On the Non-malleability of the Fiat-Shamir Transform

Sebastian Faust[1,⋆], Markulf Kohlweiss[2],
Giorgia Azzurra Marson[3,⋆⋆], and Daniele Venturi[1,⋆]

[1] Aarhus University
[2] Microsoft Research
[3] Technische Universität Darmstadt

Abstract. The Fiat-Shamir transform is a well studied paradigm for removing interaction from public-coin protocols. We investigate whether the resulting non-interactive zero-knowledge (NIZK) proof systems also exhibit non-malleability properties that have up to now only been studied for NIZK proof systems in the common reference string model: first, we formally define simulation soundness and a weak form of simulation extraction in the random oracle model (ROM). Second, we show that in the ROM the Fiat-Shamir transform meets these properties under lenient conditions. A consequence of our result is that, in the ROM, we obtain truly efficient non malleable NIZK proof systems essentially for free. Our definitions are sufficient for instantiating the Naor-Yung paradigm for CCA2-secure encryption, as well as a generic construction for signature schemes from hard relations and simulation-extractable NIZK proof systems. These two constructions are interesting as the former preserves both the leakage resilience and key-dependent message security of the underlying CPA-secure encryption scheme, while the latter lifts the leakage resilience of the hard relation to the leakage resilience of the resulting signature scheme.

1 Introduction

Zero-knowledge proof systems [26] are a powerful tool for designing cryptographic primitives and protocols. They force malicious parties to behave according to specification while allowing honest parties to protect their secrets. Non-interactive zero-knowledge (NIZK) proofs [10] consist of a single proof message passed from the prover to the verifier. They are particularly useful for designing public-key encryption and signature schemes as the proof can be added to the

⋆ The authors acknowledge support from the Danish National Research Foundation and The National Science Foundation of China (under the grant 61061130540) for the Sino-Danish Center for the Theory of Interactive Computation, and also from the CFEM research center (supported by the Danish Strategic Research Council) within which part of this work was performed.

⋆⋆ Supported by CASED. Part of the work done while a Master Student in Rome, La Sapienza.

S. Galbraith and M. Nandi (Eds.): INDOCRYPT 2012, LNCS 7668, pp. 60–79, 2012.
© Springer-Verlag Berlin Heidelberg 2012

ciphertext and signature respectively. Understanding the most efficient NIZK proofs that are sufficiently strong, i.e., sufficiently non-malleable, for building signature and encryption schemes with strong security properties is thus of fundamental importance in cryptography. It was shown by Goldreich and Oren [25] that NIZK proofs are unattainable in the standard model. To avoid this impossibility result, one must rely on additional assumptions, such as common reference strings [9] (CRS model) or idealizations of hash functions [7] (random oracle model, ROM).

With the aim of finding the "right" definition in the non-interactive case, several flavors of non-malleability [19] have been introduced for NIZK in the CRS model [37,38,24,31]. The notion of *simulation soundness*, which bridges soundness and zero knowledge, guarantees that soundness holds even after seeing accepting proofs, for both true and *false* statements, produced by the simulator. This strengthened soundness notion was first proposed by Sahai in [37], and later improved by De Santis et al. [38]. The notion of simulation extraction [38,28] in addition requires that accepting proofs allow to extract witnesses. Different variants of simulation extraction have been proposed by [15,18].

Until recently, zero-knowledge in general and NIZK in particular were considered to be primarily of theoretical interest. Significant exceptions being efficient Σ-protocols [16,17] and their non-interactive relatives based on the Fiat-Shamir (FS) transform [21]. A Σ-protocol is a three-move interactive scheme where the prover sends the first message and the verifier sends a random challenge as the second message. In a nutshell, the Fiat-Shamir transform removes the interaction by computing the challenge as the hash value of the first message and the theorem that is being proven. Σ-protocols and the Fiat-Shamir transform are widely used in the construction of efficient identification [21], anonymous credential [14], signature [35,1], e-voting schemes [8], and many other cryptographic constructions [11,6,23].

Most work on the provable security of zero-knowledge has, however, been conducted either on interactive proof systems in the plain model or on NIZK in the CRS model, while practitioners often preferred Fiat-Shamir based NIZK proofs for their simplicity and efficiency. The use of the Fiat-Shamir transform was most thoroughly explored in the security proofs of signature schemes in the random oracle model [35,1], but was otherwise often used heuristically. The question thus arises whether one can lay sound foundations for the FS transform in the light of recent research on CRS-based NIZKs. To this end, we provide non-malleability definitions for NIZK in the random oracle model that closely follow the established CRS-based definitions [28]. An earlier result oriented in the same direction, but concerning a Σ-protocol for a *specific* language,[1] was given by Fouque and Pointcheval [23]. Their proof strategy relies on the forking lemma [35] and (implicitly) on the fact that the Σ-protocol they consider has a particular property called *strong special honest-verifier zero-knowledge* (SS-HVZK). Since there exist Σ-protocols that do not satisfy the SS-HVZK property, Fouque

[1] This is the language used in the Naor-Yung transform when the underlying encryption is the ElGamal scheme.

and Pointcheval's proof cannot be immediately extended to the general case. Moreover, we make the random oracle explicit in our definition, which is crucial as definitions in the random oracle model can be brittle [40].

Our first observation is that much less is required to show simulation soundness for any FS-NIZK proof. Namely, in the random oracle model, simulation soundness simply follows from the soundness and the HVZK properties of the underlying interactive protocol. In particular, it is neither necessary to rely on the forking lemma, nor on the strong property of SS-HVZK. We also show that the proof strategy of [23], when generalized properly to any Σ-protocol, yields something more than just simulation soundness. In fact, one gets some form of simulation extractability, which we call *weak simulation extractability*. In a nutshell, *full simulation extractability* requires that even after seeing many simulated proofs, whenever an adversary outputs a new accepted proof, we can build an algorithm to extract a valid witness. Sometimes, such a strong extraction property is called *online* extraction [22] because the extractor outputs a witness directly after receiving the adversary's proof. In comparison, our notion is weaker in that it allows the extractor to fully control the adversary (i.e., rewind it).

Our contribution. Our contributions are threefold. First, we formally define the notions of zero-knowledge (which holds trivially for the Fiat-Shamir transform), simulation soundness, and simulation extractability for NIZKs in the random oracle model. Second, we show that simulation soundness and a weak form of simulation extractability come for free if one uses the FS-transform for turning Σ-protocols into NIZK proof systems. Third, we investigate the consequences of this result by showing that our definitions are sufficient for instantiating the Naor-Yung paradigm for constructing CCA2-secure encryption schemes, and generic construction for signature schemes from hard relations and simulation-extractable NIZK proof systems [18]. These two constructions are particularly interesting as the former preserves both leakage resilience and key-dependent message security of the underlying CPA-secure encryption scheme, while the latter lifts the leakage resilience of the hard relation to the leakage resilience of the resulting signature scheme. To our knowledge, these are the most efficient schemes having such properties, if one is willing to rely on the ROM.[2]

Related work. The only other efficient transform for Σ-protocols yielding simulation soundness (again in the random oracle model) is Fischlin's transform [22] which is designed with the purpose of online extraction and is less efficient than the classical Fiat-Shamir transform. Therefore, it would be interesting to investigate whether Fischlin's transform achieves a stronger form of simulation extractability. We notice that in the interactive case, a general transform from any Σ-protocol to an (unbounded) simulation-sound Σ-protocol using one-time signatures has been proposed [24]. In the common reference string model the most efficient simulation-sound or simulation extractable NIZK proof system are

[2] In particular we obtain as a special case the Alwen et al. [4] leakage-resilient signature scheme based on the Okamoto identification scheme.

based on Groth-Sahai proofs [29]. One has however to pay the price of proving a structure-preserving CCA secure encryption [18] (for true-simulation extraction) or a structure-preserving signature scheme [3] (for full simulation extraction).

2 Preliminaries

Notation. Let k be a security parameter. A function ν is called *negligible* if $\nu(k) \leq k^{-c}$ for any $c > 0$ and sufficiently large k. Given two functions f, g, we write $f \approx g$ if there exists a negligible function ν such that $|f(k) - g(k)| < \nu(k)$. Given an algorithm $\mathcal{A}$, $y \leftarrow \mathcal{A}(x)$ means that y is the output of $\mathcal{A}$ on input x; when $\mathcal{A}$ is randomized, then y is a random variable. We write $\mathcal{A}^H$ to denote the fact that $\mathcal{A}$ has oracle access to some function H. PPT stands for probabilistic-polynomial time. A *decision problem* related to a language $\mathcal{L} \subseteq \{0,1\}^*$ consists in determining whether a string x is in $\mathcal{L}$ or not. Given an instance x, we say that $\mathcal{A}$ decides (or recognizes) $\mathcal{L}$ if, after a finite number of steps, the algorithm halts and outputs $\mathcal{A}(x) = 1$ if $x \in \mathcal{L}$, otherwise $\mathcal{A}(x) = 0$. (Sometimes, we may call "theorem" a string belonging to the language at hand.) We can associate to any NP-language $\mathcal{L}$ a polynomial-time recognizable relation $\mathcal{R}_{\mathcal{L}}$ defining $\mathcal{L}$ itself, that is $\mathcal{L} = \{x : \exists w \text{ s.t. } (x, w) \in \mathcal{R}_{\mathcal{L}}\}$, where $|w| \leq poly(|x|)$. The string w is called a *witness* or *certificate* for membership of $x \in \mathcal{L}$. For NP, w corresponds to the non-deterministic choices made by $\mathcal{A}$.

Interactive protocols. An interactive proof system (IPS) for membership in $\mathcal{L}$ is a two-party protocol, where a prover wants to convince an efficient verifier that a string x belongs to $\mathcal{L}$. In a zero-knowledge interactive proof system, a prover $\mathcal{P}$ can convince a verifier $\mathcal{V}$ that $x \in \mathcal{L}$ without revealing anything beyond the fact that the statement is indeed true. Informally, this means that $\mathcal{V}$ cannot exploit the interaction with $\mathcal{P}$ for gaining extra-knowledge. Such a property is formalized by requiring the existence of an efficient algorithm $\mathcal{S}$, the *zero-knowledge simulator*, which produces messages indistinguishable from conversations between an honest prover $\mathcal{P}$ and a malicious verifier $\mathcal{V}^*$. Besides the zero-knowledge property, any proof system satisfies two standard requirements: proving true statements is always possible, while it should be infeasible to convince the verifier to accept a false statement as correct. These two conditions are called *completeness* and *soundness* respectively. Related to the concept of interactive proof systems, but even more subtle, is the notion of *proof of knowledge*. In a proof of knowledge (PoK), $\mathcal{P}$ wants to convince $\mathcal{V}$ that he *knows* a secret witness which implies the validity of some assertion, and not merely that the assertion is true. To formalize the fact that a prover actually "knows something", we require that there exists an efficient algorithm $\mathcal{E}$, called *knowledge extractor*, that when given complete access to the program of the prover can extract the witness.

An IPS or an interactive PoK is called *public-coin* when the verifier's moves consist merely of tossing coins and sending their outcomes to the prover. (In contrast, in a *private-coin* IPS the verifier does not need to show the outcome of the coins to the prover [27].) We are mainly interested in a specific class of

Prover $\mathcal{P}(x, w; 1^k)$		Verifier $\mathcal{V}(x; 1^k)$
$\alpha \leftarrow \mathcal{P}_0(x, w; \rho)$	$\xrightarrow{\alpha}$	
	$\xleftarrow{\beta}$	$\beta \xleftarrow{\$} \mathcal{V}_0(x, \alpha)$
$\gamma \leftarrow \mathcal{P}_1(\alpha, \beta, x, w; \rho)$	$\xrightarrow{\gamma}$	Accept iff $\mathcal{V}_1(x, \alpha, \beta, \gamma) = 1$

Fig. 1. A Σ-protocol for a language $\mathcal{L}$

public-coin interactive PoK systems for NP-languages, called Σ-protocols. Here, the parties involved share a string x belonging to a language $\mathcal{L} \in$ NP and the prover also holds a witness w for membership of $x \in \mathcal{L}$. Thus, the prover $\mathcal{P}$ wants to convince the verifier $\mathcal{V}$ that it "knows" a witness w for x, i.e. that x is in the language, without revealing the witness itself. Σ-protocols have a 3-move shape where the first message α, called *commitment*, is sent by the prover and then, alternatively, the parties exchange the other messages β and γ, called (respectively) *challenge* and *response*. The interaction is depicted in Figure 1. Besides the standard properties held by any IPS, Σ-protocols satisfy a flavour of zero-knowledge — called *honest-verifier zero knowledge* (HVZK) — saying that an *honest* verifier taking part in the protocol does not learn anything beyond the validity of the theorem being proven.

Definition 1 (Σ-protocols). *A Σ-protocol $\Sigma = (\mathcal{P}, \mathcal{V})$ for an* NP*-language $\mathcal{L}$ is a three-round public-coin IPS where $\mathcal{P} = (\mathcal{P}_0, \mathcal{P}_1)$ and $\mathcal{V} = (\mathcal{V}_0, \mathcal{V}_1)$ are PPT algorithms, with the following additional proprieties:*

Completeness. *If $x \in \mathcal{L}$, any proper execution of the protocol between $\mathcal{P}$ and $\mathcal{V}$ ends with the verifier accepting $\mathcal{P}$'s proof.*

Honest-Verifier Zero Knowledge (HVZK). *There exists an efficient algorithm $\mathcal{S}$, called zero-knowledge simulator, such that for any PPT distinguisher $\mathcal{D} = (\mathcal{D}_0, \mathcal{D}_1)$ and for any $(x, w) \in \mathcal{R}_\mathcal{L}$, the view of the following two experiments, real and simulated, are computationally indistinguishable:*

Experiment $\mathsf{Exp}^{\text{REAL}}_{\Sigma,\mathcal{D}}(1^k)$	**Experiment** $\mathsf{Exp}^{\text{SIM}}_{\Sigma,\mathcal{D}}(\mathcal{S}, 1^k)$
$(x, w, \delta) \leftarrow \mathcal{D}_0(1^k)$	$(x, w, \delta) \leftarrow \mathcal{D}_0(1^k)$
$\pi \leftarrow \langle \mathcal{P}(x, w; 1^k), \mathcal{V}(x; 1^k) \rangle$	$\pi \leftarrow \mathcal{S}(x, 1^k)$
Output $\mathcal{D}_1(\pi, \delta)$	*Output* $\mathcal{D}_1(\pi, \delta)$

where $\langle \mathcal{P}(x, w), \mathcal{V}(x) \rangle$ denotes the verdict returned at the end of the interaction between $\mathcal{P}$ and $\mathcal{V}$ on common input x and private input w.

Soundness. *If $x \notin \mathcal{L}$ then any malicious (even unbounded) prover $\mathcal{P}^*$ is accepted only with negligible probability.*

Special soundness. *There exists an efficient algorithm $\mathcal{E}$, called special extractor, such that given two accepting conversations (α, β, γ) and $(\alpha, \beta', \gamma')$ for a string x, where $\beta \neq \beta'$, then $w \leftarrow \mathcal{E}(\alpha, \beta, \gamma, \beta', \gamma', x)$ is such that $(x, w) \in \mathcal{R}_\mathcal{L}$.*

The special soundness property is strong enough to imply both soundness and that Σ-protocols are PoK [17]. Sometimes Σ-protocols are required to meet

stronger notions of HVZK. We discuss these notions and implications and non implications between them in the full version [20].

A non-standard condition that many Σ-protocols satisfy, introduced by Fischlin in [22], requires that responses are quasi unique, i.e. given an accepting proof it should be infeasible to find a new valid response for that proof.

Definition 2 (Quasi unique responses). *A Σ-protocol has* quasi unique responses *if for any PPT $\mathcal{A}$ and for any security parameter k it holds:*

$$\mathrm{Prob}[(x,\alpha,\beta,\gamma,\gamma') \leftarrow \mathcal{A}(1^k) \ : \ \mathcal{V}(x,\alpha,\beta,\gamma) = \mathcal{V}(x,\alpha,\beta,\gamma') = 1 \ \wedge \ \gamma \neq \gamma'] \approx 0.$$

A Σ-protocol has *unique responses* if the probability above is zero. The latter condition, defined by Unruh in [39], is also known as *strict soundness*.

Min-entropy of commitments. Following [1,2], we use the concept of min-entropy to measure how likely it is for a commitment to collide with a fixed value.

Definition 3 (Min-entropy of commitment). *Let k be a security parameter and $\mathcal{L}$ be an* NP*-language with relation $\mathcal{R}_{\mathcal{L}}$. Consider a pair $(x,w) \in \mathcal{R}_{\mathcal{L}}$ and let $(\mathcal{P},\mathcal{V})$ be an arbitrary three-round IPS. Denote with $\mathsf{Coins}(k)$ the set of coins used by the prover and consider the set $A(x,w) = \{\mathcal{P}_0(x,w;\rho) : \rho \leftarrow \mathsf{Coins}(k))\}$ of all possible commitments associated to w. The min-entropy function associated to $(\mathcal{P},\mathcal{V})$ is defined as $\varepsilon(k) = \min_{(x,w)}(-\log_2 \mu(x,w))$,where the minimum is taken over all possible (x,w) drawn from $\mathcal{R}_{\mathcal{L}}$ and $\mu(x,w)$ is the maximum probability that a commitment takes on a particular value, i.e., $\mu(x,w) = \max_{\alpha \in A(x,w)}(Prob[\mathcal{P}_0(x,w;\rho) = \alpha : \rho \leftarrow \mathsf{Coins}(k))])$.*

We say that $(\mathcal{P},\mathcal{V})$ is *non-trivial* if $\varepsilon(k) = \omega(\log(k))$ is super-logarithmic in k. Often, the commitment is drawn uniformly from some set. In order for $(\mathcal{P},\mathcal{V})$ to be non-trivial, this set must have size exponential in k. Notice that most of natural Σ-protocols meet such a condition and, in fact, non-triviality is quite easy to achieve, e.g. by appending redundant random bits to the commitment.

Forking lemma. To prove our second main result, we make use of the following version of the forking lemma, which appeared in [6].

Lemma 1 (General forking lemma). *Fix an integer Q and a set $\mathcal{H}$ of size $h \geq 2$. Let P be a randomized program that on input $y, h_1, \ldots, h_Q$ returns a pair, the first element of which is an integer in the range $0, \ldots, Q$ and the second element of which we refer to as a* side output*. Let IG be a randomized algorithm that we call the input generator. The accepting probability of P, denoted* acc*, is defined as the probability that $J \geq 1$ in the experiment $y \leftarrow \mathsf{IG}; h_1, \ldots, h_Q \leftarrow \mathcal{H}; (J,s) \leftarrow \mathsf{P}(y, h_1, \ldots, h_Q)$.*

The forking algorithm F_{P} associated to P is the randomized algorithm that on input y proceeds as follows.

Algorithm $\mathsf{F}_{\mathsf{P}}(y)$

Pick coins ρ for P at random

$h_1, \ldots, h_Q \leftarrow \mathcal{H}$
$(I, s) \leftarrow \mathsf{P}(y, h_1, \ldots, h_Q; \rho)$
If $I = 0$ *return* $(0, \perp, \perp)$
$h'_I, \ldots, h'_Q \leftarrow \mathcal{H}$
$(I', s') \leftarrow \mathsf{P}(y, h_1, \ldots, h_{I-1}, h'_I, \ldots, h'_Q; \rho)$
If $(I = I') \wedge (h_I \neq h'_I)$ *return* $(1, s, s')$ *else return* $(0, \perp, \perp)$

Let $\text{ext} = \text{Prob}[b = 1 : y \leftarrow \mathsf{IG}; (b, s, s') \leftarrow \mathsf{F}_{\mathsf{P}}(y)]$, *then* $\text{ext} \geq \text{acc}\left(\frac{\text{acc}}{Q} - \frac{1}{h}\right)$.

3 Properties of NIZKs in the Random Oracle Model

Removing interaction. The Fiat-Shamir transform was originally designed to turn three-round identification schemes into efficient signature schemes. As Σ-protocols are an extension of three-round identification schemes, it is not surprising that they can be considered as a starting point for the Fiat-Shamir transform. The Fiat-Shamir paradigm applies to any Σ-protocol (and more generally to any three-round public-coin proof system): We start from an interactive protocol $(\mathcal{P}, \mathcal{V})$ and remove the interaction between $\mathcal{P}$ and $\mathcal{V}$ by replacing the challenge, chosen at random by the verifier, with a hash value $H(\alpha, x)$ computed by the prover, where H is a hash function modeled as a random oracle. Thus, the interactive protocol $(\mathcal{P}, \mathcal{V})$ is turned into a non-interactive one: The resulting protocol, denoted $(\mathcal{P}^H, \mathcal{V}^H)$, is called *Fiat-Shamir proof system.*

Throughout this paper, we refer to the so called *explicitly programmable* random oracle model [40] (EPROM) where the simulator is allowed to program the random oracle explicitly. We model this by defining the zero-knowledge simulator $\mathcal{S}$ of a non-interactive zero-knowledge proof system as a stateful algorithm that can operate in two modes: $(h_i, st) \leftarrow \mathcal{S}(1, st, q_i)$ takes care of answering random oracle queries (usually by lazy sampling) while $(\pi, st) \leftarrow \mathcal{S}(2, st, x)$ simulates the actual proof. Note that calls to $\mathcal{S}(1, \cdots)$ and $\mathcal{S}(2, \cdots)$ share the common state st that is updated after each operation.

Definition 4 (Unbounded non-interactive zero knowledge). *Let $\mathcal{L}$ be a language in* NP*. Denote with $(\mathcal{S}_1, \mathcal{S}_2)$ the oracles such that $S_1(q_i)$ returns the first output of $(h_i, st) \leftarrow \mathcal{S}(1, st, q_i)$ and $\mathcal{S}_2(x, w)$ returns the first output of $(\pi, st) \leftarrow \mathcal{S}(2, st, x)$ if $(x, w) \in \mathcal{R}_{\mathcal{L}}$. We say a protocol $(\mathcal{P}^H, \mathcal{V}^H)$ is a NIZK proof for language $\mathcal{L}$ in the random oracle model, if there exists a PPT simulator $\mathcal{S}$ such that for all PPT distinguishers $\mathcal{D}$ we have*

$$\text{Prob}[\mathcal{D}^{H(\cdot), \mathcal{P}^H(\cdot,\cdot)}(1^k) = 1] \approx \text{Prob}[\mathcal{D}^{\mathcal{S}_1(\cdot), \mathcal{S}_2(\cdot,\cdot)}(1^k) = 1],$$

where both $\mathcal{P}$ and $\mathcal{S}_2$ oracles output $\perp$ if $(x, w) \notin \mathcal{R}_{\mathcal{L}}$.

A well known fact is that, in the random oracle model, the Fiat-Shamir transform allows to efficiently design digital signature schemes [21] and non-interactive zero-knowledge protocols. In fact, an appealing characteristic of this transform is that many properties of the starting protocol are still valid after applying

it. In particular, it has been proven that the Fiat-Shamir transform turns any three-round public-coin zero-knowledge interactive proof system into a NIZK proof system [7]. It is straightforward to prove that the same holds when the starting protocol is ZK only with respect to a honest verifier, as stated in the following Theorem.

Theorem 1 (Fiat-Shamir NIZKs). *Let k be a security parameter. Consider a non-trivial three-round public-coin honest-verifier zero-knowledge interactive proof system $(\mathcal{P}, \mathcal{V})$ for a language $\mathcal{L} \in$ NP. Let H be a function with range equal to the space of the verifier's coins. In the random oracle model the proof system $(\mathcal{P}^H, \mathcal{V}^H)$, derived from $(\mathcal{P}, \mathcal{V})$ by applying the Fiat-Shamir transform, is unbounded non-interactive zero-knowledge.*

Proof (sketch). To prove that the proof system $(\mathcal{P}^H, \mathcal{V}^H)$ is non-interactive zero-knowledge it is sufficient to show that there exists a simulator $\mathcal{S}$ as required in Definition 4. This can be done by invoking the HVZK simulator associated with the underlying interactive proof system. In particular, $\mathcal{S}$ works as follows:

- To answer query $q = (x, \alpha)$ to $\mathcal{S}_1$, $\mathcal{S}(1, st, q)$ lazily samples a lookup table $\mathcal{T}_H$ kept in state st. It checks whether $\mathcal{T}_H(q)$ is already defined. If this is the case, it returns the previously assigned value; otherwise it returns and sets a fresh random value (of the appropriate length).
- To answer query x to $\mathcal{S}_2$ (respectively $\mathcal{S}_2'$), $\mathcal{S}(2, st, x)$ calls the HVZK simulator of $(\mathcal{P}, \mathcal{V})$ on input x to obtain a proof (α, β, γ). Then, it updates $\mathcal{T}_H$ in such a way that $\beta = \mathcal{T}_H(x, \alpha)$. If $\mathcal{T}_H$ happens to be already defined on this input, $\mathcal{S}$ returns failure and aborts.

We call this simulator *canonical*. The main result of Fiat-Shamir [21] (expressed for their particular identification protocol) is that $\mathcal{S}$ is a "good" NIZK simulator. The crucial step in the proof is that the starting protocol $(\mathcal{P}, \mathcal{V})$ is non-trivial (cf. Definition 3), thus the probability of failure in each of the queries to $\mathcal{S}_2'$ is upper-bounded by $\text{Prob}[\text{failure}] \leq 2^{-\varepsilon(k)}$, which is negligible in k. □

Simulation soundness. The soundness property of a proof system ensures that no malicious prover can come up with an accepting proof for a string that does not belong to the language in question (i.e., for a false theorem). However, it is not clear whether this condition still holds *after* the attacker observes valid proofs for adaptively chosen (true or false) statements. The notion of *simulation soundness* deals with this case.

Definition 5 (Unbounded simulation soundness). *Let $\mathcal{L}$ be a language in* NP. *Consider a proof system $(\mathcal{P}^H, \mathcal{V}^H)$ for $\mathcal{L}$, with zero-knowledge simulator $\mathcal{S}$. Denote with $(\mathcal{S}_1, \mathcal{S}_2')$ the oracles such that $S_1(q_i)$ returns the first output of $(h_i, st) \leftarrow \mathcal{S}(1, st, q_i)$ and $\mathcal{S}_2'(x)$ returns the first output of $(\pi, st) \leftarrow \mathcal{S}(2, st, x)$. We say that $(\mathcal{P}^H, \mathcal{V}^H)$ is simulation sound with respect to $\mathcal{S}$ in the random oracle model, if for all PPT adversaries $\mathcal{A}$ the following holds:*

$$\text{Prob}[(x^\star, \pi^\star) \leftarrow \mathcal{A}^{\mathcal{S}_1(\cdot), \mathcal{S}_2'(\cdot)} : (x^\star, \pi^\star) \notin \mathcal{T} \ \wedge \ x^\star \notin \mathcal{L} \ \wedge \ \mathcal{V}^{\mathcal{S}_1}(x^\star, \pi^\star) = 1] \approx 0,$$

where $\mathcal{T}$ is the list of pairs (x_i, π_i), i.e., respectively queries asked to and proofs returned by the simulator.

We stress that the above definition relies crucially on the zero-knowledge property of $(\mathcal{P}^H, \mathcal{V}^H)$, as we use a probability experiment that defines a property of $\mathcal{S}$ to define a property about $(\mathcal{P}^H, \mathcal{V}^H)$. In particular the definition is most meaningful for a simulator $\mathcal{S}$ for which the simulation of the random oracle of $\mathcal{S}_1$ is consistent with a truly random oracle H. Also note that $\mathcal{S}_2'$ allows $\mathcal{A}$ to ask for simulated proofs of false statements.

The possibility to request proofs of false statements has an interesting consequence: simulation soundness holds only with respect to specific simulators and not in general for all NIZK simulators. In particular, one can construct a NIZK proof system $(\mathcal{P}^H, \mathcal{V}^H)$ that is simulation sound with respect to a simulator $\mathcal{S}$ but for which there exists a valid NIZK simulator $\hat{\mathcal{S}}$, such that $(\mathcal{P}^H, \mathcal{V}^H)$ cannot be simulation sound with respect to $\hat{\mathcal{S}}$. To see this, consider a $\mathcal{V}^H$ that accepts all proofs if $H(0) = 0$. $\hat{\mathcal{S}}$ simulates a consistent random oracle until it receives a proof of a false statement (one of which could be hard-coded in $\hat{\mathcal{S}}$ or easy to recognize) at which point it sets $\mathcal{T}_H(0) = 0$. Note that a similar counterexample exists for CRS-based NIZK [28]: $\hat{\mathcal{S}}_2$ can simply return the simulation trapdoor when queried on a false statement.

Simulation extractability. Combining simulation soundness and knowledge extraction, we may require that even after seeing (polynomially) many simulated proofs, whenever $\mathcal{A}$ makes a new proof it is possible to extract a witness. This property is called simulation extractability, and implies simulation soundness. Indeed, if we can extract a witness from the adversary's proof even with small probability, then obviously the statement must belong to the language in question. We introduce a weaker flavor of simulation extractability which we call *weak simulation extractability*. The main difference with full simulation extractability is that the extractor $\mathcal{E}_{\mathcal{A}}$ is now given complete control over the adversary $\mathcal{A}$, meaning that it is allowed to rewind $\mathcal{A}$ and gets to see the answers of $(\mathcal{S}_1, \mathcal{S}_2')$. Moreover, we require that if $\mathcal{A}$ outputs an accepting proof with some probability, then $\mathcal{E}_{\mathcal{A}}$ can extract with almost the same probability.

Definition 6 (Weak simulation extractability). *Let $\mathcal{L}$ be a language in* NP*. Consider a NIZK proof system $(\mathcal{P}^H, \mathcal{V}^H)$ for $\mathcal{L}$ with zero-knowledge simulator $\mathcal{S}$. Let $(\mathcal{S}_1, \mathcal{S}_2')$ be oracles returning the first output of $(h_i, st) \leftarrow \mathcal{S}(1, st, q_i)$ and $(\pi, st) \leftarrow \mathcal{S}(2, st, x)$ respectively. We say that $(\mathcal{P}^H, \mathcal{V}^H)$ is* weakly simulation-extractable *with* extraction error ν *and with respect to $\mathcal{S}$ in the random oracle model, if for all PPT adversaries $\mathcal{A}$ there exists an efficient algorithm $\mathcal{E}_{\mathcal{A}}$ with access to the answers $\mathcal{T}_H$, $\mathcal{T}$ of $(\mathcal{S}_1, \mathcal{S}_2')$ respectively such that the following holds. Let:*

$$\mathrm{acc} = \mathrm{Prob}\big[(x^\star, \pi^\star) \leftarrow \mathcal{A}^{\mathcal{S}_1(\cdot), \mathcal{S}_2'(\cdot)}(1^k; \rho) : (x^\star, \pi^\star) \notin \mathcal{T}; \mathcal{V}^{\mathcal{S}_1}(x^\star, \pi^\star) = 1\big]$$

$$\mathrm{ext} = \mathrm{Prob}\big[(x^\star, \pi^\star) \leftarrow \mathcal{A}^{\mathcal{S}_1(\cdot), \mathcal{S}_2'(\cdot)}(1^k; \rho); \\ w^\star \leftarrow \mathcal{E}_{\mathcal{A}}(x^\star, \pi^\star; \rho, \mathcal{T}_H, \mathcal{T}) : \ (x^\star, \pi^\star) \notin \mathcal{T}; (x^\star, w^\star) \in \mathcal{R}_{\mathcal{L}}\big],$$

where the probability space in both cases is over the random choices of $\mathcal{S}$ *and the adversary's random tape* ρ*. Then, there exist a constant* $d > 0$ *and a polynomial* p *such that whenever* $\mathrm{acc} \geq \nu$*, we have* $\mathrm{ext} \geq \frac{1}{p}(\mathrm{acc} - \nu)^d$.

The above definition is inspired by similar notions in the context of proofs of knowledge [5,30,39]. The value ν is called *extraction error* of the proof system. We omit for better readability that values $\mathrm{acc}, \mathrm{ext}, p, \nu$ all depend on the security parameter k. Note that a non-negligible extractor error can be made exponentially small by sequential repetitions (see the full version [20] for a proof).

Proposition 1 (Extraction error amplification). *Let* $(\mathcal{P}^H, \mathcal{V}^H)$ *be a weakly simulation extractable NIZK proof system with extraction error* ν*. Then, the proof system* $(\mathcal{P}'^H, \mathcal{V}'^H)$ *obtained by repeating sequentially* $(\mathcal{P}^H, \mathcal{V}^H)$ *for a number* n *of times yields a weakly simulation extractable NIZK proof system with extraction error* ν^n*.*

It is useful to look at the relation between weak simulation extractability and the following stronger property modeling online-extraction.

Definition 7 (Full Simulation extractability). *Let* $\mathcal{L}$ *be a language in* NP*. Consider a NIZK proof system* $(\mathcal{P}^H, \mathcal{V}^H)$ *for* $\mathcal{L}$ *with simulator* $\mathcal{S}$*. Let* $(\mathcal{S}_1, \mathcal{S}_2')$ *be oracles returning the first output of* $(h_i, st) \leftarrow \mathcal{S}(1, st, q_i)$ *and* $(\pi, st) \leftarrow \mathcal{S}(2, st, x)$ *respectively. We say that* $(\mathcal{P}^H, \mathcal{V}^H)$ *is* strongly simulation extractable *with respect to* $\mathcal{S}$ *in the random oracle model, if there exists an efficient algorithm* $\mathcal{E}$ *such that for all PPT adversaries* $\mathcal{A}$ *the following holds. Let:*

$$\mathrm{Prob}\big[w^\star \leftarrow \mathcal{E}(st, x^\star, \pi^\star) \ : \ (x^\star, \pi^\star) \leftarrow \mathcal{A}^{\mathcal{S}_1(\cdot), \mathcal{S}_2'(\cdot)}(1^k; \rho); \\ (x^\star, \pi^\star) \notin \mathcal{T}; \quad \mathcal{V}^{\mathcal{S}_1}(x^\star, \pi^\star) = 1; \quad (x^\star, w^\star) \notin \mathcal{R}_\mathcal{L}\big] \approx 0$$

where $\mathcal{T}$ *is the list of transcripts* (x_i, π_i) *returned by the simulator and the probability space is over the random choices of* $\mathcal{S}$ *and the adversary's randomness* ρ*.*

4 On the Non-malleability of the Fiat-Shamir Transform

4.1 Simulation Soundness

We now show that NIZK proofs obtained via the Fiat-Shamir transform from any IPS of the public-coin type additionally satisfying the HVZK property are simulation sound. Since Σ-protocols are a special class of HVZK public-coin IPSs, we get as a corollary that Fiat-Shamir NIZK proofs obtained from Σ-protocols are simulation-sound.

Theorem 2 (Simulation soundness of the Fiat-Shamir transform). *Consider a* non-trivial *three-round public-coin HVZK interactive proof system* $(\mathcal{P}, \mathcal{V})$ *for a language* $\mathcal{L} \in$ NP*, with quasi unique responses. In the random oracle model, the proof system* $(\mathcal{P}^H, \mathcal{V}^H)$ *derived from* $(\mathcal{P}, \mathcal{V})$ *via the Fiat-Shamir transform is a* simulation-sound *NIZK with respect to its canonical simulator* $\mathcal{S}$*.*

Proof. We assume that $(\mathcal{P}^H, \mathcal{V}^H)$ is a non-interactive zero-knowledge proof system with the simulator $\mathcal{S}$ described in the proof of Theorem 1, and show that $(\mathcal{P}^H, \mathcal{V}^H)$ is simulation sound. We proceed by contradiction. Suppose there exists a PPT adversary $\mathcal{A}$ that breaks the simulation soundness of the non-interactive protocol with non-negligible probability

$$\epsilon := \mathrm{Prob}\big[(x^\star, \pi^\star) \leftarrow \mathcal{A}^{\mathcal{S}_1(\cdot), \mathcal{S}_2'(\cdot)} : \ (x^\star, \pi^\star) \notin \mathcal{T} \ \wedge \ x^\star \notin \mathcal{L} \ \wedge \ \mathcal{V}^{\mathcal{S}_1}(x^\star, \pi^\star) = 1\big].$$

In such a case, we are able to build two reductions $\hat{\mathcal{P}}$ and $\mathcal{P}^*$ which, by using $\mathcal{A}$ as a black-box, violate either the quasi unique response or the soundness properties of the underlying interactive protocol $(\mathcal{P}, \mathcal{V})$ respectively, contradicting our hypothesis. Recall that $\mathcal{S}_1$ simulates answers to the RO, while $\mathcal{S}_2'$ replies with an accepting proof π. Without loss of generality we assume that whenever adversary $\mathcal{A}$ succeeds and outputs an accepting proof $(\alpha^\star, \gamma^\star)$, she has previously queried the oracle $\mathcal{S}_1$ on input $(x^\star, \alpha^\star)$. The argument for this is that it is straightforward to transform any adversary that violates this condition into an adversary that makes one additional query to $\mathcal{S}_1$ and wins with the same probability.

A simple but crucial observation is that adversary $\mathcal{A}$ may have learned $\alpha^\star$ by querying the oracle $\mathcal{S}_2'$ on input $x^\star$ or might have computed it itself. We denote the first by the event proof, the second by the event $\overline{\mathsf{proof}}$. As these events are mutually exclusive and exhaustive, we have:

$$\mathrm{Prob}[\mathcal{A} \text{ wins}] \ = \mathrm{Prob}[\mathcal{A} \text{ wins } \wedge \ \mathsf{proof}] \ + \ \mathrm{Prob}[\mathcal{A} \text{ wins } \wedge \ \overline{\mathsf{proof}}].$$

Now we have two different cases to analyze, each of them corresponding to the probability in the expression above.

In the first case (when proof happens), we assume that $x^\star$ is asked to $\mathcal{S}_2'$ and the answer is a proof of the type $(\alpha^\star, -)$. We show how to use an adversary $\mathcal{A}$ that makes use of $(x^\star, \alpha^\star)$ in its fake proof to build a reduction $\hat{\mathcal{P}}$. In this way we bound $\mathrm{Prob}[\mathcal{A} \text{ wins } \wedge \ \mathsf{proof}]$ by the probability that $\hat{\mathcal{P}}$ wins in breaking the quasi unique response property.

Consider an algorithm $\hat{\mathcal{P}}$ which runs $\mathcal{A}$ internally as a black-box. Thus, $\hat{\mathcal{P}}$ sees all queries $\mathcal{A}$ makes to the oracles $\mathcal{S}_1$ and $\mathcal{S}_2'$ and produces their answers. The internal description of $\hat{\mathcal{P}}$ follows:

- $\hat{\mathcal{P}}$ answers $\mathcal{S}_1$ and $\mathcal{S}_2'$ and keeps lists $\mathcal{T}_H$ and $\mathcal{T}$ respectively as the real simulator $\mathcal{S}$ would.
- When $\mathcal{A}$ outputs a fake proof $(\alpha^\star, \gamma^\star)$ for $x^\star$, $\hat{\mathcal{P}}$ looks through its lists $\mathcal{T}$ and $\mathcal{T}_H$ until it finds $(x^\star, (\alpha^\star, \gamma))$ and $((x^\star, \alpha^\star), \beta)$ respectively;
- It returns $(x^\star, \alpha^\star, \beta, \gamma^\star, \gamma)$.

We claim that algorithm $\hat{\mathcal{P}}$ breaks the quasi unique response property. Indeed, the proof produced by $\mathcal{A}$ is accepting by $\mathcal{V}^H$ on common input $x^\star$. On the other hand, the proof $(\alpha^\star, \gamma)$ is given by the simulator, therefore it must be accepting for $x^\star$. Given this, it holds $\mathcal{V}^H(x^\star, \alpha^\star, \gamma^\star) \ = \ \mathcal{V}^H(x^\star, \alpha^\star, \gamma) = 1$, that means

$$\mathcal{V}(x^\star, \alpha^\star, H(x^\star, \alpha^\star), \gamma^\star) \ = \ \mathcal{V}(x^\star, \alpha^\star, H(x^\star, \alpha^\star), \gamma) = 1,$$

where $H(x^\star, \alpha^\star) = \beta$. The conclusion is that either $\gamma = \gamma^\star$, that is excluded since $\mathcal{A}$ cannot win by printing a simulated proof, or algorithm $\hat{\mathcal{P}}$ succeeds in breaking the quasi unique response property. We obtain:

$$\text{Prob}[\mathcal{A} \text{ wins } \wedge \ \mathsf{proof}] \ = \ \text{Prob}[\hat{\mathcal{P}} \text{ wins}] \leq negl(k).$$

In case proof does not happen, we can use adversary $\mathcal{A}$ that does not query $\mathcal{S}_2'$ with input $x^\star$ to build a reduction $\mathcal{P}^*$ and bound $\text{Prob}[\mathcal{A} \text{ wins } \wedge \ \overline{\mathsf{proof}}]$ by the probability $\text{Prob}[\mathcal{P}^* \text{ wins}] \cdot Q$ of breaking the soundness of the underlying interactive scheme. $\mathcal{P}^*$ runs $\mathcal{A}$ as a black-box and has to simulate its environment by answering the queries to $\mathcal{S}_1$ and $\mathcal{S}_2'$ in a consistent way. More precisely, $\mathcal{P}^*$ works as follows. It guesses uniformly at random an index $j \in [Q]$ and replies to queries to $\mathcal{S}_1$ and $\mathcal{S}_2'$ in the following way:

1. Answer query (x_i, α_i) to $\mathcal{S}_1$:
   (a) Query $1 \leq i \leq j-1$: Returns $H(x_i, \alpha_i)$ if it is already defined; otherwise it samples a random value β_i and sets $H(x_i, \alpha_i) := \beta_i$.
   (b) Query j: Runs the protocol with the honest verifier $\mathcal{V}$ for statement x_j, using as a commitment the value α_j. Obtains challenge β_j from $\mathcal{V}$ and program the oracle as $H(x_j, \alpha_j) := \beta_j$. The answer to $\mathcal{A}$'s query is β_j.
   (c) Query $j+1 \leq i \leq Q$: Proceed as in Step 1a.
2. Answer query x to $\mathcal{S}_2'$: Run the HVZK simulator of the interactive protocol on input x to obtain an accepting proof (α, β, γ), and program the oracle H in such a way that $H(x, \alpha) := \beta$. If the NIZK simulator returns failure, which happens when $H(x, \alpha)$ is already defined, output failure and abort, otherwise output (α, γ).
3. Answer $\mathcal{V}$'s challenge: Let $x^\star, (\alpha^\star, \gamma^\star)$ be the instance and the proof output by $\mathcal{A}$. Return $\gamma^\star$ to $\mathcal{V}$ as the response to challenge β_j in step 1b.

We need to estimate the probability that $\mathcal{P}^*$ succeeds in breaking the soundness of the interactive scheme $(\mathcal{P}, \mathcal{V})$ in terms of the probability that $\mathcal{A}$ outputs an accepting proof $(\alpha^\star, \gamma^\star)$ for a false statement $x^\star$. Suppose that $(x^\star, \alpha^\star)$ has been asked to the random oracle as the $j^\star$-th query and we have $j = j^\star$, i.e., $\mathcal{P}^*$ guesses the correct index for which $\mathcal{A}$ outputs an accepting proof for a false statement $x^\star$. In such a case, $\mathcal{P}^*$ breaks the soundness of $(\mathcal{P}, \mathcal{V})$. Hence, we get:

$$\begin{aligned} \text{Prob}[\mathcal{P}^* \text{ wins}] \ &= \ \text{Prob}[\mathcal{A} \text{ wins } \wedge \ j = j^\star \ \wedge \ \overline{\mathsf{proof}}] \\ &= \ \text{Prob}[\mathcal{A} \text{ wins } \wedge \ \overline{\mathsf{proof}}] \cdot \text{Prob}[j = j^\star], \end{aligned}$$

where the second equality comes from the fact that $\mathcal{P}^*$ guesses $j^\star$ correctly indipendently of the event that $\mathcal{A}$ is successful and $\overline{\mathsf{proof}}$ happens. Since the index j is chosen at random in $[Q]$, we have $\text{Prob}[\mathcal{P}^* \text{ wins}] \ = \ \frac{1}{Q} \cdot \text{Prob}[\mathcal{A} \text{ wins } \wedge \ \overline{\mathsf{proof}}]$. Whenever $\mathcal{P}^*$ wins, it breaks the soundness of the interactive scheme: by hypothesis, this happens only with negligible probability. Therefore:

$$\text{Prob}[\mathcal{A} \text{ wins } \wedge \ \overline{\mathsf{proof}}] \ = \ Q \cdot \text{Prob}[\mathcal{P}^* \text{ wins}] \leq negl(k).$$

Now we can bound the probability that $\mathcal{A}$ succeeds. As we assume, $\mathcal{A}$ breaks the simulation soundness of the scheme with non-negligible probability ϵ:

$$\text{Prob}[\mathcal{A} \text{ wins}] \leq \text{Prob}[\mathcal{A} \text{ wins} \wedge \mathsf{proof}] + \text{Prob}[\mathcal{A} \text{ wins} \wedge \overline{\mathsf{proof}}] \leq negl(k),$$

thus $\epsilon \leq negl(k)$, that is a contradiction. □

On the quasi-unique responses condition. We remark that assuming $(\mathcal{P}, \mathcal{V})$ has quasi-unique responses is not an artifact of the proof. In fact, without this property, proofs would be malleable and breaking the simulation soundness would be an easy task. Consider a FS-NIZK proof system for which responses are not quasi-unique. An efficient adversary $\mathcal{A}$ can always query the simulator on input a false statement $x^\star$, obtaining a simulated proof $\mathcal{S}'_2(x^\star) \to \pi^\star = (\alpha^\star, \beta^\star, \gamma^\star)$. Given $\pi^\star$, $\mathcal{A}$ might be able to find, with non-negligible probability, a new response $\gamma^{\star\star} \neq \gamma^\star$ such that $(\alpha^\star, \beta^\star, \gamma^{\star\star})$ is also accepting. Hence, the scheme cannot be simulation sound.

4.2 Weak Simulation Extractability

The argument Fouque and Pointcheval use in [23] to show that the proof system they consider is simulation sound is roughly as follows. Assume there exists an adversary $\mathcal{A}$ which outputs a pair $(x^\star, \pi^\star)$ breaking the simulation soundness, as in the experiment of Definition 5. Then, one can invoke a suitable version of the forking lemma to show that it is possible to "extract" a witness $w^\star$ for $x^\star$ from such an adversary, contradicting the fact that $x^\star$ is false. The reduction simulates the list $\mathcal{T}$ for $\mathcal{A}$ in the simulation soundness experiment, in particular one needs to fake accepting proofs for (adaptively chosen and potentially false) theorems selected by the attacker. In order to do so, Fouque and Pointcheval (implicitly) rely on the SS-HVZK property. The next theorem is a generalization of the above strategy which does not rely on the SS-HVZK property and indeed applies to arbitrary languages. Moreover, we are able to prove a stronger statement, namely that Fiat-Shamir proofs satisfy weak simulation extractability (and not only simulation soundness). For simplicity the following theorem assumes (perfect) unique responses, but could be generalized using the same reduction as for Theorem 2.

Theorem 3 (Weak simulation extractability of the Fiat-Shamir transform). *Let $\Sigma = (\mathcal{P}, \mathcal{V})$ be a* non-trivial *Σ-protocol with unique responses for a language $\mathcal{L} \in$* NP*. In the random oracle model, the* NIZK *proof system $\Sigma_{FS} = (\mathcal{P}^H, \mathcal{V}^H)$ resulting by applying the Fiat-Shamir transform to Σ is weakly simulation extractable with* extraction error $\nu = \frac{Q}{h}$ *for the canonical simulator $\mathcal{S}$. Here, Q is the number of random oracle queries and h is the number of elements in the range of H. Furthermore, the extractor $\mathcal{E}_\mathcal{A}$ needs to run $\mathcal{A}^{\mathcal{S}_1(\cdot), \mathcal{S}'_2(\cdot)}$ twice, where $\mathcal{A}$ and $\mathcal{E}_\mathcal{A}$ are both defined in Definition 6.*

Proof. Let $\mathcal{S}$ be the canonical zero-knowledge simulator described in the proof of Theorem 2. Denote with $(x^\star, \alpha^\star, \gamma^\star)$ the pair statement/proof returned by

$\mathcal{A}^{\mathcal{S}_1(\cdot),\mathcal{S}'_2(\cdot)}$; we describe an extractor $\mathcal{E}_{\mathcal{A}}$, able to compute a witness $w^\star$ by rewinding $\mathcal{A}$ once.

We want to exploit the general forking lemma1. In order to do so, we define program $\mathsf{P}(1^k, h_1, \ldots, h_Q; \rho_{\mathsf{P}})$ as follows: P virtually splits ρ_{P} into two random tapes ρ and $\rho_{\mathcal{S}}$ (e.g. by using even bits for ρ and odd bits for $\rho_{\mathcal{S}}$) and runs internally $\mathcal{A}^{\mathcal{S}_1(\cdot),\mathcal{S}'_2(\cdot)}$ with randomness ρ. P uses values $(h_1, \ldots, h_Q)$ to simulate fresh answers of $\mathcal{S}_1$, and $\rho_{\mathcal{S}}$ to simulate answers of $\mathcal{S}_2$. If $\mathcal{A}^{\mathcal{S}_1(\cdot),\mathcal{S}'_2(\cdot)}$ outputs $(x^\star, (\alpha^\star, \gamma^\star))$, P checks that it is a valid proof and not in $\mathcal{T}$ (otherwise it returns $(0, \perp)$). Then, because of the unique response property, $(x^\star, \alpha^\star)$ must correspond to some fresh query to $\mathcal{S}_1$ and P outputs $(J, (x^\star, \alpha^\star, \gamma^\star))$, where $J > 0$ is the index corresponding to the random oracle query $(x^\star, \alpha^\star)$. We say that P is successful whenever $J \geq 1$, and we denote with acc the corresponding probability. Given program P, we consider two related runs of P with the same random tape but different hash values, as specified by the forking algorithm F_{P} of Lemma 1. Denote with $(I, (x^\star, \alpha^\star, \gamma^\star)) \leftarrow \mathsf{P}(1^k, h_1, \ldots, h_Q; \rho)$ and $(I', (x^{\star\star}, \alpha^{\star\star}, \gamma^{\star\star})) \leftarrow \mathsf{P}(1^k, h_1, \ldots, h_{I-1}, h'_I, \ldots, h'_Q; \rho)$ the two outputs of $\mathcal{A}$ in these runs. By the forking lemma we know that with probability ext $\geq$ acc$(\text{acc}/Q - 1/h)$ the forking algorithm will return indexes I, I' such that $I = I'$, $I \geq 1$ and $h_I \neq h'_I$.

Notice that since F_{P}'s forgeries are relative to the *same* random oracle query $I = I'$, we must have $x^\star = x^{\star\star}$ and $\alpha^\star = \alpha^{\star\star}$; on the other hand we have $h_I \neq h'_I$. We are thus in a position to invoke the special extractor $\mathcal{E}$ for the underlying proof system, yielding a valid witness $w^\star \leftarrow \mathcal{E}(\alpha^\star, h_I, \gamma^\star, h'_I, \gamma^{\star\star}, x^\star)$ such that $(x^\star, w^\star) \in \mathcal{R}_{\mathcal{L}}$.

Assume now that acc $\geq \nu$. By applying the general forking lemma1 we obtain that ext $\geq \text{acc}^2/Q - \text{acc}/h$. Since Q is polynomial while h is exponentially large in the security parameter, for sure $\frac{Q}{h} < 1$ (in particular, it is negligible in k). As $\nu := \frac{Q}{h}$, we have:

$$\frac{\text{acc}^2}{Q} - \frac{\text{acc}}{h} = \frac{1}{Q}(\text{acc}^2 - \text{acc} \cdot \nu).$$

Now, since acc $\geq \nu$, we have acc $\cdot \nu \geq \nu^2$, that is $\nu^2 - \text{acc} \cdot \nu \leq 0$. Hence,

$$\frac{1}{Q}(\text{acc}^2 - \text{acc} \cdot \nu) \geq \frac{1}{Q}(\text{acc}^2 - 2\text{acc} \cdot \nu + \nu^2) = \frac{1}{Q}(\text{acc} - \nu)^2.$$

The previous inequality matches the definition of weak extractability with values $p = Q$ and $d = 2$. □

5 Applications

In the literature there is a large number of applications for simulation-sound or extractable NIZKs. One of the first request for simulation soundness comes from the setting of public key encryption, for the design of encryption schemes with *chosen-ciphertext security* using the Naor-Yung (NY) paradigm [34]. At a high level, the NY works as follows: given two key pairs (sk, pk) and (sk', pk')

for CPA-secure encryption schemes Π and Π' respectively, a ciphertext consists of two encryptions c, c' of the same message m, under different keys pk, pk', and a NIZK proof π that both c and c' encrypt m. In order to achieve security against *adaptive* chosen-ciphertext attacks (CCA2), the underlying NIZK must be simulation-sound. While achieving CCA2 security is probably one of the most prominent application of simulation soundness, simulation-sound or extractable proofs have been used, e.g., to build also leakage-resilient signatures or KDM secure encryption. In this section, we review some important applications of such proof systems and show how our result provides more efficient constructions in the ROM or generalizes earlier results that use the Fiat-Shamir transform.

5.1 Leakage Resilience

Simulation-sound and simulation-extractable NIZK proofs have been very useful in constructing leakage-resilient encryption and signature schemes [18,32,33]. Here, we consider these works and show that our result immediately yields efficient leakage-resilient schemes in the random oracle model.

Leakage-resilient signatures. A signature scheme is leakage resilient if it is hard to forge a signature even given (bounded) leakage from the signing key. Obviously, this requires that the amount of leakage given to the adversary has to be smaller than the length of the secret key, as otherwise the leakage may just reveal such a key, trivially breaking the security of the signature scheme.

We instantiate the generic construction of leakage-resilient signatures based on leakage-resilient hard relations and simulation-extractable NIZKs of [18] using the Fiat-Shamir transform. Let $\mathcal{R}$ be a λ-leakage-resilient hard relation with sampling algorithm $\mathsf{Gen}_{\mathcal{R}}$ (for detailed definitions, see the full version [20]). Let $(\mathcal{P}^H, \mathcal{V}^H)$ be a NIZK argument[3] for relation $\mathcal{R}'$ defined by $\mathcal{R}'((pk, m), sk) \Leftrightarrow \mathcal{R}(pk, sk)$. Consider the following signature scheme:

$\mathsf{KeyGen}(1^k)$: Calls $(pk, sk) \leftarrow \mathsf{Gen}_{\mathcal{R}}(1^k)$ and returns the same output.
$\mathsf{Sign}(sk, m)$: Outputs $\sigma \leftarrow \mathcal{P}^H((pk, m), sk)$.[4]
$\mathsf{Vrf}(pk, m, \sigma)$: Verifies the signature by invoking $\mathcal{V}^H((pk, m), \sigma)$.

Notice that $\sigma \leftarrow \mathcal{P}^H((pk, m), sk)$ is a NIZK proof for the hard relation obtained by applying the Fiat-Shamir transform.

We chose to state the theorem below using an argument system as this is the minimal requirement under which leakage resilience of the scheme can be proven. Since our FS-based protocols are weakly simulation-extractable NIZK *proof* systems, they automatically satisfy the hypothesis of Theorem 4.

Theorem 4. *If $\mathcal{R}$ is a 2λ-leakage-resilient hard relation and $(\mathcal{P}^H, \mathcal{V}^H)$ is a weakly simulation-extractable NIZK argument with negligible extraction error*

[3] As opposed to a proof system where soundness needs to hold unconditionally, in an *argument* system it is sufficient that soundness holds with respect to a computationally bounded adversary.

[4] Note that m is part of the instance being proven.

for relation $\mathcal{R}'((pk, m), sk) \Leftrightarrow \mathcal{R}(pk, sk)$*, then the above scheme is a* λ*-leakage-resilient signature scheme in the random oracle model.*

The proof of the theorem from above follows the one of Theorem 4.3 in [18]. A couple of subtleties arise, though. The main idea of the proof is to build a reduction from an adversary $\mathcal{A}$ breaking λ-leakage resilience of the signature scheme to an adversary $\mathcal{B}$ breaking the hardness of the 2λ-leakage-resilient hard relation $\mathcal{R}$. Roughly speaking, in the reduction $\mathcal{B}$ is given some instance pk and simulates the signing queries of $\mathcal{A}$ by using the zero-knowledge simulator of the NIZK, and the leakage queries by using the leakage oracle for the relation $\mathcal{R}$. At some point $\mathcal{A}$ outputs a forgery $\sigma^\star$ and $\mathcal{B}$ invokes the extractor of Theorem 3 to get $sk^\star \leftarrow \mathcal{E}_{\mathcal{A}}(pk, \sigma^\star)$. The first issue is that we are only guaranteed *weak* simulation-extractability, whereas the proof of [18] relies on full simulation-extractability. [5] However, this is not a problem because we just need to show that $\mathcal{B}$ outputs a valid witness with non-negligible probability. A second issue involves the extractor of Theorem 3, which needs to rewind $\mathcal{A}$ once and, thus, to simulate twice its environment (including the leakage queries). This causes the loss of a factor 2 in the total amount of tolerated leakage. We refer the reader to the full version [20] for the details.

We emphasize that the leakage-resilient signature scheme of Alwen et al. [4], obtained by applying the Fiat-Shamir transform to the Okamoto identification scheme, follows essentially the above paradigm. Here, one may view the public and secret keys of the Okamoto ID scheme form an instance of a leakage-resilient hard relation, while the NIZK proof corresponds to the Fiat-Shamir transform applied to the Okamoto identification protocol.

Naor-Yung with leakage. The definition of IND-CPA and IND-CCA security of an encryption scheme can be extended to the leakage setting by giving the adversary access to a leakage oracle. Naor and Segev [33] show that the Naor-Yung paradigm instantiated with a simulation-sound NIZK allows to leverage CPA-security to CCA-security even in the presence of leakage. In other words, if Π is CPA-secure against λ-key-leakage attacks, the encryption scheme obtained by applying the Naor-Yung paradigm to (Π, Π), using a simulation-sound NIZK, is CCA2-secure against λ-key-leakage attacks. In the full version [20] we revisit their proof in the ROM, dealing with the issue that the leakage queries can potentially depend on H. We stress that for the proof only *simulation soundness* is needed (i.e., our result from Theorem 2) and not weak simulation extractability.

In what follows, we propose a concrete instantiation of the result above, relying on the BHHO encryption scheme from [12]. Let $\mathbb{G}$ be a group of prime-order q. For randomly selected generators $g_1, \ldots, g_\ell \xleftarrow{\$} \mathbb{G}$, the public key is a tuple $pk = (g_1, \ldots, g_\ell, h)$, where $h = \prod_{i=1}^{\ell} g_i^{z_i}$ for a secret key $sk = (z_1, \ldots, z_\ell) \in \mathbb{Z}_q^\ell$. To encrypt a message $m \in \mathbb{G}$, choose a random $r \xleftarrow{\$} \mathbb{Z}_q$ and output $c = (c_1, \ldots, c_{\ell+1}) = (g_1^r, \ldots, g_\ell^r, m \cdot h^r)$. The message m can be recovered by computing $m = c_{\ell+1} \cdot (\prod_{i=1}^{\ell} c_i^{z_i})^{-1}$.

[5] Actually, they rely on a weaker property called *true* simulation-extractability [18].

Assuming that the DDH problem is hard in $\mathbb{G}$, Naor and Segev [33] showed that the BHHO encryption scheme is CPA-secure against λ-key-leakage attacks for any $\ell = 2 + \frac{\lambda+\omega(\log k)}{\log q}$, where k is the security parameter. Applying the Naor-Yung paradigm, consider the language:

$$\mathcal{L} = \big\{(c, pk, c', pk') : \exists r, r' \in \mathbb{Z}_q, m \in \mathbb{G} \text{ s.t.} \\ c = (g_1^r, \ldots, g_\ell^r, h^r \cdot m),\ c' = (g_1^{r'}, \ldots, g_\ell^{r'}, h'^{r'} \cdot m)\big\},$$

where $c = (c_1, \ldots, c_{\ell+1})$ and $c' = (c'_1, \ldots, c'_{\ell+1})$ are BHHO encryptions with randomness r and r', using public keys $pk = (g_1, \ldots, g_\ell, h)$ and $pk' = (g_1, \ldots, g_\ell, h')$ respectively. The pair $w = (r, r')$ is a witness for a string $x = (c, pk, c', pk') \in \mathcal{L}$. Consider the following interactive protocol $\Sigma = (\mathcal{P}, \mathcal{V})$ for the above language:

1. $\mathcal{P}$ chooses s, s' at random from $\mathbb{Z}_q$ and computes the commitment:
$$\vec{\alpha} = ((\alpha_1, ..., \alpha_\ell), (\alpha'_1, ..., \alpha'_\ell), \alpha'') = ((g_1^s, ..., g_\ell^s), (g_1^{s'}, ..., g_\ell^{s'}), h^s \cdot (h')^{s'}).$$
2. The verifier $\mathcal{V}$ chooses a random challenge $\beta \xleftarrow{\$} \mathbb{Z}_q$.
3. The prover computes the response $\vec{\gamma} = (\gamma, \gamma') = (s - \beta \cdot r, s' + \beta \cdot r')$.
4. Given a proof $\pi = (\vec{\alpha}, \beta, \vec{\gamma})$, the verifier $\mathcal{V}$ checks that:
$$(\alpha_1, \ldots, \alpha_\ell) = (g_1^\gamma \cdot c_1^\beta, \ldots, g_\ell^\gamma \cdot c_\ell^\beta) \\ (\alpha'_1, \ldots, \alpha'_\ell) = (g_1^{\gamma'} \cdot (c'_1)^{-\beta}, \ldots, g_\ell^{\gamma'} \cdot (c'_\ell)^{-\beta}) \\ \alpha'' = h^\gamma \cdot (h')^{\gamma'} \cdot (c_{\ell+1} \cdot (c'_{\ell+1})^{-1})^\beta.$$

In the full version [20] we prove that the above protocol is a Σ-protocol for the language $\mathcal{L}$. With the Naor-Yung paradigm applied to the BHHO encryption scheme we get a ciphertext (c, c', π) consisting of $4\ell + 3$ elements in $\mathbb{G}$ plus 2 elements in $\mathbb{Z}_q$. Moreover, the fact that the BHHO encryption scheme is CPA-secure against key leakage together with the result of Naor-Segev, show that the above instantiation is CCA-secure against key-leakage attacks.

Corollary 1. *Let k be a security parameter. Assuming that the DDH problem is hard in $\mathbb{G}$, the Naor-Yung paradigm applied to the BHHO encryption scheme yields an encryption scheme that is CCA-secure against λ-key-leakage attacks in the random oracle model for $\lambda = \ell \log q(1 - \frac{2}{\ell} - \frac{\omega(\log k)}{\ell \log q}) = L(1 - o(1))$, where L is the length of the secret key. An encryption consists of $4\ell + 3$ elements in $\mathbb{G}$ plus 2 elements in $\mathbb{Z}_q$.*

5.2 Key-Dependent Message Security

Key-dependent message (KDM) security of a public-key encryption scheme requires that the scheme remains secure even against attackers allowed to see encryptions of the value $f(sk)$, where $f \in \mathcal{F}$ for some class of functions $\mathcal{F}$.

Camenisch, Chandran and Shoup [13] show that a variation of the Naor-Yung paradigm instantiated with a simulation-sound NIZK can still leverage CPA-security to CCA-security, even in the context of KDM security. We revisit their

proof in the random oracle model in the full version [20]. Also in this case, only *simulation soundness* is needed for the proof.

Roughly, for some function family $\mathcal{F}$, if Π is KDM[$\mathcal{F}$]-CPA secure and Π' is CPA-secure, the scheme Π'' obtained by applying the Naor-Yung paradigm to (Π, Π') — i.e., an encryption of $m \in \mathcal{M}$ is a tuple $c'' = (c, c', \pi)$ where c encrypts m under Π, c' encrypts m under Π' and π is a simulation-sound NIZK proof that c and c' encrypt the same message — is KDM[$\mathcal{F}$]-CCA secure.

Let $sk_i[j]$ denote the j-th bit of sk_i. The BHHO encryption scheme was the first KDM-CPA secure encryption scheme, with respect to the class of all projection functions $\mathcal{F}_{\downarrow} = \mathcal{F}_{\text{read}} \cup \mathcal{F}_{\text{flip}}$, where

$$\mathcal{F}_{\text{read}} = \left\{ f_{i,j} : \vec{sk} \to sk_i[j] \right\}_{i,j} \text{ and } \mathcal{F}_{\text{flip}} = \left\{ f_{i,j} : \vec{sk} \to 1 - sk_i[j] \right\}_{i,j}.$$

More generally, when the message space is a linear space over $\mathbb{Z}_q$, we define the function class $\mathcal{PJ}(\mathcal{F}_{\downarrow})$ as the class of all affine combinations of elements in $\mathcal{F}_{\downarrow}$.

Now we can instantiate the general transform of [13] as follows. We choose Π to be BHHO, Π' to be ElGamal (say with $pk' = h' = g_1^{z'_1}$) and we build a Σ-protocol Σ' for the Naor-Yung language relative to Π and Π'. Protocol Σ' can be easily derived from protocol Σ of the last section, by just compressing the commitment as in $\vec{\alpha} = ((g_1^s, \ldots, g_\ell^s), g_1^{s'}, h^s \cdot (h')^{s'})$ (and simplifying the verification procedure accordingly). Hence, Theorem 2 yields the following result.

Corollary 2. *Assuming that the DDH problem is hard in* $\mathbb{G}$*, the Naor-Yung paradigm instantiated with BHHO and ElGamal encryption schemes yields a* KDM[$\mathcal{PJ}(\mathcal{F}_{\downarrow})$]-CCA *secure encryption scheme in the random oracle model. An encryption consists of* $\ell + 3$ *elements in* $\mathbb{G}$ *plus* 3 *elements in* $\mathbb{Z}_q$*.*

Beyond Naor-Yung. Another paradigm that yields chosen-ciphertext security from NIZKs, based on *proving knowledge of the plaintext*, was suggested by Rackoff and Simon [36]. Such a construction is somewhat more natural and more efficient than the twin-encryption paradigm: a message m is encrypted (only once) under a CPA-secure encryption scheme, and a NIZK proof of knowledge of the plaintext is attached to the ciphertext. However, (to the best of our knowledge) truly efficient constructions for sufficiently strong NIZK proofs of knowledge are not available even using random oracles. One can hope that using the weaker from of extractability afforded by the Fiat-Shamir transform one could at least obtain NM-CPA secure encyption, and this is indeed what is aimed at in the ongoing work of [8].

Acknowledgments. We thank Marc Fischlin and Ivan Damgård for the useful feedbacks provided on earlier versions of the paper.

References

1. Abdalla, M., An, J.H., Bellare, M., Namprempre, C.: From Identification to Signatures via the Fiat-Shamir Transform: Minimizing Assumptions for Security and Forward-Security. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 418–433. Springer, Heidelberg (2002)

2. Abdalla, M., An, J.H., Bellare, M., Namprempre, C.: From identification to signatures via the Fiat-Shamir transform: Necessary and sufficient conditions for security and forward-security. IEEE Transactions on Information Theory 54(8), 3631–3646 (2008)
3. Abe, M., Fuchsbauer, G., Groth, J., Haralambiev, K., Ohkubo, M.: Structure-Preserving Signatures and Commitments to Group Elements. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 209–236. Springer, Heidelberg (2010)
4. Alwen, J., Dodis, Y., Wichs, D.: Leakage-Resilient Public-Key Cryptography in the Bounded-Retrieval Model. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 36–54. Springer, Heidelberg (2009)
5. Bellare, M., Goldreich, O.: On Defining Proofs of Knowledge. In: Brickell, E.F. (ed.) CRYPTO 1992. LNCS, vol. 740, pp. 390–420. Springer, Heidelberg (1993)
6. Bellare, M., Neven, G.: Multi-signatures in the plain public-key model and a general forking lemma. In: ACM Conference on Computer and Communications Security, pp. 390–399 (2006)
7. Bellare, M., Rogaway, P.: Random oracles are practical: A paradigm for designing efficient protocols. In: ACM Conference on Computer and Communications Security, pp. 62–73 (1993)
8. Bernhard, D., Pereira, O., Warinschi, B.: On necessary and sufficient conditions for private ballot submission. Cryptology ePrint Archive, Report 2012/236 (2012), http://eprint.iacr.org/
9. Blum, M., Feldman, P., Micali, S.: Non-interactive zero-knowledge and its applications (extended abstract). In: STOC, pp. 103–112 (1988)
10. Blum, M., De Santis, A., Micali, S., Persiano, G.: Noninteractive zero-knowledge. SIAM J. Comput. 20(6), 1084–1118 (1991)
11. Boneh, D., Boyen, X., Shacham, H.: Short Group Signatures. In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 41–55. Springer, Heidelberg (2004)
12. Boneh, D., Halevi, S., Hamburg, M., Ostrovsky, R.: Circular-Secure Encryption from Decision Diffie-Hellman. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 108–125. Springer, Heidelberg (2008)
13. Camenisch, J., Chandran, N., Shoup, V.: A Public Key Encryption Scheme Secure against Key Dependent Chosen Plaintext and Adaptive Chosen Ciphertext Attacks. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 351–368. Springer, Heidelberg (2009)
14. Camenisch, J., Lysyanskaya, A.: An Efficient System for Non-transferable Anonymous Credentials with Optional Anonymity Revocation. In: Pfitzmann, B. (ed.) EUROCRYPT 2001. LNCS, vol. 2045, pp. 93–118. Springer, Heidelberg (2001)
15. Chase, M., Lysyanskaya, A.: On Signatures of Knowledge. In: Dwork, C. (ed.) CRYPTO 2006. LNCS, vol. 4117, pp. 78–96. Springer, Heidelberg (2006)
16. Cramer, R., Damgård, I., Schoenmakers, B.: Proof of Partial Knowledge and Simplified Design of Witness Hiding Protocols. In: Desmedt, Y.G. (ed.) CRYPTO 1994. LNCS, vol. 839, pp. 174–187. Springer, Heidelberg (1994)
17. Damgård, I.: On Σ-protocols (2002), http://www.daimi.au.dk/~ivan/Sigma.ps
18. Dodis, Y., Haralambiev, K., López-Alt, A., Wichs, D.: Efficient Public-Key Cryptography in the Presence of Key Leakage. In: Abe, M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 613–631. Springer, Heidelberg (2010)
19. Dolev, D., Dwork, C., Naor, M.: Nonmalleable cryptography. SIAM J. Comput. 30(2), 391–437 (2000)
20. Faust, S., Kohlweiss, M., Marson, G.A., Venturi, D.: On the non-malleability of the Fiat-Shamir transform. Cryptology ePrint Archive (2012), http://eprint.iacr.org/

21. 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)
22. Fischlin, M.: Communication-Efficient Non-interactive Proofs of Knowledge with Online Extractors. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 152–168. Springer, Heidelberg (2005)
23. Fouque, P.-A., Pointcheval, D.: Threshold Cryptosystems Secure against Chosen-Ciphertext Attacks. In: Boyd, C. (ed.) ASIACRYPT 2001. LNCS, vol. 2248, pp. 351–368. Springer, Heidelberg (2001)
24. Garay, J.A., MacKenzie, P.D., Yang, K.: Strengthening zero-knowledge protocols using signatures. J. Cryptology 19(2), 169–209 (2006)
25. Goldreich, O., Oren, Y.: Definitions and properties of zero-knowledge proof systems. J. Cryptology 7(1), 1–32 (1994)
26. Goldwasser, S., Micali, S., Rackoff, C.: The knowledge complexity of interactive proof systems. SIAM J. Comput. 18(1), 186–208 (1989)
27. Goldwasser, S., Sipser, M.: Private coins versus public coins in interactive proof systems. In: STOC, pp. 59–68 (1986)
28. Groth, J.: Simulation-Sound NIZK Proofs for a Practical Language and Constant Size Group Signatures. In: Lai, X., Chen, K. (eds.) ASIACRYPT 2006. LNCS, vol. 4284, pp. 444–459. Springer, Heidelberg (2006)
29. Groth, J., Sahai, A.: Efficient Non-interactive Proof Systems for Bilinear Groups. In: Smart, N.P. (ed.) EUROCRYPT 2008. LNCS, vol. 4965, pp. 415–432. Springer, Heidelberg (2008)
30. Halevi, S., Micali, S.: More on proofs of knowledge. Cryptology ePrint Archive, Report 1998/015 (1998), `http://eprint.iacr.org/`
31. Jain, A., Pandey, O.: Non-malleable zero knowledge: Black-box constructions and definitional relationships. Cryptology ePrint Archive, Report 2011/513 (2011), `http://eprint.iacr.org/`
32. Katz, J., Vaikuntanathan, V.: Signature Schemes with Bounded Leakage Resilience. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 703–720. Springer, Heidelberg (2009)
33. Naor, M., Segev, G.: Public-Key Cryptosystems Resilient to Key Leakage. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 18–35. Springer, Heidelberg (2009)
34. Naor, M., Yung, M.: Public-key cryptosystems provably secure against chosen ciphertext attacks. In: STOC, pp. 427–437 (1990)
35. Pointcheval, D., Stern, J.: Security arguments for digital signatures and blind signatures. J. Cryptology 13(3), 361–396 (2000)
36. Rackoff, C., Simon, D.R.: Non-interactive Zero-Knowledge Proof of Knowledge and Chosen Ciphertext Attack. In: Feigenbaum, J. (ed.) CRYPTO 1991. LNCS, vol. 576, pp. 433–444. Springer, Heidelberg (1992)
37. Sahai, A.: Non-malleable non-interactive zero knowledge and adaptive chosen-ciphertext security. In: FOCS, pp. 543–553 (1999)
38. De Santis, A., Di Crescenzo, G., Ostrovsky, R., Persiano, G., Sahai, A.: Robust Non-interactive Zero Knowledge. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 566–598. Springer, Heidelberg (2001)
39. Unruh, D.: Quantum proofs of knowledge. To appear in CRYPTO (2012)
40. Wee, H.: Zero Knowledge in the Random Oracle Model, Revisited. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 417–434. Springer, Heidelberg (2009)