On Instance Compression, Schnorr/Guillou-Quisquater, and the Security of Classic Protocols for Unique Witness Relations*

Yi Deng[1,2], Xuyang Song[1], Jingyue Yu[1] and Yu Chen[1,2]

[1] SKLOIS, Institute of Information Engineering, CAS, P.R.China
[2] State Key Laboratory of Cryptology, P. O. Box 5159, Beijing ,100878,China
{deng, songxuyang, yujingyue, chenyu}@iie.ac.cn

Abstract. We revisit the problem of whether the witness hiding property of classic 3-round public-coin proof systems for languages/distributions with unique witnesses are still witness hiding. Though strong black-box *impossibility* results are known for them [Pas11, HRS09], we provide some less unexpected *positive* results on the witness hiding security of classic protocols:

- We develop an embedding technique and prove that the witness hiding property of the standalone Schnorr protocol based on a *weaker* version of one-more like discrete logarithm (DL) assumption asserting that, for an arbitrary constant ℓ, it is infeasible for a PPT algorithm to solve l DL instances with being restricted to query the DL oracle *only once*. Similar result holds for the Guillou-Quisquater protocol.

  This improves over the positive result of [BP02] in that when applying their technique to the standalone setting, the underlying assumption is stronger and required to hold only for $\ell = 2$.
- Following the framework of [HN10], we introduce the notion of *tailored instance compression* to capture the essence of the known one-more like assumptions, which provides new insight into the hardness of one-more DL/RSA problems and allows us to reveal some strong consequences of breaking our weaker version of one-more like assumption, including zero knowledge protocols for the AND-DL and AND-RSA languages with extremely efficient communication and non-trivial hash combiner for hash functions based on DL problem.

  These consequences can be viewed as positive evidences for the security of Schnorr and Guillou-Quisquater protocols.
- We observe that the previously known impossibility results on the witness hiding of public-coin protocols for unique witness relation make certain *restriction* on the reduction. By introducing an input-distribution-switching technique, we bypass these known impossibility results and prove that, for any hard language L, if a distribution $(\mathbb{X}, \mathbb{W})$ over *unique* witness relation R_L has an indistinguishable counterpart distribution over some *multiple* witnesses relation, then any witness indistinguishable protocols (including ZAPs and all known 3-round public-coin protocols, such as Blum protocol and GMW protocol) are indeed witness hiding for the distribution $(\mathbb{X}, \mathbb{W})$. We also show a wide range of cryptographic problems with unique witnesses satisfy the "if condition" of this result, and thus admit constant-round public-coin witness hiding proof system.

  This is the *first* positive result on the witness-hiding property of the classic protocols for non-trivial unique witness relations.

* Supported by the National Natural Science Foundation of China (Grant No. 61379141), and the Open Project Program of the State Key Laboratory of Cryptology.

1 Introduction

Witness hiding proof system, introduced by Feige and Shamir [FS90], is a relaxed yet natural notion of zero knowledge proof [GMR89]. Instead of requiring an efficient simulation for the view of the verifier as in zero knowledge proof, witness hiding property only requires that, roughly speaking, the interaction with honest prover does not help the verifier compute any new witness for the statement being proven that he did not know before. One immediate application of such a security notion is identification: Witness hiding proof allows a prover to prove his identity without leaking the associated secret key, and this security notion is sufficient for preventing impersonation attack from malicious verifiers.

The witness hiding property of some practical protocols, which are usually not zero knowledge, is often being proved via another beautiful and widely applicable notion of witness indistinguishability introduced in the same paper of [FS90]. A witness indistinguishable proof guarantees that if the statement has two independent witnesses, then the malicious verifier cannot tell which witness is being used by the prover in an execution of the protocol. The idea underlying the security proof of witness hiding via witness indistinguishability is as follows. Suppose that for a hard language, each instance has two witnesses and it is infeasible for an efficient algorithm, given one witness as input, to compute the other one, then the witness indistinguishable protocol is actually witness hiding with respect to such instances. This is because we can take one witness as input to play the role of honest prover and then use the verifier's ability of breaking witness hiding to either break witness indistinguishability of this protocol or obtain a new witness. Therefore, the parallelized version of 3-round public-coin classic protocols of [Blu86, GMW91] are witness hiding with respect to such languages.

What happens if the hard language consists of instances that have exactly one witness? This problem has turned out to be quite subtle. The Guillou-Quisquater [GQ88] and the Schnorr [Sch89] identification protocols are perhaps the best-known efficient protocols for unique witness relations, but their security has long remained open. On the positive side, Shoup [Sho97] presented positive result that the Schnorr identification protocol is secure in the generic group model, and Bellare and Palacio [BP02] showed that the security of the Guillou-Quisquater and Schnorr identification protocols can be based on the so-called one-more RSA and one-more discrete logarithm assumptions, respectively [BP02, BNPS03]. These security proofs of course imply that the Schnorr and the Guillou-Quisquater identification protocols are witness hiding in the standalone setting where there is only a single execution of the protocol. However, the underlying assumptions/models are quite strong and non-standard.

Indeed, there is an obstacle in the way of basing constant-round public-coin protocols for unique witness relations on standard assumption. As mentioned before, the basic approach to prove witness hiding of a protocol is to find an efficient way to exploit the power of the malicious verifier to break some hardness assumptions. For the instance that has exact one witness, however, to exploit the power of the malicious verifier requires the reduction itself to know the unique witness to the statement being proven in the first place (by the soundness property of the protocol), which usually does not lead to a desired contradiction even if the malicious verifier does have the ability to break witness hiding of the protocol.

Haitner, Rosen and Shaltiel [HRS09] gave the first proof that constant-round public-coin witness hiding protocols for unique witness relations cannot be based on standard assumptions via some restricted types of black-box reductions. Pass [Pas11] showed that if we further require witness hiding to hold under sequential repetition, then we can significantly strengthen the impossibility result of [HRS09]. Some similar impossibility results on the problem whether we can base the aforementioned one-more discrete logarithm assumption on standard hardness assumption were also given in [Pas11] and [ZZC+14]. We would like

to point out that these impossibility results may have some impact on other important problems. For example, in [Pas06] Pass showed a deep connection between the problem of whether the classic constant-round public-coin proofs are witness hiding for all NP languages and the longstanding problem whether we can base one-way functions on NP-complete problem.

1.1 Our Contribution

Our main contribution is an optimistic point of view on the witness hiding security of the classic public-coin proof systems.

We develop an embedding technique and prove that the witness hiding property of the standalone Schnorr (Guillou-Quisquater) protocol based on a version of one-more like DL (RSA, respectively) assumption that significantly weaker than the assumed in the proofs of [BP02]. To see the plausibility of our still-non-standard assumption, we follow the framework of [HN10] and introduce the notion of *tailored instance compression*, which captures the essence of the known one-more like assumptions, and more importantly, provides new insight into the hardness of one-more DL/RSA problems and allows us to reveal some surprising consequences of breaking our weaker version of one-more like assumptions, including zero knowledge proofs with extremely low communication complexity for the AND-DL and AND-RSA languages and non-trivial hash combiner for hash functions based on DL problem.

We observe that all previously known impossibility results [Pas11, HRS09] on the witness hiding of public-coin protocols make an *implicit* restriction (which has not been mentioned explicitly in the statements of their main results) on the black-box reduction: For a distribution $(\mathbb{X}, \mathbb{W})$ on an *unique* witness relation, for the proof of lower bound to go through, the (black-box) reduction R is restricted to invoke the adversary verifier V^* *only* on instances in $\mathbb{X}$[1].

This leaves a problem of whether one can get around these impossibility results by removing the above restriction on the black-box reduction. We provide a positive answer to this problem. Specifically, we develop an input-distribution-switching technique and prove that, for any hard language L, if a distribution $(\mathbb{X}, \mathbb{W})$ on a *unique* witness relation R_L has an indistinguishable counterpart distribution over some *multiple* witnesses relation, then any witness indistinguishable protocols (including ZAPs and all known 3-round public-coin protocols, such as Blum protocol and GMW protocol) are indeed witness hiding for the unique witness distribution $(\mathbb{X}, \mathbb{W})$. We also show a wide range of cryptographic problems with unique witnesses satisfy the "if condition" of this result, and thus admit constant-round public-coin witness hiding proof system. This is the *first* positive result on the witness-hiding property of the classic protocols for unique witness relations.

We summarize our results in the table 1. Detailed explanations follow.

Embedding technique and the instance compression problem. Before proceeding to our embedding reduction, we recall the Schnorr protocol and Bellare and Palacio's security proof for it [BP02]. Let $\mathbb{G}$ be a group of prime order q generated by g, the prover P wants to convince the verifier V of knowledge of the discrete logarithm (unique witness) $w \in \mathbb{Z}_q$ of an element $y = g^w \in \mathbb{G}$. To do so, P first sends a random element $a = g^r \in \mathbb{G}$ to V, and upon receiving the V's challenge $c \in \mathbb{Z}_q$, it answers with a value $z \in \mathbb{Z}_q$. V accepts

[1] This restriction can be seen from the last paragraph "on the role of unique witness", page 7 of the full version (see http://www.cs.cornell.edu/ rafael/papers/schnorr.pdf) of [Pas11]:"...(in the reduction) If the statement x has a unique witness w, we can ensure that the extracted witness will be identical to the witness that the oracle A (which is V^* in our setting) would have returned.."

	Security of Schnorr/GQ	Instance Incompressibility /One-more Assumptions	WH of PC Protocols for unique witness R
BB Negative Results/Evidences	[Pas11]	[Pas11] [ZZC+14]	[HRS09] [Pas11]
Positive Results /Evidences	[BP02] This work (weaker assum.)	This work	This work

Table 1: Our results for languages with unique witnesses compared to previous work. Here we refer to the impossibility results of further basing instance incompressibility/one-more assumptions on standard hard problems as "BB negative results/evidences", and refer to the surprising consequences of breaking these assumptions as "positive results/evidences" in favor of these assumptions. As we observe, the impossibility results of [HRS09, Pas11] make an *implicit* restriction on the black-box reduction.

the proof if and only if $g^z = a \cdot y^c$. Note that, if V finally outputs the witness $w \in \mathbb{Z}_q$ at the end of interaction, then we can build an algorithm R solving two random discrete logarithm instances y and a at the same time if R is allowed to make one query to the discrete logarithm solver oracle $\mathcal{O}_{\text{dlog}}$: R have y serve as the common input and a as the first prover message, after receiving V's challenge c, R queries $\mathcal{O}_{\text{dlog}}$ on $a \cdot y^c$ and forwards the response z from the oracle to the verifier; when V outputs w, R can solve the linear equation $z = r + cw \bmod q$ and obtain r. This useful observation was also exploited by Bellare and Palacio [BP02] to prove the security of the Schnorr protocol as an identification scheme under the one-more discrete logarithm problem.

We now show how to conduct embedding reduction R that leads to better security proof based on a weaker assumption.

Suppose that we are given a set of discrete logarithm instances $(y_1, y_2, \ldots, y_\ell)$ to solve. For simplicity, we assume $\ell = 2^l$ for some integer l. The first part of R is a *compressing* process. R partitions them into $\ell/2$ pairs, for each pair of instances, one serving as the the common input and the other serving as the first prover message in a session, and invokes $\ell/2$ incarnations of the verifier in parallel. After collecting $\ell/2$ challenges from the $\ell/2$ invocations of the verifier, R has to solve $\ell/2$ new instances in order to answer each verifier. At this point, rather than querying $\mathcal{O}_{\text{dlog}}$ on these new instances, R pauses all these interactions and partitions the new $\ell/2$ instances into $\ell/4$ pairs, and then repeats the above step and invokes $\ell/4$ incarnations of the verifier in parallel, and will get $\ell/8$ new instances to solve. Continuing to repeat this, by viewing each partial interaction with a verifier as a node we get a tree in which each node takes in two instances and outputs one instance. Finally, R reaches the root and has only one instance to solve.

The second part of R is an *unfolding* process. R queries $\mathcal{O}_{\text{dlog}}$ on the root instance, then by using the verifier's power of breaking witness hiding as above, R is able to solve the two instances flowing into this node. Note that, the two instances R just solved will help it solve the four instances that flows into the two nodes at the level above the root (without making queries to oracle anymore), and repeating this process R will solve all these ℓ instances $(y_1, y_2, \ldots, y_\ell)$. Observe that in the entire embedding reduction, R makes only

a single query (at the root of the tree) to $\mathcal{O}_{\text{dlog}}$ and solves all ℓ DL instances. This process is exemplified in Figure 2.

The actual embedding reduction needs to make each invocation of the verifier independent by using the random self-reducibility of the discrete logarithm problem. As we will see, the quantity ℓ can be an arbitrarily large constant, or any polynomial when the verifier's success probability is close to 1. Thus, assuming that it is infeasible for a PPT oracle algorithm to solve ℓ discrete logarithm instances at the same time when restricted to making a *single* query to the discrete logarithm solver oracle, the standalone Schnorr protocol is witness hiding. Similar results can also be obtained for the Guillou-Quisquater's protocol and some other Σ-protocols for group homomorphisms.

This improves the positive result of [BP02] in that when applying the technique of [BP02] to the standalone setting, the above result requires the corresponding assumption to hold only with respect to the case of $\ell = 2$.

Our reduction R leads to the following *tailored instance compression* problem for DL: Construct a triplet of efficient algorithms $(\mathsf{Z}, \mathsf{C}, \mathsf{U})$ such that: On input ℓ instances $(y_1, ..., y_\ell)$ of DL, the compression algorithm Z outputs a single DL instance y; on input $(y_1, ..., y_\ell)$ together with their corresponding witnesses $(w_1, ..., w_\ell)$, the witness compression algorithm C[2] outputs a witness w to the instance $y \leftarrow \mathsf{Z}(y_1, ..., y_\ell)$; given the witness w to y, the unfolding algorithm U outputs all witnesses $(w_1, ..., w_\ell)$ to these ℓ instances. In terms of the tailored instance compression, our result on Schnorr protocol can be rephrased as follows: If the tailored instance compression scheme for DL does not exist, then Schnorr protocol is secure.

What if instance compression schemes exist for DL and RSA? We observe that the existence of instance compression scheme for DL/RSA with strong parameters has somewhat surprising consequences.

The first consequence is that, assuming the existence of good instance compression scheme for DL, then for any polynomial ℓ, the AND-DL statement $\{(y_1, y_2, ...y_\ell, g, \mathbb{G}) : \exists w_1, w_2, ..., w_\ell, s.t. \wedge_{i=1}^{\ell} g^{w_i} = y_i\}$ admits a zero knowledge proof with extremely efficient communication of size $O(1)$ group elements. The existence of tailored instance compression scheme for RSA yields a similar consequence.

The second consequence is a construction of non-trivial hash combiner for hash functions based on DL problem. Recall that given a group $\mathbb{G}$, its generator g and a random element $y \in \mathbb{G}$, we have a hash function $H_{(g,y)} : (m_0, m_1) \rightarrow g^{m_0} y^{m_1}$ that is collision-resistant. The hash combiner for DL-based hash functions is of interest in the scenario where a set of mutually untrusting parties, given a group $\mathbb{G}$ and g, want to set up a single collision-resistant hash function trusted by every one.

Several previous papers [Pie07, CRS+07, Pie08] defined *universal* hash combiners (that works for arbitrary hash functions), and showed non-trivial fully black-box combiners do not exist. Note that the above hash combiner needs to take the common parameters of the group and its generator, and works only for DL-based hash functions. However, it is still inconceivable that the above hash combiner with large ℓ exists in the real world.

We view these strong consequences as positive evidences for the security of Schnorr and Guillou-Quisquater protocols.

Input-distribution-switching technique: jumping out of the box. As mentioned before, the known previously known impossibility results hold only with respect to *restricted* reduction. We introduce an *input-distribution-switching* technique to get around these impossibility results.

[2] It is easy to see that we can construct the witness compression algorithm C by making simple adaptation to the compressing part of our embedding reduction.

Suppose that, for a hard language L_1 with *unique* witness relation R_{L_1}, and a distribution ensemble $(\mathbb{X}^1, \mathbb{W}^1)$ over R_{L_1}, there exists a distribution ensemble $(\mathbb{X}^2, \mathbb{W}^2)$ over relation R_{L_2} of a language L_2 with *two or more* witnesses that is indistinguishable from $(\mathbb{X}^1, \mathbb{W}^1)$. What can we say about the security of the classic public-coin protocols for $(\mathbb{X}^1, \mathbb{W}^1)$? At least we know that such protocols are witness indistinguishable for $(\mathbb{X}^2, \mathbb{W}^2)$.

A very vague intuition behind this positive result is that, for the same malicious verifier V^*, if we invoke V^* on both instances in $\mathbb{X}^1$ and $\mathbb{X}^2$, it should have the same behavior in these two settings since these instances are indistinguishable. This vague idea leads us to introduce the *input-distribution-switching technique*, which enables us to prove that if the ensembles $(\mathbb{X}^1, \mathbb{W}^1)$ and $(\mathbb{X}^2, \mathbb{W}^2)$ further satisfy the following properties:

- Given a sample x from $\mathbb{X}^1$, it is hard to find the unique witness for x;
- For every x in the support of $\mathbb{X}^2$, witnesses in $R_{L_2}(x)$ are uniformly distributed.

Then the classic constant-round public-coin protocols are actually witness hiding for $(\mathbb{X}^1, \mathbb{W}^1)$.

The idea of considering different types of distributions $\mathbb{X}^1$ and $\mathbb{X}^2$ on the common input already appeared in Goldreich's definition of strong witness indistinguishability [Gol01], but there they do not require indistinguishability of $(\mathbb{X}^1, \mathbb{W}^1)$ and $(\mathbb{X}^2, \mathbb{W}^2)$ since such requirement on the witness distributions $\mathbb{W}^1$ and $\mathbb{W}^2$ would trivialize the definition of witness indistinguishability. We also note that it is not clear whether there exist constant-round public-coin strong witness indistinguishable proofs for non-trivial languages (see also Appendix C of [Gol04]).

In our setting, the indistinguishability requirement on witness distributions $\mathbb{W}^1$ and $\mathbb{W}^2$ is helpful in achieving significant positive results on witness hiding protocols that bypass some previously known limitations. We give several examples of such distribution ensembles $(\mathbb{X}^1, \mathbb{W}^1)$ based on standard assumptions such as DDH, the existence of lossy trapdoor functions [PW08] and subgroup decision assumptions [DN02, BGN05, GOS12], and applying the above result we show the classic protocols of [Blu86, GMW91, DN00, GOS12] are actually witness hiding under sequential repetition for a wide range of useful cryptographic problems with unique witnesses.

2 Preliminaries

In this section we present for completeness some definitions we will use throughout this paper.

Basic Notations. We write the set $\{1, 2, ..., m\}$ as $[m]$. For a distribution D over a finite set $S \subseteq \{0,1\}^*$, we denote by $x \leftarrow D$ the process that the sample $x \in S$ is drawn according to the distribution D. We say a function $\mu(\cdot)$ is negligible if for every polynomial $p(\cdot)$, we have $\mu(n) < 1/p(n)$ for sufficiently large n. We abbreviate probabilistic polynomial-time with PPT.

Let L be an NP language defined by a polynomially bounded relation $R_L = \{(x, w) : x \in L; w \text{ is a witness for } x \in L\}$, and let $R_L(x) = \{w : (x, w) \in R_L\}$ denote the set of the witnesses of $x \in L$.

Let n be security parameter. We say two ensembles, $\mathbb{X} = \{X_n\}_{n\in\mathbb{N}}$ and $\mathbb{Y} = \{Y_n\}_{n\in\mathbb{N}}$, are computationally distinguishable if for every PPT $\mathcal{D}$, there exists a negligible function $\mu(n)$ such that

$$|\Pr[\mathcal{D}(X_n) = 1] - \Pr[\mathcal{D}(Y_n) = 1]| \leq \mu(n).$$

Interactive Proofs. An *interactive proof system* $\langle P, V\rangle$ [GMR89] for a language L is a pair of interactive Turing machines in which the prover P wishes to convince the verifier V

of some statement $x \in L$. We denote by $\langle P, V \rangle(x)$ the output of V at the end of interaction on common input x, and without loss of generality, we have the verifier V outputs 1 (resp. 0) if V accepts (resp. rejects).

Definition 1 (Interactive Proofs). *A pair of interactive Turing machines $\langle P, V \rangle$ is called an interactive proof system for language L if V is a PPT machine and the following conditions hold:*

- *Completeness: For every $x \in L$, $\Pr[\langle P, V \rangle(x) = 1] = 1$.*
- *Soundness: For every $x \notin L$, and every (unbounded) prover P^*, there exists a negligible function $\mu(n)$ (where $|x| = n$) such that*

$$\Pr[\langle P^*, V \rangle(x) = 1] < \mu(n).$$

An interactive *argument* [BCC88] is an interactive proof except that for which soundness is only required to hold against PPT cheating provers. We often use "protocol" to refer to both proof system and argument system.

Witness Indistinguishability. Witness indistinguishable proof system guarantees that if the statement has two independent witnesses, then the malicious verifier cannot tell which witness is being used by the prover in an execution of the protocol.

Definition 2 (Witness Indistinguishability). *Let L be an NP language defined by R_L. We say that $\langle P, V \rangle$ is witness indistinguishable for relation R_L if for every PPT V^* and every sequence $\{(x, w, w')\}_{x \in L}$, where $(x, w), (x, w') \in R_L$ the following two probability ensembles are computationally indistinguishable:*

$$\{\langle P(w), V^* \rangle(x)\}_{x \in L} \overset{c}{\approx} \{\langle P(w'), V^* \rangle(x)\}_{x \in L}.$$

Witness Hiding. Loosely speaking, witness hiding of a protocol [FS90] refers to the following property: for an input $x \in L$ that is being proven, if a verifier can extract a witness in $R_L(x)$ after interacting with the prover, then he could have done so without such an interaction. This notion is formally defined with respect to a distribution ensemble over inputs as follows.

Definition 3 (Distribution of Hard Instances). *Let L be an NP language defined by R_L. Let $\mathbb{X} = \{X_n\}_{n \in \mathbb{N}}$ be a distribution ensemble. We say that $\mathbb{X}$ is hard for R_L if for every PPT machine M*

$$\Pr\left[M(X_n) \in R_L(X_n)\right] < \mu(n).$$

Definition 4 (Witness Hiding (under Sequential Repetition)). *Let L be an NP language defined by R_L, $(\mathbb{X}, \mathbb{W}) = \{(X_n, W_n)\}_{n \in \mathbb{N}}$ be a distribution over R_L. We say $\langle P, V \rangle$ is witness hiding for $(\mathbb{X}, \mathbb{W})$ if for every PPT machine V^**

$$\Pr\left[\langle P(W_n), V^* \rangle (X_n) \in R_L(X_n)\right] < \mu(n).$$

We say that $\langle P, V \rangle$ is witness hiding under sequential repetition if it is witness hiding for $(\mathbb{X}, \mathbb{W})$ under any polynomially number of sequential repetitions.

Remark 1. According to our definition of witness hiding, it is easy to verify that if there is witness hiding protocol for $(\mathbb{X}, \mathbb{W})$, then the distribution ensemble $\mathbb{X} = \{X_n\}_{n \in \mathbb{N}}$ on instances must be hard.

Zero Knowledge Proofs. A stronger security notion for the prover of an interactive proof system is zero knowledge, which requires the entire view of a malicious verifier can be reconstructed by a PPT algorithm efficiently.

Definition 5 (Zero Knowledge Proofs). *We say that an interactive proof system $\langle P, V\rangle$ for language L is zero knowledge if for any PPT V^*, there exists a PPT Sim such that*

$$\{View^P_{V^*}(x)\}_{x\in L} \overset{c}{\approx} \{Sim(x)\}_{x\in L}.$$

Where $\{View^P_{V^*}(x)\}_{x\in L}$ denotes the distribution of the view of the malicious verifier V^* in the real interaction.

3 Tailored Instance Compression for Search Problems

The study of instance compression was initiated by Harnik and Naor [HN10]. We tailor their definition for our purpose. Roughly speaking, a tailored instance compression scheme for a (search) NP problem can compress a long instance(s) into a shorter instance, and given the solution to the shorter instance, we can solve all the original instance(s). It should be noted that the impossibility results of [FS11, Dru15] with respect to NP-complete languages also hold for our tailored definition.

Definition 6 (Tailored Instance Compression for Search Problem). *Let L be a NP language and R_L its NP relation. A (ℓ, ε)-tailored instance compression scheme for R_L consists of three PPT algorithms $(\mathsf{Z}, \mathsf{C}, \mathsf{U})$, such that for sufficient large n:*

- $(x, st) \leftarrow \mathsf{Z}(x_1, \cdots, x_\ell)$*: On input $x_i \in L$ for $i \in [\ell]$, the PPT instances compression algorithm Z outputs a single $x \in L$ and the state st.*
- $R_L(\mathsf{Z}(x_1, \cdots, x_\ell) \setminus st) \ni w \leftarrow \mathsf{C}((x_1, w_1), \cdots, (x_\ell, w_\ell))$*: On input $(x_i, w_i) \in R_L$ for $i \in [\ell]$, the PPT witness compression algorithm C outputs a witness w to the instance x generated by $\mathsf{Z}(x_1, \cdots, x_\ell)$.*
- $(w_1, \cdots, w_\ell) \leftarrow \mathsf{U}(x, w, st)$*: On input $x \in L$, st, together with the corresponding witness $w \in R_L(x)$, the PPT unfolding algorithm U outputs the witnesses $w_i \in R_L(x_i)$ for all $i \in [\ell]$.*
- *For all $w \in R_L(x)$, the following holds:*

$$\Pr[(x, st) \leftarrow \mathsf{Z}(x_1, \ldots, x_\ell); (w_1, \ldots, w_\ell) \leftarrow \mathsf{U}(x, w, st) : \wedge_{i=1}^{\ell}\ w_i \in R_L(x_i)] > \varepsilon.$$

Remark 2. Our definition is stronger than the one of [HN10] in several respects. In the definition 2.25 of [HN10], the retrieving algorithm (that corresponds to our witness compression algorithm) does not take witnesses to $(x_1, \ldots, x_\ell)$ as input, and thus is not required to be efficient; the unfolding algorithm above is also not required in [HN10], but that is the key for our applications of instance compression scheme (if exists).

Observe that the one-more like assumptions can be rephrased in the framework of instance compression. For example, the one-more DL assumption is equivalent to assume non-existence of (ℓ, ε)-tailored instance compression scheme for DL with weaker requirements: 1) The witness compression algorithm is not required; 2) The instance compression algorithm is allowed to output $\ell - 1$ instances (which leads to much weak compression ratio) and the unfolding algorithm needs to take $\ell - 1$ witnesses correspondingly.

4 Embedding Reduction: the Security of Schnorr and Guillou-Quisquater Protocols and Instance Compression

In this section, we develop an embedding reduction technique to base the witness hiding security[3] of Schnorr protocol on non-existence of tailored instance compression scheme for discrete logarithm. Similar results can also be obtained for the Guillou-Quisquater's protocol and some other Σ-protocols for group homomorphisms. Note that, given a successful adversary V^*, our technique yields a tailored instance compression scheme with parameters much stronger than the ones in [BP02], and thus strengthens the results of [BP02].

4.1 The Security of Schnorr Protocol

Let $\mathbb{G}$ be a cyclic group of order q with the generator g, where q is a prime such that $q \mid p-1$, p is a prime $2^{n-1} \leq p \leq 2^n$. Given a common input y, the Schnorr protocol allows the prover P wants to convince the verifier V of knowledge of the unique discrete logarithm w of y (i.e., $y = g^w$). Formal description of this protocol can be found in Fig. 1.

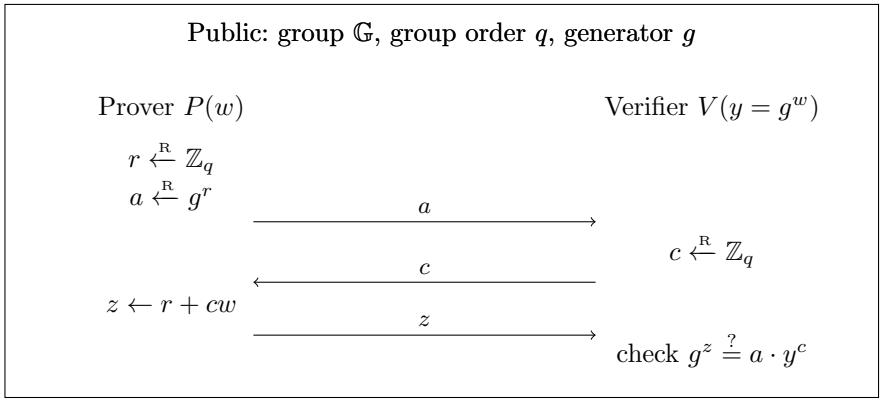

Fig. 1: Schnorr identification scheme

Given $(g, \mathbb{G})$, we define the NP relation $R_{(g,\mathbb{G})} := \{(y, w) : y = g^w\}$. We show that a successful adversarial verifier will lead to a non-trivial tailored instance compression scheme for discrete logarithm (DL) instances.

Theorem 1. *If there exists a PPT algorithm V^* that breaks witness hiding of Schnorr protocol with probability p (i.e. V^* after interaction with the prover P outputs a valid discrete logarithm w of y with probability greater than p), then there exists $(\ell, p^{\ell-1})$-tailored instance compression scheme for DL instances in $\mathbb{G}$ for any ℓ.*

Remark 3. It should be noted that for a negligible probability ε, the (ℓ, ε)-tailored instance compression scheme (if exists) is barely applicable. For achieving meaningful compression scheme from V^*, we should set ℓ to be (arbitrary) constant when p is an inverse polynomial; if p is negligibly close to 1, then ℓ can be set to be (arbitrary) polynomial. Note also that the technique of [BP02] gives us only $\ell = 2$.

[3] As mentioned, this is equivalent to the identification security in the case where the adversary communicates with $P(sk)$ in a single execution.

We first construct two efficient subroutines D and B for our embedding reduction. On input two instances (y_1, y_2), the algorithm D interacts with V^* (where y_1 serves as the common input, and y_2 serves as the first prover message) until the challenge c from V^* is received, and outputs a new instance $y_1^c y_2$; on input discrete logarithm z of $y_1^c y_2$, the algorithm B interacts with V^* until the output of V^* is received, and outputs two discrete logarithms of the two instances (y_1, y_2). Formal descriptions of D and B can be found in Algorithm $D^{V^*}$and B^{V^*}.

D^{V^*}

input : instances $y_1, y_2 \in \mathbb{G}$, random tape R_V

1: Run V^* with random tape R_V on instance y_1;
2: Send y_2 as the first prover message to V^*;

output: output: If V^* answers with a challenge $c \in \mathbb{Z}_q$, output $y = y_1^c y_2$; else output $\perp$.

B^{V^*}

input : $z \in \mathbb{Z}_q$, $y_1, y_2 \in \mathbb{G}$, random tape R_V

1: Execute the Schnorr protocol with V^* in exactly the same way as $D(y_1, y_2, R_V)$ until receiving the challenge c from V^*;
2: Send z, which is supposed to be such that $g^z = {y_1}^c y_2$, to V^* ;

output: If V^* outputs the witness w satisfying $y_1 = g^w$, output $z_1 = w$ and $z_2 = z - cw$; else output $\perp$.

The compression algorithm Z^{V^*}

input : $(y_1, y_2, \cdots, y_\ell)$

1: $st \leftarrow \{y_1, \cdots, y_\ell\}$;
2: set $y_j^0 = y_j$, for $j = 1, 2, \cdots, \ell$;
3: **for** $i \leftarrow 0$ **to** $l - 1$ **do**
4: **for** $j \leftarrow 1$ **to** 2^{l-i-1} **do**
5: $y_{2j-1}^i \leftarrow y_{2j-1}^i \cdot g^{r_{2j-1}^i}$, $y_{2j}^i \leftarrow y_{2j}^i \cdot g^{r_{2j}^i}$, where $r_{2j-1}^i, r_{2j}^i \xleftarrow{R} \mathbb{Z}_q$;
6: ${R_V}_j^i \xleftarrow{R} \{0,1\}^{poly(n)}$, where $poly(n)$ denotes the length of the random tape ${R_V}_j^i$;
7: $y_j^{i+1} \leftarrow D^{V^*}(y_{2j-1}^i, y_{2j}^i, {R_V}_j^i)$ (if D outputs $\perp$, return $\perp$);
8: Add $(y_j^{i+1}, r_{2j-1}^i, r_{2j}^i, {R_V}_j^i)$ to st;
9: **end**
10: **end**
11: set $y \leftarrow y_1^l$;
12: Return y, st;

As illustrated in Figure 2, our embedding black-box reduction naturally corresponds to a pair of efficient algorithms, a compression algorithm Z and an unfolding algorithm U. In

The unfolding algorithm U^{V^*}

input : $y \in \mathbb{G}$, $w \in \mathbb{Z}_q$, st

set $y_1^l \leftarrow y$, $z_1^l \leftarrow w$;
for $i = l-1$ **to** 0 **do**
  for $j = 1$ **to** 2^{l-i-1} **do**
    Retrieve $y_{2j-1}^i, y_{2j}^i, r_{2j-1}^i, r_{2j}^i$ and ${R_V}_j^i$ from st;
    $(z_{2j-1}^i, z_{2j}^i) \leftarrow B^{V^*}(z_j^{i+1}, y_{2j-1}^i, y_{2j}^i, {R_V}_j^i)$ (if B outputs $\perp$, return $\perp$);
    $z_{2j-1}^i \leftarrow z_{2j-1}^i - r_{2j-1}^i$, $z_{2j}^i \leftarrow z_{2j}^i - r_{2j}^i$;
  end
end

output: $(w_1, w_2, \cdots, w_\ell) = (z_1^0, z_2^0, \cdots, z_\ell^0)$

the first phase, the compression algorithm Z, taking as input discrete logarithm instances $(y_1, \ldots, y_\ell)$, invokes D recursively to generate new instance, each time D transforming two new instances into a new single one. Z outputs the final single instance $y = y_1^3$ and the corresponding st consisting of all instances input to D and the random tape of Z.

On input a witness $w = z_1^3$ to $y = y_1^3$, the unfolding algorithm U invokes B recursively, by feeding B with a discrete logarithm of an instance, to solve two instances. Finally, U will solve all instances $(y_1, y_2, ..., y_\ell)$.

For our analysis to go through, given two instances y_1, y_2, the compression algorithm Z has to choose two random strings r_1, r_2 and a fresh random tape for V^*, and then runs D on input $(y_1 g^{r_1}, y_2 g^{r_2})$. Z will store all these randomnesses in st. The formal descriptions of Z and U can be found in Algorithm $D^{V^*}$and B^{V^*}respectively. Without loss of generality, we assume that $\ell = 2^l$ for some integer l.

Proof. (of Theorem 1)

From the picture in Figure 2, we see the symmetry that, on input two instances (y_{2j-1}^i, y_{2j}^i), $D^{V^*}(y_{2j-1}^i, y_{2j}^i, {R_V}_j^i)$ generates a new instance y_j^{i+1}; whereas, on input a discrete logarithm z_j^{i+1} of y_j^{i+1}, $B^{V^*}(z_j^{i+1}, y_{2j-1}^i, y_{2j}^i, {R_V}_j^i)$ produces the two discrete logarithms (z_{2j-1}^i, z_{2j}^i) of the two instances (y_{2j-1}^i, y_{2j}^i) that are inputs to D.

We say an algorithm wins if it does not output "$\perp$". Note that all these invocations of D are independent, and that, for every i,j, the V^* success probability p is the probability that both $D^{V^*}(y_{2j-1}^i, y_{2j}^i, {R_V}_j^i)$ and $B^{V^*}(z_j^{i+1}, y_{2j-1}, y_{2j}, {R_V}_j^i)$ win, that is,

$$\Pr[D^{V^*}(y_{2j-1}^i, y_{2j}^i, {R_V}_j^i) \text{ wins} \wedge B^{V^*}(z_j^{i+1}, y_{2j-1}, y_{2j}, {R_V}_j^i) \text{ wins}] = p.$$

Observe that in the entire reduction there are exactly $(\ell - 1)$ pairs of invocations of D^{V^*} and B^{V^*}, thus we have the probability

$$\Pr[(y, st) \leftarrow \mathsf{Z}^{V^*}(y_1, y_2, \cdots, y_\ell); (w_1, w_2, \cdots, w_\ell) \leftarrow \mathsf{U}^{V^*}(y, st, w) : \wedge_{i=1}^{\ell} y_i = g^{w_i}].$$

is $p^{\ell-1}$.

Note that when given as input all the witnesses $(w_1, \cdots, w_\ell)$ of the target instances $(y_1, \cdots, y_\ell)$ to Z, Z is able to compute the witness to every instance output by D. Thus by making a straightforward adaptation of Z we get a PPT witness compression algorithm C as desired. This completes the proof.

□

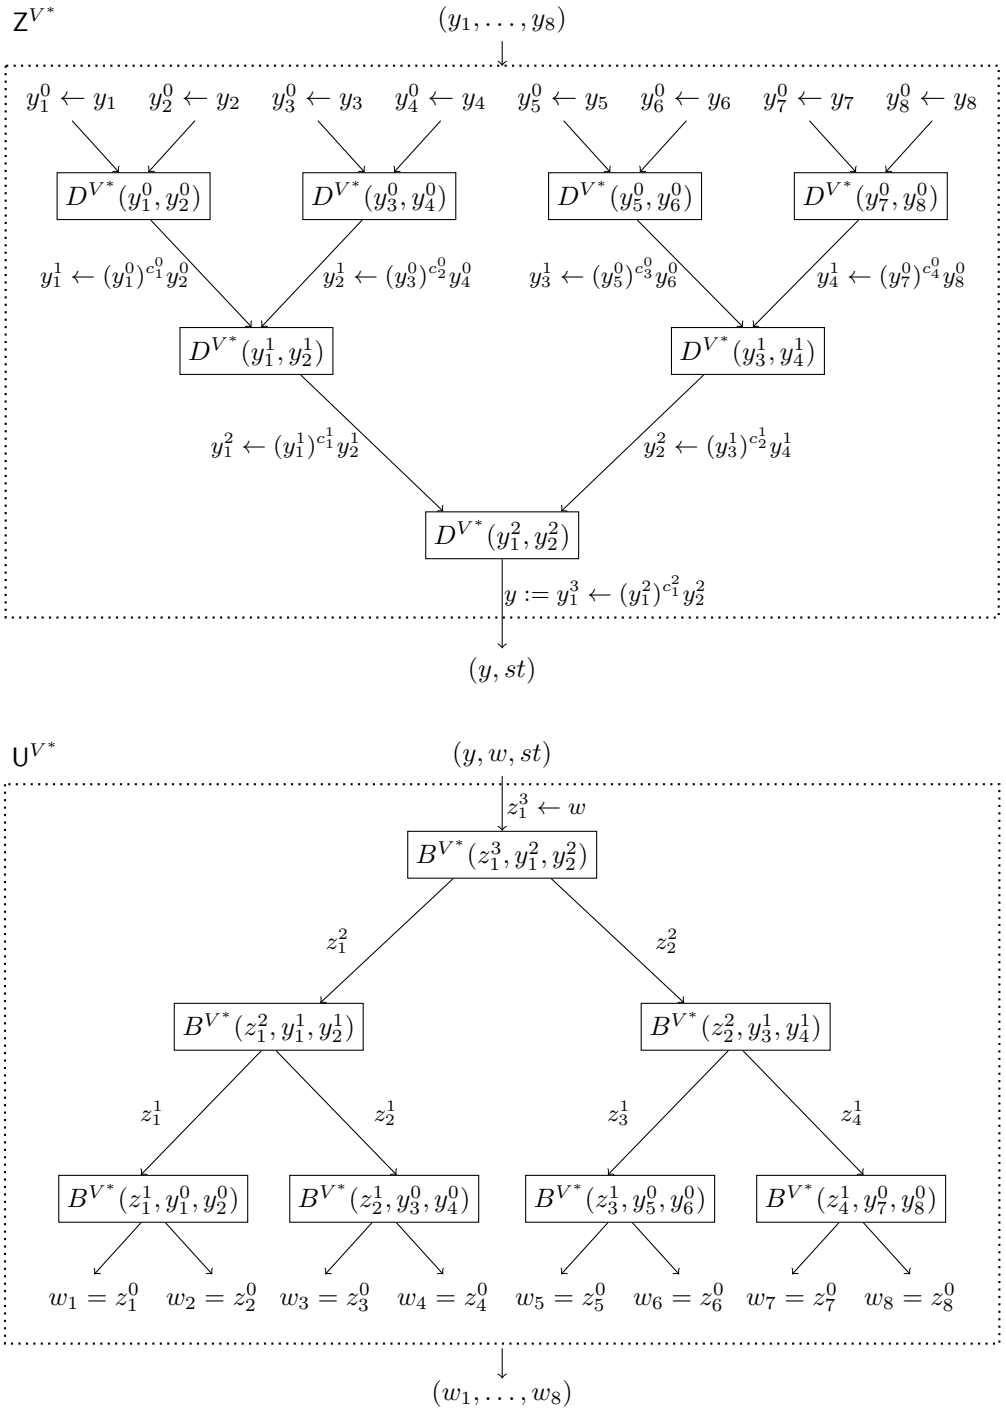

Fig. 2: Simplified reduction for $\ell = 8$. We assume that V^* is deterministic and with probability 1 it breaks witness hiding of Schnorr protocol.

4.2 Security of the Guillou-Quisquater Protocol

In this section we state a similar result on Guillou-Quisquater identification protocol [GQ88]. The reduction is essentially the same as the one for Schnorr protocol, and here we omit it.

The Guillou-Quisquater Protocol Let $N = pq$ be an RSA modulus (i.e. p and q are large distinct primes for security parameter n) and $e < \varphi(N)$ be an odd prime satisfying $\gcd(d, \varphi(N)) = 1$ and $ed \equiv 1 \mod \varphi(N)$. The Guillou-Quisquater protocol proceeds as follows (See Figure 3). The prover P wants to convince the verifier V of the unique e-th root w modulo N of a given number y. First, P chooses $r \in \mathbb{Z}_N^*$ at random and sends $a = r^e \mod N$ to the verifier V. Upon receiving the verifier's challenge c, P responses with $z = r \cdot w^c$. V accepts if and only if $z^e = a \cdot y^c$.

Given (e, N), we define the NP relation $R_{e,N} := \{(y, w) : y = w^e mod N\}$. Similar to the Schnorr protocol, we have the following theorem.

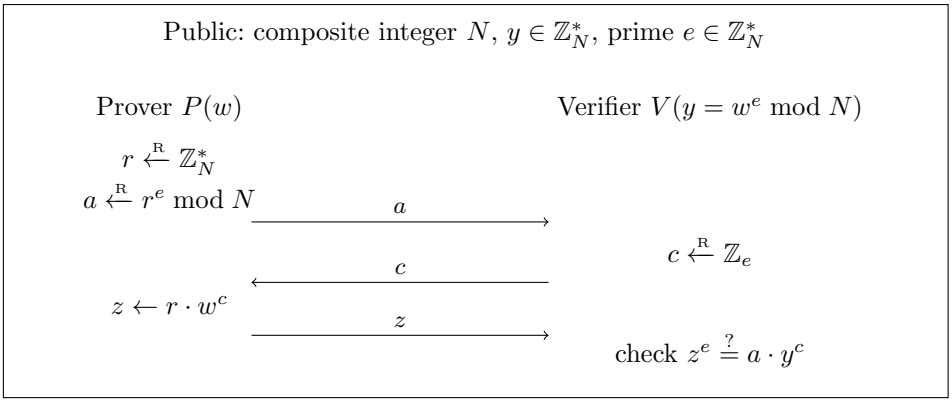

Fig. 3: GQ identification scheme

Theorem 2. *If there exists a PPT algorithm V^* that breaks witness hiding of Guillou-Quisquater protocol with probability p (i.e. V^* after interaction outputs the witness w with probability greater than p), then there exists $(\ell, p^{\ell-1})$-tailored instance compression scheme for RSA instances in $\mathbb{Z}_N^*$ for any ℓ.*

Remark 4. We also note that our reduction can also apply to $\Sigma-$protocols for group homomorphisms [Mau15, CFGV13].

5 Some Consequences of Existence of Good Tailored Instance Compression Schemes for DL and RSA

In this section, we show some strong consequences of the existence of good tailored instance compression schemes for DL and RSA problems. To simplify our presentation, we consider only $(poly(n), 1 - negl(n))$-tailored instance compression schemes, where $poly(n)$ denotes an arbitrary polynomial in security parameter n. Such an instance compression scheme can be constructed from the efficient adversary that can break the witness hiding of Schnorr/Guillou-Quisquater protocol with probability negligibly close to 1. We also stress that, as showed in [Pas11], even for such an adversary, no black-box reduction can turn it into an algorithm that breaks some standard assumptions and reach a contradiction.

5.1 Extremely Communication-Efficient Zero Knowledge Protocols for AND-DL and AND-RSA

Suppose that there is a $(poly(n), 1-negl(n))$-tailored instance compression scheme $(\mathsf{Z}, \mathsf{C}, \mathsf{U})$ for DL. In this subsection we further assume that the compression algorithm Z is deterministic without loss of generality: Since almost all possible random tapes for Z are good in the sense that on every such random tape Z will output an instance, together with some state information, for which the unfolding algorithm will succeed, we can publish a good random tape and let each party execute Z on the same random tape when needed.

The immediate consequence of such a tailored instance compression scheme is that, for an arbitrary polynomial ℓ, the AND-DL statement, $\{(y_1, y_2, \ldots, y_\ell, g, \mathbb{G}) : \exists w_1, w_2 \ldots, w_\ell, s.t. \wedge_{i=1}^{\ell} g^{w_i} = y_i\}$, have a proof of size $|w_i|$, since we can have both the prover and the verifier run Z on $(y_1, y_2, \ldots, y_\ell)$ and obtain a single instance y of the same size of y_i, and then the prover send the w (such that $g^w = y$) to the verifier, which accepts if $g^w = y$ and all w_i, obtained from the unfolding algorithm U, satisfy $g^{w_i} = y_i$.

With this succinct proof for the AND-DL statement, the Feige-Shamir zero knowledge protocol of [FS90] for AND-DL statements can be implemented in an extremely communication-efficient way (with communication of size $O(1)$ group elements).

Protocol Feige-Shamir

Common input: $y_1, y_2, \ldots, y_\ell \in \mathbb{G}$.
The prover P's input: $w_1, w_2, \ldots, w_\ell, s.t. \wedge_{i=1}^{\ell} g^{w_i} = y_i$.

First phase: The verifier chooses $w_0', w_1' \xleftarrow{R} \mathbb{Z}_q$ independently and at random, compute $y_0' = g^{w_0'}$ and $y_1' = g^{w_1'}$, and then execute the 3-round Σ_{OR} protocol (OR-composition of the Schnorr protocol[CDS94]), in which V plays the role of the prover, to prove the knowledge of the witness to the statement $(y_0' \vee y_1')$;
Second phase: Both the prover and the verifier run Z on $(y_1, y_2, \ldots, y_\ell)$ and obtain a new instance $y \in \mathbb{G}$, and then the prover runs the witness compression algorithm C on $w_1, w_2, \ldots, w_\ell$ to obtain w such that $g^w = y$, and proves to the verifier the knowledge of the witness to the statement $(y \vee y_0' \vee y_1')$ using Σ_{OR} protocol of [CDS94].

This leads to the following proposition.

Proposition 1. *If there exists $(poly(n), 1-negl(n))$-tailored instance compression scheme for AND-DL, then for an arbitrary polynomial ℓ, the AND-DL statement, $\{(y_1, y_2, \ldots, y_\ell, g, \mathbb{G}) : \exists w_1, w_2 \ldots, w_\ell, s.t. \wedge_{i=1}^{\ell} g^{w_i} = y_i\}$, has a zero knowledge protocol with communication complexity of $O(1)$ group elements.*

5.2 Special Hash Combiner

The second consequence is a construction of non-trivial hash combiner for hash functions based on DL problem, which would help a set of ℓ mutually untrusting parties set up a single trusted collision-resistant hash function from a given group.

Consider the cyclic group $\mathbb{G}$ mentioned in Section 4.1. Let $y = g^w$ for some w. $h^y : {\mathbb{Z}_q}^2 \to \mathbb{G}$ is *collision resistant hash functions* (CRHFs) based on DL problem defined as follows:

$$h^y(m_0, m_1) = g^{m_0} y^{m_1}.$$

Clearly, finding a collision for h^y is equivalent to solving the discrete logarithm problem $w = \log_g y$.

Definition 7 (Hash Combiner for CRHFs Based on DL Problem). *A non-uniform PPT Turing machine* $\mathsf{H} : \mathcal{R} \times {\mathbb{Z}_q}^2 \to \{0,1\}^v$ *is said to be a randomized* (k, ℓ)*-combiner for CRHFs based on DL, if it satisfies the following conditions:*

- *For any given* ℓ *elements of* $\mathbb{G}$ *(i.e.* $y_1, \cdots, y_\ell$*), for every* $r \in \mathcal{R}$, $\mathsf{H}^{y_1, y_2, \cdots, y_\ell}(r, \cdot, \cdot)$ *is a collision resistant hash function, if at least* k *components* y_i *can be used to construct collision resistant hash functions* $h^{y_i}(\cdot, \cdot)$*.*
- *For every PPT adversary* $\mathcal{B}$ *breaks the collision resistent hash combiner* $\mathsf{H}^{y_1, y_2, \cdots, y_\ell}(r, \cdot, \cdot)$*, there exists a PPT reduction* R*, s.t.* $R^{\mathcal{B}}$ *can find collisions for at least* $\ell - k + 1$ *hash functions* h^{y_i}*,* $i \in [\ell]$*, with overwhelming probability.*

Now we will show that the combiner for CRHFs based on the DL problem can be constructed by the compression algorithm for DL instances. The previous papers [Pie07, CRS+07, Pie08] showed that there doesn't exist "fully"[4] black-box combiners whose output length is significantly smaller than what can achieved by trivially concatenating the output of any $\ell - k + 1$ of the components. We can construct a special non-black-box $(1, \ell)$-combiner for CRHFs based on DL problem whose output length is significantly smaller using the instances compression algorithm mentioned in Corollary 2, under the discrete logarithm assumption.

Proposition 2. *Suppose there exists* $(poly(n), 1 - negl(n))$*-tailored instance compression algorithms for any given* $\ell (= poly(n))$ *DL instances* $y_1, y_2, \ldots, y_\ell$ *in* $\mathbb{G}$*. Then there exists a randomized* $(1, \ell)$*-combiner* $\mathsf{H}^{y_1, y_2, \ldots, y_\ell}(r, \cdot, \cdot)$ *for CRHFs based on DL problem, with the same output length* v *as the regular discrete logarithm hash functions* h^{y_i}*.*

Proof. Assume that there exists $(poly(n), 1 - negl(n))$-tailored instance compression algorithms for DL. That is, for any polynomial ℓ, there exists a pair of PPT algorithms (Z, U), for $w = \log_g y$, such that

$$\Pr[(y, st) \leftarrow \mathsf{Z}(y_1, \ldots, y_\ell); (w_1, \ldots, w_\ell) \leftarrow \mathsf{U}(y, w, st) : \wedge_{i=1}^{\ell}\ w_i = \log_g y_i] > 1 - negl(n).$$

The combiner has the following form:

$$\mathsf{H}^{(y_1, y_2, \cdots, y_\ell)}(r, m_0, m_1) = h^y(m_0, m_1) = g^{m_0} y^{m_1}.$$

where $y \leftarrow \mathsf{Z}(y_1, y_2, \cdots, y_\ell)$, and r is the same random tape as the compression algorithm Z used.

Note that a pair of collisions for h^y will give the discrete logarithm of y, which in turn can be used (by applying U) to solve all DL instances $y_1, \ldots, y_\ell$, and therefore we can find a pair of collisions for each hash function h^{y_i} efficiently. Thus this combiner is a $(1, \ell)$-combiner for CRHFs based on DL problem as defined in Definition 7.

□

Application of Special Hash Combiner: How to Set Up a Global Hash. Suppose in a multi-party setting, a given number of participants, $P_1, \cdots, P_\ell$, each P_i has its own hash function h^{y_i} with the same common parameter $\mathbb{G}, g$, and want to set up a *single* hash function trusted by all of them. The need for a global hash function was also addressed in [CLP13]. While we can't simple choose some participant's hash function as the global hash function for obvious reasons. we can use our special hash combiner to solve this puzzle: Each participant runs the instance compression algorithm Z on these $(y_1, \cdots, y_l)$ locally

[4] Fully black box combiners mean both constructions and security proofs are black-box.

and generates a single common $y \in \mathbb{G}$, and then they set $H_{(g,y)} : (m_0, m_1) \to g^{m_0} y^{m_1}$ to be the global hash function. This function is collision-resistant free since every collision would lead to a solution to the instance y', which will enable the unfolding algorithm U to find all discrete logarithms of these random y_i's, and thus if there is one y_i generated at random by an honest party, no PPT algorithm can find a collision for $H_{(g,y)}$.

6 Witness Hiding Protocols for Distributions on Some Hard Relations with Unique Witnesses

In this section we prove a general theorem on witness hiding of constant-round public-coin proofs systems for unique witness relations and present its applications to several cryptographic problems.

6.1 A General Theorem

Let L_1 and L_2 be NP languages (possibly the same), R_{L_1} and R_{L_2} be their corresponding witness relations. Let $(\mathbb{X}^1, \mathbb{W}^1) = \{(X_n^1, W_n^1)\}_{n\in\mathbb{N}}$ be a distribution ensemble over R_{L_1} with *unique* witnesses, and $(\mathbb{X}^2, \mathbb{W}^2) = \{(X_n^2, W_n^2)\}_{n\in\mathbb{N}}$ be a distribution ensemble over R_{L_2} with *multiple* witnesses.

Theorem 3. *If the above distribution ensembles satisfy the following conditions:*

1. $(\mathbb{X}^1, \mathbb{W}^1)$ *and* $(\mathbb{X}^2, \mathbb{W}^2)$ *are computationally indistinguishable.*
2. *For sufficiently large n, for every PPT machine M, there is negligible function $\mu(n)$, such that*

$$\Pr\left[(x, w) \leftarrow (X_n^2, W_n^2); w' \leftarrow M(x, w) : w' \in R_L(x) \wedge w \neq w'\right] < \mu(n).$$

3. *For every n and x in X_n^2, witnesses in $R_{L_2}(x)$ are uniformed distributed.*[5]

Then, any witness indistinguishable constant-round public-coin proof systems (including the parallelized version of 3-round public-coin proofs of [Blu86, GMW91] and ZAPs of [DN00, GOS12]) are witness hiding (under sequential repetition) for $(\mathbb{X}^1, \mathbb{W}^1)$.

Proof. Let $\langle P, V\rangle$ be an arbitrary witness indistinguishable proof system. In the following, we present our proof only for the standalone case. Note that the same proof works also for these protocols under sequential repetition.

Suppose, towards a contradiction, that there are infinitely many n, a polynomial p, and a PPT verifier V^* such that

$$\Pr\left[\left\langle P(W_n^1), V^*\right\rangle (X_n^1) \in R_{L_1}(X_n^1)\right] > \frac{1}{p(n)}. \tag{1}$$

Let $\mathbb{S}$ be the set of such n's. Fix an $n \in \mathbb{S}$ and consider the following two experiments:

EXP^b $(b \in \{1, 2\})$: Sample $(x, w) \leftarrow (X_n^b, W_n^b)$, play the role of honest prover $P(x, w)$ and interact with $V^*(x)$. When V^* terminates, output what V^* outputs.

[5] This condition can be significantly relaxed, but we stick to it for simplifying presentation.

Denote by WIN^b that EXP^b outputs a witness for x. By the indistinguishability of $(\mathbb{X}^1, \mathbb{W}^1)$ and $(\mathbb{X}^2, \mathbb{W}^2)$, we have that for some negligible function $\mu(n)$

$$\Pr\left[\text{WIN}^2\right] = \Pr\left[\left\langle P(W_n^2), V^*\right\rangle(X_n^2) \in R_{L_2}(X_n^2)\right] > \frac{1}{p(n)} - \mu(n). \tag{2}$$

It follows from the second property of (X_n^2, W_n^2) that

$$\Pr\left[(x,w) \leftarrow (X_n^2, W_n^2) : \langle P(w), V^*\rangle(x) = w' \in R_{L_2}(x) \wedge w' \neq w\right] < \mu(n). \tag{3}$$

Now by (2) and (3), we have

$$\Pr\left[(x,w) \leftarrow (X_n^2, W_n^2) : \langle P(w), V^*\rangle(x) = w' \wedge w' = w\right] > \frac{1}{p(n)} - \mu(n), \tag{4}$$

which can be rewritten as

$$\begin{aligned}
&\Pr\left[(x,w) \leftarrow (X_n^2, W_n^2) : \left\langle P(w), V^*\right\rangle(x) = w' \wedge w' = w\right] \\
&= \sum_x \sum_w \Pr\left[\langle P(w), V^*\rangle(x) = w' \wedge w' = w\right] \Pr\left[w \leftarrow W_n^2 | x\right] \Pr\left[x \leftarrow X_n^2\right] \\
&> \frac{1}{p(n)} - \mu(n).
\end{aligned}$$

Theorem 3 follows from the following two claims.

Claim 1. There exists x in the support of X_n^2 satisfying the following two conditions:

$$\begin{aligned}
&- \sum_w \Pr\left[\langle P(w), V^*\rangle(x) = w' \wedge w' = w\right] \Pr\left[w \leftarrow W_n^2 | x\right] > \frac{1}{2p(n)} - \mu(n). \\
&- \sum_w \Pr\left[\langle P(w), V^*\rangle(x) = w' \in R_{L_2}(x) \wedge w' \neq w\right] \Pr\left[w \leftarrow W_n^2 | x\right] < \mu(n).
\end{aligned}$$

Claim 2. There exists x in the support of X_n^2, $w_1, w_2 \in R_{L_2}(x)$ such that

$$\left|\Pr\left[\langle P(w_1), V^*\rangle(x) = w_1\right] - \Pr\left[\langle P(w_2), V^*\rangle(x) = w_1\right]\right| > \frac{1}{poly(n)}.$$

Note that Claim 2 holds for each $n \in \mathbb{S}$, and thus we conclude that V^* breaks the witness indistinguishability of $\langle P, V\rangle$ on a sequence $\{(x, w_1, w_2)\}_{x \in X_n^2, n \in \mathbb{S}}$, which contradicts the fact that $\langle P, V\rangle$ is witness indistinguishable for *multiple* witnesses relation. This proves theorem 3. □

We now give the detailed proofs of the above two claims.

*Proof (of **Claim 1**).* We define the following two events:

- EVENT_{eq}: $\langle P(w), V^*\rangle(x) = w' \wedge w' = w$.
- EVENT_{neq}: $\langle P(w), V^*\rangle(x) = w' \in R_{L_2}(x) \wedge w' \neq w$.

and two sets:

- $\mathbb{H}$: $\{x : \sum_w \Pr\left[\text{EVENT}_{eq}\right] \Pr\left[w \leftarrow W_n^2 | x\right] > \frac{1}{2p(n)} - \mu(n)\}$.
- $\mathbb{K}$: $\{x : \sum_w \Pr\left[\text{EVENT}_{neq}\right] \Pr\left[w \leftarrow W_n^2 | x\right] < \mu(n)\}$.

Observe that

$$\begin{aligned}\frac{1}{p(n)} - \mu(n) &< \Pr\left[(x,w) \leftarrow (X_n^2, W_n^2) : \langle P(w), V^* \rangle (x) = w' \wedge w' = w\right] \\ &= \sum_{x \in \mathbb{H}} \sum_{w} \Pr\left[\text{EVENT}_{eq}\right] \Pr\left[w \leftarrow W_n^2 | x\right] \Pr\left[x \leftarrow X_n^2\right] \\ &\quad + \sum_{x \notin \mathbb{H}} \sum_{w} \Pr\left[\text{EVENT}_{eq}\right] \Pr\left[w \leftarrow W_n^2 | x\right] \Pr\left[x \leftarrow X_n^2\right] \\ &= \sum_{w} \Pr\left[\text{EVENT}_{eq}\right] \Pr\left[w \leftarrow W_n^2 | x \in \mathbb{H}\right] \Pr\left[x \leftarrow X_n^2 : x \in \mathbb{H}\right] \\ &\quad + \sum_{w} \Pr\left[\text{EVENT}_{eq}\right] \Pr\left[w \leftarrow W_n^2 | x \notin \mathbb{H}\right] \Pr\left[x \leftarrow X_n^2 : x \notin \mathbb{H}\right],\end{aligned}$$

which, by the definitions of EVENT_{eq} and set $\mathbb{H}$, leads to

$$\Pr\left[x \leftarrow X_n^2 : x \in \mathbb{H}\right] > \frac{1}{2p(n)} - \mu(n). \tag{5}$$

Similarly, by (3), we have

$$\begin{aligned}\mu(n) &> \Pr\left[(x,w) \leftarrow (X_n^2, W_n^2) : \langle P(w), V^* \rangle (x) = w' \in R_{L_2}(x) \wedge w' \neq w\right] \\ &= \sum_{x \in \mathbb{K}} \sum_{w} \Pr\left[\text{EVENT}_{neq}\right] \Pr\left[w \leftarrow W_n^2 | x\right] \Pr\left[x \leftarrow X_n^2\right] \\ &\quad + \sum_{x \notin \mathbb{K}} \sum_{w} \Pr\left[\text{EVENT}_{neq}\right] \Pr\left[w \leftarrow W_n^2 | x\right] \Pr\left[x \leftarrow X_n^2\right] \\ &= \sum_{w} \Pr\left[\text{EVENT}_{neq}\right] \Pr\left[w \leftarrow W_n^2 | x \in \mathbb{K}\right] \Pr\left[x \leftarrow X_n^2 : x \in \mathbb{K}\right] \\ &\quad + \sum_{w} \Pr\left[\text{EVENT}_{neq}\right] \Pr\left[w \leftarrow W_n^2 | x \notin \mathbb{K}\right] \Pr\left[x \leftarrow X_n^2 : x \notin \mathbb{K}\right],\end{aligned}$$

which, by the definitions of EVENT_{neq} and set $\mathbb{K}$, leads to

$$\Pr\left[x \leftarrow X_n^2 : x \in \mathbb{K}\right] > 1 - \mu'(n) \tag{6}$$

for some negligible function $\mu'(n)$.

Thus, by (5) and (6), we conclude

$$\Pr\left[x \leftarrow X_n^2 : x \in \mathbb{H} \cap \mathbb{K}\right] > \frac{1}{2p(n)} - \mu(n) - \mu'(n),$$

which means there exist at least one x in the support of X_n^2 that satisfies both conditions of Claim 1, as desired. □

The proof of Claim 2 is based on Claim 1.

*Proof (of **Claim 2**).* Fix a x in the support of X_n^2 that satisfies the two conditions of Claim 1. Note that W_n^2 is uniformed distributed on $R_{L_2}(x)$, and by the first condition of Claim 1, we have a $w_1 \in R_{L_2}(x)$ such that

$$\Pr\left[\langle P(w_1), V^* \rangle (x) = w_1\right] > \frac{1}{2p(n)} - \mu(n).$$

By the second condition of Claim 1, we can obtain another witness $w_2 \in R_{L_2}(x)$, $w_2 \neq w_1$, such that

$$\Pr\left[\langle P(w_2), V^*\rangle(x) = w_1\right] < \mu(n),$$

since otherwise, we would have

$$\begin{aligned}
&\sum_{w} \Pr\left[\langle P(w), V^*\rangle(x) = w' \in R_{L_2}(x) \wedge w' \neq w\right] \Pr\left[w \leftarrow W_n^2 | x\right] \\
\geqslant &\sum_{w_2(\neq w_1)} \Pr\left[\langle P(w_2), V^*\rangle(x) = w_1\right] \Pr\left[w_2 \leftarrow W_n^2 : w_2 \neq w_1\right] \\
= &\sum_{w_2(\neq w_1)} \Pr\left[\langle P(w_2), V^*\rangle(x) = w_1\right] \frac{|R_{L_2}(x)| - 1}{|R_{L_2}(x)|} \\
> &\frac{1}{poly(n)} \cdot \frac{|R_{L_2}(x)| - 1}{|R_{L_2}(x)|},
\end{aligned}$$

which breaks the second condition of Claim 1[6]. Thus we obtain a desired tuple (x, w_1, w_2), completing the proof of Claim 2. □

6.2 Examples of Distributions on Unique Witness Relations

In this subsection, we present some examples of distribution $(\mathbb{X}^1, \mathbb{W}^1)$ on hard unique witness relations that satisfy the "if conditions" of Theorem 3, including distributions over OR-DDH tuples with unique witnesses, the images of lossy trapdoor functions and commitments with unique openings. Thus, for these distributions on unique witness relations, the classic constant-round public-coin proof systems, such as parallelized version of classic 3-round public-coin proofs of [Blu86, GMW91] and ZAPs of [DN00, GOS12], are witness hiding.

Example 1: OR-DDH Tuples with Unique Witnesses. The first example is for distribution $(\mathbb{X}^1, \mathbb{W}^1)$ on hard instances with unique witnesses based on DDH assumption.
DDH assumption: Let Gen be a randomized algorithm that on security parameter 1^n outputs $(\mathbb{G}, g, q)$, where $\mathbb{G}$ is a cyclic group of order q with generator g. Then for a randomly chosen triplet (a, b, c), for every PPT algorithm $\mathcal{A}$ and sufficient large n, there exists a negligible function $\mu(n)$ such that

$$|\Pr[\mathcal{A}(1^n, (\mathbb{G}, g, q), g^a, g^b, g^{ab}) = 1] - \Pr[\mathcal{A}(1^n, (\mathbb{G}, g, q), g^a, g^b, g^c) = 1]| < \mu(n).$$

Consider the following two distribution ensembles $(\mathbb{X}^1, \mathbb{W}^1) = \left\{(X_n^1, W_n^1)\right\}_{n\in\mathbb{N}}$ and $(\mathbb{X}^2, \mathbb{W}^2) = \left\{(X_n^2, W_n^2)\right\}_{n\in\mathbb{N}}$ based on the DDH assumption:

- $(X_n^1, W_n^1) = \{((\mathbb{G}, g, q), x, w) : (\mathbb{G}, g, q) \leftarrow \mathsf{Gen}(1^n)$, the instance x is an OR-DDH tuples $(g^{a_1}, g^{a_2}, g^{a_1 a_2})$ or (g^{b_1}, g^{b_2}, g^c) (where $c \neq b_1 b_2$) with the unique witness $w = (a_1, a_2, a_1 a_2)\}$;
- $(X_n^2, W_n^2) = \{((\mathbb{G}, g, q), x, w) : (\mathbb{G}, g, q) \leftarrow \mathsf{Gen}(1^n)$, the instance x is an OR-DDH tuples $(g^{a_1}, g^{a_2}, g^{a_1 a_2})$ or $(g^{b_1}, g^{b_2}, g^{b_1 b_2})$ with multiple witnesses $w_0 = (a_1, a_2, a_1 a_2)$, $w_1 = (b_1, b_2, b_1 b_2)\}$.

Based on Theorem 3, we have that all the witness hiding protocols for $(\mathbb{X}^2, \mathbb{W}^2)$ above are also witness hiding for $(\mathbb{X}^1, \mathbb{W}^1)$ above, under DDH assumption.

[6] Note that $|R_{L_2}(x)| > 1$.

Example 2: Lossy Trapdoor Functions. We now present another example of distribution ensembles $(\mathbb{X}^1, \mathbb{W}^1)$ based on lossy trapdoor functions.

Recall the definition of lossy trapdoor functions [PW08]. Let n be the security parameter (representing the input length of the function) and $\ell(n)$ be the lossiness of the collection.

Definition 8. *A collection of (m, k)-lossy trapdoor functions is given by a tuple of PPT algorithms $(\mathsf{Gen}, \mathsf{F}, \mathsf{F}^{-1})$. It satisfying the following property:*

- *Easy to sample an injective function with trapdoor: $\mathsf{Gen}_{inj}(\cdot) := \mathsf{Gen}(\cdot, 1)$ outputs (s, t) where s is the description of an injective function f_s and t is its trapdoor, $\mathsf{F}(s, \cdot)$ computes the function $f_s(\cdot)$ over the domain $\{0,1\}^n$, and $\mathsf{F}(t, \cdot)$ computes the function $f_s^{-1}(\cdot)$. If a value y is not in the image of f_s, then $\mathsf{F}(t, y)$ is unspecified.*
- *Easy to sample a lossy function: $\mathsf{Gen}_{lossy}(\cdot) := \mathsf{Gen}(\cdot, 0)$ outputs $(s, \perp)$ where s is the description of function f_s, and $\mathsf{F}(s, \cdot)$ computes the function $f_s(\cdot)$ over the domain $\{0,1\}^m$ whose image has size at most 2^{m-k}.*
- *Hard to distinguish injective and lossy: the first outputs of Gen_{inj} and Gen_{lossy} are computationally indistinguishable.*

Now we consider the following two distribution ensembles $(\mathbb{X}^1, \mathbb{W}^1) = \left\{(X_n^1, W_n^1)\right\}_{n\in\mathbb{N}}$ and $(\mathbb{X}^2, \mathbb{W}^2) = \left\{(X_n^2, W_n^2)\right\}_{n\in\mathbb{N}}$ based on lossy trapdoor function:

- $(X_n^1, W_n^1) := \{((s, y), w) : s \leftarrow \mathsf{Gen}_{\mathrm{inj}}(1^n); w \leftarrow \{0,1\}^n; f_s(w) = y\}$.
- $(X_n^2, W_n^2) := \{((s, y), w) : s \leftarrow \mathsf{Gen}_{\mathrm{lossy}}(1^n); w \leftarrow \{0,1\}^n; f_s(w) = y\}$.

Based on Theorem 3, we have that all the witness hiding protocols for $(\mathbb{X}^2, \mathbb{W}^2)$ above are also witness hiding for $(\mathbb{X}^1, \mathbb{W}^1)$ above, under the existence of lossy trapdoor functions.

Example 3: Commitments with Unique Openings Our third example of distribution ensembles $(\mathbb{X}^1, \mathbb{W}^1)$ is based on mixed commitments [DN02, GOS12].

A mixed commitment scheme is basically a commitment scheme has two different flavors of key generation algorithms. In the binding mode, Gen_1 generates a perfect binding commitment key, in which case a valid commitment uniquely defines one possible message. In the hiding mode, Gen_2 generates a perfect hiding commitment key, in which case the commitment reveals no information whatsoever about the message. Moreover, two kinds of keys are computationally indistinguishable.

Define the following two distribution ensembles $(\mathbb{X}^1, \mathbb{W}^1) = \{(X_n^1, W_n^1)\}_{n\in\mathbb{N}}$ and $(\mathbb{X}^2, \mathbb{W}^2) = \{(X_n^2, W_n^2)\}_{n\in\mathbb{N}}$ based on the mixed commitments:

- $(X_n^1, W_n^1) = \{((x, pk), w) : pk \leftarrow \mathsf{Gen}_1(1^n); m \xleftarrow{\mathrm{R}} M; r \xleftarrow{\mathrm{R}} R; x \leftarrow \mathsf{Com}_{pk}(m; r)\}$.
- $(X_n^2, W_n^2) = \{((x, pk), w) : pk \leftarrow \mathsf{Gen}_2(1^n); m \xleftarrow{\mathrm{R}} M; r \xleftarrow{\mathrm{R}} R; x \leftarrow \mathsf{Com}_{pk}(m; r)\}$.

Assuming the existence of mixed commitments, all the witness hiding protocols for $(\mathbb{X}^2, \mathbb{W}^2)$ above are also witness hiding for $(\mathbb{X}^1, \mathbb{W}^1)$ above.

7 Concluding Remarks and Open problems

We provide the first positive result on the witness hiding security of the classic proof systems for some hard distributions over relations with unique witnesses satisfying certain conditions, including distributions over OR-DDH triplet with unique witnesses, the images of lossy trapdoor functions and commitments with unique openings.

For some hard distributions over relations with unique witnesses that do not satisfy these conditions, such as distributions over DL and RSA problems, we introduce an embedding technique and show that the Schnorr/Guillou-Quisquater protocols are witness hiding unless (ℓ, ε)-tailored instance compression schemes for DL/RSA exist. In the standalone setting, our result improves the positive result of [BP02].

We strongly believe that for Schnorr and Guillou-Quisquater protocols there are no efficient adversary that can break their witness hiding with probability negligibly close to 1, though it is shown in [Pas11] that we cannot prove this based on standard assumption via black-box reduction. We show that such an adversary can be used to construct $(poly, 1 - negl(n))$-tailored instance compression scheme for DL/RSA, which will leads to surprising consequences.

Our results also leave several problems.

1. The first problem is to pinpoint the necessary and sufficient conditions on the hard distribution that admits constant-round public-coin witness hiding protocol.
2. It is known that instance compression scheme is impossible with respect to NP-complete languages, and that the DL and RSA problems are unlikely to be NP-complete. We wonder if tailored instance compression schemes (with moderate parameters) exist for DL/RSA. We think both positive and negative answers to this problem would have interesting consequences.

References

[BCC88] Gilles Brassard, David Chaum, and Claude Crépeau. Minimum disclosure proofs of knowledge. *J. Comput. Syst. Sci*, 37(2):156–189, 1988.

[BGN05] Dan Boneh, Eu-Jin Goh, and Kobbi Nissim. Evaluating 2-dnf formulas on ciphertexts. In *TCC*, volume 3378 of *LNCS*, pages 325–341. Springer, 2005.

[Blu86] Manuel Blum. How to prove a theorem so no one else can claim it. In *ICM*, pages 1444–1451, 1986.

[BNPS03] Mihir Bellare, Chanathip Namprempre, David Pointcheval, and Michael Semanko. The one-more-RSA-inversion problems and the security of chaum's blind signature scheme. *J. Cryptology*, 16(3):185–215, 2003.

[BP02] Mihir Bellare and Adriana Palacio. GQ and schnorr identification schemes: Proofs of security against impersonation under active and concurrent attacks. In *CRYPTO*, volume 2442 of *LNCS*, pages 162–177. Springer, 2002.

[CDS94] Ronald Cramer, Ivan Damgård, and Berry Schoenmakers. Proofs of partial knowledge and simplified design of witness hiding protocols. In *CRYPTO*, volume 839 of *LNCS*, pages 174–187. Springer, 1994.

[CFGV13] Dario Catalano, Dario Fiore, Rosario Gennaro, and Konstantinos Vamvourellis. Algebraic (trapdoor) one-way functions and their applications. In *TCC*, volume 7785 of *LNCS*, pages 680–699. Springer, 2013.

[CLP13] Ran Canetti, Huijia Lin, and Omer Paneth. Public-coin concurrent zero-knowledge in the global hash model. In *TCC'13*, volume 7785 of *LNCS*, pages 80–99. Springer, 2013.

[CRS+07] Ran Canetti, Ronald L. Rivest, Madhu Sudan, Luca Trevisan, Salil P. Vadhan, and Hoeteck Wee. Amplifying collision resistance: A complexity-theoretic treatment. In *Advances in Cryptology - CRYPTO 2007, 27th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2007, Proceedings*, pages 264–283, 2007.

[DN00] Cynthia Dwork and Moni Naor. Zaps and their applications. In *FOCS*, pages 283–293. IEEE, 2000.

[DN02] Ivan Damgård and Jesper Buus Nielsen. Perfect hiding and perfect binding universally composable commitment schemes with constant expansion factor. In *CRYPTO*, volume 2442 of *LNCS*, pages 581–596. Springer, 2002.

[Dru15] Andrew Drucker. New limits to classical and quantum instance compression. *SIAM J. Comput.*, 44(5):1443–1479, 2015.

[FS90] Uriel Feige and Adi Shamir. Witness indistinguishable and witness hiding protocols. In *STOC*, pages 416–426. ACM, 1990.

[FS11] Lance Fortnow and Rahul Santhanam. Infeasibility of instance compression and succinct pcps for NP. *J. Comput. Syst. Sci.*, 77(1):91–106, 2011.

[GMR89] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM J. Comput.*, 18(1):186–208, 1989.

[GMW91] Oded Goldreich, Silvio Micali, and Avi Wigderson. Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems. *J. ACM*, 38(3):690–728, 1991.

[Gol01] Oded Goldreich. *The Foundations of Cryptography - Volume 1, Basic Techniques*. Cambridge University Press, 2001.

[Gol04] Oded Goldreich. *The Foundations of Cryptography - Volume 2, Basic Applications*. Cambridge University Press, 2004.

[GOS12] Jens Groth, Rafail Ostrovsky, and Amit Sahai. New techniques for noninteractive zero-knowledge. *J. ACM*, 59(3):11, 2012.

[GQ88] Louis C. Guillou and Jean-Jacques Quisquater. A practical zero-knowledge protocol fitted to security microprocessor minimizing both transmission and memory. In *EUROCRYPT*, volume 330 of *LNCS*, pages 123–128. Springer, 1988.

[HN10] Danny Harnik and Moni Naor. On the compressibility of *NP* instances and cryptographic applications. *SIAM J. Comput.*, 39(5):1667–1713, 2010.

[HRS09] Iftach Haitner, Alon Rosen, and Ronen Shaltiel. On the (im)possibility of arthur-merlin witness hiding protocols. In *TCC*, volume 5444 of *LNCS*, pages 220–237. Springer, 2009.

[Mau15] Ueli Maurer. Zero-knowledge proofs of knowledge for group homomorphisms. *Designs, Codes and Cryptography*, 77(2–3):663–676, 2015.

[Pas06] Rafael Pass. Parallel repetition of zero-knowledge proofs and the possibility of basing cryptography on NP-hardness. In *IEEE CCC*, pages 96–110. IEEE, 2006.

[Pas11] Rafael Pass. Limits of provable security from standard assumptions. In *STOC*, pages 109–118. ACM, 2011.

[Pie07] Krzysztof Pietrzak. Non-trivial black-box combiners for collision-resistant hash-functions don't exist. In *Advances in Cryptology - EUROCRYPT 2007, 26th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Barcelona, Spain, May 20-24, 2007, Proceedings*, pages 23–33, 2007.

[Pie08] Krzysztof Pietrzak. Compression from collisions, or why CRHF combiners have a long output. In *Advances in Cryptology - CRYPTO 2008, 28th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2008. Proceedings*, pages 413–432, 2008.

[PW08] Chris Peikert and Brent Waters. Lossy trapdoor functions and their applications. In *STOC*, pages 187–196. ACM, 2008.

[Sch89] Claus-Peter Schnorr. Efficient identification and signatures for smart cards. In *EUROCRYPT*, volume 434 of *LNCS*, pages 688–689. Springer, 1989.

[Sho97] Victor Shoup. Lower bounds for discrete logarithms and related problems. In *EUROCRYPT*, volume 1233 of *LNCS*, pages 256–266. Springer, 1997.

[ZZC+14] Jiang Zhang, Zhenfeng Zhang, Yu Chen, Yanfei Guo, and Zongyang Zhang. Black-box separations for one-more (static) CDH and its generalization. In *ASIACRYPT*, volume 8874 of *LNCS*, pages 366–385. Springer, 2014.