

Statistical ZAP Arguments

Saikrishna Badrinarayanan* Rex Fernando† Aayush Jain† Dakshita Khurana‡
Amit Sahai†

Abstract

Dwork and Naor (FOCS'00) first introduced and constructed two message public coin witness indistinguishable proofs (ZAPs) for NP based on trapdoor permutations. Since then, ZAPs have also been obtained based on the decisional linear assumption on bilinear maps, and indistinguishability obfuscation, and have proven extremely useful in the design of several cryptographic primitives.

However, all known constructions of two-message public coin (or even publicly verifiable) proof systems only guarantee witness indistinguishability against computationally bounded verifiers. In this paper, we construct the first public coin two message witness indistinguishable (WI) arguments for NP with *statistical* privacy, assuming the learning with errors (LWE) assumption holds with an explicit, efficiently computable upper bound on the adversary's advantage. Prior to this, there were no known constructions of two-message publicly verifiable WI protocols under lattice assumptions, even satisfying the weaker notion of computational witness indistinguishability.

*VISA Research.

†University of California Los Angeles.

‡University of Illinois Urbana-Champaign.

Contents

1	Introduction	1
1.1	Our Results	2
2	Overview of Techniques	2
2.1	A Simple Two-Message Public-Coin Computational WI Argument	3
2.2	Using Correlation-Intractable Hashing with Statistically Hiding Extractable Commitments	5
2.3	Statistically Hiding Extractable Commitments	5
2.4	Statistical ZAP Arguments	6
2.5	Organization	8
3	Preliminaries	8
3.1	Correlation Intractable Hash Functions	9
3.2	Oblivious Transfer	9
3.3	Proof Systems	10
4	Extractable Commitments	11
4.1	Definitions	11
4.2	Protocol	14
5	Our Statistical WI Protocol	19
5.1	Modified Blum Protocol	19
5.2	Statistical ZAPs	21
6	Statistical ZK Arguments with Super-Polynomial Simulation	26
	References	37

1 Introduction

Witness indistinguishability (WI) is one of the most widely used notions of privacy for proof systems. Informally, WI protocols [FS90] allow a prover to convince a verifier that some statement X belongs to an NP language L , with the following privacy guarantee: if there are two witnesses w_0, w_1 that both attest to the fact that $X \in L$, then a verifier should not be able to distinguish an honest prover using witness w_0 from an honest prover using witness w_1 . WI is a relaxation of zero-knowledge and has proven to be surprisingly useful. Since WI is a relaxation, unlike zero-knowledge, there are no known lower bounds on the rounds of interaction needed to build WI protocols in the plain model.

Indeed, Dwork and Naor [DN00, DN07] introduced the notion of two-message public-coin witness indistinguishable proofs (ZAPs) without any setup assumptions, and also constructed it assuming trapdoor permutations. We observe that the public-coin feature of ZAPs yield public verifiability of the resulting proof system, since a third party can use the public coins of the verifier to determine whether or not the prover’s response constitutes a valid proof. Subsequently, Groth et al. [GOS06] constructed ZAPs assuming the decisional linear assumption, and Bitansky and Paneth [BP15] constructed ZAPs from indistinguishability obfuscation and one way functions.

Our goal: ZAPs with statistical privacy. As originally introduced, ZAPs satisfied soundness against unbounded provers (i.e. were *proofs*), and witness indistinguishability against computationally bounded verifiers. In this work, we examine whether these requirements can be reversed: can we achieve witness indistinguishability against computationally unbounded verifiers, while achieving soundness against computationally bounded cheating provers? We call such objects *statistical ZAP arguments*.

An analogue of this question has a long history of study in the context of zero-knowledge protocols. Indeed, zero-knowledge protocols for NP were originally achieved guaranteeing privacy to hold only against computationally bounded verifiers [GMW86]. In the case of zero-knowledge, the notion of *statistical* zero-knowledge arguments was achieved soon after [BCR86, Cha86], that strengthened the privacy requirement to hold against computationally unbounded verifiers, while requiring soundness to hold only against computationally bounded provers.

Because ZAPs require a single message each from the verifier and the prover, a better comparison would perhaps be to non-interactive zero-knowledge (NIZK) [BFM88]. Even in the case of NIZKs, we have had arguments for NP satisfying statistical zero-knowledge since 2006 [GOS06]. And yet, the following natural question has remained open since the introduction of ZAPs nearly two decades ago.

Do there exist statistical ZAP arguments for NP in the plain model?

Statistical witness indistinguishability, just like its zero-knowledge counterpart, guarantees *everlasting privacy* against malicious verifiers, long after protocols have completed execution. Of course, to achieve statistical privacy, we must necessarily sacrifice soundness against unbounded provers. But such a tradeoff could often be desirable, since soundness is usually necessary only in an online setting: in order to convince a verifier of a false statement, a cheating prover must find a way to cheat *during* the execution of the protocol.

The main challenge: achieving a public-coin protocol. The recent work of Kalai et al. [KKS18] constructed the first two message statistically witness indistinguishable arguments in the plain model under standard sub-exponential assumptions. However, their arguments are only *privately verifiable*.

The blueprint of [KKS18], which builds on other similar approaches in the computational witness indistinguishability setting [BGI⁺17, JKKR17], uses oblivious transfer (OT) to reduce interaction

in a Σ -protocol. In all these approaches, the verifier obtains the third message of the Σ -protocol via the output of the OT, and therefore these approaches fundamentally require the use of private coins for verification. It is also worth noting that these protocols are not sound against provers that have access to the private coins of the verifier, which restricts their applicability. Additionally, the verifier’s message is *not reusable*, which means that soundness is not guaranteed if the same verifier message is reused across multiple executions.

On the other hand, a *public coin* argument, which is the focus of this work, does not suffer from any of these limitations. In fact, where the verifier’s message only needs to be a uniformly random string. Such a string can easily be generated, for example, via an MPC protocol, and can then be reused across multiple executions with no loss in soundness.

We stress that prior to our work, even two message statistically witness indistinguishable arguments that were only publicly verifiable (and not necessarily public coin) were not known from any falsifiable assumptions.

1.1 Our Results

In this paper, we construct the first two message public coin statistically witness indistinguishable arguments for NP in the plain model. Our constructions assume *explicit* hardness of the learning with errors (LWE) problem, where there is an efficiently computable upper bound μ_{LWE} on the advantage of any polynomial-time adversary. In fact, these are the first known two-message public coin (or even publicly verifiable) arguments based on lattice assumptions, satisfying *any notion of witness indistinguishability* (computational/statistical). We provide an informal theorem below.

Informal Theorem 1. *Assuming explicit hardness of the learning with errors (LWE) assumption, there exist two message public-coin statistically witness indistinguishable arguments for NP in the plain model.*

Our results are obtained by combining two recent results in a new way: recent constructions of correlation-intractable hash functions based on LWE [CCH⁺19] and the statistically hiding extractable commitments of [KKS18] (which are built upon [KS17]). This yields a new method of using correlation intractable hash functions to instantiate the Fiat-Shamir transform, by extracting messages from *statistically hiding commitments*, instead of from statistically binding trapdoor commitments – that we believe may be of independent interest.

2 Overview of Techniques

In this section, we provide a brief overview of the techniques we use to build a two message public coin statistical WI argument (henceforth referred to as a ZAP).

Our starting point is the popular technique to construct ZAPs for NP, due to Dwork and Naor [DN02]. Their construction makes use of a statistically sound NIZK in the common random string model, and can be described as follows.

- In the first round, the verifier picks uniformly random strings $\text{crs}_1, \dots, \text{crs}_\lambda$, where λ denotes the security parameter, and sends them to the prover.
- In the second round, the prover samples a uniformly random string crs' . It computes proofs (π_1, \dots, π_ℓ) where π_i is a NIZK proof for the instance x that verifies under $\text{crs}'_i = \text{crs}' \oplus \text{crs}_i$. The prover sends crs' along with proof strings (π_1, \dots, π_ℓ) to the verifier.

The soundness of this protocol can be proven based on the statistical soundness of NIZK, in the following way. Fix an instance $x \notin L$. Statistical soundness of the NIZK implies that with probability at least $1/2$ over the choice of crs from the domain of the common random string of NIZK, *there does not exist a proof π that verifies for instance x with respect to crs* . Put another way, for fixed x , for at least $1/2$ of the strings in the domain of the common random string of the NIZK, *there does not exist a proof for x* . One can use this fact to argue combinatorially that over the choice of random $\text{crs}_1, \dots, \text{crs}_\lambda$, the probability that there exists crs' for which there exist proofs with respect to every member of the set $\{\text{crs}'_i = \text{crs}' \oplus \text{crs}_i\}_{i \in [\lambda]}$, is negligible.

The proof of witness indistinguishability follows quite simply, by switching the witness in each of the proofs one by one.

But when applied to our context, this approach immediately encounters the following problems.

1. The soundness argument outlined above crucially requires that with high probability over the CRS of the NIZK, there just should not exist a proof for any fixed false instance. This translates to requiring *statistical soundness* of the underlying NIZK.
2. One cannot hope to get a WI argument secure against unbounded verifiers via this transform, unless the underlying NIZK also satisfies privacy against unbounded verifiers, i.e. satisfies statistical zero-knowledge.
3. It is believed that statistically sound and statistical zero-knowledge NIZKs for all of NP cannot exist.
4. Even if we only desired *computational* witness indistinguishability based on lattice assumptions, no statistically sound NIZKs in the common random string model are known from lattice assumptions.

As an intermediate objective, we will first try to tackle problem #4 and build a publicly verifiable computational WI argument based on LWE.

2.1 A Simple Two-Message Public-Coin Computational WI Argument

We make a few modifications to the template above so as to obtain a publicly verifiable computational WI argument based on LWE.

Before we describe these modifications, we list a few ingredients. We will assume that there exists a dense public key encryption scheme PKE, that is, a scheme for which every string in $\{0, 1\}^{pk}$ corresponds to a valid public key (and therefore every string has a valid secret key). We will further assume the existence of a correlation intractable hash function family. Informally, a hash function family \mathcal{H} is correlation-intractable for a function family \mathcal{F} if:

- Given a fixed function $f \in \mathcal{F}$, and a randomly generated key K (that can depend on f), the probability that an adversary outputs x such that $(x, \mathcal{H}(K, x)) = (x, f(x))$ is at most ϵ .
- The hash key K statistically hides the function f , such that adversaries cannot distinguish a random key from a key for f with advantage better than ϵ .

We will set $\epsilon = 2^{-2|pk|}$. We will use Π to denote a parallel repetition of Blum's Σ -protocol for Graph Hamiltonicity, represented as $\{a_i = \text{com}(\hat{a}_i)\}_{i \in [\lambda]}$, $\{e_i\}_{i \in [\lambda]}$, $\{z_i\}_{i \in [\lambda]}$, where $\{a_i\}_{i \in [\lambda]}$ represents the first commitments sent by the prover, $\{e_i\}_{i \in [\lambda]}$ is a challenge string sent by the verifier and $\{z_i\}_{i \in [\lambda]}$ represents the corresponding third message by the prover. Let the instance be x and its witness be w . Then, the protocol is described as follows.

1. In the first round, the verifier randomly samples a key K for the correlation intractable hash function \mathcal{H} for bounded size NC_1 functions.
2. In the second round, the prover picks a key pair (pk, sk) for the scheme PKE. Then the prover uses $\text{PKE.Enc}(pk, \cdot)$ as a commitment scheme to compute the commitments $\{a_i\}_{i \in [\lambda]}$. Next, the prover computes $e = \mathcal{H}(K, x, \{a_i\}_{i \in [\lambda]}) \in \{0, 1\}^\lambda$, and uses (x, w, a, e) to compute $z = (z_1, \dots, z_\lambda)$ according to the protocol Π . It outputs $(pk, \{a_i = \text{PKE.Enc}(pk, \hat{a}_i)\}_{i \in [\lambda]}, e, z)$

While witness indistinguishability of this protocol is easy to see, arguing soundness is trickier. In order to argue soundness, the reduction will simply try to *guess* the public key pk^* that the prover will use, and will abort if this guess is not correct. Note that such a guess is correct with probability at least $2^{-|pk^*|}$.

Suppose a cheating prover convinces a verifier to accept false statements with probability $\frac{1}{p(\lambda)}$ for some polynomial $p(\cdot)$. Then, with probability at least $\frac{1}{p(\cdot)} \cdot 2^{-|pk^*|}$, the reduction guesses pk^* correctly, and the prover provides a convincing proof of a false statement using pk^* .

In the next hybrid, the challenger guesses pk^* together with the corresponding secret key sk^* , and then samples a correlation intractable hash key for a specific function $f_{sk^*}(\cdot)$. The function $f_{sk^*}(\cdot)$ on input x , along with a (the messages committed in the Σ -protocol), outputs the only possible string e_{bad} for which there exists a string z such that (a, e_{bad}, z) verifies for $x \notin L$.¹ Note that this function is in NC_1 . By ϵ -security of the correlation intractable hash family (where $\epsilon = 2^{-2|pk^*|}$), with probability at least $\left(\frac{1}{p(\cdot)} - 2^{-|pk^*|}\right) \cdot 2^{-|pk^*|}$, the reduction guesses pk^* correctly, and the prover provides a convincing proof of a false statement using pk^* .

Finally, since the correlation intractable hash function is ϵ -secure, in the final hybrid adversary cannot produce a proof for x with probability greater than ϵ , as this will mean that he output a^*, e^*, z^* such that $e^* = f_{bad}(x, a^*)$.

The protocol sketched above is public-coin, because when we instantiate the correlation-intractable hash family with the LWE-based one by [PS19], the hash keys are statistically close to uniform.

In the description above, we also relied on a dense public key encryption scheme, which is unfortunately not known to exist based on LWE. However, we note that we can instead use a scheme with the property that at least $1/2$ of the strings in $\{0, 1\}^{\ell_{\text{PKE}}}$ correspond to correct encryption keys with a valid secret key, and the property that public keys are pseudorandom. Then, the verifier sends λ public keys pk_1, \dots, pk_λ , and the prover outputs pk' , and then uses the public keys $\{(pk' \oplus pk_i)\}_{i \in [\lambda]}$ to compute λ proofs. Soundness can be obtained by arguing that with overwhelming probability, there will exist an index $i \in [\lambda]$ such that $(pk' \oplus pk_i)$ has a secret key, just like the [DN02] technique described at the beginning of this overview.

However, the construction above falls short of achieving statistical witness indistinguishability against malicious verifiers. The reason is the following: arguing that the construction described above satisfies soundness requires relying on correlation intractability of the hash function. In order to invoke the correlation intractable hash function, it is crucial that the prover be “committed” to a well-defined, unique message $\{a_i\}_{i \in [\lambda]}$, that can be extracted using the secret key sk^* of the public key encryption scheme. At first, statistical hiding, together with such extraction, may appear to be contradictory objectives.

Indeed, we will try obtain a weaker version of these contradictory objectives, and specifically, we will rely on a two-message statistically hiding extractable commitment scheme [KKS18].

¹Note that this property is satisfied by any Σ -Protocol with a $1/2$ -special soundness where the bad challenge e_{bad} can be computed efficiently from the precommitted values $\{\hat{a}_i\}$, such as Blum’s Σ -protocol.

2.2 Using Correlation-Intractable Hashing with Statistically Hiding Extractable Commitments

In the recent exciting work on using LWE-based correlation-intractable hashing [PS19, CCH⁺19] for achieving soundness, as well as in the “warm up” ZAP protocol described above, the correlation-intractable hash function is used as follows. Because the LWE-based CI-hash function is designed to avoid an *efficiently computable* function f of the prover’s first message, it is used together with a public-key encryption scheme: the prover’s first message is encrypted using the public key, and the function f is built to contain the secret key of the encryption scheme, so that it can decrypt the prover’s first message in order to calculate the challenge that must be avoided.

Our work imagines a simple modification of this strategy of using correlation-intractable hashing for arguing soundness. The main idea is that we want to replace the encryption scheme (which necessarily can only at most provide computational hiding) with an *extractable* statistically hiding commitment scheme. We will describe what this object entails in more detail very shortly, but the main observation is that such an extractable commitment in fact reveals the value being committed to with a tiny (but tunable) probability – crucially in a way that prevents a malicious prover from learning whether the commitment will reveal the committed value or not. With such a commitment scheme, the efficient function f underlying the correlation-intractable hash function will only “work” in the rare case that the commitment reveals the value being committed. But since a cheating prover can’t tell whether its committed values will be revealed or not, soundness will still hold overall, even though the actual guarantee of the correlation-intractable hash function is only invoked with a tiny probability in the proof of soundness. We now elaborate.

2.3 Statistically Hiding Extractable Commitments

Any statistically hiding commitment must lose all information about the committed message, except with negligible probability. This makes it challenging to define notions of extraction for statistically hiding commitments. In 4 rounds or more, this notion is easier to define, as extraction is possible even from statistically hiding commitments, simply by rewinding the adversary. However, traditional rewinding techniques break down completely when considering two-message commitments.

Nevertheless, the recent work of [KKS18], building on [KS17], defined and constructed two-message statistically hiding extractable commitments, which they used to construct two-message statistical WI arguments, that were *privately verifiable*. In what follows, we abstract out the properties of a statistically hiding extractable commitment. A more formal description can be found in Section 5. We point out that we only need to rely on significantly simpler definitions than the ones in [KKS18], and we give much simpler proofs that the constructions in [KKS18] according to our new definitions. This may be of independent interest.

Defining Statistically Hiding Extractable Commitments. We start with an important observation about statistically hiding commitments, which gives a hint about how one can possibly define (and construct) two-message statistically hiding extractable commitments. Namely, any statistically hiding commitment must lose all information about the committed message, *but may retain this information with some small negligible probability*. Specifically,

- A commitment that leaks the committed message with probability ϵ (where ϵ is a fixed negligible function in the security parameter) and statistically hides the message otherwise, will continue to be statistically hiding.

- At the same time, one could ensure that no matter the behavior of the committer, the message being committed *does get leaked to the honest receiver with probability at least ϵ* .
- Moreover, the committer does not know whether or not the committed message was leaked to the receiver. This property is important and will be crucially used in our proofs.

In spirit, this corresponds to establishing an erasure channel over which the committer transmits his message to the receiver. This channel almost always erases the committed message, but is guaranteed to transmit the committed message with a very small probability (ϵ). Moreover, just like cryptographic erasure channels, the committer does not know whether or not his message was transmitted. Additionally, because this is a commitment, we require computational binding: once the committer transmits his message (that is, commits), he should not be able to change his mind about the message, *even if* the message did not get transmitted. Finally, we say that “extraction occurs” whenever the message does get transmitted, and we require that extraction occur with probability at least ϵ , even against a malicious committer.

Next, we describe how we interface these commitments with correlation intractable hash functions to obtain two-message statistical ZAP arguments.

2.4 Statistical ZAP Arguments

With this tool in mind, we make the following observations:

1. We would like to replace the encryption scheme used for generating the first message a for the sigma protocol, sent by the prover in the second round, with a statistically hiding commitment.
2. The first message of this commitment will be generated by the verifier. Furthermore, because we want a public coin protocol, we require this message to be pseudorandom.
3. We will require that with some negligible probability ϵ , *all* messages committed by the prover get transmitted to the verifier, that is with probability ϵ , the verifier can recover all the messages committed by the prover in polynomial time given his secret state. Next, using an insight from the simple protocol in Section 2.1, we will set the security of the correlation intractable hash function, so that it is infeasible for any polynomially sized adversary to break correlation intractability with probability ϵ .

The protocol is then as follows:

- In the first round, the verifier samples a hash key K for the correlation intractable hash function \mathcal{H} , for the same function family \mathcal{F} as Section 2.1. The verifier also samples strings $q = \{c_{1,j}\}_{j \in [\text{poly}(\lambda)]}$ uniformly at random, where poly is a polynomial denoting the number of commitments made by the prover. The verifier sends q and K over to the prover.
- In the second round, the prover computes the first message of the sigma protocol a (where the number of parallel repetitions equals the output length of correlation intractable hash function). This message a is generated using the statistically hiding extractable commitment scheme com with q as the first message. The prover computes $e = \mathcal{H}(K, a)$ and uses e to compute the third message z of the sigma protocol, by opening some subset of the commitments made by the prover. The prover outputs (a, e, z) .

We now provide some intuition for the security of this protocol.

- **Soundness:** To argue soundness, we follow an approach that is similar to the soundness proof for the computational ZAP argument described in Section 2.1 (although with some additional technical subtleties). We discuss one such subtlety here:

Let $\ell = |e|$. For simplicity, at this point, let's assume that the correlation-intractable hash function is subexponentially secure. In other words, it is at most $2^{-\ell^\delta}$ -secure². For this reason, we require the commitments to be *jointly* extractable in polynomial time with probability at least $2^{-\ell^\delta}$. Note that the total number of commitments is $N = \ell \cdot \text{poly}(\lambda)$.

However, statistically hiding commitments, as originally constructed in [KKS18], are such that if a single commitment can be extracted with probability ϵ , then N commitments can be extracted with probability roughly ϵ^N . Setting $N = \ell \cdot \text{poly}(\lambda)$ as above implies that trivially, the probability of extraction will be roughly $O(2^{-\ell \cdot \text{poly}(\lambda)})$, which is smaller than the required probability $2^{-\ell^\delta}$.

However, we observe that the commitments constructed in [KKS18] can be modified very slightly so that the probability of extraction can be $2^{-g(\lambda)}$ for any efficiently computable function g that is bounded by any polynomial in λ . Thus, for example, the probability of extraction can be made to be $\lambda^{-\log(\lambda)}$. In other words, this extraction probability can be made to be *independent of the total number of commitments, N* . We describe this modification in additional detail in Section 4.2.

Using commitments that satisfy the property stated above, we observe that we can switch to a hybrid where the challenger samples the commitment messages on behalf of the verifier, and hardwires the secret state used for extraction inside the hash key. The function is defined such that in the event that extraction occurs (given the secret state), the verifier can use the extracted values to compute the bad challenge e_{bad} (just as in Section 2.1), by evaluating a depth bounded function f_{bad} on the extracted values, and otherwise e_{bad} is set to 0. If the adversary breaks soundness with noticeable probability ϵ , then with probability roughly at least $2^{-g(\lambda)} \cdot \epsilon$, the outputs of the adversary satisfy:

$$\mathcal{H}(K, a) = e_{bad}$$

We set the function g so that $2^{-g(\lambda)} \cdot \epsilon$ is significantly larger than the advantage of the hash function, and then the event above suffices to contradict correlation intractability. It also turns out that because g is tunable independently of the strength of any of the cryptographic primitives that we are using, the hash function does not need to be subexponentially secure. In fact, it suffices to have security with any explicit, efficiently computable advantage μ . We explain this further in Section 5.

- **Statistical Witness Indistinguishability:** Statistical witness indistinguishability composes under parallel repetition, and therefore can be proven index-by-index based on the statistical hiding property of the commitment.

Additional details about the construction and the proof can be found in Section 5.

Related Work. In a concurrent and independent work, [GJJM] also constructed a 2-message public-coin statistically witness indistinguishable argument from quasipolynomial LWE. Another concurrent and independent work is that of [LVW], who construct another 2-message statistical

²More formally, if the output of the hash function is ℓ bits long, then even if we rely on sub-exponential assumptions, we cannot hope to have the guessing advantage be smaller than $2^{-\ell^\delta}$ for a small positive constant $\delta < 1$.

witness indistinguishable argument, which is *private-coin* but *publically verifiable*, from a bilinear map assumption. The authors call this a “ZAPR”. ZAPRs were previously constructed by [Gro, Lip, GGPR] under a different name, but these three older constructions all rely on non-standard non-falsifiable assumptions.

2.5 Organization

The rest of this paper is organized as follows. In Section 3, we describe some of the preliminaries such as correlation intractability, oblivious transfer and proof systems. In Section 4, we define a simplified variant and present a slightly modified construction of extractable statistically hiding commitments, first proposed by [KKS18]. Finally, in Section 5, we construct and prove the security of our statistical ZAP argument.

3 Preliminaries

Notation. Throughout this paper, we will use λ to denote the security parameter, and $\text{negl}(\lambda)$ to denote any function that is asymptotically smaller than $\frac{1}{\text{poly}(\lambda)}$ for any polynomial $\text{poly}(\cdot)$.

The statistical distance between two distributions D_1, D_2 is denoted by $\Delta(D_1, D_2)$ and defined as:

$$\Delta(D_1, D_2) = \frac{1}{2} \sum_{v \in V} |\Pr_{x \leftarrow D_1}[x = v] - \Pr_{x \leftarrow D_2}[x = v]|.$$

We say that two families of distributions $D_1 = \{D_{1,\lambda}\}, D_2 = \{D_{2,\lambda}\}$ are statistically indistinguishable if $\Delta(D_{1,\lambda}, D_{2,\lambda}) = \text{negl}(\lambda)$. We say that two families of distributions $D_1 = \{D_{1,\lambda}\}, D_2 = \{D_{2,\lambda}\}$ are computationally indistinguishable if for all non-uniform probabilistic polynomial time distinguishers \mathcal{D} ,

$$|\Pr_{r \leftarrow D_{1,\lambda}}[\mathcal{D}(r) = 1] - \Pr_{r \leftarrow D_{2,\lambda}}[\mathcal{D}(r) = 1]| = \text{negl}(\lambda).$$

Let Π denote an execution of a protocol. We use $\text{View}_A(\Pi)$ denote the view, including the randomness and state of party A in an execution Π . We also use $\text{Output}_A(\Pi)$ denote the output of party A in an execution of Π .

Remark 1. *In what follows we define several 2-party protocols. We note that in all these protocols both parties take as input the security parameter 1^λ . We omit this from the notation for the sake of brevity.*

Definition 1 (Σ -protocols). *Let $L \in \text{NP}$ with corresponding witness relation R_L . A protocol $\Pi = \langle P, V \rangle$ is a Σ -protocol for relation R_L if it is a three-round public-coin protocol which satisfies:*

- **Completeness:** *For all $(x, w) \in R_L$, $\Pr[\text{Output}_V \langle P(x, w), V(x) \rangle = 1] = 1 - \text{negl}(\lambda)$, assuming P and V follow the protocol honestly.*
- **Special Soundness:** *There exists a polynomial-time algorithm A that given any x and a pair of accepting transcripts $(a, e, z), (a, e', z')$ for x with the same first prover message, where $e \neq e'$, outputs w such that $(x, w) \in R_L$.*
- **Honest verifier zero-knowledge:** *There exists a probabilistic polynomial time simulator \mathcal{S}_Σ such that for all $(x, w) \in R_L$, the distributions $\{\mathcal{S}_\Sigma(x, e)\}$ and $\{\text{View}_V \langle P(x, w(x)), V(x, e) \rangle\}$ are statistically indistinguishable. Here $\mathcal{S}_\Sigma(x, e)$ denotes the output of simulator \mathcal{S} upon input x and e , such that V 's random tape (determining its query) is e .*

3.1 Correlation Intractable Hash Functions

We adapt definitions of a correlation intractable hash function family from [PS19, CCH⁺19].

Definition 2. For any polynomials $k, (\cdot), s(\cdot) = \omega(k(\cdot))$ and any $\lambda \in \mathbb{N}$, let $\mathcal{F}_{\lambda, s(\lambda)}$ denote the class of NC^1 circuits of size $s(\lambda)$ that on input $k(\lambda)$ bits output λ bits. Namely, $f : \{0, 1\}^{k(\lambda)} \rightarrow \{0, 1\}^\lambda$ is in $\mathcal{F}_{\lambda, s}$ if it has size $s(\lambda)$ and depth bounded by $O(\log \lambda)$.

Definition 3. [μ_{CI} -Correlation Intractable Hash Function Family] A hash function family $\mathcal{H} = (\text{Setup}, \text{Eval})$ is μ_{CI} -correlation intractable (CI) with respect to $\mathcal{F} = \{\mathcal{F}_{\lambda, s(\lambda)}\}_{\lambda \in \mathbb{N}}$ as defined in Definition 2, if the following two properties hold:

- **Correlation Intractability:** For every $f \in \mathcal{F}_{\lambda, s}$, every non-uniform polynomial-size adversary \mathcal{A} , every polynomial s , and every large enough $\lambda \in \mathbb{N}$,

$$\Pr_{K \leftarrow \mathcal{H}.\text{Setup}(1^\lambda, f)} \left[\mathcal{A}(K) \rightarrow x \text{ such that } (x, \mathcal{H}.\text{Eval}(K, x)) = (x, f(x)) \right] \leq \mu_{CI}(\lambda).$$

- **Statistical Indistinguishability of Hash Keys:** Moreover, for every $f \in \mathcal{F}_{\lambda, s}$, for every unbounded adversary \mathcal{A} , and every large enough $\lambda \in \mathbb{N}$,

$$\left| \Pr_{K \leftarrow \mathcal{H}.\text{Setup}(1^\lambda, f)} [\mathcal{A}(K) = 1] - \Pr_{K \leftarrow \{0, 1\}^\ell} [\mathcal{A}(K) = 1] \right| \leq 2^{-\lambda^{\Omega(1)}},$$

where ℓ denotes the size of the output of $\mathcal{H}.\text{Setup}(1^\lambda, f)$.

The work of [PS19] gives a construction of correlation intractable hash functions with respect to $\mathcal{F} = \{\mathcal{F}_{\lambda, s(\lambda)}\}_{\lambda \in \mathbb{N}}$, based on polynomial LWE with polynomial approximation factors. We observe that their construction also satisfies Definition 3, assuming LWE with an explicit efficiently computable advantage upper bound.

3.2 Oblivious Transfer

Definition 4 (Oblivious Transfer). *Oblivious transfer is a protocol between two parties, a sender S with input messages (m_0, m_1) and receiver R with input a choice bit b . The correctness requirement is that R obtains output m_b at the end of the protocol (with probability 1). We let $\langle S(m_0, m_1), R(b) \rangle$ denote an execution of the OT protocol with sender input (m_0, m_1) and receiver input bit b . We require OT that satisfies the following properties:*

- **Computational Receiver Security.** For any non-uniform PPT sender S^* and any $(b, b') \in \{0, 1\}$, the views $\text{View}_{S^*}(\langle S^*, R(b) \rangle)$ and $\text{View}_{S^*}(\langle S^*, R(b') \rangle)$ are computationally indistinguishable.

We say that the OT scheme is μ_{OT} -secure if all PPT malicious senders have distinguishing advantage less than $\mu_{OT}(\lambda)$.

- **Statistical Sender Security.** This is defined using the real-ideal paradigm, and requires that for any distribution on the inputs (m_0, m_1) and any unbounded adversarial receiver R^* , there exists a (possibly unbounded) simulator Sim_{R^*} that interacts with an ideal functionality \mathcal{F}_{ot} on behalf of R^* . Here \mathcal{F}_{ot} is an oracle that obtains the inputs (m_0, m_1) from S and b from Sim_{R^*} (simulating the malicious receiver), and outputs m_b to Sim_{R^*} . Then $\text{Sim}_{R^*}^{\mathcal{F}_{\text{ot}}}$ outputs a receiver view that is statistically indistinguishable from the real view of the malicious receiver $\text{View}_{R^*}(\langle S(m_0, m_1), R^* \rangle)$. We say that the OT protocol satisfies $(1 - \delta)$ statistical sender security if the statistical distance between the real and ideal distributions is at most δ .

We use the following sender security property in our protocols (which follows from the definition of sender security in Definition 4 above).

Claim 1. *For any two-message OT protocol satisfying Definition 4, for every malicious receiver R^* and every first message m_{R^*} generated by R^* such that the sender does not reject given m_{R^*} , we require that there exists an unbounded machine E which extracts b such that either of the following statements is true:*

- *For all m_0, m_1, m_2 , $\text{View}_{R^*}\langle S(m_0, m_1), R^* \rangle$ and $\text{View}_{R^*}\langle S(m_0, m_2), R^* \rangle$ are statistically indistinguishable and $b = 0$, or,*
- *For all m_0, m_1, m_2 , $\text{View}_{R^*}\langle S(m_0, m_1), R^* \rangle$ and $\text{View}_{R^*}\langle S(m_2, m_1), R^* \rangle$ are statistically indistinguishable and $b = 1$.*

Proof. From the (unbounded) simulation property of the two-message OT, there exists a simulator that extracts a receiver input bit b from the first message of R^* , sends it to the ideal functionality, obtains m_b and generates an indistinguishable receiver view. Then, by the definition of sender security, when $b = 0$, the simulated view must be close to both $\text{View}_{R^*}\langle S(m_0, m_1), R^* \rangle$, and $\text{View}_{R^*}\langle S(m_0, m_2), R^* \rangle$. Similarly, when $b = 1$, the simulated view must be statistically close to both $\text{View}_{R^*}\langle S(m_0, m_1), R^* \rangle$, and $\text{View}_{R^*}\langle S(m_2, m_1), R^* \rangle$. \square

Throughout the paper, we focus on two-message oblivious transfer. We now discuss an additional specific property of two-message OT protocols.

Property 1. *The message sent by the receiver is pseudorandom - in particular, this means that the receiver can just sample and send a uniformly random string as a valid message to the sender.*

Two-message OT protocols with this additional property have been constructed based on the DDH assumption [NP01], LWE assumption [BD18], and a stronger variant of smooth-projective hashing, which can be realized from DDH as well as the N^{th} -residuosity and Quadratic Residuosity assumptions [Kal05, HK12]. Such two-message protocols can also be based on witness encryption or indistinguishability obfuscation (*iO*) together with one-way permutations [SW14].

3.3 Proof Systems

An n -message interactive protocol for deciding a language L with associated relation R_L proceeds in the following manner:

- At the beginning of the protocol, P and V receive the size of the instance and security parameter, and execute the first $n - 1$ messages.
- At some point during the protocol, P receives input $(x, w) \in R_L$. P sends x to V together with the last message of the protocol. Upon receiving the last message from P , V outputs 1 or 0.

An execution of this protocol with instance x and witness w is denoted by $\langle P(x, w), V(x) \rangle$. One can consider both proofs – with soundness against unbounded provers, and arguments – with soundness against computationally bounded provers.

Definition 5 (Two-Message Interactive Arguments). *A two-message delayed-input interactive protocol (P, V) for deciding a language L is an interactive argument for L if it satisfies the following properties:*

- **Completeness:** For every $(x, w) \in R_L$,

$$\Pr[\text{Output}_V\langle P(x, w), V(x) \rangle = 1] = 1 - \text{negl}(\lambda),$$

where the probability is over the random coins of P and V , and where in the protocol P receives (x, w) right before computing the last message of the protocol, and V receives x together with the last message of the protocol.

- **Non-Adaptive Soundness:** For every (non-uniform) PPT prover P^* that on input 1^λ (and without access to the verifier's message) outputs a length $1^{p(\lambda)}$ and $x \in \{0, 1\}^{p(\lambda)} \setminus L$,

$$\Pr[\text{Output}_V\langle P^*, V \rangle(x) = 1] = \text{negl}(\lambda),$$

where the probability is over the random coins of V .

Witness Indistinguishability. A proof system is witness indistinguishable if for any statement with at least two witnesses, proofs computed using different witnesses are indistinguishable. In this paper, we only consider statistical witness indistinguishability, which we formally define below.

Definition 6 (Statistical Witness Indistinguishability). *A delayed-input interactive argument (P, V) for a language L is said to be statistical witness-indistinguishable if for every unbounded verifier V^* , every polynomially bounded function $n = n(\lambda) \leq \text{poly}(\lambda)$, and every $(x_n, w_{1,n}, w_{2,n})$ such that $(x_n, w_{1,n}) \in R_L$ and $(x_n, w_{2,n}) \in R_L$ and $|x_n| = n$, the following two ensembles are statistically indistinguishable:*

$$\{\text{View}_{V^*}\langle P(x_n, w_{1,n}), V^*(x_n) \rangle\} \text{ and } \{\text{View}_{V^*}\langle P(x_n, w_{2,n}), V^*(x_n) \rangle\}$$

4 Extractable Commitments

4.1 Definitions

We take the following definition of statistically hiding extractable commitments from [KKS18]. As before, we use λ to denote the security parameter, and we let $p = \text{poly}(\lambda)$ be an arbitrary fixed polynomial such that the message space is $\{0, 1\}^{p(\lambda)}$.

We restrict ourselves to commitments with non-interactive decommitment, and where the (honest) receiver is not required to maintain any state at the end of the commit phase in order to execute the decommit phase. Our construction will satisfy this property and this will be useful in our applications to constructing statistically private arguments.

Definition 7 (Statistically Hiding Commitment Scheme). *A commitment $\langle \mathcal{C}, \mathcal{R} \rangle$ is a two-phase protocol between a committer \mathcal{C} and receiver \mathcal{R} , consisting of algorithms `Commit`, `Decommit` and `Verify`. At the beginning of the protocol, \mathcal{C} obtains as input a message $M \in \{0, 1\}^p$. Next, \mathcal{C} and \mathcal{R} execute the commit phase, and obtain a commitment transcript, denoted by τ , together with private states for \mathcal{C} and \mathcal{R} , denoted by $\text{state}_{\mathcal{C}, \tau}$ and $\text{state}_{\mathcal{R}, \tau}$ respectively. We use the notation*

$$(\tau, \text{state}_{\mathcal{C}, \tau}, \text{state}_{\mathcal{R}, \tau}) \leftarrow \text{Commit}\langle \mathcal{C}(M), \mathcal{R} \rangle.$$

Later, \mathcal{C} and \mathcal{R} possibly engage in a decommit phase, where the committer \mathcal{C} computes and sends message $y = \text{Decommit}(\tau, \text{state}_{\mathcal{C}, \tau})$ to \mathcal{R} . At the end, \mathcal{R} computes `Verify`(τ, y) to output \perp or a message $\widetilde{M} \in \{0, 1\}^p$.³

A statistically hiding commitment scheme is required to satisfy three properties:

³ We note that in our definition, \mathcal{R} does not need to use private state $\text{state}_{\mathcal{R}, \tau}$ from the commitment phase in order to execute the `Verify` algorithm in the decommitment phase.

- **(Perfect) Completeness.** If \mathcal{C}, \mathcal{R} honestly follow the protocol, then for every $M \in \{0, 1\}^p$:

$$\Pr[\text{Verify}(\tau, \text{Decommit}(\tau, \text{state}_{\mathcal{C}, \tau})) = M] = 1$$

where the probability is over $(\tau, \text{state}_{\mathcal{C}, \tau}) \leftarrow \text{Commit}\langle \mathcal{C}(M), \mathcal{R} \rangle$.

- **Statistical Hiding.** For every two messages $M_1, M_2 \in \{0, 1\}^{2p}$, every unbounded malicious receiver \mathcal{R}^* and honest committer \mathcal{C} , a commitment is $\delta(\lambda)$ -statistically hiding if the statistical distance between the distributions $\text{View}_{\mathcal{R}^*}(\text{Commit}\langle \mathcal{C}(M_1), \mathcal{R}^* \rangle)$ and $\text{View}_{\mathcal{R}^*}(\text{Commit}\langle \mathcal{C}(M_2), \mathcal{R}^* \rangle)$ is at most $\delta(\lambda)$. The scheme is statistically hiding if $\delta(\lambda) \leq \frac{1}{\text{poly}(\lambda)}$ for every polynomial $\text{poly}(\cdot)$.
- **Computational Binding.** Consider any non-uniform PPT committer \mathcal{C}^* that produces $\tau \leftarrow \text{Commit}\langle \mathcal{C}^*, \mathcal{R} \rangle$, and then outputs y_1, y_2 . Let $\widetilde{M}_1 = \text{Verify}(\tau, y_1)$ and $\widetilde{M}_2 = \text{Verify}(\tau, y_2)$. Then, we require that

$$\Pr[(\widetilde{M}_1 \neq \perp) \wedge (\widetilde{M}_2 \neq \perp) \wedge (\widetilde{M}_1 \neq \widetilde{M}_2)] = \text{negl}(\lambda),$$

over the randomness of sampling $\tau \leftarrow \text{Commit}\langle \mathcal{C}^*, \mathcal{R} \rangle$.

We also define an extractor \mathcal{E} that given black-box access to \mathcal{C}^* , and then without executing any decommitment phase with \mathcal{C}^* , outputs message \widetilde{M} committed by \mathcal{C}^* with probability at least ϵ : we require “correctness” of this extracted message \widetilde{M} . We also require that no PPT adversary can distinguish transcripts where extraction is successful from those where it is unsuccessful. This is formally described in Definition 8.

Definition 8 (ϵ -Extractable Statistically Hiding Commitment). *We say that a statistically hiding commitment scheme is ϵ -extractable if the following holds: Denote $(\tau, \text{state}_{\mathcal{C}, \tau}, \text{state}_{\mathcal{R}, \tau}) \leftarrow \text{Commit}\langle \mathcal{C}^*, \mathcal{R} \rangle$. We require that there exists a deterministic polynomial time extractor \mathcal{E} that on input $(\tau, \text{state}_{\mathcal{R}, \tau})$ outputs \widetilde{M} such that the following properties hold.*

- **Frequency of Extraction.** For every PPT committer \mathcal{C}^* ,

$$\Pr[\mathcal{E}(\tau, \text{state}_{\mathcal{R}, \tau}) \neq \perp] = \epsilon$$

where the probability is over $(\tau, \text{state}_{\mathcal{C}, \tau}, \text{state}_{\mathcal{R}, \tau}) \leftarrow \text{Commit}\langle \mathcal{C}^*, \mathcal{R} \rangle$.

- **Correctness of Extraction.** For every PPT committer \mathcal{C}^* , every $(\tau, \text{state}_{\mathcal{C}, \tau}, \text{state}_{\mathcal{R}, \tau}) \in \text{Supp}(\text{Commit}\langle \mathcal{C}^*, \mathcal{R} \rangle)$, and every y , denoting $\widetilde{M} = \mathcal{E}(\tau, \text{state}_{\mathcal{R}, \tau})$ and $M = \text{Verify}(\tau, y)$, if $\widetilde{M} \neq \perp$ and $M \neq \perp$, then $\widetilde{M} = M$.
- **Indistinguishability of Extractable Transcripts.** For every \mathcal{C}^* ,

$$|\Pr[\mathcal{C}^*(\tau) = 1 \mid \mathcal{E}(\tau, \text{state}_{\mathcal{R}, \tau}) \neq \perp] - \Pr[\mathcal{C}^*(\tau) = 1 \mid \mathcal{E}(\tau, \text{state}_{\mathcal{R}, \tau}) = \perp]| = \text{negl}(\lambda)$$

where the probability is over $(\tau, \text{state}_{\mathcal{R}, \tau}) \leftarrow \text{Commit}\langle \mathcal{C}^*, \mathcal{R} \rangle$.

We also consider a stronger definition, of ϵ -extractable statistically hiding ℓ multi-commitments, where we require that an entire sequence of ℓ commitments can be extracted with probability ϵ , that is independent of ℓ . We will also modify the `Verify` algorithm so that it obtains as input the transcript $\tau := (\tau_1, \tau_2, \dots, \tau_\ell)$ of all ℓ commitments, together with an index $i \in [\ell]$ and the decommitment $\text{state}_{\mathcal{C}, \tau, i}$ to a single commitment.

Definition 9 (ϵ -Extractable Statistically Hiding ℓ Multi-Commitments). A sequence of commitments $\langle \mathcal{C}, \mathcal{R} \rangle$ is a two-phase protocol between a committer \mathcal{C} and receiver \mathcal{R} , consisting of algorithms *Commit*, *Decommit* and *Verify*. At the beginning of the protocol, \mathcal{C} obtains as input ℓ messages $(M_1, \dots, M_\ell) \in \{0, 1\}^{p\ell}$. Next, \mathcal{C} and \mathcal{R} execute the commit phase, and obtain a commitment transcript, denoted by $\tau := (\tau_1, \dots, \tau_\ell)$, together with private states for \mathcal{C} and \mathcal{R} , denoted by $\{\text{state}_{\mathcal{C}, \tau, i}\}_{i \in [\ell]}$ and $\{\text{state}_{\mathcal{R}, \tau}\}_{i \in [\ell]}$ respectively. We use the notation

$$(\tau, \{\text{state}_{\mathcal{C}, \tau, i}\}_{i \in [\ell]}, \{\text{state}_{\mathcal{R}, \tau, i}\}_{i \in [\ell]}) \leftarrow \text{Commit}(\mathcal{C}(M), \mathcal{R}).$$

Later, \mathcal{C} and \mathcal{R} possibly engage in a decommit phase, where the committer \mathcal{C} computes and sends message $y = \text{Decommit}(\tau_i, \text{state}_{\mathcal{C}, \tau, i})$ to \mathcal{R} . At the end, \mathcal{R} computes $\text{Verify}(\tau, i, y_i)$, where $i \in [\ell]$, to output \perp or a message $\widetilde{M}_i \in \{0, 1\}^p$.⁴ A statistically hiding commitment scheme is required to satisfy three properties:

- **(Perfect) Completeness.** If \mathcal{C}, \mathcal{R} honestly follow the protocol, then for every $M_1, \dots, M_\ell \in \{0, 1\}^{p\ell}$ and every $i \in [\ell]$:

$$\Pr[\text{Verify}(\tau, i, \text{Decommit}(\tau_i, \text{state}_{\mathcal{C}, \tau, i})) = M_i] = 1$$

where the probability is over $(\tau, \{\text{state}_{\mathcal{C}, \tau, i}\}_{i \in [\ell]}) \leftarrow \text{Commit}(\mathcal{C}(M), \mathcal{R})$.

- **Statistical Hiding.** A set of commitments is $\delta(\lambda)$ -statistically hiding if for every set $I \subseteq [\ell]$, every set of messages $\{M_i\}_{i \in [I]}, \{M'_i\}_{i \in [I]} \in \{0, 1\}^{2p\ell}$ such that $(M_i = M'_i)$ for every $i \in I$, every unbounded malicious receiver \mathcal{R}^* and honest committer \mathcal{C} , the statistical distance between the distributions $\text{View}_{\mathcal{R}^*}(\text{Commit}(\mathcal{C}(M_1, \dots, M_\ell), \mathcal{R}^*), \{\text{Decommit}(\tau_i, \text{state}_{\mathcal{C}, \tau, i})\}_{i \in I})$ and $\text{View}_{\mathcal{R}^*}(\text{Commit}(\mathcal{C}(M'_1, \dots, M'_\ell), \mathcal{R}^*), \{\text{Decommit}(\tau_i, \text{state}_{\mathcal{C}, \tau, i})\}_{i \in I})$ is at most $(\ell - |I|) \cdot \delta(\lambda)$. The scheme is statistically hiding if $\delta(\lambda) \leq \frac{1}{\text{poly}(\lambda)}$ for every polynomial $\text{poly}(\cdot)$.
- **Computational Binding.** Consider any non-uniform PPT committer \mathcal{C}^* that produces $\tau \leftarrow \text{Commit}(\mathcal{C}^*, \mathcal{R})$ and then outputs y_1, y_2 . For any $i \in [\ell]$, let $\widetilde{M}_1 = \text{Verify}(\tau, i, y_1)$ and $\widetilde{M}_2 = \text{Verify}(\tau, i, y_2)$. Then, we require that

$$\Pr[(\widetilde{M}_1 \neq \perp) \wedge (\widetilde{M}_2 \neq \perp) \wedge (\widetilde{M}_1 \neq \widetilde{M}_2)] = \text{negl}(\lambda),$$

over the randomness of sampling $\tau \leftarrow \text{Commit}(\mathcal{C}^*, \mathcal{R})$.

Definition 10 (ϵ -Extractable Sequence of ℓ Statistically Hiding Commitments). We say that a sequence of ℓ statistically hiding commitments is ϵ -extractable if the following holds:

For every $i \in [\ell]$, Denote $\{(\tau_i, \text{state}_{\mathcal{C}, \tau, i}, \text{state}_{\mathcal{R}, \tau, i})\}_{i \in [\ell]} \leftarrow \text{Commit}(\mathcal{C}^*, \mathcal{R})$ and let $\tau := (\tau_1, \tau_2, \dots, \tau_\ell)$ denote the transcript of all ℓ commitments. We require that there exists a deterministic polynomial time extractor \mathcal{E} that on input $\{(\tau_i, \text{state}_{\mathcal{R}, \tau, i})\}_{i \in [\ell]}$ outputs $\{\widetilde{M}_i\}_{i \in [\ell]}$ such that the following properties hold.

- **Frequency of Extraction.** Let $(\widetilde{M}_1, \dots, \widetilde{M}_\ell) := \mathcal{E}(\{(\tau_i, \text{state}_{\mathcal{R}, \tau, i})\}_{i \in [\ell]})$, then

$$\Pr[\exists i \text{ such that } \widetilde{M}_i = \perp] \leq (1 - \epsilon)$$

where the probability is over τ and over the random coins of \mathcal{R} .

⁴ We point out that in this definition, \mathcal{R} does not need to use private state $\text{state}_{\mathcal{R}, \tau}$ from the commitment phase in order to execute the *Verify* algorithm in the decommitment phase. We also point out that the verify algorithm can be modified so that the committer can open a subset of commitments together, instead of one at a time.

- **Correctness of Extraction.** For every $\{(\tau_i, \text{state}_{\mathcal{C}, \tau, i}, \text{state}_{\mathcal{R}, \tau, i})\}_{i \in [\ell]} \in (\text{Supp}(\text{Commit}(\mathcal{C}^*, \mathcal{R})))^\ell$, denoting $\widetilde{M}_1, \dots, \widetilde{M}_\ell = \mathcal{E}(\{(\tau_i, \text{state}_{\mathcal{R}, \tau, i})\}_{i \in [\ell]})$, we require the following. For every $i \in [\ell]$ and every string y , denoting $M = \text{Verify}(\tau, i, y)$, if $M \neq \perp$ and $\widetilde{M}_i \neq \perp$, then $\widetilde{M}_i = M$.

- **Indistinguishability of Extractable Transcripts.** For every PPT adversary \mathcal{C}^* ,

$$|\Pr[\mathcal{C}^*(\tau) = 1 \mid \exists i \text{ such that } \widetilde{M}_i = \perp] - \Pr[\mathcal{C}^*(\tau) = 1 \mid \nexists i \text{ such that } \widetilde{M}_i = \perp]| = \text{negl}(\lambda)$$

where $\widetilde{M}_1, \dots, \widetilde{M}_\ell := \mathcal{E}(\tau, \text{state}_{\mathcal{R}, \tau})$ and where the probability is over $(\tau, \text{state}_{\mathcal{R}, \tau}) \leftarrow \text{Commit}(\mathcal{C}^*, \mathcal{R})$.

4.2 Protocol

In this section, we construct two-message statistically hiding, extractable commitments according to Definition 8 assuming the existence of two message oblivious transfer (OT). Our construction is described in Figure 1.

Primitives Used. Let $\text{OT} = (\text{OT}_1, \text{OT}_2)$ denote a two-message string oblivious transfer protocol according to Definition 4, also satisfying Property 1. Let $\text{OT}_1(b; r_1)$ denote the first message of the OT protocol with receiver input b and randomness r_1 , and let $\text{OT}_2(M_0, M_1; r_2)$ denote the second message of the OT protocol with sender input strings M_0, M_1 and randomness r_2 .⁵

Extraction parameter: m .

Committer Input: Message $M \in \{0, 1\}^p$.

Commit Stage:

Receiver Message.

1. Pick challenge string $\text{ch} \xleftarrow{\$} \{0, 1\}^m$.
2. Sample uniform randomness $\{r_{1,i}\}_{i \in [m]}$.
3. Compute and send $\{\text{OT}_1(\text{ch}_i, r_{1,i})\}_{i \in [m]}$ using m instances of two-message OT.

Committer Message.

1. Sample a random string $r \xleftarrow{\$} \{0, 1\}^m$.
For every $i \in [m]$ and every $b \in \{0, 1\}$, sample $M_i^b \xleftarrow{\$} \{0, 1\}^p$ subject to $\bigoplus_{i \in [m]} M_i^{r_i} = M$.
2. For every $i \in [m]$ compute $o_{2,i} = \text{OT}_2(M_i^0, M_i^1; r_{2,i})$ with uniform randomness $r_{2,i}$.
3. Send $(r, \{o_{2,i}\}_{i \in [m]})$.

Reveal Stage: The committer reveals M , and all values $\{M_i^0, M_i^1\}_{i \in [m]}$ as well as the randomness $r_{2,i}$. The receiver accepts the decommitment to message M if and only if:

1. For all $i \in [m]$, $o_{2,i} = \text{OT}_2(M_i^0, M_i^1; r_{2,i})$,
2. $\bigoplus_{i \in [m]} M_i^{r_i} = M$.

Figure 1: Extractable Commitments

⁵Note that OT_2 also depends on OT_1 . We omit this dependence in our notation for brevity.

Observe that the protocol satisfies the property mentioned in the definition that the verify algorithm in the decommitment phase does not require the private randomness used by the receiver in the commit phase. Further, observe that if the oblivious transfer protocol satisfies Property 1, the receiver's message can alternately be generated by just sampling a uniformly random string. Thus, this would give an extractable commitment protocol where the receiver's algorithms are public coin.

We will now prove the following main theorem.

Theorem 1. *Set $\epsilon = 1/2^m(1 - 2^m \cdot m \cdot \mu_{ot})$. Assuming that the underlying OT protocol is μ_{OT} -secure against malicious senders, $(1 - \delta_{OT})$ secure against malicious receivers according to Definition 4, and satisfies Property 1, the scheme in Figure 1 is a $(1 - 2^{-m} - \delta_{OT})$ statistically hiding, ϵ -extractable commitment scheme according to Definition 8. Further, the receiver's algorithms are public coin.*

Recall that such two-message OT protocols with this additional property have been constructed based on the DDH assumption [NP01], LWE assumption [BD18], and a stronger variant of smooth-projective hashing, which can be realized from DDH as well as the N^{th} -residuosity and Quadratic Residuosity assumptions [Kal05, HK12]. Further, we note that in all these OT constructions, computing the output of the protocol can be represented by an NC^1 circuit and we will use this fact later. Instantiating the OT protocol in the above construction, we get the following corollary:

Corollary 2. *Set $\epsilon = 1/2^m(1 - 2^m \cdot m \cdot \mu_{ot})$. There exists some $c > 0$ such that, assuming μ_{ot}^c secure LWE/DDH/QR/ N^{th} -residuosity, the scheme in Figure 1 is a $(1 - 2^{-m} - \delta_{OT})$ statistically hiding, ϵ -extractable commitment scheme according to Definition 8 where the receiver's algorithms are public coin.*

We now prove the above theorem by showing statistical hiding, computational binding, and extractability in Lemma 1, Lemma 2 and Lemma 3 below respectively.

Lemma 1. *Assuming the underlying OT satisfies $(1 - \delta_{OT})$ statistical sender security according to Definition 4, the scheme in Figure 1 is $(1 - 2^{-m} - \delta_{OT})$ statistically hiding according to Definition 7.*

Proof. Fix any (unbounded) malicious receiver \mathcal{R}^* . Let $m_{\mathcal{R}^*}$ be the message sent by \mathcal{R}^* during the commit phase. By Claim 1, $m_{\mathcal{R}^*}$ uniquely defines a receiver challenge ch . With probability 2^{-m} , $r \neq \text{ch}$ for r chosen uniformly at random by an honest committer. Conditioned on $r \neq \text{ch}$, there exists at least one index $j \in [m]$ such that $r_j \neq \text{ch}_j$. When $r_j \neq \text{ch}_j$, by $(1 - \delta_{OT})$ statistical sender security of OT, $M_j^{r_j}$ is $(1 - \delta_{OT})$ -statistically hidden from any malicious receiver. Since $M_j^{r_j}$ is one of the shares in an XOR secret sharing of M , the message M is $(1 - 2^{-m} - \delta_{OT})$ statistically hidden from any malicious receiver. \square

Lemma 2. *Assuming that the underlying OT satisfies receiver security according to Definition 4, the scheme in Figure 1 is computationally binding against non-uniform PPT malicious committers, according to Definition 7.*

Proof. Suppose for contradiction that there exists a non-uniform malicious PPT cheating committer \mathcal{C}^* and a polynomial $p(\cdot)$ such that outputs transcript $\tau \leftarrow \text{Commit}(\mathcal{C}^*, \mathcal{R}), y_1, y_2$ such that $\widetilde{M}_1 = \text{Verify}(\tau, y_1), \widetilde{M}_2 = \text{Verify}(\tau, y_2)$, and

$$\Pr[(\widetilde{M}_1 \neq \perp) \wedge (\widetilde{M}_2 \neq \perp) \wedge (\widetilde{M}_1 \neq \widetilde{M}_2)] = \frac{1}{p(\lambda)}$$

We will construct a reduction \mathcal{A} that has black-box access to such a committer \mathcal{C}^* , and breaks receiver OT security according to Definition 4. \mathcal{A} takes as input $o' = \text{OT}_1(b)$ and is required to distinguish the case when $b = 0$ from when $b = 1$. \mathcal{A} does the following:

1. Sample $\ell \xleftarrow{\mathbb{S}} [m]$, $\{\text{ch}_i\}_{i \in [m] \setminus \{\ell\}} \xleftarrow{\mathbb{S}} \{0, 1\}^{m-1}$.
2. Sample $r_{1,i}$ uniformly at random and compute $\{o_{1,i} = \text{OT}_1(\text{ch}_i; r_{1,i})\}_{i \in [m] \setminus \{\ell\}}$. Set $o_{1,\ell} = o'$.
3. Forward $\{o_{1,i}\}_{i \in [m]}$ as the first message of the scheme in Figure 1 to the adversary \mathcal{C}^* .
4. Obtain $r, \{o_{2,i}\}_{i \in [m]}$ from \mathcal{C}^* . This is the end of the commit phase, denote transcript by τ .
5. Obtain y_1, y_2 from \mathcal{C}^* . If $\text{Verify}(\tau, y_1) = \text{Verify}(\tau, y_2)$ or if either of them are \perp , abort and output \perp . Else continue.
6. Parse y_1 as $M, \{M_i^b, r_{2,i}\}_{i \in [m], b \in \{0,1\}}$ and y_2 as $\widetilde{M}, \{\widetilde{M}_i^b, \widetilde{r}_{2,i}\}_{i \in [m], b \in \{0,1\}}$.
7. Define $\mathbb{S} = \{j \in [m] : M_j^{r_j} \neq \widetilde{M}_j^{r_j}\}$. Since $M \neq \widetilde{M}$, $|\mathbb{S}| > 1$. If $\ell \in \mathbb{S}$, output $1 - r_\ell$, else output \perp .

We prove the following claim, which will contradict OT security according to Definition 4 and complete the proof of the lemma.

Claim 2. *Let b' denote the output of \mathcal{A} . Then, $\Pr[b' = b] \geq \frac{1}{m \cdot p(\lambda)}$.*

Proof of Claim. By assumption, with probability at least $\frac{1}{p(\lambda)}$, the following event \mathbb{E} occurs: \mathcal{A} proceeds to Step 5 and obtains $M \neq \widetilde{M}$.

After \mathcal{A} proceeds to Step 5, it first creates the set \mathbb{S} of indices $j \in [m]$ where $M_j^{r_j} \neq \widetilde{M}_j^{r_j}$. Since $M \neq \widetilde{M}$, $|\mathbb{S}| > 1$. Since \mathcal{A} samples ℓ independently and uniformly at random, $\Pr[\ell \in \mathbb{S} | \mathbb{E}] \geq \frac{1}{m}$. Thus, $\Pr[(\ell \in \mathbb{S}) \wedge \mathbb{E}] \geq \frac{1}{m \cdot p(\lambda)}$.

Let us now condition on $(\ell \in \mathbb{S}) \wedge \mathbb{E}$. For any $i \in [m]$, by correctness of the i^{th} parallel OT, the i^{th} OT statistically binds the sender to a unique input value $M_i^{\text{ch}_i}$. Thus, for any $i \in [m]$, the existence of $M_i^{r_i} \neq \widetilde{M}_i^{r_i}$ that reconstruct to $o_{2,i}$ implies that $\text{ch}_i \neq r_i$ (except with probability $\text{negl}(\lambda)$). Since we conditioned on $\ell \in \mathbb{S}$, we have that $M_\ell^{r_\ell} \neq \widetilde{M}_\ell^{r_\ell}$, and therefore $(b = \text{ch}_\ell) = (1 - r_\ell) = b'$.

Thus, $\Pr[b' = b | (\ell \in \mathbb{S}) \wedge \mathbb{E}] = 1$. This implies that $\Pr[b' = b \wedge (\ell \in \mathbb{S}) \wedge \mathbb{E}] \geq \frac{1}{m \cdot p(\lambda)}$. Since \mathcal{A} outputs \perp if either of the events $(\ell \in \mathbb{S})$ and \mathbb{E} did not occur, we have that $\Pr[b' = b] \geq \frac{1}{m \cdot p(\lambda)}$, proving the claim. \square

Lemma 3. *The scheme in Figure 1 is an ϵ -extractable commitment scheme, where $\epsilon = 2^{-m}(1 - 2^m \cdot \mu_{\text{OT}})$.*

Proof. We begin by describing the extractor \mathcal{E} from Definition 8. We denote the first message of transcript τ by τ_1 and the second message by τ_2 . In the following, we assume that if the committer's message τ_2 is malformed or blank then it is treated as if r and \widetilde{M}_i for $i \in [m]$ are the all-zeros string. \mathcal{E} obtains transcript (τ_1, τ_2) together with $\text{state}_{\mathcal{R}, \tau}$, and does the following:

- Parse $\tau_2 = (r, \{o_{2,i}\}_{i \in [m]})$.
- If $r \neq \text{ch}$, output \perp .
- Else, use $\text{state}_{\mathcal{R}, \tau}$ to obtain $\{\widetilde{M}_i^{\text{ch}_i}\}_{i \in [m]}$ from $\{o_{2,i}\}_{i \in [m]}$.
- Compute $\widetilde{M} = \bigoplus_{i \in [m]} \widetilde{M}_i^{\text{ch}_i}$ and output \widetilde{M} .

We will now analyze the extractor \mathcal{E} . First, before describing the properties, we quickly observe that the extractor can be represented by an NC^1 circuit since computing the output of the OT protocol can be represented by an NC^1 circuit.

Frequency of Extraction. Recall that extraction succeeds from a trial when the r chosen by the committer is equal to the ch chosen by the extractor in this trial. We now prove the following claim, which asserts that the event $r = \text{ch}$ occurs with probability at least ϵ .

Claim 3. *The probability (over the coins of \mathcal{R}) that $r = \text{ch}$ is at least $2^{-m} \cdot (1 - 2^m \cdot m \cdot \mu_{OT})$.*

Proof of Claim. Suppose the claim is not true. We will use this to contradict the receiver security of the OT protocol. To do so, we describe our reduction algorithm $\hat{\mathcal{A}}$ and distinguisher $\hat{\mathcal{D}}$ for the receiver security game.

The reduction $\hat{\mathcal{A}}$ picks two challenges $\text{ch}_1, \text{ch}_2 \xleftarrow{\$} \{0, 1\}^m$ at random, and creates auxiliary information consisting of these two challenges. We will use $\text{OT}_1(\text{ch}; R)$ to denote $\{\text{OT}_1(\text{ch}_i; R_i)\}_{i \in [m]}$, created using randomness R .

Now, $\hat{\mathcal{D}}$ will obtain from the OT challenger as challenge either $\tau_1 = \text{OT}_1(\text{ch}_1; R)$ or $\tau_1 = \text{OT}_1(\text{ch}_2; R)$. $\hat{\mathcal{D}}$ runs the malicious sender \mathcal{C}^* on the message τ_1 , obtaining τ_2 . It obtains r from τ_2 : if $r = \text{ch}_1$, it outputs 1. Otherwise, it aborts and outputs \perp .

We will now analyze the probability that $\hat{\mathcal{D}}$ outputs 1 in the two cases where $\text{ch} = \text{ch}_1$ or $\text{ch} = \text{ch}_2$. If $\text{ch} = \text{ch}_1$, then by assumption, we have that $\Pr[\hat{\mathcal{D}} = 1 | \text{ch} = \text{ch}_1] < 2^{-m} \cdot (1 - m \cdot 2^m \cdot m \cdot \mu_{OT})$. On the other hand, if $\text{ch} = \text{ch}_2$, then no information about ch_1 is given to the distinguisher \mathcal{D} . Therefore, $\Pr[\hat{\mathcal{D}} = 1 | \text{ch} = \text{ch}_2] = 2^{-m}$.

Thus, we have that $|\Pr[\hat{\mathcal{D}} = 1 | \text{ch} = \text{ch}_1] - \Pr[\hat{\mathcal{D}} = 1 | \text{ch} = \text{ch}_2]| \geq m \cdot \mu_{OT}$, which is a contradiction since the underlying single instance OT is μ_{OT} -secure against malicious senders. \square

Claim 4. *Let μ_{OT} be an efficiently computable negligible function. For any constant $c \in (0, 1)$, set $m = c \cdot \log_2 \mu_{OT}^{-1}$, then in the scheme described above, $\epsilon > \mu_{OT}^c / 2$.*

Proof. The previous claim shows that $\epsilon = 2^{-m} (1 - 2^m \cdot m \cdot \mu_{OT})$. Choose m such that $2^{-m} = \mu_{OT}^c$. This implies that $m = c \cdot \log_2 \mu_{OT}^{-1}$. Then, plugging this value of m above gives $\epsilon = \mu_{OT}^c (1 - m \cdot \mu_{OT}^{1-c})$. As μ_{OT}^{1-c} is a negligible function, $\epsilon > \mu_{OT}^c / 2$. \square

Correctness of Extraction. The following claim proves that the extraction is “correct”.

Claim 5. *Whenever the extraction algorithm $\mathcal{E}(\tau, \text{state}_{\tau, \mathcal{R}})$ outputs $\widetilde{M} \neq \perp$, the following holds: The transcript (τ_1, τ_2) statistically binds \mathcal{C}^* to either an invalid message, or a single message M , such that $\widetilde{M} = M$.*

Proof. First, when $r \neq \text{ch}$, the extractor outputs \perp .

When $r = \text{ch}$, the extractor obtains $\{M_i^{\text{ch}_i}\}_{i \in [m]}$ as OT output, and outputs $\widetilde{M} = \bigoplus_{i \in [m]} M_i^{\text{ch}_i}$. By correctness of OT and since $r = \text{ch}$, $\widetilde{M} = \bigoplus_{i \in [m]} M_i^{r_i}$, correctness of extracted value follows. Moreover, by correctness of OT, the committer is statistically bound to a single input for each index $i \in [m]$, and therefore to a unique message in all transcripts where $r = \text{ch}$. \square

Indistinguishability of Extractable Transcripts. Towards a contradiction, suppose there is a committer \mathcal{C}^* and a polynomial $p(\cdot)$ such that:

$$|\Pr[\mathcal{C}^*(\tau) = 1 | \widetilde{M} = \perp] - \Pr[\mathcal{C}^*(\tau) = 1 | \widetilde{M} \neq \perp]| > \frac{1}{p(\lambda)}$$

We will use this committer to contradict the receiver security of the OT protocol. We now describe our reduction algorithm $\hat{\mathcal{A}}$ and distinguisher $\hat{\mathcal{D}}$ for the receiver security game.

The reduction $\hat{\mathcal{A}}$ chooses two challenges $\text{ch}_1, \text{ch}_2 \xleftarrow{\$} \{0, 1\}^m$ at random, and creates auxiliary information consisting of these two challenges. Next, $\hat{\mathcal{A}}$ obtains as input (from the OT challenger) either $\tau_1 = \text{OT}_1(\text{ch} = \text{ch}_1; R)$ or $\tau_1 = \text{OT}_1(\text{ch} = \text{ch}_2; R)$.

\mathcal{A} now runs the malicious sender \mathcal{C}^* on the message τ_1 , obtaining τ_2, aux . It generates the joint distribution $(\tau_1, \tau_2, \text{aux})$. If $r \neq \text{ch}_1$, it outputs $\mathcal{C}^*(\tau_1, \tau_2, \text{aux})$. Otherwise, it aborts and outputs \perp .

Note that $\widetilde{M} \neq \perp \iff r = \text{ch}$, and therefore, we have that

$$|\Pr[\mathcal{C}^*(\tau) = 1 | r \neq \text{ch}] - \Pr[\mathcal{C}^*(\tau) = 1 | r = \text{ch}]| > \frac{1}{p(\lambda)}$$

We now analyze two cases.

- Suppose $\text{ch} = \text{ch}_1$. Then, $r \neq \text{ch} \implies r \neq \text{ch}_1$, and $r = \text{ch} \implies r = \text{ch}_1$. In this case,

$$|\Pr[\mathcal{C}^*(\tau) = 1 | r \neq \text{ch}_1, \text{ch} = \text{ch}_1] - \Pr[\mathcal{C}^*(\tau) = 1 | r = \text{ch}_1, \text{ch} = \text{ch}_1]| > \frac{1}{p(\lambda)}$$

which implies that

$$\begin{aligned} & \Pr[\hat{\mathcal{D}}(\tau) = 1 | \text{ch} = \text{ch}_1] = \\ & \Pr[\hat{\mathcal{D}}(\tau) = 1 \wedge r \neq \text{ch}_1 | \text{ch} = \text{ch}_1] > \\ & \Pr[\mathcal{C}^*(\tau) = 1 \wedge r \neq \text{ch}_1 | \text{ch} = \text{ch}_1] = \\ & \Pr[\mathcal{C}^*(\tau) = 1 | r \neq \text{ch}_1, \text{ch} = \text{ch}_1] \cdot \Pr[r \neq \text{ch}_1 | \text{ch} = \text{ch}_1] > \\ & \frac{1}{p(\lambda)} \cdot (1 - 2^{-m} - \epsilon) > \frac{1}{2p(\lambda)}. \end{aligned}$$

- Suppose $\text{ch} = \text{ch}_2$. In this case, no information about ch_1 is given to \mathcal{C}^* , and therefore

$$\Pr[\mathcal{D}^*(\tau) = 1 | \text{ch} = \text{ch}_2] \leq \Pr[\mathcal{C}^*(\tau) = 1 \wedge r \neq \text{ch}_1 | \text{ch} = \text{ch}_2] = 2^{-m} \cdot 1 - 2^{-m}$$

Therefore,

$$|\Pr[\mathcal{D}^*(\tau) = 1 | \text{ch} = \text{ch}_1] - \Pr[\mathcal{D}^*(\tau) = 1 | \text{ch} = \text{ch}_2]| > \frac{1}{2p(\lambda)} - 2^{-m} > \frac{1}{3p(\lambda)}$$

This contradicts receiver security, and therefore the lemma follows. \square

We will rely on a simple variant of this commitment scheme that for any $\ell = \text{poly}(\lambda)$, results in an ϵ -extractable sequence of ℓ commitments where $\epsilon > \mu_{OT}^c/2$ (independent of ℓ) for any constant $c \in (0, 1)$. This variant requires the committer to sample a *single* random string $r \xleftarrow{\$} \{0, 1\}^m$ (refer to Step 1 of the committer message), and reuse the same string r across all ℓ commitments. Since the committer outputs r in the clear, a receiver obtaining such a sequence can efficiently check, during the verify phase of decommitment, whether all ℓ commitments use the same randomness r , and output \perp if this is not the case. This sequence of ℓ commitments is statistically hiding by a simple hybrid argument, and is computationally binding for the same reason as Lemma 2.

The extractor \mathcal{E} is modified so that on input a transcript of ℓ commitments, it outputs 0^ℓ if the random string r is not identical across all commitments. Therefore, the extractor outputs $(\widetilde{M}_1, \dots, \widetilde{M}_\ell)$, where there exists an $i \in [\ell]$ such that $\widetilde{M}_i = \perp$ if and only if the committer used r

identically across all commitments, and $\text{ch} \neq r$. Therefore, frequency of extraction follows by Claim 3, correctness of extraction follows by Claim 5 together with the fact that when r is not identical across all commitments, the `Verify` algorithm outputs \perp . Finally, indistinguishability of extractable transcripts follows identically as in the case of a single commitment. Because $\text{poly}(\lambda) \cdot \mu_{OT}^c < \mu_{OT}^{c'}$ for $c' < c$ for any constant $c, c' \in (0, 1)$ and any polynomial poly we have that, setting $m = c' \log \mu_{OT}^{-1}$ for any $c' < c$ for this construction gives us the following corollary.

Corollary 3. *Let μ_{OT} be an efficiently computable negligible function. For any $\ell = \text{poly}(\lambda)$ and any $c \in (0, 1)$, there exists an ϵ -extractable sequence of ℓ statistically hiding commitments, where $\epsilon > \mu_{OT}^c/2$, assuming $\mu_{OT}^{O(1)}$ -secure LWE/DDH/QR/ N^{th} -residuosity⁶. Further, the extractor \mathcal{E} can be represented by an NC^1 circuit.*

5 Our Statistical WI Protocol

5.1 Modified Blum Protocol

We begin by describing a very simple modification to the Blum Σ -protocol for Graph Hamiltonicity. The protocol we describe will have soundness error $\frac{1}{2} - \text{negl}(\lambda)$ against adaptive PPT provers, and will satisfy *statistical* zero-knowledge. Since Graph Hamiltonicity is NP-complete, this protocol can also be used to prove any statement in NP via a Karp reduction. This protocol is described in Figure 2.

We give an overview of the protocol here. Note that the only modification to the original protocol of Blum [Blu86] is that we use two message statistically hiding, extractable commitments instead of non-interactive statistically binding commitments. The proofs of soundness and statistical honest-verifier zero-knowledge are fairly straightforward. They roughly follow the same structure as [Blu86], replacing statistically binding commitments with statistically hiding commitments.

Lemma 4. *Assuming that `extcom` is computationally binding, the protocol in Figure 2 satisfies soundness against PPT provers that may choose x adaptively in the second round of the protocol.*

Proof. The proof of soundness follows by the computational binding property of `extcom` and the soundness of the (original) Blum protocol.

Let L denote the language consisting of all graphs that have a Hamiltonian cycle. Consider a cheating prover P^* that convinces a malicious verifier about a statement $x \notin L$ with probability $\frac{1}{2} + h(n)$, where $h(\cdot) > \frac{1}{\text{poly}(\cdot)}$ for some polynomial $\text{poly}(\cdot)$. By an averaging argument, this means that there exists at least one transcript prefix τ consisting of the first two messages of the protocol, where for $G \notin L$ sent by the prover in the third message, $\Pr[V \text{ accepts} | \tau, G \notin L] > \frac{1}{2}$. This implies that there exists a cheating prover that generates a transcript prefix τ , for which it provides an accepting opening corresponding to both $b = 0$ and $b = 1$, with probability at least $h(n)$. Next, we argue that such a cheating prover must break the (computational) binding of `com`.

Since $G \notin L$, it is information theoretically impossible for any cheating prover to generate a commitment to a unique string $\pi, \pi(G)$ such that there exists a Hamiltonian cycle in $\pi(G)$. Therefore, any prover that opens a transcript prefix τ, G corresponding to both $b = 0$ and $b = 1$ for $G \notin L$, must open at least one commitment in the set $\{\text{extcom}_P, \{\text{extcom}_{i,j}\}_{i,j \in p \times p}\}$ to two different values, thereby giving a contradiction to the binding of the commitment scheme. \square

Lemma 5. *Assuming that `extcom` is statistically hiding, the protocol in Figure 2 satisfies honest-verifier statistical zero-knowledge.*

⁶The constant $O(1)$ indicates the constant required to build μ_{OT} -secure OT secure against malicious senders.

Modified Blum Argument

1. **Verifier Message:** The verifier does the following:
 - Sample the first message $\text{extcom}_{1,i,j}$ for independent instances of the extractable commitment, where $i, j \in [p(\lambda)] \times [p(\lambda)]$, *uniformly at random*.
 - Send an additional first message $\text{extcom}_{1,P}$ for another independent instance of the extractable commitment, again sampled *uniformly at random*.
2. **Prover Message:** The prover gets input graph $G \in \{0,1\}^{p(\lambda) \times p(\lambda)}$ represented as an adjacency matrix, with $(i,j)^{\text{th}}$ entry denoted by $G[i][j]$, Hamiltonian cycle $H \subseteq G$. Here $p(\cdot)$ is an a-priori fixed polynomial. The prover does the following:
 - Sample a random permutation π on $p(\lambda)$ nodes, and compute $c_P = \text{extcom}_{2,P}(\pi)$ as a commitment to π using extcom .
 - Compute $\pi(G)$, which is the adjacency matrix corresponding to the graph G when its nodes are permuted according to π . Compute $c_{i,j} = \text{extcom}_{2,i,j}(\pi(G)[i][j])$ for $(i,j) \in [p(\lambda)] \times [p(\lambda)]$.
 - Send $G, c_P, c_{i,j}$ for $(i,j) \in [p(\lambda)] \times [p(\lambda)]$.
3. **Verifier Message:** Sample and send $c \xleftarrow{\$} \{0,1\}$ to the prover.
4. **Prover Message:** The prover does the following:
 - If $c = 0$, send π and the decommitments of $\text{extcom}_P, \text{extcom}_{i,j}$ for $(i,j) \in [p(\lambda)] \times [p(\lambda)]$.
 - If $c = 1$, send the decommitment of $\text{extcom}_{i,j}$ for all (i,j) such that $\pi(H)[i][j] = 1$.
5. **Verifier Output:** The verifier does the following:
 - If $c = 0$, accept if and only if all extcom openings were accepted and $\pi(G)$ was computed correctly by applying π on G .
 - If $c = 1$, accept if and only if all extcom openings were accepted and all the opened commitments form a Hamiltonian cycle.

Remark: Observe that since the receiver's algorithms in the extractable commitment scheme are public coin, the above protocol is also public coin.

Figure 2: Modified Blum SZK Argument

Proof. The simulation strategy is identical to that of [Blu86]. The simulator Sim first guesses the challenge bit c' . It begins an interaction with the malicious verifier. On obtaining the first message from the verifier, if $c' = 0$, it samples π uniformly at random and generates a commitment to $\pi, \pi(G)$ following honest prover strategy to generate the commitment. If $c' = 1$, it samples π, H' uniformly at random where H' is an arbitrary hamiltonian cycle, and generates a commitment to $\pi, \pi(H')$ following honest prover strategy to generate the commitment. Next, it waits for the verifier to send c , and if $c \neq c'$, it aborts and repeats the experiment. If $c = c'$, then it decommits to the commitments according to honest prover strategy.

Note that when $c = c' = 1$, the resulting simulation is perfect zero-knowledge since the simulated view of the verifier is identical to the view generated by an honest prover. On the other hand when $c = c' = 0$, it follows from the statistical hiding property of the commitment `extcom` that the verifier cannot distinguish the case where `extcom` is a commitment to $\pi, \pi(G)$ and a hamiltonian cycle is opened in $\pi(G)$, from the case where `extcom` is not a commitment to $\pi(G)$, but instead to some $\pi(H')$ for a hamiltonian cycle H' . \square

Since honest-verifier zero-knowledge composes under parallel repetition, we can repeat the protocol several times in parallel to get negligible soundness error. Formally, we have the following lemma:

Lemma 6. *Assuming that `extcom` is statistically hiding, the protocol in Figure 2 satisfies honest verifier statistical zero-knowledge under parallel repetition.*

Finally, Cramer et al. [CDS94] showed that honest verifier zero knowledge where the receiver's algorithms are public coin implies witness indistinguishability even against malicious verifiers. As a result, we get the following lemma:

Lemma 7. *Assuming that `extcom` is statistically hiding, the protocol in Figure 2 satisfies statistical witness indistinguishability under parallel repetition.*

5.2 Statistical ZAPs

In this section, we prove the following theorem:

Theorem 4. *There exists a two message public-coin statistical witness indistinguishable argument system for NP in the plain model assuming that the following primitives exist:*

- *Two-message oblivious transfer (OT) that is μ_{OT} secure against malicious senders, satisfying Definition 4 and Property 1, and*
- *μ_{CI} -correlation intractable hash functions,*

where μ_{CI} and μ_{OT} are efficiently computable negligible functions in λ .

Recall from previous sections that we can use the above OT to build the extractable commitment which is then used to build a four message Σ -protocol that is a modification to Blum's protocol. As mentioned before, we can instantiate both the OT and the correlation intractable hash function assuming the learning with errors (LWE) problem is hard. More specifically, assuming any PPT adversary has advantage at most μ_{LWE} for LWE, where μ_{LWE} is an efficiently computable function, we can efficiently compute upper bounds μ_{CI} and μ_{OT} on the advantages of the CI and OT constructions, respectively, as well.

Therefore, instantiating both the primitives in the above theorem gives us the following:

Theorem 5. *Assuming μ_{LWE} -hard LWE for some efficiently computable negligible function μ_{LWE} , there exists a two message public-coin statistical witness indistinguishable argument system for NP in the plain model.*

Notations and Primitives used.

- Let λ be the security parameter.
- Let $\mu = \max\{\mu_{OT}, \mu_{CI}\}$.

- Let $\Sigma := (\Sigma_1, \dots, \Sigma_\lambda)$ denote λ parallel repetitions of the modified Blum Sigma protocol constructed in Section 5.1, where for $i \in [\ell]$, $\Sigma_i = (q_i, a_i, e_i, z_i)$. Let the underlying commitment scheme be instantiated with $m = c \log \mu$, where $0 < c < 1$.
- Let \mathcal{H} be a correlation intractable hash function with respect to $\{\mathcal{F}_{\lambda, s(\lambda)}\}_{\lambda \in \mathbb{N}}$ according to Definition 3 that outputs strings of length λ , where $s(\lambda) = 2s_1(\lambda)$ where s_1 is the size of the extractor \mathcal{E} used in the commitment scheme and \mathcal{F} denotes the class of all NC^1 circuits of size $s(\lambda)$ as defined in Definition 2. Recall the correlation-intractability advantage is assumed to be at most μ_{CI} .

Construction. Let x be any instance in $\{0, 1\}^\lambda$ and let w be the corresponding witness for the statement $x \in L$.

1. Verifier's message to the Prover:

- Sample $q := \{q_i\}_{i \in [\lambda]}$.
- Sample $K \leftarrow \mathcal{H}.\text{Setup}(1^\lambda, 0^\ell)$.
- Output (q, K) .

2. Prover's message to the Verifier:

- Compute $\{a_i\}_{i \in [\lambda]}$ as a response to $\{q_i\}_{i \in [\lambda]}$.
- Compute $e \leftarrow \mathcal{H}.\text{Eval}(K, a)$.
- Compute $\{z_i\}_{i \in [\lambda]}$ with respect to the challenge string e .
- Output (x, a, e, z) .

3. Verification: The verifier does the following:

- If $\mathcal{H}.\text{Eval}(K, a) \neq e$, output reject.
- Else if (x, q, a, e, z) does not verify according to the Σ protocol, output reject.
- Else output accept.

Completeness. Completeness of the protocol can be easily observed from the correctness of the underlying primitives: the protocol Σ and the hash function H .

Public Coin. Recall from the statistical indistinguishability of hash keys property that an honest verifier can just sample a uniformly random string as the hash key K . This, along with the fact that the underlying protocol Σ is public coin results in the above protocol also being public coin.

Soundness. We now prove computational soundness of the protocol above. Towards a contradiction, fix any adversary \mathcal{A} that breaks soundness of the protocol with probability $\frac{1}{p(\lambda)}$ for some polynomial $p(\cdot)$.

We consider a sequence of hybrids where the first hybrid corresponds to the real soundness experiment.

- **Hybrid₀** : This hybrid corresponds to the experiment where the challenger behaves identically to the verifier in the actual protocol.

- **Hybrid₁**: In this hybrid, instead of generating the verifier’s first message as uniformly random string, the challenger **Ch** now computes the first message of the extractable commitment scheme used in the underlying protocol Σ as done in the protocol description in Figure 1. In particular, the underlying OT receiver messages are not sampled as uniformly random strings but instead are computed by running the OT receiver algorithm. As a result, **Ch** now has some internal state r_{state} as part of the extractable commitment scheme that is not public.
- **Hybrid₂**: This hybrid is the same as the previous hybrid except that the hash key K is generated as follows. $K \leftarrow \mathcal{H}.\text{Setup}(1^\lambda, R_{x,q})$ where the relation $R_{x,q}$ consists of tuples of the form (a, y) where y is computed by an efficient function $f_{bad,x,q}$ described below. $f_{bad,x,q}$ has the statement x , verifier’s first messages q and secret state r_{state} corresponding to q hardwired, takes as input the prover’s message a and does the following.
 1. Run the extractor algorithm \mathcal{E} on input $(r_{\text{state}}, \tau = (q, a))$ to compute m . Note that \mathcal{E} can be represented by an NC^1 circuit of size $s_1(\lambda)$ for some polynomial s_1 .
 2. If $m \neq \perp$, this means that m is the tuple of messages committed to in the set of λ commitment tuples $(c_P, \{c_{i,j}\})$. For each $k \in [\lambda]$, check whether the message committed to by the tuple $\{c_{i,j}\}$ is indeed equal to $\pi(G)$ where π is the permutation committed to in c_P . If so, then set $e_k = 0$ and else set $e_k = 1$. Set $y = (e_1, \dots, e_\lambda)$.⁷
 3. If $m = \perp$, set $y = 0^\lambda$.

Before proving the soundness of the protocol using the hybrids, we define an event that helps us in the proof.

Event **E**: Let τ denote the transcript of an execution of the above protocol and let $\tau_{\mathcal{C}}$ denote the transcript of the commitment scheme in the execution. Let $\text{state}_{\mathcal{R}}$ denote the state of the verifier when it runs the receiver algorithm of the commitment scheme. We will say that the event **E** occurs if for any honest verifier V :

$$[V(\tau) = 1 \wedge \mathcal{E}(\tau_{\mathcal{C}}, \text{state}_{\mathcal{R}}) \neq \perp].$$

We now continue the proof of soundness with the following claims.

Lemma 8. *Assuming the pseudorandomness of receiver messages of the OT protocol used in the underlying extractable commitment scheme (Property 1),*

$$|\Pr[V(\tau) = 1 | \text{Hybrid}_1] - \Pr[V(\tau) = 1 | \text{Hybrid}_0]| = \text{negl}(\lambda)$$

Proof. The only difference between the two hybrids is that in **Hybrid₀**, the OT receiver messages in the extractable commitment scheme used in the underlying protocol Σ are generated as uniformly random strings while in **Hybrid₁**, they are generated by running the algorithm OT_1 on behalf of the OT receiver. It is easy to see that if the difference in the adversary’s success probability in breaking soundness between these two hybrids is non-negligible, we can break the pseudorandomness of receiver messages property (Property 1) of the underlying two message OT protocol, which is a contradiction. \square

Lemma 9. *Assuming the frequency of extraction property and the indistinguishability of extractable transcripts property of the extractable commitment scheme, there exists a polynomial $p(\cdot)$ such that*

$$\Pr[\mathbf{E} \text{ occurs in Hybrid}_1] \geq \epsilon \cdot \frac{1}{p(\lambda)},$$

⁷Essentially, since $x \notin L$, if the cheating prover has to succeed, it can either generate a successful response z_k for verifier’s query bit $e_k = 0$ or $e_k = 1$ and this function determines which bit it is.

where the probability is over the randomness of V , and where ϵ is the extraction probability of the underlying commitment scheme.

Proof. Fix $x \notin L$. We will consider a reduction \mathcal{B} that interacts with the adversary and relies on the frequency of extraction property and the indistinguishability of extractable transcripts property of the extractable commitment scheme to prove the lemma.

\mathcal{B} interacts with a challenger Ch for the commitment scheme and receives a first round message com_1 for the ℓ -extractable commitment scheme. It then interacts with the adversary \mathcal{A} as the verifier in the ZAP protocol, setting com_1 as its message on behalf of the receiver in the underlying commitment scheme, and sampling the hash key $K \leftarrow \mathcal{H}.\text{Setup}(1^\lambda, 0^\ell)$. After completing the protocol execution with \mathcal{A} , \mathcal{B} forwards the commitments sent by \mathcal{A} as its message com_2 of the commitment scheme to the challenger Ch . Further, \mathcal{B} outputs 1 in its interaction with Ch if the proof provided by \mathcal{A} verifies, and 0 otherwise.

Let τ denote the transcript of the ZAP protocol and $\tau_{\mathcal{C}}$ the transcript of the underlying commitment scheme. Let state_r be the state of the receiver in the commitment scheme as sampled by the challenger Ch .

First, we observe that by Lemma 8, there exists a polynomial $p(\cdot)$ such that adversary \mathcal{A} breaks the soundness property in Hybrid_1 with non-negligible probability $\frac{1}{p(\lambda)}$. This implies that $\Pr[\mathcal{B}(\tau_{\mathcal{C}}) = 1] \geq \frac{1}{p(\lambda)}$ over the random coins of \mathcal{B}, Ch . This gives us the following equation.

$$\begin{aligned} \Pr[\mathcal{B}(\tau_{\mathcal{C}}) = 1] &= (\Pr[\mathcal{B}(\tau_{\mathcal{C}}) = 1 \mid \mathcal{E}(\tau_{\mathcal{C}}, \text{state}_{\mathcal{R}}) \neq \perp] \cdot \Pr[\mathcal{E}(\tau_{\mathcal{C}}, \text{state}_{\mathcal{R}}) \neq \perp] \\ &\quad + \Pr[\mathcal{B}(\tau_{\mathcal{C}}) = 1 \mid \mathcal{E}(\tau_{\mathcal{C}}, \text{state}_{\mathcal{R}}) = \perp] \cdot \Pr[\mathcal{E}(\tau_{\mathcal{C}}, \text{state}_{\mathcal{R}}) = \perp]) \geq \frac{1}{p(\lambda)} \end{aligned} \quad (1)$$

From the indistinguishability of extractable transcripts property, we have that:

$$|\Pr[\mathcal{B}(\tau_{\mathcal{C}}) = 1 \mid \mathcal{E}(\tau_{\mathcal{C}}, \text{state}_{\mathcal{R}}) \neq \perp] - \Pr[\mathcal{B}(\tau_{\mathcal{C}}) = 1 \mid \mathcal{E}(\tau_{\mathcal{C}}, \text{state}_{\mathcal{R}}) = \perp]| = \text{negl}(\lambda) \quad (2)$$

From the frequency of extraction property, we have that :

$$\Pr[\mathcal{E}(\tau_{\mathcal{C}}, \text{state}_{\mathcal{R}}) \neq \perp] \geq \epsilon \quad (3)$$

where all equations are over the random coins of the challenger Ch and reduction \mathcal{B} . Combining Equations (1) and (2) implies that there exists a polynomial $q(\cdot)$ such that

$$\Pr[\mathcal{B}(\tau_{\mathcal{C}}) = 1 \mid \mathcal{E}(\tau_{\mathcal{C}}, \text{state}_{\mathcal{R}}) \neq \perp] \geq \frac{1}{q(\lambda)},$$

which, by Equation (3), implies that

$$\Pr[\mathcal{B}(\tau) = 1 \wedge \mathcal{E}(\tau_{\mathcal{C}}, \text{state}_{\mathcal{R}}) \neq \perp] = \Pr[\mathcal{B}(\tau_{\mathcal{C}}) = 1 \mid \mathcal{E}(\tau_{\mathcal{C}}, \text{state}_{\mathcal{R}}) \neq \perp] \cdot \Pr[\mathcal{E}(\tau_{\mathcal{C}}, \text{state}_{\mathcal{R}}) \neq \perp] \geq \frac{1}{q(\lambda)} \cdot \epsilon.$$

Thus we have

$$\Pr[\mathbf{E} \text{ occurs in Hybrid}_1] \geq \epsilon \cdot \frac{1}{q(\lambda)}.$$

This completes the proof of the Lemma. \square

Lemma 10. *Assuming the statistical indistinguishability of hash keys of the correlation intractable hash function, there exists a polynomial $p(\cdot)$ such that*

$$\Pr[\mathbf{E} \text{ occurs in Hybrid}_2] \geq \epsilon \cdot \frac{1}{p(\lambda)},$$

where the probability is over the randomness of V , and where ϵ is the extraction probability of the underlying commitment.

Proof. Assume for the sake of contradiction that the lemma is not true. We will show that we can break the statistical indistinguishability of hash keys property of the correlation intractable hash function.

We will design a reduction \mathcal{B} that interacts with \mathcal{A} , where \mathcal{B} acts as verifier in the above ZAP protocol. \mathcal{B} interacts with a challenger Ch for the correlation intractable hash function. Initially, \mathcal{B} samples the first round message q for the underlying Sigma protocol just as in Hybrid₁, along with associated receiver state $\text{state}_{\mathcal{R}}$ for the commitment scheme, and sends both to Ch. \mathcal{B} obtains a hash key K sampled either uniformly at random (as in Hybrid₁) or by running the setup algorithm of the hash function as described in Hybrid₂. \mathcal{B} uses this key K in its interaction with the adversary \mathcal{A} and completes executing the ZAP protocol. Observe that if Ch sampled a hash key uniformly at random, the interaction between \mathcal{A} and \mathcal{B} is identical to Hybrid₁ and if Ch sampled as hash key as described in Hybrid₂, the interaction between \mathcal{A} and \mathcal{B} is identical to Hybrid₂.

Now, \mathcal{B} tests if event \mathbf{E} occurs. That is, it checks if the ZAP protocol verifies and if so, runs the extractor $\mathcal{E}(\tau_{\mathcal{C}}, \text{state}_{\mathcal{R}})$ using the transcript $\tau_{\mathcal{C}}$ for the commitment scheme. If the extractor cE does not output \perp , then event \mathbf{E} occurs and \mathcal{B} guesses that the hash key was uniformly sampled in its interaction with the challenger Ch. Otherwise, it guesses that the hash key was not uniformly sampled. Thus, if the event \mathbf{E} occurs with probability $\geq \epsilon \cdot \frac{1}{p(\lambda)}$ in Hybrid₁, and occurs with probability $\epsilon \cdot \text{negl}(\lambda)$ in Hybrid₂, \mathcal{B} can distinguish between the hash keys with advantage $\frac{\epsilon}{q(\lambda)}$ for some polynomial q . This is a contradiction, and this completes the proof of the lemma. \square

Lemma 11. *Assuming the μ_{CI} -correlation intractable property of the hash function, the soundness of the underlying protocol Σ and the correctness of extraction of the extractable commitment scheme,*

$$\Pr[\mathbf{E} \text{ occurs in Hybrid}_2] \leq \epsilon \cdot \text{negl}(\lambda).$$

Proof. Suppose the claim is not true. This implies that

$$\Pr[V(\tau) = 1 \wedge \mathcal{E}(\tau_{\mathcal{C}}, \text{state}_{\mathcal{R}}) \neq \perp] = \epsilon \cdot \frac{1}{p(\lambda)}$$

for some polynomial p . Let us consider any transcript on which event \mathbf{E} occurs. Let (q, K) denote the verifier's message and (x, a, e, z) denote the prover's message. Then, from the correctness of the ZAP protocol, it must be the case that (q, a, e, z) verifies according to protocol Σ and $e = H(K, a)$. Further, since the extractor \mathcal{E} succeeds on this transcript, the commitment scheme is statistically binding. Therefore, we can invoke the special soundness of the underlying modified Blum Σ protocol (as in the case of the regular Blum protocol) to state that for the statement $x \notin L$ and prefix (q, a) there can exist at most one pair (e^*, z^*) such that (q, a, e^*, z^*) verifies successfully. Therefore, the adversary's message e must be equal to this value e^* .

Now, from the description of the relation R used in defining the hash key K in Hybrid₂, we observe that, by the correctness of extraction, $f_{bad,x,q}(a) = e^* = H(K, a)$. Thus, for any transcript that satisfies the conditions in event \mathbf{E} , $f_{bad,x,q}(a) = e^* = H(K, a)$. Fix q which maximizes the probability of \mathbf{E} occurring. We can build a reduction \mathcal{B} that on input K , using the adversary \mathcal{A} , produces a such that $f_{bad,x,q}(a) = e^* = H(K, a)$ with probability at least $\epsilon \cdot \frac{1}{p(\lambda)} > \mu^c \cdot \frac{1}{2p(\lambda)} > \mu \geq \mu_{CI}$. Since by Definition 3 the advantage of any polynomial-time adversary in this game must be at most μ_{CI} , this yields a contradiction. \square

Note that Lemma 10 and Lemma 11 contradict each other, and therefore the adversary does not break soundness in the real experiment. This completes the proof of soundness. \square

Statistical Witness Indistinguishability. Let \mathcal{A} denote the unbounded time adversarial verifier and Ch denote the challenger. Let x be the challenge instance of length λ and w_0 and w_1 be a pair

of witnesses for $x \in L$. Consider a pair of hybrids where the first hybrid Hybrid_0 corresponds to Ch running the honest prover algorithm with witness w_0 being used and the second hybrid Hybrid_1 corresponds to Ch running the honest prover algorithm with witness w_1 being used. We now show that these two hybrids are statistically indistinguishable to complete the proof.

Claim 6. *Assuming the Σ -protocol is statistically witness indistinguishable, Hybrid_0 is statistically indistinguishable from Hybrid_1 .*

Proof. We now show that if there exists an unbounded time adversary \mathcal{A} for which the two hybrids are not statistically indistinguishable, we can build a reduction \mathcal{B} that can break the witness indistinguishability of the underlying modified Blum’s Sigma protocol which is a contradiction to Lemma 7. \mathcal{B} acts as the challenger in its interaction with the adversary \mathcal{A} that is trying to distinguish between these two hybrids. Further, \mathcal{B} acts as the adversary in its interaction with a challenger \mathcal{C} in trying to break the WI property of the modified Blum Sigma protocol. Initially, \mathcal{A} sends a statement x , a pair of witnesses (w_0, w_1) and a first round message (q, K) for the above ZAP construction. \mathcal{B} forwards (x, w_0, w_1) to the challenger \mathcal{C} and sends q as its first message of the underlying protocol Σ . \mathcal{C} responds with its round two message a on behalf of the prover. \mathcal{B} computes $e \leftarrow \mathcal{H}.\text{Eval}(K, x, (q, a))$ and sends it to \mathcal{C} . Finally, \mathcal{C} responds with the last round message z on behalf of the prover. Now, \mathcal{B} sends the tuple (x, a, e, z) to \mathcal{A} as the prover message for the above ZAP protocol. Observe that if the challenger \mathcal{C} interacted using witness w_0 , then the interaction between the reduction \mathcal{B} and the adversary \mathcal{A} is identical to Hybrid_0 and if the challenger \mathcal{C} interacted using witness w_1 , then the interaction between the reduction \mathcal{B} and the adversary \mathcal{A} is identical to Hybrid_1 . Thus, if these two hybrids are not statistically indistinguishable to \mathcal{A} , \mathcal{B} can use the same guess used by \mathcal{A} to distinguish them, to break the statistical witness indistinguishability property of the protocol Σ which is a contradiction. \square

6 Statistical ZK Arguments with Super-Polynomial Simulation

In this section, we prove that the protocol given in Section 5.2, essentially unmodified, is a statistical zero-knowledge argument with super-polynomial simulation assuming subexponential LWE, meaning that the simulator runs in super-polynomial time and the protocol is zero-knowledge against unbounded verifiers. The only modification we make to the Zap protocol is that the OT protocol underlying the extractable commitment is instantiated with a smaller security parameter λ^γ relative to the security parameter λ of the CI hash function, where $\gamma < 1$. This is in order to allow a subexponential-time simulator to extract the receiver’s bit. In order to still have receiver’s security of the OT protocol, we need to make the stronger assumption of subexponential hardness of LWE.

We begin with the definition of statistical zero knowledge arguments with superpolynomial simulation.

Definition 11 (T_{Sim} -SPS Statistical Zero Knowledge Arguments). *We call an interactive protocol between a PPT prover P with input $(x, w) \in R_L$ for some language L , and PPT verifier V with input x , denoted by $\langle P, V \rangle(x, w)$, a T_{Sim} -super-polynomial simulation (SPS) statistical zero-knowledge argument for L if it satisfies the following properties:*

- **Completeness.** *For every $(x, w) \in R_L$, $\Pr[\langle P^*, V \rangle(x, w) = 1] \geq 1 - \text{negl}(\lambda)$, where the probability is over the random coins of P and V .*
- **Soundness.** *For every PPT non-uniform prover P^* , that given 1^λ chooses an input length 1^p , and then chooses $x \in \{0, 1\}^p \setminus L$, $\Pr[\langle P^*, V \rangle(x, w) = 1] \leq \text{negl}(\lambda)$, where the probability is over the random coins of V .*

- **T_{Sim} -Statistical Zero Knowledge.** There exists a (uniform) simulator Sim that runs in time T_{Sim} , such that for every $x \in L$, every unbounded verifier V^* , the two distributions $\text{View}_{V^*}[\langle P, V^* \rangle(x, w)]$ and $S^{V^*}(x, z)$ are statistically close.

For this section we use the following definition for statistical sender's security of the OT protocol.

Definition 12 (Statistical Sender's Security of OT). *An OT protocol $(\text{OT}_1, \text{OT}_2)$ satisfies statistical sender's security if there exists a (possibly exponential time) extractor OTExt which on any input ot_1 outputs 0 or 1, such that for all ot_1 and $M^0, M^1, \widetilde{M}^0, \widetilde{M}^1$ with $M^{\text{OTExt}(\text{ot}_1)} = \widetilde{M}^{\text{OTExt}(\text{ot}_1)}$, we have that $\text{OT}_2(\text{ot}_1, M^0, M^1)$ is statistically indistinguishable from $\text{OT}_2(\text{ot}_1, \widetilde{M}^0, \widetilde{M}^1)$.*

In this section we will prove the following theorem.

Theorem 6. *For any $c > 0$, there exists a two message public-coin T_{Sim} -SPS statistical zero knowledge argument system for NP in the plain model, where $T_{\text{Sim}} = 2^{\lambda^c}$, assuming that the following primitives exist:*

- Two-message oblivious transfer (OT) that is subexponentially secure against malicious senders, satisfying Definition 12 and Property 1, and
- Quasi-polynomially correlation intractable hash functions, satisfying Definition 3.

Instantiating the CI hash function and the OT protocol with subexponential LWE to achieve the desired level of security, we obtain the following:

Corollary 7. *Assuming subexponential LWE, for any $c_{\text{Sim}} > 0$, there exists a T_{Sim} -SPS statistical zero knowledge argument for NP in the plain model, where $T_{\text{Sim}} = 2^{\lambda^{c_{\text{Sim}}}}$.*

The OT Extractor. When we construct the simulator, we will need to bound the running time of OTExt . Specifically, the LWE-based OT protocol [BD18] has the receiver send a matrix A in the first round, and in the second round the sender sends two encryptions, using A as a public key in two different ways. [BD18] prove that at most one of these public keys is non-lossy; we need that there exists an exponential-time extractor that determines which one. We prove that OTExt takes at most time 2^{λ^C} , with C a constant, when we instantiate the OT scheme using [BD18]. First we describe the extractor from [BD18]. The ot_1 message is a matrix $A \in \mathbb{Z}_q^{n \times m}$ which defines a q -ary lattice $\Lambda_q(A)$. The behavior of $\text{OTExt}(\text{ot}_1 = A)$ is as follows, where γ , σ and q are all $\text{poly}(\lambda)$:

1. Calculate $r = \text{rank}(\Lambda_q(A) \cap \gamma\mathcal{B})$. If $r > n/2$ then return 0.
2. If $r \leq n/2$ then approximate $x = \rho_\sigma(\Lambda_q(A))$ to within $\pm 2^{n/2-1}$.
 - (a) If $x > 2^{n/2} + 2^{n/2-1}$ then return 0.
 - (b) Otherwise, return 1.

Theorem 8 (Theorem 5.5 in [BD18]). *The OT scheme satisfies Definition 12 using the OTExt defined above.*

In order to analyze the running time of OTExt , we import the following two theorems.

Theorem 9 (Theorem 2.2 in [BD18]). *For any lattice $\Lambda \in R^m$, parameter $\sigma > 0$ and $u \geq 1/\sqrt{2\pi}$, it holds that*

$$\rho_\sigma(\Lambda \setminus 2\sigma\sqrt{m}\mathcal{B}) \leq 2^{-c_u m} \rho_\sigma(\Lambda),$$

where $c_u = -\log(\sqrt{2\pi}eu \cdot e^{-\pi u^2})$.

Theorem 10 (from [BLR⁺18]). *There exists an algorithm that, for any lattice $\Lambda_q(A)$ of dimension m and bound $x = \text{poly}(m)$, enumerates all vectors of length at most x in $\Lambda_q(A)$ in time 2^{m^C} , where C is a constant.*

Claim 7. *There exists a constant C such that the running time of OTExt is 2^{λ^C} .*

Proof. We prove the claim by showing that there are two constants C_1 and C_2 such that calculating $\text{rank}(\lambda_q(a) \cap \gamma \lfloor \cdot \rfloor)$ takes time $2^{\lambda^{C_1}}$ and approximating $x = \rho_\sigma(\Lambda_q(A))$ to within $\pm 2^{n/2-1}$ takes time $2^{\lambda^{C_2}}$. The claim then follows taking any $C' > \max\{C_1, C_2\}$.

We note that $\text{rank}(\Lambda_q(A) \cap \gamma \mathcal{B})$ can be computed exactly by enumerating all $\Lambda_q(A) \cap \gamma \mathcal{B}$ and then finding the maximum number of linearly independent vectors in this set. By Theorem 10, this can be done in $2^{\lambda^{C_1}}$ for some constant C_1 .

To approximate $\rho_\sigma(\Lambda_q(A))$, we note that Theorem 9 and the fact that $\int_{-\infty}^{\infty} \rho_\sigma(x) d^m x = \sigma^{m/2}$ imply that, setting $u = \sigma$,

$$\begin{aligned} \rho_\sigma(\Lambda \cap 2\sigma\sqrt{m}\mathcal{B}) &\geq \rho_\sigma(\Lambda)(1 - 1/2^{c_u m}) \\ &> \rho_\sigma(\Lambda) - \frac{2^{\log \sigma \cdot m/2}}{2^{c_u m}} \\ &> \rho_\sigma(\Lambda) - \frac{1}{2^{m/2}}, \end{aligned}$$

with the last inequality following because $c_u > u$ asymptotically. This is well within our approximation requirements, so it suffices to calculate $\rho_\sigma(\Lambda \cap 2\sigma\sqrt{m}\mathcal{B})$ by enumerating all elements in $\Lambda \cap 2\sigma\sqrt{m}\mathcal{B}$. By Theorem 10, it is possible to do this in $2^{\lambda^{C_2}}$ for some constant C_2 . \square

Construction. Consider the protocol in Section 5.2, with one small modification: the extractable commitment scheme is instantiated with security parameter λ^γ , for a $\gamma > 0$ which we specify later. In the following we prove this protocol still satisfies soundness and that it also satisfies SPS statistical zero knowledge.

Soundness. We show that instantiating the extractable commitment with security parameter λ^γ preserves the security properties of the extractable commitment and thus preserves soundness of the modified protocol.

We first state a fact which is commonly used in complexity-leveraging proofs, that assuming a primitive is subexponentially secure, instantiating the primitive with a much smaller security parameter still satisfies polynomial security.

Fact 1. *Let λ be a security parameter. Assume subexponential receiver's security of the OT scheme. For any $\gamma > 0$, if we set $\lambda' = \lambda^\gamma$ and instantiate the OT scheme with security parameter λ' then the OT scheme satisfies (polynomial) receiver's OT security with respect to the original security parameter λ .*

Claim 8. *Let λ be a security parameter. Assuming subexponential receiver's security of the OT scheme, for any $\gamma > 0$, instantiating the extractable commitment scheme from Section 4 with security parameter $\lambda' = \lambda^\gamma$ satisfies computational binding, frequency of extraction, correctness of extraction and indistinguishability of extractable transcripts with respect to the original security parameter λ .*

Proof. Recall that in Section 4, Lemmas 2 and 3 prove the extractable commitment scheme satisfies computational binding, frequency of extraction, correctness of extraction and indistinguishability of extractable transcripts, assuming computational receiver's security of the underlying OT scheme. Thus by Fact 1 these lemmas apply to an extractable commitment scheme instantiated with $\lambda' = \lambda^\gamma$. \square

Theorem 11. *Assuming subexponential receiver's security of the OT scheme and quasi-polynomially correlation intractable hash functions, the protocol described above satisfies computational soundness against malicious PPT provers.*

Proof. By Claim 8, it follows that the same argument for soundness in Section 5.2 holds for the protocol described above. \square

Zero Knowledge.

In order to construct the simulator, we first show that we can use the exponential-time OT extractor OTExt to equivocate the extractable commitment scheme defined in Section 4.

We define two exponential-time algorithms Equiv_1 and Equiv_2 . Equiv_1 takes as input an extcom_1 message, and outputs an extcom_2 message along with a state σ . Equiv_2 takes the extcom_2 message and state generated by Equiv_1 , along with an arbitrary message \tilde{M} , and outputs a valid opening for the extcom_2 message generated by Equiv_1 .

More formally, $\text{Equiv}_1(\text{extcom}_1)$ performs the following steps:

1. Parse $\text{extcom}_1 = \{\text{ot}_{1,i}\}_{i \in [m]}$.
2. Run $\text{OTExt}(\text{ot}_{1,i})$ to get ch_i for each $i \in [m]$.
3. Sample a random string $r \xleftarrow{\$} \{0,1\}^m$ such that $r \neq \text{ch}$.
4. For each $i \in [m]$, compute $\text{ot}_{2,i} \leftarrow \text{OT}_2(\text{ot}_{1,i}, M_i^0, M_i^1; r_{2,i})$ where all M_i^b and $r_{2,i}$ are uniform random.
5. Output $(\text{extcom}_2 = (r, \{\text{ot}_{2,i}\}_{i \in [m]}), \sigma = (\{M_i^b\}_{i \in [m], b \in \{0,1\}}, \{\text{ch}_i\}_{i \in [m]}))$.

$\text{Equiv}_2(\text{extcom}_2, \sigma, \tilde{M})$ performs the following steps:

1. Parse extcom_2 as $(r, \{\text{ot}_{2,i}\}_{i \in [m]})$ and σ as $(\{M_i^b\}_{i \in [m], b \in \{0,1\}}, \{\text{ch}_i\}_{i \in [m]})$.
2. Generate random strings \tilde{M}_i^b for $i \in [m]$ and $b \in \{0,1\}$ such that $\tilde{M}_i^{\text{ch}_i} = M_i^{\text{ch}_i}$ and $\bigoplus_i \tilde{M}_i^{r_i} = \tilde{M}$. (Note that this is possible since there is at least one i such that $r_i \neq \text{ch}_i$.)
3. For each $i \in [m]$:
 - (a) Find a random $r'_{2,i}$ such that $\text{ot}_{2,i} = \text{OT}_2(\text{ot}_{1,i}, \tilde{M}_i^0, \tilde{M}_i^1; r'_{2,i})$.
 - (b) If no such $r'_{2,i}$ exists, set $r'_{2,i} = \perp$.
4. Output $(\tilde{M}, \{\tilde{M}_i^b\}_{i \in [m], b \in \{0,1\}}, \{r'_{2,i}\}_{i \in [m]})$.

Claim 9. Assume statistical sender's security of the OT protocol. Then for all \widetilde{M} and all (potentially malicious) receivers A ,

$$(\text{extcom}_1, \text{extcom}_2, \text{extcom}_o) \approx_s (\text{extcom}_1, \text{extcom}'_2, \text{extcom}'_o),$$

where $\text{extcom}_1 \leftarrow A$, extcom_2 is generated by an honest committer committing to \widetilde{M} , extcom_o is the corresponding honest opening, $(\text{extcom}'_2, \sigma) \leftarrow \text{Equiv}_1(\text{extcom}_1)$, and $\text{extcom}'_o \leftarrow \text{Equiv}_2(\text{extcom}'_2, \sigma, \widetilde{M})$.

We prove this claim via a sequence of hybrids:

- **Hybrid₀**: This hybrid corresponds to the case where extcom_1 , extcom_2 , and extcom_o are generated by an honest committer committing to and opening \widetilde{M} .
- **Hybrid₁**: This hybrid is identical to the previous hybrid, except that the committer uses OTExt on extcom_1 to obtain ch and sets r in extcom_2 to be a random string such that $r \neq \text{ch}$.
- **Hybrid₂**: This hybrid is identical to the previous hybrid, except that the committer uses Equiv_2 to generate extcom_o .
- **Hybrid₃**: This hybrid is identical to the previous hybrid, except that for the messages M_i^b sent in the ot_2 message in extcom_2 , the committer sets $M_i^{1-\text{ch}_i}$ to be a uniform random string.
- **Hybrid₄**: This hybrid is identical to the previous hybrid, except that the committer sets all messages M_i^b sent in the ot_2 message in extcom_2 to be uniform random strings. This corresponds to the case where the committer uses Equiv_1 and Equiv_2 to interact with the receiver.

Claim 10. Hybrid₀ is statistically indistinguishable from Hybrid₁.

Proof. Assume that there is an \widetilde{M} and an unbounded adversary A that distinguishes between Hybrid₀ and Hybrid₁ with non-negligible probability. Fixing the randomness of A , it is straightforward to construct a reduction which distinguishes with non-negligible probability between the uniform distribution over $\{0, 1\}^m$ and the uniform distribution over $\{0, 1\}^m \setminus \{\text{ch}\}$, of which the statistical distance is exponentially small in λ , a contradiction. \square

Claim 11. Assuming statistical sender's security of the OT protocol, Hybrid₁ is statistically indistinguishable from Hybrid₂.

In order to prove this claim, we first define a game \mathcal{G}_b which we will use to contradict the existence of a distinguishing adversary. In round one of the game, the adversary sends m ot_1 messages to the challenger, along with messages $\{M_i^b\}_{i \in [m], b \in \{0, 1\}}$. In round 2, the challenger first runs OTExt on each ot_1 to obtain the choice bits $\text{ch} \in \{0, 1\}^m$. Then, if $b = 0$, the challenger responds with $\{\text{OT}_2(\text{ot}_{1,i}, M_i^0, M_i^1)\}_i$. If $b = 1$, the challenger sets $M_i^{1-\text{ch}_i} \leftarrow 0^n$ before again responding with $\{\text{OT}_2(\text{ot}_{1,i}, M_i^0, M_i^1)\}_i$.

Claim 12. Assuming statistical sender's security of the OT protocol, \mathcal{G}_0 is statistically indistinguishable from \mathcal{G}_1 .

Proof. We prove this claim via a sequence of $m + 1$ subhybrids, where at hybrid Subhybrid_j , for each $i \leq j$, the challenger responds with $M_i^{1-\text{ch}_i} \leftarrow 0^n$, and for all $i > j$ the challenger responds with the original M_i^0 and M_i^1 . Subhybrid_0 is identical to the case where $b = 0$ and Subhybrid_m is identical to the case where $b = 1$, so it suffices to prove indistinguishability for each successive pair of subhybrids.

Assume that there is an unbounded adversary A which can distinguish between Subhybrid_{j-1} and Subhybrid_j . We construct a reduction R to statistical sender's security of the OT protocol. Fix the randomness of A . This fixes $\text{ot}_{1,j}$, M_j^0 and M_j^1 . R sends this $\text{ot}_{1,j}$ to the challenger C , who extracts the choice bit ch_j and responds either with $\text{OT}_2(\text{ot}_{1,j}, M_j^0, M_j^1)$ or $\text{OT}_2(\text{ot}_{1,j}, \widetilde{M}_j^0, \widetilde{M}_j^1)$, where $\widetilde{M}_j^{\text{ch}_j} = M_j^{\text{ch}_j}$ and $\widetilde{M}_j^{1-\text{ch}_j} = 0^n$. R then constructs its response to A identically to Subhybrid_{j-1} except that for $\text{ot}_{2,j}$ it uses the ot_2 message from C . R then outputs the output of A .

If C responds with $\text{OT}_2(\text{ot}_{1,j}, M_j^0, M_j^1)$ then the view of A is identical to Subhybrid_{j-1} . Moreover, if C responds with $\text{OT}_2(\text{ot}_{1,j}, \widetilde{M}_j^0, \widetilde{M}_j^1)$ then the view of A is identical to Subhybrid_j . Thus since A distinguishes these two hybrids R contradicts statistical sender's security of the OT protocol. \square

Additionally, we use the following lemma:

Lemma 12. *For each λ , let $D_\lambda : \mathcal{X}_\lambda \times \mathcal{Y}_\lambda \rightarrow \{0, 1\}$ be any function and let $\mathcal{D}_{\lambda,1}$ be a distribution on \mathcal{X}_λ and $\mathcal{D}_{\lambda,2}$ and $\mathcal{D}'_{\lambda,2}$ distributions on \mathcal{Y}_λ . Assume that there is a polynomial p and λ_0 such that for all $\lambda > \lambda_0$,*

$$\left| \Pr_{x \leftarrow \mathcal{D}_{\lambda,1}, y \leftarrow \mathcal{D}_{\lambda,2}} [D_\lambda(x, y) = 1] - \Pr_{x \leftarrow \mathcal{D}_{\lambda,1}, y' \leftarrow \mathcal{D}'_{\lambda,2}} [D_\lambda(x, y') = 1] \right| \geq 1/p(\lambda)$$

Then there exist polynomials q and r such that for all $\lambda > \lambda_0$,

$$\Pr_{x \leftarrow \mathcal{D}_{\lambda,1}} \left[\left| \Pr_{y \leftarrow \mathcal{D}_{\lambda,2}} [D_\lambda(x, y) = 1] - \Pr_{y' \leftarrow \mathcal{D}'_{\lambda,2}} [D_\lambda(x, y') = 1] \right| \geq 1/q(\lambda) \right] \geq 1/r(\lambda).$$

Proof. Assume for the sake of contradiction that there exist p , \mathcal{X}_λ , \mathcal{Y}_λ , $\mathcal{D}_{\lambda,1}$, $\mathcal{D}_{\lambda,2}$ and $\mathcal{D}'_{\lambda,2}$ such that

$$\left| \Pr_{x \leftarrow \mathcal{D}_{\lambda,1}, y \leftarrow \mathcal{D}_{\lambda,2}} [D_\lambda(x, y) = 1] - \Pr_{x \leftarrow \mathcal{D}_{\lambda,1}, y' \leftarrow \mathcal{D}'_{\lambda,2}} [D_\lambda(x, y') = 1] \right| \geq 1/p(\lambda), \quad (4)$$

but for all polynomials q and r there is some $\lambda' > \lambda_0$ such that

$$\Pr_{x \leftarrow \mathcal{D}_{\lambda',1}} \left[\left| \Pr_{y \leftarrow \mathcal{D}_{\lambda',2}} [D_{\lambda'}(x, y) = 1] - \Pr_{y' \leftarrow \mathcal{D}'_{\lambda',2}} [D_{\lambda'}(x, y') = 1] \right| \geq 1/q(\lambda') \right] < 1/r(\lambda').$$

Let q and r be such that for all λ , $1/q(\lambda) + 1/r(\lambda) < 1/p(\lambda)$. Let \mathcal{X}_0 be the set of all x such that

$$\left| \Pr_{y \leftarrow \mathcal{D}_{\lambda',2}} [D_{\lambda'}(x, y) = 1] - \Pr_{y' \leftarrow \mathcal{D}'_{\lambda',2}} [D_{\lambda'}(x, y') = 1] \right| \geq 1/q(\lambda').$$

The probability in Equation (4) can be written as

$$\begin{aligned} & \left| \sum_{\tilde{x}} \left(\Pr_{y \leftarrow \mathcal{D}_{\lambda,2}} [D_\lambda(x, y) = 1] - \Pr_{y' \leftarrow \mathcal{D}'_{\lambda,2}} [D_\lambda(x, y') = 1] \right) \Pr_{x \leftarrow \mathcal{D}_{\lambda,1}} [x = \tilde{x}] \right| \\ & \leq \sum_{\tilde{x}} \left| \Pr_{y \leftarrow \mathcal{D}_{\lambda,2}} [D_\lambda(x, y) = 1] - \Pr_{y' \leftarrow \mathcal{D}'_{\lambda,2}} [D_\lambda(x, y') = 1] \right| \Pr_{x \leftarrow \mathcal{D}_{\lambda,1}} [x = \tilde{x}] \\ & = \sum_{\tilde{x} \in \mathcal{X}_0} \left| \Pr_{y \leftarrow \mathcal{D}_{\lambda,2}} [D_\lambda(x, y) = 1] - \Pr_{y' \leftarrow \mathcal{D}'_{\lambda,2}} [D_\lambda(x, y') = 1] \right| \Pr_{x \leftarrow \mathcal{D}_{\lambda,1}} [x = \tilde{x}] \\ & + \sum_{\tilde{x} \notin \mathcal{X}_0} \left| \Pr_{y \leftarrow \mathcal{D}_{\lambda,2}} [D_\lambda(x, y) = 1] - \Pr_{y' \leftarrow \mathcal{D}'_{\lambda,2}} [D_\lambda(x, y') = 1] \right| \Pr_{x \leftarrow \mathcal{D}_{\lambda,1}} [x = \tilde{x}]. \end{aligned}$$

Setting $\lambda = \lambda'$, we have that this probability is less than $1/r(\lambda') + 1/q(\lambda') < 1/p(\lambda)$, a contradiction. \square

Proof of Claim 11. Assume that there is an \widetilde{M} and an unbounded adversary A that distinguishes between Hybrid_1 and Hybrid_2 with non-negligible probability $p(\lambda)$. Assume $A = (A_1, A_2)$, where A_1 generates ot_1 and some private state σ , and A_2 takes $(r, \text{ot}_2, \widetilde{M}, \{M_i^b\}_{i,b}, \{r_{2,i}\}_i, \sigma)$ and outputs one bit.

We build a reduction R to \mathcal{G}_b .

Let C be a challenger for \mathcal{G}_b . First, R receives (ot_1, σ) from A_1 , generates r and $\{M_i^b\}_{i,b}$ as in Hybrid_1 , and sends $(\text{ot}_1, \{M_i^b\}_{i,b})$ to C . C then responds with ot_2 .

R then computes

$$\Pr \left[\text{Exp}_c(\text{ot}_1, r, \{M_i^b\}_{i,b}, \text{ot}_2, \sigma) = 1 \right],$$

for both $c = 0$ and $c = 1$, where Exp is the experiment defined below. Note that R can compute this probability exactly because R is unbounded.

$\text{Exp}_c(\text{ot}_1, r, \{M_i^b\}_{i,b}, \text{ot}_2, \sigma)$ generates $\{\widetilde{M}_i^b\}_{i,b}$ as in Hybrid_2 and proceeds as follows:

- If $c = 0$, for each $i \in [m]$, choose a random $r'_{2,i}$ such that $\text{ot}_{2,i} = \text{OT}_2(\text{ot}_1, i, M_i^0, M_i^1, r'_{2,i})$. If no such $r'_{2,i}$ exists, R set $r'_{2,i} \leftarrow \perp$. Finally, output the result of $A_2(r, \text{ot}_2, \widetilde{M}, \{M_i^b\}_{i,b}, \{r'_{2,i}\}_i, \sigma)$.
- If $c = 1$, for each $i \in [m]$, R chooses a random $r'_{2,i}$ such that $\text{ot}_{2,i} = \text{OT}_2(\text{ot}_1, i, \widetilde{M}_i^0, \widetilde{M}_i^1, r'_{2,i})$. If no such $r'_{2,i}$ exists, R sets $r'_{2,i} \leftarrow \perp$. Finally, output the result of $A_2(r, \text{ot}_2, \widetilde{M}, \{\widetilde{M}_i^b\}_{i,b}, \{r'_{2,i}\}_i, \sigma)$.

Once R has computed these two probabilities, if $|\Pr[\text{Exp}_0 = 1] - \Pr[\text{Exp}_1 = 1]| \geq \gamma$, for a fixed γ which we specify below, R outputs 0. Otherwise, R outputs 1.

Assume C is acting according to \mathcal{G}_0 . Define $D(x, y)$, where $x = (\text{ot}_1, r, \{M_i^b\}_{i,b}, \text{ot}_2, \sigma)$ and $y = (c, r)$, to be the output of $\text{Exp}_c(\text{ot}_1, r, \{M_i^b\}_{i,b}, \text{ot}_2, \sigma; r)$ (where Exp uses randomness r). Let \mathcal{D}_1 be the distribution of x given from the interaction with R , C , and A_1 , let \mathcal{D}_2 be the distribution $(0, r)$ where r is uniform, and similarly let \mathcal{D}'_2 be the distribution $(1, r)$. Then $\Pr_{x \in \mathcal{D}_1, y \in \mathcal{D}_2} [D(x, y) = 1]$ is exactly the probability that A outputs 1 in Hybrid_1 . Similarly, $\Pr_{x \in \mathcal{D}_1, y \in \mathcal{D}'_2} [D(x, y) = 1]$ is exactly the probability that A outputs 1 in Hybrid_2 . Then by the assumption and Lemma 12, There exist polynomials q and r such that

$$\Pr_{x \leftarrow \mathcal{D}_1} \left[\left| \Pr_{y \leftarrow \mathcal{D}_2} [D(x, y) = 1] - \Pr_{y' \leftarrow \mathcal{D}'_2} [D(x, y') = 1] \right| \geq 1/q(\lambda) \right] \geq 1/r(\lambda),$$

so, setting $\gamma = q(\lambda)$, it follows that in the case of \mathcal{G}_0 R outputs 0 with probability at least $r(\lambda)$.

Now assume C is acting according to \mathcal{G}_1 . In this case, for all x , the the view of A in $\text{Exp}_0(x)$ and $\text{Exp}_1(x)$ are identical. Thus R outputs 0 with probability 0.

It follows that R distinguishes \mathcal{G}_0 from \mathcal{G}_1 with probability at least $q(\lambda)$, contradicting Claim 12. \square

Claim 13. *Assuming statistical sender's security of the OT protocol, Hybrid_2 is statistically indistinguishable from Hybrid_3 .*

Proof. We prove the claim via a sequence of subhybrids $\text{Hybrid}_{2,j}$ for $j = 0, \dots, m$, where at each subhybrid $\text{Hybrid}_{2,j}$ we switch $M_j^{1-\text{ch}_j}$ over to be a uniform random string. The beginning of this sequence is exactly Hybrid_2 and the end is exactly Hybrid_3 , so it suffices to show indistinguishability of each successive pair of subhybrids.

Assume for the sake of contradiction that there is an unbounded adversary A which can distinguish between a pair of subhybrids $\text{Hybrid}_{2,j-1}$ and $\text{Hybrid}_{2,j}$ with non-negligible advantage. We construct a reduction R to statistical sender's security of the OT protocol.

Fix the randomness of A and the choice of M_j^0 and M_j^1 . There must be a way to fix this such that A still distinguishes with non-negligible probability. Let R run $\text{Hybrid}_{2,j-1}$ with this fixed randomness until A sends ot_1 . R then forwards $\text{ot}_{1,j}$ to the OT challenger C . When C responds with ot_2 , R proceeds as in $\text{Hybrid}_{2,j-1}$ except that it uses the ot_2 response from C as $\text{ot}_{2,j}$ in its response to A . R then outputs the output of A .

If C sends $\text{ot}_2 = \text{OT}_2(\text{ot}_{1,j}, M_j^0, M_j^1)$, then the view of A is identical to that in $\text{Hybrid}_{2,j-1}$. But if C sends $\text{ot}_2 = \text{OT}_2(\text{ot}_{1,j}, \hat{M}_j^0, \hat{M}_j^1)$, where $\hat{M}_j^{\text{ch}_j} = M_j^{\text{ch}_j}$ and $\hat{M}_j^{1-\text{ch}_j}$ is uniform random, then the view of A is identical to that in $\text{Hybrid}_{2,j}$. It follows that R distinguishes between the two cases with non-negligible probability, contradicting statistical sender's OT security. \square

Claim 14. Hybrid_3 is identical to Hybrid_4 .

Proof. Since $r \neq \text{ch}$, there exists at least one j such that $r_j \neq \text{ch}_j$. Since in Hybrid_3 the receiver receives no information about $M_j^{1-\text{ch}_j}$ whenever $r_j \neq \text{ch}_j$, and the distribution $\{M_i^{r_i}\}_{i \in I}$ in Hybrid_3 is identical to the uniform random distribution assuming $i \not\subseteq [m]$, the claim follows. \square

Claim 15. There is some constant C' such that both Equiv_1 and Equiv_2 run in time $2^{\lambda^{C'}}$.

Proof. Let C_1 be the C given by Claim 7. It follows that Equiv_1 runs in time $\text{poly}(\lambda)2^{\lambda^{C_1}}$. Equiv_2 also runs in time $\text{poly}(\lambda)2^{\lambda^{C_2}}$ for some constant C_2 , so taking $C' > \max\{C_1, C_2\}$, we have that both Equiv_1 and Equiv_2 run in time at most $2^{\lambda^{C'}}$. \square

Recall that the protocol given in Section 5 relies on a variant of the extractable commitment scheme which requires the committer, when committing to multiple messages, to reuse the string r across all extcom_2 messages. Equiv_1 and Equiv_2 extend to this case naturally, since the equivocating committer can take in all ot_1 messages for every instance of the protocol, corresponding to a set $\{\text{ch}_p\}_p$ of challenger strings $\text{ch}_p \in \{0, 1\}^m$, and choose r during Equiv_1 so that $r \neq \text{ch}_p$ for all p . The rest of the equivocation procedure works unmodified.

The Simulator. Before defining the zero-knowledge simulator for the protocol, we make some notational assumptions. Assume the honest-verifier zero-knowledge simulator from Lemma 5 takes input G and a verifier challenge e and generates the corresponding set of decommitted values z that the verifier receives in the third round. That is,

$$z = \left\{ \{z_{r,i,j}\}_{i,j \in [p(\lambda)] \times [p(\lambda)]}, z_{r,P} \right\}_{r \in [\lambda]},$$

where for each $r \in [\lambda]$, if $e_r = 0$ then $\{z_{r,i,j}\}_{i,j}$ encodes a permuted adjacency matrix of G and $z_{r,P}$ encodes the permutation, and if $e_r = 1$ $\{z_{r,i,j}\}_{i,j}$ encodes only the edges of $\pi(H)$, where H is an arbitrary hamiltonian cycle, and where all other $z_{r,i,j}$ and $z_{r,P}$ are set to \perp .

We now define the zero-knowledge simulator $\text{Sim}(G, q, K)$:

1. Parse q as $\left\{ \left\{ \text{extcom}_{1,r,i,j} \right\}_{i,j \in [p(\lambda)] \times [p(\lambda)]}, \text{extcom}_{1,r,P} \right\}_{r \in [\lambda]}$
2. Run Equiv_1 on each extcom_1 in q to get a .
3. Compute $e \leftarrow \mathcal{H}.\text{Eval}(K, (G, a))$.
4. Run the honest-verifier ZK simulator from Lemma 5 with challenge e to get

$$z' = \left\{ \left\{ v_{r,i,j} \right\}_{i,j \in [p(\lambda)] \times [p(\lambda)]}, v_{r,P} \right\}_{r \in [\lambda]}$$

5. For each $v = v_{r,i,j}$ or $v_{r,P}$ that is not \perp , run $\text{Equiv}_2(v, \sigma_v)$, where σ_v is the appropriate state from step 2, to get an equivocated opening to v corresponding to the commitment generated in step 2.
6. Output (a, e, z) , where z is the family of all openings generated in step 5.

Theorem 12. *For all $x \in L$, all witnesses w for x , and all (potentially malicious) verifiers V^* ,*

$$(q, K, a, e, z) \approx_s (q, K, a', e', z'),$$

where a, e, z are generated by the honest prover from the protocol in Section 5.2 using x and w , and a', e', z' are generated by $\text{Sim}(x, q, K)$.

To prove the theorem, we use the following sequence of hybrids:

- **Hybrid₀**: This hybrid corresponds to the experiment where the challenger behaves identically to the verifier in the actual protocol.
- **Hybrid₁**: This hybrid is identical to **Hybrid₀**, except that instead of generating the commitments honestly, the prover uses Equiv_1 to generate a and Equiv_2 to generate the openings z . That is, for every extcom_1 in q , the prover uses Equiv_1 to generate a corresponding extcom_2 (which contains no information), and keeps track of the value that the prover would have committed to. If
- **Hybrid₂**: This hybrid is the same as **Hybrid₁** except that instead of generating openings honestly, the prover runs the honest-verifier zero knowledge simulator on the challenge e to generate the openings. This is identical to the behavior of Sim defined above.

Claim 16. *Assume statistical sender's security of the OT protocol. Then **Hybrid₀** is statistically indistinguishable from **Hybrid₁**.*

Proof. To prove this claim we rely on a sequence of $\lambda(p(\lambda)^2 + 1)$ subhybrids, where at each subhybrid we switch one commitment over to use Equiv_1 and Equiv_2 . The beginning of this sequence is exactly **Hybrid₀** and the end is exactly **Hybrid₁**, so it suffices to show indistinguishability of each successive pair of subhybrids.

Assume for the sake of contradiction that there is an unbounded adversary A which can distinguish between a pair of subhybrids with non-negligible advantage. Without loss of generality, assume the hybrid switches $\text{extcom}_{2,r,i,j}$ and $\text{extcom}_{o,r,i,j}$ from real to equivocated. We construct a reduction R to the indistinguishability game from Claim 9, thus contradicting this claim.

Fix all the randomness of the challenger C_{Hybrid} for the hybrid except the randomness used to generate the commitments; there must be a way to fix this randomness that preserves the advantage of A . Moreover, this fixes the values committed to by the challenger. Denote by C_{Equiv} the challenger for the equivocation game, with message \widetilde{M} the same as the fixed value committed to by C_{Hybrid} in $\text{extcom}_{2,r,i,j}$. R takes the first message q, K sent by A and passes $\text{extcom}_{1,r,i,j}$ to C_{Equiv} . After receiving $\text{extcom}'_{2,r,i,j}$ and $\text{extcom}'_{o,r,i,j}$ from C_{Equiv} , it runs C_{Hybrid} to obtain (a, e, z) and replaces the commitment $\text{extcom}_{2,r,i,j}$ in a with $\text{extcom}'_{2,r,i,j}$ and, if $\text{extcom}_{2,r,i,j}$ is opened in z , it replaces $\text{extcom}_{o,r,i,j}$ in z with $\text{extcom}'_{o,r,i,j}$. Finally, it sends this modified (a, e, z) to A and outputs the output of A . Since view of A in this reduction is identical to the first subhybrid in the case where C_{Equiv} uses real commitments and identical to the second subcase when C_{Equiv} uses equivocated commitments (with appropriate randomness fixed), we have a distinguisher which contradicts Claim 9. \square

Claim 17. *The view of the verifier in Hybrid₁ is identical to that in Hybrid₂.*

Proof. Because in both hybrids the commitments are generated using Equiv and contain no information, the view of the verifier consists entirely of the decommitted values in z . For each r , if $e_r = 0$ then the verifier learns a random permutation of G in both cases, so the views are identical in this case. If $e_r = 1$ then the verifier learns either a random permutation of the witness H or a random permutation of some arbitrary hamiltonian cycle H' , so the views are identical in this case as well. It follows that the views of the verifier are identical in Hybrid₁ and Hybrid₂. \square

Theorem 13. *There is a $\gamma > 0$ such that if the OT protocol underlying the extractable commitment is instantiated with security parameter $\lambda' = \lambda^\gamma$, Sim runs in time at most $2^{\lambda^{c_{\text{Sim}}}}$.*

Proof. The running time of Sim is bounded by $\text{poly}(\lambda)(T_{\text{Equiv}_1} + T_{\text{Equiv}_2} + T_{\mathcal{H}.\text{Eval}} + T_{\text{Sim}_h})$, where $T_{\text{Equiv}_1}, T_{\text{Equiv}_2}, T_{\mathcal{H}.\text{Eval}}$ and T_{Sim_h} are the running times of $\text{Equiv}_1, \text{Equiv}_2, \mathcal{H}.\text{Eval}$ and the honest verifier simulator, respectively.

Choose $\gamma > 0$ such that $\gamma C' < c_{\text{Sim}}$. Then, letting $\epsilon_\gamma = c_{\text{Sim}} - \gamma C'$, the running time of Sim is

$$\begin{aligned} &\leq \text{poly}(\lambda)(T_{\text{Equiv}_1} + T_{\text{Equiv}_2}) \\ &\leq \text{poly}(\lambda)(2^{(\lambda^\gamma)^{C'}}) \\ &= \text{poly}(\lambda)(2^{\lambda^{c_{\text{Sim}} - \epsilon_\gamma}}) \\ &\leq 2^{\lambda^{c_{\text{Sim}}}}. \end{aligned}$$

\square

References

- [BCR86] Gilles Brassard, Claude Crépeau, and Jean-Marc Robert. Information theoretic reductions among disclosure problems. In *FOCS*, 1986.
- [BD18] Zvika Brakerski and Nico Döttling. Two-message statistically sender-private OT from LWE. In *TCC*, 2018.
- [BFM88] Manuel Blum, Paul Feldman, and Silvio Micali. Non-interactive zero-knowledge and its applications (extended abstract). In *STOC*, pages 103–112, 1988.

- [BGI⁺17] Saikrishna Badrinarayanan, Sanjam Garg, Yuval Ishai, Amit Sahai, and Akshay Wadia. Two-message witness indistinguishability and secure computation in the plain model from new assumptions. In *ASIACRYPT*, 2017.
- [BLR⁺18] Shi Bai, Tancreède Lepoint, Adeline Roux-Langlois, Amin Sakzad, Damien Stehlé, and Ron Steinfeld. Improved security proofs in lattice-based cryptography: Using the rényi divergence rather than the statistical distance. *J. Cryptology*, 31(2):610–640, 2018.
- [Blu86] Manuel Blum. How to prove a theorem so no one else can claim it. In *Proceedings of the International Congress of Mathematicians, Berkeley, CA*, pages 1444–1451, 1986.
- [BP15] Nir Bitansky and Omer Paneth. Zaps and non-interactive witness indistinguishability from indistinguishability obfuscation. In *TCC*, 2015.
- [CCH⁺19] Ran Canetti, Yilei Chen, Justin Holmgren, Alex Lombardi, Guy N. Rothblum, Ron D. Rothblum, and Daniel Wichs. Fiat-shamir: from practice to theory. In *STOC*, 2019.
- [CDS94] Ronald Cramer, Ivan Damgård, and Berry Schoenmakers. Proofs of partial knowledge and simplified design of witness hiding protocols. In *CRYPTO*, 1994.
- [Cha86] David Chaum. Demonstrating that a public predicate can be satisfied without revealing any information about how. In *CRYPTO*, 1986.
- [DN00] Cynthia Dwork and Moni Naor. Zaps and their applications. In *FOCS*, 2000.
- [DN02] Cynthia Dwork and Moni Naor. Zaps and their applications. *Electronic Colloquium on Computational Complexity (ECCC)*, (001), 2002.
- [DN07] Cynthia Dwork and Moni Naor. Zaps and their applications. *SIAM J. Comput.*, 36(6):1513–1543, 2007.
- [FS90] Uriel Feige and Adi Shamir. Witness indistinguishable and witness hiding protocols. In *STOC*, 1990.
- [GGPR] Rosario Gennaro, Craig Gentry, Bryan Parno, and Mariana Raykova. Quadratic span programs and succinct nizks without pcps. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*.
- [GJJM] Vipul Goyal, Abhishek Jain, Zhengzhong Jin, and Giulio Malavolta. Statistical zaps and new oblivious transfer protocols. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020*.
- [GMW86] Oded Goldreich, Silvio Micali, and Avi Wigderson. Proofs that yield nothing but their validity and a methodology of cryptographic protocol design (extended abstract). In *FOCS*, 1986.
- [GOS06] Jens Groth, Rafail Ostrovsky, and Amit Sahai. Non-interactive zaps and new techniques for NIZK. In *CRYPTO*, 2006.
- [Gro] Jens Groth. Short pairing-based non-interactive zero-knowledge arguments. In Masayuki Abe, editor, *ASIACRYPT 2010*.
- [HK12] Shai Halevi and Yael Tauman Kalai. Smooth projective hashing and two-message oblivious transfer. *J. Cryptology*, 25(1):158–193, 2012.

- [JKKR17] Abhishek Jain, Yael Tauman Kalai, Dakshita Khurana, and Ron Rothblum. Distinguisher-dependent simulation in two rounds and its applications. In *CRYPTO*, 2017.
- [Kal05] Yael Tauman Kalai. Smooth projective hashing and two-message oblivious transfer. In *EUROCRYPT*, 2005.
- [KKS18] Yael Tauman Kalai, Dakshita Khurana, and Amit Sahai. Statistical witness indistinguishability (and more) in two messages. In *EUROCRYPT*, 2018.
- [KS17] Dakshita Khurana and Amit Sahai. Two-message non-malleable commitments from standard sub-exponential assumptions. *IACR Cryptology ePrint Archive*, 2017:291, 2017.
- [Lip] Helger Lipmaa. Progression-free sets and sublinear pairing-based non-interactive zero-knowledge arguments. In Ronald Cramer, editor, *TCC 2012*.
- [LVW] Alex Lombardi, Vinod Vaikuntanathan, and Daniel Wichs. Statistical ZAPR arguments from bilinear maps. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020*.
- [NP01] Moni Naor and Benny Pinkas. Efficient oblivious transfer protocols. In *SODA*, 2001.
- [PS19] Chris Peikert and Sina Shiehian. Noninteractive zero knowledge for NP from (plain) learning with errors. 2019.
- [SW14] Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: deniable encryption, and more. In *STOC*, 2014.