



NIZK from LPN and Trapdoor Hash via Correlation Intractability for Approximable Relations

Zvika Brakerski^(✉), Venkata Koppula, and Tamer Mour

Weizmann Institute of Science, Rehovot, Israel
{zvika.brakerski, venkata.koppula, tamer.mour}@weizmann.ac.il

Abstract. We present new non-interactive zero-knowledge argument systems (NIZK), based on standard assumptions that were previously not known to imply it. In particular, we rely on the hardness of both the learning parity with noise (LPN) assumption, and the existence of *trapdoor hash functions* (TDH, defined by Döttling et al., Crypto 2019). Such TDH can be based on a number of standard assumptions, including DDH, QR, DCR, and LWE.

We revisit the correlation intractability (CI) framework for converting Σ -protocols into NIZK, and present a different strategy for instantiating it by putting together two new components. First, while prior works considered the search-complexity of the relations for which CI is sought, we consider their *probabilistic representation*. Namely, a distribution over lower-complexity functions that bitwise-computes the target function with all but small (constant) probability. The second component is a new perspective for quantifying the class of relations for which CI is achieved. We show that it is instructive to consider *CI for approximable relations* (CI-Apx) which is quantified by a class of relations, but requires CI to hold against *any approximation* of any relation in this class.

We show that CI-Apx for just constant-degree polynomials suffices for NIZK, if the underlying Σ -protocol is implemented using a suitable commitment scheme. We show that such a commitment scheme can be constructed based on low noise LPN. We then show how to construct CI-Apx for constant-degree polynomials from any suitable TDH (with an enhanced correctness property that is satisfied by all existing TDH constructions).

1 Introduction

Zero-Knowledge (ZK) [17] is one of the most celebrated and widely used notions in modern cryptography. A ZK proof is a protocol in which a prover conveys the validity of a statement to a verifier in a way that reveals no additional information. In a non-interactive ZK proof system (NIZK), we wish to construct a single-message ZK proof system. Common setup is necessary for NIZK, and by default

A full version can be found at <https://eprint.iacr.org/2020/258>.

© International Association for Cryptologic Research 2020
D. Micciancio and T. Ristenpart (Eds.): CRYPTO 2020, LNCS 12172, pp. 738–767, 2020.
https://doi.org/10.1007/978-3-030-56877-1_26

(and always in this work) NIZK is considered in the common random/reference string (CRS) model. In the CRS model, a trusted string is sampled from a prescribed distribution (preferably uniform) and made available to both the prover and the verifier. Ideally, we would have liked to construct a NIZK proof system for all NP languages (or equivalently for some NP-complete language).¹ NIZK for NP turns out to be extremely useful for many applications such as CCA security [13, 23], signatures [3, 5], and numerous other applications, including recent applications in the regime of cryptocurrencies [4]. From this point and on, we use the term NIZK to refer to “NIZK for NP” unless otherwise stated.

While ZK proofs for all NP languages are known under the minimal assumption that one-way functions exist, this is far from being the case for NIZK. We focus our attention on constructions in the standard model and under standard cryptographic assumptions. For many years, NIZK under standard assumptions were only known based on Factoring [7] (or doubly enhanced trapdoor functions, which are only known to exist based on Factoring [15]) or assumptions on groups with bilinear maps [18].

More recently, constructions based on indistinguishability obfuscation were presented as well [25]. Most recently, a new line of works, starting with [10, 19, 21], focused on obtaining NIZK based on the notion of *correlation intractability* (CI) [11]. In the CI framework, it was shown that in order to construct NIZK, it suffices to construct a family of hash functions \mathcal{H} with the following property. For every efficient f , given a hash function $H \leftarrow \mathcal{H}$ from the family, it is computationally hard to find x s.t. $f(x) = H(x)$. If such correlation intractable hash (CIH) is constructed, then it can be used to securely instantiate the Fiat-Shamir paradigm [16] and derive NIZK from so-called Σ -protocols. This line of works culminated in two remarkable achievements. Canetti et al. [9] constructed NIZK based on the existence of circular secure fully homomorphic encryption. Peikert and Shiehian [24] constructed NIZK based on the hardness of the learning with errors (LWE) problem.²

These recent results opened a new avenue in the study of NIZK and raised hope that construction under additional assumptions can be presented. However, it appears that there is an inherent barrier to expanding known techniques beyond LWE-related assumptions. The current approaches for constructing CI hash from standard assumptions use the notion of *somewhere statistical CI*, in which, for any f , it is possible to sample from a distribution \mathcal{H}_f which is indistinguishable from the real \mathcal{H} , and for which the CI game is *statistically* hard to win. Roughly speaking, this is achieved, in known constructions [9, 24] by making \mathcal{H}_f perform some homomorphic evaluation of f on the input x . Thus, it appears that homomorphic evaluation of complex functions f is essential to apply these tools.

¹ In this work we only consider ZK/NIZK proof systems where the honest prover is computationally efficient given a witness to the NP language.

² To be more accurate, [24] showed how to construct a CI hash function for size s circuits, for any parameter s . This is slightly weaker than a single \mathcal{H} for all functions, but it suffices in order to instantiate the framework.

The starting point of our work is the observation that, under the learning parity with noise (LPN) assumption, we can reduce the complexity of functions for which achieving CIH implies NIZK down to functions with *probabilistic constant-degree representation*. That is, ones that can be approximated by a distribution on constant-degree polynomials.

We substantiate the usefulness of this approach by identifying a general connection between correlation intractability for a function class \mathcal{F} , which has probabilistic representation by a class \mathcal{C} (potentially of lower complexity), and CI for relations that are *approximable* by \mathcal{C} .

Correlation Intractability for relations approximable by \mathcal{C} (denoted “CI-Apx for \mathcal{C} ”) is a stronger notion than the one studied in prior works, namely CI for relations searchable by \mathcal{C} . In CI-Apx, we require that for all $C \in \mathcal{C}$ it is hard not only to find x such that $C(x) = \mathcal{H}(x)$ but, rather, that it is hard to find an x such that $\mathcal{H}(x)$ and $C(x)$ are *close* in Hamming distance.³ When the probabilistic representation \mathcal{C} of our target class \mathcal{F} is sufficiently simple, e.g. constant-degree polynomials, then the reduction from CI for \mathcal{F} to CI-Apx for \mathcal{C} opens the possibility for new constructions of CIH from standard assumptions. Specifically from assumptions that are not known to apply fully-homomorphic encryption or similarly strong primitives.

In particular, we show that CI-Apx for a function class \mathcal{C} can be constructed based on a *rate-1 trapdoor hash* scheme for \mathcal{C} . Trapdoor hash (TDH) is a fairly new cryptographic primitive which was recently introduced by Döttling et al. [14]. They also constructed rate-1 TDH for constant-degree polynomials from a number of standard assumptions, including DDH, QR, and DCR (which are not known to imply fully-homomorphic encryption) and also LWE. Consequently, we obtain CI-Apx for constant-degree polynomials from such assumptions and, therefore, CI for any class of functions with probabilistic constant-degree representation. We note that we require a slightly stronger correctness property from TDH, compared to the definition provided in [14], but it is satisfied by all known constructions.

On an interesting remark, we point out that the construction by Piekert and Shiehian [24] of CI for bounded-size circuits can be shown to satisfy the stronger notion of CI-Apx for the corresponding class of relations.

Consequences. We get non-interactive (computational) zero knowledge argument systems for NP, in the common random string model, based on the existence of any rate-1 trapdoor hash for constant degree and further assuming low-noise LPN. We stress that we can generically abstract the LPN requirement as a requirement for an extractable commitment scheme with very low-complexity approximate-extraction. By instantiating our construction using the rate-1 TDH from [14], we get, in particular, the first NIZK from low-noise LPN and DDH.

³ Note that even non-searchable relations can potentially be approximable by a class of functions. Thus via the notion CI-Apx we can extend our capabilities for constructing CIH even beyond searchable relations. This is not of direct use in this work, but may be useful for future works.

Open Questions. The main open question we leave unanswered is whether it is possible to minimize the required assumptions for constructing NIZK using CI-Apx. One may approach this problem either by constructing CI-Apx for constant degree functions based on the LPN assumption, or by further extending the CI-Apx framework to allow a more general utilization for NIZKs, possibly depending on assumptions already implying CI-Apx.

Another open question is whether we can obtain stronger notions of NIZKs, in particular NIZK proofs or NISZK, from a similar set of standard assumptions. To achieve statistical ZK using our approach simply requires the underlying commitment (with low-degree extraction) to be lossy. Getting statistically sound proof systems via CI-Apx, however, seems to be inherently more difficult, as it requires the resulting CI to be “somewhere statistical” for the approximated class of functions.

Lastly, the new constructions of ZAPs [2, 20, 22] rely on the CI framework but, unfortunately, we do not know how to extend them since the notion of commitment that is required for the ZAPs is not known to be constructible from LPN (or other assumptions with very low complexity extraction). At a high level, these works requires the public parameters of the commitment scheme to be statistically close to uniform (and this seems hard to achieve with our LPN noise regime).

1.1 Overview of Our Techniques and Results

Our construction of NIZK instantiates the general Correlation Intractability (CI) framework. The approach followed in prior work for constructing CI hash, for relations searchable by a function class \mathcal{F} , considers the straight-forward representation of \mathcal{F} as a class of circuits. In this work, we take a different angle, and tackle the CI problem for relations searchable by \mathcal{F} through its probabilistic representation by a much simpler class \mathcal{C} . Such an approach allows us to obtain CI hash for classes of relations that are sufficiently rich to imply NIZK, while avoiding the use of FHE or similar heavy machinery.

NIZK from Correlation Intractability. Our starting point for constructing NIZK is similar to the approach in previous works of applying Fiat-Shamir on ZK protocols, in a provably-sound manner, using CI hash. We start with a public-coin trapdoor Σ -protocol that follows the natural “commit-then-open” paradigm, where the prover first sends a set of commitments, then, upon receiving the verifier’s challenge bit $e \in \{0, 1\}$, he replies by opening some of the commitments. Lastly, the verifier checks that the openings are valid, and then performs an additional check over the opened values. An example of such a protocol is the ZK protocol for Hamiltonicity from [6, 15].

An important property of commit-then-open trapdoor- Σ protocols is the *unique bad challenge property*: for any instance x not in the language, if (a, e, z) is an accepting transcript, then e is uniquely determined by the first message a . This connection is characterized by a function denoted by $\text{BadChallenge} : a \mapsto e$.

In the CI paradigm, we apply Fiat-Shamir over sufficiently many repetitions of such a protocol, using a CI hash for the relation searchable by `BadChallenge`, which is defined as follows. A vector of first messages \mathbf{a} is in a relation with a vector of verifier’s challenges \mathbf{e} if on each coordinate, the corresponding \mathbf{e} entry is the unique bad challenge of that coordinate in \mathbf{a} . If a cheating prover \mathbf{P}^* succeeds in breaking the soundness of the protocol, then he must have found a `BadChallenge` correlation, i.e. vectors (\mathbf{a}, \mathbf{e}) in the relation, implying an attack against the CI of the underlying hash family.

Prior work considered protocols where the bad challenge is efficiently computable and, consequently, focused on constructing CI for all efficiently searchable relations. These contain, in particular, the relations efficiently searchable by `BadChallenge`. We deviate from this approach. We observe that `BadChallenge` can be approximated by a distribution over constant-degree polynomials when instantiating this template with an appropriate commitment scheme. This reduces our CI task to achieving CIH for functions with constant-degree *probabilistic representation*. Such CIH is implied by a special notion of correlation intractability against constant-degree functions – CI for *approximable relations*, or CI-Apx for short. Details follow.

Probabilistic Representation, Approximable Relations and CI. Assume that a class of functions $\mathcal{F} : \{0, 1\}^n \rightarrow \{0, 1\}^m$ has a *probabilistic representation* by some simpler class of functions \mathcal{C} . Namely, for any $f \in \mathcal{F}$, there exists a distribution \mathcal{C}_f over \mathcal{C} such that $\Pr[\Delta(C(x), f(x)) \leq \epsilon m] > 1 - \text{negl}(\lambda)$ for any x and a random $C \xleftarrow{\$} \mathcal{C}_f$.

Let $\mathcal{H} : \{0, 1\}^n \rightarrow \{0, 1\}^m$ be a hash family. An adversary \mathcal{A} that is able to find a correlation $\mathcal{H}(x) = f(x)$ for some f is able to find, with overwhelming probability over a random $C \leftarrow \mathcal{C}_f$, an “approximate correlation” $\Delta(\mathcal{H}(x), C(x)) \leq \epsilon m$ for some small ϵ . It follows therefore that by considering probabilistic representation, we can identify a connection between correlation intractability against f and correlation intractability against any relation that is *approximable* (or approximately searchable) by some function $C \in \mathcal{C}_f$. We denote this class of relations

$$R_C^\epsilon = \{(x, y) \in \{0, 1\}^n \times \{0, 1\}^m \mid \Delta(y, C(x)) \leq \epsilon m\} .$$

More formally, an adversary that breaks the CI of \mathcal{H} for a relation searchable by f is able to break the CI of the same hash \mathcal{H} for the relation \mathcal{R}_C^ϵ defined by some $C \in \mathcal{C}_f$. Hence, CI-Apx for \mathcal{C} (i.e. CI for all relations \mathcal{R}_C^ϵ) implies CI for \mathcal{F} .

Theorem 1.1 (CI through Probabilistic Representation, Informal). *Let \mathcal{F} be a class of functions with probabilistic representation by \mathcal{C} . Then, any CI-Apx hash family for \mathcal{C} is a CI hash for \mathcal{F} .*

Probabilistic Constant-Degree Representation of the Bad Challenge Function. Recall that in a commit-then-open trapdoor Σ -protocol, the verification is either performed over a subset of commitment openings corresponding to $e = 0$ or a subset of openings corresponding to $e = 1$. From the unique bad challenge property, it is impossible that the verification on both subsets succeed if $x \notin L$. Thus, the **BadChallenge** function can be computed in two steps: an extraction step, to extract the messages underlying the commitments of one of the aforementioned subsets, say the one corresponding to $e = 1$, followed by an efficient verification (for $e = 1$) over the extracted values. If the verification accepts, then the bad challenge must be $e = 1$ and, otherwise, the bad challenge is either $e = 0$ or does not exist (in which case a is not in the relation and the output may be arbitrary). Hence, we can split the task of probabilistically representing **BadChallenge** to two sub-tasks: extraction and post-extraction verification.

Post-extraction Verification as a 3-CNF. The post-extraction verification is an arbitrary polynomial computation and, generally, may not have probabilistic constant-degree representation as is. The first step towards a constant-degree approximation of **BadChallenge** is observing that, by relying on the Cook-Levin approach for expressing the verification procedure as a 3-CNF satisfiability problem, we may reduce the complexity of the verification to 3-CNF as follows. Let Φ_e denote the 3-CNF formula that captures the verification corresponding to challenge e ; that is, Φ_e has a satisfying witness w_e if and only if the verifier accepts the prover's second message for challenge bit e . The prover can compute w_e efficiently (using the Cook-Levin approach, this witness simply consists of all intermediate steps of the verification). Therefore, we let the prover also include commitments to w_0, w_1 in his first message. When the verifier sends challenge e , the prover also provides openings for w_e , and the verifier checks decommitments then evaluates Φ_e . By transforming the protocol as described, the bad challenge computation now consists, as before, of extraction, then an evaluation of the 3-CNF formula Φ_1 , rather than an arbitrary poly-time verification.

We can then use standard well-known randomization techniques to probabilistically approximate any 3-CNF formula by constant-degree polynomials (see Lemma 3.13).

Extraction via a Randomized Linear Function. For the extraction step, we observe that by adapting the (low-noise) LPN-based PKE scheme of Damgård and Park [12] (which is closely related to the PKE scheme by Alekhnovich [1]) we can construct an extractable commitment scheme whose extraction algorithm can be probabilistically represented by a linear function. The secret extraction key is a matrix \mathbf{S} , and the public key consists of a matrix \mathbf{A} together with $\mathbf{B} = \mathbf{A} \cdot \mathbf{S} + \mathbf{E}$. Here, \mathbf{E} is chosen from a noise distribution with suitably low noise rate. To compute a commitment for bit x , the Commit algorithm chooses a low Hamming weight vector \mathbf{r} , and outputs $\mathbf{u} = \mathbf{r}\mathbf{A}$ and $\mathbf{c} = \mathbf{r}\mathbf{B} + x^\ell$. The opening for the commitment is the randomness \mathbf{r} , and the verification algorithm simply checks that \mathbf{r} has low Hamming weight, and that the Commit algorithm, using \mathbf{r} , outputs the correct commitment. Finally, note that using \mathbf{S} , one can extract the

message underlying a commitment (\mathbf{u}, \mathbf{c}) : simply compute $\mathbf{uS} + \mathbf{c} = x^\ell + \mathbf{rE}$. By carefully setting the LPN-parameters (the noise distribution is Bernoulli with parameter $1/n^c$ for some fixed constant $c \in (1/2, 1)$), we ensure that if (\mathbf{u}, \mathbf{c}) is a valid commitment (i.e. can be opened with some x and \mathbf{r}), then \mathbf{rE} has sufficiently low Hamming weight. Therefore, by sampling a random column \mathbf{s} in \mathbf{S} , we get that $\mathbf{us} + \mathbf{c} = x$ with sufficiently high probability.

The Case of Invalid Commitments. We have shown that, using a distribution over linear functions, we can approximate extraction of valid commitments. A cheating prover, however, may chose to send invalid commitments. We claim that, in such a case, we may allow the probabilistic representation to behave arbitrarily.

Fix some $x \notin L$ and a first message a . If there exist no bad challenge for a or if the (unique) bad challenge is $e = 1$, then all commitments in a corresponding to inputs of Φ_1 must be valid (since the prover is able to open them in a way that is accepted by the verifier). Thus, we potentially have a problem only in the case where $e = 0$ is the bad challenge, i.e. the commitments of input bits to Φ_0 are valid and $\Phi_0(w_0) = 1$ on their respective openings w_0 . Our concern is that since our bad challenge function only looks at the Φ_1 locations, which may be arbitrary invalid commitments, we have no guarantee on the extraction, and therefore our bad challenge function will output $e = 1$ even though the unique bad challenge is $e = 0$. We show that this is not possible.

Let w'_1 be the arbitrary value computed by the approximate extraction algorithm on the possibly invalid commitments in the locations of the Φ_1 inputs. We will see that it still must be the case that $\Phi_1(w'_1) = 0$ and therefore the bad challenge function outputs $e = 0$ as required. The reason is that otherwise we can put together *valid commitments of both* w_0 and w'_1 , so as to create a first message a' which refutes the soundness of the original Σ -protocol, since it can be successfully opened both for $e = 0$ and for $e = 1$.

Constructing CI for Approximable Relations. The main idea behind recent constructions of CI for relations searchable by some function class \mathcal{C} [9, 24] is to construct a *somewhere statistical* CI hash family \mathcal{H} . That is, one where there exists, for any $C \in \mathcal{C}$, a distribution of hash functions \mathcal{H}_C that are indistinguishable from the real \mathcal{H} , and are statistically CI for that specific C . Namely, for any C , there exists no x such that $\mathcal{H}_C(x) = C(x)$ or, equivalently, the image of the “correlation function” $x \mapsto \mathcal{H}_C(x) + C(x) \pmod 2$ does not contain 0.

Our Approach for CI-Apx: Sparse Correlations. Our first observation is that if we are able to construct a hash family \mathcal{H} where, for every $C \in \mathcal{C}$, the function $x \mapsto \mathcal{H}_C(x) + C(x)$ actually has exponentially-sparse image (as a fraction of the entire space), then we obtain (somewhere statistical) CI-Apx for \mathcal{C} .

To see this, consider the hash function $\hat{\mathcal{H}}(x) = \mathcal{H}(x) + r \pmod 2$, where r is a uniformly random string sampled together with the hash key. The task

of breaking CI of $\hat{\mathcal{H}}(x)$ for some $C \in \mathcal{C}$ reduces to the task of finding x s.t. $\mathcal{H}_C(x) + C(x) = r \pmod 2$. Clearly, with overwhelming probability, such x does not exist when the image of $\mathcal{H}_C(x) + C(x)$ is sufficiently small. We can push our statistical argument even further to claim CI-Apx for \mathcal{C} : an adversary that breaks the CI-Apx of $\hat{\mathcal{H}}$ for \mathcal{C} finds x s.t. $\mathcal{H}_C(x)$ is in a small Hamming-ball around $C(x)$, i.e. $\mathcal{H}_C(x) + C(x) + z = r \pmod 2$, where z is a vector with relative Hamming weight at most ϵ . If $x \mapsto \mathcal{H}_C(x) + C(x)$ has exponentially-sparse image, then (for properly set parameters) so does $(x, z) \mapsto \mathcal{H}_C(x) + C(x) + z$, and therefore it is very unlikely that r is in the image.

Our goal is thus reduced to constructing a hash family \mathcal{H} , with indistinguishable distributions \mathcal{H}_C as described above, such that, for every $C \in \mathcal{C}$, the function $x \mapsto \mathcal{H}_C(x) + C(x)$ has exponentially-sparse image.

Construction from Trapdoor Hash. Our construction of CI-Apx is based on trapdoor hash (TDH) [14]. At a high level, trapdoor hash allows us to “encrypt” any function $C : x \mapsto y$ to an encoding $E : x \mapsto e$ such that C is computationally hidden given a description of E and yet, for any input x , $y = C(x)$ is *almost* information-theoretically determined by $e = E(x)$. More accurately, the range of the correlation $e + y \pmod 2$ is *sparse*. The idea is then to use such an encoding as the hash function \mathcal{H}_C described above.

More specifically, in a rate-1 TDH for a function class \mathcal{C} , we can generate, for any $C \in \mathcal{C}$, an encoding key ek_C that comes with a trapdoor td_C . Using the encoding key ek_C , one can compute a value $e \leftarrow E(\text{ek}_C, x)$ which is essentially a rate-1 encoding of $C(x)$ (i.e. $|e| = |C(x)|$). There exists also a decoding algorithm D which determines the value $C(x)$ as $C(x) = e + D(\text{td}_C, h, e)$, i.e. given e and “little additional information” about x in the form of a hash value $h = H(x)$ whose length is independent of the length of x . The security property we are interested in is *function privacy*: for any $C, C' \in \mathcal{C}$, the encoding keys ek_C and $\text{ek}_{C'}$ are indistinguishable.

We use rate-1 TDH to construct, for every $C \in \mathcal{C}$, a hash family \mathcal{H}_C such that: (i) the “correlation function” $x \mapsto \mathcal{H}_C(x) + C(x)$ has exponentially-sparse image for all $C \in \mathcal{C}$, and (ii) \mathcal{H}_C and $\mathcal{H}_{C'}$ are indistinguishable, for all $C \neq C'$. This suffices to construct CI hash for any class of functions \mathcal{F} with probabilistic representation in \mathcal{C} , as outlined above.

In the heart of our construction is the following simple observation: from the correctness of the TDH, it holds that $E(\text{ek}_C, x) + D(\text{td}_C, H(x), e) = C(x)$. Put differently, if we define $\mathcal{H}_C(x) = E(\text{ek}_C, x)$, then it holds that $\mathcal{H}_C(x) + C(x) = D(\text{td}_C, H(x), e)$. This value depends on x only through its hash $H(x)$. If the hash function H is sufficiently compressing, i.e. the length of the hash is much smaller than $|C(x)|$, then we obtain an exponentially-sparse image for $\mathcal{H}_C(x) + C(x)$ and, essentially, requirement (i) from above. Property (ii) follows from the function privacy of the underlying TDH. Overall, we get the following result.

Theorem 1.2 (CI-Apx from TDH, Informal). *Assume there exists a rate-1 TDH for \mathcal{C} . Then, there exists a CI hash for relations approximable by \mathcal{C} (CI-Apx for \mathcal{C}).*

We note that the notion of TDH that we require deviates slightly from the one defined in [14]. On one hand, they require properties that we do not, such as input privacy, and they require that the decoding algorithm is efficiently computable, whereas for our purposes inefficient decoding would have sufficed. On the other hand, we require that the underlying TDH satisfies an enhanced notion of correctness, which is satisfied by all known constructions of TDH.

We obtain CI-Apx for constant degree from standard assumptions by instantiating Theorem 1.2 based on the work of Döttling et al. [14]. They construct rate-1 TDH scheme for linear functions from various standard assumptions, including QR, DCR and LWE. Such a scheme can be easily bootstrapped to support polynomials of constant degree $d > 1$. For the DDH assumption, they construct TDH for a stricter class of “index functions”. We show in the full version [8] that their construction can be slightly adjusted, based on existing ideas, to capture also constant-degree functions and, hence, get an instantiation also from DDH.

1.2 Paper Organization

In Sect. 2, we provide some essential preliminaries. In Sect. 3, we present the framework which allows using our CI constructions to obtain NIZK, starting with the generic paradigm laid out by prior work. In Sect. 4, we show how to exploit a simple probabilistic representation of a function class for obtaining CI hash and, lastly, in Sect. 5, we show our construction of CI-Apx from TDH.

2 Preliminaries

Notation. For an integer $n \in \mathbb{N}$, $[n]$ denotes the set $\{1, \dots, n\}$. We use λ for the security parameter and $\text{negl}(\lambda)$ and $\text{poly}(\lambda)$ for a negligible function and, resp., a polynomial in λ . We use $\stackrel{c}{\equiv}$ and $\stackrel{s}{\equiv}$ to denote computational and, resp., statistical indistinguishability between two distribution ensembles. For a distribution (or a randomized algorithm) D we use $x \stackrel{\$}{\leftarrow} D$ to say that x is sampled according to D and use $x \in D$ to say that x is in the support of D . For a set S we overload the notation to use $x \stackrel{\$}{\leftarrow} S$ to indicate that x is chosen uniformly at random from S .

2.1 Learning Parity with Noise

We hereby define the standard *Decisional Learning Parity with Noise (DLPN)* assumption, which we use in this paper.

Definition 2.1 (Decisional LPN Assumption). *Let $\tau : \mathbb{N} \rightarrow \mathbb{R}$ be such that $0 < \tau(\lambda) < 0.5$ for all λ , and let $n := n(\lambda)$ and $m := m(\lambda)$ be polynomials such that $m(\lambda) > n(\lambda)$ for all λ . The (n, m, τ) -Decisional LPN ((n, m, τ) -DLPN) assumption states that for any PPT adversary \mathcal{A} , there exists a negligible function $\text{negl} : \mathbb{N} \rightarrow \mathbb{R}$, such that*

$$|\Pr[\mathcal{A}(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}) = 1] - \Pr[\mathcal{A}(\mathbf{A}, \mathbf{b}) = 1]| < \text{negl}(\lambda)$$

where $\mathbf{A} \stackrel{\$}{\leftarrow} \mathbb{Z}_2^{m \times n}$, $\mathbf{s} \stackrel{\$}{\leftarrow} \mathbb{Z}_2^n$, $\mathbf{e} \stackrel{\$}{\leftarrow} \text{Ber}_\tau^m$ and $\mathbf{b} \stackrel{\$}{\leftarrow} \mathbb{Z}_2^m$.

It is well-known that DLPN remains secure even given polynomially many samples of independent secrets and error vectors.

Proposition 2.2. *Let τ , n and m be as in Definition 2.1 above, and let $k := k(\lambda)$ be an arbitrary polynomial in the security parameter. Then, under the (n, m, τ) -DLPN assumption, for any PPT adversary \mathcal{A} , there exists a negligible function negl such that*

$$|\Pr[\mathcal{A}(\mathbf{A}, \mathbf{AS} + \mathbf{E}) = 1] - \Pr[\mathcal{A}(\mathbf{A}, \mathbf{B}) = 1]| < \text{negl}(\lambda)$$

where $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_2^{m \times n}$, $\mathbf{S} \xleftarrow{\$} \mathbb{Z}_2^{n \times k}$, $\mathbf{E} \xleftarrow{\$} \text{Ber}_\tau^{m \times k}$ and $\mathbf{B} \xleftarrow{\$} \mathbb{Z}_2^{m \times k}$.

2.2 Trapdoor Hash

We hereby recall the definition of *trapdoor hash functions* (TDH) from Döttling et al. [14], with few minor modifications. First, we are fine with weakly correct trapdoor hash schemes (as defined in [14]), where we allow the error in correctness to be two-sided. This modification further allows us to simplify the syntax of decoding for rate-1 schemes. Second, to construct correlation intractable hash, we do not require the trapdoor hash scheme to be input-private (i.e. that the hash of an input x hides x) and, consequently, we assume w.l.o.g. that the hash and encoding functions, H and E , are deterministic (in the original definition, H and E share the same randomness - this was necessary for achieving both input privacy and correctness).

Definition 2.3 (Rate-1 Trapdoor Hash). *A rate-1 trapdoor hash scheme (TDH) for a function class $\mathcal{C} = \{C_n : \{0, 1\}^n \rightarrow \{0, 1\}\}$ is a tuple of five PPT algorithms $\text{TDH} = (S, G, H, E, D)$ with the following properties.*

- **Syntax:**
 - $\text{hk} \leftarrow S(1^\lambda, 1^n)$. The sampling algorithm takes as input a security parameter λ and an input length n , and outputs a hash key hk .
 - $(\text{ek}, \text{td}) \leftarrow G(\text{hk}, C)$. The generating algorithm takes as input a hash key hk a function $C \in \mathcal{C}_n$, and outputs a pair of an encoding key ek and a trapdoor td .
 - $\mathbf{h} \leftarrow H(\text{hk}, \mathbf{x})$. The hashing algorithm takes as input a hash key hk and a string $\mathbf{x} \in \{0, 1\}^n$, and deterministically outputs a hash value $\mathbf{h} \in \{0, 1\}^n$.
 - $\mathbf{e} \leftarrow E(\text{ek}, \mathbf{x})$. The encoding algorithm takes as input an encoding key ek and a string $\mathbf{x} \in \{0, 1\}^n$, and deterministically outputs an encoding $\mathbf{e} \in \{0, 1\}$.
 - $\mathbf{e}' \leftarrow D(\text{td}, \mathbf{h})$. The decoding algorithm takes as input a trapdoor td , a hash value $\mathbf{h} \in \{0, 1\}^n$, and outputs a 0-encoding $\mathbf{e}' \in \{0, 1\}$.
- **Correctness:** *TDH is (weakly) $(1 - \tau)$ -correct (or has two-sided τ error probability), for $\tau := \tau(\lambda) < 1$, if there exists a negligible function $\text{negl}(\lambda)$ such that the following holds for any $\lambda, n \in \mathbb{N}$, any $\mathbf{x} \in \{0, 1\}^n$ and any function $C \in \mathcal{C}_n$.*

$$\Pr[\mathbf{e} + \mathbf{e}' = C(\mathbf{x}) \pmod 2] \geq 1 - \tau - \text{negl}(\lambda)$$

where $hk \leftarrow S(1^\lambda, 1^n)$, $(ek, td) \leftarrow G(hk, C)$, $h \leftarrow H(hk, x)$, $e \leftarrow E(ek, x)$, and $e' \leftarrow D(td, h)$. When $\tau = 0$ we say that the scheme is fully correct.

- **Function Privacy:** TDH is function-private if for any polynomial-length $\{1^{n_\lambda}\}_{\lambda \in \mathbb{N}}$ and any $\{f_n\}_{n \in \mathbb{N}}$ and $\{f'_n\}_{n \in \mathbb{N}}$ such that $f_n, f'_n \in \mathcal{F}_n$ for all $n \in \mathbb{N}$, it holds that

$$\{(hk_\lambda, ek_\lambda)\}_{\lambda \in \mathbb{N}} \stackrel{c}{\equiv} \{(hk_\lambda, ek'_\lambda)\}_{\lambda \in \mathbb{N}}$$

where $hk_\lambda \stackrel{\S}{\leftarrow} S(1^\lambda, 1^{n_\lambda})$, $(ek_\lambda, td_\lambda) \stackrel{\S}{\leftarrow} G(hk_\lambda, f_{n_\lambda})$ and $(ek'_\lambda, td'_\lambda) \stackrel{\S}{\leftarrow} G(hk_\lambda, f'_{n_\lambda})$.

- **Compactness:** we require that the image length of the hash function, η , is independent of n , and is bounded by some polynomial in the security parameter λ .

As pointed in [14] (Remark 4.2), we may consider a natural extension of trapdoor hash for a general class of functions $\mathcal{C} = \{C_n : \{0, 1\}^n \rightarrow \{0, 1\}^m\}$ (where $m := m(\lambda) > 1$ is a fixed polynomial). Further, if any $C \in \mathcal{C}_n$ can be represented as m parallel computations in some class $\mathcal{C}'_n : \{0, 1\}^n \rightarrow \{0, 1\}$, then a trapdoor hash scheme for $\mathcal{C}' = \{\mathcal{C}'_n\}$ directly implies a trapdoor hash scheme for \mathcal{C} with hash length independent in m .

2.3 Extractable Commitments

We hereby provide the definition of an extractable commitment scheme.⁴

Definition 2.4 (Extractable Commitment). An extractable (bit) commitment scheme is a tuple of four PPT algorithms $\text{Com} = (\text{Gen}, \text{Commit}, \text{Verify}, \text{Extract})$ with the following properties.

- **Syntax:**
 - $(pk, td) \leftarrow \text{Gen}(1^\lambda)$: The key generation algorithm takes as input the security parameter 1^λ and outputs a pair of a public key pk and trapdoor td .
 - $\text{com} \leftarrow \text{Commit}(pk, x; r)$: The committing algorithm takes as input a public key pk , a bit $x \in \{0, 1\}$ and randomness r , and outputs a commitment com .
 - $\{0, 1\} \leftarrow \text{Verify}(pk, \text{com}, x; r)$: The verification algorithm takes as input a public key pk , a commitment com , a bit $x \in \{0, 1\}$ and randomness $r \in \{0, 1\}^*$, then either accepts or rejects.
 - $x' \leftarrow \text{Extract}(td, \text{com})$: The extraction algorithm takes as input a trapdoor td and a commitment com and outputs a bit $x' \in \{0, 1\}$ or \perp .
- **Correctness:** Com is correct if there exists a negligible function negl , such that for any $x \in \{0, 1\}$,

$$\Pr[\text{Verify}(pk, \text{Commit}(pk, x; r), x; r)] > 1 - \text{negl}(\lambda)$$

where $(pk, \cdot) \stackrel{\S}{\leftarrow} \text{Gen}(1^\lambda)$ and $r \stackrel{\S}{\leftarrow} \{0, 1\}^*$.

⁴ The notion of extractable commitment is equivalent to standard public-key encryption. We use the commitment terminology since it is more natural for our setting.

- **Hiding:** Com is (computationally) hiding if it holds that

$$\{\text{Commit}(\text{pk}, 0; r)\}_\lambda \stackrel{c}{\equiv} \{\text{Commit}(\text{pk}, 1; r)\}_\lambda$$

where $(\text{pk}, \cdot) \stackrel{\$}{\leftarrow} \text{Gen}(1^\lambda)$ and $r \stackrel{\$}{\leftarrow} \{0, 1\}^*$ for all $\lambda \in \mathbb{N}$.

- **Binding:** Com is (statistically) binding if there exists a negligible function negl such that

$$\Pr[\exists \text{com}, r_0, r_1 \text{ s.t. } \text{Verify}(\text{pk}, \text{com}, 0, r_0) = \text{Verify}(\text{pk}, \text{com}, 1, r_1) = 1] < \text{negl}(\lambda)$$

where $(\text{pk}, \cdot) \stackrel{\$}{\leftarrow} \text{Gen}(1^\lambda)$.

- **Extraction:** Com has correct extraction if there exists a negligible function negl , such that for any $x \in \{0, 1\}$ and $r \in \{0, 1\}^*$, if $\text{Verify}(\text{pk}, \text{Commit}(\text{pk}, x; r), x; r)$

$$\Pr[\text{Verify}(\text{pk}, \text{com}, x; r) = 1 \wedge \text{Extract}(\text{td}, \text{com}) \neq x] < \text{negl}(\lambda)$$

where $(\text{pk}, \text{td}) \stackrel{\$}{\leftarrow} \text{Gen}(1^\lambda)$ and $\text{com} = \text{Commit}(\text{pk}, x; r)$.

Remark 2.5. Throughout the paper, we will implicitly assume that if $\text{Commit}(\text{pk}, x; r) \neq \text{com}$ then $\text{Verify}(\text{pk}, x, r) \neq 1$. This is achieved by any commitment scheme with a natural verification function (that possibly performs additional verification). Notice that in such a case correct extraction implies statistical binding.

2.4 Non-interactive Zero-Knowledge Arguments

We formally define non-interactive zero knowledge arguments as follows.

Definition 2.6 (Non-interactive Zero Knowledge). Let $n := n(\lambda)$ be a polynomial in the security parameter. A non-interactive zero knowledge (NIZK) argument Π for an NP language L , with a corresponding instance-witness relation R , consists of three PPT algorithms $\Pi = (\text{Setup}, \text{P}, \text{V})$ with the following properties.

- **Syntax:**
 - $\text{crs} \leftarrow \text{Setup}(1^\lambda)$: the setup algorithm takes a security parameter 1^λ and outputs a common reference string crs .
 - $\pi \leftarrow \text{P}(\text{crs}, x, w)$: the prover takes as input the common reference string crs , a statement $x \in \{0, 1\}^n$ and a witness w such that $(x, w) \in R$, and outputs a proof π .
 - $\{0, 1\} \leftarrow \text{V}(\text{crs}, x, \pi)$: the verifier takes as input the common reference string crs , a statement $x \in \{0, 1\}^n$ and a proof π , and either accepts (outputs 1) or rejects (outputs 0).
- **Completeness:** Π is complete if for every $\lambda \in \mathbb{N}$ and $(x, w) \in R$, it holds that

$$\Pr[\text{V}(\text{crs}, x, \text{P}(\text{crs}, x, w))] = 1$$

where $\text{crs} \stackrel{\$}{\leftarrow} \text{Setup}(1^\lambda)$.

- **Soundness:** Π is sound if for every PPT cheating prover P^* , there exists a negligible function negl , such that for every $\{x_\lambda \notin L\}_\lambda$ where $x_\lambda \in \{0, 1\}^n$ for all λ , it holds that

$$\Pr[V(\text{crs}, x_\lambda, P^*(\text{crs})) = 1] < \text{negl}(\lambda)$$

where $\text{crs} \xleftarrow{\$} \text{Setup}(1^\lambda)$.

- **Zero Knowledge:** Π is zero knowledge if there exists a PPT simulator Sim such that for every $\{(x_\lambda, w_\lambda) \in R\}_\lambda$, where $x_\lambda \in \{0, 1\}^n$ for all $\lambda \in \mathbb{N}$, it holds that

$$\{(\text{crs}, P(\text{crs}, x_\lambda, w_\lambda))\}_\lambda \stackrel{c}{\equiv} \{\text{Sim}(1^\lambda, x_\lambda)\}_\lambda$$

where $\text{crs} \xleftarrow{\$} \text{Setup}(1^\lambda)$.

We further consider few optional stronger properties that a NIZK system can satisfy:

- **Adaptive Soundness:** Π is adaptively sound if for every PPT cheating prover P^* , there exists a negligible function negl , such that

$$\Pr[x \notin L \wedge V(\text{crs}, x, \pi) = 1] < \text{negl}(\lambda)$$

where $\text{crs} \xleftarrow{\$} \text{Setup}(1^\lambda)$ and $(x, \pi) \leftarrow P^*(\text{crs})$.

- **Adaptive Zero Knowledge:** Π is adaptively zero knowledge if there exist a (stateful) PPT simulator Sim such that for every PPT adversary \mathcal{A} , it holds that

$$\{\text{Real}_{\text{Sim}, \mathcal{A}}(1^\lambda)\}_\lambda \stackrel{c}{\equiv} \{\text{Ideal}_{\text{Sim}, \mathcal{A}}(1^\lambda)\}_\lambda$$

where $\text{Real}_{\text{Sim}, \mathcal{A}}$ and $\text{Ideal}_{\text{Sim}, \mathcal{A}}$ are as defined in Fig. 1.

$\text{Real}_{\text{Sim}, \mathcal{A}}(1^\lambda)$: <ol style="list-style-type: none"> 1. $\text{crs} \leftarrow \text{Setup}(1^\lambda)$ 2. $(x, w) \leftarrow \mathcal{A}(\text{crs})$, s.t. $(x, w) \in R$ 3. $\pi \leftarrow P(\text{crs}, x, w)$ 4. Output (crs, x, π) 	$\text{Ideal}_{\text{Sim}, \mathcal{A}}(1^\lambda)$: <ol style="list-style-type: none"> 1. $\text{crs} \leftarrow \text{Sim}(1^\lambda)$ 2. $(x, w) \leftarrow \mathcal{A}(\text{crs})$ 3. $\pi \leftarrow \text{Sim}(\text{crs}, x)$ 4. Output (crs, x, π)
------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Fig. 1. $\text{Real}_{\text{Sim}, \mathcal{A}}$ and $\text{Ideal}_{\text{Sim}, \mathcal{A}}$

2.5 Correlation Intractability

Correlation intractable hash [11] constitutes one of the main building blocks in our work. We hereby provide a formal definition.

Definition 2.7 (Correlation Intractable Hash). Let $\mathcal{R} = \{\mathcal{R}_\lambda\}$ be a relation class. A hash family $\mathcal{H} = (\text{Sample}, \text{Hash})$ is said to be correlation intractable for \mathcal{R} if for every non-uniform polynomial-time adversary $\mathcal{A} = \{\mathcal{A}_\lambda\}$, there exists a negligible function $\text{negl}(\lambda)$, such that for every $R \in \mathcal{R}_\lambda$, it holds that

$$\Pr[(x, \text{Hash}(k, x)) \in R] \leq \text{negl}(\lambda)$$

where $k \stackrel{\$}{\leftarrow} \text{Sample}(1^\lambda)$ and $x = \mathcal{A}_\lambda(k)$.

We further define an essential property for utilizing CI hash for obtaining NIZK protocols.

Definition 2.8 (Programmable Hash Family). A hash family $\mathcal{H} = (\text{Sample}, \text{Hash})$, with input and output length $n := n(\lambda)$ and, resp., $m := m(\lambda)$, is said to be programmable if the following two conditions hold:

- **1-Universality:** For every $\lambda \in \mathbb{N}$, $x \in \{0, 1\}^n$ and $y \in \{0, 1\}^m$,

$$\Pr[\text{Hash}_k(x) = y] = 2^{-m}$$

where $k \stackrel{\$}{\leftarrow} \text{Sample}(1^\lambda)$.

- **Programmability:** There exists a PPT algorithm $\widetilde{\text{Sample}}(1^\lambda, x, y)$ that samples from the conditional distribution $\text{Sample}(1^\lambda) \mid \text{Hash}_k(x) = y$.

3 Non-interactive Zero Knowledge from Correlation Intractability

In this section, we provide the formal framework for constructing NIZK for NP from the following building blocks:

- (i) An extractable commitment scheme where the extraction function can be probabilistically presented by constant-degree polynomials.
- (ii) A correlation intractable hash function for relations probabilistically searchable by constant-degree polynomials.

Our framework is essentially a special case of a more general paradigm that was extensively investigated in prior works [9, 10, 21] for constructing NIZKs from general correlation intractability. Our contribution in this part of the paper is relaxing the requirement for correlation intractability, assuming a commitment scheme with the above property exists.

3.1 A Generic Framework

We first recall the generic framework from Canetti et al. [9] for achieving non-interactive zero knowledge systems from correlation intractable hash.

In its most general form, the paradigm applies the Fiat-Shamir transform [16] over Σ -protocols, which are special honest-verifier ZK protocols (possibly in the CRS model), using correlation intractable hash, in a provably-sound manner.

Roughly speaking, in Σ -protocols, for every prover’s first message a there exists (if any) a unique verifier’s challenge e that may allow a cheating prover to cheat. Thus, if we instantiate Fiat-Shamir using a hash family \mathcal{H} that is CI for the relation between such pairs (a, e) , then the soundness of the transform can be reduced to the correlation intractability of \mathcal{H} : any prover who finds a first message a where $\mathcal{H}(a)$ is the “bad challenge” e , essentially breaks \mathcal{H} .

Therefore, the type of relations we target in the above outline is formally specified as follows.

Definition 3.1 (Unique-Output Relations). *We say that a class of relations $\mathcal{R} \subset \{0, 1\}^n \times \{0, 1\}^m$ is unique-output if for every $x \in \{0, 1\}^n$ there exists at most one value $y \in \{0, 1\}^m$ such that $(x, y) \in \mathcal{R}$. We sometimes use function notation to describe such an \mathcal{R} where every $R \in \mathcal{R}$ is denoted by a function $R : \{0, 1\}^n \rightarrow \{0, 1\}^m \cup \perp$ with $R(x) = y$ for $(x, y) \in R$ and $R(x) = \perp$ if there exists no such y .*

As observed in [9], we can reduce the class of relations we target in the CI to relations that are efficiently searchable, i.e. unique-output relations where the unique output is efficiently computable. It is not the case, however, that any Σ -protocol defines such a corresponding relation. This leads us to define *trapdoor Σ -protocols* [9], which are Σ -protocol where the relation between a prover’s first message and its unique “bad challenge” is efficiently computable given a trapdoor. We formalize below.

Definition 3.2 (Searchable Relations). *Let $\mathcal{R} : \{0, 1\}^n \rightarrow \{0, 1\}^m \cup \perp$ be a unique-output class of relations. We say that \mathcal{R} is searchable by a function class $\mathcal{F} : \{0, 1\}^n \rightarrow \{0, 1\}^m \cup \perp$ if for every $R \in \mathcal{R}$, there exists $f_R \in \mathcal{F}$ such that*

$$\forall x \text{ s.t. } R(x) \neq \perp, \quad (x, f_R(x)) \in R$$

We say that \mathcal{R} is efficiently searchable if \mathcal{F} is efficiently computable.

Definition 3.3 (Trapdoor Σ -Protocol [9]). *Let $\Pi = (\text{Setup}, \text{P}, \text{V})$ be a public-coin three-message honest-verifier zero knowledge proof system for a language L in the common reference string model. Define the relation class $\mathcal{R}_\Sigma(\Pi) = \{R_{\text{crs},x} \mid \text{crs} \in \text{Setup}(1^\lambda), x \notin L\}$ where*

$$R_{\text{crs},x} = \{(\mathbf{a}, \mathbf{e}) \mid \exists \mathbf{z} \text{ s.t. } \forall (\text{crs}, x, \mathbf{a}, \mathbf{e}, \mathbf{z}) = 1\}$$

We say that Π for L is a trapdoor Σ -protocol if $R_{\text{crs},x}$ is a unique-output relation (see Definition 3.1) and there exist two PPT algorithms, tdSetup and BadChallenge , with the following properties:

- **Syntax:**
 - $(\text{crs}, \text{td}) \leftarrow \text{tdSetup}(1^\lambda)$: The trapdoor setup algorithms takes as input a security parameter 1^λ and outputs a common reference string crs and a trapdoor td .
 - $e \leftarrow \text{BadChallenge}(\text{crs}, \text{td}, x, a)$: The bad challenge algorithm takes as input a common reference string crs and its trapdoor td , an instance x , and a first message a , and outputs a second message e or \perp .
- **CRS Indistinguishability:** We require that a common reference string $\text{crs} \xleftarrow{\$} \text{Setup}(1^\lambda)$ is computationally indistinguishable from a random reference string crs' sampled with a trapdoor by $(\text{crs}', \text{td}) \xleftarrow{\$} \text{tdSetup}(1^\lambda)$.
- **Correctness:** We require that for all $\lambda \in \mathbb{N}$ and any instance $x \notin L$, first message a , and (crs, td) , such that $R_{\text{crs},x}(a) \neq \perp$, it holds

$$\text{BadChallenge}(\text{crs}, \text{td}, x, a) = R_{\text{crs},x}(a)$$

Equivalently, we require that $\mathcal{R}_\Sigma(\Pi)$ is searchable by

$$\mathcal{F}_\Sigma(\Pi) = \{f_{\text{crs},\text{td},x}(a) = \text{BadChallenge}(\text{crs}, \text{td}, x, \cdot) \mid (\text{crs}, \text{td}) \in \text{Setup}(1^\lambda), x \notin L\}$$

We recall the following theorem from [9].

Theorem 3.4 ([9]). Assume:

- (i) Π is a trapdoor Σ -protocol for L .
- (ii) \mathcal{H} is a programmable correlation intractable hash family for relation searchable by $\mathcal{F}_\Sigma(\Pi)$.

Then, the Fiat-Shamir [16] transform over Π using \mathcal{H} , $\text{FS}(\Pi, \mathcal{H})$, is an NIZK argument system for L with adaptive soundness and adaptive zero-knowledge.

Canetti et al. [9] show that any correlation intractable hash family for a reasonable class of relations can be easily transformed to a programmable hash family while preserving correlation intractability. We stress, however, that our Construction of correlation intractable hash in Sect. 2.5 directly satisfies programmability.

3.2 Special Case: Commit-then-Open Protocols

Equipped with the generic framework laid by prior work, we may now present a special case that comprises the starting point of our work.

Commit-then-Open Protocols. We consider protocols of a special form called *commit-then-open Σ -protocols*. This notion captures a natural approach for constructing ZK protocols. In particular, a variant of the ZK protocol for Graph Hamiltonicity from [6, 15] is a commit-then-open Σ -protocol.

Roughly speaking, commit-then-open Σ -protocols are protocols that use a commitment scheme (possibly in the CRS model), where the prover’s first message is a commitment on some proof string π , and his second message is always a

decommitment on a subset of π , which depends on the verifier’s challenge. Upon receiving the decommitments, the verifier checks that they are valid, then runs some verification procedure on the opened values. We hereby provide a formal definition.

Definition 3.5 (Commit-then-Open Σ -Protocols). *A commit-then-open Σ -protocol is a Σ -protocol $\Pi^{\text{Com}} = (\text{Setup}^{\text{Com}}, \text{P}^{\text{Com}}, \text{V}^{\text{Com}})$, with black-box access to a commitment scheme Com (possibly in the CRS model), such that there exist four PPT algorithms:*

- $\text{crs}' \leftarrow \text{Setup}'(1^\lambda, \text{pk})$: Takes as input a security parameter 1^λ and a commitment key pk , and outputs a common reference string crs' .
- $(\pi, \text{state}) \leftarrow \text{P}_1(\text{crs}, \text{x}, \text{w})$: Takes as input a common reference string crs , an instance x and its witness w and outputs a proof $\pi \in \{0, 1\}^\ell$ (for some polynomial $\ell := \ell(\lambda)$) and a local state state .
- $I \leftarrow \text{P}_2(\text{crs}, \text{x}, \text{w}, e, \text{state})$: Takes as input crs , x , w and state as above, and a verifier’s challenge $e \in \{0, 1\}^*$, and outputs a subset $I \subseteq [\ell]$.
- $\{0, 1\} \leftarrow \text{V}'(\text{crs}, \text{x}, e, (I, \pi_I))$: Takes as input crs , x , $e \in \{0, 1\}^*$, $I \subseteq [\ell]$ as above, and a substring of the proof $\pi_I \in \{0, 1\}^{|I|}$.

using which Π^{Com} is defined as follows:

- $\text{Setup}^{\text{Com}}(1^\lambda)$: Sample a commitment key $\text{pk} \leftarrow \text{Com.Gen}(1^\lambda)$ and possibly additional output $\text{crs}' \leftarrow \text{Setup}'(1^\lambda, \text{pk})$, and output

$$\text{crs} = (\text{crs}', \text{pk})$$

- $\text{P}^{\text{Com}}(\text{crs}, \text{x}, \text{w})$: The prover computes $(\pi, \text{state}) \leftarrow \text{P}_1(\text{crs}, \text{x}, \text{w})$, keeps the local state state , and sends a commitment on the proof π to the verifier,

$$a = \text{Com.Commit}(\text{pk}, \pi)$$

- $\text{P}^{\text{Com}}(\text{crs}, \text{x}, \text{w}, e)$: The prover’s second message consists of a decommitment on the proof bits corresponding to locations $I \leftarrow \text{P}_2(\text{crs}, \text{x}, \text{w}, e, \text{state})$,

$$z = (I, \text{Com.Decommit}(a_I))$$

- $\text{V}^{\text{Com}}(\text{crs}, \text{x}, a, e, z)$: The verifier verifies that z contains a valid decommitment to π_I and outputs

$$\text{V}'(\text{crs}, \text{x}, e, (I, \pi_I))$$

We sometimes override notation and denote $\Pi^{\text{Com}} = (\text{Setup}', \text{P}_1, \text{P}_2, \text{V}')$.

Proposition 3.6 ([6, 15]). *There exists a commit-then-open Σ -protocol with soundness $1/2$ for an NP-complete language L .*

It turns out that commit-then-open Σ -protocols allow us to relax the CI requirement for a sound Fiat-Shamir to CI for relations that are *probabilistically searchable* by constant-degree polynomials. We elaborate in the following.

3.3 Probabilistically Searchable Relations

We consider a standard notion of approximation, which we refer to as *probabilistic representation*. Roughly speaking, a function f is probabilistically represented by a function class \mathcal{C} if there exists a randomized $C \in \mathcal{C}$ that computes f with high probability, on any input.

Definition 3.7 (Probabilistic Representation). *Let $n, m \in \mathbb{N}$ and $0 < \epsilon < 1$. Let $f : \{0, 1\}^n \rightarrow \{0, 1\}^m \cup \perp$ be a function and denote $f(x) = (f_1(x), \dots, f_m(x))$ where $f_i : \{0, 1\}^n \rightarrow \{0, 1\} \cup \perp$ for all $i \in [m]$. A (bit-by-bit) ϵ -probabilistic representation of f by a class of functions $\mathcal{C} : \{0, 1\}^n \rightarrow \{0, 1\}$ consists of m distributions $\mathfrak{C}_1, \dots, \mathfrak{C}_m \subseteq \mathcal{C}$ such that*

$$\forall i \in [m], \forall x \text{ s.t. } f(x) \neq \perp, \quad \Pr_{C_i \leftarrow \mathfrak{C}_i} [f_i(x) = C_i(x)] > 1 - \epsilon$$

The following simple lemma connects between probabilistic representation and approximation. Its proof follows immediately from Chernoff’s tail bound.

Lemma 3.8 (From Probabilistic Representation to Approximation). *Let $n \in \mathbb{N}$, $\epsilon := \epsilon(\lambda) > 0$, and $m := m(\lambda)$ be a sufficiently large polynomial. For any $\lambda \in \mathbb{N}$, let $f : \{0, 1\}^n \rightarrow \{0, 1\}^m \cup \perp$, and let $\mathfrak{C} = (\mathfrak{C}_1, \dots, \mathfrak{C}_m)$ be an ϵ -probabilistic representation of f by $\mathcal{C} : \{0, 1\}^n \rightarrow \{0, 1\}$. Then, there exists a negligible function negl , such that*

$$\forall x \text{ s.t. } f(x) \neq \perp, \quad \Pr_{C \leftarrow \mathfrak{C}} [\Delta(f(x), C(x)) > 2\epsilon m] < \text{negl}(\lambda)$$

If a class of functions \mathcal{R} is searchable by functions with probabilistic representation by \mathcal{C} , we say that \mathcal{R} is *probabilistically searchable* by \mathcal{C} .

Definition 3.9 (Probabilistically-Searchable Relations). *Let $\mathcal{R} : \{0, 1\}^n \rightarrow \{0, 1\}^m \cup \perp$ be a unique-output class of relations. We say that \mathcal{R} is ϵ -probabilistically searchable by $\mathcal{C} : \{0, 1\}^n \rightarrow \{0, 1\}$ if it is searchable by \mathcal{F} and, for every $R \in \mathcal{R}$, letting $f_R \in \mathcal{F}$ be the corresponding search function (see Definition 3.2), $f_R \in \mathcal{F}$ has an ϵ -probabilistic representation by \mathcal{C} .*

Notice that CI for relations searchable by \mathcal{F} is a weaker notion than relation probabilistically searchable by \mathcal{F} . Our hope is to be able to probabilistically represent \mathcal{F} by a much simpler class of functions \mathcal{C} such that the CI task is actually simplified.

3.4 CI for Probabilistic Constant-Degree Is Sufficient for NIZK

Lastly, we show that through commit-then-open protocols, we can reduce our task to achieving CI for relations probabilistically searchable by constant-degree polynomials. More specifically, we show that any commit-then-open Σ -protocol Π^{Com} can be transformed to a slightly different commit-then-open Σ -protocol $\tilde{\Pi}^{\text{Com}}$ such that:

- Assuming Com is extractable, $\tilde{\Pi}^{\text{Com}}$ is a trapdoor Σ -protocol.
- Assuming, further, that the extraction function $f_{\text{td}}(a) = \text{Com.Extract}(\text{td}, a)$ has probabilistic constant-degree representation, then so does the trapdoor function BadChallenge , corresponding to $\tilde{\Pi}^{\text{Com}}$ and, therefore $\mathcal{R}_{\Sigma}(\tilde{\Pi}^{\text{Com}})$ is probabilistically searchable by constant-degree polynomials.

We formalize below.

Theorem 3.10. *Let Π^{Com} be a commit-then-open Σ -protocol for L with soundness $1/2$ where the output of P_1 is of length $\ell := \ell(\lambda)$. Let Com be a statistically-binding extractable commitment scheme where, for any td , the function $f_{\text{td}}(x) = \text{Com.Extract}(\text{td}, x)$ has an ϵ -probabilistic representation by c -degree polynomials, for a constant $c \in \mathbb{N}$ and $0 < \epsilon(\lambda) < 1/\ell$. Then, for any polynomial $m := m(\lambda)$, there exists a trapdoor Σ -protocol $\tilde{\Pi}^{\text{Com}}$ for L with soundness 2^{-m} such that $\mathcal{R}_{\Sigma}(\tilde{\Pi}^{\text{Com}})$ (see Definition 3.3) is ϵ' -probabilistically searchable by $6cc'$ -degree polynomials, where $c' \in \mathbb{N}$ is an arbitrary constant and $\epsilon' = \ell \cdot \epsilon + 2^{-c'}$.*

Combining Proposition 3.6, Theorem 3.10, and Theorem 3.4, we obtain the following.

Corollary 3.11 (Sufficient Conditions for NIZK for NP). *The following conditions are sufficient to obtain a NIZK argument system for NP (with adaptive soundness and adaptive zero-knowledge):*

- (i) *A statistically-binding extractable commitment scheme where, for any td , the function $f_{\text{td}}(x) = \text{Extract}(\text{td}, x)$ has an ϵ -probabilistic representation by c -degree polynomials, for a constant $c \in \mathbb{N}$ and $0 < \epsilon(\lambda) < 1/\ell(\lambda)$ for an arbitrarily large polynomial ℓ .*
- (ii) *A programmable correlation intractable hash family for relations ϵ -probabilistically searchable by c' -degree polynomials, for some constant $\epsilon > 0$ and arbitrarily large constant $c' \in \mathbb{N}$.*

For instantiating Corollary 3.11 based on standard assumptions, we may use a variant of the LPN-based PKE scheme of Damgård and Park [12] to construct a suitable extractable commitment scheme as required in (i). We discuss the details of the commitment scheme in the full version [8] and, in the following section, focus on how to obtain CI hash schemes satisfying (ii).

We now proceed and prove Theorem 3.10.

Proof of Theorem 3.10. We start by presenting the transformation from Π^{Com} to $\tilde{\Pi}^{\text{Com}}$. In fact, for simplicity, we first show how to construct a protocol $\tilde{\Pi}_1^{\text{Com}}$ which has soundness $\frac{1}{2}$. The final protocol $\tilde{\Pi}^{\text{Com}}$ with amplified soundness simply consists of m parallel repetitions of $\tilde{\Pi}_1^{\text{Com}}$. We later show that all required properties are preserved under parallel repetition and, therefore, we now focus on $\tilde{\Pi}_1^{\text{Com}}$.

Using the Cook-Levin approach, we represent any (poly-size) circuit C as a (poly-size) 3-CNF formula Φ_C such that for any input x , $C(x) = 1$ if and only if there exists an assignment w for which $\Phi_C(x, w) = 1$. We call such an assignment w a *Cook-Levin witness* for $C(x)$.

Construction 3.1. Let $\Pi^{\text{Com}} = (\text{Setup}', P_1, P_2, V')$ be a commit-then-open Σ -protocol with soundness $1/2$, i.e. where the verifier's challenge e consists of a single public coin. We construct a commit-then-open Σ -protocol $\tilde{\Pi}_1^{\text{Com}} = (\text{Setup}', \tilde{P}_1, \tilde{P}_2, \tilde{V}')$ as follows.⁵

- $\tilde{P}_1(\text{crs}, x, w)$: The prover generates a proof $\pi \leftarrow P_1(\text{crs}, x, w)$ and computes $I_0 \leftarrow P_2(\text{crs}, x, w, 0)$ and $I_1 \leftarrow P_2(\text{crs}, x, w, 1)$. Without loss of generality, we assume that subsets $I_0, I_1 \subseteq [\ell]$ are represented, in the natural way, as matrices over \mathbb{Z}_2 such that $I_e \cdot \pi = \pi_{I_e}$ (for $e \in \{0, 1\}$). It then generates, for every $e \in \{0, 1\}$, a Cook-Levin witness w_e for the computation $C_{\text{crs}, x, e}(I_e, \pi_{I_e}) = 1$ where

$$C_{\text{crs}, x, e}(I_e, \pi_{I_e}) := V'(\text{crs}, x, e, I_e, \pi_{I_e})$$

The prover then outputs

$$\tilde{\pi} = (\pi, I_0, I_1, w_0, w_1)$$

- $\tilde{P}_2(\text{crs}, x, w, e)$: Outputs the subset \tilde{I}_e , which corresponds to the locations of π_{I_e} , I_e , and w_e in z .
- $\tilde{V}'(\text{crs}, x, e, (\tilde{I}, \tilde{\pi}_{\tilde{I}}))$: The verifier parses $\tilde{\pi}_{\tilde{I}} = (I_e, \pi_{I_e}, w_e)$ then verifies that

$$\Phi_{\text{crs}, x, e}(I_e, \pi_{I_e}, w_e) = 1$$

where $\Phi_{\text{crs}, x, e}$ is the Cook-Levin 3-CNF formula for $C_{\text{crs}, x, e}$ verification.

We begin by showing that, if the underlying commitment scheme is extractable, then $\tilde{\Pi}_1^{\text{Com}}$ is a trapdoor Σ -protocol.

Lemma 3.12. Let $\text{Com} = (\text{Gen}, \text{Commit}, \text{Verify}, \text{Extract})$ be a statistically binding extractable commitment scheme, and let $\Pi^{\text{Com}} = (\text{Setup}', P_1, P_2, V')$ be commit-then-open Σ -protocol with soundness $1/2$. Then, $\tilde{\Pi}_1^{\text{Com}}$ from Construction 3.1 is a trapdoor Σ -protocol with:

- $\text{tdSetup}(1^\lambda)$: Sample $(\text{pk}, \text{td}) \leftarrow \text{Com.Gen}(1^\lambda)$ and $\text{crs}' \leftarrow \text{Setup}'(1^\lambda, \text{pk})$, then output

$$((\text{crs}', \text{pk}), \text{td})$$
- $\text{BadChallenge}(\text{crs}, \text{td}, x, a)$: Compute $\tilde{\pi}' \leftarrow \text{Extract}(\text{td}, a)$, and parse $\tilde{\pi}' = (\pi', I_0, I_1, w_0, w_1) \in \{0, 1, \perp\}^*$. For every $e \in \{0, 1\}$, if $I_e \in \{0, 1\}^*$, set $\tilde{\pi}'_e = (I_e, \pi'_{I_e}, w_e)$ and otherwise set $\tilde{\pi}'_e = \perp$.
 1. If $\tilde{\pi}'_0 \in \{0, 1\}^*$ and $\Phi_{\text{crs}, x, 0}(\tilde{\pi}'_0) = 1$, output 0.
 2. If $\tilde{\pi}'_1 \in \{0, 1\}^*$ and $\Phi_{\text{crs}, x, 1}(\tilde{\pi}'_1) = 1$, output 1.
 3. Otherwise, output \perp .

⁵ Recall that the algorithms for the actual setup, prover and verifier, are obtained by combining the algorithms in the construction with the commitment scheme Com , as described in Definition 3.5.

Proof. It is evident that, based on the statistical binding of Com , the transformation preserves the soundness of the protocol and that, based on the computational hiding of Com , it also preserves honest-verifier zero knowledge (the honest-verifier uses the simulator of Π^{Com} in a straight-forward manner and generates random commitments where necessary).

It is also clear that $\text{tdSetup}(1^\lambda)$ outputs a common reference string identical to $\text{crs} \leftarrow \text{Setup}(1^\lambda)$. We therefore focus on proving correctness of BadChallenge .

Let $x \notin L$, and crs , td and a be such that $R_{\text{crs},x}(a) = e \neq \perp$ (where $R_{\text{crs},x} \in \mathcal{R}_\Sigma(\tilde{H}_1^{\text{Com}})$ as defined in Definition 3.3). From definition of $R_{\text{crs},x}$, there exists $(\tilde{I}, \tilde{\pi}_{\tilde{I}})$ such that $\tilde{V}(\text{crs}, x, a, e, (\tilde{I}, \tilde{\pi}_{\tilde{I}})) = 1$. From the statistical binding and correct extraction of Com , it necessarily holds that $\tilde{\pi}_{\tilde{I}} = \text{Extract}(a_{\tilde{I}}) = \tilde{\pi}'_e$. Further, we have $V'(\text{crs}, x, e, (\tilde{I}, \tilde{\pi}_{\tilde{I}})) = 1$ and, therefore, $\Phi_{\text{crs},x,e}(\tilde{\pi}_{\tilde{I}}) = 1$ implying

$$\Phi_{\text{crs},x,e}(\tilde{\pi}'_e) = 1 \tag{1}$$

On the other hand, since $R_{\text{crs},x}$ is a unique-output relation (due to Lemma 3.12 and Definition 3.3), then there exists no $(\tilde{I}, \tilde{\pi}_{\tilde{I}})$ such that $\tilde{V}(\text{crs}, x, a, 1 - e, (\tilde{I}, \tilde{\pi}_{\tilde{I}})) = 1$ and, in particular, this holds for $\tilde{\pi}'_{1-e}$. Therefore, if $\tilde{\pi}'_{1-e}$ is a valid opening of $a_{\tilde{I}}$ (with \tilde{I} being the set of locations supposedly corresponding to $(I_{1-e}, \pi'_{I_{1-e}}, w_{1-e})$ in a), i.e. $\tilde{\pi}'_{1-e} = \text{Extract}(a_{\tilde{I}})$, then

$$\Phi_{\text{crs},x,1-e}(\tilde{\pi}'_{1-e}) = 0 \tag{2}$$

By combining (1) and (2), we obtain that $\text{BadChallenge}(\text{crs}, \text{td}, x, a) = e = R_{\text{crs},x}(a)$ and we finish. \square

Having shown that the protocol is a trapdoor Σ -protocol, our goal now is to show that the trapdoor function BadChallenge , which is specified in Lemma 3.12, has probabilistic representation as constant degree polynomials. Observe that, roughly speaking, BadChallenge is a composition of the extraction function, which we assume has a probabilistic constant-degree representation, and an evaluation of two CNF formulas. Since the protocol is a Σ -protocol, we show that, in fact, the randomized polynomials need to (probabilistically) evaluate only one of these formulas on the extracted value.

Thus, as a first step towards constructing efficient probabilistic constant-degree representation for BadChallenge , we seek to evaluate CNF formulas using randomized polynomials. This is done through the following lemma using standard randomization techniques. We refer the reader to the full version [8] for a full proof.

Lemma 3.13 (*k*-CNF via Probabilistic Polynomials). *Let $\ell, k, c \in \mathbb{N}$. For any k -CNF formula $\Phi : \{0, 1\}^\ell \rightarrow \{0, 1\}$, there exists a 2^{-c} -probabilistic representation by $c(k + 1)$ -degree polynomials \mathfrak{P}_Φ .*

We now use Lemma 3.13, and the assumption that Extract has probabilistic constant-degree representation, to obtain such a representation for BadChallenge .

Lemma 3.14. *Let $c, c' \in \mathbb{N}$ be arbitrary constants, and let $0 < \epsilon(\lambda) < 1/\ell(\lambda)$. Let Com be an extractable commitment scheme where, for any td , the extraction function $\text{Extract}(\text{td}, \cdot)$ has an ϵ -probabilistic representation by c -degree polynomials. Consider the protocol $\tilde{\Pi}^{\text{Com}}$ from Construction 3.1. Then, the function*

$$f_{\text{crs}, \text{td}, x}(a) = \text{BadChallenge}(\text{crs}, \text{td}, x, a),$$

as defined in Lemma 3.12, has ϵ' -probabilistic representation by $6cc'$ -degree polynomials, with $\epsilon' = \ell \cdot \epsilon + 2^{-c'}$.

Proof. Let \mathfrak{P}_{td} be the efficient ϵ -probabilistic representation of $\text{Extract}(\text{td}, \cdot)$ by c -degree polynomials. We now show a probabilistic representation of $f_{\text{crs}, \text{td}, x}$ by c' -degree polynomials, denoted by $\mathfrak{P}_{\text{crs}, \text{td}, x}$. For simplicity, we describe $\mathfrak{P}_{\text{crs}, \text{td}, x}$ as a randomized algorithm.

$\mathfrak{P}_{\text{crs}, \text{td}, x}(a)$:

1. Sample $P_{\text{td}} \xleftarrow{\$} \mathfrak{P}_{\text{td}}^\ell$, and compute $\tilde{z} = P_{\text{td}}(a)$.
2. Parse $\tilde{z} = (z, I_0, I_1, w_0, w_1)$ and compute $\tilde{z}_1 = (I_1, z_{I_1}, w_1)$.
3. Denote by \mathfrak{P}_Φ the $2^{-c'}$ -probabilistic representation of $\Phi_{\text{crs}, x, 1}$ by $3c'$ -degree polynomials (due to Lemma 3.13). Sample $P_\Phi \xleftarrow{\$} \mathfrak{P}_\Phi$, then output $b = P_\Phi(\tilde{z}_e)$.

We know that P_{td} and P_Φ are random polynomials of degrees c and $3c'$, respectively. It is also clear that, from the representation of I_1 as a matrix, then the transformation $(I_1, z) \mapsto z_{I_1}$ and, therefore, step 2 of $\mathfrak{P}_{\text{crs}, \text{td}, x}$, can be described using a fixed 2-degree polynomial. We conclude that $\mathfrak{P}_{\text{crs}, \text{td}, x}$ can be described as a distribution over $6cc'$ -degree polynomials.

It remains to show that \mathfrak{P} probabilistically computes $f_{\text{crs}, \text{td}, x}$. From the correctness of \mathfrak{P}_{td} and following Definition 3.7, if $\tilde{\pi}'_1 = \text{Extract}_{\text{td}}(a_{I_1}) \in \{0, 1\}^*$, then

$$\forall i \in I_1, \Pr[\tilde{\pi}'_i \neq \tilde{z}_i] \leq \epsilon$$

Applying union bound on the above, we get that $\Pr[\tilde{\pi}' \neq \tilde{z}] \leq |I_1| \cdot \epsilon \leq \ell \cdot \epsilon$.

Now, conditioning on $\tilde{\pi}' = \tilde{z}$, and from the correctness of \mathfrak{P}_Φ , we get that $\Pr[b \neq \Phi_{\text{crs}, x, 1}(\tilde{\pi}'_1)] \leq 2^{-c'}$ and, therefore, overall, we get that

$$\begin{aligned} & \Pr [P(a) \neq \Phi_{\text{crs}, x, 1}(\tilde{\pi}'_1)] \\ & \stackrel{P \xleftarrow{\$} \mathfrak{P}_{\text{crs}, \text{td}, x}}{\leq} \Pr[\tilde{\pi}' \neq \tilde{z}] + \Pr[P(a) \neq \Phi_{\text{crs}, x, 1}(\tilde{\pi}'_1) \mid \tilde{\pi}' = \tilde{z}] \\ & \leq \ell \cdot \epsilon + 2^{-c'} \end{aligned} \tag{3}$$

Now, if $f_{\text{crs}, \text{td}, x}(a) = 1$, then it must be the case that $\tilde{\pi}'_1 \in \{0, 1\}^*$ and $\Phi_{\text{crs}, x, 1}(\tilde{\pi}'_1) = 1$, and therefore, from (3), $P(a) = 1$ with the required probability. Otherwise, if $f_{\text{crs}, \text{td}, x}(a) = 0$, then $\tilde{\pi}'_0 \in \{0, 1\}^*$ and $\Phi_{\text{crs}, x, 0}(\tilde{\pi}'_0) = 1$. Since $\tilde{\Pi}^{\text{Com}}$ is a Σ -protocol and $\mathcal{R}_\Sigma(\tilde{\Pi}^{\text{Com}})$ is unique output (Lemma 3.12), then there exist no $\tilde{z}_1 \in \{0, 1\}^*$ such that $\Phi_{\text{crs}, x, 1}(\tilde{z}_1) = 1$ and, therefore, $P(a) = 0$ with the required probability. This completes the proof. \square

Combining Lemmas 3.12 and 3.14, we have so far proven Theorem 3.10 for the special case of $m = 1$. To derive the theorem for the general case, consider the protocol $\tilde{\Pi}^{\text{Com}}$ that consists of m parallel repetitions of $\tilde{\Pi}_1^{\text{Com}}$. Parallel repetition preserves honest-verifier zero knowledge and the Σ -protocol property (\mathcal{R}_Σ being unique-output) and, consequently, amplifies soundness to 2^{-m} . Further, if $\tilde{\Pi}_1^{\text{Com}}$ is a trapdoor Σ -protocol with `tdSetup` and `BadChallenge`, then $\tilde{\Pi}^{\text{Com}}$ is a trapdoor Σ -protocol with `tdSetup` and `BadChallengem`, where `BadChallengem`(`crs`, `td`, `x`, a_1, \dots, a_m) computes $e_i = \text{BadChallenge}(\text{crs}, \text{td}, x, a_i)$ for all $i \in [m]$ then outputs (e_1, \dots, e_m) if $\forall i e_i \in \{0, 1\}$ and outputs \perp otherwise. By Definition 3.7, if `BadChallenge` has ϵ' -probabilistic $6cc'$ -degree representation, then so does `BadChallengem`.

Hence, the proof of Theorem 3.10 is complete.

4 CI Through Probabilistic Representation

In this section, we show that if a function class \mathcal{F} has a probabilistic representation by a potentially simpler class \mathcal{C} (see Definition 3.7) then CI for relations searchable by \mathcal{F} can be reduced to CI for a class of relations that are “approximated” by \mathcal{C} . This is the first step we make towards constructing CI hash, as required by Corollary 3.11, from standard assumptions.

4.1 Approximable Relations and CI-Apx

We start by defining the notion of *approximable relations* and a related special case of correlation intractability, CI-Apx.

Definition 4.1 (CI-Apx). Let $\mathcal{C} = \{C_\lambda : \{0, 1\}^{n(\lambda)} \rightarrow \{0, 1\}^{m(\lambda)}\}$ be a function class and let $0 < \epsilon < 1$. For every $C \in \mathcal{C}$, we define the relation ϵ -approximable by C as follows

$$\mathcal{R}_C^\epsilon = \{(x, y) \in \{0, 1\}^n \times \{0, 1\}^m \mid \Delta(y, C(x)) \leq \epsilon m\}$$

A hash family that is CI for all relations $\{\mathcal{R}_C^\epsilon \mid C \in \mathcal{C}\}$ is said to be CI-Apx $_\epsilon$ for \mathcal{C} .

4.2 From CI-Apx for \mathcal{C} to CI for \mathcal{F}

We now state and prove the following general theorem.

Theorem 4.2. Let \mathcal{F} be a function class that has an ϵ -probabilistic representation by \mathcal{C} . If \mathcal{H} is CI-Apx $_{2\epsilon}$ hash for \mathcal{C} , then \mathcal{H} is CI for relations searchable by \mathcal{F} (i.e. ϵ -probabilistically searchable by \mathcal{C}).

Proof of Theorem 4.2. Suppose \mathcal{R} is searchable by $\mathcal{F} : \{0, 1\}^n \rightarrow \{0, 1\}^m$. Fix some $R \in \mathcal{R}$ and consider its corresponding search function $f \in \mathcal{F}$. Let \mathcal{C}_f be the ϵ -probabilistic representation of f by \mathcal{C} .

We start by defining a game $\text{Game}_0(\mathcal{A})$ against an adversary \mathcal{A} as follows.

$\text{Game}_0(\mathcal{A})$:

1. $k \xleftarrow{\$} \text{Sample}(1^\lambda)$.
2. $x \leftarrow \mathcal{A}(k)$.
3. Output 1 if and only if $f(x) \neq \perp$ and $\text{Hash}_k(x) = f(x)$.

It is clear that the probability of an adversary \mathcal{A} to win Game_0 upper bounds the probability he breaks the correlation intractability of \mathcal{H} for R (immediate from Definition 3.2). Our goal, then, is to show that for any PPT adversary \mathcal{A} , there exists a negligible function negl such that $\Pr[\text{Game}_0(\mathcal{A}) = 1] < \text{negl}(\lambda)$.

We now reduce Game_0 to Game_1 , which is defined below.

$\text{Game}_1(\mathcal{A})$:

1. $C \xleftarrow{\$} \mathcal{C}_f$.
2. $k \xleftarrow{\$} \text{Sample}(1^\lambda)$.
3. $x \leftarrow \mathcal{A}(k)$.
4. Output 1 if and only if $\Delta(\text{Hash}_k(x), C(x)) \leq 2\epsilon m$.

Lemma 4.3. *For any (possibly unbounded) adversary \mathcal{A} , there exists a negligible function negl , such that*

$$\Pr[\text{Game}_0(\mathcal{A}) = 1] \leq \Pr[\text{Game}_1(\mathcal{A}) = 1] + \text{negl}(\lambda)$$

Proof. The proof is derived from the fact that C in Game_1 is sampled independently of the adversary’s choice x and from Lemma 3.8, as follows.

$$\begin{aligned} \Pr[\text{Game}_0(\mathcal{A}) = 1] &= \Pr[f(x) \neq \perp \wedge \text{Hash}_k(x) = f(x)] \\ &\leq \Pr_{C \xleftarrow{\$} \mathcal{C}_f} [f(x) \neq \perp \wedge \Delta(f(x), C(x)) > 2\epsilon m] \\ &\quad + \Pr_{C \xleftarrow{\$} \mathcal{C}_f} [f(x) \neq \perp \wedge \Delta(\text{Hash}_k(x), f(x)) \leq 2\epsilon m] \\ &\leq \Pr[\text{Game}_1(\mathcal{A}) = 1] + \text{negl}(\lambda) \end{aligned}$$

□

To complete the proof of Theorem 4.2, we show that Game_1 is hard to win with non-negligible probability, based on the correlation intractability of \mathcal{H} for relations 2ϵ -approximable \mathcal{C} .

Lemma 4.4. *If \mathcal{H} is CI-Apx $_{2\epsilon}$ for \mathcal{C} then, for any $f \in \mathcal{F}$ and any PPT adversary \mathcal{A} , there exists a negligible function such that*

$$\Pr[\text{Game}_1(\mathcal{A}) = 1] < \text{negl}(\lambda)$$

Proof. Assume towards contradiction there exists $f \in \mathcal{F}$ and \mathcal{A} for which the above does not hold, namely $\Pr[\text{Game}_1(\mathcal{A}) = 1] > 1/\text{poly}(\lambda)$. Then, there exists some fixed $C \in \mathcal{C}_f$ such that $\Pr[\text{Game}_1^C(\mathcal{A}) = 1] > 1/\text{poly}(\lambda)$, where Game_1^C is defined as Game_1 with C being fixed (rather than sampled from \mathcal{C}_f). From definition, such an adversary breaks the CI-Apx $_{2\epsilon}$ of \mathcal{H} for C . \square

We conclude the proof of the theorem by combining Lemmas 4.3 and 4.4.

5 CI-Apx from Trapdoor Hash

Having shown in the previous section that CI-Apx is a useful notion to obtain CI for a function class that has a simple probabilistic representation, we now show how to construct, from rate-1 trapdoor hash for any function class \mathcal{C} [14], an CI-Apx hash for \mathcal{C} . In fact, in our proof of CI, we require that the underlying TDH scheme satisfies the following stronger notion of correctness.

Definition 5.1. (Enhanced Correctness for TDH). *We say that a (rate-1) trapdoor hash scheme TDH for $\mathcal{C} = \{\mathcal{C}_n : \{0, 1\}^n \rightarrow \{0, 1\}\}$ has enhanced $(1 - \tau)$ -correctness for $\tau := \tau(\lambda) < 1$ if it satisfies the following property:*

- **Enhanced Correctness:** *There exists a negligible function $\text{negl}(\lambda)$ such that the following holds for any $\lambda, n, \in \mathbb{N}$, any $\mathbf{h} \in \{0, 1\}^{\eta(\lambda)}$, any $\text{hk} \in \mathcal{S}(1^\lambda, 1^n)$, and any function $C \in \mathcal{C}_n$:*

$$\Pr[\forall x \in \{0, 1\}^n : \text{H}(\text{hk}, x) = \mathbf{h}, \mathbf{e} + \mathbf{e}' = C(x) \pmod 2] \geq 1 - \tau - \text{negl}(\lambda)$$

where $(\text{ek}, \text{td}) \leftarrow \text{G}(\text{hk}, C)$, $\mathbf{e} = \text{E}(\text{ek}, x)$, $\mathbf{e}' = \text{D}(\text{td}, \mathbf{h})$ and the probability is over the randomness used by G .

Theorem 5.2. *Assume there exists rate-1 trapdoor hash scheme TDH for $\mathcal{C} = \{\mathcal{C}_n : \{0, 1\}^n \rightarrow \{0, 1\}^m\}$ with enhanced $(1 - \tau)$ -correctness where the hash length is $\eta := \eta(\lambda)$. Then, for any ϵ s.t. $\epsilon + \tau < \epsilon_0$ (for some fixed universal constant ϵ_0), there exists a polynomial $m_{\epsilon, \eta, \tau}(\lambda) = O((\eta + \lambda)/\tau + \log(1/\epsilon))$ such that, for every polynomial $m > m_\epsilon$, there exists a CI-Apx $_\epsilon$ hash family for \mathcal{C} with output length $m(\lambda)$.⁶*

Recalling Corollary 3.11, and using the result from Section 4, obtaining CI-Apx for constant-degree functions is sufficient for our purpose of constructing NIZK. To instantiate Theorem 5.2 for constant-degree functions from standard assumption, we use the following result of Döttling et al. [14].

Theorem 5.3. (TDH from Standard Assumptions [14]). *For any constant $c \in \mathbb{N}$ and arbitrarily small $\tau := \tau(\lambda) = 1/\text{poly}(\lambda)$, there exists a rate-1 trapdoor hash scheme, for c -degree polynomials over \mathbb{Z}_2 , with enhanced $(1 - \tau)$ -correctness and function privacy under the DDH/QR/DCR/LWE assumption⁷.*

⁶ In fact, as implicitly implied by the proof of the theorem, our construction satisfies the stronger notion of somewhere statistical CI for the corresponding hamming-ball relations [9]. However, applying Theorem 4.2 on the construction does not preserve this property.

⁷ The error probability in the QR, DCR, and LWE constructions is even negligible.

We note some gaps between the result from [14] and the theorem above. First, the aforementioned work considers only linear functions (i.e. degree-1 polynomials) over \mathbb{Z}_2 . Second, their DDH-based construction supports even a stricter class of functions, namely only “index functions” of the form $f_i(x) = x_i$. Third, all known constructions are not proven to have enhanced correctness. In the full version [8], we show how to close these gaps by simple adjustments to the constructions and proofs from [14]. Combining Theorems 5.2 and 5.3, we obtain.

Corollary 5.4. *Let $c \in \mathbb{N}$. There exists a constant $\epsilon > 0$ such that, for any sufficiently large polynomial $m := m(\lambda)$, there exists a programmable correlation intractable hash family with output length m for all relations ϵ -approximable by c -degree polynomials over \mathbb{Z}_2 .*

5.1 The Hash Family

We now present our construction of CI-Apx from rate-1 TDH. We note that we do not use the full power of a TDH. Specifically, the decoding algorithm need not be efficient and, further, we do not use input privacy (as defined in [14]).

Construction 5.1 (Correlation Intractability from TDH). *Let $n := n(\lambda)$ and $m := m(\lambda)$ be polynomials in the security parameter, and let $\epsilon := \epsilon(\lambda) < 0.32$. Let $\mathcal{C} : \{0, 1\}^n \rightarrow \{0, 1\}$ be a function class and let $\text{TDH} = (\text{S}, \text{G}, \text{H}, \text{E}, \text{D})$ be a rate-1 trapdoor hash scheme for \mathcal{C} . Our construction of CI-Apx $_\epsilon$ hash for \mathcal{C} consists of the following algorithms.*

- **Sample(1^λ):** *Sample $\text{hk} \xleftarrow{\$} \text{S}(1^\lambda)$ and, for all $i \in [m]$, $(\text{ek}_i, \text{td}_i) \xleftarrow{\$} \text{G}(\text{hk}, C_0)$ for an arbitrary fixed $C_0 \in \mathcal{C}$, and a uniform $r \xleftarrow{\$} \{0, 1\}^m$, then output*

$$\mathbf{k} = ((\text{ek}_1, \dots, \text{ek}_m), r)$$

- **Hash(\mathbf{k}, x):** *The hash of an input $x \in \{0, 1\}^n$ under key $\mathbf{k} = ((\text{ek}_i)_{i \in [m]}, r)$ is computed as follows*

$$\mathbf{h} = \text{E}((\text{ek}_1, \dots, \text{ek}_m), x) + r \pmod 2$$

5.2 Proof of Theorem 5.2

Programmability of the construction is trivial and, thus, we focus on proving CI.

Fix some $C = (C_1, \dots, C_m) \in \mathcal{C}^m$ and consider the relation ϵ -probabilistically searchable by C , R_C^ϵ . The advantage of an adversary \mathcal{A} in breaking the CI for R_C^ϵ is demonstrated in his advantage in winning in the following game.

Game $_0(\mathcal{A})$:

1. $\mathbf{k} \xleftarrow{\$} \text{Sample}(1^\lambda)$.
2. $x \leftarrow \mathcal{A}(\mathbf{k})$.
3. Output 1 if and only if $\Delta(\text{Hash}_\mathbf{k}(x), C(x)) \leq 2\epsilon m$.

To show $\Pr[\text{Game}_0(\mathcal{A}) = 1] < \text{negl}(\lambda)$, we define a different game, Game_1 , in which we switch the encoding keys (ek_1, \dots, ek_m) in k to encoding keys corresponding to the functions C_1, \dots, C_m (rather than C_0).

$\text{Game}_1(\mathcal{A})$:

1. Sample $hk \leftarrow S(1^\lambda, 1^n)$ and $(ek'_i, td'_i) \leftarrow G(hk, C_i)$ for every $i \in [m]$.
 Sample a uniform $r \xleftarrow{\$} \{0, 1\}^m$, then set $k = ((ek'_1, \dots, ek'_m), r)$.
2. $x \leftarrow \mathcal{A}(k)$.
3. Output 1 if and only if $\Delta(\text{Hash}_k(x), C(x)) \leq 2\epsilon m$.

We claim that, based on the function privacy of the underlying trapdoor hash, we may reduce Game_0 to Game_1 .

Lemma 5.5. *Under the function privacy of TDH, for any PPT adversary \mathcal{A} , there exists a negligible function negl such that*

$$\Pr[\text{Game}_0(\mathcal{A}) = 1] \leq \Pr[\text{Game}_1(\mathcal{A}) = 1] + \text{negl}(\lambda)$$

Proof. Assume towards contradiction there exists an adversary \mathcal{A} for which the above does not hold.

We use \mathcal{A} to construct an adversary \mathcal{A}_{TDH} that distinguishes between $(hk, (ek_1, \dots, ek_m))$ and $(hk, (ek'_1, \dots, ek'_m))$, where $hk \leftarrow S(1^\lambda, 1^n)$, $ek_i \leftarrow G(hk, C_0)$ and $ek'_i \leftarrow G(hk, C_{f_i})$ (for every $i \in [m]$), with non-negligible advantage. Such an adversary breaks the function privacy of TDH via a standard hybrid argument.

On input (ek_1, \dots, ek_m, C) , \mathcal{A}_{TDH} simply calls $x \leftarrow \mathcal{A}((ek_1, \dots, ek_m), r)$, and outputs 1 iff $\Delta(\text{Hash}_k(x), C(x)) \leq 2\epsilon m$. It holds that

$$\begin{aligned} & |\Pr[\mathcal{A}_{\text{TDH}}(hk, (ek_1, \dots, ek_m)) = 1] - \Pr[\mathcal{A}_{\text{TDH}}(hk, (ek'_1, \dots, ek'_m)) = 1]| \\ &= |\Pr[\text{Game}_1(\mathcal{A}) = 1] - \Pr[\text{Game}_2(\mathcal{A}) = 1]| \geq 1/\text{poly}(\lambda) \end{aligned} \quad \square$$

Lastly, we show that Game_1 is statistically hard to win. This, together with Lemma 5.6, implies Theorem 5.2.

Lemma 5.6. *For any (possibly unbounded) adversary \mathcal{A} , there exists a negligible function negl s.t.*

$$\Pr[\text{Game}_1(\mathcal{A}) = 1] < \text{negl}(\lambda)$$

Proof. It suffices to show that there exists a negligible function negl such that

$$\Pr_k[\exists x : \Delta(\text{Hash}_k(x), C(x)) \leq 2\epsilon m] < \text{negl}(\lambda)$$

where k is sampled as in Game_1 . We denote the above event by Bad and observe that

$$\Pr[\text{Bad}] = \Pr_k[\exists x, z \in \{0, 1\}^m : |z| \leq 2\epsilon m \wedge C(x) + z = \text{Hash}_k(x) \pmod 2]$$

For any $hk \in S(1^\lambda, 1^n)$, let Bad_{hk} be the event Bad where the hash key is fixed to hk , and the probability space is over random (ek_i, td_i) and r . It is sufficient, then, to show that for all $hk \in S(1^\lambda, 1^n)$, $\Pr[\text{Bad}_{hk}] \leq \text{negl}(\lambda)$.

For any $hk \in S(1^\lambda, 1^n)$, let $TDHCor_{hk}$ denote the following event:

$$TDHCor_{hk} = [\forall x, \Delta(E(ek, x) + C(x), D(td, h)) \leq 2\tau m].$$

Then $\Pr[Bad_{hk}] \leq \Pr[\neg TDHCor_{hk}] + \Pr[Bad_{hk} \wedge TDHCor_{hk}]$. We will separately show that both $\Pr[\neg TDHCor_{hk}]$ and $\Pr[Bad_{hk} \wedge TDHCor_{hk}]$ are negligible in λ .

First, we bound $\Pr[\neg TDHCor_{hk}]$ based on the enhanced $(1 - \tau)$ -correctness of TDH and Chernoff bound: for every fixed $h \in \{0, 1\}^\eta$ and $hk \in S(1^\lambda, 1^n)$,

$$\Pr[\exists x : H(hk, x) = h, \quad \Delta(E(ek, x) + C(x), D(td, h)) > 2\tau m] \leq e^{-\tau m/3}$$

Applying union bound over all $h \in \{0, 1\}^\eta$ gives

$$\Pr[\neg TDHCor_{hk}] = \Pr[\exists x, \quad \Delta(E(ek, x) + C(x), D(td, H(hk, x))) > 2\tau m] < e^{\eta - \tau m} = \text{negl}(\lambda).$$

Second, note that $\Pr[Bad_{hk} \wedge TDHCor_{hk}] \leq \Pr[\exists x : \Delta(r, D(td, H(hk, x))) \leq 2(\tau + \epsilon)]$ where the probability is over choice of td and r . Let $\epsilon' = 2(\epsilon + \tau)$ and (for fixed hk, td) let

$$Y = \{D(td, h_x) + z' \pmod 2 \mid x \in \{0, 1\}^n, h_x = H(hk, x), z' \in \{0, 1\}^m \text{ s.t. } |z'| \leq \epsilon' m\}$$

For fixed hk, td , $\Pr_r[\exists x : \Delta(r, D(td, h_x)) < \epsilon' m] = 2^{-m}|Y|$. Thus, it suffices to show that $2^{-m}|Y|$ is negligible. Clearly, $|\{D(td, h_x) : x \in \{0, 1\}^n, h_x = H(hk, x)\}| \leq 2^\eta$. Further, we can bound

$$|\{z' \in \{0, 1\}^m \mid |z'| \leq \epsilon' m\}| = \sum_{i=1}^{\epsilon' m} \binom{m}{i} \leq \sum_{i=1}^{\epsilon' m} \left(\frac{me}{i}\right)^i \leq (e/\epsilon')^{\epsilon' m + 1}$$

and consequently, $|Y| \leq 2^\eta \cdot (e/\epsilon')^{\epsilon' m + 1}$. If ϵ' is a (universally) sufficiently small constant, and $m \geq (\lambda + \eta + \log(e/\epsilon')) / (1 - \epsilon' \log(e/\epsilon')) = O((\eta + \lambda) / \tau + \log(1/\epsilon))$,

$$2^{-m}|Y| \leq 2^{-m}(e/\epsilon)^{\epsilon' m + 1} 2^\eta < 2^{(\epsilon' \log(e/\epsilon') - 1)m + \log(e/\epsilon') + \eta} < 2^{-\lambda}.$$

□

Acknowledgments. We thank Brent Waters for pointing out an inaccuracy in a proof in a previous version of this work, and for other comments. We thank the anonymous reviewers of CRYPTO 2020 for their helpful feedback and pointers.

Research supported by the Binational Science Foundation (Grant No. 2016726), and by the European Union Horizon 2020 Research and Innovation Program via ERC Project REACT (Grant 756482) and via Project PROMETHEUS (Grant 780701). Venkata Koppula was also supported by the Simons-Berkeley Fellowship. This work was done in part while the author was visiting the Simons Institute for the Theory of Computing.

References

1. Alekhnovich, M.: More on average case vs approximation complexity. In: 44th Symposium on Foundations of Computer Science (FOCS 2003), Proceedings, Cambridge, MA, USA, 11–14 October 2003, pp. 298–307. IEEE Computer Society (2003). <https://doi.org/10.1109/SFCS.2003.1238204>
2. Badrinarayan, S., Fernando, R., Jain, A., Khurana, D., Sahai, A.: Statistical zap arguments. Cryptology ePrint Archive, Report 2019/780 (2019)
3. Bellare, M., Micciancio, D., Warinschi, B.: Foundations of group signatures: formal definitions, simplified requirements, and a construction based on general assumptions. In: Biham, E. (ed.) EUROCRYPT 2003. LNCS, vol. 2656, pp. 614–629. Springer, Heidelberg (2003). https://doi.org/10.1007/3-540-39200-9_38
4. Ben-Sasson, E., et al.: Zerocash: decentralized anonymous payments from bitcoin. In: 2014 IEEE Symposium on Security and Privacy, pp. 459–474, May 2014
5. Bender, A., Katz, J., Morselli, R.: Ring signatures: stronger definitions, and constructions without random oracles. In: Halevi, S., Rabin, T. (eds.) TCC 2006. LNCS, vol. 3876, pp. 60–79. Springer, Heidelberg (2006). https://doi.org/10.1007/11681878_4
6. Blum, M.: How to prove a theorem so no one else can claim it. In: Proceedings of the International Congress of Mathematicians, pp. 1444–1451 (1987)
7. Blum, M., Feldman, P., Micali, S.: Non-interactive zero-knowledge and its applications (extended abstract). In: 20th Annual ACM Symposium on Theory of Computing, 2–4 May 1988, pp. 103–112. ACM Press, Chicago (1988)
8. Brakerski, Z., Koppula, V., Mour, T.: NIZK from LPN and trapdoor hash via correlation intractability for approximable relations. Cryptology ePrint Archive, Report 2020/258 (2020). <https://eprint.iacr.org/2020/258>
9. Canetti, R., et al.: Fiat-Shamir: from practice to theory. In: Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, pp. 1082–1090. Association for Computing Machinery, New York (2019). <https://doi.org/10.1145/3313276.3316380>
10. Canetti, R., Chen, Y., Reyzin, L., Rothblum, R.D.: Fiat-Shamir and correlation intractability from strong KDM-secure encryption. In: Nielsen, J.B., Rijmen, V. (eds.) EUROCRYPT 2018, Part I. LNCS, vol. 10820, pp. 91–122. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-78381-9_4
11. Canetti, R., Goldreich, O., Halevi, S.: The random oracle methodology, revisited. *J. ACM* **51**(4), 557–594 (2004). <https://doi.org/10.1145/1008731.1008734>
12. Damgård, I., Park, S.: How practical is public-key encryption based on LPN and ring-LPN? Cryptology ePrint Archive, Report 2012/699 (2012). <http://eprint.iacr.org/2012/699>
13. Dolev, D., Dwork, C., Naor, M.: Non-malleable cryptography (extended abstract). In: 23rd Annual ACM Symposium on Theory of Computing, 6–8 May 1991, pp. 542–552. ACM Press, New Orleans (1991)
14. Döttling, N., Garg, S., Ishai, Y., Malavolta, G., Mour, T., Ostrovsky, R.: Trapdoor hash functions and their applications. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019. LNCS, vol. 11694, pp. 3–32. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-26954-8_1
15. Feige, U., Lapidot, D., Shamir, A.: Multiple noninteractive zero knowledge proofs under general assumptions. *SIAM J. Comput.* **29**(1), 1–28 (1999). <https://doi.org/10.1137/S0097539792230010>

16. Fiat, A., Shamir, A.: How to prove yourself: practical solutions to identification and signature problems. In: Odlyzko, A.M. (ed.) CRYPTO 1986. LNCS, vol. 263, pp. 186–194. Springer, Heidelberg (1987). https://doi.org/10.1007/3-540-47721-7_12
17. Goldwasser, S., Micali, S., Rackoff, C.: The knowledge complexity of interactive proof-systems (extended abstract). In: 17th Annual ACM Symposium on Theory of Computing, pp. 291–304, 6–8 May 1985. ACM Press, Providence (1995)
18. Groth, J., Ostrovsky, R., Sahai, A.: New techniques for noninteractive zero-knowledge. *J. ACM* **59**(3) (2012). <https://doi.org/10.1145/2220357.2220358>
19. Holmgren, J., Lombardi, A.: Cryptographic hashing from strong one-way functions (or: one-way product functions and their applications), pp. 850–858, October 2018. <https://doi.org/10.1109/FOCS.2018.00085>
20. Jain, A., Jin, Z.: Statistical zap arguments from quasi-polynomial LWE. *Cryptology ePrint Archive*, Report 2019/839 (2019)
21. Kalai, Y.T., Rothblum, G.N., Rothblum, R.D.: From obfuscation to the security of Fiat-Shamir for proofs. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017, Part II. LNCS, vol. 10402, pp. 224–251. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-63715-0_8
22. Lombardi, A., Vaikuntanathan, V., Wichs, D.: 2-message publicly verifiable WI from (subexponential) LWE. *Cryptology ePrint Archive*, Report 2019/808 (2019). <https://eprint.iacr.org/2019/808>
23. Naor, M., Yung, M.: Public-key cryptosystems provably secure against chosen ciphertext attacks. In: 22nd Annual ACM Symposium on Theory of Computing, 14–16 May 1990, pp. 427–437. ACM Press, Baltimore (1990)
24. Peikert, C., Shiehian, S.: Noninteractive zero knowledge for NP from (plain) learning with errors. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019. LNCS, vol. 11692, pp. 89–114. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-26948-7_4
25. Sahai, A., Waters, B.: How to use indistinguishability obfuscation: deniable encryption, and more. In: Shmoys, D.B. (ed.) 46th Annual ACM Symposium on Theory of Computing, 31 May – 3 June 2014, pp. 475–484. ACM Press, New York (2014)