously computed) pk , the commitment of bit $b \in \{0, 1\}$ is $c = g^{\alpha^T \cdot \mathbf{M} + [0,b]}$. Since the matrix \mathbf{M} is full rank, and α is uniformly random, we have that c is uniformly random. Hence, the commitment statistically hides b .

Instantiating SHC from Statistical Sender-Private 2-round OT. We next show a construction of SHC from any statistical sender-private 2-round OT protocol $(\text{OT}_1, \text{OT}_2, \text{OT}_3)$, where OT_3 denotes the receiver output computation algorithm.

- In the first round, the receiver \mathcal{R} samples a random string r of length ℓ . Then for each bit $r[i]$, it invokes OT_1 to generate a first round OT message $(\text{ot}_{1,i}, \text{st}_i) \leftarrow \text{OT}_1(1^\lambda, r[i])$. The public key pk is set to be the tuple of messages $\{\text{ot}_{1,i}\}_{i \in [\ell]}$, while the private key \mathbf{k} is set to be the tuple of private states $\{\text{st}_i\}_{i \in [\ell]}$.
- The committer \mathcal{C} receives pk , and its input is a bit b . It first samples a random string r' of length ℓ . For each position $i \in [\ell]$, it generates the second round OT messages $\text{ot}_{2,i} = \text{OT}_2(\text{ot}_{1,i}, r'[i], r'[i] \oplus b)$. The commitment c is set to be the tuple of second round OT messages $\{\text{ot}_{2,i}\}_{i \in [\ell]}$, and the opening $\rho = r'$.

- The verification process first computes $H(k, c, b)$ as follows: parse k as $\{\text{st}_i\}_{i \in [\ell]}$, and the commitment c as $\{\text{ot}_{2,i}\}_{i \in [\ell]}$. Then, compute $\rho_{0,i} \leftarrow \text{OT}_3(\text{ot}_{2,i}, \text{st}_i)$, set $\rho_{1,i} = \rho_{0,i} \oplus r[i]$ for each $i \in [\ell]$, and set $\{\rho_{b,i}\}_{i \in [\ell]}$ to be the output of $H(k, c, b)$. If this output equals ρ , accept, otherwise, reject.

To show the completeness of this protocol, from the construction of the committer, we know that $\rho_{0,i} = r'[i] \oplus (r[i] \cdot b)$. From the computation of $H(k, c, b)$, we have that $\rho_{b,i} = \rho_{0,i} \oplus (r[i] \cdot b) = (r'[i] \oplus (r[i] \cdot b)) \oplus (r[i] \cdot b) = r'[i] = \rho$. The statistical hiding property follows from the statistical hiding property of the underlying OT. Finally, to show the construction is computational binding, our observation is that the construction of H always satisfies $H(k, c, 0) \oplus H(k, c, 1) = r$. Hence, any adversary breaking the computational binding property can also find $\rho_0 \oplus \rho_1 = H(k, c, 0) \oplus H(k, c, 1) = r$, given only the first round messages $\text{ot}_{1,i}$. This breaks the computational receiver privacy of the OT.

3 Preliminaries

For any two (discrete) probability distributions P and Q , let $\text{SD}(P, Q)$ 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.

3.1 Learning with Errors

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

Definition 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 (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

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

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 1, 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. We defer the proof to the full version.

Lemma 1. 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

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

3.2 Computational Diffie-Hellman Assumption

Definition 3. Let G be a cyclic group of order q generated by g , where each element of G can be represented in a polynomial $n = n(\lambda)$ number of bits. The CDH assumption states that for any non-uniform PPT adversary \mathcal{A} , there exists a negligible function $\nu(\lambda)$ such that

$$\Pr[x \leftarrow \mathbb{Z}_q, y \leftarrow \mathbb{Z}_q, z \leftarrow \mathcal{A}(1^\lambda, g^x, g^y) : z = g^{xy}] < \nu(\lambda)$$

3.3 Goldreich-Levin Hardcore Predicate

Definition 4. Let f be an one-way function from $\{0, 1\}^n \rightarrow \{0, 1\}^m$, where $n = n(\lambda)$ and $m = m(\lambda)$ are polynomials of λ . The Goldreich-Levin hardcore predicate hc is defined as $\text{hc}(x, r) = \langle x, r \rangle_2$, where $x, r \in \{0, 1\}^n$, and $\langle \cdot, \cdot \rangle_2$ is the inner product function modulo 2.

Theorem 3 (Goldreich-Levin Theorem [21], modified). If there exists an PPT adversary \mathcal{A} such that

$$\Pr[x \leftarrow \{0, 1\}^n, r \leftarrow \{0, 1\}^n, b \leftarrow \mathcal{A}(1^\lambda, (f(x), r)) : b = \text{hc}(x, r)] > 1/2 + \epsilon(\lambda)$$

where $\epsilon(\lambda)$ is a non-negligible function of λ , then there exists a PPT inverter \mathcal{A}' s.t.

$$\Pr[x \leftarrow \{0, 1\}^n, x' \leftarrow \mathcal{A}'(1^\lambda, f(x)) : x' = x] > \epsilon'(\lambda)$$

where $\epsilon'(\lambda)$ is also a non-negligible function of λ .

3.4 Statistical Zap Arguments

Zaps [15] 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 $\text{Trans}(\mathcal{P}(1^\lambda, x, \omega) \leftrightarrow \mathcal{V}(1^\lambda, x))$ 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 \mathcal{P} also has a witness ω for x .

Definition 5. 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

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

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

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

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

3.5 Statistical Sender-Private Oblivious Transfer

Definition 6. A statistical sender-private oblivious transfer (OT) is a tuple of algorithms $(\text{OT}_1, \text{OT}_2, \text{OT}_3)$:

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

$\text{OT}_2(1^\lambda, \text{ot}_1, m_0, m_1)$: 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 .

$\text{OT}_3(1^\lambda, \text{ot}_2, \text{st})$: 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 $b, m_0, m_1 \in \{0, 1\}$,

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

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 , $\text{OTExt}(1^\lambda, \text{ot}_1)$ outputs a bit $b \in \{0, 1\}$. Then for any $m_0, m_1 \in \{0, 1\}$, we have $\text{SD} (\text{OT}_2(1^\lambda, \text{ot}_1, m_0, m_1), \text{OT}_2(1^\lambda, \text{ot}_1, m_b, m_b)) < \nu(\lambda)$.

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

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

Lemma 2. Assuming quasi-polynomial hardness of LWE, there exists a statistical sender private oblivious transfer scheme.

A statistical sender-private OT scheme from LWE was recently constructed by [7]. Their construction satisfies correctness and statistical sender-privacy. Further, the receiver's message in their scheme is pseudorandom, assuming LWE. We observe that assuming quasi-polynomial LWE and using Lemma 1, their scheme also satisfies quasi-polynomially pseudorandom receiver's message property.

3.6 Correlation Intractable Hash Function

The following definition is taken verbatim from [38].

Definition 7 (Searchable Relation [38]). We say that a relation $R \subseteq \mathcal{X} \times \mathcal{Y}$ is searchable in size S if there exists a function $f : \mathcal{X} \rightarrow \mathcal{Y}$ 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 the 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 8 (Correlation Intractable Hash Function, slightly modified from [38]). Correlation Intractable Hash Function (CIH) is a triple of algorithms $(\text{KGen}, \text{FakeGen}, H_{(\cdot)}(\cdot))$, with the following properties:

Let $s = s(\lambda), \ell = \ell(\lambda), d = d(\lambda)$ be $\text{poly}(\lambda)$ -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 $\Pr[k \leftarrow \text{FakeGen}(1^{\lambda}, 1^{|C_{\lambda}|}, C_{\lambda}) : \exists x \text{ s.t. } (x, H_k(x)) \in R] < \nu(\lambda)$.

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

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

Theorem 4. Assuming quasi-polynomial hardness of LWE, there exists a construction of correlation intractable hash function with quasi-polynomial pseudorandom fake key.

The construction of such a function is given in [38,8]. Specifically, we use the construction of [38], which satisfies *statistical correlation intractability*. Moreover, the 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 Lemma 3 that by using parallel repetition, we can construct a CIH with the above property from any CIH.

Lemma 3 (Amplification of Statistical Correlation Intractability). There exists a correlation intractable hash function $(\text{KGen}, \text{FakeGen}, H_{(\cdot)}(\cdot))$ 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 $\overrightarrow{C_{\lambda,\ell}}(c_1, c_2, \dots, c_\ell) = (C_\lambda(c_1), C_\lambda(c_2), \dots, C_\lambda(c_\ell))$, then for large enough λ ,

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

The CIH in [38] already satisfies the above property. In the full version, we describe a generic transformation from any CIH to one that achieves the above property.

4 Statistical Zap Arguments

4.1 Public Coin Statistical-Hiding Extractable Commitments

In this section, we start by defining and constructing a key building block in our construction of statistical Zaps, namely, a statistical-hiding extractable commitment scheme. The notion and its construction are adapted from [31], with some slight modifications to fit in our application. The main difference between our definition and that of [31] is that we require the 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 $(\text{Com}_1, \text{FakeCom}_1, \text{Com}_2, \text{Dec})$, 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) 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 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 9. A public coin statistical-hiding extractable commitment is a tuple $(\text{Com}_1, \text{FakeCom}_1, \text{Com}_2, \text{Dec})$. 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 $\text{com}_2 \leftarrow \text{Com}_2(1^\lambda, 1^\mu, \text{com}_1, \mathbf{b}', m; r)$ with randomness r . \mathcal{C} sends $(\mathbf{b}', \text{com}_2)$ to \mathcal{R} .

Opening Phase

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

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 , $\text{ComExt}(1^\lambda, 1^\mu, \text{com}_1)$ outputs $\mathbf{b} \in \{0, 1\}^\mu$, and for any $\mathbf{b}' \neq \mathbf{b}$, and $m \in \{0, 1\}$, we have*

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

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

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

Extractable FakeCom_1 and Dec satisfy the following property. For any $\mathbf{b} \in \{0, 1\}^\mu$, we have

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

Lemma 4. *Assuming quasi-polynomial hardness of LWE, there exists a public coin statistical-hiding extractable commitment scheme.*

In the full version, we construct a public coin statistical hiding extractable commitment by slightly modifying the commitment scheme of [31]. 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 3.5), 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 $\text{Com}(1^\lambda, 1^\mu)$ simply computes a uniform random string, while FakeCom_1 corresponds to the receiver algorithm in the construction of [31].

4.2 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 $(\text{Com}_1, \text{FakeCom}_1, \text{Com}_2, \text{Dec})$ be a public coin statistical-hiding extractable commitment scheme (Definition 9). We set the parameter μ of the commitment scheme

as $\Theta(\log^2 \lambda)$. Let $(\text{KGen}, \text{FakeGen}, \text{H}_{(\cdot)}(\cdot))$ be a family of CIH (Definition 8). We choose the polynomial $\ell = \ell(\lambda)$ in Lemma 3 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 $G[s, t] = \text{Dec}(1^\lambda, 1^\mu, \text{st}, c_{s,t})$.
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 $(c_1, c_2, \dots, c_\ell)$. Specifically, $C_{\text{st}}(c_1, c_2, \dots, c_\ell)$ is defined as $(C_{\text{st}}(c_1), C_{\text{st}}(c_2), \dots, C_{\text{st}}(c_\ell))$.

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 $(\text{com}_1 \leftarrow \text{Com}_1(1^\lambda, 1^\mu), k \leftarrow \text{KGen}(1^\lambda, 1^{C_{\text{st}}}))$, 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 $c_i[s, t] \leftarrow \text{Com}_2(1^\lambda, 1^\mu, \text{com}_1, \mathbf{b}', G_i[s, t]; r_i^{(s,t)})$ using randomness $r_i^{(s,t)}$. Finally let $\mathbf{c}_i = (c_i[s, t])_{s,t \in [n]}$.
3. Compute $(b_1, b_2, \dots, b_\ell) = \text{H}_k(\mathbf{c}_1, \dots, \mathbf{c}_\ell)$.
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 $\phi : G \rightarrow G_i$ 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 $\Pi = (\mathbf{b}', (\mathbf{c}_i)_{i \in [\ell]}, (z_i)_{i \in [\ell]})$ to the verifier.

Verification Upon receiving the proof $\Pi = (\mathbf{b}', (\mathbf{c}_i)_{i \in [\ell]}, (z_i)_{i \in [\ell]})$, the verifier first computes $(b_1, b_2, \dots, b_\ell) = \text{H}_k(\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_\ell)$, and then verifies each copy (\mathbf{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 $\mathbf{c}_i = (\text{Com}_2(1^\lambda, 1^\mu, \text{com}_1, \mathbf{b}', G_i[s, t]; r_i^{(s,t)}))_{s,t \in [n]}$ 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 $c_i[s, t] = \text{Com}_2(1^\lambda, 1^\mu, \text{com}_1, \mathbf{b}', 0; r_i^{(s,t)})$. If all of the checks succeed, then accept the proof, otherwise reject.

This completes the description of our construction. We defer the proof of completeness and statistical witness indistinguishability to the full version. We next prove that our construction satisfies computational soundness.

Theorem 5. *The construction in Section 4.2 satisfies computational soundness.*

Suppose $G \notin L_{\text{HAM}}$ 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 $(\text{Com}_1(1^\lambda, 1^\mu), k \leftarrow \text{KGen}(1^\lambda, 1^{|C_{\text{st}}|}))$ to \mathcal{P}^* . Let $(\mathbf{b}'^{(j)}, (c_i^{(j)})_{i \in [\ell]}, (z_i^{(j)})_{i \in [\ell]})$ be the response of \mathcal{P}^* . Execute the verifier algorithm; if verification succeeds, 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 $p_\lambda = \Pr[\mathcal{P}^* \text{ succeeds}]$.

Lemma 5. *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 Lemma 5 to the full version. Now we use the extractor Ext to build the following hybrids.

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

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

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

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

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

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

This completes the description of the hybrids. We now prove Lemmas 6 and 7 to establish the indistinguishability of the hybrids.

Lemma 6. $|\Pr[H_0 = 1] - \Pr[H_1 = 1]| < 2^{-\Omega(\log^4 \lambda)}$.

Proof. We prove this Lemma by relying on *quasi-polynomial pseudorandom receiver's message* property of the commitment scheme (Definition 9). 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 $(1^\lambda, 1^\mu, \text{com}_1)$. Firstly, \mathcal{D} computes $\mathbf{b}'_0 \leftarrow \text{Ext}(\mathcal{P}^*)$. Then, it generates $k \leftarrow \text{KGen}(1^\lambda, 1^{|\mathcal{C}_{\text{st}}|})$ and sends (com_1, k) to \mathcal{P}^* . Let $(\mathbf{b}', (c_i)_{i \in [\ell]}, (z_i)_{i \in [\ell]})$ be the response of \mathcal{P}^* . If $\mathbf{b}' = \mathbf{b}'_0$ and $(\mathbf{b}, (c_i)_{i \in [\ell]}, (z_i)_{i \in [\ell]})$ passes the verification, then output 1. Otherwise output 0.

Now $\mathcal{D}(1^\lambda, 1^\mu, \text{Com}_1(1^\lambda, 1^\mu))$ simulates the environment of H_0 for \mathcal{P}^* . Hence, $\Pr[\mathcal{D}(1^\lambda, 1^\mu, \text{Com}_1(1^\lambda, 1^\mu)) = 1] = \Pr[H_0 = 1]$ Also, $\mathcal{D}(1^\lambda, 1^\mu, \text{FakeCom}(1^\lambda, 1^\mu, \mathbf{b}'_0))$ simulates the environment of H_1 . Hence, $\Pr[\mathcal{D}(1^\lambda, 1^\mu, \text{FakeCom}_1(1^\lambda, 1^\mu, \mathbf{b}'_0)) = 1] = \Pr[H_1 = 1]$.

From Lemma 5, \mathcal{D} runs in time $2^{O(\log^2 \lambda)}$. Since the distributions $\text{Com}(1^\lambda, 1^\mu)$ and $\text{FakeCom}(1^\lambda, 1^\mu, \mathbf{b}'_0)$ are quasi-polynomially indistinguishable,

$$\begin{aligned} & |\Pr[\mathcal{D}(1^\lambda, 1^\mu, \text{Com}_1(1^\lambda, 1^\mu)) = 1] - \\ & \Pr[\mathcal{D}(1^\lambda, 1^\mu, \text{FakeCom}_1(1^\lambda, 1^\mu, \mathbf{b}'_0)) = 1]| < 2^{-\Omega(\log^4 \lambda)} \end{aligned}$$

Thus, we derive that $|\Pr[H_0 = 1] - \Pr[H_1 = 1]| \leq 2^{-\Omega(\log^4 \lambda)}$. \square

Lemma 7. $|\Pr[H_1 = 1] - \Pr[H_2 = 1]| < 2^{-\Omega(\log^4 \lambda)}$.

Proof. We prove this lemma by relying on *quasi-polynomial pseudorandom 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 $\mathbf{b}'_0 \leftarrow \text{Ext}(\mathcal{P}^*)$. Next, it generates $\text{com}_1 \leftarrow \text{FakeCom}_1(1^\lambda, 1^\mu, \mathbf{b}'_0)$ and sends (com_1, k) to \mathcal{P}^* . Let $(\mathbf{b}', (c_i)_{i \in [\ell]}, (z_i)_{i \in [\ell]})$ be the response of \mathcal{P}^* . If $\mathbf{b}' = \mathbf{b}'_0$ and $(\mathbf{b}, (c_i)_{i \in [\ell]}, (z_i)_{i \in [\ell]})$ passes the verification, then output 1. Otherwise output 0.

Now $\mathcal{D}(1^\lambda, 1^{|\mathcal{C}_{\text{st}}|}, k \leftarrow \text{KGen}(1^\lambda, 1^{|\mathcal{C}_{\text{st}}|}))$ simulates the environment of H_1 for \mathcal{P}^* . Hence, $\Pr[\mathcal{D}(1^\lambda, 1^{|\mathcal{C}_{\text{st}}|}, k \leftarrow \text{KGen}(1^\lambda, 1^{|\mathcal{C}_{\text{st}}|})) = 1] = \Pr[H_1 = 1]$.

Also, $\mathcal{D}(1^\lambda, 1^{|\mathcal{C}_{\text{st}}|}, k \leftarrow \text{FakeGen}(1^\lambda, 1^{|\mathcal{C}_{\text{st}}|}, C_{\text{st}}))$ simulates the environment of H_2 . Hence, $\Pr[\mathcal{D}(1^\lambda, 1^{|\mathcal{C}_{\text{st}}|}, k \leftarrow \text{FakeGen}(1^\lambda, 1^{|\mathcal{C}_{\text{st}}|}, C_{\text{st}})) = 1] = \Pr[H_2 = 1]$.

From Lemma 5, \mathcal{D} runs in time $2^{O(\log^2 \lambda)}$. Since the distributions $\text{KGen}(1^\lambda, 1^{|\mathcal{C}_{\text{st}}|})$ and $\text{FakeGen}(1^\lambda, 1^{|\mathcal{C}_{\text{st}}|}, C_{\text{st}})$ are quasi-polynomially indistinguishable, we have

$$\begin{aligned} & |\Pr[\mathcal{D}(1^\lambda, 1^{|\mathcal{C}_{\text{st}}|}, k \leftarrow \text{KGen}(1^\lambda, 1^{|\mathcal{C}_{\text{st}}|})) = 1] - \\ & \Pr[\mathcal{D}(1^\lambda, 1^{|\mathcal{C}_{\text{st}}|}, k \leftarrow \text{FakeGen}(1^\lambda, 1^{|\mathcal{C}_{\text{st}}|}, C_{\text{st}})) = 1]| < 2^{-\Omega(\log^4 \lambda)} \end{aligned}$$

Thus, we derive $|\Pr[H_1 = 1] - \Pr[H_2 = 1]| \leq 2^{-\Omega(\log^4 \lambda)}$. \square

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

Lemma 8. $\Pr[H_2 = 1] \geq \lambda^{-c} 2^{-\mu-2} - 2 \cdot 2^{-\Omega(\log^4 \lambda)}$

Proof. From Lemma 5, we have

$$\begin{aligned}
\Pr[H_0 = 1] &= \Pr[\mathbf{b}'_0 \leftarrow \text{Ext}(\mathcal{P}^*) : \mathcal{P}^* \text{ succeeds} \wedge \mathbf{b}' = \mathbf{b}'_0] \\
&\geq \Pr[\mathbf{b}'_0 \leftarrow \text{Ext}(\mathcal{P}^*) : \mathcal{P}^* \text{ succeeds} \wedge \mathbf{b}' = \mathbf{b}'_0 \wedge \\
&\quad \Pr[\mathcal{P}^* \text{ succeeds} \wedge \mathbf{b}' = \mathbf{b}'_0] > p_\lambda 2^{-\mu-1}] \\
&= \Pr[\mathcal{P}^* \text{ succeeds} \wedge \mathbf{b}' = \mathbf{b}'_0] \Pr[\mathcal{P}^* \text{ succeeds} \wedge \mathbf{b}' = \mathbf{b}'_0] > p_\lambda 2^{-\mu-1}] \\
&\cdot \Pr[\mathbf{b}'_0 \leftarrow \text{Ext}(\mathcal{P}^*) : \Pr[\mathcal{P}^* \text{ succeeds} \wedge \mathbf{b}' = \mathbf{b}'_0] > p_\lambda 2^{-\mu-1}] \\
&\geq \lambda^{-c} 2^{-\mu-1} \cdot (1 - \exp(-\Omega(2^{0.5\mu} p_\lambda))) \geq \lambda^{-c} 2^{-\mu-2}
\end{aligned}$$

Combining the above with the Lemma 6 and Lemma 7, we have $\Pr[H_2 = 1] \geq \lambda^{-c} 2^{-\mu-2} - 2 \cdot 2^{-\Omega(\log^4 \lambda)}$. \square

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 $H_2 = 1$ 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 9. *If hybrid H_2 outputs 1, denote $\text{COM} = (c_1, c_2, \dots, c_\ell)$ in the accepting proof. Then $H_k(\text{COM}) = C_{\text{st}}(\text{COM})$.*

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

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 $c_i[s, t] = \text{Com}_2(1^\lambda, 1^\mu, \text{com}_1, \mathbf{b}', G_i[s, t]; r_i^{(s,t)})$ for each $s, t \in [n]$. From the extractability property of the commitment scheme and $\mathbf{b}' = \mathbf{b}'_0$, we have $\text{Dec}(1^\lambda, 1^\mu, \text{st}, c_i[s, t]) = G_i[s, t]$. Since G_i is a cycle graph, $C_{\text{st}}(c_i) = 0$. 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 $c_i[s, t] = \text{Com}_2(1^\lambda, 1^\mu, \text{com}_1, \mathbf{b}', 0; r_i^{(s,t)})$ for each $e \notin E, (s, t) = \phi(e)$. Let $G_i[s, t] = \text{Dec}(1^\lambda, 1^\mu, \text{st}, c_i[s, t])$ for each $s, t \in [n]$. Since $C_{\text{st}}(c_i) = 0$, 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 $c_i[s', t'] = \text{Com}_2(1^\lambda, 1^\mu, \text{com}_1, \mathbf{b}', 0; r_i^{(s', t')})$. From extractable property of commitment scheme, $\text{Dec}(1^\lambda, 1^\mu, \text{st}, c_i[s', t']) = 0$, 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 $G \notin L_{\text{HAM}}$. \square

Combining Lemmas 8 and Lemma 9, we derive that

$$\begin{aligned}
&\Pr[k \leftarrow \text{FakeGen}(1^\lambda, 1^{|C_{\text{st}}|}, C_{\text{st}}) : \exists \text{COM}, H_k(\text{COM}) = C_{\text{st}}(\text{COM})] \\
&\geq \lambda^{-c} 2^{-\mu-2} - 2 \cdot 2^{-\Omega(\log^4 \lambda)}
\end{aligned}$$

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

5 Statistical Hash Commitments

Intuitively speaking, a statistical hash commitment (SHC) scheme is a two-round *statistical hiding* commitment scheme, where the verification of the decommitment is a simple equality check with a hash output (computed w.r.t. a hashing algorithm associated with the scheme).

Definition 10. *A statistical hash commitment scheme is a tuple of algorithms $(\text{KGen}, \text{Com}, \text{H}, \mathcal{C}, \mathcal{R})$. It proceeds as follows.*

Round 1 \mathcal{R} executes $(\text{pk}, k) \leftarrow \text{KGen}(1^\lambda)$, and sends pk to \mathcal{C} .

Round 2 \mathcal{C} 's input is a bit $b \in \{0, 1\}$. Compute $(c, \rho) \leftarrow \text{Com}(\text{pk}, b)$ and send c to \mathcal{R} .

Opening \mathcal{C} sends (b, ρ) to the \mathcal{R} .

Verification \mathcal{R} accepts iff ρ is equal to $\text{H}(k, c, b)$.

We require the scheme to satisfy the following properties.

Completeness For any $b \in \{0, 1\}$, we have

$$\Pr [(\text{pk}, k) \leftarrow \text{KGen}(1^\lambda), (c, \rho) \leftarrow \text{Com}(\text{pk}, b) : \rho = \text{H}(k, c, b)] = 1$$

Computational Binding We say that the commitment scheme is computational binding, if for any non-uniform probabilistic polynomial time adversary \mathcal{A} , there exists a negligible function $\nu(\cdot)$ such that

$$\text{Adv}(\mathcal{A}) \triangleq \Pr \left[(\text{pk}, k) \leftarrow \text{KGen}(1^\lambda), (c, \rho_0, \rho_1) \leftarrow \mathcal{A}(1^\lambda, \text{pk}) : \begin{matrix} \rho_0 = \text{H}(k, c, 0) \\ \rho_1 = \text{H}(k, c, 1) \end{matrix} \right] < \nu(\lambda)$$

Statistical Hiding For any (maliciously generated) pk , there exists a negligible function $\nu(\lambda)$ such that $\text{SD}(c_0, c_1) \leq \nu(\lambda)$, where $(c_b, \rho_b) \leftarrow \text{Com}(\text{pk}, b)$ for every $b \in \{0, 1\}$. If $\nu(\lambda) = 0$, then we say that the scheme is perfectly hiding.

5.1 Construction from CDH

Let q be an integer, and $G = \langle g \rangle$ be a cyclic group generated by g of order q .

Construction. We describe our construction of the SHC scheme.

$\text{KGen}(1^\lambda)$ Randomly sample $s, t \leftarrow \mathbb{Z}_q$, and $x \leftarrow G$. Output $(\text{pk} = (x, g^s, x^s \cdot g^t), k = (s, t))$.

$\text{Com}(\text{pk}, b)$ Parse pk as $(x, a_1, a_2) \in G \times G$. Randomly sample $u, v \leftarrow \mathbb{Z}_q$. Output $(c = (g^u \cdot x^v, g^v \cdot g^b), \rho = a_1^u \cdot a_2^v)$.

$\text{H}(k, c, b)$ Parse c as $(z_1, z_2) \in G \times G$, and parse k as (s, t) . Output $z_1^s \cdot (z_2 \cdot g^{-b})^t$.

We now prove the properties of this construction. We defer the proof of completeness to the full version.

Lemma 10 (Computational Binding). *Assuming CDH, the above construction of SHC is computational binding.*

Proof. For any n.u. probabilistic polynomial time adversary \mathcal{A} , we construct the following adversary \mathcal{A}' for CDH problem.

Adversary $\mathcal{A}'(1^\lambda, g^s, g^y)$ Sample $u \leftarrow \mathbb{Z}_q$ uniformly at random. Set $x = g^y, \text{pk} = (x, g^s, g^u)$. Execute $(c, \rho_0, \rho_1) \leftarrow \mathcal{A}(1^\lambda, \text{pk})$. Output $g^u \cdot \rho_0^{-1} \cdot \rho_1$.

We now prove that $\Pr[a \leftarrow \mathcal{A}'(1^\lambda, g^s, g^y) : a = g^{sy}] \geq \text{Adv}(\mathcal{A})$. Since in our construction, $\text{pk} = (x, g^s, x^s \cdot g^t)$, where t is uniformly random. The second component of pk is uniformly random over G . Hence, the distributions of pk in real execution and the adversary \mathcal{A}' are identical.

Now for any $u \in \mathbb{Z}_q$, there exists a unique $t' \in \mathbb{Z}_q$ such that $x^s \cdot g^{t'} = g^u$. Then, for adversary \mathcal{A}' , we have

$$\begin{aligned} \Pr[a = g^{sy}] &= \Pr[g^u \cdot \rho_0^{-1} \cdot \rho_1 = g^{sy}] = \Pr[g^{t'} = \rho_0 \cdot \rho_1^{-1}] \\ &\geq \Pr[\rho_0 = \text{H}(k, c, 0) \wedge \rho_1 = \text{H}(k, c, 1)] = \text{Adv}(\mathcal{A}) \end{aligned}$$

where $k = (s, t')$. By the hardness of CDH, we conclude that $\text{Adv}(\mathcal{A})$ is negligible. \square

Lemma 11 (Perfect Hiding). *The Construction 5.1 is perfect hiding.*

Proof. For any fixed $\text{pk} = (x, a_1, a_2)$, since v is uniformly random, $g^v \cdot g^b$ is uniformly random. Furthermore, conditioned on $g^v \cdot g^b$, since u is uniformly random, $g^u \cdot x^v$ is also uniformly random. Hence, c is uniformly random over $G \times G$. \square

5.2 Construction from any 2-round Statistical Sender-Private OT

We now describe our construction of SHC from statistical sender-private OT. Let $\ell = \ell(\lambda)$ be a polynomial in λ , and let $(\text{OT}_1, \text{OT}_2, \text{OT}_3)$ be any statistical sender private 2-round OT scheme.

KGen (1^λ) Randomly sample $r \leftarrow \{0, 1\}^\ell$.

For $i \in [\ell]$, execute $(\text{ot}_{1,i}, \text{st}_i) \leftarrow \text{OT}_1(1^\lambda, r[i])$.

Output $\text{pk} = ((\text{ot}_{1,i})_{i \in [\ell]}, k = (\text{st}_i)_{i \in [\ell]})$.

Com $(\text{pk}, b \in \{0, 1\})$ Parse pk as $(\text{ot}_{1,i})_{i \in [\ell]}$. Randomly sample $r' \leftarrow \{0, 1\}^\ell$.

For $i \in [\ell]$, execute $\text{ot}_{2,i} \leftarrow \text{OT}_2(\text{ot}_{1,i}, r'[i], r'[i] \oplus b)$.

Output $(c = (\text{ot}_{2,i})_{i \in [\ell]}, \rho = r')$.

H (k, c, b) Parse $k = (\text{st}_i)_{i \in [\ell]}, c = (\text{ot}_{2,i})_{i \in [\ell]}$.

For $i \in [\ell]$, Let $\rho_{0,i} \leftarrow \text{OT}_3(\text{st}_i, \text{ot}_{2,i})$.

Let $\rho_b = (\rho_{0,i} \oplus (r[i] \cdot b))_{i \in [\ell]}$.

Output ρ_b .

We defer the proof of completeness and statistical hiding property to the full version. Below, we prove computational binding.

Lemma 12 (Computational Binding). *Assuming computational indistinguishability of OT_1 , the above construction of SHC is computational binding.*

Proof. For any PPT adversary \mathcal{A} trying to break the computational binding property, we construct the following hybrids.

Hybrid H_0 Randomly sample $r \leftarrow \{0, 1\}^\ell$. For $i \in [\ell]$, execute $(\text{ot}_{1,i}, \text{st}_i) \leftarrow \text{OT}_1(1^\lambda, r[i])$.

Let $\text{pk} = (\text{ot}_{1,i})_{i \in [\ell]}$. Execute $(c, \rho_0, \rho_1) \leftarrow \mathcal{A}(1^\lambda, \text{pk})$. If $\rho_0 \oplus \rho_1 = r$, then output 1, otherwise output 0.

Hybrid $H_{0.5}^{i^*}$ Randomly sample $r \leftarrow \{0, 1\}^\ell$. For $1 \leq i \leq i^*$, execute $(\text{ot}_{1,i}, \text{st}_i) \leftarrow \text{OT}_1(1^\lambda, 0)$.

For $i^* < i \leq \ell$, execute $(\text{ot}_{1,i}, \text{st}_i) \leftarrow \text{OT}_1(1^\lambda, r[i])$. Let $\text{pk} = (\text{ot}_{1,i})_{i \in [\ell]}$. Execute $(c, \rho_0, \rho_1) \leftarrow \mathcal{A}(1^\lambda, \text{pk})$. If $\rho_0 \oplus \rho_1 = r$, then output 1, otherwise output 0.

Hybrid H_1 Randomly sample $r \leftarrow \{0, 1\}^\ell$. For $i \in [\ell]$, execute $(\text{ot}_{1,i}, \text{st}_i) \leftarrow \text{OT}_1(1^\lambda, 0)$.

Let $\text{pk} = (\text{ot}_{1,i})_{i \in [\ell]}$. Execute $(c, \rho_0, \rho_1) \leftarrow \mathcal{A}(1^\lambda, \text{pk})$. If $\rho_0 \oplus \rho_1 = r$, then output 1, otherwise output 0.

Lemma 13. $\Pr[H_0 = 1] \geq \text{Adv}(\mathcal{A})$.

Proof. From the construction of H , we now that $H(k, c, 0) \oplus H(k, c, 1) = r$. Hence, when \mathcal{A} wins the security game, $(c, \rho_0, \rho_1) \leftarrow \mathcal{A}(1^\lambda, \text{pk})$ with $\rho_0 = H(k, x, 0) \wedge \rho_1 = H(k, x, 1)$ implies $\rho_0 \oplus \rho_1 = H(k, x, 0) \oplus H(k, x, 1) = r$. \square

Lemma 14. Hybrid H_0 and Hybrid $H_{0.5}^0$ are identical. Furthermore, there exists a negligible function $\nu(\lambda)$ such that for each $i = 0, \dots, \ell - 1$, $|\Pr[H_{0.5}^{i^*} = 1] - \Pr[H_{0.5}^{i^*+1} = 1]| < \nu(\lambda)$.

Proof. When $i^* = 0$, all $\text{ot}_{1,i}$ are generated in the same way as in Hybrid H_0 , for all $i \in [\ell]$. Hence, Hybrid H_0 and Hybrid $H_{0.5}^0$ are identical.

To show $H_{0.5}^{i^*} \approx H_{0.5}^{i^*+1}$, we consider the following adversary \mathcal{D} for receiver's computational privacy.

$\mathcal{D}(1^\lambda, \text{ot}_1)$ Randomly sample $r \leftarrow \{0, 1\}^\ell$. For $i \in [\ell] \setminus \{i^* + 1\}$, let $(\text{ot}_{1,i}, \text{st}_i) \leftarrow \text{OT}_1(1^\lambda, r[i])$. If $r[i^* + 1] = 0$, then let $(\text{ot}_{1,i^*+1}, \text{st}_{i^*+1}) \leftarrow \text{OT}_1(1^\lambda, 0)$, otherwise let $\text{ot}_{1,i^*+1} = \text{ot}_1$. Let $\text{pk} = (\text{ot}_{1,i})_{i \in [\ell]}$. Execute $(c, \rho_0, \rho_1) \leftarrow \mathcal{A}(1^\lambda, \text{pk})$. If $\rho_0 \oplus \rho_1 = r$, then output 1, otherwise output 0.

If ot_1 is generated from $\text{OT}_1(1^\lambda, 0)$, then \mathcal{D} simulates the environment of $H_{0.5}^{i^*+1}$ for \mathcal{A} . Hence, $\Pr[H_{0.5}^{i^*+1} = 1] = \Pr[(\text{ot}_1, \text{st}) \leftarrow \text{OT}_1(1^\lambda, 0) : \mathcal{D}(1^\lambda, \text{ot}_1) = 1]$.

If ot_1 is generated from $\text{OT}_1(1^\lambda, 1)$, then \mathcal{D} simulates the environment of $H_{0.5}^{i^*}$ for \mathcal{A} . Hence, $\Pr[H_{0.5}^{i^*} = 1] = \Pr[(\text{ot}_1, \text{st}) \leftarrow \text{OT}_1(1^\lambda, 1) : \mathcal{D}(1^\lambda, \text{ot}_1) = 1]$.

From the indistinguishability of ot_1 , we know that the right hand ot_1^0 generated by $\text{OT}_1(1^\lambda, 0)$ and ot_1^1 generated by $\text{OT}_1(1^\lambda, 1)$ are indistinguishable. Hence, there exists a negligible function $\nu(\lambda)$ such that $|\Pr[H_{0.5}^{i^*} = 1] - \Pr[H_{0.5}^{i^*+1} = 1]| < \nu(\lambda)$. \square

Lemma 15. Hybrid $H_{0.5}^\ell$ is identical to H_1 . Furthermore, $\Pr[H_1 = 1] = 1/2^\ell$.

Proof. When $i^* = \ell$, we know that all $\text{ot}_{1,i}$ are generated in the same way as in Hybrid H_1 . Hence, $H_{0.5}^\ell$ and H_1 are identical.

In Hybrid H_1 , pk is completely independent of r . Hence, $\Pr[H_1 = 1] = \Pr[\rho_0 \oplus \rho_1 = r] = 1/2^\ell$. \square

By the hybrid argument, combining Lemma 13, Lemma 14, and Lemma 15, we have $\text{Adv}(\mathcal{A}) < \text{neg}(\lambda)$. \square

We defer the proof of statistical hiding property to the full version.

6 Three Round Statistical Receiver-Private Oblivious Transfer

We start by presenting the definition for 3-round statistical receiver-private oblivious transfer. We capture statistical receiver privacy via a game-based definition. We consider two definitions to capture computational sender privacy: a game-based definition that intuitively requires that any malicious receiver who interacts with an honest sender can only learn one of its two inputs, and a distinguisher-dependent simulation based definition. We defer the formal treatment of the latter as well as the proof of implication from the former to the latter definition to the full version.

Definition 11 (3-round Statistical Receiver-Private Oblivious Transfer). *A 3-round oblivious transfer is a tuple of algorithms $(\text{OT}_1, \text{OT}_2, \text{OT}_3, \text{OT}_4)$, which specify the following protocol.*

Round 1 *The sender \mathcal{S} computes $(\text{ot}_1, \text{st}_S) \leftarrow \text{OT}_1(1^\lambda)$ and sends ot_1 to the receiver \mathcal{R} .*

Round 2 *The receiver \mathcal{R} with input $\beta \in \{0, 1\}$, computes $(\text{ot}_2, \text{st}_R) \leftarrow \text{OT}_2(1^\lambda, \text{ot}_1, \beta)$ and sends ot_2 to \mathcal{S} .*

Round 3 *\mathcal{S} with input $(m_0, m_1) \in \{0, 1\}^2$ computes $\text{ot}_3 \leftarrow \text{OT}_3(1^\lambda, \text{ot}_2, \text{st}_S, m_0, m_1)$ and sends ot_3 to the receiver.*

Message Decryption *The receiver computes $m' \leftarrow \text{OT}_4(1^\lambda, \text{ot}_1, \text{ot}_3, \text{st}_R)$.*

We require the protocol to satisfy the following properties.

Correctness ⁵ *For any $\beta \in \{0, 1\}$, $(m_0, m_1) \in \{0, 1\}^2$, we have*

$$\Pr \left[\begin{array}{l} (\text{ot}_1, \text{st}_S) \leftarrow \text{OT}_1(1^\lambda) \\ (\text{ot}_2, \text{st}_R) \leftarrow \text{OT}_2(1^\lambda, \text{ot}_1, \beta) \\ \text{ot}_3 \leftarrow \text{OT}_3(1^\lambda, \text{ot}_2, \text{st}_S, m_0, m_1) \\ m' \leftarrow \text{OT}_4(1^\lambda, \text{ot}_1, \text{ot}_3, \text{st}_R) \end{array} : m' = m_\beta \right] = 1$$

Game-Based Statistical Receiver-Privacy *For any (potentially maliciously generated) ot_1^* , denote $(\text{ot}_2^{(0)}, \text{st}_R^{(0)}) \leftarrow \text{OT}_2(1^\lambda, \text{ot}_1^*, 0)$, and $(\text{ot}_2^{(1)}, \text{st}_R^{(1)}) \leftarrow \text{OT}_2(1^\lambda, \text{ot}_1^*, 1)$.*

Then we have $\text{SD}(\text{ot}_2^{(0)}, \text{ot}_2^{(1)}) < \nu(\lambda)$, where $\nu(\cdot)$ is a negligible function.

Game-Based Computational Sender-Privacy *For any probabilistic polynomial time distinguisher $\mathcal{A}_0, \mathcal{A}_1$, and any probabilistic polynomial time malicious receiver \mathcal{R}^* , we define the following games.*

Interact with \mathcal{R}^* *The challenger plays the role of an honest sender for the first round and the second round with the malicious receiver \mathcal{R}^* . Specifically, the challenger executes $(\text{ot}_1, \text{st}_S) \leftarrow \text{OT}_1(1^\lambda)$. Then send ot_1 to \mathcal{R}^* . Then the receiver \mathcal{R}^* sends ot_2^* to the challenger.*

Game $G_0(m_0, m_1)$ *This game interact with adversary \mathcal{A}_0 . In the beginning, the adversary \mathcal{A}_0 is given input $\text{View}(\mathcal{R}^*)$. Then the challenger samples $b_0 \leftarrow \{0, 1\}$ at random, and send $\text{ot}_3 \leftarrow \text{OT}_3(1^\lambda, \text{ot}_2^*, \text{st}_S, m_b, m_1)$ to \mathcal{A}_0 . Finally \mathcal{A}_0 outputs a bit b'_0 . If $b_0 = b'_0$, then we say \mathcal{A}_0 wins the game.*

⁵ We can relax the definition to be statistical correctness, which only requires the probability to be $1 - \text{negl}(\lambda)$

Game $G_1(m_0, m_1)$ This game interact with adversary \mathcal{A}_1 . In the beginning, the adversary \mathcal{A}_1 is given input $\text{View}(\mathcal{R}^*)$. Then the challenger samples $b_1 \leftarrow \{0, 1\}$ at random, and send $\text{ot}_3 \leftarrow \text{OT}_3(1^\lambda, \text{ot}_2^*, \text{st}_S, m_0, m_b)$ to \mathcal{A}_1 . Finally \mathcal{A}_1 outputs a bit b'_1 . If $b_1 = b'_1$, then we say \mathcal{A}_1 wins the game. We define the following advantage

$$\text{Adv}(\mathcal{A}_0, \mathcal{A}_1, \mathcal{R}^*) \triangleq \mathbb{E}_{\text{View}(\mathcal{R}^*)} \left[\min \left\{ \begin{aligned} & \max_{m_0, m_1 \in \{0, 1\}} \left(\left| \Pr[\mathcal{A}_0(\text{View}(\mathcal{R}^*)) \text{ wins } G_0(m_0, m_1)] - \frac{1}{2} \right| \right), \\ & \max_{m_0, m_1 \in \{0, 1\}} \left(\left| \Pr[\mathcal{A}_1(\text{View}(\mathcal{R}^*)) \text{ wins } G_1(m_0, m_1)] - \frac{1}{2} \right| \right) \right\} \right]$$

We say the oblivious transfer scheme is game-based computational sender-secure, if for any probabilistic polynomial time distinguisher $\mathcal{A}_0, \mathcal{A}_1$, and any probabilistic polynomial time malicious receiver \mathcal{R}^* , there exist a negligible function $\nu(\cdot)$ such that $\text{Adv}(\mathcal{A}_0, \mathcal{A}_1, \mathcal{R}^*) < \nu(\lambda)$.

6.1 Our Construction

We now describe a generic transformation from SHC scheme to three-round statistical receiver-private oblivious transfer.

Construction. Let $(\text{KGen}, \text{Com}, \text{H}, \mathcal{C}, \mathcal{R})$ be an SHC scheme. Let hc denote the Goldreich-Levin hardcore predicate [21]. The 3-round statistical receiver-private oblivious transfer proceeds as follows.

- $\text{OT}_1(1^\lambda)$ Execute $(\text{pk}, \text{k}) \leftarrow \text{KGen}(1^\lambda)$. Let $\text{ot}_1 = \text{pk}, \text{st}_S = \text{k}$.
- $\text{OT}_2(1^\lambda, \text{ot}_1, \beta)$ Parse $\text{ot}_1 = \text{pk}$. Run $(c, \rho) \leftarrow \text{Com}(\text{pk}, \beta)$. Output $\text{ot}_2 = c, \text{st}_R = \rho$.
- $\text{OT}_3(1^\lambda, \text{ot}_2, \text{st}_S, m_0, m_1)$ Parse $\text{ot}_2 = c$, and $\text{st}_S = \text{k}$. For any $b \in \{0, 1\}$, sample $r_b \leftarrow \{0, 1\}^\lambda$, encrypt m_b as $c_b = (\text{hc}(\text{H}(\text{k}, c, b), r_b) \oplus m_b, r_b)$. Output $\text{ot}_3 = (c_0, c_1)$.
- $\text{OT}_4(1^\lambda, \text{ot}_1, \text{ot}_3, \text{st}_R)$ Parse $\text{ot}_1 = \text{pk}, \text{ot}_3 = (c_0, c_1)$, and $\text{st}_R = \rho$. Parse c_β as $c_\beta = (u_\beta, r_\beta)$. Output $m' = u_\beta \oplus \text{hc}(\rho, r_\beta)$.

We now prove the required properties of the protocol. We defer the proof of correctness to the full version.

Lemma 16 (Statistical Receiver-Privacy). *If the underlying SHC is statistical (resp. perfect) hiding, then the construction above is statistical (resp. perfect) receiver-private.*

Proof. From the statistical hiding property of the SHC scheme, for any pk , we have $\text{SD}(\text{ot}_2^0, \text{ot}_2^1) \leq \text{neg}(\lambda)$, where $(\text{ot}_2^b, \rho^b) \leftarrow \text{Com}(\text{pk}, b)$ for any $b \in \{0, 1\}$. Hence, for any ot_1 , $\text{OT}_2(1^\lambda, \text{ot}_1, 0)$ and $\text{OT}_2(1^\lambda, \text{ot}_1, 1)$ are statistically (resp. perfectly) close. \square

Lemma 17 (Game-based Computational Sender-Privacy). *If the underlying SHC scheme is computational binding, then the 3-round oblivious transfer constructed above is game-based computational sender-private.*

Proof. For any probabilistic polynomial time adversary $\mathcal{A}_0, \mathcal{A}_1$ and any probabilistic polynomial time malicious receiver \mathcal{R}^* with $\text{Adv}(\mathcal{A}_0, \mathcal{A}_1, \mathcal{R}^*) > \delta$, where δ is a non-negligible function of λ . Then, with probability at least $\delta/2$ over $\text{View}(\mathcal{R}^*)$,

$$\exists \mathbf{m}_0 \in \{0, 1\}^2, \mathbf{m}_1 \in \{0, 1\}^2 : \left| \Pr[\mathcal{A}_0(\text{View}(\mathcal{R}^*)) \text{ wins } G_0(\mathbf{m}_0)] - \frac{1}{2} \right| > \frac{\delta}{2} \wedge \left| \Pr[\mathcal{A}_1(\text{View}(\mathcal{R}^*)) \text{ wins } G_1(\mathbf{m}_1)] - \frac{1}{2} \right| > \frac{\delta}{2}$$

Denote this fraction of $\text{View}(\mathcal{R}^*)$ as GOOD. Randomly sample $\bar{\mathbf{m}}_0, \bar{\mathbf{m}}_1 \leftarrow \{0, 1\}^2$. With probability $1/16$, we have $\bar{\mathbf{m}}_0 = \mathbf{m}_0 \wedge \bar{\mathbf{m}}_1 = \mathbf{m}_1$.

From Goldreich-Levin Theorem [21], there exists two inverters $\mathcal{A}'_0, \mathcal{A}'_1$ such that \mathcal{A}'_0 takes input $(\text{View}(\mathcal{R}^*), r_0, \text{hc}(\text{H}(\mathbf{k}, c, 1), r_1) \oplus m_1, r_1)$, output x'_0 . \mathcal{A}'_1 takes input $(\text{View}(\mathcal{R}^*), r_1, \text{hc}(\text{H}(\mathbf{k}, c, 0), r_0) \oplus m_0, r_0)$, output x'_1 . Furthermore, the inverters $\mathcal{A}'_0, \mathcal{A}'_1$ satisfy the property that for any $\mathbf{v} \in \text{GOOD}$ and $\bar{\mathbf{m}}_0 = \mathbf{m}_0 \wedge \bar{\mathbf{m}}_1 = \mathbf{m}_1$, $\Pr[x'_0 = \text{H}(\mathbf{k}, c, 0)] > \delta'$ and $\Pr[x'_1 = \text{H}(\mathbf{k}, c, 1)] > \delta'$, where $\delta' = \delta'(\lambda)$ is a non-negligible function. We construct the following adversary \mathcal{A} to attack the computational binding property of the SHC scheme.

Adversary $\mathcal{A}(1^\lambda, \text{pk})$ Set random coins and execute \mathcal{R}^* . Send \mathcal{R}^* the first round message $\text{ot}_1 = \text{pk}$, then \mathcal{R}^* replies ot_2^* . Sample $r_0 \leftarrow \{0, 1\}^\lambda, b_1 \leftarrow \{0, 1\}, r_1 \leftarrow \{0, 1\}^\lambda$, then execute $x'_0 \leftarrow \mathcal{A}'_0(\text{View}(\mathcal{R}^*), r_0, b_1, r_1)$. Sample $r'_1 \leftarrow \{0, 1\}^\lambda, b_0 \leftarrow \{0, 1\}, r'_0 \leftarrow \{0, 1\}^\lambda$, then execute $x'_1 \leftarrow \mathcal{A}'_1(\text{View}(\mathcal{R}^*), r'_1, b_0, r'_0)$. Output $(c = \text{ot}_2^*, x'_0, x'_1)$. We now prove that the advantage of \mathcal{A} satisfies

$$\text{Adv}(\mathcal{A}) = \Pr \left[(\text{pk}, \mathbf{k}) \leftarrow \text{KGen}(1^\lambda), (c, \rho_0, \rho_1) \leftarrow \mathcal{A}(1^\lambda, \text{pk}) : \begin{matrix} \rho_0 = \text{H}(\mathbf{k}, c, 0) \wedge \\ \rho_1 = \text{H}(\mathbf{k}, c, 1) \end{matrix} \right] \geq \frac{\delta \cdot \delta'^2}{128}$$

Hybrids H_0 $(\text{pk}, \mathbf{k}) \leftarrow \text{KGen}(1^\lambda)$. Set random coins and execute \mathcal{R}^* . \mathcal{R}^* replies ot_2^* . Sample $r_0 \leftarrow \{0, 1\}^\lambda, r_1 \leftarrow \{0, 1\}^\lambda$. Let $b_1 = \text{hc}(\text{H}(\mathbf{k}, c, 1), r_1) \oplus m_1$. Execute $x'_0 \leftarrow \mathcal{A}'_0(\text{View}(\mathcal{R}^*), r_0, b_1, r_1)$. Sample $r'_0 \leftarrow \{0, 1\}^\lambda, r'_1 \leftarrow \{0, 1\}^\lambda$. Let $b_0 = \text{hc}(\text{H}(\mathbf{k}, c, 0), r'_0) \oplus m_0$. Execute $x'_1 \leftarrow \mathcal{A}'_1(\text{View}(\mathcal{R}^*), r'_1, b_0, r'_0)$. If $\rho_0 = \text{H}(\mathbf{k}, c, 0) \wedge \rho_1 = \text{H}(\mathbf{k}, c, 1)$, then output 1; else output 0.

Hybrids H_1 $(\text{pk}, \mathbf{k}) \leftarrow \text{KGen}(1^\lambda)$. Set random coins and execute \mathcal{R}^* . \mathcal{R}^* replies ot_2^* . Sample $r_0 \leftarrow \{0, 1\}^\lambda, r_1 \leftarrow \{0, 1\}^\lambda$. Let $b_1 \leftarrow \{0, 1\}$. Execute $x'_0 \leftarrow \mathcal{A}'_0(\text{View}(\mathcal{R}^*), r_0, b_1, r_1)$. Sample $r'_0 \leftarrow \{0, 1\}^\lambda, r'_1 \leftarrow \{0, 1\}^\lambda$. Let $b_0 \leftarrow \{0, 1\}$. Execute $x'_1 \leftarrow \mathcal{A}'_1(\text{View}(\mathcal{R}^*), r'_1, b_0, r'_0)$. If $\rho_0 = \text{H}(\mathbf{k}, c, 0) \wedge \rho_1 = \text{H}(\mathbf{k}, c, 1)$, then output 1; else output 0.

Hybrids H_2 $(\text{pk}, \mathbf{k}) \leftarrow \text{KGen}(1^\lambda), (c, \rho_0, \rho_1) \leftarrow \mathcal{A}(1^\lambda, \text{pk})$. If $\rho_0 = \text{H}(\mathbf{k}, c, 0) \wedge \rho_1 = \text{H}(\mathbf{k}, c, 1)$, then output 1; else output 0.

From the construction of \mathcal{A} , the hybrids H_1 and H_2 are identical. Hence, $\text{Adv}(\mathcal{A}) = \Pr[\text{H}_2 = 1] = \Pr[\text{H}_1 = 1]$. Furthermore, in hybrids H_1 , with probability $1/4$, $b_1 = \text{hc}(\text{H}(\mathbf{k}, c, 1), r_1) \oplus m_1 \wedge b_0 = \text{hc}(\text{H}(\mathbf{k}, c, 0), r'_0) \oplus m_0$. Conditioned on such event, H_0 and H_1 are identical. Hence, $\Pr[\text{H}_1 = 1] \geq \Pr[\text{H}_0 = 1]/4$. In hybrid H_0 , the fraction of $\text{View}(\mathcal{R}^*) \in \text{GOOD}$ is at least $\delta/2$. With probability $1/16$, the guess of $\mathbf{m}_0, \mathbf{m}_1$ is correct. With probability δ'^2 , both \mathcal{A}'_0 and \mathcal{A}'_1 invert correctly. Hence, $\text{Adv}(\mathcal{A}) \geq$

$\frac{\delta}{2} \cdot \frac{1}{16} \cdot \delta'^2 \cdot \frac{1}{4} = \delta \cdot \delta'^2 / 128$. If $\delta(\lambda)$ is non-negligible, then $\text{Adv}(\mathcal{A})$ is also non-negligible. This contradicts with the computational binding property of the SHC scheme. \square

Acknowledgement. The first author was supported in part by the NSF award 1916939, a gift from Ripple, a JP Morgan Faculty Fellowship, a PNC center for financial services innovation award, and a Cylab seed funding award. The second and third author were supported in part by NSF SaTC award 1814919 and DARPA Safeware W911NF-15-C-0213. The last author conducted part of the research while at the Simons Institute for the Theory of Computing.

References

1. Aiello, W., Ishai, Y., Reingold, O.: Priced oblivious transfer: How to sell digital goods. In: Pfitzmann, B. (ed.) EUROCRYPT 2001. LNCS, vol. 2045, pp. 119–135. Springer, Heidelberg, Germany, Innsbruck, Austria (May 6–10, 2001). https://doi.org/10.1007/3-540-44987-6_8
2. Babai, L.: Trading group theory for randomness. In: 17th ACM STOC. pp. 421–429. ACM Press, Providence, RI, USA (May 6–8, 1985). <https://doi.org/10.1145/22145.22192>
3. Badrinarayanan, S., Fernando, R., Jain, A., Khurana, D., Sahai, A.: Statistical ZAP arguments. EUROCRYPT (2020)
4. Badrinarayanan, S., Garg, S., Ishai, Y., Sahai, A., Wadia, A.: Two-message witness indistinguishability and secure computation in the plain model from new assumptions. In: Takagi, T., Peyrin, T. (eds.) ASIACRYPT 2017, Part III. LNCS, vol. 10626, pp. 275–303. Springer, Heidelberg, Germany, Hong Kong, China (Dec 3–7, 2017). https://doi.org/10.1007/978-3-319-70700-6_10
5. Barak, B., Ong, S.J., Vadhan, S.P.: Derandomization in cryptography. In: Boneh, D. (ed.) CRYPTO 2003. LNCS, vol. 2729, pp. 299–315. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 17–21, 2003). https://doi.org/10.1007/978-3-540-45146-4_18
6. Bitansky, N., Paneth, O.: ZAPs and non-interactive witness indistinguishability from indistinguishability obfuscation. In: Dodis, Y., Nielsen, J.B. (eds.) TCC 2015, Part II. LNCS, vol. 9015, pp. 401–427. Springer, Heidelberg, Germany, Warsaw, Poland (Mar 23–25, 2015). https://doi.org/10.1007/978-3-662-46497-7_16
7. Brakerski, Z., Döttling, N.: Two-message statistically sender-private OT from LWE. In: Beimel, A., Dziembowski, S. (eds.) TCC 2018, Part II. LNCS, vol. 11240, pp. 370–390. Springer, Heidelberg, Germany, Panaji, India (Nov 11–14, 2018). https://doi.org/10.1007/978-3-030-03810-6_14
8. Canetti, R., Chen, Y., Holmgren, J., Lombardi, A., Rothblum, G.N., Rothblum, R.D., Wichs, D.: Fiat-Shamir: from practice to theory. In: Charikar, M., Cohen, E. (eds.) 51st ACM STOC. pp. 1082–1090. ACM Press, Phoenix, AZ, USA (Jun 23–26, 2019). <https://doi.org/10.1145/3313276.3316380>
9. Canetti, R., Goldreich, O., Halevi, S.: The random oracle methodology, revisited (preliminary version). In: 30th ACM STOC. pp. 209–218. ACM Press, Dallas, TX, USA (May 23–26, 1998). <https://doi.org/10.1145/276698.276741>
10. Cramer, R., Damgård, I., Schoenmakers, B.: Proofs of partial knowledge and simplified design of witness hiding protocols. In: Desmedt, Y. (ed.) CRYPTO'94. LNCS, vol. 839, pp. 174–187. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 21–25, 1994). https://doi.org/10.1007/3-540-48658-5_19

11. Cramer, R., Shoup, V.: Universal hash proofs and a paradigm for adaptive chosen ciphertext secure public-key encryption. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 45–64. Springer, Heidelberg, Germany, Amsterdam, The Netherlands (Apr 28 – May 2, 2002). https://doi.org/10.1007/3-540-46035-7_4
12. De Santis, A., Micali, S., Persiano, G.: Non-interactive zero-knowledge proof systems. In: Pomerance, C. (ed.) CRYPTO'87. LNCS, vol. 293, pp. 52–72. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 16–20, 1988). https://doi.org/10.1007/3-540-48184-2_5
13. Deng, Y., Goyal, V., Sahai, A.: Resolving the simultaneous resettability conjecture and a new non-black-box simulation strategy. In: 50th FOCS. pp. 251–260. IEEE Computer Society Press, Atlanta, GA, USA (Oct 25–27, 2009). <https://doi.org/10.1109/FOCS.2009.59>
14. Döttling, N., Garg, S., Hajiabadi, M., Masny, D., Wichs, D.: Two-round oblivious transfer from CDH or LPN. IACR Cryptology ePrint Archive **2019**, 414 (2019)
15. Dwork, C., Naor, M.: Zaps and their applications. In: 41st FOCS. pp. 283–293. IEEE Computer Society Press, Redondo Beach, CA, USA (Nov 12–14, 2000). <https://doi.org/10.1109/SFCS.2000.892117>
16. Dwork, C., Naor, M., Reingold, O., Stockmeyer, L.J.: Magic functions. In: 40th FOCS. pp. 523–534. IEEE Computer Society Press, New York, NY, USA (Oct 17–19, 1999). <https://doi.org/10.1109/SFCS.1999.814626>
17. Even, S., Goldreich, O., Lempel, A.: A randomized protocol for signing contracts. Commun. ACM **28**(6), 637–647 (1985)
18. Feige, U., Lapidot, D., Shamir, A.: Multiple non-interactive zero knowledge proofs based on a single random string (extended abstract). In: 31st FOCS. pp. 308–317. IEEE Computer Society Press, St. Louis, MO, USA (Oct 22–24, 1990). <https://doi.org/10.1109/FSCS.1990.89549>
19. Fiat, A., Shamir, A.: How to prove yourself: Practical solutions to identification and signature problems. In: Odlyzko, A.M. (ed.) CRYPTO'86. LNCS, vol. 263, pp. 186–194. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 1987). https://doi.org/10.1007/3-540-47721-7_12
20. Goldreich, O., Krawczyk, H.: On the composition of zero-knowledge proof systems. SIAM J. Comput. **25**(1), 169–192 (1996)
21. Goldreich, O., Levin, L.A.: A hard-core predicate for all one-way functions. In: 21st ACM STOC. pp. 25–32. ACM Press, Seattle, WA, USA (May 15–17, 1989). <https://doi.org/10.1145/73007.73010>
22. Goldreich, O., Micali, S., Wigderson, A.: How to play any mental game or A completeness theorem for protocols with honest majority. In: Aho, A. (ed.) 19th ACM STOC. pp. 218–229. ACM Press, New York City, NY, USA (May 25–27, 1987). <https://doi.org/10.1145/28395.28420>
23. Goldreich, O., Oren, Y.: Definitions and properties of zero-knowledge proof systems. Journal of Cryptology **7**(1), 1–32 (Dec 1994). <https://doi.org/10.1007/BF00195207>
24. Goldwasser, S., Micali, S., Rackoff, C.: The knowledge complexity of interactive proof-systems (extended abstract). In: 17th ACM STOC. pp. 291–304. ACM Press, Providence, RI, USA (May 6–8, 1985). <https://doi.org/10.1145/22145.22178>
25. Groth, J., Ostrovsky, R., Sahai, A.: Non-interactive zaps and new techniques for NIZK. In: Dwork, C. (ed.) CRYPTO 2006. LNCS, vol. 4117, pp. 97–111. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 20–24, 2006). https://doi.org/10.1007/11818175_6
26. Groth, J., Ostrovsky, R., Sahai, A.: Perfect non-interactive zero knowledge for NP. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 339–358. Springer, Heidelberg, Germany, St. Petersburg, Russia (May 28 – Jun 1, 2006). https://doi.org/10.1007/11761679_21

27. Halevi, S., Hazay, C., Polychroniadou, A., Venkatasubramanian, M.: Round-optimal secure multi-party computation. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018, Part II. LNCS, vol. 10992, pp. 488–520. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 19–23, 2018). https://doi.org/10.1007/978-3-319-96881-0_17
28. Halevi, S., Kalai, Y.T.: Smooth projective hashing and two-message oblivious transfer. *Journal of Cryptology* **25**(1), 158–193 (Jan 2012). <https://doi.org/10.1007/s00145-010-9092-8>
29. Jain, A., Kalai, Y.T., Khurana, D., Rothblum, R.: Distinguisher-dependent simulation in two rounds and its applications. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017, Part II. LNCS, vol. 10402, pp. 158–189. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 20–24, 2017). https://doi.org/10.1007/978-3-319-63715-0_6
30. Kalai, Y.T.: Smooth projective hashing and two-message oblivious transfer. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 78–95. Springer, Heidelberg, Germany, Aarhus, Denmark (May 22–26, 2005). https://doi.org/10.1007/11426639_5
31. Kalai, Y.T., Khurana, D., Sahai, A.: Statistical witness indistinguishability (and more) in two messages. In: Nielsen, J.B., Rijmen, V. (eds.) EUROCRYPT 2018, Part III. LNCS, vol. 10822, pp. 34–65. Springer, Heidelberg, Germany, Tel Aviv, Israel (Apr 29 – May 3, 2018). https://doi.org/10.1007/978-3-319-78372-7_2
32. Khurana, D., Sahai, A.: How to achieve non-malleability in one or two rounds. In: Umans, C. (ed.) 58th FOCS. pp. 564–575. IEEE Computer Society Press, Berkeley, CA, USA (Oct 15–17, 2017). <https://doi.org/10.1109/FOCS.2017.58>
33. Kilian, J.: Founding cryptography on oblivious transfer. In: 20th ACM STOC. pp. 20–31. ACM Press, Chicago, IL, USA (May 2–4, 1988). <https://doi.org/10.1145/62212.62215>
34. Lombardi, A., Vaikuntanathan, V., Wichs, D.: 2-message publicly verifiable WI from (subexponential) LWE. *IACR Cryptology ePrint Archive* **2019**, 808 (2019)
35. Lombardi, A., Vaikuntanathan, V., Wichs, D.: Statistical zap arguments from bilinear maps. EUROCRYPT (2020)
36. Naor, M., Pinkas, B.: Efficient oblivious transfer protocols. In: Kosaraju, S.R. (ed.) 12th SODA. pp. 448–457. ACM-SIAM, Washington, DC, USA (Jan 7–9, 2001)
37. Pass, R.: Simulation in quasi-polynomial time, and its application to protocol composition. In: Biham, E. (ed.) EUROCRYPT 2003. LNCS, vol. 2656, pp. 160–176. Springer, Heidelberg, Germany, Warsaw, Poland (May 4–8, 2003). https://doi.org/10.1007/3-540-39200-9_10
38. Peikert, C., Shiehian, S.: Noninteractive zero knowledge for NP from (plain) learning with errors. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019, Part I. LNCS, vol. 11692, pp. 89–114. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 18–22, 2019). https://doi.org/10.1007/978-3-030-26948-7_4
39. Rabin, M.O.: How to exchange secrets by oblivious transfer. Technical Report TR-81, Harvard University (1981)
40. Wolf, S., Wullschleger, J.: Oblivious transfer is symmetric. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 222–232. Springer, Heidelberg, Germany, St. Petersburg, Russia (May 28 – Jun 1, 2006). https://doi.org/10.1007/11761679_14
41. Yao, A.C.C.: How to generate and exchange secrets (extended abstract). In: 27th FOCS. pp. 162–167. IEEE Computer Society Press, Toronto, Ontario, Canada (Oct 27–29, 1986). <https://doi.org/10.1109/SFCS.1986.25>