

Public-Key Generation with Verifiable Randomness

Olivier Blazy¹, Patrick Towa^{2,3}, Damien Vergnaud⁴

¹ Universite de Limoges
olivier.blazy@unilim.fr

² IBM Research Zurich
tow@zurich.ibm.com

³ ENS and PSL Research University

⁴ Sorbonne Universite, LIP6 and Institut Universitaire de France
damien.vergnaud@lip6.fr

Abstract. We revisit the problem of proving that a user algorithm selected and correctly used a truly random seed in the generation of her cryptographic key. A first approach was proposed in 2002 by Juels and Guajardo for the validation of RSA secret keys. We present a new security model and general tools to efficiently prove that a private key was generated at random according to a prescribed process, without revealing any further information about the private key.

In addition to formalizing randomness verifiability in key generation, which turns out to be highly non-trivial, we give a generic protocol for all key-generation algorithms based on probabilistic circuits and prove its security. We also propose a new protocol for factoring-based cryptography that we prove secure in the aforementioned model, as well as a practical instantiation. This latter relies on a new efficient zero-knowledge argument for the double discrete logarithm problem that achieves an exponential improvement in communication complexity compared to the state of the art, and is of independent interest.

1 Introduction

Cryptographic protocols are commonly designed under the assumption that the protocol parties have access to perfect (i.e., uniform) randomness. However, random sources used in practical implementations rarely meet this assumption and provide only a stream of bits with a certain “level of randomness”. The quality of the random numbers directly determines the security strength of the systems that use them. Following preliminary work by Juels and Guajardo [37] and Corrigan-Gibbs, Mu, Boneh and Ford [24], we revisit the problem of proving that a cryptographic user algorithm has selected and correctly used a truly random seed in the generation of her cryptographic public-secret key pair.

Related Work. A prominent example that the use of randomness in public-key cryptography (and especially in key-generation protocols) is error-prone is the

recent randomness failure known as the *ROCA vulnerability* [43]. This weakness allows a private key to be recovered efficiently from the public key only (in factoring-based cryptography) for keys generated by some devices. The flawed key-generation algorithm selects specific prime numbers as part of the private key instead of generating uniformly random primes. Many certified devices were shown vulnerable (e.g., Estonian and Slovakian smartcards and standard cryptographic libraries). This kind of weaknesses is not new as in 2012, Lenstra, Hughes, Augier, Bos, Kleinjung and Wachter [38] did a sanity check of factoring-based public keys collected on the web. They showed that a significant percentage of public keys (0.5%) share a common prime factor, and this fact was explained [35] by the generation of these low entropy keys during booting. Since cryptographic failures due to weak randomness can be dire [43,38,35,50], designers should build schemes that can withstand deviations of the random sources from perfect randomness.

Following seminal works by Simmons on the threat of covert channels (also called subliminal channels) in cryptography [47], the concept of *kleptography* was proposed by Young and Yung in [51]. It models the fact that an adversary may subvert cryptographic algorithms by modifying their implementations in order to covertly leak secrets using for instance covert channels present in the randomized key-generation or signing algorithm. Following the revelations of Snowden, several sources have recently revealed that cryptographic algorithms have effectively been subverted to undermine the security of users. This raises the concern of guaranteeing a user's security even when she may be using a compromised machine or algorithm. Motivated by the (in)famous potential backdoor on the Dual Elliptic Curve Deterministic Random Bit Generator (Dual EC DRBG) [22], Bellare, Paterson, and Rogaway [9] recently initiated a formal analysis of kleptographic attacks on symmetric key encryption algorithms. For factoring-based public-key cryptography, in light of the known shortcomings of implemented key generators, a line of research has focused on proving that RSA moduli satisfy certain properties [31,18,4], or on attesting that RSA prime factors were generated with a specified prime generator [12]. This line of work is only concerned with the structure of the keys, not with the fact that they are generated with enough entropy. Juels and Guajardo [37] suggested as early as in 2002 an approach for users to prove to another party (which is typically a trusted certificate authority or CA) that her public-secret key pair was generated honestly using proper randomness. In their setting, the CA provides an additional source of randomness in an interactive process, and the user algorithm proves that it has not weakened, whether intentionally or unintentionally, the key generation procedure. The security goal of such a primitive is threefold.

- 1. Maintain User Privacy:** if the user uses a randomness source with high entropy, then an adversary (possibly the CA itself) has no additional information on the secret-key compared to a key generated by the real key-generation algorithm on uniform randomness.
- 2. Improve Randomness Quality:** if the user *or* the CA use a randomness source with high entropy, then, an adversary (other than the CA) has no

additional information on the secret-key compared to a key generated by the real key-generation algorithm on uniform randomness.

3. Resist Information Exfiltration: the generated public key leaks no information whatsoever to the outer world. In particular, a faulty user algorithm cannot use it to convey any information. In this sense, the CA certifies to the end user, that she can securely use to the generated key.

Juels and Guajardo proposed a formal security model for *verifiable random key generation* with the goal to achieve these three security objectives. Their model is unfortunately not strong enough to capture real-world threats since

- it is restricted to public-key cryptosystems where a given public-key corresponds to a *unique* secret key (and cannot be used for many recent schemes);
- it considers only a stand-alone or independent key-generation instances (and therefore does not prevent attacks such as the one considered in [38,35] where several public-keys are generated with correlated randomness sources);
- it only bounds the distance that a dishonest user can generate a given key to that of an honest user executing the key generation protocol.

As a simple example, consider the problem of generating an ElGamal public key g^x in a group $\mathbb{G} = \langle g \rangle$ of prime order p . Juels and Guajardo outlined a protocol for generating such a key with verifiable randomness. The natural idea to generate a public-key g^x in this (illusorily) simple setting is to share the secret key x as $x = x_U + x_{CA} \bmod p$ where x_U denotes the user randomness and x_{CA} denotes the CA randomness. However, this protocol fails to achieve (3) as the user algorithm can choose x_U to match a specify value after seeing x_{CA} . To overcome this issue, a simple idea would be to make the user first commit to x_U and then prove its knowledge. However, the hiding and zero-knowledge properties of commitment schemes and proof systems inherently rely on perfect randomness, which the user algorithm is assumed not to have at its disposal.

Juels and Guajardo also proposed a protocol for the generation of RSA keys where the distance in (3) increases by a factor which is polynomial in the security parameter λ (assuming some number-theoretic conjecture). Therefore, their protocol does not rule out the existence of covert channels with $O(\log \lambda)$ bits capacity.

Juels-Guajardo model was reconsidered by Corrigan-Gibbs, Mu, Boneh and Ford [24] in a weaker setting that guarantee (1) and (2) but not (3), and does not even prevent a malicious user algorithm from generating malformed keys.

Contributions. The goal of this paper is to revisit this verifiable key-generation primitive by providing the first strong formal security models and efficient, provably secure constructions.

Game-Based Security Model. We first propose a game-based model that covers concurrent protocol executions with different instances of protocol algorithms. It is inspired by the Bellare-Pointcheval-Rogaway model for authenticated key exchange [10]. The communication between the user and the CA is assumed to be carried over an insecure channel. Messages can be tapped and modified by an

adversary, and the communication between the user and the CA is asynchronous. The adversary can provide the randomness sources to the user and the CA (for multiple instances of the protocol) and should not be able to gain any information on a generated public key under the assumption that the entropy of one random sources is high enough.

The main difficulty to define the security model for this primitive is to formalize the third security objective. A dishonest user algorithm can indeed always execute several instances of the protocol with the CA until it obtains a public-key which has some specific property which allows to exfiltrate information. This is similar to the “halting attack” subliminal channel [27] and cannot be avoided.⁵ We manage to take this *narrow-band* subliminal channel into consideration in our security models while capturing the fact that in a secure protocol, this should be the only possible covert channel for a dishonest user algorithm. In practical applications, this covert channel can be prevented easily if the CA charges an important fee for a user that performs too many key generation procedures, or if an increasing time-gating mechanism for repeating queries is introduced.

This game-based security model does not suffer from the shortcomings of the model proposed by Juels and Guajardo as it allows for multiple dependent runs of the protocol and captures the resistance to exfiltration of secret information (with only the narrow-band subliminal channel from the “halting attack”).

Generic Protocol for Probabilistic Circuits. We then present a generic approach for key generation based on (families of) probabilistic circuits and we prove its security in our stringent security model. It relies on two-source randomness extractors, pseudo-random-function families and extractable commitments with associated zero-knowledge proofs.

Efficient Protocol for RSA Keys. We also propose a new generic protocol for factoring-based cryptography and prove it secure in our model. It relies on classical cryptographic tools (namely commitments, pseudo-random functions and zero-knowledge proofs). We provide an instantiation of this generic protocol based on the Dodis–Yampolskiy PRF [28] in the group of quadratic residue modulo a safe prime which outputs group elements. The main technical difficulty is to convert the outputs of this PRF into integers while proving that the RSA prime factors are outputs of the PRF. In the process, we propose a new efficient zero-knowledge proof system for the so-called *double discrete logarithm problem* (in groups of public order). A double discrete logarithm of an element $y \neq 1_{\mathbb{G}}$ in a cyclic group \mathbb{G} of prime order p with respect to bases $g \in \mathbb{G}$ and $h \in \mathbb{Z}_p^*$ (generators of \mathbb{G} and \mathbb{Z}_p^* respectively) is an integer $x \in \{0, \dots, p-1\}$ such that $y = g^{hx}$. Stadler introduced this computational problem for verifiable secret-sharing [48] and it was used subsequently to design numerous cryptographic protocols (e.g. group signatures [19], blind signatures in [3] and e-cash systems [20] and credential systems [21]). All these constructions rely on a proof

⁵ At least without adding a third party to the model as in access control encryption [26,36].

system proposed by Stadler which has $\Omega(\log p)$ computational and communication complexity. Our new proof system achieves $O(\log \log p)$ communication complexity only and permits an efficient instantiation of our generic protocol for factoring-based cryptography. As a by-product, our new protocol can be used directly in all the aforementioned applications in a public-order setting to exponentially reduce their communication complexity.

2 Preliminaries

This section introduces the notation, the hardness assumptions and the building blocks used throughout this paper (see the appendices for more details).

Notation. For $n \in \mathbb{N}$, the set of n -bit strings is denoted by $\{0, 1\}^n$ and the set of integers $\{1, \dots, n\}$ is denoted $[n]$. The set of prime numbers is denoted \mathbb{P} . The security parameter is denoted λ , and input lengths are always assumed to be bounded by some polynomial in λ . A Probabilistic algorithm is said to run in Polynomial-Time (it is said to be a PPT algorithm) if it runs in time that is polynomial in λ . A function μ is negligible if $\mu(\lambda) = \lambda^{-\omega(1)}$.

The random variable defined by the value returned by a PPT \mathcal{A} on input x is denoted $\mathcal{A}(x)$. The value returned by \mathcal{A} on input x and random string r is denoted $\mathcal{A}(x; r)$. Given a probability distribution S , a PPT algorithm that samples a random element according to S is denoted by $x \leftarrow_{\S} S$. For a finite set X , $x \leftarrow_{\S} X$ denotes a PPT algorithm that samples an element uniformly at random from X . Given a group \mathbb{G} with neutral element $1_{\mathbb{G}}$, \mathbb{G}^* denotes $\mathbb{G} \setminus \{1_{\mathbb{G}}\}$. For any two sets \mathcal{X} and \mathcal{Y} , denote by $\mathcal{Y}^{\mathcal{X}}$ the set of functions from \mathcal{X} to \mathcal{Y} .

Vectors are denoted in bold font. For two vectors \mathbf{a} and \mathbf{b} in R^n where R is a ring and n a positive integer, $\mathbf{a} \circ \mathbf{b}$ denotes the Hadamard product of \mathbf{a} and \mathbf{b} , i.e., $\mathbf{a} \circ \mathbf{b} := [a_1 b_1 \cdots a_n b_n]$.

Group Families. A *group-family generator* \mathbb{G} is a PPT which takes as input a security parameter λ and returns a tuple (\mathbb{G}, ℓ, g) , with \mathbb{G} a finite cyclic multiplicative group of prime order ℓ , and $g \in \mathbb{G}$ a generator of \mathbb{G} (i.e. $g \in \mathbb{G}^*$).

Hardness Assumptions. This section recalls classical assumptions on group-family generators.

Definition 2.1 (Discrete-Logarithm Assumption). *Let \mathbb{G} be a group-family generator. The (T, ε) -discrete logarithm assumption on \mathbb{G} states that for any adversary \mathcal{A} that runs in time at most $T(\lambda)$, the probability*

$$\Pr \left[x \leftarrow \mathcal{A}(\mathbb{G}, \ell, g, h) : (\mathbb{G}, \ell, g) \leftarrow \mathbb{G}(\lambda); x \leftarrow_{\S} \mathbb{Z}_{\ell}^*; h \leftarrow g^x \right]$$

is at most $\varepsilon(\lambda)$.

Definition 2.2 (Decisional Diffie-Hellman Assumption). *Let \mathbb{G} be a group family generator. The (T, ε) -Decisional Diffie-Hellman assumption (DDH) over*

\mathbb{G} states that the advantage (function of λ)

$$\left| \Pr \left[b = \mathcal{A}(\mathbb{G}, \ell, g, g^x, g^y, g^{\alpha b}) : \begin{array}{l} (\mathbb{G}, \ell, g) \leftarrow \mathbb{G}(\lambda) \\ (x, y) \leftarrow_{\mathfrak{s}} \mathbb{Z}_{\ell}^{*2}, b \leftarrow_{\mathfrak{s}} \{0, 1\} \\ \alpha_0 \leftarrow xy \bmod \ell, \alpha_1 \leftarrow_{\mathfrak{s}} \mathbb{Z}_{\ell}^* \end{array} \right] - \frac{1}{2} \right|$$

of any adversary \mathcal{A} that runs in time at most $T(\lambda)$ is at most $\varepsilon(\lambda)$.

Definition 2.3 (Decisional Diffie-Hellman-Inversion Assumption [13,17]).

Let \mathbb{G} be a group family generator. The (T, q, ε) -Decisional Diffie-Hellman-Inversion (DDHI) assumption over \mathbb{G} states that the advantage (function of λ)

$$\left| \Pr \left[b \stackrel{?}{=} \mathcal{A}(\mathbb{G}, \ell, g, y, z) : \begin{array}{l} (\mathbb{G}, \ell, g) \leftarrow \mathbb{G}(\lambda) \\ x \leftarrow_{\mathfrak{s}} \mathbb{Z}_{\ell}^*, b \leftarrow_{\mathfrak{s}} \{0, 1\} \\ y_i \leftarrow g^{x^i} \text{ for } i \in \{1, \dots, q(\lambda)\} \\ \alpha_0 \leftarrow 1/x \bmod \ell, \alpha_1 \leftarrow_{\mathfrak{s}} \mathbb{Z}_{\ell}^* \\ z \leftarrow g^{\alpha b} \end{array} \right] - \frac{1}{2} \right|$$

of any adversary \mathcal{A} that runs in time at most $T(\lambda)$ is at most $\varepsilon(\lambda)$.

Randomness sources and min-entropy. Imperfect randomness is modeled as arbitrary probability distributions with a certain amount of *entropy*. The *min-entropy* notion is used to measure the randomness in such an imperfect random source. A source is said to have k bits of min-entropy if its distribution has the property that each outcome occurs with probability at most 2^{-k} .

Pseudo-Random Functions. A Pseudo Random Function (PRF) [33] is an efficiently computable function of which the values are computationally indistinguishable from uniformly random values.

Formally, a function $\text{PRF}: \mathcal{K}(\lambda) \times \mathcal{X}(\lambda) \rightarrow \mathcal{Y}(\lambda)$ is a (T, q, ε) -secure PRF with key space \mathcal{K} , input space \mathcal{X} and range \mathcal{Y} (all assumed to be finite) if the advantage

$$\left| \Pr \left[1 \leftarrow \mathcal{A}^{\text{PRF}(K, \cdot)} : K \leftarrow_{\mathfrak{s}} \mathcal{K} \right] - \Pr \left[1 \leftarrow \mathcal{A}^{f(\cdot)} : f \leftarrow_{\mathfrak{s}} \mathcal{Y}^{\mathcal{X}} \right] \right|$$

of every adversary \mathcal{A} that runs in time at most $T(\lambda)$ is at most $\varepsilon(\lambda)$.

Dodis–Yampolskiy Pseudo-Random Function. Let \mathbb{G} be a group family generator. The Dodis–Yampolskiy pseudo-random function [28] in an ℓ -order group $(\mathbb{G}, \ell, g) \leftarrow_{\mathfrak{s}} \mathbb{G}$ is the map $F: (K, x) \in \mathcal{K} \times \mathcal{X} \mapsto g^{1/(K+x)} \in \mathbb{G}^*$, with $\mathcal{K} = \mathbb{Z}_{\ell}^*$ and $\mathcal{X} \subset \mathbb{Z}_{\ell}^*$. They proved that it is $(T/(q\ell^{O(1)}), q, \varepsilon q)$ -secure under the (T, q, ε) -DDHI assumption (Definition 2.3) for \mathbb{G} , where $q(\lambda) = O(\log \lambda)$ is an upper-bound on the bit-size of \mathcal{X} for all λ [28, Section 4.2].

3 Model

This section formalizes key-generation protocols for arbitrary, predetermined key-generation algorithms. Such a protocol is executed between a *user* \mathcal{U} and a

certification authority \mathcal{CA} . At the end of the protocol, \mathcal{U} obtains a pair of public–secret keys that \mathcal{CA} certifies to be indistinguishable from keys generated by a fixed algorithm KeyGen , and to have been generated with proper randomness. These requirements are formally captured by a model for randomness verifiability given below. The security definition of the model ensures that a protocol satisfying its conditions fulfills the following properties:

1. \mathcal{CA} can infer no more information about the secret key than it would from a public key generated by KeyGen if \mathcal{U} 's randomness source has high entropy
2. no external attacker can distinguish a public key generated via the protocol from a public key generation with KeyGen if the randomness source of either \mathcal{U} or \mathcal{CA} has high entropy
3. \mathcal{U} cannot bias the generation of the keys if the randomness source of \mathcal{CA} has high entropy. In particular, \mathcal{U} cannot use the public key as a subliminal channel to convey information.

3.1 Syntax

An interactive asymmetric-key-generation protocol is a triple $\text{IKG} = (\text{Setup}, \mathcal{U}, \text{CA})$ of algorithms such that $\text{Setup}(1^\lambda) \rightarrow pp$ is a probabilistic algorithm which returns public parameters and

$$\langle \mathcal{U}(pp; r_{\mathcal{U}}) \rightrightarrows \text{CA}(pp; r_{\mathcal{CA}}) \rangle \rightarrow \langle (pk_{\mathcal{U}}, sk), pk_{\mathcal{CA}} \rangle$$

are interactive algorithms. At the end of the protocol, the user key-generation algorithm \mathcal{U} returns a pair of public–secret keys, and the certificate-authority key-generation algorithm CA returns a public key.

Algorithm Setup may require some randomness, but the parameters it generates can be fixed once for all and used across multi sessions and by several users and authorities. Once parameters are fixed, high-entropy randomness is still needed to securely generate keys, and this is formalized in Section 3.2.

Definition 3.1 (Correctness). *In the \mathcal{O} -oracle model, a key-generation protocol IKG is δ -correct w.r.t. a class \mathcal{A} of algorithms if for all $\lambda \in \mathbb{N}$, for every $\mathcal{A} \in \mathcal{A}$,*

$$\Pr \left[\begin{array}{l} pp \leftarrow_{\S} \text{Setup}(1^\lambda) \\ (D_{\mathcal{U}}, D_{\mathcal{CA}}) \leftarrow_{\S} \mathcal{A}^{\mathcal{O}(\cdot)}(pp) \\ r_{\mathcal{U}} \leftarrow_{\S} D_{\mathcal{U}}, r_{\mathcal{CA}} \leftarrow_{\S} D_{\mathcal{CA}} \\ \langle (pk_{\mathcal{U}}, sk), pk_{\mathcal{CA}} \rangle \leftarrow \langle \mathcal{U}(pp; r_{\mathcal{U}}) \rightrightarrows \text{CA}(pp; r_{\mathcal{CA}}) \rangle \end{array} \right] \geq \delta.$$

The above definition is given in model in which \mathcal{A} has oracle access to \mathcal{O} . This latter may be a random oracle, but it could also simply be an oracle which returns a fixed value (i.e., the common-reference-string model) or no value at all (the standard model). Note that the last line of the probability event implicitly implies that \mathcal{U} and CA *must terminate*.

Init ($1^\lambda, U, CA, I$)	$h \leftarrow_{\S} \Omega; pp \leftarrow_{\S} \text{Setup}(1^\lambda)$ $ID \leftarrow U \cup CA$ for $i \in \llbracket I \rrbracket$ and $id \in ID$ do $st_{id}^i \leftarrow r_{id}^i \leftarrow \perp$ $used_{id}^i \leftarrow \text{FALSE}$ $acc_{id}^i \leftarrow term_{id}^i \leftarrow flag_{id}^i \leftarrow \text{FALSE}$ $sid_{id}^i \leftarrow pid_{id}^i \leftarrow \perp$ $sk_{id}^i \leftarrow pk_{id}^i \leftarrow \perp$ $Q_{\text{Reveal}} \leftarrow Q_{\text{Corrupt}} \leftarrow \emptyset$ return (pin, sin)
Oracle(M)	return $h(M)$
Dist ($id, i, \mathcal{D}_{id}^i$)	$r_{id}^i \leftarrow_{\S} \mathcal{D}_{id}^i$ // r_{id}^i is simply generated and <i>not</i> returned to \mathcal{A}_1
Exec($\mathcal{U}, i, C\mathcal{A}, j$)	if ($\mathcal{U} \notin U$ or $C\mathcal{A} \notin CA$ or $used_{\mathcal{U}}^i$ or $used_{C\mathcal{A}}^j$) return \perp if $r_{\mathcal{U}}^i \neq \perp$ and $r_{C\mathcal{A}}^j \neq \perp$ return $\langle U_i(pp, r_{\mathcal{U}}^i), CA_j(pp, r_{C\mathcal{A}}^j) \rangle$ return \perp // \mathcal{A}_1 must specify distributions beforehand
Send(id, i, M)	if $r_{id}^i = \perp$ return \perp // \mathcal{A}_1 must specify a distribution beforehand if $term_{id}^i$ return \perp $used_{id}^i \leftarrow \text{TRUE}$ $\langle m_{out}, acc, term, sid, pid, pk, sk, st_{id}^i \rangle \leftarrow \langle \text{IKG}(id, st_{id}^i, M; r_{id}^i) \rangle$ if acc and $\neg acc_{id}^i$ $sid_{id}^i \leftarrow sid; pid_{id}^i \leftarrow pid$ $acc_{id}^i \leftarrow acc$ if $term$ and $\neg term_{id}^i$ // Set keys only after termination $pk_{id}^i \leftarrow pk; sk_{id}^i \leftarrow sk$ return ($m_{out}, sid, pid, pk, sk, acc, term_{id}^i$)
Reveal(id, i)	$Q_{\text{Reveal}} \leftarrow Q_{\text{Reveal}} \cup \{(id, i)\}$ return (pk_{id}^i, sk_{id}^i)
Corrupt(id)	$Q_{\text{Corrupt}} \leftarrow Q_{\text{Corrupt}} \cup \{id\}$ for $i \in \llbracket I \rrbracket$ {if $\neg acc_{id}^i$ then $flag_{id}^i \leftarrow \text{TRUE}$ } return $\{st_{id}^i\}_{i \in \llbracket I \rrbracket}$
Test $_b(id^*, i^*)$	if ($\exists(id_0, id_1, i, j): pid_{id_0}^i = id_1$ and $pid_{id_1}^j = id_0$ and $acc_{id_0}^i$ and $\neg term_{id_1}^j$) return \perp // Once an instance accepts, its partner must eventually terminate if $\neg term_{id^*}^{i^*}$ return \perp if $flag_{id^*}^{i^*}$ return \perp // Reject if id^* was corrupt before (id^*, i^*) accepted if $(id^*, i^*) \in Q_{\text{Reveal}}$ or ($\exists(id', j): pid_{id^*}^{i^*} = id'$ and $pid_{id'}^j = id^*$ and $(id', j) \in Q_{\text{Reveal}}$) return \perp // Reject if the key of (id^*, i^*) or of its partner has been revealed if $pk_{id^*}^{i^*} \neq \perp$ if $b = 0$ $(pk, sk) \leftarrow_{\S} \text{KeyGen}(1^\lambda)$ return pk return $pk_{id^*}^{i^*}$ return \perp // Reject if (id^*, i^*) does not have a key

Fig. 1. Oracles for the Key-Generation Indistinguishability Experiment.

3.2 Security

This section gives a game-based security model for key-generation protocols with verifiable randomness. It covers concurrent protocol executions with different instances of protocol algorithms. It is inspired by the BPR model for authenticated key exchange [10].

Protocol Participants. A set of user identities U and a set of certificate-authority identities CA are assumed to be fixed. The union of the those sets form the overall identity space ID . For readability, it is implicitly assumed that during protocol executions, the messages exchanged are always prepended with the instance identifier of the receiving party. Note that several instances of the same algorithm may concurrently run during the game.

Adversaries. The game features a two-stage adversary $(\mathcal{A}_1, \mathcal{A}_2)$. Adversaries \mathcal{A}_1 and \mathcal{A}_2 may agree on a common strategy before the beginning of the game. That is to say, the strategy may be part of their code, and it may dictate which queries to make (possibly depending on the oracle answers), the order of the queries and so forth. All but the challenge query can only be made by \mathcal{A}_1 . The role of \mathcal{A}_2 is essentially only to guess whether a public key was generated with KeyGen or with the protocol, while \mathcal{A}_1 can make arbitrary queries according to the pre-established strategy.

However, \mathcal{A}_1 and \mathcal{A}_2 cannot communicate after the beginning of the game. It reflects the fact that in practice, an implementer may distribute its key generator, but does not necessarily wiretap the execution of the key-generation protocol for a particular user. From a technical viewpoint, the reason is that in a key-generation protocol, a user has to prove to the authority that she correctly performed her computation. However, the randomness used in these proofs can be used as a subliminal channel to convey information about the secret key. For instance, an engineer could in practice implement a bogus key generator which only terminates the protocol if the first bits of the proof and secret key match. The proof then serves as subliminal channel to leak information about the secret key. Later on, when a user wants to generate a certified public key, if the engineer could wiretap the protocol execution, he could infer some information about her secret key through the proof of correct computation. It is the reason why communication between the adversaries cannot be allowed.

The restriction that \mathcal{A}_1 and \mathcal{A}_2 cannot communicate after the beginning of the game means that the attacks in which the protocol executions are listened to are excluded, but as explained above, it seems to be a minimal requirement.

Game Overview. At the beginning of the game, the challenger first runs an initialization algorithm. After that, \mathcal{A}_1 can make several queries to the algorithm instances. It can in particular

- * specify distributions from which randomness is drawn and given as an input to the instances,
- * ask for the protocol to be executed between different instances of the protocol algorithms without its intervention, i.e., perform passive attacks,

- * perform active attacks by sending messages to algorithm instances of its choice,
- * later on reveal the keys that were generated by a particular instance,
- * corrupt a party (user or certificate authority), and thereby gain access to the state of all its algorithm instances.

As for \mathcal{A}_2 , it can reveal keys or make a test query that returns either (with probability 1/2 each) keys freshly generated by the key-generation algorithm or keys generated by instances of its choice via queries made by \mathcal{A}_1 . Adversary \mathcal{A}_2 must eventually return a guess for the origin of the keys it was returned, and $(\mathcal{A}_1, \mathcal{A}_2)$ wins the game if the guess of \mathcal{A}_2 is correct.

Initialization & Game Variables. During the initialization phase, game variables are declared for every instance of the protocol algorithms. Assume that there are at most $I = I(\lambda)$ instance of any participant id . Each instance $i \in I$ of a participant id maintains a state st_{id}^i . A session identity sid_{id}^i and a partner identity pid_{id}^i allow to match instances together in protocol executions. It is assumed that for each sid_{id}^i there can be at most one partner instance, i.e., one pair (id', j) such that $pid_{id}^i = id'$ and $sid_{id}^i := (id, i, id', j, sid_{id'}^j)$.

Public–secret-key variables (respectively denoted pk_{id}^i and sk_{id}^i) hold the public and secret keys that were output, if any, by the i th instance of the algorithm of participant id at that step of the computation. For certificate authorities, the secret keys are always set to \perp .

A variable $used_{id}^i$ indicates whether the adversary has performed an active attack on the i th algorithm instance of participant id .

Variables acc_{id}^i and $term_{id}^i$ respectively indicate whether the algorithm of the i th instance of participant id has accepted and terminated. As in the BPR model [10], termination and acceptance are distinguished. When an instance terminates, it does not output any further message. However, it may accept at a certain point of the computation, and terminate later. In the present context, it may for instance occur when an instance expect no further random input from its partner instance, and the rest of its computation is purely deterministic. It may then only terminate after finishing its computation. This distinction is crucial for the security definition. It is important to exclude the trivial case in which, although every computation was honestly performed, a user discards the public key if it does not follow a certain pattern, thereby influencing the distribution of the output public key (i.e., perform rejection sampling), and possibly using it a subliminal channel to convey information about the secret key.

Another variable $flag_{id}^i$ indicates whether party id was corrupted before its i th instance had accepted. Recall that acceptance intuitively means that an instance expects no further random input from its partner instance. As long as $flag_{id}^i$ is set to FALSE, the only information the adversary has about r_{id}^i is its distribution and therefore, if this distribution has high min-entropy, the adversary cannot bias the generation of the keys.

A variable r_{id}^i holds the random string to be used the i th instance of the algorithm of id .

The challenger maintains a set (initially empty) Q_{Reveal} of identity–instance pairs of which the keys were revealed. It also maintains a set (initially empty) Q_{Corrupt} of corrupt identities.

At the end of the initialization phase, the encryption parameters, the sets of participants and the user public keys are returned in a public input pin , and the rest is set in a secret input sin . That is, $pin \leftarrow (pp, U, CA, I, (pk_{id})_{id})$ and $sin \leftarrow (pin, (sk_{id})_{id}, (st_{id}^i, sid_{id}^i, pid_{id}^i, acc_{id}^i, term_{id}^i, used_{id}^i)_{i,id}, Q_{\text{Corrupt}}, Q_{\text{Reveal}})$. The secret input sin is later made available to all oracles.

Oracles. Throughout the game, adversary \mathcal{A}_1 is given access to the oracles summarized below and defined in Figure 1. It can query them *one at a time*.

- * **Oracle** : gives access to a function h chosen uniformly at random from a probability space Ω . The adversary and the protocol may depend on h . The probability space Ω specifies the model in which the protocol is considered. If it is empty, then it is the standard model. If it a space of random functions, then it is the random oracle model. As for the Common-Reference String (CRS) model, Ω is a space of constant functions (i.e., the output of Oracle is the same for all input M).
- * **Dist** : via this oracle, the adversary specifies the distribution \mathcal{D}_{id}^i from which the randomness of the i th instance of id is drawn. These distributions are always assumed to be independent of oracle Oracle. However, the distributions specified by the adversary for different instances can be correlated in any way. Oracle Dist then generates a bit string r_{id}^i according to the input distribution and does *not* return it to the adversary. Whenever oracle Exec or Send is queried on (id, i) , it uses randomness r_{id}^i for its computation. This new (compared to the BPR model) oracle is essential to express requirements on the minimal entropy used by the instances, and also to express reasonable winning conditions. It allows to properly capture properties like the fact that (1) the authority cannot infer any information about the secret key if the randomness of the user algorithm has high entropy, (2) that the output keys are indistinguishable from keys generated with the key-generation algorithm if the randomness used by the algorithm of either of the parties has high entropy, or (3) that a user algorithm cannot bias the distribution of the output keys if the randomness of the authority algorithm has high entropy.
- * **Exec** : returns the transcript of an honest (i.e., without the interference of the adversary) protocol execution between the i th instance of U and the j th instance of CA . The protocol is executed with the random strings generated for these instances by oracle Dist on the input of adversarial distributions. The notations U_i and CA_j mean that algorithms U and CA are executed using the state of the i th instance of U and the j th instance of CA respectively. It is implicitly assumed that the states $acc_{U_i}^i, term_{U_i}^i, acc_{CA_j}^j$ and $term_{CA_j}^j$ are set to TRUE after an honest protocol execution. Moreover, if the termination variable of either party is set to TRUE, the protocol is not executed and \perp is returned.

In essence, by querying oracle `Exec`, adversary \mathcal{A}_1 performs a passive eavesdropping attack.

- * `Send` : adversary \mathcal{A}_1 can perform active attack via this oracle. \mathcal{A}_1 can send any message to an instance of its choice, e.g., the i th instance of a user algorithm, which runs the honest protocol algorithm of the corresponding party on the input of the message chosen by the adversary.

To prompt the i th instance of id to initiate a protocol execution with the j th instance of id' , adversary \mathcal{A}_1 can make a `Send` query on $(id, i, (id', j))$.

$\text{IKG}(id, *)$ denotes the `IKG` algorithm of party id , i.e., either `U` or `CA`. The algorithm is executed using the randomness generated by oracle `Dist` for that instance. (Note that the input random string may be used only at certain steps of the computation.) The oracle then returns the output of the instance to the adversary. It also specifies if this instance accepted and/or terminated, and returns the session identifier and the identity of its partner in the protocol execution, as well as the public and secret keys returned by this instance, if any. Note that if the instance is that of a certificate-authority algorithm, the secret key is always set to \perp .

- * `Reveal` : on input (id, i) , returns the keys held by the i th instance of the algorithm of id . The couple (id, i) is added to the set Q_{Reveal} of revealed keys.
- * `Corrupt` : on input id , returns the states of all the instances of the algorithm of id . The identity id is added to the set Q_{Corrupt} of corrupt identities.

As for adversary \mathcal{A}_2 , it is given access to oracles `Oracle`, `Reveal` and to oracle

- * `Testb` : on input (id^*, i^*) , it returns the public key $pk_{id^*}^{i^*}$ generated via `IKG` (with an `Exec` query or `Send` queries) if $b = 0$ or a fresh public key generated via `KeyGen` if $b = 1$.

An important restriction on this query is that the following condition must be satisfied: for any instance i of the algorithm of a party id_0 , once it has accepted, i.e., once $acc_{id_0}^i$ is set to `TRUE`, the partner instance algorithm, say the j th instance of id_1 , must eventually terminate, i.e., $term_{id_1}^j$ must have been set to `TRUE` as well by the time of query `Test`. It prevents \mathcal{A}_1 from biasing the distribution of the keys by prematurely aborting the protocol although it was followed, if the resulting key does not follow a certain pattern, and which would allow \mathcal{A}_2 to guess b with a non-negligible advantage.

The other restrictions are simply that i^* -th instance of id^* must have terminated, that id^* was not corrupt before (id^*, i^*) had accepted⁶, that neither the key of the i^* th instance of id^* nor of its partner instance has been revealed, and that the i^* th instance of id^* must already hold a key.

Note that \mathcal{A}_2 can query `Test` only once. A definition in which multiple queries were allowed would be asymptotically equivalent via a standard hybrid argument.

Adversary \mathcal{A}_2 must eventually return a bit b' as a guess for the origin (i.e., either `IKG` or `KeyGen`) of the key returned by oracle `Testb`.

⁶ To understand why it is necessary for id^* not to be corrupt before (id^*, i^*) accepts even though \mathcal{A}_1 and \mathcal{A}_2 do not communicate, suppose that this condition were not imposed and consider the following strategy which allows $(\mathcal{A}_1, \mathcal{A}_2)$ to trivially win:

To achieve any form of indistinguishability from a key-generation algorithm, it is clear that either the distribution $\mathcal{D}_{id^*}^{i^*}$ or the distributions $\mathcal{D}_{id'}^j$ for the partner instance (j, id') of (i^*, id^*) must have high entropy. Indeed, if distributions with low entropy were allowed, \mathcal{A}_1 and \mathcal{A}_2 could agree on these identities, instances and distributions beforehand. Adversary \mathcal{A}_2 could then simply return 1 if and only if the challenge key is the most likely key w.r.t. $\mathcal{D}_{id^*}^{i^*}$ and $\mathcal{D}_{id'}^j$, and thereby win the game with a non-negligible advantage.

A parameter κ for the maximal min-entropy of $\mathcal{D}_{id^*}^{i^*}$ and $\mathcal{D}_{id'}^j$, specified by \mathcal{A}_1 is therefore introduced. If the adversary modified any message from the partner (j, id') of (id^*, i^*) before (id^*, i^*) accepts, then $\mathcal{D}_{id'}^j$ is set to be the Dirac mass at the zero bit-string by convention (and it thus has no entropy). The underlying idea is that as long as at least one of the two parties has a randomness source with high entropy, the key returned at the end of the protocol should be indistinguishable from a key generated by the key-generation algorithm. In particular, if the entropy of the user's source is high enough, (1) the authority no more information about the secret key than it would from a public key generated by KeyGen; (2) if either of the parties uses randomness with high entropy, no external attacker can distinguish a key generated via IKG from a key generated with KeyGen; and (3) if the authority's source has high entropy, the user algorithm cannot bias the generation of the key. The security of a key-generation protocol is then defined for adversaries that specify challenge distributions with min-entropy at least κ .

Definition 3.2 (Indistinguishability). *An interactive key-generation protocol IKG is $(T, q_{\text{Oracle}}, q_{\text{Dist}}, q_{\text{Exec}}, q_{\text{Send}}, q_{\text{Reveal}}, q_{\text{Corrupt}}, \kappa, \varepsilon)$ -indistinguishable from a key-generation algorithm KeyGen (running on uniform randomness) if for all $\lambda \in \mathbb{N}$, for every adversary $(\mathcal{A}_1, \mathcal{A}_2)$ that runs in time at most $T(\lambda)$ and makes at most $q_{\mathcal{O}}$ queries for $\mathcal{O} \in \{\text{Oracle}, \text{Dist}, \text{Exec}, \text{Send}, \text{Reveal}, \text{Corrupt}\}$, and such that $\max(H_{\infty}(\mathcal{D}_{id^*}^{i^*}), H_{\infty}(\mathcal{D}_{id'}^j)) \geq \kappa$ for query Test, the advantage (function of λ)*

$$\Pr \left[\begin{array}{l} (pin, sin) \leftarrow \text{Init}(1^\lambda, U, CA, I) \\ \mathcal{O}_1 \leftarrow \{\text{Oracle}, \text{Dist}, \text{Exec}, \text{Send}, \text{Reveal}, \text{Corrupt}\} \\ \mathcal{A}_1^{O_1(sin, \cdot)}(pin) \\ b = b': b \leftarrow_{\S} \{0, 1\} \\ \mathcal{O}_2 \leftarrow \{\text{Oracle}, \text{Reveal}, \text{Test}_b\} \\ b' \leftarrow \mathcal{A}_2^{O_2(sin, \cdot)}(pin) \\ \text{return}(b, b') \end{array} \right] - 1/2$$

of $(\mathcal{A}_1, \mathcal{A}_2)$ is at most $\varepsilon(\lambda)$.

\mathcal{A}_1 and \mathcal{A}_2 agree on (id^*, i^*) and on a distribution $\mathcal{D}_{id^*}^{i^*}$. Adversary \mathcal{A}_1 prompts (id^*, i^*) to initiate a protocol execution by making a Send query. It then corrupts id^* and obtains $st_{id^*}^{i^*}$, from which it can read $r_{id^*}^{i^*}$. Adversary \mathcal{A}_1 could then play the role of its partner and adapt the messages it sends to make sure that the resulting public key follows a certain pattern known to \mathcal{A}_2 . This latter would then be able to win the game with a non-negligible advantage.

From a practical perspective, this definition (which implies requirement 3 as it enforces indistinguishability from fresh keys generated by IKG) means that keys generated via a protocol satisfying the definition above are not subject to randomness vulnerabilities such as the ROCA vulnerabilities [43] mentioned in introduction (in which only specific primes were selected by user key generation algorithms) or such as those [38,35] in which several public keys are generated with correlated randomness sources.

4 Generic Constructions

This section presents two generic protocols for key generation with verifiable randomness: a protocol that covers a wide class of key-generation algorithms, namely those that can be represented as probabilistic circuits, and another protocol specific to the generation of RSA keys. The first protocol is of theoretical interest and shows that randomness verifiability can be achieved for wide class of key-generation algorithms, whereas the second protocol is a solution that can actually be used in practice.

4.1 Key-Generation Protocol with Verifiable Randomness for Probabilistic Circuits

This section gives a key-generation protocol with verifiable randomness for a large class of key-generation algorithms. The focus is here on the class of key-generation algorithms that can be modeled as *probabilistic circuits*.

The advantage of probabilistic circuits compared to general Turing Machines for this purpose is that the running time of a probabilistic circuit is independent of the random inputs. In a key-generation protocol with verifiable randomness, the user has to prove to the authority that she correctly performed her computation. Having a constant running time then ensures that no one can infer any information about the secret key from the statement proved by the user or the proof itself. It prevents malicious user algorithms from using the proofs as subliminal channels to pass information about the secret key.

To understand why it is important for the running time to be constant, consider the artificial random number generator described on Algorithm 1. To generate a k -bit string, it essentially consists in flipping a random coin several times for each bit and to set this bit to the parity of the number of flipped coins to obtain the first “Head”. It produces a k -bit string uniformly distributed within expected time complexity $O(k)$ and it could be used as a secret-key generation algorithm (and the public key would then be a deterministic function of the generated secret key). As a simple example, consider the ElGamal key-generation algorithm defined in a cyclic group $\langle g \rangle$ of prime order p as $sk \leftarrow_{\$} \mathbb{Z}_p^*$, $pk \leftarrow g^{sk}$. In this example, $k = \lceil \log_2 p \rceil + 1$.

For a user to prove that she correctly generated the random bit string x , she would have to commit to the x_i values and compute a proof on the s values. However, each x_i is simply the parity of the number of trials before $s = 0$.

Algorithm 1 RNG with Non-Constant Running Time.

Require: Integer k .

Ensure: Uniformly random k bit-string x .

```
1: for  $i = 0$  to  $k - 1$  do
2:    $c \leftarrow 0$ 
3:   while TRUE do
4:      $c \leftarrow (c + 1 \bmod 2)$ 
5:      $s \leftarrow_{\S} \{0, 1\}$ 
6:     if  $s = 0$  then
7:        $x_i \leftarrow c$ 
8:       break
9:     end if
10:  end while
11: end for
12:  $x \leftarrow x_0 \parallel \dots \parallel x_{k-1}$ 
13: return  $x$ 
```

Therefore, from the number of s values for which the user has to perform a proof, the authority can infer x_i . For example, if the user generated two s values for x_1 , the authority knows that $x_1 = 0$.

In other words, the statement of the proof itself reveals some information about the secret key to the certification authority; and the issue is here that the running time changes from one random run of the algorithm to the other. Restricting to probabilistic circuits eliminates this issue.

The restriction to circuits comes at a cost though. It for instance excludes the class of algorithms for which there is no known circuit that can represent them. It is for instance the case of algorithms that must efficiently generate primes during the process. Indeed, there is no known circuit that can efficiently generate prime numbers (without resorting to number-theoretic assumptions). On this ground, the generic protocol for probabilistic circuits of Section 4.1 does not apply to the RSA-key generation for instance⁷. See rather Section 4.2 for an RSA key-generation protocol with verifiable randomness for arbitrary properties that the keys must satisfy.

Before describing our protocol, we first formally define probabilistic circuits.

Probabilistic Circuits. A probabilistic circuit (as defined by Belaïd et al. [6]) is essentially a deterministic circuit augmented with uniformly random gates.

⁷ One can construct families of probabilistic circuits which output an RSA key with overwhelming probability by relying on the prime number theorem and Chernoff's bound. However, to obtain a b -bit RSA public key with probability $1 - 2^{-\lambda}$, such constructions would have $O(\lambda^2 b^4)$ gate complexity and $O(\lambda^2 b^2)$ randomness complexity (based on Miller-Rabin's primality test) or $O(\lambda b^7)$ gate complexity and $O(\lambda b^2)$ randomness complexity (based on the most efficient variant of deterministic Agrawal-Kayal-Saxena primality test [2]). Applying our generic construction to such circuits family would result in schemes with prohibitive efficiency.

The random gates produce independent and uniform random bits that are sent along their output wires.

We equivalently define a probabilistic circuit as a uniform random variable over a finite collection of deterministic boolean circuits. These boolean circuits are restricted to have the same amount n of input variables, and r fixed inputs. The number r of fixed inputs depends on the security parameter 1^λ . Denote such a circuit as $\Gamma_{b_1 \dots b_r}(x_1, \dots, x_n)$, with x_1, \dots, x_n the input variables and b_1, \dots, b_r the fixed inputs. To each element in $\{0, 1\}^r$ corresponds a circuit in the collection with the bit string as fixed inputs, so that there are 2^r circuits in the collection. However, these circuits are not required to form a uniform family (i.e., they are not required to be output by a single Turing machine); the circuit families here considered can be non-uniform.

A probabilistic circuit Γ is then defined as a uniform random variable over the set (of circuits) $\{\Gamma_b\}_{b \in \{0,1\}^r}$. Namely, for input variables x_1, \dots, x_n , the evaluation $\Gamma(x_1, \dots, x_n)$ is a uniform random variable over the set (of values) $\{\Gamma_b(x_1, \dots, x_n)\}_{b \in \{0,1\}^r}$. If $\omega \in \{0, 1\}^r$ denotes the random input to the probabilistic circuit Γ , the evaluation $\Gamma(x_1, \dots, x_n; \omega)$ is then $\Gamma_\omega(x_1, \dots, x_n)$.

The advantage of this second definition is that randomness is invoked only once instead of invoking it for each of the r random gates. To generate keys, PRFs are often used to provide long random bit strings from small secret seeds. As the goal is to build a key-generation protocol which allows the authority to certify that the keys are generated with high-entropy randomness, the user will have to prove in zero-knowledge that she correctly performed the pseudo-random evaluations. Invoking randomness only once then allows to invoke the PRF only once in the protocol.

Generic Protocol. We now give a two-party protocol in the CRS model to generate, with verifiable randomness, keys computed by probabilistic circuits. Requiring that keys are generated with verifiable randomness here means that the random inputs to the circuits must be uniformly generated in a verifiable manner. The deterministic inputs to the circuits can simply be considered as public parameters.

Building Blocks. The protocol involves

- a function family $\mathcal{H} = \{H_{hk}\}_{hk \in \{0,1\}^{d(\lambda)}}$ which is a universal computational extractor w.r.t. unpredictable sources (App. A.5)
- a two-source extractor Ext (App. A.4) with key space $\{0, 1\}^{\delta(\lambda)}$
- an extractable commitment scheme $\mathcal{C} = (\text{Setup}, \text{Com}, \text{ComVf}, \text{TSetup}, \text{ExtCom})$ (App. A.1) for the user algorithm to commit to its random string *before receiving any input from the CA*, thereby preventing it from biasing the distribution of the keys. The parameters returned by Setup are implicit inputs to the other algorithms of \mathcal{C}

- a non-interactive, extractable, zero-knowledge proof system (App. A.3) $\Pi = (\text{Setup}, \text{Prove}, \text{Verf}, \text{TSetup}^{zk}, \text{Sim}, \text{TSetup}^{ext}, \text{Ext})$ for relation

$$\mathcal{R}_\Pi := \left\{ \left((x_i)_i, k, \mathcal{C}, r_{\mathcal{CA}}, pk; r'_U, d, sk \right) : \text{ComVf}(\mathcal{C}, r'_U, d) = 1 \right. \\ \left. \wedge (pk, sk) = \Gamma(x_1, \dots, x_n; \text{Ext}_k(r'_U, r_{\mathcal{CA}})) \right\},$$

- a pseudo-random function PRF to generate the randomness necessary to compute $\Pi.\text{Prove}$.

Parameters. Given a circuit Γ with deterministic inputs x_1, \dots, x_n , to generate public parameters for the protocol on the input of a security parameter 1^λ , run $pp_{\mathcal{C}} \leftarrow \mathcal{C}.\text{Setup}(1^\lambda)$ ($pp_{\mathcal{C}}$ is a tacit input to the algorithms of \mathcal{C}), $crs \leftarrow \Pi.\text{Setup}(1^\lambda)$, and generate $hk \leftarrow_{\S} \{0, 1\}^{d(\lambda)}$ and $k \leftarrow_{\S} \{0, 1\}^{\delta(\lambda)}$. Return $pp \leftarrow (crs, pp_{\mathcal{C}}, hk, k, x_1, \dots, x_n)$.

Formal Description. Consider the interactive protocol IKG_Γ on Figure 2 between a user \mathcal{U} and a certification authority \mathcal{CA} . Each algorithm maintains acceptance and termination variables acc_{id} and $term_{id}$, for $id \in \{\mathcal{U}, \mathcal{CA}\}$, initially set to FALSE. On the input of pp and of their respective random strings r_U and $r_{\mathcal{CA}}$, the party algorithms proceed as follows:

1. \mathcal{U} separates the domain of H_{hk} in two and applies it to its randomness. It commits to the first output with the second output as randomness, and sends the resulting commitment \mathcal{C} to \mathcal{CA}
2. \mathcal{CA} , upon receiving the commitment from \mathcal{U} , sets $acc_{\mathcal{CA}} \leftarrow \text{TRUE}$ and sends its random string $r_{\mathcal{CA}}$ to \mathcal{U}
3. \mathcal{U} , upon receiving $r_{\mathcal{CA}}$ from \mathcal{CA} , sets $acc_{\mathcal{U}} \leftarrow \text{TRUE}$. Next, it extracts a seed s with Ext from the joint randomness. It evaluates Γ on x_1, \dots, x_n and s , and obtains a key pair (pk, sk) . It generates another seed s' with H_{hk} . Algorithm \mathcal{U} then evaluates PRF mode on s' to generate the randomness necessary to compute $\Pi.\text{Prove}$ since \mathcal{U} has no other random string than r_U available, i.e., it computes $r_\Pi \leftarrow \text{PRF}(s', 0)$. Algorithm \mathcal{U} then proves that it followed the protocol and correctly evaluated Γ at x_1, \dots, x_n , i.e., it computes a proof $\pi \leftarrow \Pi.\text{Prove}(crs, ((x_i)_i, k, \mathcal{C}, r_{\mathcal{CA}}, pk), (r'_U, d, sk); r_\Pi)$. After that, it erases all variables but pk, sk, π , sends pk and π to \mathcal{CA} , returns (pk, sk) and sets $term_{\mathcal{U}} \leftarrow \text{TRUE}$
4. \mathcal{CA} , upon receiving (pk, π) from \mathcal{U} , verifies the proof. If the proof is valid, it returns pk , otherwise it returns \perp . It then sets $term_{\mathcal{CA}} \leftarrow \text{TRUE}$.

Theorem 4.1 (Correctness). *Protocol IKG_Γ is 1-correct w.r.t. all algorithms if \mathcal{C} is correct and if Π is complete.*

Proof. The theorem immediately follows from the correctness of \mathcal{C} and the completeness of Π . \square

Theorem 4.2 (Indistinguishability). *Suppose that \mathcal{H} is $(T_{\mathcal{H}}^{\text{uce}}, 3, \varepsilon_{\mathcal{H}}^{\text{uce}})$ -UCE secure w.r.t. simply unpredictable sources of min-entropy at least $\kappa_{\mathcal{H}}$. Suppose*




Fig. 2. Key-Generation Protocol with Verifiable Randomness for Probabilistic Circuits.

also that Ext is a $(\kappa_{\text{Ext}}, \varepsilon_{\text{Ext}})$ -extractor for $\kappa_{\text{Ext}} \leq \min(|r'_{\mathcal{U}}|, |r_{C\mathcal{A}}|)$. If \mathcal{C} is $(T_{\mathcal{C}}^{\text{hide}}, \varepsilon_{\mathcal{C}}^{\text{hide}})$ -hiding, and satisfies $(T_{\mathcal{C}}^{\text{setup-ind}}, \varepsilon_{\mathcal{C}}^{\text{setup-ind}})$ -setup indistinguishability and $(T_{\mathcal{C}, \text{ExtCom}}^{\text{biding}}, T_{\mathcal{C}, \mathcal{A}}^{\text{biding}}, 0, \varepsilon_{\mathcal{C}}^{\text{biding}})$ -biding extractability, if Π is $(T_{\text{Ext}}^{\text{ext}}, T_{\mathcal{A}, \Pi}^{\text{ext}}, \varepsilon^{\text{ext}})$ -extractable and $(T_{\Pi}^{\text{zk}}, \varepsilon^{\text{zk}})$ -composable zero-knowledge, and if PRF is $(T_{\text{PRF}}, 1, \varepsilon_{\text{PRF}})$ -secure, then protocol IKG_{Γ} is $(T, q_{\mathcal{O}}, \kappa, \varepsilon)$ -indistinguishable from Γ in the CRS model, for $T = \min(T_{\mathcal{H}}^{\text{uce}}, T_{\mathcal{H}}^{\text{hide}}, T_{\mathcal{C}}^{\text{setup-ind}}, T_{\mathcal{C}, \text{ExtCom}}^{\text{biding}}, T_{\mathcal{C}, \mathcal{A}}^{\text{biding}}, T_{\Pi, \text{Ext}}^{\text{ext}}, T_{\Pi, \mathcal{A}}^{\text{ext}}, T_{\Pi}^{\text{zk}}, T_{\text{PRF}})$, $\mathcal{O} \in \{\text{Dist}, \text{Exec}, \text{Send}, \text{Reveal}, \text{Corrupt}\}$, $\kappa = \max(\kappa_{\mathcal{H}}, \kappa_{\text{Ext}})$ and $\varepsilon := 5\varepsilon_{\mathcal{H}}^{\text{uce}} + 5\varepsilon_{\text{Ext}} + \varepsilon_{\text{PRF}} + \varepsilon_{\Pi}^{\text{zk}} + \varepsilon_{\Pi}^{\text{ext}} + \varepsilon_{\mathcal{C}}^{\text{setup-ind}} + \varepsilon_{\mathcal{C}}^{\text{bind-ext}} + \varepsilon_{\mathcal{C}}^{\text{hide}}$.

Proof. First note that if $(\mathcal{A}_1, \mathcal{A}_2)$ wins the indistinguishability game with a probability at least ε , then there exists an adversary $(\mathcal{A}'_1, \mathcal{A}'_2)$ which wins with probability at least $\varepsilon/|ID||I|$ a variant of the game in which the adversary is required to specify the pair (id^*, i^*) of its Test query before being given access to the game oracles, i.e., a *selective* variant of the indistinguishability game. Indeed, $(\mathcal{A}'_1, \mathcal{A}'_2)$ can simply run $(\mathcal{A}_1, \mathcal{A}_2)$ as a subroutine and guess the pair (id^*, i^*) at the beginning of the game. If $(\mathcal{A}_1, \mathcal{A}_2)$ later makes its Test query on a different pair, $(\mathcal{A}'_1, \mathcal{A}'_2)$ simply aborts and sends to the challenger a bit chosen uniformly at random.

A message is subsequently said to be *oracle-generated* if it was computed by the challenger as a response to a Send query and it was not altered. Otherwise, the message is said to be *adversarially generated*.

Further distinguish two cases

1. $id^* \in CA$, or $id^* \in U$ and the random string $r_{C\mathcal{A}}$ the i^* th instance of id^* receives is oracle-generated.
2. $id^* \in U$ and the random string $r_{C\mathcal{A}}$ the i^* th instance of id^* receives is adversarially generated.

In the first case, as the commitment scheme satisfies binding extractability and as the proof system satisfies simulation extractability, the randomness used to generate the keys is really drawn from $\mathcal{D}_{id^*}^{i^*}$ and $\mathcal{D}_{id'}^j$. Moreover, since either $\mathcal{D}_{id^*}^{i^*}$ or $\mathcal{D}_{id'}^j$ has min-entropy at least κ , family \mathcal{H} and extractor Ext guarantee that the public key generated during the protocol execution is indistinguishable from a one from the key-generation algorithm.

To prove it, consider the following sequence of games.

Game 0. This is the real selective game.

Game 1. In this game, to generate the common-reference string, the challenger runs $(crs, \tau_{ext}) \leftarrow \Pi.\text{TSetup}^{ext}(1^\lambda)$ instead of $\Pi.\text{Setup}$. This game is perfectly indistinguishable from the previous one since Π is extractable.

Game 2. To answer an Exec query on $(id^*, i^*, *, *)$ or on $(*, *, id^*, i^*)$, the challenger generates a transcript of an honest protocol execution. However, instead of setting $pk_{id^*}^{i^*} \leftarrow pk$, it generates a uniformly random string σ , runs $(pk', sk') \leftarrow \Gamma(x_1, \dots, x_n; \sigma)$ and sets $pk_{id^*}^{i^*} \leftarrow pk'$.

Recall that \mathcal{A}'_1 and \mathcal{A}'_2 share no state, and that in the event in which the Test query of \mathcal{A}'_2 is not replied to with \perp , key $pk_{id^*}^{i^*}$ is not in $\mathcal{Q}_{\text{Reveal}}$ (nor is the key of the partner instance of (id^*, i^*)). Denoting by (id, j') the partner instance of (id^*, i^*) , since $\max(H_\infty(\mathcal{D}_{id^*}^{i^*}), H_\infty(\mathcal{D}_{id'}^j)) \geq \kappa$, adversary $(\mathcal{A}'_1, \mathcal{A}'_2)$ can distinguish this game from the previous one with an advantage of at most $\varepsilon_{\mathcal{H}}^{\text{uce}} + \varepsilon_{\text{Ext}}$.

Indeed, first suppose that $id^* \in U$. If $H_\infty(\mathcal{D}_{id^*}^{i^*}) \geq \kappa \geq \kappa_{\mathcal{H}}$, then $r'_{\mathcal{U}}$ is $\varepsilon_{\mathcal{H}}^{\text{uce}}$ -computationally indistinguishable from a uniformly random value. Since $|r'_{\mathcal{U}}| \geq \kappa_{\text{Ext}}$, $\max(H_\infty(r'_{\mathcal{U}}), H_\infty(r'_{id'}^j)) \geq \kappa_{\text{Ext}}$ and s is then ε^{ext} -computationally indistinguishable from a uniformly random value. On the other hand, if $H_\infty(\mathcal{D}_{id'}^j) \geq \kappa \geq \kappa_{\text{Ext}}$ then $\max(H_\infty(r'_{\mathcal{U}}), H_\infty(r_{\mathcal{C}\mathcal{A}})) \geq \kappa_{\text{Ext}}$ and s is thus ε^{ext} -computationally indistinguishable from a uniformly random source. In either sub-case, $(\mathcal{A}'_1, \mathcal{A}'_2)$ can distinguish this game from the previous one with an advantage of at most $\varepsilon_{\mathcal{H}}^{\text{uce}} + \varepsilon_{\text{Ext}}$.

In case $id^* \in CA$, a symmetric argumentation implies that $(\mathcal{A}'_1, \mathcal{A}'_2)$ can distinguish this game from the previous one with an advantage of at most $\varepsilon_{\mathcal{H}}^{\text{uce}} + \varepsilon_{\text{Ext}}$.

Game 3. If $id^* \in U$, the challenger answers a Send query on $(id^*, i^*, r_{\mathcal{C}\mathcal{A}})$ as follows. It computes (pk, π) honestly, but instead of setting $pk_{id^*}^{i^*} \leftarrow pk$, it generates a uniformly random string σ , runs $(pk', sk') \leftarrow \Gamma(x_1, \dots, x_n; \sigma)$ and sets $pk_{id^*}^{i^*} \leftarrow pk'$.

Recall that in case 1), since $r_{\mathcal{C}\mathcal{A}}$ is always oracle-generated. Therefore, the same indistinguishability between the last two games still apply and $(\mathcal{A}'_1, \mathcal{A}'_2)$ can distinguish this game from the previous one with an advantage of at most $\varepsilon_{\mathcal{H}}^{\text{uce}} + \varepsilon_{\text{Ext}}$.

Game 4. If $id^* \in CA$, the challenger answers a Send query on $(id^*, i^*, (pk, \pi))$ as follows. If (pk, π) is adversarially generated and valid, the challenger runs

$(r'_{\mathcal{U}}, d, sk) \leftarrow \Pi.\text{Ext}(crs, \tau_{ext}, ((x_i)_i, k, C, r_{C\mathcal{A}}, pk), \pi)$ and checks whether $((x_i)_i, k, C, r_{C\mathcal{A}}, pk; r'_{\mathcal{U}}, d, sk)$ is in \mathcal{R}_Π . If not, it aborts.

This game can be distinguished from the one only if the extractability of Π is contradicted. Therefore, $(\mathcal{A}'_1, \mathcal{A}'_2)$ can distinguish this game from the one with an advantage of at most $\varepsilon_\Pi^{\text{ext}}$.

Game 5. The challenger now runs $(pp_{\text{Com}}, \tau_{\text{Com}}) \leftarrow \mathcal{C}.\text{TSetup}(1^\lambda)$ instead of running $\mathcal{C}.\text{Setup}$. The setup indistinguishability of \mathcal{C} implies that adversary $(\mathcal{A}'_1, \mathcal{A}'_2)$ can distinguish this game from the previous one with an advantage of at most $\varepsilon_{\mathcal{C}}^{\text{setup-ind}}$.

Game 6. If $id^* \in CA$, the challenger of this game answers a `Send` query on $(id^*, i^*, (pk, \pi))$ as follows. If (pk, π) is adversarially generated and valid, the challenger runs $(r'_{\mathcal{U}}, \rho_{\mathcal{U}}, sk) \leftarrow \Pi.\text{Ext}(crs, \tau_{ext}, ((x_i)_i, k, C, r_{C\mathcal{A}}, pk), \pi)$. If $((x_i)_i, k, C, r_{C\mathcal{A}}, pk; r'_{\mathcal{U}}, d, sk)$ is in \mathcal{R}_Π , it extracts $r''_{\mathcal{U}} \leftarrow \text{ExtCom}(\tau_{\text{Com}}, C)$ and if $r'_{\mathcal{U}} \neq r''_{\mathcal{U}}$, the challenger aborts.

This game can be distinguished from the previous one only if the binding extractability of \mathcal{C} is contradicted. It follows that adversary $(\mathcal{A}'_1, \mathcal{A}'_2)$ can distinguish this game from the previous with an advantage of at most $\varepsilon_{\mathcal{C}}^{\text{bind-ext}}$.

Game 7. If $id^* \in CA$, to answer a `Send` query on $(id^*, i^*, (pk, \pi))$, the challenger proceeds as the challenger of the previous game, and if it does not aborts, it generates a uniformly random string σ , computes $(pk', sk') \leftarrow \Gamma(x_1, \dots, x_n; \sigma)$ and sets $pk_{id^*}^{i^*} \leftarrow pk'$.

Note that if C is adversarially generated, then $H_\infty(\mathcal{D}_{id^*}^{i^*}) \geq \kappa$ as the distribution of the partner instance is set to a Dirac mass since (id^*, i^*) has not yet accepted, by definition of oracle `Test`. If C is oracle-generated, then $\max(H_\infty(\mathcal{D}_{id^*}^{i^*}), H_\infty(\mathcal{D}_{id^*}^j)) \geq \kappa$. The same arguments for the computational indistinguishability of Game 2 and Game 1 imply that $(\mathcal{A}'_1, \mathcal{A}'_2)$ can distinguish this game from the previous one with an advantage of at most $\varepsilon_{\mathcal{H}}^{\text{uce}} + \varepsilon_{\text{Ext}}$.

In the last game, the condition that if an instance accepts then its partner must eventually terminate implies that $pk_{id^*}^{i^*}$ is computed by generating a uniformly random string σ and evaluating Γ at $(x_1, \dots, x_n; \sigma)$. The advantage of the adversary in the last game is thus nil. It follows that in case 1), the advantage of $(\mathcal{A}'_1, \mathcal{A}'_2)$ is at most

$$|ID| |I| \left(3 \left(\varepsilon_{\mathcal{H}}^{\text{uce}} + \varepsilon_{\text{Ext}} \right) + \varepsilon_\Pi^{\text{ext}} + \varepsilon_{\mathcal{C}}^{\text{setup-ind}} + \varepsilon_{\mathcal{C}}^{\text{bind-ext}} \right).$$

In the second case, $id^ \in U$ and since $r_{C\mathcal{A}}$ is adversarially generated, $\mathcal{D}_{id^*}^{i^*}$ has min-entropy at least κ . The security of \mathcal{H} and of `Ext` ensure that the commitment scheme is hiding and that the proof is zero-knowledge, which ensures that the protocol execution leaks no information about seed s . In addition to that, the high min-entropy of $\mathcal{D}_{id^*}^{i^*}$ also guarantees that s is indistinguishable from a uniformly random string. As a consequence, the resulting key is indistinguishable from one generated by the key-generation algorithm.*

To prove it, consider the following sequence of games.

Game 0. This is the real selective game.

Game 1. In this game, to generate the common-reference string, the challenger runs $(crs, \tau_{zk}) \leftarrow \Pi.\text{Setup}^{zk}(1^\lambda)$ instead of $\Pi.\text{Setup}$. Adversary $(\mathcal{A}'_1, \mathcal{A}'_2)$ can distinguish this game from the previous one with an advantage of at most ε_{Π}^{zk} .

Game 2. This game is defined as in the previous case, and $(\mathcal{A}'_1, \mathcal{A}'_2)$ can distinguish this game from the previous one with an advantage of at most $\varepsilon_{\mathcal{H}}^{\text{uce}} + \varepsilon_{\text{Ext}}$.

Game 3. To answer a prompting Send query on $(id^*, i^*, (*, *))$, the challenger now generates $r'_{\mathcal{U}}$, $\rho_{\mathcal{U}}$ and s' uniformly at random. Since $r_{C\mathcal{A}}$ is adversarially generated in case 2) and that \mathcal{U} accepts only after receiving it, the only information $(\mathcal{A}'_1, \mathcal{A}'_2)$ has about $r_{\mathcal{U}}$ is that it has a distribution $\mathcal{D}_{id^*}^{i^*}$ such that $H_\infty(\mathcal{D}_{id^*}^{i^*}) \geq \kappa \geq \kappa_{\mathcal{H}}$. In the event in which the Test query of \mathcal{A}'_2 is not replied to with \perp , instance (id^*, i^*) is not corrupt before it accepts, so the only information \mathcal{A}'_1 has about $r_{id^*}^{i^*}$ is that its distribution is $\mathcal{D}_{id^*}^{i^*}$. Moreover, if id^* is later corrupt before the end of the protocol execution, (id^*, i^*) will have already erased $r'_{\mathcal{U}}$ and $\rho_{\mathcal{U}}$ and s' .

Consequently, the UCE security of \mathcal{H} can be reduced to distinguishing this game from the previous one, and $(\mathcal{A}'_1, \mathcal{A}'_2)$ can thus distinguish this game from the previous one with an advantage of at most $\varepsilon_{\mathcal{H}}^{\text{uce}}$.

Game 4. In this game, the challenger answers a Send query on $(id^*, i^*, r_{C\mathcal{A}})$, with $r_{C\mathcal{A}}$ adversarially generated, by generating uniformly random values instead of evaluating PRF at $(s', 0)$. If id^* is later corrupt before the end of the protocol execution, (id^*, i^*) will have already erased s' and r_{Π} . Adversary $(\mathcal{A}'_1, \mathcal{A}'_2)$ can distinguish this game from the previous one with an advantage of at most ε_{PRF} .

Game 5. To answer a Send query on $(id^*, i^*, r_{C\mathcal{A}})$ such that $r_{C\mathcal{A}}$ is adversarially generated, the challenger simulates a proof $\pi \leftarrow \Pi.\text{Sim}(crs, \tau_{zk}, ((x_i)_i, k, C, r_{C\mathcal{A}}, pk))$. By the composable zero-knowledge property of Π , this game is perfectly indistinguishable from the previous one.

Game 6. To answer a prompting Send query on $(id^*, i^*, (*, *))$, the challenger runs $(C, d) \leftarrow \text{Com}(0^{|r'_{\mathcal{U}}|}; \rho_{\mathcal{U}})$ and sends C . In the event in which the Test query of \mathcal{A}'_2 is not replied to with \perp , adversary \mathcal{A}'_1 does not corrupt id^* before (id^*, i^*) accepts, so not before $\rho_{\mathcal{U}}$ is erased. As \mathcal{C} is $\varepsilon_{\mathcal{C}}^{\text{hide}}$ -hiding, $(\mathcal{A}'_1, \mathcal{A}'_2)$ can distinguish this game from the previous one with an advantage of at most $\varepsilon_{\mathcal{C}}^{\text{hide}}$.

Game 7. To answer a Send query on $(id^*, i^*, r_{C\mathcal{A}})$ such that $r_{C\mathcal{A}}$ is adversarially generated, the challenger generates s uniformly at random. As $|r'_{\mathcal{U}}| \geq \kappa_{\text{Ext}}$, s is ε^{ext} -computationally indistinguishable from a uniformly random value. $(\mathcal{A}'_1, \mathcal{A}'_2)$ can then distinguish this game from the previous one with an advantage of at most ε_{Ext} .

In the last game, the condition that if an instance accepts then its partner must eventually terminate then implies that $pk_{id^*}^{i^*}$ is computed by evaluating Γ at $(x_1, \dots, x_n; \sigma)$, where σ is a uniformly random string. The advantage of $(\mathcal{A}'_1, \mathcal{A}'_2)$ in that game is then nil. As a result, the advantage of $(\mathcal{A}_1, \mathcal{A}_2)$ in the

second case is at most

$$|ID| |I| \left(\varepsilon_{\Pi}^{zk} + 2\varepsilon_{\mathcal{H}}^{uce} + 2\varepsilon_{\text{Ext}} + \varepsilon_{\text{PRF}} + \varepsilon_{\mathcal{C}}^{\text{hide}} \right).$$

□

Discrete-Logarithm Keys. App. B presents a simple illustration of this generic protocol (but rather in the random-oracle model for better efficiency) applied to discrete-logarithm keys.

4.2 RSA-Key Generation Protocol with Verifiable Randomness for Arbitrary Relations

This section gives a two-party protocol for RSA-key generation with verifiable randomness between a user \mathcal{U} and a certification authority \mathcal{CA} . The resulting keys can be used in any RSA cryptosystem. The protocol attests that the resulting keys were generated with high-entropy randomness and that they satisfy (fixed) arbitrary properties. These properties are captured by a relation

$$\mathcal{R}_W := \{(N, e \in \mathbb{Z}; p, q \in W \subseteq \mathbb{P}) : p \neq q \wedge N = pq \wedge \gcd(e, \varphi(N)) = 1\}$$

to which the keys generated should belong, where W is a set that defines the predicates p and q must satisfy, e.g., $p = q = 3 \pmod{4}$ or p and q are safe primes. Its relative language is denoted \mathcal{R}_W . Efficient proof systems for such properties exist [49,18,4], though none of them aims at proving that the keys were generated with proper randomness.

In comparison, the protocol by Juels and Guajardo [37] only guarantees the first two properties, and does *not* ensure that the user algorithm cannot bias the distribution of the keys. However, it is precisely what lead to the ROCA vulnerabilities [43], showing that the third property is crucial for the security of any public-key cryptosystem. Moreover, without the third property, an interactive key-generation protocol is only beneficial if the user does not have high-entropy randomness locally whereas the CA does, otherwise it is only a burden for the user. On the other hand, the third property additionally guarantees the end user that if the CA has high-entropy randomness, her keys are not faulty.

As for the attestation scheme of Benhamouda et al. [12], it allows to prove that the RSA primes were generated with an arbitrary generator; and the protocols of Camenisch and Michels [18], of Auerbach and Poettering [4], and of Goldberg et al. [32], only allow to prove that RSA primes satisfy certain properties, not that they were generated with high entropy. In a sense, our goal is complementary to that of proving that RSA moduli satisfy certain properties without proving that the keys were generated with high-entropy randomness.

RSA Key-Generation Algorithm. The NIST standard [44] for the RSA [46] key-generation algorithm, further denoted $\text{KeyGen}_{\text{RSA}}$, is the following:

- choose at random two distinct large primes p and q

- compute $N \leftarrow pq$ and $\varphi(N) \leftarrow (p-1)(q-1)$
- choose an integer $2^{16} < e < 2^{256}$ such that $\gcd(e, \varphi(N)) = 1$ (e may be chosen deterministically or at random)
- compute $d \leftarrow e^{-1} \pmod{\varphi(N)}$
- Return $pk \leftarrow (N, e)$ and $sk \leftarrow (N, d)$.

Equivalently, the secret key sk can be set to (p, q, e) instead of (N, d) as one can compute (N, d) from (p, q, e) and vice-versa. It is this variant that is hereafter considered. To formally capture the requirement on p and q to be large, a parameter $b = b(\lambda)$ that specifies the bit-length of p and q is introduced.

Interpretation. There is some ambiguity as to how p and q are generated. The interpretation (which follows how the algorithm would be implemented in practice) of $\text{KeyGen}_{\text{RSA}}$ in the rest of the paper is first that there exists a PPT primality-test algorithm $\text{PrimeTest}_W(\lambda, b, e, p) \rightarrow \zeta \in \{0, 1\}$ (parameter λ is further omitted from its syntax) which tests whether an integer p is in W , b -bit long and such that $\gcd(e, (p-1)) = 1$. Algorithm $\text{KeyGen}_{\text{RSA}}$ then generates, uniformly at random, integers in $\llbracket 2^{b-1}, 2^b - 1 \rrbracket$ until it finds an integer p such that $\text{PrimeTest}_W(b, e, p) = 1$, and continues until it finds a second one $q \neq p$ such that $\text{PrimeTest}_W(b, e, q) = 1$. If no such two integers are found in a specified number of iterations $T_{\text{RSA}}(\lambda)$, the algorithm aborts and returns an invalid pair, e.g., $(0, 0)$. The random variable with values in $\{0, 1, 2\}$ that counts the number of distinct primes found in at most $T_{\text{RSA}}(\lambda)$ iterations is further denoted ctr_{RSA} .

Protocol. We now describe our protocol, further denoted IKG_{RSA} , to generate RSA keys with verifiable randomness. The protocol is given in the random-oracle to allow for practical efficiency.

Building Blocks. The protocol builds on

- the same primality-test algorithm PrimeTest_W as the one run by $\text{KeyGen}_{\text{RSA}}$. It is said to be δ -correct if with probability at most $1 - \delta$, $\text{PrimeTest}_W(b, e, p) = 0$ for $p \in W \cap \llbracket 2^{b-1}, 2^b - 1 \rrbracket$ such that $\gcd(e, (p-1)) = 1$, or $\text{PrimeTest}_W(b, e, p) = 1$ for $p \notin W \cap \llbracket 2^{b-1}, 2^b - 1 \rrbracket$ or such that $\gcd(e, (p-1)) > 1$ (i.e., it is an upper-bound on the probability that it returns a false negative or a false positive)
- a random oracle of which the domain is separated to obtain pairwise independent random oracles \mathcal{H} , \mathcal{H}_C , \mathcal{H}_{Π} and \mathcal{H}_{Π_W}
- a commitment scheme $\mathcal{C} = (\text{Setup}, \text{Com}, \text{ComVf})$ (App. A.1) for the user algorithm to commit to its random string *before* receiving any input from the CA. The parameters returned by Setup are tacit inputs to the other algorithms of \mathcal{C}
- a pseudo-random function PRF with range (non-empty) $R_{\text{PRF}} \subseteq \mathbb{N}$ for \mathcal{U} to generate the RSA primes from the seed extracted with \mathcal{H}
- a non-interactive, extractable zero-knowledge argument system $\Pi_C = (\text{Setup}, \text{Prove}, \text{Verf}, \text{Sim}, \text{Ext})$ for the relation $\{(C; r'_u, d) : \text{ComVf}(C, r'_u, d) = 1\}$ with random oracle \mathcal{H}_C , i.e., for the user to prove knowledge of an opening to her committed randomness

- a non-interactive, extractable zero-knowledge argument system $\Pi = (\text{Setup}, \text{Prove}, \text{Verf}, \text{Sim}, \text{Ext})$ with random oracle \mathcal{H}_Π for the relation

$$\left\{ (C, r_{C\mathcal{A}}, N, (a_\gamma)_{\gamma \neq i, j}; r'_{\mathcal{U}}, d, a_i, a_j) : \text{ComVf}(C, r'_{\mathcal{U}}, d) = 1, s = r'_{\mathcal{U}} \oplus \mathcal{H}(r_{C\mathcal{A}}), \right. \\ \left. \forall \gamma \in [j] a_\gamma = \text{PRF}(s, \gamma), 2^{b(\lambda)-1} \leq a_i, a_j \leq 2^{b(\lambda)} - 1, N = a_i a_j \text{ in } \mathbb{N} \right\},$$

i.e., for the user to prove the RSA primes are really the *first two primes* generated with the seed derived from the committed randomness and the randomness of the CA. This relation is further denoted \mathcal{R}_Π

- a non-interactive zero-knowledge argument system $\Pi_W = (\text{Setup}, \text{Prove}, \text{Verf}, \text{Sim})$ with random oracle \mathcal{H}_{Π_W} for relation \mathcal{R}_W
- another pseudo-random function PRF' with variable output length (encoding in unary as last input) for \mathcal{U} to generate the randomness necessary⁸ to compute $\Pi_C.\text{Prove}$, PrimeTest_W , $\Pi.\text{Prove}$ and $\Pi_W.\text{Prove}$, as the only available randomness to the parties are their input random bit strings.

Throughout the protocol, e is assumed (without loss of generality) to be a fixed⁹, hard-coded value in \mathcal{U} . For the sake of simplicity, e is further assumed to be prime, e.g., $e = 65537$ (it is a value commonly used in practice).

Parameters. Given a security parameter 1^λ and a function $T: \mathbb{N} \rightarrow \mathbb{N}_{>1}$ that gives an upper bound on the number of iterations in Algorithm 2 (and thus the running time of \mathcal{U}), to generate public parameters for IKG_{RSA} , run $pp_{\mathcal{C}} \leftarrow \mathcal{C}.\text{Setup}(1^\lambda)$, $pp_{\Pi_C} \leftarrow \Pi_C.\text{Setup}(1^\lambda)$, $pp_\Pi \leftarrow \Pi.\text{Setup}(1^\lambda)$ and $pp_{\Pi_W} \leftarrow \Pi_W.\text{Setup}(1^\lambda)$. Set and return $pp \leftarrow (b(\lambda), T(\lambda), pp_{\mathcal{C}}, pp_{\Pi_C}, pp_\Pi, pp_{\Pi_W})$.

Formal Description. Consider the interactive protocol on Figure 3 between a user \mathcal{U} and a certification authority \mathcal{CA} . Each algorithm maintains acceptance and termination variables acc_{id} and $term_{id}$ for $id \in \{\mathcal{U}, \mathcal{CA}\}$ initially set to FALSE. The party algorithms proceed as follows:

1. \mathcal{U} applies the random oracle \mathcal{H} twice to its randomness $r_{\mathcal{U}}$ to compute $r'_{\mathcal{U}} \leftarrow \mathcal{H}(0||r_{\mathcal{U}})$ and $\rho_{\mathcal{U}} \leftarrow \mathcal{H}(1||r_{\mathcal{U}})$, commits to $r'_{\mathcal{U}}$ with $\rho_{\mathcal{U}}$ as random string. Next, a seed $s' \leftarrow \mathcal{H}(2||r_{\mathcal{U}})$ from which it derives the randomness necessary to compute $\Pi_C.\text{Prove}$, and computes of proof of knowledge of an opening to the commitment. \mathcal{U} sends the commitment and the proof to \mathcal{CA}
2. \mathcal{CA} , upon receiving the commitment and the proof from \mathcal{U} , sets $acc_{\mathcal{CA}} \leftarrow \text{TRUE}$. It verifies the proof and if it holds, sends its randomness to \mathcal{U} , and otherwise returns \perp and sets $term_{\mathcal{CA}} \leftarrow \text{TRUE}$

⁸ Juels and Guajardo's protocol [37] allows for probabilistic primality-test algorithms and makes uses of proof systems, but does not specify the origin of the randomness necessary for their computation or the zero-knowledge property of the proof systems.

⁹ Alternatively, in the protocol on Figure 3, after N is computed, \mathcal{U} could continue to generate pseudo-random values until it finds one that is coprime with $\varphi(N)$ and then sets it as e . Algorithm \mathcal{U} would then also have to reveal the values that did not satisfy this property and prove that they did not, and also to prove that the chosen e and $\varphi(N)$ are coprime. Assuming e to be fixed in advance avoids this complication.

3. \mathcal{U} , upon receiving $r_{C\mathcal{A}}$ from $C\mathcal{A}$, sets $acc_{\mathcal{U}} \leftarrow \text{TRUE}$. It extracts a seed s with \mathcal{H} from the joint randomness. It continues by generating by running $((a_\gamma)_{\gamma=1}^j, i) \leftarrow \text{Algorithm 2}$.

Algorithm 2

Require: PrimeTest_W , integers T, b, e , pseudo-random function PRF , seed s .

Ensure: Pseudo-random numbers a_γ and integer i .

```

1:  $ctr, i, j \leftarrow 0$ 
2: while  $ctr < 2$  and  $j < T$  do
3:    $j \leftarrow j + 1$ ;  $a_j \leftarrow \text{PRF}(s, j)$ 
4:   if  $\text{PrimeTest}_W(b, e, a_j; \text{PRF}(s, j))$  then
5:     if  $ctr = 0$  then
6:        $i \leftarrow j$ 
7:     end if
8:      $ctr \leftarrow ctr + 1$ 
9:   end if
10: end while
11: if  $ctr < 2$  then
12:   return  $((a_\gamma)_{\gamma=1}^j, \perp)$ 
13: else
14:   return  $((a_\gamma)_{\gamma=1}^j, i)$ 
15: end if

```

- (a) if $i = \perp$ (i.e., Algorithm 2 did not find 2 primes such that $\text{PrimeTest}_W(b, e, a_j; \text{PRF}(s, j)) = 1$ in T iterations; this case is not depicted on Figure 3), \mathcal{U} sends $(r_{\mathcal{U}}, (a_\gamma)_{\gamma=1}^j)$ to $C\mathcal{A}$, returns $(0, 0)$ and sets $term_{\mathcal{U}} \leftarrow \text{TRUE}$
- (b) if $i \neq \perp$, \mathcal{U} computes a proof π that it correctly performed its computation with Π , and a proof π_W that the RSA public key is in \mathcal{L}_W with Π_W . After computing the proofs, \mathcal{U} erases all variables but $N, e, p, q, i, \pi, \pi_W$ and $(a_\gamma)_{\gamma \neq i, j}$. It sends these latter to $C\mathcal{A}$, except p and q , returns $(pk_{\mathcal{U}} \leftarrow (N, e), sk \leftarrow (p, q, e))$, and sets $term_{\mathcal{U}} \leftarrow \text{TRUE}$
- 4a. $C\mathcal{A}$, upon receiving $(r_{\mathcal{U}}, (a_\gamma)_{\gamma=1}^j)$ from \mathcal{U} , computes $r'_{\mathcal{U}}, \rho_{\mathcal{U}}$ and s as \mathcal{U} , computes $(C', d') \leftarrow \text{Com}(r'_{\mathcal{U}}, \rho_{\mathcal{U}})$, and verifies that $C' = C$ and that $\text{PRF}(s, \gamma) = a_\gamma$ for all $\gamma \in \llbracket j \rrbracket$. If all verification succeeds, $C\mathcal{A}$ returns 0, otherwise it returns \perp . It sets $term_{C\mathcal{A}} \leftarrow \text{TRUE}$
- 4b. $C\mathcal{A}$, upon receiving $(N, e, \pi, \pi_W, i, (a_\gamma)_{\gamma \neq i, j})$ from \mathcal{U} , generates a seed s'' with \mathcal{H} from its randomness, and uses it to generate the randomness necessary to compute PrimeTest_W . The resulting random string is denoted r'_W . It verifies that for all $\gamma \in \llbracket j - 1 \rrbracket \setminus \{i\}$, $\text{PrimeTest}_W(b, e, a_\gamma; r'_W) = 0$, and that π and π_W are valid. If one of the verifications did not succeed, $C\mathcal{A}$ returns \perp , otherwise it returns $pk_{C\mathcal{A}} \leftarrow (N, e)$. It sets $term_{C\mathcal{A}} \leftarrow \text{TRUE}$.

Discussion. \mathcal{U} proves among other things that p and q are really the first two random primes in W such that $\gcd(e, p - 1) = \gcd(e, q - 1) = 1$, and therefore

cannot have chosen primes with additional conditions to these. It is a subtle but crucial aspect of the protocol which ensures that \mathcal{U} cannot bias the distribution of the keys; and not doing so is precisely what allowed for the ROCA vulnerabilities [43] in which specific primes were chosen by the user algorithm.

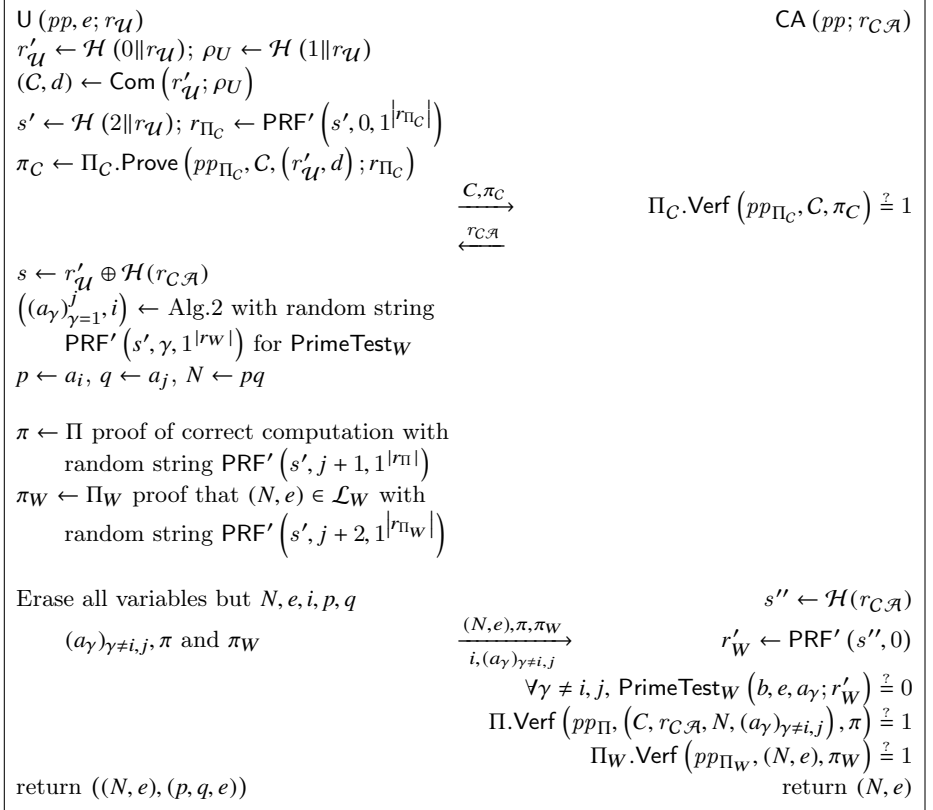


Fig. 3. RSA-Key Generation Protocol with Verifiable Randomness for an Arbitrary Relation $\mathcal{R}_{\mathcal{W}}$.

Theorem 4.3 (Correctness). *Let j be the number of iterations of Algorithm 2, and suppose that $\text{PrimeTest}_{\mathcal{W}}$ is δ -correct and that PRF' is $(T_{\text{PRF}'}, j+3, \epsilon_{\text{PRF}'})$ -secure for $T_{\text{PRF}'} = \Omega(T \cdot T_{\text{PrimeTest}})$. If \mathcal{C} is correct and if Π_C , Π and $\Pi_{\mathcal{W}}$ are complete, protocol IKG_{RSA} is $\max(1 - j(1 - \delta) - 2\epsilon_{\text{PRF}'} - q_{\mathcal{H}}(2^{-\kappa_{\mathcal{U}}} + 2^{-\kappa_{C\mathcal{A}}}), 0)$ -correct in the random-oracle model w.r.t. to the class of algorithms that run in time at most $T_{\text{PRF}'}$, make at most $q_{\mathcal{H}}$ queries to \mathcal{H} and return distributions $\mathcal{D}_{\mathcal{U}}$ and $\mathcal{D}_{C\mathcal{A}}$ independent of the random oracle and with min-entropy at least $2^{-\kappa_{\mathcal{U}}}$ and $2^{-\kappa_{C\mathcal{A}}}$, respectively.*

Proof. Assuming \mathcal{C} to be correct and Π_C , Π and Π_W to be complete, a protocol in which uniformly random values are generated instead of s' and s'' is $\max(1 - j(1 - \delta), 0)$ -correct as an honest protocol execution fails only if PrimeTest does.

Since \mathcal{D}_u and $\mathcal{D}_{C\mathcal{A}}$ are assumed to be independent of the random oracle and to have min-entropy at least with min-entropy at least $2^{-\kappa u}$ and $2^{-\kappa c\mathcal{A}}$, respectively, a protocol in which s and s' are uniformly random is $q(2^{-\kappa u} + 2^{-\kappa c\mathcal{A}})$ -statistically indistinguishable from the previous one.

Consider now the following algorithm which interacts with the PRF challenger for PRF'. Given user and certification-authority distributions, the algorithm runs the key-generation protocol as a subroutine with random strings drawn from the specified distributions, and queries the challenger when PRF' is to be evaluated on s' (which is exactly $j+3$ times). The algorithm returns 1 if the protocol execution fails and 0 otherwise. The security of PRF' implies that a protocol execution with PRF' evaluations on s' is $\max(1 - j(1 - \delta) - \varepsilon_{\text{PRF}'}, q(2^{-\kappa u} + 2^{-\kappa c\mathcal{A}}), 0)$ -correct.

Likewise, consider another algorithm which runs the key-generation protocol as a subroutine with random strings drawn from the specified distributions, and queries the challenger when PRF' is to be evaluated on s'' (which is exactly once). The algorithm returns 1 if the protocol execution fails and 0 otherwise. The security of PRF' implies that a protocol execution with PRF' evaluations on s' and on s'' is $\max(1 - j(1 - \delta) - 2\varepsilon_{\text{PRF}'}, q(2^{-\kappa u} + 2^{-\kappa c\mathcal{A}}), 0)$ -correct. \square

The following theorem shows that IKG_{RSA} is indistinguishable from $\text{KeyGen}_{\text{RSA}}$ in the random oracle if \mathcal{C} is hiding and binding, if Π_C and Π are zero-knowledge and extractable, if Π_W is zero-knowledge and sound, if PRF and PRF' are secure PRFs, if the probability that Algorithm 2 fails although IKG_{RSA} does not is small, and if the adversary makes few random-oracle queries compared to the min-entropy of the test distributions.

Theorem 4.4 (Indistinguishability). *Suppose that \mathcal{C} is $(T_{\mathcal{C}}^{\text{hide}}, \varepsilon_{\mathcal{C}}^{\text{hide}})$ -hiding and $(T_{\mathcal{C}}^{\text{bind}}, \varepsilon_{\mathcal{C}}^{\text{bind}})$ -binding, that PRF is $(T_{\text{PRF}}, j, \varepsilon_{\text{PRF}})$ -secure, that PRF' is $(T_{\text{PRF}'}, j + 3, \varepsilon_{\text{PRF}'})$ -secure, that Π_C is $(T_{\Pi_C}^{\text{ext}}, q_{\mathcal{H}_C}, \varepsilon_{\Pi_C}^{\text{ext}})$ -extractable and $(T_{\Pi_C}^{\text{zk}}, q_{\mathcal{H}_C}, 1, \varepsilon_{\Pi_C}^{\text{zk}})$ -zero-knowledge, that Π is $(T_{\Pi}^{\text{ext}}, q_{\mathcal{H}_{\Pi}}, \varepsilon_{\Pi}^{\text{ext}})$ -extractable and $(T_{\Pi}^{\text{zk}}, q_{\mathcal{H}_{\Pi}}, 1, \varepsilon_{\Pi}^{\text{zk}})$ -zero-knowledge, and that Π_W is $(T_{\Pi_W}^{\text{sound}}, q_{\mathcal{H}_W}, \varepsilon_{\Pi_W}^{\text{sound}})$ -extractable and $(T_{\Pi_W}^{\text{zk}}, q_{\mathcal{H}_W}, 1, \varepsilon_{\Pi_W}^{\text{zk}})$ -zero-knowledge. Protocol IKG_{RSA} is $(T, q_O, \kappa, \varepsilon)$ -indistinguishable from $\text{KeyGen}_{\text{RSA}}$ in the random-oracle model, for $T = \min(T_{\mathcal{C}}^{\text{hide}}, T_{\mathcal{C}}^{\text{bind}}, T_{\text{PRF}}, T_{\text{PRF}'}, T_{\Pi_C}^{\text{ext}}, T_{\Pi_C}^{\text{zk}}, T_{\Pi}^{\text{ext}}, T_{\Pi}^{\text{zk}}, T_{\Pi_W}^{\text{sound}}, T_{\Pi_W}^{\text{zk}})$, $O \in \{\text{Oracle, Dist, Exec, Send, Reveal, Corrupt}\}$, $q_{\text{Oracle}} \geq q_{\mathcal{H}} + q_{\mathcal{H}_C} + q_{\mathcal{H}_{\Pi}} + q_{\mathcal{H}_W}$, and*

$$\varepsilon := |ID| |I| \left(2^{-\kappa} q_{\mathcal{H}} + 5 \left(2^{-\kappa} q_{\mathcal{H}} + \varepsilon_{\text{PRF}} + |R_{\text{PRF}}|^{-1} + \Pr[\text{ctr} < 2, \text{ctr}_{\text{RSA}} = 2] \right) + \varepsilon_{\text{PRF}'} + \varepsilon_{\Pi_C}^{\text{ext}} + \varepsilon_{\Pi_C}^{\text{zk}} + \varepsilon_{\Pi}^{\text{ext}} + \varepsilon_{\Pi}^{\text{zk}} + \varepsilon_{\Pi_W}^{\text{sound}} + \varepsilon_{\Pi_W}^{\text{zk}} + \varepsilon_{\mathcal{C}}^{\text{bind}} + \varepsilon_{\mathcal{C}}^{\text{hide}} \right).$$

Proof. First note that if $(\mathcal{A}_1, \mathcal{A}_2)$ wins the indistinguishability game with a probability at least ε , then there exists an adversary $(\mathcal{A}'_1, \mathcal{A}'_2)$ which wins with

probability at least $\varepsilon/|ID||I|$ a variant of the game in which the adversary is required to specify the pair (id^*, i^*) of its Test query before being given access to the game oracles, i.e., a *selective* variant of the indistinguishability game. Indeed, $(\mathcal{A}'_1, \mathcal{A}'_2)$ can simply run $(\mathcal{A}_1, \mathcal{A}_2)$ as a subroutine and guess the pair (id^*, i^*) at the beginning of the game. If $(\mathcal{A}_1, \mathcal{A}_2)$ later makes its Test query on a different pair, $(\mathcal{A}'_1, \mathcal{A}'_2)$ simply aborts and sends to the challenger a bit chosen uniformly at random.

A message is subsequently said to be *oracle-generated* if it was computed by the challenger as a response to a Send query and it was not altered. Otherwise, the message is said to be *adversarially generated*.

Further distinguish two cases

1. $id^* \in CA$, or $id^* \in U$ and the random string $r_{C\mathcal{A}}$ the i^* th instance of id^* receives is oracle-generated.
2. $id^* \in U$ and the random string $r_{C\mathcal{A}}$ the i^* th instance of id^* receives is adversarially generated.

In the first case, the situation is similar to the first case of the proof of Theorem 4.2, except that the arguments are in the random-oracle model, and that \mathcal{C} is not assumed to be extractable, so the extraction is rather done via Π_C . Consider then the following sequence of games.

Game 0. This is the real selective game.

Game 1. To answer an Exec query on $(id^*, i^*, *, *)$ or on $(*, *, id^*, i^*)$, the challenger generate an honest transcript of the protocol. However, instead of setting $pk_{id^*}^* \leftarrow pk$, it generates a uniformly random string σ , runs $(pk', sk') \leftarrow \text{KeyGen}_{\text{RSA}}(b, e, W; \sigma)$ and sets $pk_{id^*}^* \leftarrow pk$.

Recall that \mathcal{A}'_1 and \mathcal{A}'_2 share no state, and that in the event in which the Test query of \mathcal{A}'_2 is not replied to with \perp , key $pk_{id^*}^*$ is not in $\mathcal{Q}_{\text{Reveal}}$ (nor is the key of the partner of (id^*, i^*)). Denoting by (id, j') the partner instance of (id^*, i^*) , since $\max(H_\infty(\mathcal{D}_{id^*}^{i^*}), H_\infty(\mathcal{D}_{id'}^{j'})) \geq \kappa$, adversary $(\mathcal{A}'_1, \mathcal{A}'_2)$ can distinguish this game from the previous one with an advantage of at most $2^{-\kappa}q_{\mathcal{H}} + \varepsilon_{\text{PRF}} + |R_{\text{PRF}}|^{-1} + \Pr[ctr < 2, ctr_{\text{RSA}} = 2]$.

Indeed, in Algorithm 2, if $ctr < 2$ after T iterations (with ctr depending on $\mathcal{D}_{id^*}^{i^*}$ and $\mathcal{D}_{id'}^{j'}$), then the key pair generated by IKG_{RSA} is $(0, 0)$, i.e., invalid RSA keys, although the key pair generated by $\text{KeyGen}_{\text{RSA}}$ is valid. Moreover, since $s = \mathcal{H}(0\|r_{\mathcal{U}}) \oplus \mathcal{H}(r_{C\mathcal{A}})$, it is $2^{-\kappa}q_{\mathcal{H}}$ -statistically indistinguishable from a uniformly random value if $H_\infty(\mathcal{D}_{id^*}^{i^*}) \geq \kappa$ or $H_\infty(\mathcal{D}_{id'}^{j'}) \geq \kappa$, for $\mathcal{D}_{id^*}^{i^*}$ and $\mathcal{D}_{id'}^{j'}$ are independent of \mathcal{H} .

For a uniformly random seed s , the security of PRF implies that its $j + 3$ evaluations on s are ε_{PRF} -computationally indistinguishable from uniformly random values.

Lastly, if $p = q$, then the key generated is not a valid RSA key, but it only occurs with probability at most $|R_{\text{PRF}}|^{-1}$.

Therefore, $(\mathcal{A}'_1, \mathcal{A}'_2)$ can distinguish this game from the previous one with an advantage of at most $2^{-\kappa}q_{\mathcal{H}} + \varepsilon_{\text{PRF}} + |R_{\text{PRF}}|^{-1} + \Pr[ctr < 2, ctr_{\text{RSA}} = 2]$.

Game 2. If $id^* \in U$, the challenger answers a Send query on $(id^*, i^*, r_{C\mathcal{A}})$ as follows. It computes (pk, π) honestly, but instead of setting $pk_{id^*}^{i^*} \leftarrow pk$, it generates a uniformly random string σ , runs $(pk', sk') \leftarrow \text{KeyGen}_{\text{RSA}}(b, e, W; \sigma)$ and sets $pk_{id^*}^{i^*} \leftarrow pk$.

Recall that in case 1), since $r_{C\mathcal{A}}$ is always oracle-generated. Therefore, the same indistinguishability between the last two games still apply and $(\mathcal{A}'_1, \mathcal{A}'_2)$ can distinguish this game from the previous one with an advantage of at most $2^{-\kappa}q_{\mathcal{H}} + \varepsilon_{\text{PRF}} + |\mathcal{R}_{\text{PRF}}|^{-1} + \Pr[\text{ctr} < 2, \text{ctr}_{\text{RSA}} = 2]$.

Game 3. If $id^* \in CA$, the challenger answers a Send query on $(id^*, i^*, (N, e, \pi, \pi_W, i, (a_\gamma)_{\gamma \neq i, j}))$ as follows. If $(N, e, \pi, \pi_W, i, (a_\gamma)_{\gamma \neq i, j})$ is adversarially generated, for all $\gamma \neq i, j$, $\text{PrimeTest}_W(b, e, a_\gamma; r'_W) = 0$, and proofs π and π_W are valid, then the challenger runs $\Pi.(\text{Ext}_0, \text{Ext}_1)$ on \mathcal{A}'_1 . If extraction fails, the challenger aborts; otherwise it obtains a tuple $(r'_{\mathcal{U}}, d, a_i, a_j)$, and checks whether $(C, r_{C\mathcal{A}}, N, (a_\gamma)_{\gamma \neq i, j}; r'_{\mathcal{U}}, d, a_i, a_j)$ is in \mathcal{R}_Π . If not, it aborts. If so, $N = a_i a_j$ in \mathbb{N} , and the challenger checks whether $(N, e; a_i, a_j)$ of $(N, e; a_j, a_i)$ is in \mathcal{R}_{Π_W} . If not, it aborts. Indeed, Since $W \subseteq \mathbb{P}$, if a_i and a_j are not prime, then the soundness of Π_W is contradicted. If a_i and a_j are prime, then unless the soundness of Π_W is contradicted, the only possible witness for (N, e) is (a_i, a_j) or (a_j, a_i) by the fundamental theorem of arithmetic. Therefore, if neither $(N, e; a_i, a_j)$ nor $(N, e; a_j, a_i)$ is in \mathcal{R}_{Π_W} , the soundness of Π_W is contradicted.

This game can be distinguished from the one only if the extractability of Π or the soundness of Π_W is contradicted. Therefore, $(\mathcal{A}'_1, \mathcal{A}'_2)$ can distinguish this game from the one with an advantage of at most $\varepsilon_{\Pi}^{\text{ext}} + \varepsilon_{\Pi_W}^{\text{sound}}$.

Game 4. If $id^* \in CA$, the challenger answers a Send query on $(id^*, i^*, (N, e, \pi, \pi_W, i, (a_\gamma)_{\gamma \neq i, j}))$ the challenger proceeds as the one of the previous game, but if $(C, r_{C\mathcal{A}}, N, (a_\gamma)_{\gamma \neq i, j}; r'_{\mathcal{U}}, d, a_i, a_j)$ is in the relation, it runs $\Pi_C.(\text{Ext}_0, \text{Ext}_1)$. If extraction fails, the challenger aborts; otherwise it obtains a tuple $(r''_{\mathcal{U}}, d')$. If $r'_{\mathcal{U}} \neq r''_{\mathcal{U}}$, the challenger aborts.

This game can be distinguished from the previous one only if the extractability of Π_C or the binding property of \mathcal{C} is contradicted. It follows that adversary $(\mathcal{A}'_1, \mathcal{A}'_2)$ can distinguish this game from the previous with an advantage of at most $\varepsilon_{\Pi_C}^{\text{ext}} + \varepsilon_{\mathcal{C}}^{\text{bind}}$.

Game 5. If $id^* \in CA$, to answer a Send query on $(id^*, i^*, (N, e, \pi, \pi_W, i, (a_\gamma)_{\gamma \neq i, j}))$, the challenger proceeds as the one of the previous game, and if it does not aborts, it generates a uniformly random string σ , runs $(pk', sk') \leftarrow \text{KeyGen}_{\text{RSA}}(b, e, W; \sigma)$ and sets $pk_{id^*}^{i^*} \leftarrow pk$.

Note that if C is adversarially generated, then $H_\infty(\mathcal{D}_{id^*}^{i^*}) \geq \kappa$ as the distribution of the partner instance is set to a Dirac mass since (id^*, i^*) has not yet accepted, by definition of oracle Test. If C is oracle-generated, then $\max(H_\infty(\mathcal{D}_{id^*}^{i^*}), H_\infty(\mathcal{D}_{id^*}^j)) \geq \kappa$. The same arguments for the computational indistinguishability of Game 2 and Game 1 imply that $(\mathcal{A}'_1, \mathcal{A}'_2)$ can distinguish this game from the previous one with an advantage of at most $2^{-\kappa}q_{\mathcal{H}} + \varepsilon_{\text{PRF}} + |\mathcal{R}_{\text{PRF}}|^{-1} + \Pr[\text{ctr} < 2, \text{ctr}_{\text{RSA}} = 2]$.

In the last game, the condition that if an instance accepts then its partner must eventually terminate implies that $pk_{id^*}^{i^*}$ is computed by generating a uniformly random string σ and running $\text{KeyGen}_{\text{RSA}}$ on input $(b, e, W; \sigma)$. The advantage of the adversary in the last game is thus nil. It follows that in case 1), the advantage of $(\mathcal{A}_1, \mathcal{A}_2)$ is at most

$$|ID| |I| \left(3 \left(2^{-\kappa} q_{\mathcal{H}} + \varepsilon_{\text{PRF}} + |\mathcal{R}_{\text{PRF}}|^{-1} + \Pr[\text{ctr} < 2, \text{ctr}_{\text{RSA}} = 2] \right) + \varepsilon_{\Pi}^{\text{ext}} + \varepsilon_{\Pi_W}^{\text{sound}} + \varepsilon_{\Pi_C}^{\text{ext}} + \varepsilon_{\mathcal{C}}^{\text{bind}} \right).$$

In the second case, the situation is similar to the second case of the proof of Theorem 4.2, except that the arguments are in the random-oracle model, and that Π_W is only assumed to be sound, not extractable. Consider then the following sequence of games.

Game 0. This is the real selective game.

Game 1. This game is defined as in the previous case, and $(\mathcal{A}'_1, \mathcal{A}'_2)$ can distinguish this game from the previous one with an advantage of at most $2^{-\kappa} q_{\mathcal{H}} + \varepsilon_{\text{PRF}} + |\mathcal{R}_{\text{PRF}}|^{-1} + \Pr[\text{ctr} < 2, \text{ctr}_{\text{RSA}} = 2]$.

Game 2. To answer a prompting Send query on $(id^*, i^*, (*, *))$, the challenger now generates $r'_{\mathcal{U}}, \rho_{\mathcal{U}}$ and s' uniformly at random. Since $r_{C\mathcal{A}}$ is adversarially generated in case 2) and that \mathcal{U} accepts only after receiving it, the only information $(\mathcal{A}'_1, \mathcal{A}'_2)$ has about $r_{\mathcal{U}}$ is that it has a distribution $\mathcal{D}_{id^*}^{i^*}$ such that $H_{\infty}(\mathcal{D}_{id^*}^{i^*}) \geq \kappa$. In the event in which the Test query of \mathcal{A}'_2 is not replied to with \perp , instance (id^*, i^*) is not corrupt before it accepts, so the only information \mathcal{A}'_1 has about $r_{id^*}^{i^*}$ is that its distribution is $\mathcal{D}_{id^*}^{i^*}$. Moreover, if id^* is later corrupt before the end of the protocol execution, (id^*, i^*) will have already erased $r'_{\mathcal{U}}$ and $\rho_{\mathcal{U}}$ and s' .

Consequently, $(\mathcal{A}'_1, \mathcal{A}'_2)$ can thus distinguish this game from the previous one with an advantage of at most $2^{-\kappa} q_{\mathcal{H}}$.

Game 3. In this game, the challenger answers a Send query on $(id^*, i^*, r_{C\mathcal{A}})$, with $r_{C\mathcal{A}}$ adversarially generated, if $\text{ctr} < 2$, the challenger aborts. $(\mathcal{A}'_1, \mathcal{A}'_2)$ can thus distinguish this game from the previous one with an advantage of at most $\Pr[\text{ctr} < 2, \text{ctr}_{\text{RSA}} = 2]$.

Game 4. In this game, the challenger answers a Send query on $(id^*, i^*, r_{C\mathcal{A}})$, with $r_{C\mathcal{A}}$ adversarially generated, by generating uniformly random values instead of evaluating PRF' at $(s', 0)$. If id^* is later corrupt before the end of the protocol execution, (id^*, i^*) will have already erased s' and r_{Π} . Adversary $(\mathcal{A}'_1, \mathcal{A}'_2)$ can distinguish this game from the previous one with an advantage of at most $\varepsilon_{\text{PRF}'}$.

Game 5. Denoting by $Q_{\mathcal{H}_{\Pi}}$ the of queries to \mathcal{H} and their responses, and the challenger now runs $\Pi.\text{Sim}(0, Q_{\mathcal{H}_{\Pi}}, \cdot)$ to answer random-oracle queries. To answer a Send query on $(id^*, i^*, r_{C\mathcal{A}})$ such that $r_{C\mathcal{A}}$ is adversarially generated, the challenger simulates a proof $\pi \leftarrow \Pi.\text{Sim}(1, Q_{\mathcal{H}_{\Pi}}, (C, r_{C\mathcal{A}}, N, (a_{\gamma})_{\gamma \neq i, j}; r'_{\mathcal{U}}, d, a_i, a_j))$. Likewise, the challenger simulates a proof π_W .

In the event in which the Test query of \mathcal{A}'_2 is not replied to with \perp , adversary \mathcal{A}'_1 does not corrupt id^* before (id^*, i^*) accepts, so not before r_Π and r_{Π_W} are erased. By the zero-knowledge property of Π and Π_W , adversary $(\mathcal{A}'_1, \mathcal{A}'_2)$ can distinguish this game from the previous one with an advantage of at most $\varepsilon_{\Pi}^{\text{zk}} + \varepsilon_{\Pi_W}^{\text{zk}}$.

Game 6. To answer a prompting Send query on $(id^*, i^*, (*, *))$, the challenger simulates a proof Π_C . In the event in which the Test query of \mathcal{A}'_2 is not replied to with \perp , adversary \mathcal{A}'_1 does not corrupt id^* before (id^*, i^*) accepts, so not before r_{Π_C} is erased.

By the zero-knowledge property of Π_C adversary $(\mathcal{A}'_1, \mathcal{A}'_2)$ can distinguish this game from the previous one with an advantage of at most $\varepsilon_{\Pi_C}^{\text{zk}}$.

Game 7. To answer a prompting Send query on $(id^*, i^*, (*, *))$, the challenger runs $(C, d) \leftarrow \text{Com}(0^{|r_{\mathcal{U}}|}; \rho_{\mathcal{U}})$ and sends C . In the event in which the Test query of \mathcal{A}'_2 is not replied to with \perp , adversary \mathcal{A}'_1 does not corrupt id^* before (id^*, i^*) accepts, so not before $\rho_{\mathcal{U}}$ is erased. As \mathcal{C} is $\varepsilon_{\mathcal{C}}^{\text{hide}}$ -hiding, $(\mathcal{A}'_1, \mathcal{A}'_2)$ can distinguish this game from the previous one with an advantage of at most $\varepsilon_{\mathcal{C}}^{\text{hide}}$.

Game 8. To answer a Send query on $(id^*, i^*, r_{C\mathcal{A}})$ such that $r_{C\mathcal{A}}$ is adversarially generated, the challenger generates s uniformly at random. As $H_\infty(\mathcal{D}_{id^*}^i) \geq \kappa$, seed s is $2^{-\kappa} q_{\mathcal{H}}$ -statistically indistinguishable from a uniformly random value. $(\mathcal{A}'_1, \mathcal{A}'_2)$ can then distinguish this game from the previous one with an advantage of at most $2^{-\kappa} q_{\mathcal{H}}$.

Game 9. To answer a Send query on $(id^*, i^*, r_{C\mathcal{A}})$ such that $r_{C\mathcal{A}}$ is adversarially generated, the challenger generates uniformly random numbers instead of evaluating PRF on s . Algorithm $(\mathcal{A}'_1, \mathcal{A}'_2)$ can distinguish this game from the previous one with an advantage of at most ε_{PRF} .

Game 10. To answer a Send query on $(id^*, i^*, r_{C\mathcal{A}})$ such that $r_{C\mathcal{A}}$ is adversarially generated, the challenger aborts if $p = q$. Algorithm $(\mathcal{A}'_1, \mathcal{A}'_2)$ can distinguish this game from the previous one with an advantage of at most $|R_{\text{PRF}}|^{-1}$.

In the last game, the condition that if an instance accepts then its partner must eventually terminate then implies that $pk_{id^*}^i$ is computed by running $\text{KeyGen}_{\text{RSA}}(b, e, W; \sigma)$, where σ is a uniformly random string. The advantage of $(\mathcal{A}'_1, \mathcal{A}'_2)$ in that game is then nil. As a result, the advantage of $(\mathcal{A}_1, \mathcal{A}_2)$ in the second case is at most

$$|ID| |I| \left(3q_{\mathcal{H}} 2^{-\kappa} + 2\varepsilon_{\text{PRF}} + \varepsilon_{\text{PRF}'} + \varepsilon_{\Pi}^{\text{zk}} + \varepsilon_{\Pi_W}^{\text{zk}} + \varepsilon_{\Pi_C}^{\text{zk}} + \varepsilon_{\mathcal{C}}^{\text{hide}} + 2|R_{\text{PRF}}|^{-1} + 2\Pr[\text{ctr} < 2, \text{ctr}_{\text{RSA}} = 2] \right).$$

□

Theorems 4.5 and 4.7 show that by tuning the running time of Algorithm 2 depending on the number of primes in the range of PRF that satisfy the condi-

tions on p and q , the probability that Algorithm 2 fails although IKG_{RSA} does not is small.

Theorem 4.5 (Running Time of IKG_{RSA}). *Let $n(b, W, e, R_{\text{PRF}})$ denote the number of primes p in $\llbracket 2^{b-1}, 2^b - 1 \rrbracket \cap W \cap R_{\text{PRF}}$ such that $\gcd(e, p-1) = 1$, and let β denote $n(b, W, e, R_{\text{PRF}}) |R_{\text{PRF}}|^{-1}$. If $n(b, W, e, R_{\text{PRF}}) \geq 1$, then for numbers chosen uniformly at random in R_{PRF} , the number of trials necessary to obtain two primes that satisfy the above conditions is $2/\beta$ in expectation. Moreover, for any $\delta > 0$, setting $J \leftarrow \lceil 2(1 + \delta)/\beta \rceil$, if PRF is $(T_{\text{PRF}}, J, \varepsilon_{\text{PRF}})$ -secure and $\mathcal{D}_{\mathcal{U}}$ and $\mathcal{D}_{\mathcal{CA}}$ are independent of the random oracle, have respective min-entropy at least $2^{-\kappa_{\mathcal{U}}}$ and $2^{-\kappa_{\mathcal{CA}}}$, and are returned by an algorithm that runs in time at most T_{PRF} and makes at most $q_{\mathcal{H}}$ queries to \mathcal{H} , then, in Algorithm 2, $\Pr[j > J] \leq \exp(- (2\delta/\beta J)^2 \beta J/2) + q_{\mathcal{H}} (2^{-\kappa_{\mathcal{U}}} + 2^{\kappa_{\mathcal{CA}}}) + \varepsilon_{\text{PRF}}$.*

Proof. Let $(X_{\mu})_{\mu \geq 1}$ be a family of independent random variables with the same binomial distribution of parameter β . In essence, X_{μ} indicates whether the μ th trial, if it were done with a uniformly random function, would result in a prime number. For an integer $\nu \geq 1$, let $T_{\nu} := \min\{\mu \geq 1 : X_1 + \dots + X_{\mu} = \nu\}$. It is a stopping time which indicates the number of trials before getting ν primes numbers. The variable $T_{\nu} - \nu$, which indicates the number of non-primes before ν primes are found, has a negative binomial distribution with parameters ν and $1 - \beta$. Its expectation is then $\nu(1 - \beta)/\beta$. Therefore, by $\mathbb{E}[T_{\nu}] = \nu/\beta$. In the case of RSA keys, $\nu = 2$.

Moreover, the Chernoff bound (App. A.6) implies that the probability that the number of trials is higher than this expectation decreases exponentially fast. Formally, consider a real number $\delta > 0$, and set $J \leftarrow \lceil 2(1 + \delta)/\beta \rceil$. Note that $T_2 > J$ if and only if $X_1 + \dots + X_J < 2$, which is equivalent to $X_1 + \dots + X_J - J\beta < 2 - J\beta \leq -2\delta$. Note also that $0 \leq 2\delta/\beta J < 1$ for all $\delta > 0$. The Chernoff bound implies that $\Pr(T_2 > J) \leq P(X_1 + \dots + X_J - J\beta < -2\delta) \leq \exp(- (2\delta/\beta J)^2 \beta J/2)$.

In Algorithm 2, integers are not generated uniformly at random, but rather with PRF. However, for any $\delta > 0$, setting $J \leftarrow \lceil 2(1 + \delta)/\beta \rceil$, if PRF is $(T_{\text{PRF}}, J, \varepsilon_{\text{PRF}})$ -secure and $\mathcal{D}_{\mathcal{U}}$ and $\mathcal{D}_{\mathcal{CA}}$ are independent of the random oracle, have respective min-entropy at least $2^{-\kappa_{\mathcal{U}}}$ and $2^{-\kappa_{\mathcal{CA}}}$, and are returned by an algorithm that runs in time at most T_{PRF} and makes at most $q_{\mathcal{H}}$ queries to \mathcal{H} , then, the inequality of the theorem statement holds.

To show it, first notice that an execution of protocol IKG_{RSA} is $q_{\mathcal{H}} (2^{\kappa_{\mathcal{U}}} + 2^{\kappa_{\mathcal{CA}}})$ -statistically indistinguishable from one in which s is generated uniformly at random. Consider now an adversary for the PRF game with the DY PRF which makes oracle queries until it either obtains two primes (possibly at the J th query) or reaches J query with at most one prime returned by the oracle. The adversary returns 1 in the first case and 0 in the second. The number of trials with PRF is then at most the number of trials with a uniformly random function plus the advantage of this adversary in the PRF game, and the theorem follows. \square

Corollary 4.6. *In Algorithm 2, if $T > 2/\beta$, setting $\delta \leftarrow \beta T/2 - 1$ and $J \leftarrow \lceil 2(1 + \delta)/\beta \rceil$, $\Pr[\text{ctr} < 2] \leq \exp(- (2\delta/\beta J)^2 \beta J/2) + q_{\mathcal{H}} (2^{-\kappa_{\mathcal{U}}} + 2^{\kappa_{\mathcal{CA}}}) + \varepsilon_{\text{PRF}}$.*

Theorem 4.7 (Running Time of $\text{KeyGen}_{\text{RSA}}$). Let $n(b, W, e)$ denote the number of primes p in $\llbracket 2^{b-1}, 2^b - 1 \rrbracket \cap W$ such that $\gcd(e, p-1) = 1$, and let β_{RSA} denote $n(b, W, e)2^{-b+1}$. If $n(b, W, e) \geq 1$, then the expected running time of $\text{KeyGen}_{\text{RSA}}$ is $2/\beta_{\text{RSA}}$. Moreover, for $T_{\text{RSA}} > 2/\beta_{\text{RSA}}$, setting $\delta_{\text{RSA}} \leftarrow \beta_{\text{RSA}}T_{\text{RSA}}/2 - 1$ and $J_{\text{RSA}} \leftarrow \lceil 2(1 + \delta_{\text{RSA}})/\beta_{\text{RSA}} \rceil$,

$$\Pr[ctr_{\text{RSA}} = 2] \geq 1 - \exp\left(-\left(\frac{2\delta_{\text{RSA}}}{\beta_{\text{RSA}}J_{\text{RSA}}}\right)^2 \frac{\beta_{\text{RSA}}J_{\text{RSA}}}{2}\right).$$

Proof. It follows by the exact same analysis as in the first part of the proof of Theorem 4.5. \square

As ctr and ctr_{RSA} are independent random variables (the randomness of by $\text{KeyGen}_{\text{RSA}}$ in query Test is independent of the distributions given by \mathcal{A}'_1), $\Pr[ctr < 2, ctr_{\text{RSA}} = 2]$ is upper-bounded by the product of the upper-bounds of Corollary 4.6 and Theorem 4.7, and this upper bound mainly (but not only) depends on T , T_{RSA} and R_{PRF} .

5 Instantiation of the RSA-Key Generation Protocol

In this section, we instantiate the protocol of Section 4.2 for RSA key-generation with verifiable randomness. To do so, we provide efficient instantiations for each of the building blocks.

Recently, several important advancements have been made on the efficiency of the commit-and-prove paradigm on committed values which combine algebraic and non-algebraic statements [21,16,5]. These improvements for *cross-domains* statements allow to prove efficiently for instance that some committed value corresponds to a pre-image of some value of a given hash function such as SHA-256 or that some value is the output of some non-algebraic PRF (i.e. HMAC-SHA-256 or AES) using some committed key. To generate an RSA modulus of 3072 bits (for 128-bit security) using the generic protocol from Section 4.2, the PRF must return 1536-bit integers and the use of non-algebraic PRF with the technique from [21,16,5] would result in prohibitive schemes.

On this account, we present an instantiation based on an algebraic PRF, namely the Dodis–Yampolskiy PRF, and use techniques [16] due to Bünz, Bootle, Boneh, Poelstra, Wuille and Maxwell for range proofs and arithmetic-circuit satisfiability to obtain short proofs of correct computation (i.e., Π in Section 4.2).

In the process, we give the first logarithmic-size (in the bit-length of the group order) argument of knowledge of double discrete logarithms, and argument of equality of a discrete logarithm in a group and a double discrete logarithm in another related group. In contrast, the protocol of Camenisch and Stadler [19] for the first relation, and the protocol of Chase et al. [21] for the second are linear in the security parameter.

Parameters. We consider two related group-family generators GroupGen_1 and GroupGen_2 . Given a security parameter λ , to generate an RSA modulus which

is the product of two $b(\lambda)$ -bit prime numbers, let ℓ be the smallest prime of binary length equal to $b(\lambda)$ such that $2\ell + 1$ is also a prime number (i.e. ℓ is a Sophie Germain prime number, or equivalently $2\ell + 1$ is a $b(\lambda) + 1$ -bit *safe prime*). GroupGen_2 returns, on input λ , the group \mathbb{G}_2 of quadratic residues modulo $2\ell + 1$ (which is of prime order ℓ). The group-family generator GroupGen_1 returns on input λ some group \mathbb{G}_1 of prime order Λ such that ℓ divides $\Lambda - 1$ and $\Lambda > (2\ell + 1)^2$. In practice¹⁰, \mathbb{G}_1 can be taken as a prime order subgroup Λ of \mathbb{Z}_r^* for some prime number r such that Λ divides $r - 1$.

The restriction to quadratic residues is necessary for assumptions like the DDH and the q -DDHI assumptions to hold over \mathbb{G}_2 . However, it introduces a bias by design (not from the user algorithm) in the RSA keys generated: p and q are necessarily quadratic residues modulo $2\ell + 1$. The reason is that the values returned by the DY PRF are not actually integers but \mathbb{G}_2 elements. Nonetheless, it is already the case for 1/4 of all RSA moduli since the factors p and q returned by $\text{KeyGen}_{\text{RSA}}$.

Commitment Scheme. Scheme \mathcal{C} is the Pedersen commitment scheme [45] in \mathbb{G}_2 for the user to commit her randomness.

used to commit to the secret RSA primes p and q , and the same Pedersen scheme.

Pseudo-Random Functions. PRF is the Dodis–Yampolskiy (DY) PRF (see Section 2) in the group $\mathbb{G}_2 = QR_{2\ell+1}$ of quadratic residues modulo $2\ell + 1$. It is used to generate the secret RSA primes p and q . Since $2\ell + 1$ is $b(\lambda) + 1$ bits long, p and q are $b(\lambda)$ bits long with probability close to 1/2. The reason $2\ell + 1$ is chosen to be one bit larger than p and q is to ensure that *all* primes of $b(\lambda)$ bits can be returned by the PRF so as not to introduce a bias. As for PRF', it can be any efficient pseudo-random function, e.g., HMAC [7].

Argument for R_W . The argument system Π_W depends on the properties that the prime factors of N must satisfy, e.g., they must be congruent to 3 modulo 4 or be safe primes. To prove that $p = q = 3 \pmod{4}$, one can prove that N is of the form $p^r q^s$ with $p = q = 3 \pmod{4}$ using the protocol of van de Graaf and Peralta [49], and run in parallel the protocol of Boyar et al. [15] to prove that N is square-free. To prove that p and q are safe primes, there exist proof systems in the literature such as Camenisch and Michel's [18]. Besides, Goldberg et al. [32] recently built a protocol to prove that $\text{gcd}(e, \phi(N)) = 1$.

Argument of Correct Computation. The last component is an extractable zero-knowledge argument system Π in the random-oracle model for the user algorithm to prove that it correctly performed its computation, i.e., an argument system for \mathcal{R}_Π . Section 5.1 presents a perfectly honest-verifier zero-knowledge interactive protocol for \mathcal{R}_Π that also satisfies witness-extended emulation. Lemma A.1 implies that it is extractable in the random-oracle model.

¹⁰ To generate RSA moduli which are products of two 1536-bit primes, the instantiation with the Dodis–Yampolskiy PRF uses $\ell = 2^{1535} + 554415$ which is a Sophie Germain prime, $\Lambda = (4\ell + 18)\ell + 1$ and $r = 1572 \cdot \Lambda + 1$.

5.1 Zero-Knowledge Argument with the Dodis–Yampolskiy PRF

This section gives a zero-knowledge argument Π in the case of the DY PRF in $\mathbb{G}_2 = QR_{2\ell+1}$.

Formally, let $2\ell + 1$ be a $b(\lambda) + 1$ -bit (i.e., $b(\lambda) + 1 = \lfloor \log(2\ell + 1) \rfloor + 1$) safe prime (i.e., ℓ is a Sophie Germain prime) and let Λ be a prime integer such that ℓ divides $\Lambda - 1$ and $\Lambda > (2\ell + 1)^2$. Consider $\mathbb{G}_1 = \langle G_1 \rangle$ a group of prime order Λ (in which p and q will be committed) and $\mathbb{G}_2 = \langle G_2 \rangle = QR_{2\ell+1}$ the group of quadratic residues modulo $2\ell + 1$, which is a cyclic group of order ℓ . Recall that the DY PRF is defined as the map $(K, x) \mapsto G_2^{1/(K+x)}$.

Proof Strategy. To prove knowledge of a witness for the membership of $(C, r_{C\mathcal{A}}, N, (a_\gamma)_{\gamma \neq i, j})$ to the language relative to \mathcal{R}_Π , the user algorithm commits to $p = a_i$ and $q = a_j$ in \mathbb{G}_1 with the Pedersen commitment scheme and respective randomness r_p and r_q . The commitments are denoted P and Q .

The user algorithm then proves knowledge of a witness for $\mathcal{R}_0 \cap \mathcal{R}_1$, with

$$\begin{aligned} \mathcal{R}_0 := & \left\{ (C, r_{C\mathcal{A}}, N, P, Q, (a_\gamma)_{\gamma \neq i, j}; r'_u, \rho_u, a_i, a_j, r_p, r_q) : \right. \\ & \text{ComVf}(C, r'_u, \rho_u) = 1, s = r'_u + \mathcal{H}(r_{C\mathcal{A}}) \bmod \ell \\ & \left. \forall \gamma \in \llbracket j \rrbracket, a_\gamma = \text{PRF}(s, \gamma), \text{ComVf}(P, a_i, r_p) = \text{ComVf}(Q, a_j, r_q) = 1 \right\} \end{aligned}$$

and

$$\begin{aligned} \mathcal{R}_1 := & \left\{ (C, r_{C\mathcal{A}}, N, P, Q, (a_\gamma)_{\gamma \neq i, j}; r'_u, \rho_u, a_i, a_j, r_p, r_q) : \text{ComVf}(P, a_i, r_p) = 1 \right. \\ & \left. \text{ComVf}(Q, a_j, r_q) = 1, 2^{b(\lambda)-1} \leq a_i, a_j \leq 2^{b(\lambda)} - 1, N = a_i a_j \text{ in } \mathbb{N} \right\}. \end{aligned}$$

To prove knowledge of a witness for relation \mathcal{R} , it then suffices to prove in parallel knowledge of a witness for \mathcal{R}_0 and of a witness for \mathcal{R}_1 on the same public inputs. Note that the binding property of the Pedersen commitment scheme in \mathbb{G}_1 (relying on the DLOG assumption) guarantees that the a_i and a_j values used in both proofs are the same (up to a relabeling).

Relation \mathcal{R}_0 . We start by giving two preliminary protocols:

- a logarithmic-size zero-knowledge argument of knowledge of a double-discrete logarithm (Section 5.2) using Bulletproof techniques [16]. The resulting proofs are of size logarithmic in the bit-length of the group order. In comparison, the protocol of Camenisch and Stadler [19] has proofs of size linear in the security parameter
- a logarithmic-size argument of equality of a discrete logarithm in a group and a double discrete logarithm in another related group (Section 5.3). In contrast, the protocol of Chase et al. [21, Section 4.3] for this relation uses the techniques of Camenisch and Stadler and therefore has proofs of size linear in the security parameter.

We then combine the latter proof with the proof in Section 5.4 to obtain a proof for relation \mathcal{R}_0 .

Relation \mathcal{R}_1 . The aggregated logarithmic range proof of Bünz et al. [16, Section 4.2] is sufficient to prove that the values committed in P and Q modulo Λ are in $\llbracket 2^{b-1}, 2^b - 1 \rrbracket$ (which is equivalent to proving that the values committed in $PG_1^{-2^{b-1}}$ and $QG_1^{-2^{b-1}}$ are in $\{0, \dots, 2^{b-1} - 1\}$). With the hypotheses on the parameters Λ and ℓ , the verifier is convinced that the equation $N = a_i a_j$ holds in \mathbb{N} . Indeed, the equation $N = a_i a_j \pmod{\Lambda}$ implies that there exists $m \in \mathbb{Z}$ such that $N = a_i a_j + m\Lambda$. Integer m cannot be strictly positive as otherwise N would be strictly greater than Λ . Besides, m cannot be strictly negative since $\Lambda > (2\ell + 1)^2 > a_i a_j$; it is therefore nil and the equation $N = a_i a_j$ holds in \mathbb{N} .

5.2 Logarithmic-Size Argument of Double Discrete Logarithm

This section gives a zero-knowledge argument with logarithmic communication size for proving knowledge of a double discrete logarithm. It uses as a sub-argument the logarithmic-size inner-product argument for arithmetic-circuit satisfiability of Bünz et al. [16, Section 5.2], which is complete, perfectly honest-verifier zero-knowledge, and satisfies witness-extended emulation (App. A.2).

Following the ideas of Bootle et al. [14], Bünz et al. convert any arithmetic circuit with n multiplications gates into a Hadamard product $\mathbf{a}_L \circ \mathbf{a}_R = \mathbf{a}_O$ and $Q \leq 2n$ linear constraints of the form

$$\langle \mathbf{w}_{L,q}, \mathbf{a}_L \rangle + \langle \mathbf{w}_{R,q}, \mathbf{a}_R \rangle + \langle \mathbf{w}_{O,q}, \mathbf{a}_O \rangle = c_q$$

for $q \in \llbracket Q \rrbracket$, with $\mathbf{w}_{L,q}, \mathbf{w}_{R,q}, \mathbf{w}_{O,q} \in \mathbb{Z}_p^n$ and $c_q \in \mathbb{Z}_p$. The vectors $\mathbf{a}_L, \mathbf{a}_R$ respectively denote the vectors of left and right inputs to the multiplications gates, and \mathbf{a}_O the vector of outputs. The linear constraints ensure the consistency between the outputs and the inputs of two consecutive depth level of the circuit. Bootle et al. [14, App. A] give a general method to find such linear constraints, though it may not always result in the most compact ones for a specific circuit.

Bünz et al. actually give an argument for a more general relation which includes Pedersen commitments of which the openings are included in the linear consistency constraints. Concretely, given a group \mathbb{G} of prime order p and positive integers n, m and Q , Bünz et al. give a zero-knowledge argument for the relation

$$\begin{aligned} & \{(g, h \in \mathbb{G}, \mathbf{g}, \mathbf{h} \in \mathbb{G}^n, \mathbf{V} \in \mathbb{G}^m, \mathbf{W}_L, \mathbf{W}_R, \mathbf{W}_O \in \mathbb{Z}_p^{Q \times n}, \mathbf{W}_V \in \mathbb{Z}_p^{Q \times m}, \\ & \mathbf{c} \in \mathbb{Z}_p^Q; \mathbf{a}_L, \mathbf{a}_R, \mathbf{a}_O \in \mathbb{Z}_p^n, \mathbf{v}, \gamma \in \mathbb{Z}_p^m) : V_j = g^{v_j} h^{\gamma_j} \forall j \in \llbracket m \rrbracket \\ & \wedge \mathbf{a}_L \circ \mathbf{a}_R = \mathbf{a}_O \wedge \mathbf{W}_L \mathbf{a}_L + \mathbf{W}_R \mathbf{a}_R + \mathbf{W}_O \mathbf{a}_O = \mathbf{W}_V \mathbf{v} + \mathbf{c}\}. \end{aligned}$$

The soundness of their argument relies on the discrete logarithm assumption over the generator of \mathbb{G} . Concerning the proof size, the prover sends $2\lceil \log_2 n \rceil + 8$ group elements and $5 \mathbb{Z}_p$ elements.

The main difficulty in the case of a proof of a double discrete logarithm relation is to re-write the problem in a way that is suitable to apply the proof for arithmetic circuits. Formally, the goal is to give a zero-knowledge argument for the relation

$$\mathcal{R}_{2\text{DLOG}} := \left\{ (G_1, H_1, G_2, Y; x \in \mathbb{Z}_\ell, r \in \mathbb{Z}_\Lambda) : Y = G_1^{G_2^x} H_1^r \right\}.$$

First, let $n(\lambda) + 1 := b(\lambda)$ be the bit-length of ℓ . Given the bit representation $(x_i)_{i=0}^n$ of x , $G_2^x = G_2^{\sum_{i=0}^n x_i 2^i} = \prod_i (G_2^{2^i})^{x_i}$. An important observation is that for $x_i \in \{0, 1\}$, $(G_2^{2^i})^{x_i} = x_i G_2^{2^i} + (1 - x_i) = x_i (G_2^{2^i} - 1) + 1$. The addition here is over \mathbb{Z}_Λ , although the notation is purely formal since $x_i \in \{0, 1\}$. It thus follows that an argument for $\mathcal{R}_{2\text{DLOG}}$ is equivalent to an argument for

$$\left\{ (G_1, H_1, G_2, Y; (x_i)_{i=0}^n \in \{0, 1\}^n, r \in \mathbb{Z}_\Lambda) : Y = G_1^{\prod_i (x_i (G_2^{2^i} - 1) + 1)} H_1^r \right\},$$

which is also equivalent to an argument for

$$\left\{ (G_1, H_1, G_2, Y; (a_i)_{i=0}^n, r \in \mathbb{Z}_\Lambda) : Y = G_1^{\prod_i a_i} H_1^r \wedge a_i \in \{1, G_2^{2^i}\} \right\}.$$

To this end, consider the following array

$$\begin{array}{cccccc} a_0 & a_1 & a_2 & \cdots & a_n & \\ 1 & a_0 & a_0 a_1 & \cdots & a_0 a_1 \cdots a_{n-1} & \\ a_0 & a_0 a_1 & a_0 a_1 a_2 & \cdots & a_1 \cdots a_n. & \end{array}$$

Notice that its third row is the product of the first two. In other words, if $\mathbf{a} \leftarrow (a_0, a_1, \dots, a_n) \in \mathbb{Z}_\Lambda^{n+1}$ and $\mathbf{b} \leftarrow (b_0, b_1, \dots, b_{n-1}) \in \mathbb{Z}_\Lambda^n$, then $\mathbf{a} \circ (1 \ \mathbf{b}) = (\mathbf{b} \ y)$ for $y := G_2^x$.

Moreover, for $\mathbf{a}_L := [\mathbf{a} \ \mathbf{a} - \mathbf{1}^{n+1}]^T$, $\mathbf{a}_R := [1 \ \mathbf{b} \ \mathbf{a} - \mathbf{G}_2^{2^{n+1}}]^T$ and $\mathbf{a}_O := [\mathbf{b} \ y \ \mathbf{0}^{n+1}]^T \in \mathbb{Z}_\Lambda^{2(n+1)}$ where $\mathbf{G}_2^{2^{n+1}}$ denotes the vector $(G_2, G_2^2, G_2^{2^2}, \dots, G_2^{2^n})$, one has $\mathbf{a}_L \circ \mathbf{a}_R = \mathbf{a}_O$. If one can prove knowledge of scalars $y, r \in \mathbb{Z}_\Lambda$ and of vectors $\mathbf{a}_L, \mathbf{a}_R$ and \mathbf{a}_O such that $Y = G_1^y H_1^r$ and $\mathbf{a}_L \circ \mathbf{a}_R = \mathbf{a}_O$, and such that the vectors are of the form above, then one can prove knowledge of $(a_i)_{i=0}^n \in \prod_i \{1, G_2^{2^i}\}$ and $(b_i)_{i=0}^{n-1}$ such that $y = a_n b_{n-1} = a_n a_{n-1} b_{n-2} = \dots = a_n a_{n-1} \cdots a_1 b_0 = a_n \cdots a_0$ and $Y = G_1^y H_1^r$. That is to say, one can prove knowledge of a double discrete logarithm.

To prove such a relation, one can use the argument of Bünz et al. [16] for arithmetic circuits with the right linear constraints to ensure that the vectors

are of the appropriate form. To express these constraints, consider matrices

$$\mathbf{W}_L := \begin{bmatrix} \mathbf{0}_{(n+2) \times 2(n+1)} \\ \mathbf{I}_{n+1} & -\mathbf{I}_{n+1} \\ \mathbf{I}_{n+1} & \mathbf{0}_{(n+1) \times (n+1)} \\ \mathbf{0}_{(n+1) \times 2(n+1)} \end{bmatrix}, \mathbf{W}_R := \begin{bmatrix} 0 & 1 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 1 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 & 1 & 0 & \mathbf{0}_{(n+2) \times (n+1)} \\ 0 & 0 & \cdots & 0 & 0 & 1 & \\ 0 & 0 & \cdots & 0 & 0 & 0 & \\ 1 & 0 & \cdots & 0 & 0 & 0 & \\ & & & & & & \mathbf{0}_{(n+1) \times 2(n+1)} \\ & & & & & & \mathbf{0}_{(n+1) \times (n+1)} & -\mathbf{I}_{n+1} \\ & & & & & & \mathbf{0}_{(n+1) \times 2(n+1)} \end{bmatrix},$$

$$\mathbf{W}_O := \begin{bmatrix} 0 \\ -\mathbf{I}_n & \vdots \\ 0 & \mathbf{0}_{(n+2) \times (n+1)} \\ 0 & \cdots & 0 & 1 \\ 0 & \cdots & 0 & 0 \\ & & & \mathbf{0}_{2(n+1) \times 2(n+1)} \\ \mathbf{0}_{(n+1) \times (n+1)} & \mathbf{I}_{n+1} \end{bmatrix}, \mathbf{W}_V := \begin{bmatrix} \mathbf{0}_{n \times 2} \\ 0 & 1 \\ 1 & 0 \\ \mathbf{0}_{3(n+1) \times 2} \end{bmatrix},$$

and vectors $\mathbf{v} := \begin{bmatrix} 1 \\ y \end{bmatrix}$, $\mathbf{c}^T := [\mathbf{0}_{1 \times (n+2)} \mathbf{1}^{n+1} \mathbf{G}_2^{2^{n+1}} \mathbf{0}_{1 \times (n+1)}]$.

Three vectors $\mathbf{a}_L, \mathbf{a}_R$ and $\mathbf{a}_O \in \mathbb{Z}_\Lambda^{2(n+1)}$ satisfy the equation $\mathbf{W}_L \mathbf{a}_L + \mathbf{W}_R \mathbf{a}_R + \mathbf{W}_O \mathbf{a}_O = \mathbf{W}_V \mathbf{v} + \mathbf{c}$ if and only if there exists $\mathbf{a} \in \mathbb{Z}_\Lambda^{n+1}$ and $\mathbf{b} \in \mathbb{Z}_\Lambda^n$ such that $\mathbf{a}_L^T := [\mathbf{a} \ \mathbf{a} - \mathbf{1}^{n+1}]$, $\mathbf{a}_R^T := [1 \ \mathbf{b} \ \mathbf{a} - \mathbf{G}_2^{2^{n+1}}]$ and $\mathbf{a}_O^T := [\mathbf{b} \ y \ \mathbf{0}^{n+1}] \in \mathbb{Z}_\Lambda^{2(n+1)}$.

Indeed,

- * the first n rows of the equation guarantee that for $i = 0, \dots, n-1$, $\mathbf{a}_{O,i} = \mathbf{a}_{R,i+1}$,
- * the $n+1$ th line ensures that $\mathbf{a}_{O,n+1} = y$,
- * the $n+2$ th line imposes $\mathbf{a}_{R,1} = 1$ (G_1 is here used as a commitment to 1),
- * the next $n+1$ lines are satisfied if and only if $\mathbf{a}_{L,[n+1]} = \mathbf{a}_{L,[n+1:]} + \mathbf{1}^{n+1}$,
- * the next $n+1$ lines guarantee that $\mathbf{a}_{R,[n+1]} = \mathbf{a}_{L,[n+1:]} + \mathbf{G}_2^{2^{n+1}}$,
- * the last $n+1$ lines ensure that $\mathbf{a}_{O,[2(n+1):]} = \mathbf{0}^{n+1}$.

If vectors $\mathbf{a}_L, \mathbf{a}_R$ and \mathbf{a}_O additionally satisfy $\mathbf{a}_L \circ \mathbf{a}_R = \mathbf{a}_O$, then $a_i = \mathbf{a}_{L,i} \in \{1, G_2^{2^i}\}$ for $i = 0, \dots, n$ and $\mathbf{a} \circ (1 \ \mathbf{b}) = (\mathbf{b} \ y)$.

The argument of Bünz et al. is therefore sufficient to prove in zero-knowledge knowledge of a double discrete logarithm. The soundness of the proof relies on the discrete logarithm assumption over \mathbb{G}_1 .

Regarding the proof size, the prover sends $(2\lceil \log_2 2(n+1) \rceil + 8) \mathbb{G}_1$ elements and $5 \mathbb{Z}_\Lambda$ elements. Notice that the argument of Bünz et al. requires $4(n+1)$ elements of \mathbb{G}_1^* in addition to G_1 and H_1 . To guarantee its soundness, no discrete logarithm relation between these elements, G_1 and H_1 must be known to the prover. They can then be chosen uniformly at random during set-up.

5.3 Logarithmic-Size Argument of Discrete-Logarithm Equality in two Groups

Building on the argument of double discrete logarithm of Section 5.2, this section gives a zero-knowledge argument for the relation

$$\mathcal{R}_{\text{DLOG-2}} := \left\{ (G_1, H_1, G_2, H_2, Y, X; x \in \mathbb{Z}_\ell, r_1, r_2 \in \mathbb{Z}_\Lambda) : Y = G_1^{G_2^x} H_1^{r_1}, \right. \\ \left. X = G_2^x H_2^{r_2} \right\}.$$

As in Section 5.2, write G_2^x as $\prod_i \overbrace{(x_i(G_2^{2^i} - 1) + 1)}^{a_i}$ for $x_i \in \{0, 1\}$, and $H_2^{r_2}$ as $\prod_i \overbrace{(r_{2,i}(H_2^{2^i} - 1) + 1)}^{c_i}$ for $r_{2,i} \in \{0, 1\}$. Note that $G_2^x = G_1^{\prod_i a_i \prod_i c_i}$. An argument for \mathcal{R} is then equivalent to an argument for

$$\left\{ (G_1, H_1, G_2, H_2, Y, X; (a_i)_{i=0}^n, (c_i)_{i=0}^n, r \in \mathbb{Z}_\Lambda) : Y = G_1^{\prod_i a_i} H_1^r, \right. \\ \left. G_1^x = G_1^{\prod_i a_i \prod_i c_i} \wedge a_i, c_i \in \{1, G_2^{2^i}\} \right\}.$$

To give an argument for this latter relation, consider the following array (written over several lines)

$$\begin{array}{cccccc} a_0 & a_1 & \cdots & a_n & c_0 & \\ 1 & b_0 = a_0 & \cdots & b_{n-1} = a_0 \cdots a_{n-1} & b_{n-1} a_n & \\ a_0 & b_1 = a_0 a_1 & \cdots & G_2^x = a_0 \cdots a_n & a_0 \cdots a_n c_0 & \\ & & & & & \\ \cdots & c_1 & \cdots & c_n & & \\ \cdots & d_0 = a_0 \cdots a_n c_0 & \cdots & d_{n-1} = a_0 \cdots a_n c_0 \cdots c_{n-1} & & \\ \cdots & a_0 \cdots a_n c_0 c_1 & \cdots & X = \prod_i a_i \prod_i c_i & & \end{array}$$

Its third row is the product of the first two. It follows that for

$$\mathbf{a}_L^T := [\mathbf{a} \ \mathbf{a} - \mathbf{1}^{n+1} \ \mathbf{c} \ \mathbf{c} - \mathbf{1}^{n+1}], \\ \mathbf{a}_R^T := [1 \ \mathbf{b} \ \mathbf{a} - \mathbf{G}_2^{2^{n+1}} \ \prod_i a_i \ \mathbf{d} \ \mathbf{c} - \mathbf{H}_2^{2^{n+1}}], \\ \mathbf{a}_O^T := [\mathbf{b} \ \prod_i a_i \ \mathbf{0}^{n+1} \ \mathbf{d} \ X \ \mathbf{0}^{n+1}] \in \mathbb{Z}_\Lambda^{4(n+1)},$$

the equality $\mathbf{a}_L \circ \mathbf{a}_R = \mathbf{a}_O$ holds.

As in Section 5.2, it suffices to find linear constraints, i.e., matrices \mathbf{M}_L , \mathbf{M}_R , \mathbf{M}_O and \mathbf{M}_V , and a vector \mathbf{c} , to enforce that three vectors \mathbf{a}_L , \mathbf{a}_R and \mathbf{a}_O that satisfy the equation $\mathbf{M}_L \mathbf{a}_L + \mathbf{M}_R \mathbf{a}_R + \mathbf{M}_O \mathbf{a}_O = \mathbf{M}_V \mathbf{v} + \mathbf{c}$, where $\mathbf{v} = \begin{bmatrix} 1 \\ X \end{bmatrix}$ (G_1 and G_1^X are respectively used as commitments to 1 and X), are of the form above.

To express such constraints, consider the matrices \mathbf{W}_L , \mathbf{W}_R and \mathbf{W}_O of Section 5.2, and let \mathbf{W}'_L , \mathbf{W}'_R and \mathbf{W}'_O be their respective sub-matrices obtained by

removing the $n+1$ th line (the one that enforced that $\mathbf{a}_{O,n+1} = y$ therein) and \mathbf{W}'_L , \mathbf{W}'_R and \mathbf{W}''_O their sub-matrices obtained by removing the $n+2$ th line (the one which imposed that $\mathbf{a}_{R,1} = 1$). Define also \mathbf{W}'_V as the sub-matrix of \mathbf{W}_V obtained by removing its $n+1$ th line and its second column and \mathbf{W}''_V as the sub-matrix of \mathbf{W}'_V obtained by removing its $n+2$ th line and its first column.

Consider then the following matrices

$$\begin{aligned} \mathbf{M}_L &:= \begin{bmatrix} \mathbf{W}'_L & \\ & \mathbf{W}''_L \end{bmatrix}, \quad \mathbf{M}_R := \begin{bmatrix} \mathbf{W}'_R & \\ & \mathbf{W}''_R \end{bmatrix} \\ \mathbf{M}_O &:= \begin{bmatrix} \mathbf{W}'_O & \\ & \mathbf{W}''_O \end{bmatrix}, \quad \mathbf{M}_V := \begin{bmatrix} \mathbf{W}'_V & \\ & \mathbf{W}''_V \end{bmatrix}. \end{aligned}$$

By definitions of \mathbf{M}_L , \mathbf{M}_R , \mathbf{M}_O , \mathbf{M}_V , \mathbf{v} and \mathbf{c} , three vectors \mathbf{a}_L , \mathbf{a}_R and $\mathbf{a}_O \in \mathbb{Z}_\Lambda^{(n+1)(n+8)}$ satisfy the equations $\mathbf{M}_L \mathbf{a}_L + \mathbf{M}_R \mathbf{a}_R + \mathbf{M}_O \mathbf{a}_O = \mathbf{M}_V \mathbf{v} + \mathbf{d}$ and $\mathbf{a}_L \circ \mathbf{a}_R = \mathbf{a}_O$ if and only if they are of the form above.

The argument of Bünz et al. is therefore sufficient to prove the relation $\mathcal{R}_{\text{DLOG-2}}$. The prover sends $2 \lceil \log_2 4(n+1) \rceil + 8 \mathbb{G}_1$ elements and $5\mathbb{Z}_\Lambda$ elements.

5.4 An Intermediate Protocol in \mathbb{G}_2

This section gives an perfect honest verifier zero-knowledge protocol for relation

$$\begin{aligned} \mathcal{R}'_0 &:= \left\{ (G_2, H_2, X_p, X_q, U, K_u, (K_\gamma)_{\gamma \neq i,j}, (a_\gamma)_\gamma; x_p, x_q, r_p, r_q, u, \rho_u) : \right. \\ &\quad \forall \pi \in \{p, q\}, X_\pi = G_2^{x_\pi} H_2^{r_\pi}, U = G_2^u H_2^{\rho_u}, \\ &\quad \left. \forall \gamma, a_\gamma = G_2^{x_\gamma}, x_\pi(u + K_\pi) = x_\gamma(u + K_\gamma) = 1 \pmod{\ell} \right\}. \end{aligned}$$

To this end, note that for $\pi \in \{p, q\}$, $(UG_2^{K_\pi})^{x_\pi} H_2^{-x_\pi \rho_u} = G_2$, and that $\forall \gamma$, $a_\gamma^u = G_2 a_\gamma^{-K_\gamma}$. One can then apply to the $G_2 a_\gamma^{-K_\gamma}$ values techniques for proving simultaneous equality of discrete logarithms [23]. In the present case, they consist in letting the verifier choose $j-2$ uniformly random values $r_1, \dots, r_{i-1}, r_{i+1}, \dots, r_{j-1}$ of bit-length λ . The prover then computes a proof of knowledge of the discrete logarithm or $\prod_\gamma a_\gamma^{r_\gamma}$ to the base $\prod_\gamma (G_2 a_\gamma^{-K_\gamma})^{r_\gamma}$ (see Figure 4).

As the proof system is public-coin, it can be made non-interactive in the random-oracle model via the Fiat-Shamir heuristic by computing $(r_\gamma)_\gamma$ as $\mathcal{H}_0(G_2, H_2, (X_\pi), U, (K_\pi), (K_\gamma)_\gamma, (a_\gamma)_\gamma)$ and c as $\mathcal{H}_1(G_2, H_2, (X_\pi), U, (K_\pi), (K_\gamma)_\gamma, (a_\gamma)_\gamma, (r_\gamma)_\gamma, (Y_\pi), V, (H_\pi), B)$ for two random oracles \mathcal{H}_0 and \mathcal{H}_1 with appropriate ranges. The proof then consists of $(c, (z_\pi, t_\pi)_{\pi \in \{p,q\}}, w, \tau_u, (\tau_\pi)_\pi)$, i.e., $9 \mathbb{Z}_\ell$ elements.

The protocol is complete and perfectly honest-verifier zero-knowledge, and satisfies witness-extended emulation under the discrete logarithm assumption over \mathbb{G}_2 . The completeness of the protocol and the fact that it is zero-knowledge are straightforward. To prove that the protocol satisfies the witness-extended-emulation property, note that from two distinct accepting transcripts $((r_\gamma)_\gamma, (Y_\pi),$

$V, (H_\pi), B, c, ((z_\pi), (t_\pi), w, \tau_u, (\tau_\pi))$ and $((r_\gamma)_\gamma, ((Y_\pi), V, (H_\pi), B), c', ((z'_\pi), (t'_\pi)), w', \tau'_u, (\tau'_\pi))$, one has

$$\begin{aligned} G_2^{(z_\pi - z'_\pi)/(c' - c)} H_2^{(t_\pi - t'_\pi)/(c' - c)} &= X_\pi, \\ G_2^{(w - w')/(c' - c)} H_2^{(\tau_u - \tau'_u)/(c' - c)} &= U, \\ (UG_2^{K_u})^{(z_\pi - z'_\pi)/(c' - c)} H_2^{(\tau_\pi - \tau'_\pi)/(c' - c)} &= G_2, \\ \prod_\gamma (a_\gamma^{(w - w')/(c' - c)})^{r_\gamma} &= \prod_\gamma (G_2 a_\gamma^{-K_\gamma})^{r_\gamma}. \end{aligned}$$

Replacing U in the third line with the expression in the second, one has $G_2^{(w - w') + K_u(z - z')/(c' - c)} H_2^{(\tau_u - \tau'_u)(z - z')/(c' - c)^2 + (\tau_\pi - \tau'_\pi)/(c' - c)} = G_2$. Under the discrete-logarithm assumption over \mathbb{G}_2 , $(\tau_u - \tau'_u)(z_\pi - z'_\pi)/(c' - c)^2 + (\tau_\pi - \tau'_\pi)/(c' - c) = 0 \pmod{\ell}$ and $(w - w')/(c' - c) + K_u(z_\pi - z'_\pi)/(c' - c) = 1 \pmod{\ell}$.

Moreover, unless $a_\gamma^{(w - w')/(c' - c)} = G_2 a_\gamma^{-K_\gamma}$ for all $\gamma \neq i, j$, the last equality holds with probability at most $2^{-\lambda}$ since the r_γ values are pairwise independent and independent of $(c, ((z_\pi), (t_\pi), w, \tau_u, (\tau_\pi))$ and $(c', ((z'_\pi), (t'_\pi), w', \tau'_u, (\tau'_\pi)))$.

It follows that setting $x_\pi \leftarrow (z_\pi - z'_\pi)/(c' - c)$, $r \leftarrow (t_\pi - t'_\pi)/(c' - c)$, $u \leftarrow (w - w')/(c' - c)$ and $\rho_u \leftarrow (\tau_u - \tau'_u)/(c' - c)$, the tuple $((x_\pi), (r_\pi), u, \rho_u)$ is a valid witness for relation \mathcal{R}'_0 .

To extract a witness from a prover with a fixed random string, it suffices to repeat the following procedure:

1. run the protocol once, rewind the protocol to the computation step right after the first challenge was sent and run it a second time with a fresh verifier randomness
2. if the verifier accepts both executions, and the second challenges of these latter are different, extract a witness as above; otherwise restart.

The expected running time of this procedure is at most $1/(\psi - 2^{-\lambda} - \ell^{-1})^2 \sim 1/\psi^2$, with ψ the probability that the verifier accepts.

To produce an indistinguishable transcript, the emulator simply runs the protocol with uniform challenges as the protocol is honest verifier.

5.5 Protocol for \mathcal{R}_0

To prove knowledge of a witness for \mathcal{R}_0 , the prover starts setting by setting $K_p := \mathcal{H}(r_{C\mathcal{A}}) + i$, $K_q := \mathcal{H}(r_{C\mathcal{A}}) + j$, and $u := r'_u$. It then

- computes two commitments $X_p = G_2^{x_p} H_2^{r_p}$ and $X_q = G_2^{x_q} H_2^{r_q}$, for $x_p = (u + K_p)^{-1} \pmod{\ell}$ and $x_q = (u + K_q)^{-1} \pmod{\ell}$
- computes a proof $\pi_{\text{DLOG-2},p}$ that the double discrete-logarithm of P is the discrete logarithm of X_p , and similarly a proof $\pi_{\text{DLOG-2},q}$ for Q and X_q
- computes a proof π' for relation \mathcal{R}'_0 with X_p and X_q .

The final proof π_0 for \mathcal{R}_0 consists of $(X_p, X_q, \pi_{\text{DLOG-2},p}, \pi_{\text{DLOG-2},q}, \pi'_0)$.




Fig. 4. Honest-Verifier Zero-Knowledge Protocol for Relation \mathcal{R}_1 .

Security. It is important to note that the security of the generated key is weakened compared to an RSA-key of the same size since the CA can recover seed s (and thus the prime factors) by solving a discrete logarithm problem in \mathbb{G}_2 . For 3072-bit RSA moduli, this protocol therefore only provides 96 bits of security (with respect to the CA) instead of the expected 128-bit security level. To avoid this issue, one can increase the bit size of the prime numbers to 3072 bits (but at the cost of generating RSA moduli of twice this size). Another possibility is to use other groups for \mathbb{G}_2 with (alleged) harder discrete logarithm problem, e.g., the group of points of an elliptic curve over \mathbb{F}_p or an algebraic torus defined over \mathbb{F}_{p^2} (with compact representation of group elements in \mathbb{F}_p) for a 1536-bit prime p . This may however introduce a new bias for the generated primes and require to adapt the zero-knowledge proofs.

Efficiency. The asymptotic complexity of the communication size depends on the number of trials to obtain two primes in W since the prover has to send $(a_\gamma)_{\gamma \neq i, j}$. However, even though the communication is asymptotically linear in the number of trials, the overhead incurred by the proof of correct computation should in practice be small.

	\mathbb{Z}_ℓ	\mathbb{Z}_N	\mathbb{Z}_Λ	\mathbb{G}_1	\mathbb{G}_2	Total (kB)
\mathcal{R}_0	9	0	10	$4 \lceil \log_2 4b \rceil + 16$	2	346
\mathcal{R}_1	0	0	5	$2 \lceil \log_2 2(b-1) \rceil + 4$	0	142

Fig. 5. Size of the Zero-Knowledge Arguments (for a 96-bit security level and 3072-bit RSA moduli).

Total Proof Size. As discussed in Section 5.3, proofs $\pi_{\text{DLOG-2},p}$ and $\pi_{\text{DLOG-2},q}$ both consists of $2\lceil\log_2 4(n+1)\rceil + 8 \mathbb{G}_1$ elements and $5\mathbb{Z}_\Lambda$ elements.

Proof π' consists of $9 \mathbb{Z}_\ell$ elements (see Section 5.4). Proof π_0 for \mathcal{R}_0 therefore consists of $2 \mathbb{G}_2$ elements, $4\lceil\log_2 4(n+1)\rceil + 16 \mathbb{G}_1$ elements, $10 \mathbb{Z}_\Lambda$ elements and $9 \mathbb{Z}_\ell$ elements. As for the proof for \mathcal{R}_1 , the aggregated proof that 2 values committed in \mathbb{G}_1 are in $\llbracket 0, 2^{b-1} - 1 \rrbracket$ consists of $2\lceil\log_2 2(b-1)\rceil + 4 \mathbb{G}_1$ elements (recall that $n+1 = b$) and $5 \mathbb{Z}_\Lambda$ elements.

Running Time. An important question about the protocol is the number of necessary PRF trials to obtain two primes that satisfy the conditions required for the factors of N (captured by $W \subseteq \mathbb{P}$). We estimate the number j of necessary trials in the case $W = \mathbb{P} \cap \llbracket 2^{b-1}, 2^b - 1 \rrbracket$, i.e., when \mathcal{U} simply has to prove that p and q are prime of $b(\lambda)$ bits. The following analysis shows (using a number-theoretic heuristic) that the number of trials exceeds $17b(\lambda) = O(\log \lambda)$ (so the DY PRF remains secure), and that the probability that it is larger than that decreases exponentially fast.

For an integer x , denote by $\pi(x)$ the number of primes smaller or equal to x . For $x \geq 30$, Chebyshev's theorem [29, p. 9] ensures that $0.92x/\ln(x) < \pi(x) < 1.11x/\ln(x)$. Therefore, if $b > \log_2 30 + 1$, $\pi(2^b) > 0.92 \cdot 2^b/b \ln(2)$ and $\pi(2^{b-1}) < 1.11 \cdot 2^b/2(b-1) \ln(2)$.

For two integers q and a and an integer x , denote by $\pi(x; q, a)$ the number of primes congruent to a modulo q and smaller or equal to x . For all $x > q$, then the Brun–Titchmarsh theorem [41] implies that $\pi(x; q, a) \leq \frac{2x}{\varphi(q) \ln(x/q)}$.

The number of primes in $\llbracket 2^{b-1}, 2^b - 1 \rrbracket$ which are not congruent to 1 modulo e where e is a prime number greater than 17 is then at least

$$0.92 \cdot \frac{2^b}{b \ln(2)} - \frac{2 \cdot 2^b}{(e-1)(b \ln(2) - \ln(e))} - 1.11 \cdot \frac{2^b}{2(b-1) \ln(2)} \geq 0.12 \frac{2^b}{(b - \ln(e)/\ln(2))}$$

for $b \geq 6 \ln(e)/\ln(2)$.

To apply Theorem 4.5, it remains to estimate the number of b -bit primes p in $\mathcal{QR}_{2\ell+1}$ and such that $\gcd(e, p-1) = 1$.

Assuming that around half of the primes in $\llbracket 2^{b-1}, 2^b - 1 \rrbracket$ that are not congruent to 1 modulo e are quadratic residues modulo $2\ell+1$, the factor β in Theorem 4.5 is at least $0.12/(b - \ln(e)/\ln(2)) \geq 0.12/b$ (the range of the PRF is has size ℓ).

This heuristic is supported by the fact that half of the integers in $\llbracket 2\ell \rrbracket$ are quadratic residues, and that primes chosen uniformly at random often have properties similar to those of general integers chosen uniformly at random.

Overall Communication Size. In the last flow of the protocol, the prover then sends an integer N , two commitments in \mathbb{G}_1 , $17b(\lambda) - 2$ integers in $\llbracket 0, 2\ell \rrbracket$ with high probability, i.e., the $(a_\gamma)_{\gamma \neq i, j}$ values which are integers returned by the PRF and not in W , and the proof of correct computation of which the size is summarized in Table 5.

Acknowledgements

The authors thank Jonathan Bootle for fruitful discussions. This work was supported by the French ANR ALAMBIC Project (ANR-16-CE39-0006), the CHIST-ERA USEIT project and the EU H2020 Research and Innovation Program under Grant Agreement No. 786725 (OLYMPUS).

References

1. Abdalla, M., Benhamouda, F., Blazy, O., Chevalier, C., Pointcheval, D.: SPHF-friendly non-interactive commitments. In: Sako, K., Sarkar, P. (eds.) ASIACRYPT 2013, Part I. LNCS, vol. 8269, pp. 214–234. Springer, Heidelberg (Dec 2013)
2. Agrawal, M., Kayal, N., Saxena, N.: PRIMES is in P. *Ann. Math. (2)* 2004(2), 781–793 (2004)
3. Au, M.H., Susilo, W., Mu, Y.: Proof-of-knowledge of representation of committed value and its applications. In: Steinfeld, R., Hawkes, P. (eds.) ACISP 10. LNCS, vol. 6168, pp. 352–369. Springer, Heidelberg (Jul 2010)
4. Auerbach, B., Poettering, B.: Hashing solutions instead of generating problems: On the interactive certification of RSA moduli. In: Abdalla, M., Dahab, R. (eds.) PKC 2018, Part II. LNCS, vol. 10770, pp. 403–430. Springer, Heidelberg (Mar 2018)
5. Backes, M., Hanzlik, L., Herzberg, A., Kate, A., Pryvalov, I.: Efficient non-interactive zero-knowledge proofs in cross-domains without trusted setup. In: Lin, D., Sako, K. (eds.) PKC 2019, Part I. LNCS, vol. 11442, pp. 286–313. Springer, Heidelberg (Apr 2019)
6. Belaïd, S., Benhamouda, F., Passelègue, A., Prouff, E., Thillard, A., Vergnaud, D.: Private multiplication over finite fields. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017, Part III. LNCS, vol. 10403, pp. 397–426. Springer, Heidelberg (Aug 2017)
7. Bellare, M., Canetti, R., Krawczyk, H.: Keying hash functions for message authentication. In: Kobitz, N. (ed.) CRYPTO'96. LNCS, vol. 1109, pp. 1–15. Springer, Heidelberg (Aug 1996)
8. Bellare, M., Hoang, V.T., Keelveedhi, S.: Instantiating random oracles via UCEs. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part II. LNCS, vol. 8043, pp. 398–415. Springer, Heidelberg (Aug 2013)
9. Bellare, M., Paterson, K.G., Rogaway, P.: Security of symmetric encryption against mass surveillance. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014, Part I. LNCS, vol. 8616, pp. 1–19. Springer, Heidelberg (Aug 2014)
10. Bellare, M., Pointcheval, D., Rogaway, P.: Authenticated key exchange secure against dictionary attacks. In: Preneel, B. (ed.) EUROCRYPT 2000. LNCS, vol. 1807, pp. 139–155. Springer, Heidelberg (May 2000)
11. Bellare, M., Rogaway, P.: Random oracles are practical: A paradigm for designing efficient protocols. In: Denning, D.E., Pyle, R., Ganesan, R., Sandhu, R.S., Ashby, V. (eds.) ACM CCS 93. pp. 62–73. ACM Press (Nov 1993)
12. Benhamouda, F., Ferradi, H., Géraud, R., Naccache, D.: Non-interactive provably secure attestations for arbitrary RSA prime generation algorithms. In: Foley, S.N., Gollmann, D., Sneekenes, E. (eds.) ESORICS 2017, Part I. LNCS, vol. 10492, pp. 206–223. Springer, Heidelberg (Sep 2017)

13. Boneh, D., Boyen, X.: Short signatures without random oracles. In: Cachin, C., Camenisch, J. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 56–73. Springer, Heidelberg (May 2004)
14. Bootle, J., Cerulli, A., Chaidos, P., Groth, J., Petit, C.: Efficient zero-knowledge arguments for arithmetic circuits in the discrete log setting. In: Fischlin, M., Coron, J.S. (eds.) EUROCRYPT 2016, Part II. LNCS, vol. 9666, pp. 327–357. Springer, Heidelberg (May 2016)
15. Boyar, J., Friedl, K., Lund, C.: Practical zero-knowledge proofs: Giving hints and using deficiencies. In: Quisquater, J.J., Vandewalle, J. (eds.) EUROCRYPT'89. LNCS, vol. 434, pp. 155–172. Springer, Heidelberg (Apr 1990)
16. Büinz, B., Bootle, J., Boneh, D., Poelstra, A., Wuille, P., Maxwell, G.: Bulletproofs: Short proofs for confidential transactions and more. In: 2018 IEEE Symposium on Security and Privacy. pp. 315–334. IEEE Computer Society Press (May 2018)
17. Camenisch, J., Hohenberger, S., Lysyanskaya, A.: Compact e-cash. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 302–321. Springer, Heidelberg (May 2005)
18. Camenisch, J., Michels, M.: Proving in zero-knowledge that a number is the product of two safe primes. In: Stern, J. (ed.) EUROCRYPT'99. LNCS, vol. 1592, pp. 107–122. Springer, Heidelberg (May 1999)
19. Camenisch, J., Stadler, M.: Efficient group signature schemes for large groups (extended abstract). In: Kaliski Jr., B.S. (ed.) CRYPTO'97. LNCS, vol. 1294, pp. 410–424. Springer, Heidelberg (Aug 1997)
20. Canard, S., Gouget, A.: Divisible e-cash systems can be truly anonymous. In: Naor, M. (ed.) EUROCRYPT 2007. LNCS, vol. 4515, pp. 482–497. Springer, Heidelberg (May 2007)
21. Chase, M., Ganesh, C., Mohassel, P.: Efficient zero-knowledge proof of algebraic and non-algebraic statements with applications to privacy preserving credentials. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016, Part III. LNCS, vol. 9816, pp. 499–530. Springer, Heidelberg (Aug 2016)
22. Checkoway, S., Maskiewicz, J., Garman, C., Fried, J., Cohney, S., Green, M., Heninger, N., Weinmann, R.P., Rescorla, E., Shacham, H.: A systematic analysis of the juniper dual EC incident. In: Weippl, E.R., Katzenbeisser, S., Kruegel, C., Myers, A.C., Halevi, S. (eds.) ACM CCS 2016. pp. 468–479. ACM Press (Oct 2016)
23. Chow, S.S.M., Ma, C., Weng, J.: Zero-knowledge argument for simultaneous discrete logarithms. *Algorithmica* 64(2), 246–266 (2012)
24. Corrigan-Gibbs, H., Mu, W., Boneh, D., Ford, B.: Ensuring high-quality randomness in cryptographic key generation. In: Sadeghi, A.R., Gligor, V.D., Yung, M. (eds.) ACM CCS 2013. pp. 685–696. ACM Press (Nov 2013)
25. Cramer, R., Shoup, V.: A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In: Krawczyk, H. (ed.) CRYPTO'98. LNCS, vol. 1462, pp. 13–25. Springer, Heidelberg (Aug 1998)
26. Damgård, I., Haagh, H., Orlandi, C.: Access control encryption: Enforcing information flow with cryptography. In: Hirt, M., Smith, A.D. (eds.) TCC 2016-B, Part II. LNCS, vol. 9986, pp. 547–576. Springer, Heidelberg (Oct / Nov 2016)
27. Desmedt, Y.: Simmons' protocol is not free of subliminal channels. In: Ninth IEEE Computer Security Foundations Workshop, March 10 - 12, 1996, Dromquinna Manor, Kenmare, County Kerry, Ireland. pp. 170–175 (1996)
28. Dodis, Y., Yampolskiy, A.: A verifiable random function with short proofs and keys. In: Vaudenay, S. (ed.) PKC 2005. LNCS, vol. 3386, pp. 416–431. Springer, Heidelberg (Jan 2005)

29. Dusart, P.: Autour de la fonction qui compte le nombre de nombres premiers. Ph.D. thesis, Université de Limoges (1998)
30. Fiat, A., Shamir, A.: How to prove yourself: Practical solutions to identification and signature problems. In: Odlyzko, A.M. (ed.) CRYPTO'86. LNCS, vol. 263, pp. 186–194. Springer, Heidelberg (Aug 1987)
31. Gennaro, R., Micciancio, D., Rabin, T.: An efficient non-interactive statistical zero-knowledge proof system for quasi-safe prime products. In: Gong, L., Reiter, M.K. (eds.) ACM CCS 98. pp. 67–72. ACM Press (Nov 1998)
32. Goldberg, S., Reyzin, L., Sagga, O., Baldimtsi, F.: Efficient noninteractive certification of RSA moduli and beyond. In: Galbraith, S.D., Moriai, S. (eds.) ASIACRYPT 2019, Part III. LNCS, vol. 11923, pp. 700–727. Springer, Heidelberg (Dec 2019)
33. Goldreich, O., Goldwasser, S., Micali, S.: How to construct random functions (extended abstract). In: 25th FOCS. pp. 464–479. IEEE Computer Society Press (Oct 1984)
34. Groth, J., Ostrovsky, R., Sahai, A.: Non-interactive zaps and new techniques for NIZK. In: Dwork, C. (ed.) CRYPTO 2006. LNCS, vol. 4117, pp. 97–111. Springer, Heidelberg (Aug 2006)
35. Heninger, N., Durumeric, Z., Wustrow, E., Halderman, J.A.: Mining your ps and qs: Detection of widespread weak keys in network devices. In: Kohno, T. (ed.) USENIX Security 2012. pp. 205–220. USENIX Association (Aug 2012)
36. Izabachène, M., Pointcheval, D., Vergnaud, D.: Mediated traceable anonymous encryption. In: Abdalla, M., Barreto, P.S.L.M. (eds.) LATINCRYPT 2010. LNCS, vol. 6212, pp. 40–60. Springer, Heidelberg (Aug 2010)
37. Juels, A., Guajardo, J.: RSA key generation with verifiable randomness. In: Naccache, D., Paillier, P. (eds.) PKC 2002. LNCS, vol. 2274, pp. 357–374. Springer, Heidelberg (Feb 2002)
38. Lenstra, A.K., Hughes, J.P., Augier, M., Bos, J.W., Kleinjung, T., Wachter, C.: Public keys. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 626–642. Springer, Heidelberg (Aug 2012)
39. Li, X.: Improved two-source extractors, and affine extractors for polylogarithmic entropy. In: Dinur, I. (ed.) 57th FOCS. pp. 168–177. IEEE Computer Society Press (Oct 2016)
40. Lindell, Y.: Parallel coin-tossing and constant-round secure two-party computation. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 171–189. Springer, Heidelberg (Aug 2001)
41. Montgomery, H.L., Vaughan, R.C.: The large sieve. *Mathematika* 20(2), 119–134 (1973)
42. Motwani, R., Raghavan, P.: Randomized Algorithms. Cambridge International Series on Parallel Computation, Cambridge University Press (1995)
43. Nemeč, M., Šýs, M., Svenda, P., Klinec, D., Matyas, V.: The return of copper-smith's attack: Practical factorization of widely used RSA moduli. In: Thuraisingham, B.M., Evans, D., Malkin, T., Xu, D. (eds.) ACM CCS 2017. pp. 1631–1648. ACM Press (Oct / Nov 2017)
44. NIST: National Institute of Standards and Technology – Digital signature standard (dss). <https://csrc.nist.gov/publications/detail/fips/186/4/final> (2013)
45. Pedersen, T.P.: Non-interactive and information-theoretic secure verifiable secret sharing. In: Feigenbaum, J. (ed.) CRYPTO'91. LNCS, vol. 576, pp. 129–140. Springer, Heidelberg (Aug 1992)

46. Rivest, R.L., Shamir, A., Adleman, L.M.: A method for obtaining digital signatures and public-key cryptosystems. *Communications of the Association for Computing Machinery* 21(2), 120–126 (1978)
47. Simmons, G.J.: The prisoners’ problem and the subliminal channel. In: Chaum, D. (ed.) *CRYPTO’83*. pp. 51–67. Plenum Press, New York, USA (1983)
48. Stadler, M.: Publicly verifiable secret sharing. In: Maurer, U.M. (ed.) *EUROCRYPT’96*. LNCS, vol. 1070, pp. 190–199. Springer, Heidelberg (May 1996)
49. van de Graaf, J., Peralta, R.: A simple and secure way to show the validity of your public key. In: Pomerance, C. (ed.) *CRYPTO’87*. LNCS, vol. 293, pp. 128–134. Springer, Heidelberg (Aug 1988)
50. Yilek, S., Rescorla, E., Shacham, H., Enright, B., Savage, S.: When private keys are public: results from the 2008 debian openssl vulnerability. In: Feldmann, A., Mathy, L. (eds.) *Proceedings of the 9th ACM SIGCOMM Internet Measurement Conference, IMC 2009, Chicago, Illinois, USA, November 4-6, 2009*. pp. 15–27. ACM (2009), <https://doi.org/10.1145/1644893.1644896>
51. Young, A., Yung, M.: Kleptography: Using cryptography against cryptography. In: Fumy, W. (ed.) *EUROCRYPT’97*. LNCS, vol. 1233, pp. 62–74. Springer, Heidelberg (May 1997)

A Additional Preliminaries

This section gives further preliminary material.

A.1 Non-interactive Commitments

The definitions in this section are derived from those of Abdalla et al. [1]. A (non-interactive) commitment scheme consists of the following algorithms.

Setup (1^λ) $\rightarrow pp$: generates public parameters (or common-reference string) on the input of a security parameter 1^λ . These parameters are implicit inputs to the other algorithms.

Com (x) $\rightarrow (C, d)$: computes a commitment C to a value and an opening or decommitment information d .

ComVf (C, x, d) $\rightarrow d \in \{0, 1\}$: a deterministic returns a bit indicating whether the decommitment d is valid (bit 1) for C and x , or not (bit 0). It is assumed that if $x = \perp$, then it returns 0.

A commitment scheme is *correct* if for all $\lambda \in \mathbb{N}$, for all $pp \leftarrow \text{Setup}(1^\lambda)$ and for all x $\Pr[\text{ComVf}(C, x, d) = 1 : (C, d) \leftarrow \text{Com}(x)] = 1$.

A commitment scheme $(T, \varepsilon^{\text{hide}})$ -*hiding* (statistically hiding) if for all $\lambda \in \mathbb{N}$, for every (computationally unbounded) adversary \mathcal{A} that runs it times at most

$T(\lambda)$,

$$\Pr \left[b = b' : \begin{array}{l} pp \leftarrow \text{Setup}(1^\lambda) \\ (x_0, x_1, st) \leftarrow \mathcal{A}(pp) \\ b \leftarrow_{\S} \{0, 1\} \\ (C, d) \leftarrow \text{Com}(x_b) \\ b' \leftarrow \mathcal{A}(st, C) \\ \text{return } (b, b') \end{array} \right] - 1/2 \leq \varepsilon^{\text{hide}}(\lambda).$$

A 0-statistically hiding scheme is said to be *perfectly hiding*.

A commitment scheme is $(T, \varepsilon^{\text{bind}})$ -*binding* if for all $\lambda \in \mathbb{N}$, for every adversary \mathcal{A} that runs in time at most $T(\lambda)$,

$$\Pr \left[\begin{array}{l} \text{ComVf}(C, x_0, d_0) = \text{ComVf}(C, x_1, d_1) = 1 \\ \wedge x_0 \neq x_1 \end{array} : \begin{array}{l} pp \leftarrow \text{Setup}(1^\lambda) \\ (C, x_0, d_0, x_1, d_1) \leftarrow \mathcal{A}(pp) \end{array} \right] \leq \varepsilon^{\text{bind}}(\lambda).$$

Pedersen Commitments. The protocol in Section 5 uses the Pedersen commitment scheme which is homomorphic and allows to algebraic prove properties of a committed values. Given G a group-family generator, let \mathbb{G} be a cyclic group of prime order ℓ generated by G . Let g_1, g_2 denote two generators of \mathbb{G} . The Pedersen-commitment [45] common reference string is $(\mathbb{G}, \ell, g_1, g_2)$. For a value $x \in \mathbb{Z}_\ell$, the pedersen-commitment algorithm computes generates $r \leftarrow_{\S} [0, \ell - 1]$ and computes $h \leftarrow g_1^x g_2^r$. Note that the commitments are perfectly hiding. To verify a commitment h to a with decommitment information r , the opening algorithm simply verifies that $h = g_1^a g_2^r$. Under the discret-logarithm assumption over G , the scheme is binding.

Extractable Commitments. A commitment scheme $(\text{Setup}, \text{Com}, \text{ComVf})$ is extractable if there exists an algorithm $\text{TSetup}(1^\lambda) \rightarrow (pp, \tau)$ which generates public parameters and a trapdoor, and an algorithm $\text{ExtCom}(\tau, C) \rightarrow x/\perp$ which, on the input of a trapdoor and of a commitment, returns a message or \perp if the commitment is invalid.

The scheme then satisfies *trapdoor correctness* if for all 1^λ , for all $(pp, \tau) \leftarrow \text{TSetup}(1^\lambda)$, for all x , $\Pr[\text{ExtCom}(\tau, C) = x : (C, \delta) \leftarrow \text{Com}(x)] = 1$. The commitment scheme satisfies $(T, \varepsilon^{\text{setup-ind}})$ -*setup indistinguishability* if no algorithm running in time at most $T(\lambda)$ can distinguish the output of Setup from the first component of the output of TSetup with an advantage greater than $\varepsilon^{\text{setup-ind}}(\lambda)$. Lastly, an extractable commitment scheme satisfies $(T_{\text{ExtCom}}, T_{\mathcal{A}}, q, \varepsilon^{\text{bind-ext}})$ -*binding extractability* if ExtCom runs in time at most $T_{\text{ExtCom}}(\lambda)$ and if for all 1^λ , for every adversary \mathcal{A} that runs in time at most $T(\lambda)$ and makes at most q oracle queries,

$$\Pr \left[\begin{array}{l} \text{ComVf}(C, x, d) = 1 \\ \wedge x \neq x' \end{array} : \begin{array}{l} (pp, \tau) \leftarrow \text{TSetup}(1^\lambda) \\ (C, x, d) \leftarrow \mathcal{A}^{\text{ExtCom}(\tau, \cdot)}(pp) \\ x' \leftarrow \text{ExtCom}(\tau, C) \end{array} \right] \leq \varepsilon^{\text{bind-ext}}(\lambda).$$

Public-key encryption schemes secure against chosen-ciphertext attacks such as Cramer and Shoup's [25] are perfectly binding commitment schemes which are additionally extractable since the secret key can be used as a trapdoor.

A.2 Interactive Argument Systems

An argument system for a language $\mathcal{L} = \mathcal{L}_{pp}$ (with corresponding relation \mathcal{R}) consists of a triple $\Pi = (\text{Setup}, \text{Prove}, \text{Verf})$ such that $\text{Setup}(1^\lambda) \rightarrow pp$ returns public parameters on the input of a security parameter and $\langle \text{Prove}(pp, x, w) \rightleftharpoons \text{Verf}(pp, x) \rangle \rightarrow (\tau, b) \in \{0, 1\}^* \times \{0, 1\}$ are interactive algorithms (τ denotes the transcript of the interaction and b the decision bit of Verf).

Π is *complete* if for all $\lambda \in \mathbb{N}$, for all $pp \leftarrow \text{Setup}(1^\lambda)$, for all $(x, w) \in \mathcal{R}$ $\Pr[(*, 1) \leftarrow \langle \text{Prove}(pp, x, w) \rightleftharpoons \text{Verf}(pp, x) \rangle] = 1$.

Π satisfies $(T, \varepsilon^{\text{sound}})$ -*soundness* if for all $\lambda \in \mathbb{N}$, for every adversary \mathcal{A} that runs in time at most $T(\lambda)$

$$\Pr \left[\begin{array}{l} pp \leftarrow \text{Setup}(1^\lambda) \\ b = 1 \wedge x \notin \mathcal{L} : (st, x) \leftarrow \mathcal{A}(pp) \\ (\tau, b) \leftarrow \langle \mathcal{A}(st, x) \rightleftharpoons \text{Verf}(pp, x) \rangle \end{array} \right] \leq \varepsilon^{\text{sound}}(\lambda).$$

Π satisfies $(T_{\text{Prove}^*}, T_{\mathcal{E}}, T_{\mathcal{A}}, \varepsilon^{\text{w}2e})$ -*witness-extended emulation* [40] if for all $\lambda \in \mathbb{N}$, for every deterministic algorithm Prove^* running in time at most $T_{\text{Prove}^*}(\lambda)$, there exists an algorithm (emulator) \mathcal{E} running in expected time $T_{\mathcal{E}}(\lambda)$ and such that for every adversary \mathcal{A} that runs in time at most $T_{\mathcal{A}}(\lambda)$,

$$\left| \Pr \left[\begin{array}{l} pp \leftarrow \text{Setup}(1^\lambda) \\ b = 1 : (st, x, s) \leftarrow \mathcal{A}(pp) \\ (\tau, \beta) \leftarrow \langle \text{Prove}^*(pp, x, s) \rightleftharpoons \text{Verf}(pp, x) \rangle \\ b \leftarrow \mathcal{A}(st, (\tau, \beta)) \end{array} \right] - \Pr \left[\begin{array}{l} b = 1 \wedge \\ (\beta = 0 \vee (x, w) \in \mathcal{R}) : (st, x, s) \leftarrow \mathcal{A}(pp) \\ ((\tau, \beta), w) \leftarrow \mathcal{E}^{\langle \text{Prove}^*(pp, x, s) \rightleftharpoons \text{Verf}(pp, x) \rangle}(pp, x) \\ b \leftarrow \mathcal{A}(st, (\tau, \beta)) \end{array} \right] \right| \leq \varepsilon^{\text{w}2e}(\lambda).$$

Algorithm \mathcal{E} is given access to a transcript oracle $\langle \text{Prove}^*(pp, x, s) \rightleftharpoons \text{Verf}(pp, x) \rangle$ which can be rewound to any step of its computation and run anew on fresh verifier-randomness.

This definition means that if Π satisfies witness-extended emulation, then whenever Verf accepts its interaction with Prove^* , emulator \mathcal{E} can produce a closely distributed transcript, but also extract a valid witness. Π is thus also an argument of knowledge, which implies that it is sound. The string s given to Prove^* can be considered as an internal state which includes its random string.

Π is *T-perfectly honest-verifier zero-knowledge* if there exists an algorithm Sim such that for all $\lambda \in \mathbb{N}$, for every adversary \mathcal{A} running in time at most $T(\lambda)$,

$$\Pr \left[(x, w) \in \mathcal{R} \wedge b = 1 : \begin{array}{l} pp \leftarrow \text{Setup}(1^\lambda) \\ (st, x, w, r) \leftarrow \mathcal{A}(pp) \\ (\tau, \beta) \leftarrow \langle \text{Prove}(pp, x, w) \rangle \Rightarrow \text{Verf}(pp, x; r) \\ b \leftarrow \mathcal{A}(st, (\tau, \beta)) \end{array} \right] \\ = \Pr \left[\begin{array}{l} pp \leftarrow \text{Setup}(1^\lambda) \\ (x, w) \in \mathcal{R} \wedge b = 1 : \begin{array}{l} (st, x, w, r) \leftarrow \mathcal{A}(pp) \\ (\tau, \beta) \leftarrow \text{Sim}(pp, x, r) \\ b \leftarrow \mathcal{A}(st, (\tau, \beta)) \end{array} \end{array} \right].$$

Π is said to be *public coin* if all messages sent by Verf are chosen uniformly at random and independently of the messages sent by Prove .

Fiat–Shamir Heuristic. The Fiat–Shamir heuristic [30] can be used to turn a public-coin, interactive argument system into a non-interactive one in the random-oracle model [11]. Given a random oracle \mathcal{H} , the messages of the verifier are computed by evaluating \mathcal{H} at the word and the transcript of the interactive until that point of the computation of the prover. With oracle access to \mathcal{H} , the prover can then compute a full transcript (or argument), further denoted π instead of τ , without interacting with the verifier, and this latter can also verify the transcript without any interaction.

The non-interactive argument system derived from an interactive one $\Pi = (\text{Setup}, \text{Prove}, \text{Verf})$ via the Fiat–Shamir heuristic with a random oracle \mathcal{H} is denoted $\Pi^{\mathcal{H}} := (\text{Setup}^{\mathcal{H}}, \text{Prove}^{\mathcal{H}}, \text{Verf}^{\mathcal{H}})$.

A non-interactive argument system $\Pi^{\mathcal{H}} := (\text{Setup}^{\mathcal{H}}, \text{Prove}^{\mathcal{H}}, \text{Verf}^{\mathcal{H}})$ is said to be *complete* if for all $\lambda \in \mathbb{N}$, for all $pp \leftarrow \text{Setup}(1^\lambda)$, for all $(x, w) \in \mathcal{R}$ $\Pr [\text{Verf}^{\mathcal{H}}(pp, x, \text{Prove}^{\mathcal{H}}(crs, x, w)) = 1] = 1$.

$\Pi^{\mathcal{H}}$ satisfies $(T, q_{\mathcal{H}}, \varepsilon^{\text{sound}})$ -*soundness* if for all $\lambda \in \mathbb{N}$, for every adversary \mathcal{A} that runs in time at most $T(\lambda)$ and makes at most $q_{\mathcal{H}}$ queries to \mathcal{H} ,

$$\Pr \left[\text{Verf}^{\mathcal{H}}(pp, x, \pi) = 1 \wedge x \notin \mathcal{L} : \begin{array}{l} pp \leftarrow \text{Setup}(1^\lambda) \\ (x, \pi) \leftarrow \mathcal{A}^{\mathcal{H}(\cdot)}(crs) \end{array} \right] \leq \varepsilon^{\text{sound}}(\lambda).$$

$\Pi^{\mathcal{H}}$ is $(T_{\text{Prove}^*}, T_{\text{Ext}}, T_{\mathcal{A}}, q_{\mathcal{H}}, \varepsilon^{\text{ext}})$ -*extractable* if for every deterministic algorithm Prove^* running in time at most $T_{\text{Prove}^*}(\lambda)$, there exists an algorithm $(\text{Ext}_0(Q, q) \rightarrow Q', \text{Ext}_1(pp, x) \rightarrow w)$ running in expected time at most $T_{\text{Ext}}(\lambda)$ and such that for all $\lambda \in \mathbb{N}$, for every adversary \mathcal{A} that runs in time at most $T_{\mathcal{A}}(\lambda)$ and

such that Prove^* and \mathcal{A} together make at most $q_{\mathcal{H}}$ oracle queries to \mathcal{H} ,

$$\left[\Pr \left[\begin{array}{l} pp \leftarrow \text{Setup}(1^\lambda) \\ (x, s) \leftarrow \mathcal{A}^{\mathcal{H}(\cdot)}(pp) \\ \pi \leftarrow \text{Prove}^{*\mathcal{H}(\cdot)}(pp, x, s) \\ b \leftarrow \text{Verf}^{\mathcal{H}}(pp, x, \pi) \end{array} \right] \right. \\ \left. - \Pr \left[\begin{array}{l} pp \leftarrow \text{Setup}(1^\lambda); Q \leftarrow \emptyset \\ (x, s) \leftarrow \mathcal{A}^{\text{Ext}_0(Q, \cdot)}(pp) \\ (b, w) \leftarrow \text{Ext}_1^{\text{Prove}^* \text{Ext}_0(Q, \cdot)}(pp, x, s) \end{array} \right] \right] \leq \varepsilon^{\text{ext}}(\lambda).$$

Algorithm $(\text{Ext}_0, \text{Ext}_1)$ is given access to a Prove^* oracle which can be rewound to any step of its computation and run anew with fresh Ext_0 randomness.

This definition means that if $\Pi^{\mathcal{H}}$ is extractable, then whenever Prove^* is able to compute a valid proof, extractor $(\text{Ext}_0, \text{Ext}_1)$ can extract a valid witness. Π is thus an argument of knowledge, which implies that it is sound. The string s given to Prove^* can be considered as an internal state which includes its random string.

Lemma A.1. *Let $\Pi = (\text{Setup}, \text{Prove}, \text{Verf})$ be an argument system for a language \mathcal{L} and $\Pi^{\mathcal{H}} := (\text{Setup}^{\mathcal{H}}, \text{Prove}^{\mathcal{H}}, \text{Verf}^{\mathcal{H}})$ be the non-interactive argument system for \mathcal{L} derived from $\Pi = (\text{Setup}, \text{Prove}, \text{Verf})$ via the Fiat–Shamir heuristic with a random oracle \mathcal{H} . If Π is a public-coin protocol which satisfies $(T_{\text{Prove}^*}, T_{\mathcal{E}}, T_{\mathcal{A}}, \varepsilon^{w2e})$ -witness-extended emulation, assuming that there are m coins, then $\Pi^{\mathcal{H}}$ is $(T_{\text{Prove}^*}, T_{\text{Ext}}, T_{\mathcal{A}}, q_{\mathcal{H}}, \varepsilon^{\text{ext}})$ -extractable, with*

$$q_{\mathcal{H}} \leq T_{\text{Prove}^*} + T_{\mathcal{E}} + T_{\mathcal{A}}, \quad T_{\text{Ext}} \leq T_{\mathcal{E}} + q_{\mathcal{H}} \quad \text{and} \quad \varepsilon^{\text{ext}} \geq \varepsilon^{w2e} - \frac{T_{\mathcal{E}} \cdot m \cdot q_{\mathcal{H}}}{2^c},$$

where $c(\lambda)$ denotes the bit-length of the smallest public coin sent by algorithm Verf .

Proof. The proof follows readily from the definitions. Assuming that Π satisfies $(T_{\text{Prove}^*}, T_{\mathcal{E}}, T_{\mathcal{A}}, \varepsilon^{w2e})$ -witness-extended emulation, there exists an algorithm (emulator) \mathcal{E} running in expected time $T_{\mathcal{E}}$ as described above.

The extractor $(\text{Ext}_0, \text{Ext}_1)$ for $\Pi^{\mathcal{H}}$ can be simply constructed as follows. The first algorithm Ext_0 is the classical simulation for the random oracle: Ext_0 lazily samples a lookup table for the random oracle \mathcal{H} using a state Q . Each time it is queried on q , it checks whether $\mathcal{H}(q)$ is already defined. If this is the case, it returns the previously assigned value in Q , otherwise it returns and sets a fresh random value (of appropriate length).

The second algorithm Ext_1 consists simply in running \mathcal{E} using the random oracle \mathcal{H} to simulate oracle $\langle \text{Prove}^*(pp, x, s) \Rightarrow \text{Verf}(pp, x) \rangle$. More precisely, since Π is public-coin, the execution of $\text{Verf}(pp, x)$ corresponds to queries to Ext_0 . The simulation aborts only if a previous query to Ext_0 was made on the same (random) input. The number of such queries made by \mathcal{E} is clearly upper-bounded by $T_{\mathcal{E}}$ and for each coin in each of them, the probability that the simulation

aborts is upper-bounded by $q_{\mathcal{H}}2^{-c}$. The random oracle is then simply modified when Prove^* is rewound to any step of its computation and run anew on fresh verifier-randomness.

When \mathcal{E} eventually returns $((\tau, \beta), w)$, Ext_1 simply returns (β, w) (i.e., transcript τ is discarded). If Ext_1 does not abort, then by the simulation and the definition of \mathcal{E} ,

$$\left| \Pr \left[b = 1 : \begin{array}{l} pp \leftarrow \text{Setup}(1^\lambda) \\ (x, s) \leftarrow \mathcal{A}^{\mathcal{H}(\cdot)}(pp) \\ \pi \leftarrow \text{Prove}^{*\mathcal{H}(\cdot)}(pp, x, s) \\ b \leftarrow \text{Verf}^{\mathcal{H}}(pp, x, \pi) \end{array} \right] \right. \\ \left. - \Pr \left[b = 0 \vee (x, w) \in \mathcal{R} : \begin{array}{l} pp \leftarrow \text{Setup}(1^\lambda); Q \leftarrow \emptyset \\ (x, s) \leftarrow \mathcal{A}^{\text{Ext}_0(Q, \cdot)}(pp) \\ (b, w) \leftarrow \text{Ext}_1^{\text{Prove}^* \text{Ext}_0(Q, \cdot)}(pp, x, s)(pp, x) \end{array} \right] \right| \leq \varepsilon^{w2e},$$

and the claimed inequalities are satisfied. \square

$\Pi^{\mathcal{H}}$ is $(T_{\mathcal{A}}, q_{\mathcal{H}}, q_{\text{Prove}}, \varepsilon^{\text{zk}})$ -zero-knowledge if there exists an algorithm Sim such that $\text{Sim}(0, Q, q) \rightarrow (h, Q')$ and $\text{Sim}(1, Q, x) \rightarrow (\pi, Q')$, and such that for all $\lambda \in \mathbb{N}$, for every adversary \mathcal{A} that runs in time at most $T_{\mathcal{A}}(\lambda)$ and makes at most $q_{\mathcal{H}}$ queries to oracle \mathcal{H} and q_{Prove} queries to oracle Prove ,

$$\left| \Pr [b = 1 : \begin{array}{l} pp \leftarrow \text{Setup}(1^\lambda) \\ b \leftarrow \mathcal{A}^{\mathcal{H}(\cdot), \text{Prove}^{\mathcal{H}}(pp, \cdot)}(pp) \end{array} \right] \\ - \Pr \left[b = 1 : \begin{array}{l} pp \leftarrow \text{Setup}(1^\lambda); Q \leftarrow \emptyset \\ b \leftarrow \mathcal{A}^{\mathcal{O}_{\text{Sim}_0}(Q, \cdot), \mathcal{O}_{\text{Sim}_1}(Q, \cdot)}(pp) \end{array} \right] \right| \leq \varepsilon^{\text{zk}}(\lambda),$$

with $\mathcal{O}_{\text{Sim}_0}$ an oracle that computes $(h, Q') \leftarrow \text{Sim}(0, Q, q)$ on input (Q, q) and returns h , and $\mathcal{O}_{\text{Sim}_1}$ an oracle that computes $(\pi, Q) \leftarrow \text{Sim}(1, Q, x, w)$ if $(x, w) \in \mathcal{R}$ and returns \perp if $(x, w) \notin \mathcal{R}$. Set Q can be considered as a state which stores all pairs (q, h) of queries and responses.

Whenever Π and $\Pi^{\mathcal{H}}$ are *not* mentioned in the same context and it is clear that it is the non-interactive variant which is considered (e.g., in Section 4.2), superscript \mathcal{H} is omitted.

A.3 Non-interactive Proof Systems

A non-interactive proof system Π for a language $\mathcal{L} = \mathcal{L}_{\text{crs}}$ (with corresponding relation \mathcal{R}) consists of an algorithm $\text{Setup}(1^\lambda) \rightarrow \text{crs}$ which returns a common reference string, an algorithm $\text{Prove}(\text{crs}, x, w) \rightarrow \pi$ which computes a proof on the input of a word x and of a witness w , and an algorithm $\text{Verf}(\text{crs}, x, \pi) \rightarrow b \in \{0, 1\}$ which returns a bit indicated whether the proof is considered valid.

Π is *complete* if for all $\lambda \in \mathbb{N}$, for all $\text{crs} \leftarrow \text{Setup}(1^\lambda)$, for all $(x, w) \in \mathcal{R}$, $\Pr [\text{Verf}(\text{crs}, x, \text{Prove}(\text{crs}, x, w)) = 1] = 1$.

Π satisfies $(T, \varepsilon^{\text{sound}})$ -*soundness* if for all $\lambda \in \mathbb{N}$, for every adversary \mathcal{A} running in time at most $T(\lambda)$,

$$\Pr \left[\text{Verf}(crs, x, \pi) = 1 \wedge x \notin \mathcal{L} : \begin{array}{l} crs \leftarrow \text{Setup}(1^\lambda) \\ (x, \pi) \leftarrow \mathcal{A}(crs) \end{array} \right] \leq \varepsilon^{\text{sound}}(\lambda).$$

Π is $(T_{\text{Ext}}, T_{\mathcal{A}}, \varepsilon^{\text{ext}})$ -*extractable* if there exists an algorithm $\text{TSetup}(1^\lambda) \rightarrow (crs, \tau)$ and an algorithm $\text{Ext}(crs, \tau, x, \pi) \rightarrow w$ running in time at most $T_{\text{Ext}}(\lambda)$ such that the distribution of the first component of TSetup is the same as that of Setup , and such that for all $\lambda \in \mathbb{N}$, for every adversary \mathcal{A} running in time at most $T_{\mathcal{A}}(\lambda)$,

$$\Pr \left[\text{Verf}(crs, x, \pi) = 1 \wedge (x, \text{Ext}(crs, \tau, x, \pi)) \notin \mathcal{R} : \begin{array}{l} (crs, \tau) \leftarrow \text{TSetup}(1^\lambda) \\ (x, \pi) \leftarrow \mathcal{A}(crs) \end{array} \right] \leq \varepsilon^{\text{ext}}(\lambda).$$

Π is $(T, \varepsilon^{\text{zk}})$ -*composable zero-knowledge* if there exist two algorithms $\text{TSetup}(1^\lambda) \rightarrow (crs, \tau)$ and $\text{Sim}(crs, \tau, x) \rightarrow \pi$ such that for all $\lambda \in \mathbb{N}$, for every adversary \mathcal{A} running in time at most $T(\lambda)$,

$$\left| \Pr[1 \leftarrow \mathcal{A}(crs) : crs \leftarrow \text{Setup}(1^\lambda)] - \Pr[1 \leftarrow \mathcal{A}(crs) : (crs, \tau) \leftarrow \text{TSetup}(1^\lambda)] \right| \leq \varepsilon^{\text{zk}}(\lambda)$$

and

$$\begin{aligned} \Pr[1 \leftarrow \mathcal{A}(st, crs, \pi) : & \begin{array}{l} (crs, \tau) \leftarrow \text{TSetup}(1^\lambda) \\ (st, x, w) \leftarrow \mathcal{A}(crs, \tau) \\ \pi \leftarrow \text{Prove}(crs, x, w) \end{array} \\ & = \Pr \left[1 \leftarrow \mathcal{A}(st, crs, \pi) : \begin{array}{l} (crs, \tau) \leftarrow \text{TSetup}(1^\lambda) \\ (st, x, w) \leftarrow \mathcal{A}(crs, \tau) \\ \pi \leftarrow \text{Sim}(crs, \tau, x) \end{array} \right]. \end{aligned}$$

Proof Systems for Circuit Satisfiability. Groth, Ostrovsky and Sahai designed [34, Section 4.2] a pairing-based non-interactive, extractable, composable zero-knowledge proof system for circuit satisfiability. The reader is referred to their paper for more details.

A.4 Randomness Extractors

For a discrete distribution X over a set S , its *min-entropy* is denoted $\mathbf{H}_\infty(X) = \min_{x \leftarrow S} \{-\log \Pr[X = x]\}$. A distribution X is called a κ -*source* if $\mathbf{H}_\infty(X) \geq \kappa$.

A *randomness extractor* is a function which, applied to the output of one or several discrete distributions (possibly together with an additional uniformly random seed), returns a value indistinguishable from a uniformly random one. It is well known that there does not exist good deterministic extractors for a single

random source (i.e. without the additional random seed). In this paper, we use *two-source extractors* defined as follows.

Let $\mathcal{H} = \{H_x: \{0, 1\}^n \times \{0, 1\}^m \rightarrow \{0, 1\}^d\}_{x \in \{0, 1\}^d}$ be a hash function family. \mathcal{H} is said to be a (κ, ε) -*two-source extractor* if for any pair of discrete distributions I, J over $\{0, 1\}^n$ such that $\max(H_\infty(I), H_\infty(J)) \geq \kappa$, for X uniformly random over $\{0, 1\}^d$ and U is uniformly random over $\{0, 1\}^m$, the distributions of $(X, H_X(I, J))$ and (X, U) are ε -close. For some parameter sets, there exist deterministic two-source extractors (i.e., with $d = 0$) (e.g. [39]). Using the random-oracle heuristic, two-source extractors with high security and efficient parameters can be efficiently constructed. Alternatively, one can also rely on the weaker notion of *universal computational extractor* [8].

A.5 Universal Computational Extractors

Bellare, Hoang and Keelveedhi introduced Universal Computational Extractors [8] (UCEs) as a standard-model security notion for keyed hash functions which relates them to random-oracles. Their definition features a two-stage adversary, i.e., a *source* \mathcal{S} and a *distinguisher* \mathcal{D} . The source is given access to an oracle which either computes the hash function or is a random oracle, whereas the distinguisher is only given the hash-function key and some leakage information from the source. \mathcal{D} is then suppose to guess whether the oracle computed a hash function or was a random oracle.

The formal definition of a UCE is actually given w.r.t. source classes instead of a single source. Formally, given a class of sources \mathcal{S} and maps $d: \mathbb{N} \rightarrow \mathbb{N}$ and $N, M: \mathbb{N} \rightarrow 2^{\mathbb{N}}$, a family $\mathcal{H} = \left\{ H_x \in \bigcup_{n \in N(\lambda), m \in M(\lambda)} (\{0, 1\}^m)^{\{0, 1\}^n \times \{1\}^m} \right\}_{x \in \{0, 1\}^{d(\lambda)}}$ of (variable-input-length and variable-output-length) functions is (T, q, ε) -UCE secure w.r.t. \mathcal{S} if for all $\lambda \in \mathbb{N}$, for every $T(\lambda)$ -time adversary $(\mathcal{S}, \mathcal{D})$ such that $\mathcal{S} \in \mathcal{S}$ and such \mathcal{S} makes at most q oracle queries,

$$\left| \Pr \left[\begin{array}{l} x \leftarrow_{\S} \{0, 1\}^{d(\lambda)}; Q \leftarrow \emptyset \\ b \leftarrow_{\S} \{0, 1\} \\ b = b': L \leftarrow \mathcal{S}^{O_b(x, Q, \cdot)}(1^\lambda) \\ b' \leftarrow \mathcal{D}(1^\lambda, x, L) \\ \text{return } (b, b') \end{array} \right] - 1/2 \right| \leq \varepsilon(\lambda),$$

with $O_b(x, Q, \cdot)$ an oracle which, on input $(a, 1^m)$ such that $a \in \bigcup_{n \in N(\lambda)} \{0, 1\}^n$ and $m \in M(\lambda)$, proceeds as follows:

- if there exists $h \in \{0, 1\}^m$ such that $(a, 1^m, h) \in Q$, return h
- else
 - * if $b = 1$, return $H_x(a, 1^m)$
 - * else, generate $h \leftarrow_{\S} \{0, 1\}^m$, add $(a, 1^m, h)$ to Q and return h .

As \mathcal{S} could simply return one of its query and the response from the oracle in the leakage information, restrictions must imposed on the source for the security

notion to be achievable. A classical requirement is then that \mathcal{S} should be a class of *unpredictable* sources. A source \mathcal{S} is unpredictable if it is computationally hard to determine its hash queries even given its leakage information L , in case it interacts with a random oracle. Formally, a source \mathcal{S} is (T, q, ε) -simply unpredictable if for any $\lambda \in \mathbb{N}$, for every adversary or *simple*¹¹ predictor \mathcal{P} running in time at most $T(\lambda)$ and making at most q oracle queries,

$$\Pr \left[\begin{array}{l} x \leftarrow_{\mathcal{S}} \{0, 1\}^{d(\lambda)}; Q \leftarrow \emptyset \\ L \leftarrow \mathcal{S}^{\mathcal{O}(x, Q, \cdot)}(1^\lambda) \\ Q' \leftarrow \mathcal{P}(1^\lambda, L) \\ \text{return } (Q, Q') \end{array} \right] \leq \varepsilon(\lambda),$$

with $\mathcal{O}(x, Q, \cdot)$ an oracle which, on input $(a, 1^m)$ such that $a \in \bigcup_{n \in N(\lambda)} \{0, 1\}^n$ and $m \in M(\lambda)$, returns h if there exists $h \in \{0, 1\}^m$ such that $(a, 1^m, h) \in Q$, and otherwise generates $h \leftarrow_{\mathcal{S}} \{0, 1\}^m$, adds $(a, 1^m, h)$ to Q and returns h .

A.6 Chernoff's Bound

The Chernoff bound gives bound on the tail distribution of sums of independent Bernoulli random variables.

Theorem A.1 ([42, Theorem 4.2]). *Let X_1, X_2, \dots, X_n be independent Bernoulli random variables such that, for $1 \leq i \leq n$, $\Pr[X_i = 1] = p_i$, with $0 < p_i < 1$. Then, for $X = \sum_{i=1}^n X_i$, $\mu = \mathbb{E}[X] = \sum_{i=1}^n p_i$, and any $0 < \delta \leq 1$,*

$$\Pr[X < (1 - \delta)\mu] < \exp(-\mu\delta^2/2).$$

B Protocol for Discrete-Logarithm Keys

This section instantiate the generic protocol of Figure 2 in the case of discrete-logarithm keys. For simplicity, we present a scheme based on Pedersen commitments and classical Schnorr-like zero-knowledge proofs.

Let \mathbb{G} be a group family generator, λ be an integer and $(\mathbb{G}, \ell, g) \leftarrow \mathbb{G}(\lambda)$. Let g_1, g_2 denote two generators of \mathbb{G} used for the Pedersen commitment scheme. The key-generation protocol for discrete-logarithm keys is then the following.

1. \mathcal{U} applies the random oracle \mathcal{H} twice to its randomness $r_{\mathcal{U}}$ to compute $r'_{\mathcal{U}} \leftarrow \mathcal{H}(0 \| r_{\mathcal{U}})$ and $\rho_{\mathcal{U}} \leftarrow \mathcal{H}(1 \| r_{\mathcal{U}})$, commits to $r'_{\mathcal{U}}$ (with randomness $\rho_{\mathcal{U}}$) using Pedersen's commitment, computes a proof $\pi_{\mathcal{C}}$ of knowledge of an opening to the commitment with a Schnorr-type proof $\Pi_{\mathcal{C}}$ in the random-oracle model (with an oracle $\mathcal{H}_{\mathcal{C}}$), and sends the resulting commitment to \mathcal{CA} .

¹¹ The predictor is qualified as simple as it is not given access to oracle \mathcal{O} . Bellare, Hoang and Keelveedhi proved [8, Lemma 4.3] that a source is unpredictable (i.e., with a predictor which is given oracle access) if and only if it is simply unpredictable.

2. CA sets $acc_{C\mathcal{A}} \leftarrow \text{TRUE}$, verifies the proof, and it is correct, sends its randomness $r_{C\mathcal{A}}$ to \mathcal{U} ; otherwise it returns \perp and sets $term_{C\mathcal{A}} \leftarrow \text{TRUE}$
3. \mathcal{U} extracts a seed from the joint randomness by computing

$$s = r'_{\mathcal{U}} + \mathcal{H}(r_{C\mathcal{A}}) \bmod \ell.$$

It then computes a non-interactive zero-knowledge proof that it correctly performed its computation, i.e., it computes, in the random oracle model (with a different oracle \mathcal{H}_{Π}), a proof

$$\pi \leftarrow \Pi.\text{Prove} \left\{ r'_{\mathcal{U}} : C = g_1^{r'_{\mathcal{U}}} g_2^{\rho_{\mathcal{U}}} \wedge pk = g_1^{r'_{\mathcal{U}}} \cdot g_1^{\mathcal{H}(r_{C\mathcal{A}})} \right\}$$

which is a standard Schnorr-like proof of representation. Since the only randomness available to \mathcal{U} is $r_{\mathcal{U}}$, it is also used to generate the randomness necessary to compute the zero-knowledge proof. For this reason, \mathcal{U} derives another seed $s' \leftarrow \mathcal{H}(0\|s)$ and then uses \mathcal{H} with a counter on s' .

After computing π , \mathcal{U} erases $r_{\mathcal{U}}$, s and all temporary variables used to compute π . Algorithm \mathcal{U} then sends $pk \leftarrow (g_1, g_1^s)$ and π to $C\mathcal{A}$, returns $(pk_{\mathcal{U}} \leftarrow pk, sk_{\mathcal{U}} \leftarrow s)$ and sets $term_{\mathcal{U}} \leftarrow acc_{\mathcal{U}} \leftarrow \text{TRUE}$

4. CA verifies π . If this verification did not succeed, it returns \perp , otherwise it returns $pk_{C\mathcal{A}} \leftarrow pk$. It sets $term_{C\mathcal{A}} \leftarrow \text{TRUE}$.




Fig. 6. Discrete-Logarithm Key-Generation Protocol with Verifiable Randomness.

This protocol is 1-correct as Schnorr proofs are. In the random-oracle model and under the discrete-logarithm assumption over \mathbb{G} , it also satisfies the indistinguishability security notion of Section 3.2 w.r.t. randomness sources with $\Omega(\lambda)$ min-entropy which are independent of the random-oracle.