

Statistical ZAPR Arguments from Bilinear Maps

Alex Lombardi* Vinod Vaikuntanathan† Daniel Wichs‡

February 24, 2020

Abstract

Dwork and Naor (FOCS '00) defined ZAPs as 2-message witness-indistinguishable proofs that are public-coin. We relax this to *ZAPs with private randomness* (ZAPRs), where the verifier can use private coins to sample the first message (independently of the statement being proved), but the proof must remain publicly verifiable given only the protocol transcript. In particular, ZAPRs are *reusable*, meaning that the first message can be reused for multiple proofs without compromising security.

Known constructions of ZAPs from trapdoor permutations or bilinear maps are only computationally WI (and statistically sound). Two recent results of Badrinarayanan-Fernando-Jain-Khurana-Sahai and Goyal-Jain-Jin-Malavolta [EUROCRYPT '20] construct the first *statistical ZAP arguments*, which are statistically WI (and computationally sound), from the quasipolynomial LWE assumption. Here, we construct *statistical ZAPR arguments* from the quasipolynomial decision-linear (DLIN) assumption on groups with a bilinear map. Our construction relies on a combination of several tools, including the Groth-Ostrovsky-Sahai NIZK and NIWI [EUROCRYPT '06, CRYPTO '06, JACM '12], “sometimes-binding statistically hiding commitments” [Kalai-Khurana-Sahai, EUROCRYPT '18] and the “MPC-in-the-head” technique [Ishai-Kushilevitz-Ostrovsky-Sahai, STOC '07].

*MIT. Email: alexjl@mit.edu. Research supported in part by an NDSEG fellowship. Research supported in part by NSF Grants CNS-1350619 and CNS-1414119, and by the Defense Advanced Research Projects Agency (DARPA) and the U.S. Army Research Office under contracts W911NF-15-C-0226 and W911NF-15-C-0236.

†MIT. Email: vinodv@mit.edu. Research was supported in part by NSF Grants CNS-1350619 and CNS-1414119, an NSF-BSF grant CNS-1718161, the Defense Advanced Research Projects Agency (DARPA) and the U.S. Army Research Office under contracts W911NF-15-C-0226 and W911NF-15-C-0236, an IBM-MIT grant and a Microsoft Trustworthy and Robust AI grant.

‡Northeastern and NTT Research Inc. Email: wichs@ccs.neu.edu. Research supported by NSF grants CNS-1314722, CNS-1413964, CNS-1750795 and the Alfred P. Sloan Research Fellowship.

Contents

1	Introduction	1
1.1	Technical Overview	3
1.2	Organization	6
2	Preliminaries	6
2.1	Witness Indistinguishable Arguments	6
3	Tools for the Main Construction	8
3.1	Non-Interactive Statistical Zero Knowledge Arguments	8
3.2	Locally Zero Knowledge Proofs	9
3.3	Sometimes-Binding Statistically Hiding (SBSH) Commitments	10
4	Construction of Statistical ZAPR Arguments	15
4.1	Description	15
4.2	Proof of Theorem 4.2	17
	References	19

1 Introduction

Zero-Knowledge and Witness-Indistinguishability. Zero-knowledge (ZK) proofs, introduced in the ground-breaking paper of Goldwasser, Micali, and Rackoff [GMR85], have found countless uses in cryptography. Unfortunately, such protocols are known to require at least 3 rounds of interaction [GO94] in the plain model without additional setup, which is the model that we consider throughout this work. Witness indistinguishable (WI) proofs [FS90] are a natural relaxation of zero-knowledge, which has turned out to be extremely useful. A WI proof generated using any witnesses w for an NP statement x is indistinguishable from a proof generated with any other possible witness w' for x . Unlike in the case of ZK, there are no lower bounds on the round complexity of WI proofs.

ZAPs and Non-Interactive WI (NIWI). The work of Dwork and Naor [DN00, DN07] constructed two-message public-coin WI proofs, which they called *ZAPs*. By now, we have constructions of ZAPs under any of: trapdoor permutations (factoring) [FLS99, DN00]; the decision-linear assumption (DLIN) in bilinear maps [GOS06a]; indistinguishability obfuscation [BP15]; or learning with errors [BFJ⁺20, GJJM20, LVW19]. In fact, we can even get completely non-interactive WI proofs (NIWI) assuming either trapdoor permutations and a mild complexity-theoretic derandomization assumption [BOV03] or the bilinear DLIN assumption [GOS06a].

ZAPs and ZAPRs The original definition of ZAPs from [DN00, DN07] required that they are public coin, meaning that the first message from the verifier to the prover consists of uniform randomness. The main advantage of such protocols is that they are *publicly verifiable*, meaning that anybody can decide whether the proof is accepting or rejecting by only looking at the protocol transcript. Moreover, in such publicly verifiable protocols, the first message is inherently *reusable* for multiple different proofs of different statements, and security holds even if the cheating prover learns whether the verifier accepts or rejects various proofs with the same first message (since this decision only depends on the public transcript). This is in contrast to *secret-coin* two-message WI proofs, which may be insecure under such reuse.

In this work, we introduce an intermediate notion that we call *ZAPs with private randomness* (ZAPRs). ZAPRs allow the verifier to use secret coins to generate the first message, but we still require the proofs to be *publicly verifiable*, and we require that the first message is sampled independently of the statement being proved. Therefore, ZAPRs have essentially the same advantages as ZAPs, and the two can be used interchangeably in most applications.¹

Statistical WI. Most prior constructions of ZAPs (and 2-message WI protocols in general) only achieve computational WI security, often with statistical soundness [DN00, GOS06a, BP15]. However, it is arguably more important for WI security to hold statistically than it is for soundness. In particular, we want privacy to be preserved long into the future after the protocols have finished executing, despite the potential that computational assumptions may become broken in the long term. On the other hand, soundness is only relevant during the protocol execution itself and, even

¹One notable exception where the “public coin” nature of ZAPs is used essentially is for *derandomization* of the verifier message [BOV03]; however, this seems to require ZAPs satisfying statistical soundness, while we focus on computationally sound, statistically WI protocols in this work.

if the underlying assumptions are broken after the protocol finished executing, it is too late for the adversary to take advantage of this.

Interestingly, 2-message statistically WI protocols were unknown until recently. The first progress on this problem was only made by Kalai, Khurana and Sahai [KKS18], who constructed a *secret-coin* 2-message statistical WI protocol under standard quasi-polynomial assumptions (DDH or QR or N th residuosity). Unfortunately, their protocol is not publicly verifiable and the first message is not reusable (a simple attack breaks soundness under such reuse). Even more recently, Badrinarayanan et al. [BFJ⁺20] along with Goyal et al. [GJJM20]² constructed the first statistical ZAP arguments under the quasi-polynomial LWE assumption. These last two results rely on recent constructions of NIZKs from LWE [CLW18, CCH⁺19, PS19] via *correlation-intractable hash functions*, which in turn rely on fully homomorphic encryption/commitments from LWE. This left open the question of whether we can achieve such statistical ZAP or ZAPR arguments under other assumptions, without relying on LWE or “fully homomorphic cryptography”.

Our Results. In this work, we construct statistical ZAPR arguments from the quasi-polynomial decision-linear (DLIN) assumption in groups with a bilinear map. More generally, we construct ZAPR arguments using three generic ingredients:

- *Non-interactive statistical ZK (NISZK) arguments in the common-reference string (CRS) model.* We need the scheme to have the additional property that every *valid* CRS in the support of the setup algorithm ensures that the resulting arguments are statistically WI. This is guaranteed, for example, if the NISZK argument system satisfies perfect zero knowledge, as in [GOS06b, GOS12]. One can think of this property as ensuring WI security even if the CRS is chosen “semi-maliciously” using adversarial randomness *but still from the support of the setup algorithm*.
- *Non-interactive WI proofs (NIWI) in the plain model,* where the WI property is computational and soundness is statistical. As mentioned above, we know how to construct such NIWI proofs assuming either trapdoor permutations and a mild complexity-theoretic derandomization assumption [BOV03] or the bilinear DLIN assumption [GOS06a].
- *Sometimes binding, statistically hiding (SBSH) commitments.* This is a relaxation of a notion introduced recently by [KKS18].³ It is a 2-round commitment protocol where the receiver chooses a random α in the first round, and the sender sends a random β and uses $\text{ck} = (\alpha, \beta)$ as a commitment key to create a commitment $\text{Com}(\text{ck}, m)$ to his message m in the second round. Even if the receiver chooses α maliciously, the commitment key ck is statistically hiding with overwhelming probability over a random choice of β . However, there is some inverse quasi-polynomial probability ϵ such that, even if the sender chooses β maliciously after seeing α , the commitment key $\text{ck} = (\alpha, \beta)$ makes the commitment statistically binding. Furthermore, the sender cannot tell whether this rare event occurs or not.

The first two primitives can be constructed under the bilinear DLIN assumption using the techniques of [GOS06a]. (We will require that the primitives satisfy quasi-polynomial security and

²The conference paper [GJJM20] subsumes the construction of statistical ZAP arguments in a preprint of Jain and Jin [JJ19].

³The main difference is that their commitment needed to be “sometimes extractable” whereas ours only needs to be “sometimes statistically binding”.

therefore need to rely on quasi-polynomial DLIN.) The last primitive can be constructed under a variety of quasi-polynomial assumptions such as DDH or QR or N 'th residuosity [KKS18], and we show it can also be done under quasi-polynomial DLIN.

Our construction broadens the set of assumptions from which we can build statistical ZAPR arguments (previously only quasi-polynomial LWE was known) and gives an alternate approach for achieving them without relying on correlation intractability.

What about adaptive soundness? We show that our statistical ZAPR arguments, under the quasi-polynomial bilinear DLIN assumption, satisfy *non-adaptive soundness*: for any false statement x , a (quasi-poly time) cheating prover P^* cannot find proof π^* for x that the verifier would accept. One could potentially ask for the stronger security notion of *adaptive soundness*: informally, a protocol is adaptively sound if a cheating prover P^* cannot find *any false statement* $x^* \notin L$ along with an accepting proof π^* for x^* .

As is standard for adaptive security notions, if we strengthen our assumption to the *subexponential* security of bilinear DLIN, we can make use of complexity leveraging [BB04] and obtain a statistical ZAPR argument that is adaptively sound for statements of a priori bounded length. More formally, for every length $\ell(\lambda)$, there is a statistical ZAPR argument $\Pi^{(\ell)}$ that is adaptively sound for statements of length $\ell(\lambda)$.

One would ideally hope for a protocol satisfying adaptive soundness for unbounded (poly-length) statements. However, there is some evidence that such a protocol would be difficult to obtain. In particular, in the context of *NISZK arguments*, a result of Pass [Pas16] shows that there is no black-box reduction from the adaptive soundness of a NISZK protocol to a “falsifiable assumption” [Nao03]. There is additionally no known non-black-box construction overcoming this impossibility result (without relying on non-falsifiable assumptions, as in [AF07]).

Given the similarity between NISZK arguments and statistical ZAPR arguments (if anything, the latter seem harder to achieve), we consider this to be a barrier to constructing adaptively sound statistical ZAPR arguments. However, no formal impossibility result is known; indeed, we do not even know how to rule out the existence of *statistical ZAP proofs* (ZAPs satisfying both statistical soundness and statistical WI) for all of NP.

1.1 Technical Overview

We now describe our construction using the above primitives. We start with a very simple construction, which already gives a 2-message (publicly verifiable) statistical WI protocol for $\text{NP} \cap \text{coNP}$ and conveys some of the intuition.

Interestingly, our warm-up protocol relies on only the *polynomial hardness* of bilinear DLIN (rather than quasi-polynomial hardness), yielding a 2-message statistical WI protocol for a broad class of languages without relying on super-polynomial assumptions.

We then describe our more complex construction, which works for all of NP.

1.1.1 Warm-Up: A simple protocol for $\text{NP} \cap \text{coNP}$

As a warm up, we describe a very simple 2-message statistical WI argument for languages $L \in \text{NP} \cap \text{coNP}$. In this warm-up construction, the first message depends on the statement x being proved, but we remove this in the full construction. The construction makes use of NISZK arguments and

NIWI as above (but does not require SBSH commitments). The main ideas behind the construction are that:

1. The prover uses the [GOS12] NISZK argument system to prove that $x \in L$, where we let the verifier choose the CRS. This already provides “semi-malicious” WI security. To get full WI, we need to ensure that the CRS is valid (in the support of the setup algorithm).
2. The verifier uses a NIWI to prove that the CRS is valid. The challenge is to only rely on WI security rather than full ZK. To do so, we let the verifier prove that either the CRS is valid or $x \notin L$.

In more detail, the protocol proceeds as follows.

Verifier \rightarrow **Prover**: The verifier samples a CRS of a NISZK argument. He then uses a NIWI to prove that either the CRS is valid (i.e., in the support of the setup algorithm, using the random coins of the setup algorithm as a witness) or $x \notin L$. The first message consists of the CRS along with the NIWI proof.

Prover \rightarrow **Verifier**: The prover verifies the NIWI proof (aborting if it does not accept) and then uses the NISZK argument with the received CRS to prove that $x \in L$.

For $x \in L$, the statistical WI security of the ZAPR follows from the statistical soundness of the NIWI, which ensures that the CRS is valid, together with the statistical WI of the NISZK, which holds for all valid CRS.

For $x \notin L$, the computational soundness of the ZAPR follows by first relying on the computational WI security of the NIWI to argue that the prover cannot notice if we modify the NIWI proof to use the witness for $x \notin L$ instead of the randomness of the setup algorithm. With this change, we can then rely on the computational soundness of the NISZK argument to argue that the prover cannot produce a valid NISZK proof for $x \in L$.

1.1.2 The Full Construction

The full construction is more involved. In addition to the three primitives mentioned previously (NISZK, NIWI, and SBSH commitments), we also rely on an additional information-theoretic tool that we now describe.

Locally-ZK Proofs (LZK) via “MPC in the Head”. We introduce a new tool called *locally ZK proofs* (LZK). An LZK proof consists of a probabilistic encoding that maps a witness w for a statement x into a proof string $\pi \in \Sigma^\ell$ for some alphabet Σ . There is also a polynomial size set $\{S_1, \dots, S_Q\}$ of “queries” $S_i \subseteq [\ell]$ and a verification algorithm $\text{Verify}(x, i, \pi[S_i])$ that locally verifies that π is consistent on the positions S_i . The proof satisfies two statistical security properties:

- **Global Soundness**: If there exists some proof $\pi \in \Sigma^\ell$ such that $\text{Verify}(x, i, \pi[S_i]) = 1$ for all $i \in [Q]$ then $x \in L$.
- **t -Local-ZK**: For any t queries S_{a_1}, \dots, S_{a_t} the values $\pi[S_{a_1}], \dots, \pi[S_{a_t}]$ can be simulated without knowing the witness.

We can think of LZK proofs as a relaxation of ZK-PCPs [KPT97] where the verifier needs to make *all* the queries to be convinced of soundness but ZK holds locally. We construct such LZK proofs for any Q and $t < Q/2$ using the “MPC in the head” technique [IKOS07]. In particular, to construct the proof π , the encoding algorithm runs a (semi-honest information-theoretic) MPC protocol with Q parties and security against t corruptions. Each party has as input a secret share (in an additive secret sharing) of the witness w and the MPC outputs 1 to each party iff the shares add up to a valid witness for x . The proof π is of length $\ell = Q + Q(Q - 1)/2$ and contains the view of each party $i \in [Q]$ in the protocol, as well as the contents of the $Q(Q - 1)/2$ communication channels between each pair of parties $\{i, j\}$. Each query set S_i contains locations that correspond to the view of party i and all of the communication channels that involve party i . The verification algorithm for i checks that the view of the party i and the communication channels involving party i correspond to an honest execution of the protocol and that the output of the protocol is 1. It is easy to check that this satisfies global soundness and t -local ZK.

ZAPR Construction. We now describe our ZAPR construction using NIWIs, NISZKs, sometimes binding statistically hiding commitments, and LZK proofs. To rely on quasi-polynomial assumptions, we choose the parameter Q of the LZK proof to be $\text{poly}(\log \lambda)$.

Verifier \rightarrow Prover: The verifier samples $3Q$ CRS’s of the NISZK. We interpret this as Q bundles of 3 CRS’s each. The verifier then gives a NIWI proof that, in each bundle, at least 2 out of 3 of the CRS’s are valid. He does so by choosing a random 2 of the 3 CRS’s in each bundle and using the corresponding randomness of the setup algorithm for them as the witness. Lastly, the verifier also sends the first message α of the SBSH commitment scheme.

Prover \rightarrow Verifier: The prover verifies the NIWI proofs and aborts if any of them do not accept. The prover then samples an LZK proof $\pi \in \Sigma^\ell$ for the statement $x \in L$. It samples the SBSH commitment component β and uses the commitment key $\text{ck} = (\alpha, \beta)$ to commit to each of the ℓ blocks of π separately. Lastly, it chooses a random CRS in each bundle $i \in [Q]$ and uses it to give an NISZK argument showing that the LZK verifier outputs $\text{Verify}(x, i, \pi[S_i]) = 1$, where $\pi[S_i]$ is contained in the committed values. It sends back β , all the commitments, and the NISZK arguments.

We first argue that the above construction is statistically WI. By the statistical hiding of the commitment scheme, the commitments do not reveal anything about the committed values. By the statistical soundness of the NIWI, we know that at least 2 of the 3 CRS’s in each bundle are valid. Since the prover chooses a random CRS in each bundle, on expectation at least $2Q/3$ of the chosen CRS’s are valid and, by Chernoff, at least $Q/2$ of them are valid with overwhelmingly probability. The NISZK arguments for the valid CRS’s are statistically WI and hence do not reveal any information about the committed values. The remaining $t < Q/2$ NISZK arguments may reveal some information about the committed values $\pi[S_i]$. But, by the locally-ZK property of the proof π , this does not reveal anything about w .

Next, we argue that the construction is computationally sound. Assume that the adversarial prover succeeds in proving a false statement with non-negligible probability δ . The commitment scheme ensures that there is a ϵ probability that $\text{ck} = (\alpha, \beta)$ is binding and, because the prover cannot tell whether this occurred or not, the probability that (1) *the commitment is binding* and (2) *the prover succeeds in proving a false statement* is $\epsilon \cdot \delta$, which is inverse quasi-polynomial. Next,

we rely on the (quasi-polynomial) computational WI security of the NIWI argument to argue that the prover cannot learn which 2 of the 3 CRS's in each bundle had their setup randomness used as a witness in the NIWI. Therefore, even if we condition on (1) and (2), there is an inverse quasi-polynomial $(1/3)^Q$ chance that (3) *in each bundle, the prover chooses the one CRS whose setup randomness was not used in the NIWI*. Altogether there is an inverse quasi-poly probability of (1), (2) and (3) occurring simultaneously. But if this happens, then (as guaranteed by the global soundness of the LZK proof) at least one of the statements proved via the NISZK is false and therefore the prover breaks the (quasi-polynomial) soundness of the NISZK arguments.

In our presentation, we assume quasi-polynomial hardness of the underlying primitives, but only ensure that the statistical WI holds with a quasi-polynomial error. We could analogously assume sub-exponential hardness and ensure that statistical WI holds with a sub-exponentially small error.

1.2 Organization

The rest of the paper is organized as follows. In Section 2, we describe basic preliminaries on witness indistinguishability and ZAPRs. In Section 3, we introduce and discuss some of the main tools used in our construction: NISZK arguments, locally zero knowledge proofs, and sometimes-binding statistically hiding commitments. Finally, in Section 4, we present our construction of statistical ZAPR arguments from these building blocks.

2 Preliminaries

We say that a function $\mu(\lambda)$ is *negligible* if $\mu(\lambda) = O(\lambda^{-c})$ for every constant c , and that two distribution ensembles $X = \{X_\lambda\}$ and $Y = \{Y_\lambda\}$ are computationally indistinguishable ($X \approx_c Y$) if for all polynomial-sized circuit ensembles $\{\mathcal{A}_\lambda\}$,

$$\left| \Pr[\mathcal{A}_\lambda(X_\lambda) = 1] - \Pr[\mathcal{A}_\lambda(Y_\lambda) = 1] \right| = \text{negl}(\lambda).$$

More generally, for any function $\delta(\lambda)$, we say that X and Y are δ -computationally indistinguishable ($X \approx_{c,\delta} Y$) if for all polynomial-sized circuit ensembles $\{\mathcal{A}_\lambda\}$,

$$\left| \Pr[\mathcal{A}_\lambda(X_\lambda) = 1] - \Pr[\mathcal{A}_\lambda(Y_\lambda) = 1] \right| = O(\delta(\lambda)).$$

2.1 Witness Indistinguishable Arguments

Definition 2.1. A witness indistinguishable argument system Π for an NP relation R consists of ppt interactive algorithms (P, V) with the following syntax.

- $P(x, w)$ is an interactive algorithm that takes as input an instance x and witness w that $(x, w) \in R$.
- $V(x)$ is an interactive algorithm that takes as input an instance x . At the end of an interaction, it outputs a bit b . If $b = 1$, we say that V **accepts**, and otherwise we say that V **rejects**.

The proof system Π must satisfy the following requirements for every polynomial function $n = n(\lambda)$. Recall that $\mathcal{L}(R)$ denotes the language $\{x : \exists w \text{ s.t. } (x, w) \in R\}$ and R_n denotes the set $R \cap (\{0, 1\}^n \times \{0, 1\}^*)$.

- **Completeness.** For every $(x, w) \in R$, it holds with probability 1 that V accepts at the end of an interaction $\langle P(x, w), V(x) \rangle$.
- **Soundness.** For every $\{x_{n(\lambda)} \in \{0, 1\}^{n(\lambda)} \setminus \mathcal{L}(R)\}_\lambda$ and every polynomial size $P^* = \{P_\lambda^*\}$, there is a negligible function ν such that V accepts with probability $\nu(\lambda)$ at the end of an interaction $\langle P^*(x), V(x) \rangle$.
- **Witness Indistinguishability.** For every ppt (malicious) verifier V^* and every ensemble $\{(x_n, (w_{0,n}, w_{1,n}), z_n) : (x_n, w_{0,n}), (x_n, w_{1,n}) \in R_n\}_\lambda$, the distribution ensembles

$$\text{view}_{V^*} \langle P(x, w_0), V^*(x, w_0, w_1, z) \rangle$$

and

$$\text{view}_{V^*} \langle P(x, w_1), V^*(x, w_0, w_1, z) \rangle$$

are computationally indistinguishable.

In the work, we focus on obtaining two message WI arguments for NP. A (two message) WI argument system can also satisfy various stronger properties. We describe the variants relevant to this work below.

- **Public Verification:** A WI argument system is publicly verifiable if the verifier's accept/reject algorithm is an efficiently computable function of the transcript (independent of the verifier's internal state).
- **Delayed Input:** A *two-message* WI argument system is *delayed input* if the (honestly sampled) verifier message $\alpha \leftarrow V(1^\lambda, x) = V(1^\lambda, 1^n)$ depends only on the length $n = |x|$.
- **Statistical Soundness.** For every $\{x_n \in \{0, 1\}^n \setminus \mathcal{L}(R)\}$ and every (*unbounded*) $P^* = \{P_\lambda^*\}$, there is a negligible function ν such that V accepts with probability $\nu(\lambda)$ at the end of an interaction $\langle P^*(x), V(x) \rangle$.
- **Statistical Witness Indistinguishability.** For every polynomial function $n(\lambda)$, every (*unbounded*) (malicious) verifier V^* , and every ensemble $\{(x_n, (w_{0,n}, w_{1,n}), z_n) : (x_n, w_{0,n}), (x_n, w_{1,n}) \in R_n\}_\lambda$, the distribution ensembles

$$\text{view}_{V^*} \langle P(x, w_0), V^*(x, w_0, w_1, z) \rangle$$

and

$$\text{view}_{V^*} \langle P(x, w_1), V^*(x, w_0, w_1, z) \rangle$$

are *statistically* indistinguishable.

Our goal is to construct a 2-message argument system that is publicly verifiable, delayed input, and satisfies statistical witness indistinguishability. We call such protocols *statistical ZAPR arguments*.

Definition 2.2 (Statistical ZAPR Arguments). *A 2-message argument system (P, V) is a statistical ZAPR argument system if it is a delayed-input, publicly verifiable protocol satisfying statistical witness indistinguishability.*

As a tool towards our construction, we make use of another variant of WI arguments: non-interactive witness indistinguishable proofs (NIWIs).

Definition 2.3 (NIWI Proofs). *A one-message proof system is a non-interactive witness indistinguishable proof system if it satisfies statistical soundness and (computational) witness indistinguishability.*

By [GOS06a], we know that NIWIs exist based on the decision linear assumption on groups with bilinear maps.

Lemma 2.4 ([GOS06a]). *Under the DLIN assumption, there exists a NIWI proof system for NP.*

3 Tools for the Main Construction

3.1 Non-Interactive Statistical Zero Knowledge Arguments

We make use of non-interactive statistical zero knowledge arguments in the *common reference string model*, as constructed by [GOS06b] under the DLIN assumption on bilinear groups. Moreover, we make use of the fact that the GOS protocol satisfies *statistical witness indistinguishability* in the presence of semi-malicious setup, which we describe below.

Definition 3.1. *A non-interactive statistical zero knowledge (NISZK) argument system Π for an NP relation R consists of three ppt algorithms (Setup, P , V) with the following syntax.*

- **Setup**($1^n, 1^\lambda$) takes as input a statement length n and a security parameter λ . It outputs a common reference string crs .
- $P(\text{crs}, x, w)$ takes as input the common reference string, as well as x and w such that $(x, w) \in R$. It outputs a proof π .
- $V(\text{crs}, x, \pi)$ takes as input the common reference string, a statement x , and a proof π . It outputs a bit b . If $b = 1$, we say that V **accepts**, and otherwise we say that V **rejects**.

The proof system Π must satisfy the following requirements for every polynomial function $n = n(\lambda)$.

- **Completeness.** For every $(x, w) \in R$, it holds with probability 1 that $V(\text{crs}, x, \pi) = 1$ in the probability space defined by sampling $\text{crs} \leftarrow \text{Setup}(1^{|x|}, 1^\lambda)$ and $\pi \leftarrow P(\text{crs}, x, w)$.
- **(Non-adaptive) Soundness.** For every $\{x_n \in \{0, 1\}^n \setminus \mathcal{L}(R)\}$ and every polynomial size $P^* = \{P_\lambda^*\}$, there is a negligible function ν such that

$$\Pr_{\substack{\text{crs} \leftarrow \text{Setup}(1^n, 1^\lambda) \\ \pi \leftarrow P_\lambda^*(\text{crs})}} [V(\text{crs}, x_n, \pi) = 1] \leq \nu(\lambda).$$

- **Statistical Zero Knowledge.** There is a ppt simulator Sim such that for every ensemble $\{(x_n, w_n) \in R_n\}$, the distribution ensembles

$$\left\{ (\text{crs}_{\lambda, n}, P(\text{crs}_{\lambda, n}, x_n, w_n)) \right\}_\lambda$$

and

$$\left\{ \text{Sim}(x_n, 1^\lambda) \right\}_\lambda$$

are statistically indistinguishable in the probability space defined by sampling $\text{crs}_{\lambda, n} \leftarrow \text{Setup}(1^n, 1^\lambda)$ (and evaluating P and Sim with independent and uniform randomness).

In this work, we consider a strengthening of statistical zero knowledge⁴ to a setting where the CRS is chosen in a semi-malicious way.

Definition 3.2 (Semi-Malicious Statistical Witness Indistinguishability). *We say that a NISZK argument system (Setup, P, V) is statistically witness indistinguishable in the presence of semi-malicious setup if for every polynomial function $n(\lambda)$ and every ensemble $\{(\text{crs}_{\lambda,n}, x_n, (w_{0,n}, w_{1,n}), z_n) : \text{crs}_{\lambda,n} \in \text{Supp}(\text{Setup}(1^\lambda, 1^n)) \text{ and } (x_n, w_{0,n}), (x_n, w_{1,n}) \in R_n\}_\lambda$, the distribution ensembles*

$$\left\{ (\text{crs}_{\lambda,n}, P(\text{crs}_n, x_n, w_{0,n})), z_n \right\}_\lambda$$

and

$$\left\{ (\text{crs}_{\lambda,n}, P(\text{crs}_n, x_n, w_{1,n})), z_n \right\}_\lambda$$

are statistically indistinguishable.

In other words, witness indistinguishability is guaranteed for *any* CRS that can be output by the $\text{Setup}(1^\lambda, 1^n)$ algorithm. Moreover, we have the following:

Remark 3.1. Any NISZK argument system satisfying perfect zero knowledge (or perfect WI) satisfies semi-malicious statistical (and even perfect) WI.

Therefore, we obtain the following conclusion from [GOS12]:

Lemma 3.3. *Under the DLIN assumption on groups with a bilinear map, there exists an NISZK argument system for NP satisfying semi-malicious statistical WI.*

3.2 Locally Zero Knowledge Proofs

In this section, we define “locally zero knowledge proofs”, which one can think of as a weak kind of zero-knowledge PCP [KPT97] that captures the “MPC in the head” paradigm [IKOS07].

Definition 3.4 (t -Local Zero Knowledge Proof). *For an NP language L (with witness relation R), a t -local zero-knowledge proof $\text{lzkp} = (\text{Prove}, \text{Verify})$ is a pair of PPT algorithms with the following syntax.*

- $\text{Prove}(x, w)$ takes as input a statement $x \in L$ and witness $w \in R_x$; it outputs a proof $\pi = (\pi_1, \dots, \pi_\ell) \in \Sigma^\ell$ for some alphabet Σ .
- $\text{Queries} = \{S_1, \dots, S_Q\} \subset \{0, 1\}^{[\ell]}$ is a set of “allowable queries”; we require that it is possible to enumerate Queries in time $\text{poly}(n, Q)$.
- $\text{Verify}(x, i, \pi_{S_i})$ takes as input a statement x , index i (describing some set $S_i \in \text{Queries}$), and string $\pi_{S_i} \in \Sigma^{|S_i|}$; it outputs a bit $b \in \{0, 1\}$.

We say that lzkp has $Q = |\text{Queries}|$ possible queries and block length Σ . Moreover, we require that the following properties hold.

⁴Technically, it is only a strengthening of witness indistinguishability.

- **Completeness:** for any valid pair (x, w) and any index $i \in [Q]$, we have that $\text{Verify}(x, i, \pi_{S_i}) = 1$ with probability 1 over the randomness of $\pi \leftarrow \text{Prove}(x, w)$.
- **Soundness:** for any $x \notin L$ and any proof π , there exists some index $i \in Q$ such that $\text{Verify}(x, i, \pi_{S_i}) = 0$.
- **Perfect Zero Knowledge for t Queries:** there exists a PPT simulator $\text{Sim}(x, i_1, \dots, i_t) \rightarrow \tilde{\pi}_{S^*}$ such that for every valid pair (x, w) and every collection of t indices $i_1, \dots, i_t \in [Q]$, the distribution on $\tilde{\pi}_{S^*}$ is identical to the marginal distribution of an honestly generated proof π on the subset $S^* = S_{i_1} \cup \dots \cup S_{i_t}$.

Lemma 3.5. For any $t > 0$, there exists a t -local zero knowledge proof for Circuit-SAT with $Q = 2t + 1$ possible queries.

Proof. (sketch) Let Π denote an MPC protocol for distributed Circuit-SAT (that is, the functionality $(w_1, \dots, w_T) \mapsto C(\bigoplus w_i)$ for an arbitrary input circuit C) for $T = 2t + 1$ parties satisfying information theoretic security against a collection of t semi-honest parties. Following [IKOS07], we define the following proof system:

- **Prove** (x, w) : interpret $x = C$ as a circuit; set $(w_i)_{i=1}^T$ to be a T -out-of- T secret sharing of w , and let $\pi = \left((\text{view}_i)_{i=1}^T, (\tau_{ij})_{i \neq j} \right)$ denote the following information regarding an honest execution of Π (evaluating $C(\bigoplus w_i)$): view_i denotes the view of party i in this execution, and τ_{ij} denotes the communication transcript between party i and party j .
- **Queries:** for every $i \in [T]$, we define the set $S_i \subset [T + \binom{T}{2}]$ to be $\{\text{view}_i\} \cup \{\tau_{i,j}\}_{j=1}^T$.
- **Verify** (x, i, π_{S_i}) outputs 1 if and only if (for $S_i = \{\text{view}_i\} \cup \{\tau_{i,j}\}_{j=1}^T$):
 - view_i is internally consistent and outputs 1.
 - For every j , view_i is consistent with $\tau_{i,j}$.

It was implicitly shown in [IKOS07] that this protocol satisfies the desired properties. Completeness holds assuming that Π is perfectly complete; soundness holds because if $x \notin L$, then there is no valid witness for x , and hence any consistent collection of views and transcripts $\left((\text{view}_i)_{i=1}^T, (\tau_{ij})_{i \neq j} \right)$ for Π must correspond to a global execution of Π outputting 0. Perfect zero knowledge for t joint queries holds by the perfect security of Π against t semi-honest parties. \square

3.3 Sometimes-Binding Statistically Hiding (SBSH) Commitments

For simplicity, we focus on two-message commitment schemes with the following form:

- **Key Agreement:** The sender and receiver execute a two-message protocol in which they publicly agree on a commitment key ck (the transcript of the protocol). We require that the sender message be *public-coin*⁵ (i.e., it simply outputs a string β). In other words,
 - The receiver $R(\rho) \rightarrow \alpha$ outputs a message α using randomness ρ .

⁵Equivalently, we require that the commitment scheme is hiding even given a “partial opening”, i.e., the randomness used in this phase.

- The (honest) sender S samples and sends a uniformly random string $\beta \leftarrow \{0, 1\}^\ell$.
- The commitment key is defined to be $\text{ck} = (\alpha, \beta)$.

- **Non-Interactive Commitment:** the sender commits to a message m using a (non-interactive) PPT algorithm $\text{Com}(\text{ck}, m)$.

We call these schemes “non-interactive commitment schemes with key agreement.” We will denote a transcript of this commitment scheme $(\alpha, \beta, \text{com})$.

We say that a commitment key ck is **binding** if the non-interactive commitment scheme Com with hardwired key ck is perfectly binding.

Definition 3.6 (Sometimes-Binding Statistically Hiding (SBSH) Commitments). *A non-interactive commitment scheme with key agreement (R, S, Com) is a sometimes-binding statistically hiding (SBSH) commitment scheme with parameters (ϵ, δ) if the following three properties hold.*

- **Statistical hiding:** for any malicious PPT receiver R^* (using randomness ρ and outputting message α), the view of R^* in an interaction with an honest sender statistically hides the sender’s message m ; that is,

$$\{(\rho, \alpha, \beta, \text{Com}(\text{ck}, 0))\} \approx_s \{(\rho, \alpha, \beta, \text{Com}(\text{ck}, 1))\}$$

for $\alpha = R^*(\rho)$, $\beta \leftarrow \{0, 1\}^\ell$, and $\text{ck} = (\alpha, \beta)$.

- **Sometimes statistical binding:** for any malicious PPT sender $S^*(\alpha) \rightarrow (\beta^*, \text{st})$ for the key agreement phase, and for any PPT distinguisher $D(\text{st}) \rightarrow b \in \{0, 1\}$, we have that

$$\Pr[D(\text{st}) = 1 \wedge \text{ck} := (\alpha, \beta^*) \text{ is binding}] = \epsilon \cdot \Pr[D(\text{st}) = 1] \pm \delta \cdot \text{negl}(\lambda),$$

where the probability is taken over $\alpha \leftarrow R(1^\lambda)$, $(\beta^*, \text{st}) \leftarrow S^*(\alpha)$, and the randomness of D .

In other words, it is a statistically hiding commitment scheme such that, even for malicious PPT senders S^* , the commitment key ck is binding with probability roughly ϵ , and moreover any event that S^* produces (with sufficiently high probability) occurs “independently” of the event that ck is binding.

Constructions. The works [KKS18, BFJ⁺20, GJJM20] construct variants of SBSH commitment schemes (for ϵ and δ both inverse quasi-polynomial in the security parameter) from (quasi-polynomially secure) 2-message OT satisfying IND-based security against PPT senders and statistical sender privacy against unbounded receivers.⁶ This leads to instantiations based on DDH [NP05], QR/DCR [HK12] and LWE [BD18]. In fact, the [NP05] oblivious transfer scheme can be generalized to a variant that relies on the DLIN assumption (rather than DDH) on (not necessarily bilinear) cryptographic groups, which then yields SBSH commitments based on DLIN as well.

⁶All three of these works use slightly different security definitions than we use here, but the [BFJ⁺20, GJJM20] instantiations can easily be shown to satisfy our variant of the security property.

3.3.1 Extending Naor-Pinkas OT to DLIN

Definition 3.7 (DLIN [BBS04]). Let \mathbb{G} a group of prime order q with generator g (all parametrized by the security parameter λ), where the tuple (\mathbb{G}, g, q) is public. The DLIN assumption states that

$$(g^a, g^b, g^c, g^{ar_1}, g^{ar_2}, g^{c(r_1+r_2)}) : a, b, c, r_1, r_2 \leftarrow \mathbb{Z}_q$$

is computationally indistinguishable from a uniformly random distribution over \mathbb{G}^6 .

It will be convenient for us to work with “matrix in the exponent” notation, where for a matrix $M \in \mathbb{Z}_q^{n \times m}$ we let g^M denote the matrix of group elements $(g^{M_{i,j}})$. We define the set \mathcal{D} of matrices

$$\mathcal{D} = \left\{ \begin{bmatrix} a & 0 & c \\ 0 & b & c \end{bmatrix} : a, b, c \in \mathbb{Z}_q^* \right\}$$

Then the DLIN assumption can be equivalently written as

$$\left((g^{\mathbf{D}}, g^{\mathbf{rD}}) : \mathbf{D} \leftarrow \mathcal{D}, \mathbf{r} \leftarrow \mathbb{Z}_q^2 \right) \approx_c \left((g^{\mathbf{D}}, g^{\mathbf{u}}) : \mathbf{D} \leftarrow \mathcal{D}, \mathbf{u} \leftarrow \mathbb{Z}_q^3 \right)$$

We also define $g^{\mathcal{D}}$ to be the set $\{g^{\mathbf{D}} : \mathbf{D} \in \mathcal{D}\}$. Membership in $g^{\mathcal{D}}$ can be checked efficiently.

OT Construction and Security. We define a 2-round oblivious transfer scheme $(\text{OT}_1, \text{OT}_2, \text{Rec})$ where the receiver computes $(\text{ot}_1, \text{st}) \leftarrow \text{OT}_1(b)$ with the choice bit $b \in \{0, 1\}$, the sender computes $\text{ot}_2 \leftarrow \text{OT}_2(\text{ot}_1, m_0, m_1)$ and receiver recovers $m_b = \text{Rec}(\text{ot}_2, \text{st})$. We define the functions as follows:

- $\text{ot}_1 \leftarrow \text{OT}_1(b)$: Sample $\mathbf{D} \leftarrow \mathcal{D}$, $\mathbf{r} \leftarrow \mathbb{Z}_q^2$ and define $\mathbf{v}_b = \mathbf{rD}$, $\mathbf{v}_{1-b} = (0, 0, 1) - \mathbf{v}_b$. Output $\text{ot}_1 = (g^{\mathbf{D}}, g^{\mathbf{v}_0}, g^{\mathbf{v}_1}), \text{st} = (b, \mathbf{r})$.
- $\text{OT}_2(\text{ot}_1, m_0, m_1)$: Parse $\text{ot}_1 = (g^{\mathbf{D}}, g^{\mathbf{v}_0}, g^{\mathbf{v}_1})$ and $m_0, m_1 \in \mathbb{G}$. Check that $g^{\mathbf{D}} \in g^{\mathcal{D}}$ and that $g^{\mathbf{v}_0 + \mathbf{v}_1} = g^{(0,0,1)}$; if not then abort. Sample $\mathbf{a}_0 \leftarrow \mathbb{Z}_q^3, \mathbf{a}_1 \leftarrow \mathbb{Z}_q^3$ and output $\text{ot}_2 = (g^{\mathbf{D}\mathbf{a}_0^T}, g^{\mathbf{D}\mathbf{a}_1^T}, g^{\mathbf{v}_0 \cdot \mathbf{a}_0^T} \cdot m_0, g^{\mathbf{v}_1 \cdot \mathbf{a}_1^T} \cdot m_1)$.
- $\text{Rec}(\text{ot}_2, \text{st})$: Parse $\text{ot}_2 = (g^{\mathbf{z}_0}, g^{\mathbf{z}_1}, h_0, h_1)$ and $\text{st} = (b, \mathbf{r})$. Output $h_b \cdot g^{-\mathbf{r} \cdot \mathbf{z}_b^T}$.

We now show that this scheme satisfies the same properties as Naor-Pinkas OT.

- *Correctness*: For any b, m_0, m_1 it holds that if $(\text{ot}_1, \text{st}) \leftarrow \text{OT}_1(b), \text{ot}_2 \leftarrow \text{OT}_2(\text{ot}_1, m_0, m_1), m = \text{Rec}(\text{ot}_2, \text{st})$ then $m = m_b$ with probability 1.

Proof. This is because, using the notation of the scheme, we have $g^{\mathbf{v}_b} = g^{\mathbf{rD}}, g^{\mathbf{z}_b} = g^{\mathbf{D}\mathbf{a}_b^T}$ and hence

$$h_b = g^{\mathbf{v}_b \cdot \mathbf{a}_b^T} \cdot m_b = g^{\mathbf{rD} \cdot \mathbf{a}_b^T} \cdot m_b = g^{\mathbf{r} \cdot \mathbf{z}_b^T} \cdot m_b.$$

So $h_b \cdot g^{-\mathbf{r} \cdot \mathbf{z}_b^T} = m_b$. □

- *Computational Receiver Security*: We have

$$(\text{ot}_1 : (\text{ot}_1, \text{st}) \leftarrow \text{OT}_1(0)) \approx (\text{ot}_1 : (\text{ot}_1, \text{st}) \leftarrow \text{OT}_1(1)).$$

Proof. This follows from DLIN. In particular, we can modify the OT_1 algorithm to sample $\mathbf{v}_b \leftarrow \mathbb{Z}_q^3$ instead of $\mathbf{v}_b \leftarrow \mathbf{rD}$ and the distribution of ot_1 is indistinguishable. But in this case the bit b is statistically hidden since in either case the vectors $\mathbf{v}_0, \mathbf{v}_1$ are just uniformly random subject to $\mathbf{v}_0 + \mathbf{v}_1 = (0, 0, 1)$. \square

- *Statistical Sender Security:* There exists an inefficient function Extract such that, for any ot_1 , if $b = \text{Extract}(\text{ot}_1)$ then $\text{OT}_2(\text{ot}_1, m_0, m_1)$ statistically hides m_{1-b} : for any m_0, m_1, m'_0, m'_1 such that $m_b = m'_b$ we have $\text{OT}_2(\text{ot}_1, m_0, m_1)$ is statistically close to $\text{OT}_2(\text{ot}_1, m'_0, m'_1)$.

Proof. We define $\text{Extract}(\text{ot}_1 = (g^{\mathbf{D}}, g^{\mathbf{v}_0}, g^{\mathbf{v}_1}))$ to output 0 if \mathbf{v}_0 is in the row-space of \mathbf{D} and 1 otherwise. If it does not hold that $g^{\mathbf{D}} \in g^{\mathcal{D}}$ and that $g^{\mathbf{v}_0 + \mathbf{v}_1} = g^{(0,0,1)}$ then $\text{OT}_2(\text{ot}_1, \dots)$ aborts and we are done. Otherwise, at most one of $\mathbf{v}_0, \mathbf{v}_1$ is in the row-space of \mathbf{D} since $(0, 0, 1)$ is not in the row space. Therefore \mathbf{v}_{1-b} is not in the row-space of \mathbf{D} . But this means that $g^{\mathbf{D}\mathbf{a}_{1-b}^T}, g^{\mathbf{v}_{1-b} \cdot \mathbf{a}_{1-b}^T}$ are mutually random and independent over the choice of \mathbf{a}_{1-b} and therefore the message m_{1-b} is perfectly hidden. \square

This completes the construction of statistically sender private (2-message) OT from DLIN. Moreover, quasi-polynomial security of the scheme is inherited from the (quasi-polynomial) DLIN assumption, so we additionally obtain SBSH commitments from quasi-polynomial DLIN.

3.3.2 SBSH Commitments via NIWI

In this section, we present another construction of SBSH commitments from bilinear DLIN using a proof technique similar to that of our main construction in Section 4.

The OT-based commitment schemes above satisfy a stronger security property than “sometimes statistical binding”: informally, they are “sometimes extractable”. We write down a construction that does not involve any extraction using two generic building blocks (both instantiable based on DLIN): NIWI proofs along with a slight strengthening of dual-mode commitments in the CRS model.

Definition 3.8 (Semi-Malicious Secure Dual-Mode Commitment). *A non-interactive commitment scheme $\text{Com}(\text{ck}, m)$ in the CRS model is a semi-malicious secure dual-mode commitment if there are two additional algorithms $(\text{BindingSetup}, \text{HidingSetup})$ satisfying the following properties.*

- $\text{BindingSetup}(1^\lambda) \rightarrow \text{ck}$ and $\text{HidingSetup}(1^\lambda) \rightarrow \text{ck}$ both output a commitment key.
- **Key Indistinguishability:** Commitment keys output by BindingSetup and HidingSetup are computationally indistinguishable.
- **Honest Binding:** $(\text{BindingSetup}, \text{Com})$ is a statistically binding commitment scheme in the CRS model.
- **Semi-Malicious Hiding:** For any commitment key ck in the support of HidingSetup , the commitment distribution $\text{Com}(\text{ck}, m)$ (with ck hardwired) statistically hides the message m .

That is, a semi-malicious secure dual-mode commitment satisfies the property that commitments using semi-maliciously chosen “hiding keys” still statistically hide the underlying message. We say that a key ck “is a hiding key” if ck is in the support of HidingSetup .

Remark 3.2. The [GOS06a] homomorphic commitment scheme based on DLIN is a semi-malicious secure dual-mode commitment scheme. It was explicitly shown to be a dual-mode commitment, but by inspection, we see that it is statistically hiding for an arbitrary (hardwired) key from the “hiding” distribution.

We now show how to construct a sometimes-binding statistically hiding commitment scheme using NIWI proofs and a semi-malicious secure dual-mode commitment; this in particular yields such a scheme based on the DLIN assumption on bilinear groups. Our construction is inspired by the construction of [KKS18, BFJ⁺20, GJJM20].

Construction 3.9. Let $(\text{BindingSetup}, \text{HidingSetup}, \text{Com})$ denote a semi-malicious secure dual-mode commitment scheme, and let $(\text{niwi.Prove}, \text{niwi.Verify})$ denote a NIWI proof system. We then define the following two-message commitment scheme:

- **Receiver message:** for $\ell = \log(\frac{1}{\epsilon})$, the receiver samples a random string $r \leftarrow \{0, 1\}^\ell$ along with ℓ pairs of commitment keys $\{\text{ck}_{i,b}\}_{i \in [\ell], b \in \{0,1\}}$, such that

- ck_{i,r_i} is sampled using $\text{BindingSetup}(1^\lambda)$; and
- $\text{ck}_{i,1-r_i}$ is sampled using $\text{HidingSetup}(1^\lambda)$ with randomness $\text{tk}_{i,1-r_i}$.

The receiver then outputs $\{\text{ck}_{i,b}\}_{i \in [\ell], b \in \{0,1\}}$ along with a NIWI proof that for every $i \in [\ell]$, at least one out of $(\text{ck}_{i,0}, \text{ck}_{i,1})$ is a hiding key (using witness $\text{tk}_{i,1-r_i}$).

- **Sender Key Selection:** the sender first verifies the NIWI above and aborts if the check fails. The sender then samples and outputs a uniformly random string $s \leftarrow \{0, 1\}^\ell$.
- **Non-Interactive Commitment:** to commit to a bit m , the sender samples 2ℓ uniformly random bits $\{\sigma_{i,b}\}$. The sender then outputs $\{\text{com}_{i,b} \leftarrow \text{Com}(\text{ck}_{i,b}, \rho_{i,b})\}$ along with $c := m \oplus \bigoplus_i \sigma_{i,s_i}$.

It now remains to show that this commitment scheme satisfies the desired security properties.

- **Statistical hiding:** without loss of generality, consider a fixed first message $(\{\text{ck}_{i,b}\}, \pi)$ sent by a (potentially malicious) receiver R^* . In order for hiding to be broken, this proof π must be accepted by the sender S , so by the soundness of our NIWI, we know that there exists a string r^* such that $\text{ck}_{i,1-r_i^*}$ is in the support of $\text{HidingSetup}(1^\lambda)$. Now, we note that if the sender S picks any $s \neq r^*$, the commitment $(\{\text{com}_{i,b}\}, c)$ statistically hides the underlying message m ; this is because for any i such that $s_i \neq r_i^*$, we have that com_{i,s_i} statistically hides σ_{i,s_i} and hence $c = m \oplus \bigoplus_i \sigma_{i,s_i}$ statistically hides m . Since S only picks $s = r^*$ with probability $2^{-\ell} = \epsilon$, we conclude that this commitment is statistically hiding.
- **Sometimes statistical binding:** we claim that (ϵ, δ) sometimes statistical binding holds assuming (1) the dual-mode commitment satisfies $\delta \cdot \text{negl}(\lambda)$ -key indistinguishability, and (2) the NIWI is $\delta \cdot \text{negl}(\lambda)$ -witness indistinguishable.

Equivalently, we want to show that the following two distributions are $\delta \cdot \text{negl}(\lambda)$ -computationally indistinguishable for any malicious PPT sender S^* :

$$\{(\alpha, S^*(\alpha), r)\} \approx_{c, \delta \cdot \text{negl}(\lambda)} \{(\alpha, S^*(\alpha), r')\}$$

where $r, r' \leftarrow \{0, 1\}^\ell$ are i.i.d. and α is computed using r . To prove the above indistinguishability, consider the following sequence of hybrids.

- H_0 : This is the LHS, $\{(\alpha, S^*(\alpha), r)\}$.
- H_1 : Same as H_0 , except that the receiver samples ck_{i,r_i} using `HidingSetup` (instead of `BindingSetup`). In other words, in H_1 , all keys $ck_{i,b}$ are sampled from `HidingSetup`. We have that $H_0 \approx_{c,\delta \cdot \text{negl}(\lambda)} H_1$ by the key indistinguishability of the dual-mode commitment.
- H_2 : Same as H_1 , except that the proof π is sampled using a random ℓ -tuple of witnesses (as opposed to witnesses $\{tk_{i,1-r_i}\}$). We have that $H_1 \approx_{c,\delta \cdot \text{negl}(\lambda)} H_2$ by the witness indistinguishability of the NIWI.
- H_3 : Same as H_2 , except that r is replaced by r' in the third slot. We have that $H_2 \equiv H_3$ because r and r' are i.i.d. conditioned on $(\alpha, S^*(\alpha))$ as computed in H_2/H_3 .
- H_4 : Same as H_3 , except that π is sampled using witnesses $\{tk_{i,1-r_i}\}$; indistinguishability is the same as H_1/H_2 .
- H_5 : Same as H_4 , except that the receiver samples ck_{i,r_i} using `BindingSetup` (instead of `HidingSetup`); indistinguishability is the same as H_0/H_1 . This is the RHS.

This completes the proof of indistinguishability.

4 Construction of Statistical ZAPR Arguments

We now give our construction of statistical ZAPR arguments, which are proven sound under the quasi-polynomial DLIN assumption in bilinear groups.

4.1 Description

Our construction uses the following ingredients. Let $\epsilon = \epsilon(\lambda)$ denote a fixed negligible function.

- Let $\text{lzkp} = (\text{lzkp.Prove}, \text{lzkp.Queries}, \text{lzkp.Verify})$ denote a t -local zero knowledge proof with $Q = 2t + 1 = \log_3(\frac{1}{\epsilon})$.
- Let $\text{sbsb} = (\text{sbsb.R}, \text{sbsb.S}, \text{sbsb.Com})$ denote a SBSH commitment scheme with parameters (ϵ, ϵ^2) .
- Let $\text{niwi} = (\text{niwi.Prove}, \text{niwi.Verify})$ denote a NIWI proof system for NP that satisfies $\epsilon(\lambda)^3 \cdot \text{negl}(\lambda)$ -witness indistinguishability.
- Let $\text{nizk} = (\text{nizk.Setup}, \text{nizk.Prove}, \text{nizk.Verify})$ denote a NISZK argument system with $\epsilon(\lambda)^3 \cdot \text{negl}(\lambda)$ (computational) soundness error along with semi-malicious statistical witness indistinguishability.

Construction 4.1. *With $\text{niwi}, \text{nizk}, \text{lzkp}, \text{sbsb}$ as above, we define the following two-message argument system $\text{zapr} = (\text{zapr.V}, \text{zapr.Prove}, \text{zapr.Verify})$ as follows*

- **Verifier message:** $\text{zapr.V}(1^n, 1^\lambda)$ does the following.

- Sample a commitment first message $\alpha \leftarrow \text{sbsbsh.R}(1^\lambda)$.
- Sample $3Q$ common reference strings $\text{crs}_{i,a} \leftarrow \text{nizsk.Setup}(1^n, 1^\lambda; \rho_{i,a})$ (using randomness $\rho_{i,a}$).
- Sample a random string $r \leftarrow \{0, 1, 2\}^Q$.
- Sample a proof

$$\text{niwi.}\pi \leftarrow \text{niwi.Prove}(\varphi, \{\text{crs}_{i,a}\}_{i \in [Q], a \in [3]}, \{\rho_{i,r_i+1}, \rho_{i,r_i+2}\}_{i \in [Q]}),$$

where sums $r_i + 1, r_i + 2$ are computed mod 3, and $\varphi(\{\text{crs}_{i,a}\}_{i \in [Q], a \in [3]})$ denotes the statement “for every $i \in [Q]$, at least two out of $\{\text{crs}_{i,0}, \text{crs}_{i,1}, \text{crs}_{i,2}\}$ are in the support of $\text{nizsk.Setup}(1^n, 1^\lambda)$.”

- Output $(\alpha, \{\text{crs}_{i,a}\}_{i \in [Q], a \in [3]}, \text{niwi.}\pi)$.

- **Prover message:** Given a verifier message $(\alpha, \{\text{crs}_{i,a}\}_{i \in [Q], a \in [3]}, \text{niwi.}\pi)$ and an instance-witness pair $(x, w) \in R_L$, zapr.Prove does the following.

- Verify the proof $\text{niwi.}\pi$ with respect to $\{\text{crs}_{i,a}\}_{i \in [Q], a \in [3]}$ and abort if the check fails.
- Sample a (uniformly random) sbsbsh second message β and set $\text{ck} = (\alpha, \beta)$.
- Sample a locally zero knowledge proof

$$(\text{lzkp.}\pi_1, \dots, \text{lzkp.}\pi_\ell) \leftarrow \text{lzkp.Prove}(x, w).$$

- For $j \in [\ell]$, sample commitments $\text{com}_j \leftarrow \text{sbsbsh.Com}(\text{ck}, \text{lzkp.}\pi_j; \sigma_j)$ to the symbol $\text{lzkp.}\pi_j$.
- Sample a random string $s \leftarrow \{0, 1, 2\}^Q$.
- For every $i \in [Q]$ sample a NISZK proof

$$\text{nizsk.}\pi_i \leftarrow \text{nizsk.Prove}(\text{crs}_{i,s_i}, \psi, i, \text{ck}, \text{com}_{S_i}, \sigma_{S_i})$$

for the statement $\psi(\text{ck}, i, \text{com}_{S_i})$ denoting “ com_{S_i} is a commitment (under ck) to a string π_{S_i} such that $\text{lzkp.Verify}(x, i, \pi_{S_i})$ outputs 1.”

- Output $(\beta, \{\text{com}_j\}_{j \in [\ell]}, s, \{\text{nizsk.}\pi_i\}_{i \in [Q]})$.

- **Proof Verification:** given a statement x and transcript

$$\tau = \left(\alpha, \{\text{crs}_{i,a}\}_{i \in [Q], a \in [3]}, \text{niwi.}\pi, \beta, \{\text{com}_j\}_{j \in [\ell]}, s, \{\text{nizsk.}\pi_i\}_{i \in [Q]} \right),$$

zapr.Verify does the following: for every $i \in [Q]$, verify the proof $\text{nizsk.}\pi_i$ using crs_{i,s_i} ; output 1 if all Q proofs are accepted.

We now proceed to prove the following theorem about Construction 4.1.

Theorem 4.2. *If $\text{lzkp}, \text{sbsbsh}, \text{niwi}$, and nizsk satisfy the hypotheses stated in Section 4.1, then zapr is a ZAPR argument system with $\epsilon^{\Omega(1)}$ (computational) soundness error and $\epsilon^{\Omega(1)}$ -statistical witness indistinguishability.*

This has the following implication for bilinear DLIN-based statistical ZAPR arguments.

Corollary 4.3. *Under the bilinear DLIN assumption (ruling out inverse quasi-polynomial advantage), there exist statistical ZAPR arguments for NP with inverse quasi-polynomial soundness error and satisfying inverse quasi-polynomial statistical indistinguishability.*

Under the (inverse) subexponential bilinear DLIN assumption, there exist statistical ZAPR arguments for NP with inverse subexponential soundness error and satisfying inverse subexponential statistical indistinguishability.

4.2 Proof of Theorem 4.2

Completeness of our protocol follows from the completeness of `niwi`, `nizsk`, `lzkp`, and the correctness of `sbsh`. Moreover, the protocol is delayed input and publicly verifiable by construction. In the rest of this section, we prove that the protocol is computationally sound and statistically witness indistinguishable.

4.2.1 Statistical Witness Indistinguishability

Let (x, w_0, w_1) denote a statement x along with two witnesses w_0, w_1 for $x \in L$. Let V^* denote a malicious (unbounded) verifier, which without loss of generality we may assume to be deterministic and outputs a message $m_1 = (\alpha, \{\text{crs}_{i,a}\}_{i \in [Q], a \in [3]}, \text{niwi}.\pi)$. We want to show that a proof $\text{zapr.Prove}(m_1, x, w_0)$ is statistically indistinguishable from a proof $\text{zapr.Prove}(m_1, x, w_1)$.

To do so, we first note that if $\text{niwi.Verify}(\varphi, \{\text{crs}_{i,a}\}_{i \in [Q], a \in [3]}, \text{niwi}.\pi)$ outputs 0, then the `zapr` prover aborts and hence indistinguishability trivially holds. Hence, we assume that the NIWI verification passes.

In this case, the perfect soundness of `niwi` implies that there exists a string $r^* \in \{0, 1, 2\}^Q$ such that for all $i \in [Q]$, crs_{i,r_i^*+1} and crs_{i,r_i^*+2} are in the support of $\text{nizsk.Setup}(1^n, 1^\lambda)$. Since the prover samples $s \leftarrow \{0, 1, 2\}^Q$ uniformly at random, we know that the agreement between s and r^* is at most $t = \frac{Q-1}{2}$ with probability $\geq 1 - 2^{-\Omega(Q)} = 1 - \epsilon^{\Omega(1)} = 1 - \text{negl}(\lambda)$ by a Chernoff bound. Therefore, we assume that this event holds in the following analysis.

We now consider the following sequence of hybrids; let `USim` denote the unbounded simulator for `nizsk` corresponding to the semi-malicious witness indistinguishability property. For $s \in \{0, 1, 2\}^Q$, let $\text{Good}(s) \subset [Q]$ denote the set of $j \in [Q]$ such that $s_j \neq r_j^*$, and let $\text{Bad}(s)$ denote the remaining set.

- $H_{0,b}$: this is an honest proof $\text{zapr.Prove}(m_1, x, w_b)$.
- $H_{1,b}$: this is the same as $H_{0,b}$, except that for all $j \in \text{Good}(s)$, we sample $\text{nizsk}.\pi_i \leftarrow \text{USim}(\text{crs}_{i,s_i}, \psi, \text{ck}, \text{com}_{S_i})$. We have that $H_{1,b} \approx_s H_{0,b}$ by the semi-malicious witness indistinguishability of `nizsk` (and the fact that crs_{s_i} is in the support of $\text{nizsk.Setup}(1^n, 1^\lambda)$ for all $i \in \text{Good}(s)$).
- $H_{2,b}$: this is the same as $H_{1,b}$, except that for all $j \notin \bigcup_{i \in \text{Bad}(s)} S_i$, we sample $\text{com}_j \leftarrow \text{sbsh.Com}(\text{ck}, 0)$ to be a commitment to an all 0s string. We have that $H_{1,b} \approx_s H_{2,b}$ by the statistical hiding of `sbsh` (which can be invoked because the commitment randomness used to sample com_j is not used anywhere in these hybrids).
- $H_{3,b}$: this is the same as $H_{2,b}$, except that for all $j \in \bigcup_{i \in \text{Bad}(s)} S_i$, we instead sample $\text{lzkp}.\pi_j \leftarrow \text{lzkp.Sim}(x, \text{Bad}(s))$ using the `lzkp` simulator. We have that $H_2 \approx_s H_3$ by the perfect zero knowledge of `lzkp` (which can be invoked because the symbols $\text{lzkp}.\pi_j$ for $j \notin \bigcup_{i \in \text{Bad}(s)} S_i$ do not appear in these hybrids).

Finally, we note that H_3 is defined independently of the bit b ; hence, statistical witness indistinguishability holds.

4.2.2 Computational Soundness

We claim that our argument system has computational soundness error at most ϵ .

To see this, let $x \notin L$ be a false statement, and suppose that an efficient cheating prover $P^*(\alpha, \{\text{crs}_{i,a}\}_{i \in [Q], a \in [3]}, \text{niwi}_\pi)$ successfully breaks the soundness of `zapr` with probability at least ϵ . We then make the following sequence of claims about P^* .

- $P^*(\alpha, \{\text{crs}_{i,a}\}_{i \in [Q], a \in [3]}, \text{niwi}_\pi)$ breaks the soundness of `zapr` and outputs a message β^* such that $\text{ck} = (\alpha, \beta^*)$ is binding with probability $\epsilon^2(1 - \text{negl}(\lambda))$. This follows directly from the $(\epsilon, \epsilon^2 \cdot \text{negl}(\lambda))$ “sometimes statistical binding” property of `sbsh`.
- $P^*(\alpha, \{\text{crs}_{i,a}\}_{i \in [Q], a \in [3]}, \text{niwi}_\pi)$ simultaneously:
 - breaks the soundness of `zapr`,
 - outputs β^* such that ck is a binding key, and
 - outputs $s = r$ (the verifier’s random string)

with probability $\epsilon^3(1 - \text{negl}(\lambda))$. This holds by the $\epsilon^3 \cdot \text{negl}(\lambda)$ -witness indistinguishability of `niwi`, using the following argument. Consider an alternative experiment in which the verifier samples $r, r' \leftarrow \{0, 1, 2\}^Q$ i.i.d. and uses the r' -witness when computing `niwi`. π instead of the r -witness; in this experiment, P^* indeed satisfies the above three conditions with probability $\epsilon^3(1 - \text{negl}(\lambda))$, since here, r is independent of the rest of the experiment (and so $s = r$ with probability ϵ conditioned on the rest of the experiment). Then, the same holds true in the real soundness experiment by the $\epsilon^3 \cdot \text{negl}(\lambda)$ -witness indistinguishability of `niwi`.

This last claim about P^* contradicts the $\epsilon^3 \cdot \text{negl}(\lambda)$ -soundness of `nizsk`. This is because when ck is a binding key, the soundness of `lzkp` implies that for any collection of commitments $(\text{com}_1, \dots, \text{com}_\ell)$, there exists some index i such that the statement $\psi(\text{ck}, i, \text{com}_{S_i})$ is false. By randomly guessing which of the Q statements is false, P^* can therefore be used to contradict the $\epsilon^3 \cdot \text{negl}(\lambda)$ -soundness of `nizsk`.

Acknowledgements

We thank the anonymous reviewers for their helpful comments and suggestions.

References

- [AF07] Masayuki Abe and Serge Fehr, *Perfect nizk with adaptive soundness*, Theory of Cryptography Conference, Springer, 2007, pp. 118–136.
- [BB04] Dan Boneh and Xavier Boyen, *Efficient selective-id secure identity-based encryption without random oracles*, International conference on the theory and applications of cryptographic techniques, Springer, 2004, pp. 223–238.
- [BBS04] Dan Boneh, Xavier Boyen, and Hovav Shacham, *Short group signatures*, Annual International Cryptology Conference, Springer, 2004, pp. 41–55.
- [BD18] Zvika Brakerski and Nico Döttling, *Two-message statistically sender-private ot from lwe*, Theory of Cryptography Conference, Springer, 2018, pp. 370–390.
- [BFJ⁺20] Saikrishna Badrinarayan, Rex Fernando, Aayush Jain, Dakshita Khurana, and Amit Sahai, *Statistical zap arguments*, Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT 2020), Springer, 2020, <https://eprint.iacr.org/2019/780>.
- [BOV03] Boaz Barak, Shien Jin Ong, and Salil Vadhan, *Derandomization in cryptography*, Annual International Cryptology Conference, Springer, 2003, pp. 299–315.
- [BP15] Nir Bitansky and Omer Paneth, *ZAPs and Non-Interactive Witness Indistinguishability from Indistinguishability Obfuscation*, Theory of Cryptography - TCC 2015, 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*, Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, ACM, 2019.
- [CLW18] Ran Canetti, Alex Lombardi, and Daniel Wichs, *Fiat-shamir: From practice to theory, part ii (non-interactive zero knowledge and correlation intractability from circular-secure fhe)*, IACR Cryptology ePrint Archive **2018** (2018).
- [DN00] Cynthia Dwork and Moni Naor, *Zaps and their applications*, Proceedings 41st Annual Symposium on Foundations of Computer Science, IEEE, 2000, pp. 283–293.
- [DN07] ———, *Zaps and their applications*, SIAM Journal on Computing **36** (2007), no. 6, 1513–1543.
- [FLS99] Uriel Feige, Dror Lapidot, and Adi Shamir, *Multiple noninteractive zero knowledge proofs under general assumptions*, SIAM Journal on Computing **29** (1999), no. 1, 1–28.
- [FS90] Uriel Feige and Adi Shamir, *Witness indistinguishable and witness hiding protocols*, Proceedings of the twenty-second annual ACM symposium on Theory of computing, Citeseer, 1990, pp. 416–426.
- [GJJM20] Vipul Goyal, Abhishek Jain, Zhengzhong Jin, and Giulio Malavolta, *Statistical zaps and new oblivious transfer protocols*, Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT 2020), Springer, 2020, Subsumes [JJ19].

- [GMR85] S Goldwasser, S Micali, and C Rackoff, *The knowledge complexity of interactive proof systems*, Proceedings of the seventeenth annual ACM symposium on Theory of computing, ACM, 1985, pp. 291–304.
- [GO94] Oded Goldreich and Yair Oren, *Definitions and properties of zero-knowledge proof systems*, Journal of Cryptology **7** (1994), no. 1, 1–32.
- [GOS06a] Jens Groth, Rafail Ostrovsky, and Amit Sahai, *Non-interactive zaps and new techniques for nizk*, Annual International Cryptology Conference, Springer, 2006, pp. 97–111.
- [GOS06b] ———, *Perfect non-interactive zero knowledge for np*, Annual International Conference on the Theory and Applications of Cryptographic Techniques, Springer, 2006, pp. 339–358.
- [GOS12] Jens Groth, Rafail Ostrovsky, and Amit Sahai, *New techniques for noninteractive zero-knowledge*, J. ACM **59** (2012), no. 3, 11:1–11:35.
- [HK12] Shai Halevi and Yael Tauman Kalai, *Smooth projective hashing and two-message oblivious transfer*, Journal of Cryptology **25** (2012), no. 1, 158–193.
- [IKOS07] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai, *Zero-knowledge from secure multiparty computation*, Proceedings of the thirty-ninth annual ACM symposium on Theory of computing, ACM, 2007, pp. 21–30.
- [JJ19] Abhishek Jain and Zhengzhong Jin, *Statistical zap arguments from quasi-polynomial lwe*, Cryptology ePrint Archive, Report 2019/839, 2019, <https://eprint.iacr.org/2019/839>.
- [KKS18] Yael Tauman Kalai, Dakshita Khurana, and Amit Sahai, *Statistical witness indistinguishability (and more) in two messages*, Annual International Conference on the Theory and Applications of Cryptographic Techniques, Springer, 2018, pp. 34–65.
- [KPT97] Joe Kilian, Erez Petrank, and Gábor Tardos, *Probabilistically checkable proofs with zero knowledge*, STOC, vol. 97, Citeseer, 1997, pp. 496–505.
- [LVW19] Alex Lombardi, Vinod Vaikuntanathan, and Daniel Wichs, *2-message publicly verifiable wi from (subexponential) lwe*, Cryptology ePrint Archive, Report 2019/808, 2019, <https://eprint.iacr.org/2019/808>.
- [Nao03] Moni Naor, *On cryptographic assumptions and challenges*, Annual International Cryptology Conference, Springer, 2003, pp. 96–109.
- [NP05] Moni Naor and Benny Pinkas, *Computationally secure oblivious transfer*, Journal of Cryptology **18** (2005), no. 1, 1–35.
- [Pas16] Rafael Pass, *Unprovable security of perfect nizk and non-interactive non-malleable commitments*, computational complexity **25** (2016), no. 3, 607–666.
- [PS19] Chris Peikert and Sina Shiehian, *Noninteractive zero knowledge for np from (plain) learning with errors*, Tech. report, IACR Cryptology ePrint Archive, 2019.