

On Round Optimal Statistical Zero Knowledge Arguments

Nir Bitansky^{1*} and Omer Paneth^{2 †}

¹ Tel Aviv University

² MIT and Northeastern University

Abstract. We construct the first three message statistical zero knowledge arguments for all of NP, matching the known lower bound. We do so based on keyless multi-collision resistant hash functions and the Learning with Errors assumption — the same assumptions used to obtain round optimal computational zero knowledge.

The main component in our construction is a statistically witness indistinguishable argument of knowledge based on a new notion of statistically hiding commitments with subset opening.

1 Introduction

Since their introduction three decades ago [GMR89], the concept of zero knowledge protocols has played a central role in the development of modern cryptography. Different flavors of zero knowledge protocols have been studied according to the level of soundness and zero knowledge achieved. Either property can be *statistical* or *computational*, meaning that it holds against unbounded or computationally bounded adversaries, respectively. Protocols that satisfy both properties statistically, known as *statistical zero knowledge proofs*, are only possible for languages in $\mathbf{AM} \cap \mathbf{coAM}$ [For89, AH91]; however, once either property is relaxed to computational, protocols for all of \mathbf{NP} can be constructed assuming one way functions [GMW91, Nao91, BCC88, NOVY98, HNO⁺09].

In this work, we focus on *statistical zero knowledge arguments* for \mathbf{NP} ; namely, computationally sound protocols where the view of any efficient verifier can be efficiently simulated up to a negligible statistical difference. Such protocols are especially appealing due to their *everlasting zero knowledge guarantee* — even if the verifier stores conversations with the prover and post-processes

*Member of the Check Point Institute of Information Security. Supported by ISF grant 18/484, the Alon Young Faculty Fellowship, and by Len Blavatnik and the Blavatnik Family Foundation.

†Supported by NSF Grants CNS-1413964, CNS-1350619 and CNS-1414119, and 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.

them for as long as it wants, it does not learn anything that cannot be simulated efficiently.

A foundational question is the round complexity of such protocols. Four message protocols can be constructed based on collision resistant hash functions [BJY97]. In terms of lower bounds, at least three messages are necessary even for computational zero knowledge [GO94]. Constructing zero knowledge arguments with matching round complexity has so far appeared more difficult in the statistical setting than in the computational one. Computational zero knowledge in three messages has been long known under unfalsifiable knowledge assumptions [HT98, BP04, CD09, BCC⁺14] and recently also under a falsifiable multi-collision resistance assumption on keyless hash functions [BKP18]. In contrast, three message statistical zero knowledge has been out of reach (even under knowledge assumptions).

The difference between the statistical and computational settings seems to run somewhat deeper, and manifests itself even in *witness indistinguishability*, a relaxation of zero knowledge. While computational witness indistinguishability has been long known in three [GMW91, FLS99], two [DN07], or even one message [BOV07, GOS12], statistical witness indistinguishability in less than four messages has only been very recently achieved [KKS18]. In fact, *witness indistinguishable arguments of knowledge*, which are essential in the constructions of three message computational zero knowledge, are still not known in less than four messages in the statistical setting.

All and all, we are faced with the question:

What is the round complexity of statistical zero knowledge arguments?

1.1 Results

We construct the first round optimal statistical zero knowledge argument under the same assumptions on which computational zero knowledge arguments are currently known.

Theorem 1.1 (Informal). *Assuming keyless multi-collision resistant hash functions and LWE , both quasi-polynomially hard, there exist three message statistical zero knowledge arguments.*³

Keyless multi-collision resistant hash functions, introduced in [BKP18], are functions $H : \{0, 1\}^{2\lambda} \rightarrow \{0, 1\}^\lambda$ guaranteeing that no efficient adversary with non-uniform description of polynomial size S can find more than $\text{poly}(S)$ elements that collide under H . Here poly is some fixed polynomial and the adversary's running time may be an arbitrarily larger polynomial, or even quasipolynomial (as required in the above theorem). While at this point non-standard, multi-collision resistance of keyless hash functions is a falsifiable and relatively simple assumption, which plausibly holds for existing keyless hash functions (see discussion in [BKP18]).

³Here LWE is used to realize several generic primitives, which we address in the body.

Four message protocols from weaker assumptions. When considering keyed hash functions, multi-collision resistance becomes a standard assumption that relaxes the classical notion of collision resistance. A recent line of work explores such hash functions demonstrating that their power goes beyond one-way functions to achieve some of the applications of collision resistance [BDRV18, BKP18, KNY17, KNY18, BHKY19]. Our second result is along this vein showing that four message statistical zero knowledge arguments can be based on (keyed) multi-collision resistance instead of full fledged collision resistance.

Theorem 1.2 (Informal). *Assuming (keyed) multi-collision resistant hash functions, there exist four message statistical zero knowledge arguments.*

A main building block in both of the above results is a new statistically witness indistinguishable argument of knowledge, which was so far known from collision resistance in four messages.

Theorem 1.3 (Informal). *Assuming multi-collision resistant hash functions, there exist four message statistically witness indistinguishable arguments of knowledge. If the hash functions are keyless the arguments have three messages.*

Most of the technical effort in this work is devoted toward proving this theorem.

1.2 Technical Overview

We now provide an overview of the main ideas and techniques behind our results. We start with our construction of statistically witness indistinguishable arguments of knowledge and then move on to explain how they are used to construct statistical zero knowledge arguments.

Classical witness indistinguishable protocols from bit commitments. To understand the challenge, let us recall how classical witness indistinguishable arguments of knowledge are designed. Such protocols traditionally involve three basic steps: a prover commitment, a verifier challenge, and a prover decommitment. The prover commitment consists of multiple bit commitments, a subset of which are opened in the decommitment step according to the challenge given by the verifier. For instance, in Blum’s Hamiltonicity protocol [Blu86], the prover commits to the entries of the adjacency matrix of a graph and then, according to the challenge, either opens a subset of edges corresponding to a Hamiltonian cycle or the entire graph. This is repeated in parallel to decrease the soundness error to negligible. (The protocol contains additional details that we omit.)

Indeed, given statistically hiding bit commitments, such protocols can be shown to be statistically witness indistinguishable. However, focusing on round complexity, such commitments inherently require at least two messages — if the commitment was non-interactive, a cheating prover could have equivocal openings for some commitments non-uniformly hardwired into its code; such openings always exist due to statistical hiding.

Weakly binding commitments. While standard binding cannot be achieved non-interactively for statistically hiding bit commitments, keyless multi-collision resistant hash functions are known to imply non-interactive statistically hiding commitments with *weak binding* [BKP18]. Weak binding says that an attacker could only ever open a given commitment to values from some polynomial-size set. Intuitively, this means that even if it has equivocal openings of some commitment hardwired in its code, these cannot be used to sample more openings (except with negligible probability). Weak binding, however, is only meaningful in commitments for long strings and is completely meaningless for bit commitments, where the prover can open commitments to both zero and one. Accordingly, it is not clear how to use it in classical witness indistinguishable protocols that are all based on bit commitments.

An analogous problem is encountered in the work of [BKP18]. They define commitments for long strings with a subset opening property that enables to open only a given subset of bits, without having to open the entire string. While traditionally we require that every individual bit is fixed by the commitment, they suggest a relaxed definition suited for the setting of weak binding. They require that for any fixed adversary, a commitment to a long string $X \in \{0, 1\}^L$ fixes a *global* set of strings $\mathbf{X} \subseteq \{0, 1\}^L$ of polynomial size K , so that whenever the adversary opens some subset of bits $I \subseteq [L]$, they must all be *simultaneously consistent* with one string in the set \mathbf{X} (except with negligible probability over the adversary’s coins).

Our first observation is that commitments with subset opening and the above weak, but global, binding guarantee is sufficient for establishing soundness and also knowledge extraction in classical protocols. Roughly speaking, this is because for any prover that convinces the verifier of accepting with noticeable probability ε , there must be a single string $X \in \mathbf{X}$, such that with probability ε/K , the prover convinces the verifier *while answering consistently with X* . Since ε/K is still noticeable, soundness and knowledge extraction essentially reduce to those of the original protocol in the fully binding setting.

The work of [BKP18] constructs commitments with subset opening that fall short of achieving statistical hiding (indeed they focus on succinctness rather than hiding). In their construction, an opening of a given subset also reveals information regarding unopened bits, thus making them unfit for instantiating witness indistinguishable protocols. At high level, the reason they do not achieve statistical hiding is that to enforce consistency, subsets of bits are always opened in a correlated way and the correlations also pertain to unopened bits (this will become more clear below when we describe our construction).

Statistically hiding commitments with subset opening. We provide a construction that also achieves statistical hiding. Specifically, commitments to any two vectors X, X' are statistically indistinguishable, even given an opening of any subset of entries where X and X' agree. Our construction combines ideas from [BKP18] along with new ideas aimed toward statistical hiding.

The construction is based on statistically hiding weakly binding string commitments (with no subset opening) and Shamir secret sharing [Sha79]. At ab-

stract level, Shamir secret sharing with parameters (n, d) allows to sample an encoding $\hat{b} \in \Sigma^n$ over some alphabet Σ of a secret bit $b \in \{0, 1\}$, so that two properties are guaranteed. First, any two encodings $\hat{0}$ and $\hat{1}$ of zero and one agree on at most d entries. Second, any d entries of a random encoding \hat{b} perfectly hide the bit b .

The commitment scheme works as follows:

- **Commitment:** to commit to a string $X \in \{0, 1\}^L$, we sample encodings \hat{X}_i for all bits X_i of X . Then, considering the matrix $\hat{X} = \left(\hat{X}_{i,j} \right)$ whose rows are the encodings \hat{X}_i , we compute a statistically hiding weakly-binding commitment to each row and each column of the matrix \hat{X} .
- **Opening:** opening a subset of entries $I \subseteq [L]$ consists of two messages: a receiver challenge C , and a corresponding decommitment. The challenge consists of d random indices $C \subseteq [n]$. A decommitment involves opening the commitments to all rows $i \in I$ and all columns $j \in C$. The receiver verifies that the individual decommitments are valid and are consistent (on the intersection of corresponding rows and columns).

Let us explain at high level how the scheme satisfies statistical hiding and (global) weak binding. To show that any unopened entry $i \notin I$ remains statistically hidden, we rely on the fact that the commitment corresponding to its row \hat{X}_i is statistically hiding and never opened. Since only d column commitments are opened, only d entries of \hat{X}_i are revealed and thus the Shamir hiding property guarantees that the bit X_i remains statistically hidden.

Proving weak (global) binding is inspired by ideas from [BKP18]. Roughly speaking, the weak binding of individual row commitments guarantees that for every row i , the sender can only ever open the corresponding commitment to encodings \hat{X}_i from some fixed set \mathbf{S}_i of polynomial size $K = \text{poly}(\lambda)$ in the security parameter λ . We then use the fact that encodings of two distinct bits are far apart to argue that with overwhelming probability the answers to the random challenge (columns) uniquely fix all the bits of the string X . We further show that due to the individual weak binding of column commitments, these answers come from a polynomial set, and accordingly the string X determined by these answers must also come from a polynomial-size set \mathbf{X} .

In a bit more detail, we choose the parameters (n, d) of Shamir secret sharing so that the relative agreement of encodings of distinct bits is polynomially small $d/n < \lambda^{-\Omega(1)}$. Then choosing a large enough constant τ such that

$$L \cdot K^2 \cdot (d/n)^\tau \ll 1$$

guarantees that for τ random locations $T \subseteq [n]$, for all bits $i \in [L]$, any two encodings in \mathbf{S}_i that agree on T must encode the same bit.

By the weak binding of column commitments, any opened column $j \in T$ is taken from a fixed set \mathbf{S}'_j of polynomial size K , which means that all opened bits must simultaneously be consistent with some X taken from a set \mathbf{X}_T of polynomial size K^τ . The actual construction chooses d random challenges C for

super-logarithmic d , so that with overwhelming probability they include a fixing set $T \subseteq C$, and the set \mathbf{X} is the union of all corresponding sets \mathbf{X}_T (the number of which is at most d^τ and thus polynomial).

Statistical zero knowledge. We now explain how the statistical zero knowledge protocols behind Theorems 1.1 and 1.2 is obtained. The four message protocol from (keyed) multi-collision resistance is obtained in a black-box way by replacing the statistically witness indistinguishable argument of knowledge (from collision resistance) in the protocol of [BJY97] with our witness indistinguishable argument from multi-collision resistance. We focus on explaining how three-message statistical zero knowledge is obtained from keyless multi-collision resistant hashing.

Our starting point is the three-message *computational* zero knowledge argument of [BBK⁺16] and its subsequent extension in [BKP18] based on multi-collision resistant hash functions. At high level, their protocol follows the recipe of Barak’s non-black-box simulation technique [Bar01]. To prove an **NP** statement $x \in \mathcal{L}$, the prover sends a shrinking commitment cmt to the code of some (potentially long) program Π and the verifier responds with a random string r . Then, the prover gives a succinct witness indistinguishable argument of knowledge proving that either $x \in \mathcal{L}$ or that the committed program $\Pi(\text{cmt})$ outputs r . At high level, by committing to the code of the verifier itself, a non-black-box simulator is able to produce an accepting transcript without using the witness, while a cheating prover, who does not know the verifier’s randomness r , can only commit to such a program with negligible probability.

In [BBK⁺16, BKP18], the succinct witness indistinguishable argument of knowledge is constructed from a (non-succinct) witness indistinguishable argument of knowledge, a *secure function evaluation scheme*, and a *weak memory delegation scheme*, which they construct based on keyless multi-collision resistant hashing. Upgrading the protocol from computational zero knowledge to statistical zero knowledge requires two main changes. First, the prover commitment cmt is replaced with a statistically hiding commitment that is weakly binding, which as already observed in [BKP18] is sufficient. Second, the succinct witness indistinguishable argument of knowledge is replaced with a statistically witness indistinguishable one.

To obtain the succinct witness indistinguishable argument, we replace the (computationally) witness indistinguishable argument with our statistically witness indistinguishable argument. In addition, need the secure function evaluation scheme to satisfy statistical function hiding, which can be achieved assuming LWE [OPP14, BD18]. The actual construction requires that the witness indistinguishable argument possesses additional properties, such as adaptive witness indistinguishability and adaptive argument of knowledge, when the statement proven is adversarially chosen after the first two messages of the protocol. The protocol we construct is a slightly tweaked version of the protocol described in this introduction that satisfies these properties.

1.3 More Related Work

We next address additional related work in more detail.

More on statistical zero knowledge arguments. From the early constructions of zero knowledge protocols [GMW91, BCC88] it was evident that statistically hiding commitments are sufficient to obtain statistically zero knowledge arguments in a super logarithmic number of rounds. (See [HNO⁺09] for a survey on statistically hiding commitments.) Early constant round constructions [BCY91] were based on specific number theoretic assumptions. The work of Bellare, Jakobson, and Yung constructed computational zero knowledge arguments in four messages from one-way functions; however, their construction in fact uses a four message witness indistinguishable argument of knowledge in a generic manner. Using two-message statistically hiding commitments [DPP93, HM96], such arguments can be obtained from collision resistant hashing.

The round complexity of zero knowledge proofs. This paper focuses on the notion of (statistical zero knowledge) arguments. The round complexity of zero-knowledge proofs (which are statistically sound) for \mathbf{NP} has also been studied extensively. Four-message proofs are impossible to achieve via black-box simulation, except for languages in $\mathbf{NP} \cap \mathbf{coAM}$ [Kat12]. Four message proofs with non-black-box simulation are only known assuming multi-collision-resistance keyless hash functions [BKP18]. Recent evidence [FGJ18] suggests that, differently from zero-knowledge arguments, zero-knowledge proofs may be impossible to achieve in three messages (even with non-black-box simulation).

The black box barrier. Goldreich and Krawczyk show that three message computational (let alone, statistical) zero knowledge arguments cannot be achieved with black box simulation [GK96]. The seminal work of Barak was the first to show that *non-black-box simulation* could potentially cross such black box barriers [Bar01]. Works of Bitansky et al. [BCPR14, BBK⁺16] obtain three message (computational) zero knowledge arguments in case where either the (adversarial) verifier or prover have an a-priori bounded description (and arbitrary polynomial running time). Following, the work of [BKP18] obtains such arguments also against non-uniform verifiers and provers relying on keyless multi-collision resistance.

A stronger notion of statistical zero knowledge. The literature (e.g. [HNO⁺09]) also considers a stronger form of statistical zero knowledge than the one presented in this introduction where the simulator is not only required to statistically simulate the view of efficient verifiers, but also of inefficient ones, given oracle access to the verifier. We note that this notion is outright impossible in three messages where black box simulation is impossible [GK96]. Our four message protocol, in fact, does achieve this stronger notion.

2 Preliminaries

We rely on the standard computational concepts and notation:

- A PPT is a probabilistic polynomial-time algorithm.
- A uniform algorithm is T -time if it runs in time polynomial in T . (T may be super-polynomial in its input size.)
- We follow the standard habit of modeling any efficient adversary strategy as a family of polynomial-size circuits. For an adversary \mathcal{A} corresponding to a family of polynomial-size circuits $\{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$, we often omit the subscript λ , when it is clear from the context.
- We say that a function $f : \mathbb{N} \rightarrow \mathbb{R}$ is negligible if for all constants $c > 0$, there exists $N \in \mathbb{N}$ such that for all $n > N$, $f(n) < n^{-c}$. We sometimes denote negligible functions by *negl*.
- We say that a function $f : \mathbb{N} \rightarrow \mathbb{R}$ is noticeable if there exists a constant $c > 0$ and $N \in \mathbb{N}$ such that for all $n > N$, $f(n) \geq n^{-c}$.
- We denote statistical distance by **SD**.
- For two random variables X, Y and $\varepsilon \in [0, 1]$, we write $X \approx_\varepsilon Y$ to denote the fact that $\mathbf{SD}(X, Y) \leq \varepsilon$.
- For two ensembles $\mathcal{X} = \{X_\lambda\}_{\lambda \in \mathbb{N}}$ and $\mathcal{Y} = \{Y_\lambda\}_{\lambda \in \mathbb{N}}$, we say that \mathcal{X} and \mathcal{Y} are statistically indistinguishable if there exists a negligible $\mu(\cdot)$, such that for all λ , $X_\lambda \approx_{\mu(\lambda)} Y_\lambda$. We denote this by $\mathcal{X} \approx_s \mathcal{Y}$.
- For a string X of length n , and a subset $I \subseteq [n]$, we denote by $X|_I$ its restriction to the entries in I .

2.1 Statistical Zero-Knowledge Arguments

In what follows, we denote by $\langle \mathsf{P} \rightleftharpoons \mathsf{V} \rangle$ a protocol between two parties P and V . For input w for P , and common input x , we denote by $\langle \mathsf{P}(w) \rightleftharpoons \mathsf{V} \rangle(x)$ the output of V in the protocol. For honest verifiers this output will be a single bit indicating acceptance (or rejection), whereas malicious verifiers output their entire view. Throughout, we assume that honest parties in all protocols are uniform PPT algorithms.

Definition 2.1. A protocol $\langle \mathsf{P} \rightleftharpoons \mathsf{V} \rangle$ for an **NP** relation $\mathcal{R}(x, w)$ is a statistical zero-knowledge argument if it satisfies:

Completeness: For any $\lambda \in \mathbb{N}, x \in \mathcal{L}(\mathcal{R}) \cap \{0, 1\}^\lambda, w \in \mathcal{R}(x)$:

$$\Pr [\langle \mathsf{P}(w) \rightleftharpoons \mathsf{V} \rangle(x) = 1] = 1 .$$

Computational Soundness: For every polynomial-size circuit family of provers $\mathsf{P}^* = \{\mathsf{P}_\lambda^*\}_\lambda$, there exists a negligible function μ , such that for any $x \in \{0, 1\}^\lambda \setminus \mathcal{L}(\mathcal{R})$,

$$\Pr [\langle \mathsf{P}_\lambda^* \rightleftharpoons \mathsf{V} \rangle(x) = 1] \leq \mu(\lambda) .$$

Statistical Zero-Knowledge: There exists a PPT simulator S such that for every polynomial-size circuit family $\mathsf{V}^* = \{\mathsf{V}_\lambda^*\}_\lambda$:

$$\{\langle \mathsf{P}(w) \rightleftharpoons \mathsf{V}_\lambda^* \rangle(x)\}_{\substack{(x,w) \in \mathcal{R} \\ |x|=\lambda}} \approx_s \{\mathsf{S}(x, \mathsf{V}_\lambda^*)\}_{\substack{(x,w) \in \mathcal{R} \\ |x|=\lambda}} .$$

2.2 Weakly Binding Commitments and Multi-Collision Resistant Hash Functions

We define weakly-binding statistically-hiding commitments [BKP18]. The definition addresses both the setting of keyed hash functions as well as keyless ones.

Syntax: A commitment scheme is associated with an input length function $\ell(\lambda)$ and polynomial-time algorithms $\text{SHC} = (\text{SHC.Gen}, \text{SHC.Com})$ with the following syntax:

- $\text{pk} \leftarrow \text{SHC.Gen}(1^\lambda)$: a probabilistic algorithm that takes the security parameter 1^λ and outputs a key $\text{pk} \in \{0, 1\}^\lambda$. In the keyless setting, this algorithm is deterministic and outputs a fixed key $\text{pk} \equiv 1^\lambda$.
- $\text{cmt} \leftarrow \text{SHC.Com}(X; \text{pk})$: a probabilistic algorithm that takes the key pk and an input $X \in \{0, 1\}^{\ell(\lambda)}$ and outputs a commitment cmt . When we want to be explicit about the randomness r used by the algorithm, we may write $\text{SHC.Com}(X; \text{pk}, r)$.

Definition 2.2 (Weakly-Binding Statistically-Hiding Commitments). For a polynomial $\ell(\cdot)$, a weakly-binding statistically-hiding commitment $\text{SHC} = (\text{SHC.Gen}, \text{SHC.Com})$, for messages of length ℓ , satisfies:

Statistical Hiding: For any key and any two plaintexts, the corresponding commitments are statistically close:

$$\left\{ \text{SHC.Com}(X; \text{pk}) \right\}_{\substack{\lambda \in \mathbb{N}, \text{pk} \in \{0, 1\}^\lambda \\ X, X' \in \{0, 1\}^{\ell(\lambda)}}} \approx_{2^{-\lambda}} \left\{ \text{SHC.Com}(X'; \text{pk}) \right\}_{\substack{\lambda \in \mathbb{N}, \text{pk} \in \{0, 1\}^\lambda \\ X, X' \in \{0, 1\}^{\ell(\lambda)}}} .$$

Weak Binding: For any non-uniform polynomial-size probabilistic $\mathcal{A} = \{\mathcal{A}_\lambda^1, \mathcal{A}_\lambda^2\}_\lambda$ there exists a polynomial $K(\cdot)$ and a negligible $\mu(\cdot)$, such that for all $\lambda \in \mathbb{N}$,

$$\Pr_{\substack{\text{pk} \leftarrow \text{SHC.Gen}(1^\lambda) \\ (\text{cmt}, \text{st}) \leftarrow \mathcal{A}_\lambda^1(\text{pk})}} \left[\begin{array}{l} \exists \mathbf{X} \text{ of size } K(\lambda) : \\ \Pr_{(X, r) \leftarrow \mathcal{A}_\lambda^2(\text{st})} \left[\begin{array}{l} \text{cmt} = \text{SHC.Com}(X; \text{pk}, r) \\ X \notin \mathbf{X} \end{array} \right] \leq \mu(\lambda) \end{array} \right] \geq 1 - \mu(\lambda) .$$

Multi-collision resistance. We also define multi-collision resistant hash functions [BKP18], which are similar to weakly binding commitments, only that the hiding requirement is replaced with the requirement that they shrink their input (accordingly, they are also deterministic).

Syntax: A hashing scheme is associated with an input length function $\ell(\lambda)$ and polynomial-time algorithms $\text{H} = (\text{H.Gen}, \text{H.Hash})$ with the following syntax:

- $\text{pk} \leftarrow \text{H.Gen}(1^\lambda)$: a probabilistic algorithm that takes the security parameter 1^λ and outputs a key $\text{pk} \in \{0, 1\}^\lambda$. In the keyless setting, this algorithm is deterministic and outputs a fixed key $\text{pk} \equiv 1^\lambda$.
- $Y \leftarrow \text{H.Hash}(X; \text{pk})$: a deterministic algorithm that takes the key pk and an input $X \in \{0, 1\}^{\ell(\lambda)}$ and outputs a hash value Y .

Definition 2.3 (Multi-Collision Resistant Hash). For a polynomial $\ell(\cdot)$, a multi-collision resistant hash $H = (H.Gen, H.Hash)$, for messages of length ℓ , satisfies:

Compression: $\ell(\lambda) > \lambda$ and $|H.Hash(X)| = \lambda$ for all $\lambda \in \mathbb{N}$, key $pk \in \{0, 1\}^\lambda$, and $X \in \{0, 1\}^{\ell(\lambda)}$. If $\ell = \lambda \cdot (1 + \Omega(1))$ we say that H is linearly compressing and if $\ell = \lambda^{1+\Omega(1)}$ we say that H is polynomially compressing.

Multi-Collision Resistance: For any non-uniform polynomial-size probabilistic $\mathcal{A} = \{\mathcal{A}_\lambda^1, \mathcal{A}_\lambda^2\}_\lambda$ there exists a polynomial K and a negligible $\mu(\cdot)$, such that for all $\lambda \in \mathbb{N}$,

$$\Pr_{\substack{pk \leftarrow \text{SHC.Gen}(1^\lambda) \\ (Y, st) \leftarrow \mathcal{A}_\lambda^1(pk)}} \left[\exists \mathbf{X} \text{ of size } K(\lambda) : \Pr_{X \leftarrow \mathcal{A}_\lambda^2(st)} [Y = H.Hash(X; pk) \wedge X \notin \mathbf{X}] \leq \mu(\lambda) \right] \geq 1 - \mu(\lambda) .$$

We also consider a generalized notion of T -secure multi-collision resistant hashing that allows addressing attackers that run in super-polynomial time. The constructions in this paper all rely on the above polynomial notion. Quasipolynomial security is used in [BKP18] to construct weak memory delegation as defined in Section 2.3. We state the definition here for completeness.

T -Secure Multi-Collision Resistance: For any non-uniform polynomial-size probabilistic $\mathcal{A} = \{\mathcal{A}_\lambda^1\}$ and any uniform T -time \mathcal{A}_2 there exists a polynomial K and a negligible $\mu(\cdot)$, such that for all $\lambda \in \mathbb{N}$,

$$\Pr_{\substack{pk \leftarrow \text{SHC.Gen}(1^\lambda) \\ (Y, st) \leftarrow \mathcal{A}_\lambda^1(pk)}} \left[\exists \mathbf{X} \text{ of size } K(\lambda) : \Pr_{X \leftarrow \mathcal{A}_2(st)} [Y = H.Hash(X; pk) \wedge X \notin \mathbf{X}] \leq \mu(\lambda) \right] \geq 1 - \mu(\lambda) .$$

Remark 2.1. Note that for polynomial $T(\lambda) = \text{poly}(\lambda)$, T -security coincides with (plain) security. Indeed, the non-uniformity of \mathcal{A}^2 can always be pushed to \mathcal{A}^1 who passes a state to \mathcal{A}^2 .

In [BKP18], it is shown that multi-collision resistant hashing implies weakly binding string commitments.

Theorem 2.1 ([BKP18]). Assuming a multi-collision-resistant keyless hash that is either:

- polynomially compressing
- or, linearly compressing and quasipoly(λ)-secure

there exist, for every polynomial $L(\cdot)$, a weakly-binding statistically-hiding commitment and multi-collision-resistant keyless hash, both for messages of length L .

2.3 Weak Memory Delegation

The notion of weak memory delegation was defined in [BKP18] as a relaxation of memory delegation [CKLR11, KP16]. In a two-message memory delegation scheme, an untrusted server provides the client a short commitment or *digest* dig of a large memory D . The client can then delegate any arbitrary deterministic computation M to be executed over the memory. The server responds with the computation's output y , as well as a short proof of correctness that can be verified by the client in time that is independent of that of the delegated computation and the size of the memory.

In the definition of memory delegation in [KP16], the soundness requirement says that having provided the digest dig , a cheating prover should not be able to prove that a given computation M results in more than a single outcome y . In weak memory delegation, the prover should not be able to prove consistency with *too many* outcomes y .

Syntax: A two-message memory delegation scheme is associated with polynomial-time algorithms $(\text{MD.Mem}, \text{MD.Query}, \text{MD.Prove}, \text{MD.Ver})$ with the following syntax:

- $\text{dig} \leftarrow \text{MD.Mem}(1^\lambda, D)$: a deterministic polynomial-time algorithm that given the security parameter 1^λ and memory D , outputs a digest $\text{dig} \in \{0, 1\}^\lambda$ of the memory.
- $(\mathbf{q}, \text{vst}) \leftarrow \text{MD.Query}(1^\lambda)$: a randomized polynomial-time algorithm that given the security parameter 1^λ , outputs a query \mathbf{q} and a secret state vst .
- $\pi \leftarrow \text{MD.Prove}(1^\lambda, D, (M, 1^t, y), \mathbf{q})$: a deterministic algorithm that takes the security parameter 1^λ , a memory string D , a (deterministic) Turing machine M , an output string y , and time bound 1^t such that $|D| \leq t \leq 2^\lambda$ and $M(D)$ outputs y within t steps. It outputs a proof π .
- $b \leftarrow \text{MD.Ver}(1^\lambda, \text{dig}, (M, t, y), \text{vst}, \pi)$: a deterministic polynomial time oracle algorithm that takes the security parameter 1^λ , a digest dig , a (deterministic) Turing machine M , a time bound t , an output string y , a secret state vst and a proof π . It outputs an acceptance bit b .

Definition 2.4 (Entropic Distribution Ensemble). We say that an efficiently samplable distribution ensemble $\{Y_\lambda\}_{\lambda \in \mathbb{N}}$ is entropic if

$$H_\infty(Y_\lambda) := -\log \max_{y \in \text{supp}(Y_\lambda)} \Pr[Y_\lambda = y] = \Omega(\lambda) .$$

Definition 2.5 (Weak Memory Delegation). A two-message delegation scheme

$\text{MD} = (\text{MD.Mem}, \text{MD.Query}, \text{MD.Prove}, \text{MD.Ver})$ satisfies:

Efficiency: There exists a polynomial p such that for every $\lambda \in \mathbb{N}$ and D such that $|D| \leq 2^\lambda$, $\text{MD.Mem}(1^\lambda, D)$ outputs a digest dig of length at most $p(\lambda)$.

Correctness: For every security parameter $\lambda \in \mathbb{N}$, every $(M, t, y) \in \{0, 1\}^\lambda$, and every D such that $M(D)$ outputs y within t steps, and $|D| \leq t \leq 2^\lambda$:

$$\Pr \left[\text{MD.Ver}(1^\lambda, \text{dig}, (M, t, y), \text{vst}, \pi) = 1 \mid \begin{array}{l} \text{dig} \leftarrow \text{MD.Mem}(1^\lambda, D) \\ (\text{q}, \text{vst}) \leftarrow \text{MD.Query}(1^\lambda) \\ \pi \leftarrow \text{MD.Prove}(1^\lambda, D, (M, 1^t, y), \text{q}) \end{array} \right] = 1 .$$

Weak Soundness for Time- T : For every non-uniform polynomial-size probabilistic $(\mathcal{A}_1, \mathcal{A}_2)$, there exists a negligible function μ , such that for every ensemble of samplable entropic distributions $\{Y_\lambda\}_{\lambda \in \mathbb{N}}$, $\lambda \in \mathbb{N}$ and $t \leq T(\lambda)$:

$$\Pr \left[\text{MD.Ver}(1^\lambda, \text{dig}, (M, t, y), \text{vst}, \pi) = 1 \mid \begin{array}{l} (\text{dig}, M, \text{st}) \leftarrow \mathcal{A}_1(1^\lambda) \\ (\text{q}, \text{vst}) \leftarrow \text{MD.Query}(1^\lambda) \\ y \leftarrow Y_\lambda \\ \pi \leftarrow \mathcal{A}_2(\text{q}, y; \text{st}) \end{array} \right] \leq \mu(\lambda) .$$

Theorem 2.2 ([BKP18]). Assuming a linearly compressing quasipoly(λ)-secure multi-collision-resistant keyless hash and quasipoly(λ)-secure fully-homomorphic encryption, there exists a two-message memory-delegation scheme with weak soundness for time-quasipoly(λ).

2.4 Function Hiding Secure Function Evaluation

We define two-message secure function evaluation protocols with statistical function hiding.

Syntax: Let $\mathcal{C} = \{\mathcal{C}_\lambda\}_\lambda$ be a family of circuits. A two-message secure function evaluation protocol for \mathcal{C} is associated with polynomial-time algorithms (SFE.Enc, SFE.Eval, SFE.Dec) with the following syntax:

- $(\text{sk}, \text{ct}) \leftarrow \text{SFE.Enc}(1^\lambda, x)$: a probabilistic algorithm that takes a security parameter 1^λ and a string $x \in \{0, 1\}^*$ and outputs a secret key sk and a ciphertext ct .
- $\widehat{\text{ct}} \leftarrow \text{SFE.Eval}(C, \text{ct})$: a probabilistic algorithm that takes a circuit $C \in \mathcal{C}$ and a ciphertext ct and outputs an evaluated ciphertext $\widehat{\text{ct}}$.
- $\widehat{x} \leftarrow \text{SFE.Dec}(\widehat{\text{ct}}; \text{sk})$: a deterministic algorithm that takes a ciphertext $\widehat{\text{ct}}$ and the secret key sk and outputs a string \widehat{x} .

Definition 2.6. A two-message secure function evaluation protocol (SFE.Enc, SFE.Eval, SFE.Dec) for a family of circuits $\mathcal{C} = \{\mathcal{C}_\lambda\}_\lambda$ satisfies:

- **Perfect Correctness:** For any $\lambda \in \mathbb{N}$, $x \in \{0, 1\}^*$ and circuit $C \in \mathcal{C}_\lambda$

$$\Pr \left[\text{SFE.Dec}(\widehat{\text{ct}}; \text{sk}) = C(x) \mid \begin{array}{l} (\text{sk}, \text{ct}) \leftarrow \text{SFE.Enc}(x; 1^\lambda) \\ \widehat{\text{ct}} \leftarrow \text{SFE.Eval}(C, \text{ct}) \end{array} \right] = 1 .$$

- **Semantic Security:** For any polynomial $\ell(\lambda)$ and non-uniform polynomial-size probabilistic $\mathcal{A} = \{\mathcal{A}_\lambda\}_\lambda$, there exists a negligible function ν such that every $\lambda \in \mathbb{N}$, and pair of messages $x_0, x_1 \in \{0, 1\}^{\ell(\lambda)}$:

$$\Pr \left[\mathcal{A}_\lambda(\text{ct}) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ (\text{sk}, \text{ct}) \leftarrow \text{SFE.Enc}(x_b; 1^\lambda) \end{array} \right] \leq \frac{1}{2} + \mu(\lambda) .$$

- **Statistical Circuit Privacy:** There exist unbounded algorithms Sim, Ext such that:

$$\left\{ \text{SFE.Eval}(C, \text{ct}^*) \right\}_{\substack{\lambda \in \mathbb{N}, C \in \mathcal{C}_\lambda \\ \text{ct}^* \in \{0, 1\}^{\text{poly}(\lambda)}}} \approx_s \left\{ \text{Sim}(C(\text{Ext}(\text{ct}^*; 1^\lambda)); 1^\lambda) \right\}_{\substack{\lambda \in \mathbb{N}, C \in \mathcal{C}_\lambda \\ \text{ct}^* \in \{0, 1\}^{\text{poly}(\lambda)}}} .$$

Such secure function evaluation schemes are known based on LWE [OPP14, BV11, BD18].

2.5 Shamir Secret Sharing

We define Shamir secret sharing schemes.

Syntax: A Shamir secret sharing scheme is associated with functions $\delta(\lambda), n(\lambda)$, a field \mathbb{F}_λ and a probabilistic polynomial-time encoding algorithm $\hat{S} \leftarrow \text{SSS.Enc}(S; 1^\lambda)$ that takes a secret $S \in \mathbb{F}$ and a parameter 1^λ and outputs an encoding $\hat{S} \in \mathbb{F}^n$.

Definition 2.7 (Shamir Secret Sharing). For polynomials $\delta(\cdot), n(\cdot)$ A Shamir secret sharing encoding SSS.Enc satisfies:

Perfect Hiding: Any δ coordinates in the encoding are perfectly hiding:

$$\left\{ \text{SSS.Enc}(S_0; 1^\lambda)|_I \right\}_{\substack{\lambda \in \mathbb{N}, \\ S_0, S_1 \in \mathbb{F}, \\ I \in \binom{[n]}{\delta}}} \equiv \left\{ \text{SSS.Enc}(S_1; 1^\lambda)|_I \right\}_{\substack{\lambda \in \mathbb{N}, \\ S_0, S_1 \in \mathbb{F}, \\ I \in \binom{[n]}{\delta}}} ,$$

where $\binom{[n]}{\delta}$ denotes the collection of subsets $I \subseteq [n]$ of size δ .

Distance: For any $\lambda \in \mathbb{N}$ and distinct secrets $S_0, S_1 \in \mathbb{F}$,

$$\Delta(\text{SSS.Enc}(S_0; 1^\lambda), \text{SSS.Enc}(S_1; 1^\lambda)) \geq 1 - \delta/n ,$$

where Δ denotes the relative hamming distance over \mathbb{F}^n .

Shamir secret sharing schemes are known to exist unconditionally [Sha79].

3 Weakly-Binding Commitments with Subset Opening

In this section, we define and construct weakly-binding statistically hiding commitments with subset opening.

3.1 Definition

The definition addresses both the setting of keyed hash functions as well as keyless ones.

Syntax: A commitment with subset opening is associated with a length function $L(\lambda)$ and polynomial-time algorithms $\text{SHC} = (\text{CSO.Gen}, \text{CSO.Com}, \text{CSO.Chal}, \text{CSO.Open}, \text{CSO.Ver})$ with the following syntax:

- $\text{pk} \leftarrow \text{CSO.Gen}(1^\lambda)$: a probabilistic algorithm that takes the security parameter 1^λ and outputs a key $\text{pk} \in \{0, 1\}^\lambda$. In the keyless setting, this algorithm is deterministic and outputs a fixed key $\text{pk} \equiv 1^\lambda$.
- $(\text{cmt}, \text{st}) \leftarrow \text{CSO.Com}(X; \text{pk})$: a probabilistic algorithm that takes the key pk and a string $X \in \{0, 1\}^{L(\lambda)}$ and outputs a commitment cmt and private state st .
- $C \leftarrow \text{CSO.Chal}(\text{pk})$: a probabilistic algorithm that takes the key pk and outputs a challenge C .
- $d \leftarrow \text{CSO.Open}(I, C, \text{st})$: a deterministic algorithm that takes an index set I , a challenge C and private state st and outputs a decommitment d .
- $b \leftarrow \text{CSO.Ver}(\text{cmt}, \alpha, I, C, d)$: a deterministic algorithm that takes a commitment cmt , an index set I , an assignment $\alpha : I \rightarrow \{0, 1\}$, and decommitment d and outputs an acceptance bit b .

Definition 3.1 (Weakly-Binding Statistically-Hiding Commitments with Subset Opening). For a polynomial $L(\cdot)$, a weakly-binding statistically-hiding commitment with subset opening $\text{CSO} = (\text{CSO.Gen}, \text{CSO.Com}, \text{CSO.Chal}, \text{CSO.Open}, \text{CSO.Ver})$, for strings of length L , satisfies:

Subset Statistical Hiding: There exists a negligible $\mu(\cdot)$ such that for all $\lambda \in \mathbb{N}$, no unbounded adversary \mathcal{A} wins the following game with probability greater than $1/2 + \mu(\lambda)$:

1. \mathcal{A} submits to a challenger $\text{pk} \in \{0, 1\}^\lambda$, $X_0, X_1 \in \{0, 1\}^{L(\lambda)}$.
2. The challenger samples a random $b \leftarrow \{0, 1\}$ and $(\text{cmt}_b, \text{st}_b) \leftarrow \text{CSO.Com}(X_b; \text{pk})$, and gives cmt_b to \mathcal{A} .
3. \mathcal{A} submits a commitment challenge C .
4. The challenger computes $d = \text{CSO.Open}(I, C, \text{st})$ where $I = \{i \in [L] : X_0[i] = X_1[i]\}$ is the set of indices on which the strings X_0, X_1 agree.
5. \mathcal{A} wins if it correctly guesses the bit b .

Weak Binding: For any non-uniform polynomial-size probabilistic $\mathcal{A} = \{\mathcal{A}_\lambda^1, \mathcal{A}_\lambda^2\}_\lambda$ there exists a polynomial $K(\cdot)$ and a negligible $\mu(\cdot)$, such that for all $\lambda \in \mathbb{N}$,

$$\Pr_{\substack{\text{pk} \leftarrow \text{CSO.Gen}(1^\lambda) \\ (\text{cmt}, \text{st}) \leftarrow \mathcal{A}_\lambda^1(\text{pk})}} \left[\begin{array}{l} \exists \mathbf{X} \text{ of size } K(\lambda) : \\ \Pr_{\substack{C \leftarrow \text{CSO.Chal}(1^\lambda) \\ (\alpha, I, d) \leftarrow \mathcal{A}_\lambda^2(C; \text{st})}} \left[\begin{array}{l} \text{CSO.Ver}(\text{cmt}, \alpha, I, C, d) = 1 \\ \alpha \notin \mathbf{X}_I \end{array} \right] \leq \mu(\lambda) \end{array} \right] \geq 1 - \mu(\lambda) ,$$

where \mathbf{X}_I denotes the set of assignments $X|_I : I \rightarrow \{0, 1\}$ to indices in I induced by every $X \in \mathbf{X}$.

Succinct Commitment: $|\text{cmt}| = \lambda$ for any $\text{pk} \in \{0, 1\}^\lambda$ and any cmt in the support of $\text{CSO.Com}(\cdot; \text{pk})$.

In the remainder of this section we construct weakly-binding statistically hiding commitments with subset opening.

Theorem 3.1. *Assuming a polynomially compressing multi-collision-resistant keyless hash there exists a weakly-binding statistically-hiding commitment with subset opening for strings of length L for any polynomial L .*

3.2 Construction

We provide a construction of weakly-binding statistically-hiding commitments with subset opening from (plain) weakly-binding statistically-hiding commitments and multi-collision resistant hash functions.

Ingredients:

- SSS.Enc a Shamir secret sharing encoding with parameters $\delta(\lambda) = \sqrt{\lambda}$, $n(\lambda) = \lambda$, and field $\mathbb{F} = \{\mathbb{F}_\lambda\}_\lambda$.
- SHC a weakly binding statistically hiding commitment for strings of length $\ell = |\mathbb{F}_\lambda| \cdot \max(L, n)$. We denote the size of commitment in this scheme by $t(\lambda)$.
- H a multi-collision resistant hash for strings of length $\ell' = (L + n) \cdot t$.

The scheme CSO:

- CSO.Gen(1^λ) : Runs the key generator for the underlying commitment $\text{pk}_1 \leftarrow \text{CSO.Gen}(1^\lambda)$ and hash $\text{pk}_2 \leftarrow \text{H.Gen}(1^\lambda)$, and outputs $\text{pk} = (\text{pk}_1, \text{pk}_2)$.
- CSO.Com($X; \text{pk}$) :
 - For $i \in [L]$, compute an encoding $\widehat{X}_i \leftarrow \text{SSS.Enc}(X_i; 1^\lambda)$.
 - Consider the matrix $\widehat{X} = (\widehat{X}_{i,j})_{i \in [L], j \in [n]}$.
 - Compute commitments to the rows $\text{rowcmt}_i \leftarrow \text{SHC.Com}(\widehat{X}_{i,1}, \dots, \widehat{X}_{i,n}; \text{pk}_1)$.
 - Compute commitments to the columns $\text{colcmt}_j \leftarrow \text{SHC.Com}(\widehat{X}_{1,j}, \dots, \widehat{X}_{L,j}; \text{pk}_1)$.
 - Compute a hash of all of the above commitments $Y \leftarrow \text{H.Hash}((\text{rowcmt}_i)_i, (\text{colcmt}_j)_j; \text{pk}_2)$ and output $\text{cmt} = Y$ as the commitment.
 - Output as the state st all the commitments $(\text{rowcmt}_i)_i, (\text{colcmt}_j)_j$, all the randomness $(\hat{r}_i)_i, (\text{rr}_i)_i, (\text{rc}_j)_j$ used to generate the encodings \widehat{X}_i and commitments $\text{rowcmt}_i, \text{colcmt}_j$, respectively, and the encoding \widehat{X}_i themselves.
- CSO.Chal(pk_1) : sample δ random column indices $j_1, \dots, j_\delta \leftarrow [n]$ and output $C = (j_i)_{i \in [\delta]}$.

- $\text{CSO.Open}(I, C, \text{st})$: output as the decommitment information d all the commitments $(\text{rowcmt}_i)_i, (\text{colcmt}_j)_j$, the randomness $(\hat{r}_i, \text{rr}_i)_{i \in I}$ used to compute the encodings and commitments corresponding to all rows $i \in I$, the randomness $(\text{rc}_j)_{j \in C}$ and the columns $(\hat{X}_{1,j}, \dots, \hat{X}_{L,j})_{j \in C}$ themselves, corresponding to all challenge columns $j \in C$.
- $\text{CSO.Ver}(\text{cmt}, \alpha, I, C, d)$:
 - Parse d as $(\text{rowcmt}_i)_i, (\text{colcmt}_j)_j, (\text{rr}_i, \hat{r}_i)_{i \in I}, (\text{rc}_j, \hat{X}_{1,j}, \dots, \hat{X}_{L,j})_{j \in C}$.
 - Verify that $\text{cmt} = \text{H.Hash}(((\text{rowcmt}_i)_i, (\text{colcmt}_j)_j); \text{pk}_2)$.
 - For every $i \in I$, compute $\hat{\alpha}_i := \text{SSS.Enc}(\alpha(i); \hat{r}_i)$ and verify that $\text{rowcmt}_i = \text{SHC.Com}(\hat{\alpha}_i; \text{pk}_1, \text{rr}_i)$.
 - For every $j \in C$, verify that $\text{colcmt}_j = \text{SHC.Com}(\hat{X}_{1,j}, \dots, \hat{X}_{L,j}; \text{pk}_1, \text{rc}_j)$ and that for every $i \in I$, $\alpha_{i,j} = \hat{X}_{i,j}$.

3.3 Analysis

We now analyze the construction. We first prove subset statistical hiding and then prove weak binding.

Proposition 3.1. *The construction is subset statistically hiding.*

Proof. We first claim that statistical hiding holds for any fixed challenge set C .

Claim 3.2. *Fix any pk and $C \in [n]^\delta$, and set of indices $I \subseteq [n]$. There exists a simulator S such that for any $X \in \{0, 1\}^L$,*

$$\text{cmt}, d \approx_{2^{-\Omega(\lambda)}} S(X|_I) ,$$

where $(\text{cmt}, \text{st}) \leftarrow \text{CSO.Com}(X; \text{pk})$ and $d \leftarrow \text{CSO.Open}(I, C, \text{st})$.

Proof. We describe how the simulator samples the commitments $\text{rowcmt}_i, \text{colcmt}_j$:

- For every $i \in I$, compute an encoding \hat{X}_i of X_i , and sample a commitment rowcmt_i to this encoding.
- For every $i \notin I$, sample a commitment rowcmt_i to the all-zero string. Also sample an encoding \hat{Y}_i of for an arbitrary bit Y , and store \hat{Y}_i .
- For every $j \in C$, sample a commitment colcmt_j to the j -th column of the matrix whose rows are given by the encodings \hat{X}_i for $i \in I$ and by the encodings \hat{Y}_i for $i \notin I$.
- For every $j \notin C$ sample a commitment colcmt_j to the all-zero string.

We now argue that the commitments $\text{rowcmt}_i, \text{colcmt}_j$ produced above along with an opening with respect to I, C are $2^{-\Omega(\lambda)}$ -close to their distribution in a commitment to X .

Consider a hybrid distribution cmt^*, d^* where we change the distribution of commitments to X as follows. For all unopened rows $i \notin I$, we change the commitment rowcmt_i from a commitment to \hat{X}_i to a commitment to an all-zero

string, and for all unopened columns $j \notin C$, we change the commitment colcmt_j from a commitment to the j -column of the matrix $(\widehat{X}_{i,j})_{i,j}$ to a commitment to the all zero string.

Then by the statistical hiding of the underlying commitment SHC,

$$\text{cmt}, d \approx_{2^{-\lambda \cdot (L+n)}} \text{cmt}^*, d^* ,$$

where $2^{-\lambda(L+n)} \leq 2^{-\Omega(\lambda)}$.

Next, note that the only difference between the hybrid distribution cmt^*, d^* and the simulated distribution $S(X|I)$ is in the commitments to the columns $j \in C$. In the first, the plaintexts are the columns of the matrix $(\widehat{X}_{ij})_{i \in [L], j \in C}$, whereas in the second its the concatenation of $(\widehat{X}_{ij})_{i \in I, j \in C}$ and $(\widehat{Y}_{ij})_{i \in [L] \setminus I, j \in C}$. However, since $|C| \leq \delta$, the perfect hiding of the Shamir secret sharing implies that these two matrices are identically distributed.

Now fix any key pk and any two $X, X' \in \{0, 1\}^L$, and let $I \subseteq [L]$ be the set of indices on they agree. Then by Claim 3.2, for any fixed challenge C , the distributions cmt, d and cmt', d' corresponding to X and X' , respectively, are $2^{-\Omega(\lambda)}$ -close.

To complete the proof, we now show that they remain close also when C is chosen adaptively.

Claim 3.3. *For any (unbounded) adversary \mathcal{A} ,*

$$\text{cmt}, C, d \approx_{2^{-\Omega(\lambda)}} \text{cmt}', C', d' ,$$

where $(\text{cmt}, \text{st}) \leftarrow \text{CSO.Com}(X; \text{pk})$, $C \leftarrow \mathcal{A}(\text{cmt})$ and $d \leftarrow \text{CSO.Open}(I, C, \text{st})$, and cmt', C', d' is sampled similarly with respect to X' .

Proof.

$$\text{SD}((\text{cmt}, C, d), (\text{cmt}', C', d')) =$$

$$\sum_{\alpha, \beta, \gamma} |\Pr[(\text{cmt}, C, d) = (\alpha, \beta, \gamma)] - \Pr[(\text{cmt}', C', d') = (\alpha, \beta, \gamma)]| \leq \sum_{\beta \in [n]^\delta} 2^{-\Omega(\lambda)} = \lambda^{\sqrt{\lambda}} \cdot 2^{-\Omega(\lambda)} \leq 2^{-\Omega(\lambda)} ,$$

where the first inequality follows from Claim 3.2.

This completes the proof of Proposition 3.1

Proposition 3.2. *The construction is weakly binding.*

Proof. Fix a polynomial-size adversary $\mathcal{A} = \{\mathcal{A}_\lambda^1, \mathcal{A}_\lambda^2\}_\lambda$ against the commitment. The proof is divided to two main claims. We first prove a (computational) claim attesting that with overwhelming probability \mathcal{A} 's commitment fixes

a polynomial-size set of strings \mathbf{S} , such that any valid opening of the underlying weakly-binding commitment is to a string from \mathbf{S} . This claim relies on the weak binding of the underlying commitment and the multi-collision resistance of the hash function. Then we prove an information-theoretic claim that shows that provided the restriction to the set \mathbf{S} there also exists a polynomial-size global set of strings \mathbf{X} , such that any opening of some subset must be consistent with one of the strings in \mathbf{X} .

The following claim asserts that a commitment from the adversary \mathcal{A} , fixes a polynomial-size set of strings $\mathbf{S} \subseteq \{0, 1\}^\ell$, such that the adversary can only open any commitments $\text{rowcmt}_i, \text{colcmt}_j$ to a string from \mathbf{S} .

Claim 3.4. *There exist a polynomial $K(\cdot)$ and a negligible function $\mu(\cdot)$ such that for all $\lambda \in \mathbb{N}$, except with probability $\mu(\lambda)$ over $\text{pk} \leftarrow \text{CSO.Gen}(1^\lambda)$ and $(\text{cmt}, \text{st}) \leftarrow \mathcal{A}_\lambda^1(\text{pk})$ there exists a set $\mathbf{S} \subseteq \{0, 1\}^\ell$ of size $K(\lambda)$ such that*

$$\Pr_{\substack{C \leftarrow \text{CSO.Chal}(1^\lambda) \\ (\alpha, I, d) \leftarrow \mathcal{A}_\lambda^2(C; \text{st})}} \left[\text{CSO.Ver}(\text{cmt}, \alpha, I, C, d) = 1 \wedge \widehat{X}(\alpha, I, d) \not\subseteq \mathbf{S} \right] \leq \mu(\lambda) ,$$

where $\widehat{X}(\alpha, I, d) \subseteq \{0, 1\}^\ell$ is the set of rows and columns of the matrix $(\widehat{X})_{ij}$, which \mathcal{A} opens in its decommitment.

Proof. Assume toward contradiction that for any polynomial K , there exists a noticeable function ε , such that for infinitely many $\lambda \in \mathbb{N}$, with probability $\varepsilon(\lambda)$ over $\text{pk} \leftarrow \text{CSO.Gen}(1^\lambda)$ and $(\text{cmt}, \text{st}) \leftarrow \mathcal{A}_\lambda^1(\text{pk})$ for any set $\mathbf{S} \subseteq \{0, 1\}^\ell$ of size $K(\lambda)$

$$\Pr_{\substack{C \leftarrow \text{CSO.Chal}(1^\lambda) \\ (\alpha, I, d) \leftarrow \mathcal{A}_\lambda^2(C; \text{st})}} \left[\text{CSO.Ver}(\text{cmt}, \alpha, I, C, d) = 1 \wedge \widehat{X}(\alpha, I, d) \not\subseteq \mathbf{S} \right] \geq \varepsilon(\lambda) .$$

We consider two complementary cases:

1. For infinitely many λ , except with probability $\varepsilon/2$ over $\text{pk}, \text{cmt}, \text{st}$ there exists a set $\mathbf{S}' \subseteq \{0, 1\}^{\ell'}$ of size \sqrt{K} such that except with probability $\varepsilon/2$, \mathcal{A} never opens the hash value cmt to $S = ((\text{rowcmt}_i)_i, (\text{colcmt}_j)_j) \notin \mathbf{S}'$.
2. For infinitely many λ , with probability at least $\varepsilon/2$ over $\text{pk}, \text{cmt}, \text{st}$ for any set $\mathbf{S}' \subseteq \{0, 1\}^{\ell'}$ of size \sqrt{K} , with probability $\varepsilon/2$, \mathcal{A} opens the hash value cmt to $S = ((\text{rowcmt}_i)_i, (\text{colcmt}_j)_j) \notin \mathbf{S}'$.

First, note that the second case implies that \mathcal{A} breaks the multi-collision resistance of the underlying hash \mathbf{H} . Thus, we can assume that the first case holds. It follows that

Claim 3.5. *For infinitely many λ , with probability at least $\varepsilon/2$ over $\text{pk}, \text{cmt}, \text{st}$, there exists a set $\mathbf{S}' \subseteq \{0, 1\}^{\ell'}$ of size \sqrt{K} as required by the first condition, but for any set \mathbf{S} of size K and with probability $\varepsilon/2$ over the decommitment phase $\widehat{X}(\alpha, I, d) \not\subseteq \mathbf{S}$ whereas the opened $S = ((\text{rowcmt}_i)_i, (\text{colcmt}_j)_j) \in \mathbf{S}'$.*

Proof. This follows directly from our assumption toward contradiction and the fact that the first case above holds.

From hereon fix $\text{pk}, \text{cmt}, \text{st}$ such that Claim 3.5 holds. We next argue that

Claim 3.6. *There exists $S' = ((\text{rowcmt}_i)_i, (\text{colcmt}_j)_j) \in \mathbf{S}'$ and $\text{cmt}' \in S'$ such that for any set $\mathbf{X} \subseteq \{0, 1\}^\ell$ of size $\sqrt{K}/(L+n)$, with probability $\varepsilon/2(L+n)\sqrt{K}$ over the decommitment phase cmt is opened to S' , but cmt' is opened to $X \notin \mathbf{X}$.*

Proof. Otherwise, for each $S' \in \mathbf{S}'$ and $\text{cmt}' \in S'$, we can choose $\mathbf{X}(S', \text{cmt}')$ such that the above does not hold, and obtain

$$\mathbf{S} = \bigcup_{S' \in \mathbf{S}, \text{cmt}' \in S'} \mathbf{X}(S', \text{cmt}')$$

of size K , which violates Claim 3.5.

We now obtain an adversary $\mathcal{B} = \{\mathcal{B}_\lambda^1, \mathcal{B}_\lambda^2\}_\lambda$ that breaks the weak binding of the commitment SHC. $\mathcal{B}_\lambda^1(\text{pk}_1)$ first samples $\text{pk}_2 \leftarrow \text{H.Gen}(1^\lambda)$, sets $\text{pk} = (\text{pk}_1, \text{pk}_2)$, runs $\mathcal{A}_\lambda^1(\text{pk})$ and obtains a state st , it then simulates a random challenge $C \leftarrow C(1^\lambda)$ and runs $\mathcal{A}_\lambda^2(C; \text{st})$. It then takes the set of commitments $S' = (\text{rowcmt}_i)_{i \in L}, (\text{colcmt}_j)_{j \in [n]}$ and outputs a random commitment $\text{cmt}' \in S'$ along with the state st . $\mathcal{B}_\lambda^2(\text{st})$ samples a new random challenge C' and runs $\mathcal{A}_\lambda^2(C'; \text{st})$, and outputs whatever opening \mathcal{A}_λ^2 outputs for cmt' , or \perp if there is no such opening.

Claim 3.7. *\mathcal{B} breaks weak binding.*

Proof. By construction and Claims 3.5, 3.6, with probability at least $\frac{\varepsilon}{2} \cdot \frac{\varepsilon}{2(L+n)\sqrt{K}}$ over the commitment phase of \mathcal{B}_λ^1 , for any set $\mathbf{X} \subseteq \{0, 1\}^\ell$ of size $\sqrt{K}/(L+n)$, \mathcal{B}_λ^2 opens that commitment to a value $X \notin \mathbf{X}$ with probability at least $\varepsilon/2(L+n)\sqrt{K}$.

This complete the proof of Claim 3.4.

We now proceed to prove the binding property of the scheme. The following claim asserts that whenever all strings \hat{X} are consistent with some set \mathbf{S} , there exists a polynomial-size set of strings $\mathbf{X} \subseteq \{0, 1\}^L$ such that all openings are consistent with \mathbf{X} , which will conclude the proof.

Claim 3.8. *Let K be the polynomial given by Claim 3.4 and let τ be a constant such that $\tau > 2 \log_\lambda(2K^2)$. Fix $\lambda, \text{pk}, \text{cmt}, \text{st}$ such that Claim 3.4 holds with respect to a set \mathbf{S} . Then there exists $\mathbf{X} \subseteq \{0, 1\}^L$ of size $K' = (nK)^\tau$ such that*

$$\Pr_{\substack{C \leftarrow \text{CSO.Chal}(1^\lambda) \\ (\alpha, I, d) \leftarrow \mathcal{A}_\lambda^2(C; \text{st})}} [\text{CSO.Ver}(\text{cmt}, \alpha, I, C, d) = 1 \wedge \alpha \notin \mathbf{X}|_I] \leq \mu(\lambda) .$$

Proof. We first argue that random τ locations fix at most a single encoding in \mathbf{S} .

Claim 3.9. *Let $T \subseteq [n]$ be a set of τ indices chosen independently at random, then*

$$\Pr_T \left[\exists \widehat{S}_0, \widehat{S}_1 \in \mathbf{S} \text{ such that } \widehat{S}_0|_T = \widehat{S}_1|_T \mid \begin{array}{l} \widehat{S}_b \in \text{SSS.Enc}(S_b) \\ S_0 \neq S_1 \end{array} \right] \leq 1/2 .$$

We call any set T as above fixing.

Proof. The proof follows directly from the distance of Shamir encodings, the bound K on the size of \mathbf{S} , and the definition of τ . Specifically the above can be bounded by

$$|\mathbf{S}|^2 \left(\frac{\delta}{n} \right)^\tau \leq K^2 \left(\frac{\sqrt{\lambda}}{\lambda} \right)^{2 \log_\lambda(2K^2)} = \frac{1}{2} .$$

Now, let T be any fixing set and consider any set of τ columns

$$Q = \{(Q_{1,j}, \dots, Q_{L,j}) : j \in [T]\} .$$

Then, T, Q fix a unique string $X(T, Q) \in \{0, 1, \perp\}^L$ defined as follows:

- If T is not fixing, set $X(T, Q) = \perp$. Otherwise, proceed to the following.
- For any $i \in [L]$, if for some $b \in \{0, 1\}$, there exists a Shamir encoding $\widehat{X}_i \in \text{SSS.Enc}(b; 1^\lambda)$ such that $\widehat{X}_i \in \mathbf{S}$ and $\widehat{X}_{i,j} = Q_{i,j}$ for all $j \in T$, set $X_i(T, Q) = b$. Otherwise, set $X_i(T, Q) = \perp$.
(Since T is fixing, there exists at most a single $b \in \{0, 1\}$ that satisfies the condition.)

We now define the set \mathbf{X} as follows:

$$\mathbf{X} := \{X(T, Q) \mid T \in [n]^\tau, Q \in \mathbf{S}^\tau\} .$$

First, note that \mathbf{X} is of size at most $(n \cdot |\mathbf{S}|)^\tau = (nK)^\tau$, which is polynomial in λ . We now argue that

$$\Pr_{\substack{C \leftarrow \text{CSO.Chal}(1^\lambda) \\ (\alpha, I, d) \leftarrow \mathcal{A}_\lambda^2(C; \text{st})}} [\text{CSO.Ver}(\text{cmt}, \alpha, I, C, d) = 1 \wedge \alpha \notin \mathbf{X}|_I] \leq \mu(\lambda) .$$

Recall that $C \subseteq [n]$ consists of λ indices chosen independently at random. In particular, by Claim 3.9, C contains a fixing set T except with probability $(1/2)^{\delta/\tau} = 2^{-\tilde{\Omega}(\sqrt{\lambda})}$. Recall that except with negligible probability all row and column commitments opened by \mathcal{A} are consistent with some string in \mathbf{S} . From hereon we assume that the latter occurs and that C contains a fixing T .

Then when $\text{CSO.Ver}(\text{cmt}, \alpha, I, C, d) = 1$, \mathcal{A} opens the columns in T to some $Q = (Q_{1,j}, \dots, Q_{L,j})_{j \in [T]} \in \mathbf{S}^\tau$. Also the assignment $\alpha : I \rightarrow \{0, 1\}$ is such that for every $i \in I$, there exists an encoding $\widehat{\alpha}_i \in \text{SSS.Enc}(\alpha_i; 1^\lambda) \in \mathbf{S}$ and $\widehat{\alpha}$ is consistent with Q . By definition, it follows that for every $i \in I$, $\alpha_i = X_i(T, Q)$; namely, $\alpha \in \mathbf{X}|_I$ as required.

This completes the proof of Claim 3.8.

This completes the proof of Proposition 3.2

4 Offline-Online Statistically WI Arguments of Knowledge

In this section, we define and construct offline-online statistically witness indistinguishable arguments of knowledge.

4.1 Definition

In offline-online statistically WI arguments of knowledge, the protocol $\langle P \rightleftharpoons V \rangle$ can be divided to:

- an offline protocol $\langle \text{OffP} \rightleftharpoons \text{OffV} \rangle(1^\lambda, 1^\ell)$, where the parties take as common input the security parameter 1^λ and an input size 1^ℓ and output each a state st_P, st_V .
- an online protocol $\langle \text{OnP}(\text{st}_P, w) \rightleftharpoons \text{OnV}(\text{st}_V) \rangle(x)$, where the parties, in addition to their previous state, take as common input an instance $x \in \mathcal{L} \cap \{0, 1\}^\ell$ and the prover obtains also obtain as input a witness $w \in \mathcal{R}_\mathcal{L}(x)$.

We now formally define the properties that such systems are required to satisfy.

Definition 4.1 (Offline-Online SWIAOK). $\langle P \rightleftharpoons V \rangle$ is an offline-online statistically witness indistinguishable argument of knowledge for \mathcal{L} if it satisfies:

Completeness: For any $\ell, \lambda \in \mathbb{N}$, $x \in \mathcal{L} \cap \{0, 1\}^\ell$, and $w \in \mathcal{R}_\mathcal{L}(x)$:

$$\Pr \left[\begin{array}{l} \langle \text{OffP} \rightleftharpoons \text{OffV} \rangle(1^\lambda, 1^\ell) = (\text{st}_P, \text{st}_V) \\ \langle \text{OnP}(\text{st}_P, w) \rightleftharpoons \text{OnV}(\text{st}_V) \rangle(x) = 1 \end{array} \right] = 1 .$$

Adaptive statistical witness indistinguishability: For any polynomial $\ell(\cdot)$ and unbounded verifier V^* , there exists a negligible function $\mu(\cdot)$ such that for all $\lambda \in \mathbb{N}$:

$$\Pr \left[\begin{array}{l} \langle \text{OnP}(\text{st}_P, w_b) \rightleftharpoons \text{OnV}^*(\text{st}_V) \rangle(x) = b \mid \\ \langle \text{OffP} \rightleftharpoons \text{OffV}^* \rangle(1^\lambda, 1^{\ell(\lambda)}) = (\text{st}_P, (\text{st}_V, x, w_0, w_1)) \\ b \leftarrow \{0, 1\} \end{array} \right] \leq \frac{1}{2} + \mu(\lambda) ,$$

where $x \in \mathcal{L} \cap \{0, 1\}^{\ell(\lambda)}$ and $w_0, w_1 \in \mathcal{R}_\mathcal{L}(x)$.

Adaptive proof of knowledge: there is a uniform PPT extractor \mathcal{E} such that for any polynomial $\ell(\cdot)$ and any non-uniform polynomial-size prover $P^* = \{P_\lambda^*\}_{\lambda \in \mathbb{N}}$ there is a polynomial $K(\cdot)$ and a negligible $\mu(\cdot)$ such that for all $\lambda \in \mathbb{N}$:

$$\begin{array}{l}
\text{if} \\
\text{then}
\end{array}
\Pr \left[\begin{array}{l}
\langle \text{OnP}_\lambda^*(\text{st}_P) \Leftrightarrow \text{OnV}(\text{st}_V) \rangle(x) = 1 \mid \\
\langle \text{OffP}^* \Leftrightarrow \text{OffV} \rangle(1^\lambda, 1^{\ell(\lambda)}) = ((\text{st}_P, x), \text{st}_V)] = \varepsilon , \\
\Pr \left[\begin{array}{l}
\langle \text{OnP}_\lambda^*(\text{st}_P) \Leftrightarrow \text{OnV}(\text{st}_V) \rangle(x) = 1 \mid \\
w \leftarrow \mathcal{E}^{\text{P}^*}_\lambda(x, \text{st}_P, \text{st}_V) \\
w \in \mathcal{R}_\mathcal{L}(x)
\end{array} \right. \\
\left. \langle \text{OffP}^* \Leftrightarrow \text{OffV} \rangle(1^\lambda, 1^{\ell(\lambda)}) = ((\text{st}_P, x), \text{st}_V) \right] \\
\geq \text{poly} \left(\frac{\varepsilon}{K(\lambda)} \right) - \mu(\lambda) ,
\end{array}$$

where $x \in \{0, 1\}^{\ell(\lambda)}$.

Offline succinctness: All messages sent by OffP in the offline stage are of length λ (independently of ℓ).

In the remainder of this section we construct offline-online statistically witness indistinguishable arguments of knowledge.

Theorem 4.1. *Assuming a polynomially-compressing multi-collision-resistant keyless hash there exists an offline-online statistically witness indistinguishable argument of knowledge for NP with two messages in the offline part and one message in the online part.*

4.2 A Protocol for Hamiltonicity

We now give an offline-online protocol for the NP complete problem of Hamiltonicity. The protocol is essentially the Lapidot-Shamir protocol [LS90a] whereas instead of using standard (binding) commitments, we rely on the notion of weakly-binding commitments with subset opening from the previous section.

Ingredients and notation:

- Let $n(\cdot)$ be a polynomial and let $(\text{CSO.Gen}, \text{CSO.Com}, \text{CSO.Chal}, \text{CSO.Open})$ be a weakly binding statistically hiding commitment with subset opening for strings of length $n^2 \cdot \lambda$.
- A graph G with n nodes will be represented by its $n \times n$ adjacency matrix. Sometimes we may think of G as a string in the natural way.
- Let G, H be two graphs on the same set of nodes $[n]$ and let $\varphi : [n] \rightarrow [n]$ be a permutation. We write $H \subseteq G$ to denote the fact that H 's set of edges is contained in G 's set of edges. We write $\varphi(G) = H$ to denote the fact that $H_{\varphi(i), \varphi(j)} = G_{i, j}$ for every $i, j \in [n]$.
- A Hamiltonian cycle graph H is a graph that consists of a Hamiltonian cycle (and no additional edges).




Fig. 1: A 3-message SWI argument of knowledge for Hamiltonicity.

4.3 Analysis

The offline succinctness property follows directly from the succinct commitment property of the underlying commitment with subset opening. We focus on proving the argument of knowledge and the statistical witness indistinguishability properties.

Proposition 4.1. *Protocol 1 is an adaptive argument of knowledge.*

Proof. Fix a non-uniform polynomial-size prover $P^* = \{P_\lambda^*\}_{\lambda \in \mathbb{N}}$ such that

$$\Pr [\langle \text{On}P_\lambda^*(\text{st}_P) \rightrightarrows \text{On}V(\text{st}_V) \rangle(x) = 1 \mid \langle \text{Off}P^* \rightrightarrows \text{Off}V \rangle(1^\lambda, 1^\ell) = ((\text{st}_P, x), \text{st}_V)] = \varepsilon .$$

We now describe how the witness extractor $\mathcal{E}(P_\lambda^*, x, \text{st}_P, \text{st}_V)$ operates. \mathcal{E} first emulates the last prover message corresponding to the state of the offline phase it is given. It obtains $(\mathbf{I}, \mathbf{H}|_{\mathbf{I}}, \varphi)$. \mathcal{E} then rewinds the prover P_λ^* back to the offline phase, and sends it a fresh random challenge σ', C' . It then obtains in the online phase corresponding $(\mathbf{I}', \mathbf{H}'|_{\mathbf{I}'}, \varphi')$. \mathcal{E} now looks for an $i \in [\lambda]$ such that $\sigma_i = \text{antiedges}$ and $\sigma'_i = \text{cycle}$, and returns the cycle $\varphi_i^{-1}(H'_i)$. If any of the above fail, it aborts.

Claim 4.2. *There exists a polynomial $K(\cdot)$ and a negligible $\mu(\cdot)$ such that*

$$\Pr \left[\begin{array}{l} \langle \text{On}P_\lambda^*(\text{st}_P) \rightrightarrows \text{On}V(\text{st}_V) \rangle(x) = 1 \\ w \leftarrow \mathcal{E}(P_\lambda^*, x, \text{st}_P, \text{st}_V) \\ w \in \mathcal{R}_{\mathcal{L}}(x) \end{array} \mid \langle \text{Off}P^* \rightrightarrows \text{Off}V \rangle(1^\lambda, 1^{\ell(\lambda)}) = ((\text{st}_P, x), \text{st}_V) \right] \geq \frac{\varepsilon^3}{8K^2(\lambda)} - \mu(\lambda) .$$

Proof. First, by an averaging argument with probability at least $\varepsilon/2$ over the choice of the first two messages in the protocol, namely the key pk and commitment cmt , it holds that with probability at least $\varepsilon/2$ over the rest of the protocol the prover convinces the verifier of accepting. In addition, by the weak binding property of the underlying commitment with subset opening, there exists a polynomial K , such that except with negligible probability $\nu(\lambda)$ over the first two messages, there exists a set \mathbf{X} of size at most K , consisting of strings $\mathbf{H} \in \{0,1\}^{n^2 \times \lambda}$, such that except with negligible probability $\nu(\lambda)$, any valid opening of the commitment by the prover is consistent with some $\mathbf{H} \in \mathbf{X}$. Also, note that for two random challenges σ, σ' there exists some i such that $\sigma_i = \text{antiedges}$ and $\sigma'_i = \text{cycle}$ except with probability $(3/4)^\lambda$.

It follows that with probability at least $\varepsilon/2 - \nu$ over the choice of the first two messages, there exists a single string $\mathbf{H}^* \in \mathbf{X}$, such that with probability at least $(\varepsilon/2K)^2 - (3/4)^\lambda$, in both executions performed by \mathcal{E} :

1. The prover convinces the verifier.
2. Both openings $\mathbf{H}|_{\mathbf{I}}$ and $\mathbf{H}'|_{\mathbf{I}'}$ obtained by \mathcal{E} are consistent with \mathbf{H} .
3. For some i , $\sigma_i = \text{antiedges}$ and $\sigma'_i = \text{cycle}$.

Since the verifier accepts in the second execution, H_i is a Hamiltonian cycle graph. Since the verifier accepts in the first execution, and H_i is a Hamiltonian cycle, $\varphi_i^{-1}(H_i)$ is a Hamiltonian cycle in G , since all anti-edges in G are mapped to anti-edges in H_i .

Overall, the extractor succeeds with probability at least $\varepsilon^3/8K^2 - \lambda^{-\omega(1)}$.

Proposition 4.2. *Protocol 1 satisfies adaptive statistical witness indistinguishability.*

The proof of witness indistinguishability is similar to that of the original Lapidot-Shamir protocol [LS90b]. The main difference is that there it is convenient to first prove WI for a single instance (with a single challenge σ) and then rely on a generic hybrid argument, whereas in our protocol the commitment with subset opening correlates the λ instances. The proof accordingly proceeds via a slightly less generic hybrid argument. The proof the the proposition can be found in the full version of this work.

5 A Three Message Statistical Zero Knowledge Argument

In this section, we construct a three message statistical zero knowledge argument.

Theorem 5.1. *Assuming a linearly compressing quasipoly(λ)-secure multi-collision-resistant keyless hash, a quasipoly(λ)-secure fully-homomorphic encryption, and a two-message secure function evaluation protocol with statistical function hiding, there exists a statistical zero knowledge argument for **NP** with three messages.*

Ingredients and notation:

- A two-message weak memory delegation scheme (MD.Mem, MD.Query, MD.Prove, MD.Ver) with weak soundness for time- T for $T = \text{quasipoly}(\lambda)$, as in Definition 2.5.
- A two-message secure function evaluation protocol (SFE.Enc, SFE.Eval, SFE.Dec) with statistical function hiding as in Definition 2.6.
- An offline-online statistically witness indistinguishable argument of knowledge $\langle P \leftrightarrow V \rangle$ where the offline part of the protocol $\langle \text{OffP} \leftrightarrow \text{OffV} \rangle$ consists of two messages and the online part of the protocol $\langle \text{OnP} \leftrightarrow \text{OnV} \rangle$ consists of one message. Such a protocol is defined and constructed in Section 4.
- A weakly-binding statistically-hiding keyless commitment SHC.Com as in Definition 2.2.
- For a string x , denote by \mathcal{M}_x a Turing machine that given memory $D = V^*$, emulates the Turing machine encoded by V^* on the input x , parses the result as $(u, w_2, q, \widehat{ct}_\tau)$, and outputs u .
- Denote by $\mathcal{V}_{\text{param}}$ a circuit that has the string **param** hardcoded and operates as follows. Given as input a secret state **vst** for the delegation scheme:

- parse $\text{param} = (x, \mathbf{q}, u, \text{dig}, t, \pi)$,
- return 1 (“accept”) if either of the following occurs:
 - * the delegation verifier accepts: $\text{MD.Ver}(1^\lambda, \text{dig}, (\mathcal{M}_x, t, u), \text{vst}, \pi) = 1$,
 - * the query and secret state are inconsistent: $(\mathbf{q}, \text{vst}) \notin \text{MD.Query}(1^\lambda)$.
(We can assume without loss of generality that the state vst contains the random coins of MD.Query and, therefore, consistency can be tested efficiently.)

In words, $\mathcal{V}_{\text{param}}$, given the secret state vst , first verifies the proof π that “ $\mathcal{M}_x(D) = (u, \dots)$ ” where D is the database corresponding to the digest dig . In addition, it verifies that \mathbf{q} is truly consistent with the coins vst .

- Denote by $\mathbf{1}$ a circuit of the same size as $\mathcal{V}_{\text{param}}$ that always returns 1.

We describe our three-message zero-knowledge protocol in Figure 2.

Proposition 5.1. *Protocol 2 is a statistical zero-knowledge argument.*

Proof (Sketch). As explained in the introduction, the protocol is based on the zero-knowledge protocols in [BBK⁺16, BKP18]. In particular, Bitansky et al. [BBK⁺16] show a three-message computational zero knowledge protocol in the *global hash model*, where parties have access to a collision-resistant hash function sampled during a setup phase. The main differences between our protocol and theirs is:

- Their two-message memory delegation scheme has full soundness instead of weak soundness.
- Their secure function evaluation has computational function-hiding instead of statistical.
- Their offline-online argument of knowledge is computationally witness indistinguishable instead of statistically.
- The non-interactive commitment is perfectly binding and computationally hiding instead of weakly binding and statistically hiding.

Next we outline the analysis of [BBK⁺16] and explain how to modify it for our protocol.

Soundness. Assuming that $x \notin \mathcal{L}$, in order to pass the witness indistinguishable argument of knowledge with respect to an evaluated cipher $\widehat{\text{ct}}$ that decrypts to 1, the prover must know a proof $\pi \in \{0, 1\}^\lambda$ and an opening of cmt to a digest $\text{dig} \in \{0, 1\}^\lambda$ and a time bound $t \leq T(\lambda)$ such that $\mathcal{V}_{\text{param}}(\text{vst}) = 1$. This, by definition, means that (dig, π, t) are such that the delegation verifier MD.Ver is convinced that the digest dig corresponds to a machine V^* such that $V^*(\text{wi}_1, \text{cmt}) = (u, \dots)$.

By the weak binding of cmt , the prover can only open the commitment to a polynomial number of different digests. Therefore, there must exist one digest dig for for which the prover can convince delegation verifier MD.Ver with high probability for an output u with high entropy, contradicting the weak soundness of the delegation scheme. In order to break the underlying delegation scheme we also rely on the semantic security of the encryption scheme to hide the secret verification state vst from the prover.

Protocol 2

Common Input: an instance $x \in \mathcal{L} \cap \{0, 1\}^\lambda$, for security parameter λ .
P: a witness $w \in \mathcal{R}_{\mathcal{L}}(x)$.

1. P computes
 - w_{i_1} , the first message of the offline prover $\text{OffP}(1^\lambda, 1^{\ell_\Psi(\lambda)})$ where ℓ_Ψ is the length of the statement Ψ defined in Step 3 below,
 - $\text{cmt} \leftarrow \text{SHC.Com}(0^{2\lambda}; 1^\lambda)$, a commitment to the all zero string,
 and sends (w_{i_1}, cmt) .
2. V computes
 - st_V and w_{i_2} , the state and the second message of the offline verifier $\text{OffV}(1^\lambda, 1^{\ell_\Psi(\lambda)})$ after receiving the message w_{i_1} .
 - $(q, \text{vst}) \leftarrow \text{MD.Query}(1^\lambda)$, a query and secret state for the delegation scheme,
 - $(\text{ct}_{\text{vst}}, \text{sk}) \leftarrow \text{SFE.Enc}(1^\lambda, \text{vst})$, an encryption of the secret state,
 - $u \leftarrow \{0, 1\}^\lambda$, a uniformly random string,
 and sends $(u, w_{i_2}, q, \text{ct}_{\text{vst}})$.
3. P computes
 - $\widehat{\text{ct}} \leftarrow \text{SFE.Eval}(\mathbf{1}, \text{ct}_{\text{vst}})$, an evaluation of the constant one function,
 - st_P , the state of the offline prover $\text{OffP}(1^\lambda, 1^{\ell_\Psi(\lambda)})$ after receiving the message w_{i_2} .
 - w_{i_3} , the message for online prover OnP given the state st_P , the statement $\Psi = \Psi_1(x) \vee \Psi_2(w_{i_1}, \text{cmt}, q, u, \text{ct}_{\text{vst}}, \widehat{\text{ct}})$ of length $\ell_\Psi(\lambda)$ given by:

$$\left\{ \exists w \mid (x, w) \in \mathcal{R}_{\mathcal{L}} \right\} \vee \left\{ \exists \begin{array}{l} \text{dig}, \pi, r_{\text{cmt}}, r_{\widehat{\text{ct}}} \in \{0, 1\}^{\text{poly}(\lambda)} \\ t \leq T(\lambda) \end{array} \mid \begin{array}{l} \text{cmt} = \text{Com}(\text{dig}, t; r_{\text{cmt}}) \\ \text{param} = ((w_{i_1}, \text{cmt}), q, u, \text{dig}, t, \pi) \\ \widehat{\text{ct}} = \text{SFE.Eval}(\mathcal{V}_{\text{param}}, \text{ct}_{\text{vst}}, r_{\widehat{\text{ct}}}) \end{array} \right\},$$

- and the witness $w \in \mathcal{R}_{\mathcal{L}}(x)$ for Ψ_1 ,
 and sends $(\widehat{\text{ct}}, w_{i_3})$.
4. V verifies that $\text{SFE.Dec}(\widehat{\text{ct}}; \text{sk}) = 1$ and that the online verifier OnP with state st_V and the statement Ψ accepts after receiving the message w_{i_3} .

Fig. 2: A three-message statistical ZK argument of knowledge for **NP**.

Statistical zero knowledge. To show statistical zero knowledge, we construct a non-black-box simulator following the simulator of Barak [Bar01]. At high-level, the simulator uses the code of the (malicious) verifier V^* as the memory for the delegation scheme, and completes the witness indistinguishable argument of knowledge using a witness for the trapdoor statement Ψ_2 . The witness consists of (dig, π, t) where dig is the digest corresponding to V^* , $t \approx |V^*|$ and π is the corresponding delegation proof that $V^*(w_1, \text{cmt}) = (u, \dots)$, which is now true by definition.

By the perfect completeness of the delegation scheme, we know that for any encrypted secret state vst , given a query \mathbf{q} that is consistent with vst , the delegation verifier MD.Ver will accept the corresponding proof. Thus, the perfect function hiding of the secure function evaluation (which holds also if the verifier produces a malformed ciphertext) guarantees that the evaluated ciphertext $\hat{\text{ct}}$ in the simulated proof is statistically close to that computed in the real proof where the prover actually evaluates the constant $\mathbf{1}$ circuit.

Relying also on the statistical witness indistinguishability of the argument of knowledge and the statistical hiding of cmt we deduce that V^* 's view in the real proof and the simulated view are statistically close.

References

- [AH91] William Aiello and Johan Håstad. Statistical zero-knowledge languages can be recognized in two rounds. *J. Comput. Syst. Sci.*, 42(3):327–345, 1991.
- [Bar01] Boaz Barak. How to go beyond the black-box simulation barrier. In *42nd Annual Symposium on Foundations of Computer Science, FOCS 2001, 14-17 October 2001, Las Vegas, Nevada, USA*, pages 106–115, 2001.
- [BBK⁺16] Nir Bitansky, Zvika Brakerski, Yael Tauman Kalai, Omer Paneth, and Vinod Vaikuntanathan. 3-message zero knowledge against human ignorance. In *Theory of Cryptography - 14th International Conference, TCC 2016-B, Beijing, China, October 31 - November 3, 2016, Proceedings, Part I*, pages 57–83, 2016.
- [BCC88] Gilles Brassard, David Chaum, and Claude Crépeau. Minimum disclosure proofs of knowledge. *J. Comput. Syst. Sci.*, 37(2):156–189, 1988.
- [BCC⁺14] Nir Bitansky, Ran Canetti, Alessandro Chiesa, Shafi Goldwasser, Huijia Lin, Aviad Rubinfeld, and Eran Tromer. The hunting of the SNARK. *IACR Cryptology ePrint Archive*, 2014:580, 2014.
- [BCPR14] Nir Bitansky, Ran Canetti, Omer Paneth, and Alon Rosen. On the existence of extractable one-way functions. In *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 505–514, 2014.
- [BCY91] Gilles Brassard, Claude Crépeau, and Moti Yung. Constant-round perfect zero-knowledge computationally convincing protocols. *Theor. Comput. Sci.*, 84(1):23–52, 1991.
- [BD18] Zvika Brakerski and Nico Döttling. Two-message statistically sender-private OT from LWE. In *Theory of Cryptography - 16th International Conference, TCC 2018, Panaji, India, November 11-14, 2018, Proceedings, Part II*, pages 370–390, 2018.

- [BDRV18] Itay Berman, Akshay Degwekar, Ron D. Rothblum, and Prashant Nalini Vasudevan. Multi-collision resistant hash functions and their applications. In *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part II*, pages 133–161, 2018.
- [BHKY19] Nir Bitansky, Iftach Hainer, Ilan Komargodski, and Eylon Yogev. Distributional collision resistance beyond one-way functions. In *Advances in Cryptology - EUROCRYPT 2019 - 38th Annual International Conference on the Theory and Applications of Cryptographic Techniques*, 2019.
- [BJY97] Mihir Bellare, Markus Jakobsson, and Moti Yung. Round-optimal zero-knowledge arguments based on any one-way function. In *Advances in Cryptology - EUROCRYPT '97, International Conference on the Theory and Application of Cryptographic Techniques, Konstanz, Germany, May 11-15, 1997, Proceeding*, pages 280–305, 1997.
- [BKP18] Nir Bitansky, Yael Tauman Kalai, and Omer Paneth. Multi-collision resistance: a paradigm for keyless hash functions. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 671–684, 2018.
- [Blu86] Manuel Blum. How to prove a theorem so no one else can claim it. 1986.
- [BOV07] Boaz Barak, Shien Jin Ong, and Salil P. Vadhan. Derandomization in cryptography. *SIAM J. Comput.*, 37(2):380–400, 2007.
- [BP04] Mihir Bellare and Adriana Palacio. The knowledge-of-exponent assumptions and 3-round zero-knowledge protocols. In *Proceedings of the 24th Annual International Cryptology Conference*, pages 273–289, 2004.
- [BV11] Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011, Palm Springs, CA, USA, October 22-25, 2011*, pages 97–106, 2011.
- [CD09] Ran Canetti and Ronny Ramzi Dakdouk. Towards a theory of extractable functions. In *TCC*, pages 595–613, 2009.
- [CKLR11] Kai-Min Chung, Yael Tauman Kalai, Feng-Hao Liu, and Ran Raz. Memory delegation. In *Advances in Cryptology - CRYPTO 2011 - 31st Annual Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2011. Proceedings*, pages 151–168, 2011.
- [DN07] Cynthia Dwork and Moni Naor. Zaps and their applications. *SIAM J. Comput.*, 36(6):1513–1543, 2007.
- [DPP93] Ivan Damgård, Torben P. Pedersen, and Birgit Pfitzmann. On the existence of statistically hiding bit commitment schemes and fail-stop signatures. In *Advances in Cryptology - CRYPTO '93, 13th Annual International Cryptology Conference, Santa Barbara, California, USA, August 22-26, 1993, Proceedings*, pages 250–265, 1993.
- [FGJ18] Nils Fleischhacker, Vipul Goyal, and Abhishek Jain. On the existence of three round zero-knowledge proofs. *IACR Cryptology ePrint Archive*, 2018:167, 2018.
- [FLS99] Uriel Feige, Dror Lapidot, and Adi Shamir. Multiple noninteractive zero knowledge proofs under general assumptions. *SIAM J. Comput.*, 29(1):1–28, 1999.
- [For89] Lance Fortnow. The complexity of perfect zero-knowledge. *Advances in Computing Research*, 5:327–343, 1989.

- [GK96] Oded Goldreich and Ariel Kahan. How to construct constant-round zero-knowledge proof systems for NP. *J. Cryptology*, 9(3):167–190, 1996.
- [GMR89] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM J. Comput.*, 18(1):186–208, 1989.
- [GMW91] Oded Goldreich, Silvio Micali, and Avi Wigderson. Proofs that yield nothing but their validity for all languages in NP have zero-knowledge proof systems. *J. ACM*, 38(3):691–729, 1991.
- [GO94] Oded Goldreich and Yair Oren. Definitions and properties of zero-knowledge proof systems. *Journal of Cryptology*, 7(1):1–32, December 1994.
- [GOS12] Jens Groth, Rafail Ostrovsky, and Amit Sahai. New techniques for non-interactive zero-knowledge. *J. ACM*, 59(3):11, 2012.
- [HM96] Shai Halevi and Silvio Micali. Practical and provably-secure commitment schemes from collision-free hashing. In *Advances in Cryptology - CRYPTO '96, 16th Annual International Cryptology Conference, Santa Barbara, California, USA, August 18-22, 1996, Proceedings*, pages 201–215, 1996.
- [HNO⁺09] Iftach Haitner, Minh-Huyen Nguyen, Shien Jin Ong, Omer Reingold, and Salil P. Vadhan. Statistically hiding commitments and statistical zero-knowledge arguments from any one-way function. *SIAM J. Comput.*, 39(3):1153–1218, 2009.
- [HT98] Satoshi Hada and Toshiaki Tanaka. On the existence of 3-round zero-knowledge protocols. In *Proceedings of the 18th Annual International Cryptology Conference*, pages 408–423, 1998.
- [Kat12] Jonathan Katz. Which languages have 4-round zero-knowledge proofs? *J. Cryptology*, 25(1):41–56, 2012.
- [KKS18] Yael Tauman Kalai, Dakshita Khurana, and Amit Sahai. Statistical witness indistinguishability (and more) in two messages. In *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part III*, pages 34–65, 2018.
- [KNY17] Ilan Komargodski, Moni Naor, and Eylon Yogev. White-box vs. black-box complexity of search problems: Ramsey and graph property testing. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 622–632, 2017.
- [KNY18] Ilan Komargodski, Moni Naor, and Eylon Yogev. Collision resistant hashing for paranoids: Dealing with multiple collisions. In *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part II*, pages 162–194, 2018.
- [KP16] Yael Tauman Kalai and Omer Paneth. Delegating RAM computations. In *Theory of Cryptography - 14th International Conference, TCC 2016-B, Beijing, China, October 31 - November 3, 2016, Proceedings, Part II*, pages 91–118, 2016.
- [LS90a] Dror Lapidot and Adi Shamir. Publicly verifiable non-interactive zero-knowledge proofs. In *CRYPTO*, pages 353–365, 1990.
- [LS90b] Dror Lapidot and Adi Shamir. Publicly verifiable non-interactive zero-knowledge proofs. In *Advances in Cryptology - CRYPTO '90, 10th Annual*

- International Cryptology Conference, Santa Barbara, California, USA, August 11-15, 1990, Proceedings*, pages 353–365, 1990.
- [Nao91] Moni Naor. Bit commitment using pseudorandomness. *J. Cryptology*, 4(2):151–158, 1991.
- [NOVY98] Moni Naor, Rafail Ostrovsky, Ramarathnam Venkatesan, and Moti Yung. Perfect zero-knowledge arguments for *NP* using any one-way permutation. *J. Cryptology*, 11(2):87–108, 1998.
- [OPP14] Rafail Ostrovsky, Anat Paskin-Cherniavsky, and Beni Paskin-Cherniavsky. Maliciously circuit-private FHE. In *Advances in Cryptology - CRYPTO 2014 - 34th Annual Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2014, Proceedings, Part I*, pages 536–553, 2014.
- [Sha79] Adi Shamir. How to share a secret. *Commun. ACM*, 22(11):612–613, 1979.