

Compactly Hiding Linear Spans

Tightly Secure Constant-Size Simulation-Sound QA-NIZK Proofs and Applications

Benoît Libert¹, Thomas Peters², Marc Joye³, and Moti Yung⁴

¹ Ecole Normale Supérieure de Lyon (France)

² Ecole Normale Supérieure (France)

³ Technicolor (USA)

⁴ Google Inc. and Columbia University (USA)

Abstract. Quasi-adaptive non-interactive zero-knowledge (QA-NIZK) proofs is a recent paradigm, suggested by Jutla and Roy (ASIACRYPT '13), which is motivated by the Groth-Sahai seminal techniques for efficient non-interactive zero-knowledge (NIZK) proofs. In this paradigm, the common reference string may depend on specific language parameters, a fact that allows much shorter proofs in important cases. It even makes certain standard model applications competitive with the Fiat-Shamir heuristic in the Random Oracle idealization. Such QA-NIZK proofs were recently optimized to constant size by Jutla and Roy (CRYPTO '14) and Libert *et al.* (EUROCRYPT '14) for the important case of proving that a vector of group elements belongs to a linear subspace. While the QA-NIZK arguments of Libert *et al.* provide unbounded simulation-soundness and constant proof length, their simulation-soundness is only loosely related to the underlying assumption (with a gap proportional to the number of adversarial queries) and it is unknown how to alleviate this limitation without sacrificing efficiency. In this paper, we deal with the question of whether we can simultaneously optimize the proof size and the tightness of security reductions, allowing for important applications with tight security (which are typically quite lengthy) to be of shorter size. We resolve this question by designing a novel simulation-sound QA-NIZK argument showing that a vector $\mathbf{v} \in \mathbb{G}^n$ belongs to a subspace of rank $t < n$ using a constant number of group elements. Unlike previous short QA-NIZK proofs of such statements, the unbounded simulation-soundness of our system is nearly tightly related (i.e., the reduction only loses a factor proportional to the security parameter) to the standard Decision Linear assumption. To show simulation-soundness in the constrained context of tight reductions, we explicitly point at a technique—which may be of independent interest—of hiding the linear span of a vector defined by a signature (which is part of an OR proof). As an application, we design a public-key cryptosystem with almost tight CCA2-security in the multi-challenge, multi-user setting with improved length (asymptotically optimal for long messages). We also adapt our scheme to provide CCA security in the key-dependent message scenario

(KDM-CCA2) with ciphertext length reduced by 75% when compared to the best known tightly secure KDM-CCA2 system so far.

Keywords: Security tightness, constant-size QA-NIZK proofs, simulation soundness, CCA2 security.

1 Introduction

In this paper, we consider the problem of achieving (almost) tight security in short simulation-sound non-interactive zero-knowledge proofs and chosen-ciphertext-secure encryption. While tight security results are known in both cases [36,39], they incur quite long proofs and ciphertexts. A natural question is to develop tools and techniques to make them short and, in the process, develop deeper understanding of this highly constrained setting. As an answer in this direction, we describe space-efficient methods and constructions with almost tight security. For the specific problem of proving that a vector of group elements belongs to a linear subspace, our main result is the first constant-size NIZK arguments whose simulation-soundness tightly relates to a standard assumption.

TIGHT AND ALMOST TIGHT SECURITY. Any public-key system must rely on some hardness assumption. To provide concrete guarantees, the security proof should preferably give a tight reduction from a well-established assumption. Namely, a successful adversary should imply a probabilistic polynomial time (PPT) algorithm breaking the assumption with nearly the same advantage. Tightness matters because the loss in the reduction may necessitate the use of a larger (at times prohibitively larger) security parameter to counteract the loss. The importance of tightness was first advocated by Bellare and Rogaway [10] in the context of digital signatures 18 years ago. Since then, it received a continuous attention with a flurry of positive and negative results in the random oracle model [25,26,47,24,11,59,2,45] and in the standard model [59,40,14,6,41].

A highly challenging problem has been to obtain tight security under standard assumptions in the standard model. For many primitives, satisfactory solutions have remained elusive until very recently. Bellare, Boldyreva and Micali [7] raised the problem of constructing a chosen-ciphertext-secure public-key cryptosystem based on a standard assumption and whose exact security does not degrade with the number of users or the number of challenge ciphertexts. The first answer to this question was only given more than a decade later by Hofheinz and Jager [39] and it was more a feasibility result than a practical solution. In the context of identity-based encryption (IBE), Chen and Wee [23] designed the first “almost tightly” secure system —meaning that the degradation factor only depends on the security parameter λ , and not on the number q of adversarial queries— based on a simple assumption in the standard model,⁵ which resolved an 8-year-old open problem [61].

⁵ Using random oracles, Katz and Wang [47] previously gave a tightly secure variant of the Boneh-Franklin IBE [17].

NIZK PROOFS AND SIMULATION-SOUNDNESS. Non-interactive zero-knowledge proofs [15] are crucial tools used in the design of countless cryptographic protocols. In the standard model, truly efficient constructions remained lacking until the last decade, when Groth and Sahai [37] gave nearly practical non-interactive witness indistinguishable (NIWI) and zero-knowledge (NIZK) proof systems for a wide class of languages in groups endowed with a bilinear map. While quite powerful, their methods remain significantly more costly than the non-interactive proof heuristics enabled by the Fiat-Shamir paradigm [30] in the idealized random oracle model [9]. recently, Jutla and Roy [43] showed that important efficiency improvements are possible for *quasi-adaptive* NIZK (QA-NIZK) proofs, i.e., where the common reference string (CRS) may depend on the specific language for which proofs are being generated but a single CRS simulator works for the entire class of languages. For the specific task of proving that a vector of n group elements belongs to a linear subspace of rank t , Jutla and Roy [43] gave computationally sound QA-NIZK proofs of length $\Theta(n - t)$ where the Groth-Sahai (GS) techniques entail $\Theta(n + t)$ group elements per proof. They subsequently refined their techniques, reducing the proof’s length to a constant [44], regardless of the number of equations or the number of variables. Libert *et al.* [50] independently obtained similar improvements using different techniques. Other constructions were recently given by Abdalla *et al.* [1] and Kiltz and Wee [48] who gave a general methodology for building short QA-NIZK arguments.

The design of non-malleable protocols, primarily IND-CCA2-secure encryption schemes, at times appeals to NIZK proofs endowed with a property named *simulation-soundness* by Sahai [58]: informally, an adversary should remain unable to prove a false statement by itself, even with the help of an oracle generating simulated proofs for (possibly false) adversarially-chosen statements. Groth [36] and Camenisch *et al.* [19] extended the Groth-Sahai techniques so as to obtain simulation-sound NIZK proofs. Their techniques incur a substantial overhead due to the use of quadratic pairing product equations, OR proofs or IND-CCA2-secure encryption schemes. It was shown [46,52,42] that one-time simulation-soundness —where the adversary obtains only one simulated proof— is much cheaper to achieve than unbounded simulation-soundness (USS). When it comes to proving membership of linear subspaces, Libert, Peters, Joye and Yung [50] gave very efficient unbounded simulation-sound quasi-adaptive NIZK proofs which do not require quadratic pairing product equations or IND-CCA2-secure encryption. As in the improved solution of Kiltz and Wee [48], their USS QA-NIZK arguments have constant size, regardless of the dimensions of the considered subspace. Unfortunately, the simulation-soundness of their proof system does not tightly reduce to the underlying assumption. The multiplicative gap between the reduction’s probability of success and the adversary’s advantage depends on the number q of simulated proofs observed by the adversary. As a consequence, the results of [50,48] do not imply tight chosen-ciphertext security [39] in a scenario —first envisioned by Bellare, Boldyreva and Micali [7]— where the adversary obtains polynomially many challenge ciphertexts. As of now, USS proof systems based on OR proofs [36,39] are the only ones to enable

tight multi-challenge security and it is unclear how to render them as efficient as [50] for linear equations.

TIGHTNESS AND CHOSEN-CIPHERTEXT SECURITY. Bellare, Boldyreva and Micali [7] showed that, if a public-key cryptosystem is secure in the sense of the one-user, one-challenge security definition [57], it remains secure in a more realistic multi-user setting where the adversary obtains polynomially many challenge ciphertexts. Their reduction involves a loss of exact security which is proportional to the number of users *and* the number of challenge ciphertexts. They also showed that, in the Cramer-Shoup encryption scheme [28], the degradation factor only depends on the number of challenges per user. Hofheinz and Jager [39] used a tightly secure simulation-sound proof system to build the first encryption system whose IND-CCA2 security tightly reduces to a standard assumption in the multi-user, multi-challenge setting. Due to very large ciphertexts, their scheme was mostly a feasibility result and the same holds for the improved constructions of Abe *et al.* [5]. Until recently, the only known CCA2-secure encryption schemes with tight security in the multi-challenge, multi-user setting either relied on non-standard q -type assumptions [38] —where the number of input elements depends on the number of adversarial queries— or incurred long ciphertexts [39,5] comprised of hundreds of group elements (or both). One of the reasons is that solutions based on standard assumptions [39,5,51] build on simulation-sound proof systems relying on OR proofs. Libert *et al.* [51] gave an almost tightly IND-CCA2 system in the multi-challenge setting where, despite their use of OR proofs, ciphertexts only require 69 group elements under the Decision Linear assumption. Unfortunately, their result falls short of implying constant-size simulation-sound QA-NIZK proofs of linear subspace membership since each vector coordinate would require its own proof elements. In particular, the technique of [51] would result in long proofs made of $O(\lambda)$ group elements in the setting of key-dependent message CCA2 security, where $O(1)$ group elements per proof suffices [44, Section 6] if we accept a loose reduction.

Very recently, Hofheinz *et al.* [41] put forth an almost tightly secure IBE scheme in the multi-challenge, multi-instance scenario. While their result implies an almost tightly CCA2 secure public-key encryption scheme via the Canetti-Halevi-Katz paradigm [21], it relies on composite order groups. In [41], it was left as an open problem to apply the same technique under standard assumptions in the (notoriously much more efficient) prime order setting.

OUR CONTRIBUTIONS. As a core technical innovation, this paper presents short QA-NIZK proofs of linear subspace membership (motivated by those in [50,44]) where the unbounded simulation-soundness property can be *almost tightly* —in the terminology of Chen and Wee [23]— related to the standard Decision Linear (DLIN) assumption [16]. As in [23], the loss of concrete security only depends on the security parameter, and not on the number of simulated proofs obtained by the adversary, which solves a problem left open in [50]. Our construction only lengthens the QA-NIZK proofs of Libert *et al.* [50] by a factor of 2 and thus retains the constant proof length of [50], independently of the dimensions of

the subspace. In particular, it does not rely on an IND-CCA2-secure encryption scheme —which, in this context, would require a tightly secure CCA2 cryptosystem to begin with— and it does not even require quadratic equations.

Building on our QA-NIZK proofs and the Naor-Yung paradigm [56], we obtain a new public-key encryption scheme which is proved IND-CCA2-secure in the multi-challenge, multi-user setting under the Decision Linear assumption via an almost tight reduction. While the reduction is slightly looser than those of [39,5], our security bound does not depend on the number of users or the number of challenges, so that our scheme is as secure in the multi-challenge, multi-user scenario as in the single-challenge, single-user setting. Like [39,5], our construction features publicly recognizable well-formed ciphertexts, which makes it suitable for non-interactive threshold decryption. Moreover, our ciphertexts are much shorter than those of [39,5] as they only consist of 48 group elements under the DLIN assumption, whereas the most efficient construction based on the same assumption [51] entails 69 group elements per ciphertext.

Our constant-size proofs offer more dramatic savings when it comes to encrypting long messages without affecting the compatibility with zero-knowledge proofs. We can encrypt N group elements at once while retaining short proofs, which only takes $2N + 46$ group elements per ciphertext. The asymptotic expansion ratio of 2 —which is inherent to the Naor-Yung technique— is thus optimal. To our knowledge, all prior results on tight CCA2 security would incur $\Theta(N)$ elements per proof and thus a higher expansion rate in this situation. In turn, our encryption schemes imply tightly secure non-interactive universally composable (UC) commitments [27,20] with adaptive security in the erasure model. In particular, using the same design principle as previous UC commitments [53,31,43] based on CCA2-secure cryptosystems, our scheme for long messages allows committing to N group elements at once with a two-fold expansion rate.

Using our QA-NIZK proof system, we also construct an almost tightly secure encryption scheme with key-dependent message chosen-ciphertext security (KDM-CCA2) [12,18] —in the sense of [19]— with shorter ciphertexts. Analogously to the Jutla-Roy construction [44, Section 6], our system offers substantial savings w.r.t. [19] as it allows for constant-size proofs even though, due to the use of the Boneh *et al.* approach [18] to KDM security, the dimension of underlying vectors of group elements depends on the security parameter. Like the Jutla-Roy construction [44], our KDM-CCA2 system only lengthens the ciphertexts of its underlying KDM-CPA counterpart by a constant number of group elements. Unlike [44], however, the KDM-CCA2 security of our scheme is almost tightly related to the DLIN assumption. So far, the most efficient tightly KDM-CCA2 system was implied by the results of Hofheinz-Jager [39] and Abe *et al.* [5], which incur rather long proofs. Our QA-NIZK proofs yield ciphertexts that are about 75% shorter, as we show in the full version of the paper.

OUR TECHNIQUES. Our QA-NIZK arguments (as the construction in [50]) build on linearly homomorphic structure-preserving signatures (LHSPS) [49]. In [50], each proof of subspace membership is a Groth-Sahai NIWI proof of knowledge of a homomorphic signature on the vector \mathbf{v} whose membership is being proved.

The security analysis relies on the fact that, with some probability, all simulated proofs take place on a perfectly NIWI Groth-Sahai CRS while the adversary’s fake proof pertains to a perfectly binding CRS. Here, in order to do this without applying Waters’ partitioning method [61] to the CRS space as in [54], we let the prover generate a Groth-Sahai CRS $\mathbf{F} = (\mathbf{f}_1, \mathbf{f}_2, \mathbf{F})$ of its choice (a similar technique was used by Escala and Groth [29] in a different context), for vectors of group elements $\mathbf{f}_1, \mathbf{f}_2, \mathbf{F} \in \mathbb{G}^3$, and first prove that this CRS is perfectly binding (i.e., \mathbf{F} lives in $\text{span}\langle \mathbf{f}_1, \mathbf{f}_2 \rangle$). This seemingly additional “freedom” that we give the prover ends up allowing a stronger simulator (tight simulation-soundness).

Simulation-soundness is, in fact, obtained by having the prover demonstrate that either: (i) The prover’s CRS \mathbf{F} is perfectly binding; or (ii) The prover knows a signature which only the NIZK simulator would be able to compute using some simulation trapdoor. One key idea is that, since the latter OR proof involves a relatively short statement (namely, the membership of a two-dimensional subspace) which the adversary has no control on, it can be generated using a constant number of group elements and using *only* linear pairing product equations.

In order to efficiently prove the above OR statement, we leverage the algebraic properties of a variant of the Chen-Wee signature scheme [23], which was proved almost tightly secure under the DLIN assumption, recently proposed by Libert *et al.* [51]. In short, the real prover computes a pseudo-signature σ (without knowing the signing key) on the verification key of a one-time signature and uses the real witnesses to prove that \mathbf{F} is a perfectly binding CRS. In contrast, the simulator computes a real signature σ using the private key instead of the real witnesses. In order to make sure that simulated proofs will be indistinguishable from real proofs, we apply a technique —implicitly used in [51]— consisting of hiding the linear subspace from where a partially committed vector of group elements defined by the signature σ is chosen: while a pseudo-signature fits within a proper subspace of a linear space specified by the public key, real signatures live in the full linear space. A difference between our approach and the one of [51] is our non-modular and more involved use of the signature scheme, yet the technique we point at above may be useful elsewhere. Our QA-NIZK CRS actually contains the description of a linear subspace which mixes the public key components of the signature and vectors used to build the prover’s Groth-Sahai CRS \mathbf{F} . In order to implement the OR proof, our idea is to make sure that the only way to prove a non-perfectly-binding CRS \mathbf{F} is to compute the committed σ as a real signature for a legally modified public key. By “legally modified key,” we mean that some of its underlying private components may be scaled by an adversarially-chosen factor $x \in \mathbb{Z}_p$ as long as the adversary also outputs g^x . While we rely on an unusual security property of the signature which allows the adversary to tamper with the public key, this property can be proved under the standard DLIN assumption in the scheme of [51]. This unusual property is a crucial technique allowing us to prove the OR statement about the ephemeral CRS \mathbf{F} without using quadratic equations.

In turn, the simulation-soundness relies on the fact that, unless some security property of the signature of [51] is broken, the adversary still has to generate

its fake proof on a perfectly binding CRS. If this condition is satisfied, we can employ the arguments as in [50] to show that the reduction is able to extract a non-trivial homomorphic signature, thus breaking the DLIN assumption.

FULL VERSION. The full version of this paper is available as Cryptology ePrint Archive, Report 2015/242 at URL <http://eprint.iacr.org/2015/242>.

2 Background and Definitions

2.1 Hardness Assumptions

We consider groups $(\mathbb{G}, \mathbb{G}_T)$ of prime-order p endowed with a bilinear map $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$. In this setting, we rely on the standard Decision Linear assumption.

Definition 1 ([16]). *The Decision Linear Problem (DLIN) in \mathbb{G} , is to distinguish the distributions $(g^a, g^b, g^{ac}, g^{bd}, g^{c+d})$ and $(g^a, g^b, g^{ac}, g^{bd}, g^z)$, with $a, b, c, d \xleftarrow{R} \mathbb{Z}_p, z \xleftarrow{R} \mathbb{Z}_p$. The DLIN assumption asserts the intractability of DLIN for any PPT distinguisher.*

We also use the following problem, which is at least as hard as DLIN [22].

Definition 2. *The Simultaneous Double Pairing problem (SDP) in $(\mathbb{G}, \mathbb{G}_T)$ is, given group elements $(g_z, g_r, h_z, h_u) \in \mathbb{G}^4$, to find a non-trivial triple $(z, r, u) \in \mathbb{G}^3 \setminus \{(1_{\mathbb{G}}, 1_{\mathbb{G}}, 1_{\mathbb{G}})\}$ such that $e(z, g_z) \cdot e(r, g_r) = 1_{\mathbb{G}_T}$ and $e(z, h_z) \cdot e(u, h_u) = 1_{\mathbb{G}_T}$.*

2.2 Quasi-Adaptive NIZK Proofs and Simulation-Soundness

Quasi-Adaptive NIZK (QA-NIZK) proofs are NIZK proofs where the CRS is allowed to depend on the specific language for which proofs have to be generated. The CRS is divided into a fixed part Γ , produced by an algorithm K_0 , and a language-dependent part ψ . However, there should be a single simulator for the entire class of languages.

Let λ be a security parameter. For public parameters $\Gamma \leftarrow K_0(\lambda)$, let \mathcal{D}_Γ be a probability distribution over a collection of relations $\mathcal{R} = \{R_\rho\}$ parametrized by a string ρ with an associated language $\mathcal{L}_\rho = \{x \mid \exists w : R_\rho(x, w) = 1\}$.

We consider proof systems where the prover and the verifier both take a label lbl as additional input. For example, this label can be the message-carrying part of an ElGamal-like encryption. Formally, a tuple of algorithms (K_0, K_1, P, V) is a QA-NIZK proof system for \mathcal{R} if there exists a PPT simulator (S_1, S_2) such that, for any PPT adversaries $\mathcal{A}_1, \mathcal{A}_2$ and \mathcal{A}_3 , we have the following properties:

Quasi-Adaptive Completeness:

$$\begin{aligned} & \Pr[\Gamma \leftarrow K_0(\lambda); \rho \leftarrow \mathcal{D}_\Gamma; \psi \leftarrow K_1(\Gamma, \rho); \\ & \quad (x, w, \text{lbl}) \leftarrow \mathcal{A}_1(\Gamma, \psi, \rho); \pi \leftarrow P(\psi, x, w, \text{lbl}) : \\ & \quad V(\psi, x, \pi, \text{lbl}) = 1 \text{ if } R_\rho(x, w) = 1] = 1 . \end{aligned}$$

Quasi-Adaptive Soundness:

$$\Pr[\Gamma \leftarrow \mathsf{K}_0(\lambda); \rho \leftarrow \mathcal{D}_\Gamma; \psi \leftarrow \mathsf{K}_1(\Gamma, \rho); (x, \pi, \text{lbl}) \leftarrow \mathcal{A}_2(\Gamma, \psi, \rho) : \\ \mathsf{V}(\psi, x, \pi, \text{lbl}) = 1 \wedge \neg(\exists w : R_\rho(x, w) = 1)] \in \text{negl}(\lambda) .$$

Quasi-Adaptive Zero-Knowledge:

$$\Pr[\Gamma \leftarrow \mathsf{K}_0(\lambda); \rho \leftarrow \mathcal{D}_\Gamma; \psi \leftarrow \mathsf{K}_1(\Gamma, \rho) : \mathcal{A}_3^{\mathsf{P}(\psi, \dots)}(\Gamma, \psi, \rho) = 1] \\ \approx \Pr[\Gamma \leftarrow \mathsf{K}_0(\lambda); \rho \leftarrow \mathcal{D}_\Gamma; (\psi, \tau_{sim}) \leftarrow \mathsf{S}_1(\Gamma, \rho) : \\ \mathcal{A}_3^{\mathsf{S}(\psi, \tau_{sim}, \dots)}(\Gamma, \psi, \rho) = 1] ,$$

where

- $\mathsf{P}(\psi, \dots)$ emulates the actual prover. It takes as input (x, w) and lbl and outputs a proof π if $(x, w) \in R_\rho$. Otherwise, it outputs \perp .
- $\mathsf{S}(\psi, \tau_{sim}, \dots)$ is an oracle that takes as input (x, w) and lbl . It outputs a simulated proof $\mathsf{S}_2(\psi, \tau_{sim}, x, \text{lbl})$ if $(x, w) \in R_\rho$ and \perp if $(x, w) \notin R_\rho$.

We assume that the CRS ψ contains an encoding of ρ , which is thus available to V . The definition of Quasi-Adaptive Zero-Knowledge requires a single simulator for the entire family of relations \mathcal{R} .

The property called *simulation-soundness* [58] requires that the adversary remain unable to prove false statements even after having seen simulated proofs for potentially false statements. We consider the strongest form, called *unbounded simulation-soundness* (USS) as opposed to one-time simulation-soundness, where the adversary is allowed to see polynomially many simulated proofs.

In order to use QA-NIZK proofs in a modular manner without degrading the exact security of our constructions, we will require simulation-soundness to hold *even* if the adversary \mathcal{A}_4 has a trapdoor τ_m that allows deciding membership in the language \mathcal{L}_ρ . We thus assume that the algorithm \mathcal{D}_Γ outputs a language parameter ρ and a trapdoor τ_m that allows recognizing elements of \mathcal{L}_ρ . This trapdoor τ_m is revealed to \mathcal{A}_4 and should not help prove false statements.

Enhanced Unbounded Simulation-Soundness: For any PPT adversary \mathcal{A}_4 ,

$$\Pr[\Gamma \leftarrow \mathsf{K}_0(\lambda); (\rho, \tau_m) \leftarrow \mathcal{D}_\Gamma; (\psi, \tau_{sim}) \leftarrow \mathsf{S}_1(\Gamma, \rho); \\ (x, \pi, \text{lbl}) \leftarrow \mathcal{A}_4^{\mathsf{S}_2(\psi, \tau_{sim}, \dots)}(\Gamma, \psi, \rho, \tau_m) : \\ \mathsf{V}(\psi, x, \pi, \text{lbl}) = 1 \wedge \neg(\exists w : R_\rho(x, w) = 1) \wedge (x, \pi, \text{lbl}) \notin Q] \in \text{negl}(\lambda) ,$$

where the adversary is allowed unbounded access to an oracle $\mathsf{S}_2(\psi, \tau, \dots)$ that takes as input statement-label pairs (x, lbl) (where x may be outside \mathcal{L}_ρ) and outputs simulated proofs $\pi \leftarrow \mathsf{S}_2(\psi, \tau_{sim}, x, \text{lbl})$ before updating the set $Q = Q \cup \{(x, \pi, \text{lbl})\}$, which is initially empty.

The standard notion of soundness can be enhanced in a similar way, by handing the membership testing trapdoor τ_m to \mathcal{A}_2 . In the weaker notion of one-time simulation-soundness, only one query to the S_2 oracle is allowed.

In order to achieve tight security in the multi-user setting, we also consider a notion of unbounded simulation-soundness in the multi-CRS setting. Namely,

the adversary is given a set of μ reference strings $\{\psi_\kappa\}_{\kappa=1}^\mu$ for language parameters $\{\rho_\kappa\}_{\kappa=1}^\mu$ and should remain unable to break the soundness of one these after having seen multiple simulated proofs for each CRS ψ_κ . A standard argument shows that (enhanced) unbounded simulation-soundness in the multi CRS setting is implied by the same notion in the single CRS setting. However, the reduction is far from being tight as it loses a factor μ . In our construction, the random self-reducibility of the underlying hard problems fortunately allows avoiding this security loss in a simple and natural way.

Enhanced Unbounded Simulation-Soundness in the multi-CRS setting: For any PPT adversary \mathcal{A}_4 , we have

$$\begin{aligned} & \Pr[\Gamma \leftarrow \mathsf{K}_0(\lambda); \{\rho_\kappa, \tau_{m,\kappa}\}_{\kappa=1}^\mu \leftarrow \mathcal{D}_\Gamma; (\{\psi_\kappa, \tau_{sim,\kappa}\}_{\kappa=1}^\mu) \leftarrow \mathsf{S}_1(\Gamma, \{\rho_\kappa\}_{\kappa=1}^\mu); \\ & \quad (\kappa^*, x, \pi, \text{lbl}) \leftarrow \mathcal{A}_4^{\mathsf{S}_2(\{\psi_\kappa\}_{\kappa=1}^\mu, \{\tau_{sim,\kappa}\}_{\kappa=1}^\mu, \dots)}(\Gamma, \{\psi_\kappa, \rho_\kappa, \tau_{m,\kappa}\}_{\kappa=1}^\mu) : \\ & \mathsf{V}(\psi_{\kappa^*}, x, \pi, \text{lbl}) = 1 \wedge \neg(\exists w : R_{\rho_{\kappa^*}}(x, w) = 1) \wedge (\kappa^*, x, \pi, \text{lbl}) \notin Q] \in \text{negl}(\lambda) . \end{aligned}$$

Here, \mathcal{A}_4 has access to an oracle $\mathsf{S}_2(\{\psi_\kappa\}_{\kappa=1}^\mu, \{\tau_{sim,\kappa}\}_{\kappa=1}^\mu, \dots)$ that takes as input tuples (j, x, lbl) (where x may be outside \mathcal{L}_{ρ_j}) and outputs simulated proofs $\pi \leftarrow \mathsf{S}_2(\{\psi_\kappa\}_{\kappa=1}^\mu, \{\tau_{sim,\kappa}\}_{\kappa=1}^\mu, j, x, \text{lbl})$ for \mathcal{L}_{ρ_j} before updating the set $Q = Q \cup \{(j, x, \pi, \text{lbl})\}$, which is initially empty.

The standard notion of soundness extends to the multi-CRS setting in a similar way and it can be enhanced by giving $\{\psi_\kappa\}_{\kappa=1}^\mu$ and the membership trapdoors $\{\tau_{m,\kappa}\}_{\kappa=1}^\mu$ to the adversary. The definition of quasi-adaptive zero-knowledge readily extends as well, by having S_1 output $\{\psi_\kappa, \tau_{sim,\kappa}\}_{\kappa=1}^\mu$ while the oracle S and the simulator S_2 both take an additional index $j \in \{1, \dots, \mu\}$ as input.

2.3 Linearly Homomorphic Structure-Preserving Signatures

Structure-preserving signatures [4,3] are signature schemes where messages and public keys consist of elements in the group \mathbb{G} of a bilinear configuration $(\mathbb{G}, \mathbb{G}_T)$.

Libert *et al.* [49] considered structure-preserving with linear homomorphic properties (see the full version of the paper for formal definitions). This section reviews the one-time linearly homomorphic structure-preserving signature (LHSPS) of [49].

Keygen(λ, n): given a security parameter λ and the subspace dimension $n \in \mathbb{N}$, choose bilinear group $(\mathbb{G}, \mathbb{G}_T)$ of prime order $p > 2^\lambda$. Then, choose $g_z, g_r, h_z, h_u \xleftarrow{R} \mathbb{G}$. For $i = 1$ to n , choose $\chi_i, \gamma_i, \delta_i \xleftarrow{R} \mathbb{Z}_p$ and compute $g_i = g_z^{\chi_i} g_r^{\gamma_i}$, $h_i = h_z^{\chi_i} h_u^{\delta_i}$. The private key is $\text{sk} = \{(\chi_i, \gamma_i, \delta_i)\}_{i=1}^n$ and the public key is $\text{pk} = (g_z, g_r, h_z, h_u, \{(g_i, h_i)\}_{i=1}^n) \in \mathbb{G}^{2n+4}$.

Sign($\text{sk}, (M_1, \dots, M_n)$): to sign a vector $(M_1, \dots, M_n) \in \mathbb{G}^n$ using $\text{sk} = \{(\chi_i, \gamma_i, \delta_i)\}_{i=1}^n$, output $\sigma = (z, r, u) = (\prod_{i=1}^n M_i^{-\chi_i}, \prod_{i=1}^n M_i^{-\gamma_i}, \prod_{i=1}^n M_i^{-\delta_i})$.

SignDerive($\text{pk}, \{(\omega_i, \sigma^{(i)})\}_{i=1}^\ell$): given pk as well as ℓ tuples $(\omega_i, \sigma^{(i)})$, parse $\sigma^{(i)}$ as $\sigma^{(i)} = (z_i, r_i, u_i)$ for $i = 1$ to ℓ . Return the triple $\sigma = (z, r, u) \in \mathbb{G}^3$, where $z = \prod_{i=1}^\ell z_i^{\omega_i}$, $r = \prod_{i=1}^\ell r_i^{\omega_i}$, $u = \prod_{i=1}^\ell u_i^{\omega_i}$.

Verify($\mathbf{pk}, \sigma, (M_1, \dots, M_n)$): given $\sigma = (z, r, u) \in \mathbb{G}^3$ and (M_1, \dots, M_n) , return 1 if and only if $(M_1, \dots, M_n) \neq (1_{\mathbb{G}}, \dots, 1_{\mathbb{G}})$ and (z, r, u) satisfy

$$1_{\mathbb{G}_T} = e(g_z, z) \cdot e(g_r, r) \cdot \prod_{i=1}^n e(g_i, M_i) = e(h_z, z) \cdot e(h_u, u) \cdot \prod_{i=1}^n e(h_i, M_i) . \quad (1)$$

Our simulation-sound proof system will rely on the fact that the above scheme provides tight security under the DLIN assumption, as implicitly shown in [49].

3 Constant-Size QA-NIZK Proofs of Linear Subspace Membership with Tight Simulation-Soundness

At a high level, our proof system can be seen as a variant of the construction of Libert *et al.* [50] with several modifications allowing to tightly relate the simulation-soundness property to the DLIN assumption. The construction also uses the tightly signature scheme of [51].

3.1 Intuition

Like [50], we combine linearly homomorphic signatures and Groth-Sahai proofs for pairing product equations. Each QA-NIZK proof consists of a Groth-Sahai NIWI proof of knowledge of a homomorphic signature on the candidate vector⁶ \mathbf{v} . By making sure that all simulated proofs take place on a perfectly WI CRS, the simulator is guaranteed to leak little information about its simulation trapdoor, which is the private key of the homomorphic signature. At the same time, if the adversary’s proof involves a perfectly binding CRS, the reduction can extract a homomorphic signature that it would have been unable to compute and solve a DLIN instance. To implement this approach, the system of [50] uses Waters’ partitioning technique [61] in the fashion of [54], which inevitably [40] affects the concrete security by a factor proportional to the number q of queries.

Our first main modification is that we let the prover compute the Groth-Sahai NIWI proof on a CRS \mathbf{F} of his own and append a proof π_F that the chosen CRS is perfectly binding, which amounts to proving the membership of a two-dimensional linear subspace $\text{span}\langle \mathbf{f}_1, \mathbf{f}_2 \rangle$. At first, it appears that π_F has to be simulation-sound itself since, in all simulated proofs, the reduction must trick the adversary into believing that the ephemeral CRS \mathbf{F} is perfectly sound. Fortunately, the reduction only needs to do this for vectors of its choice —rather than adversarially chosen vectors— and this scenario can be accommodated by appropriately mixing the subspace of Groth-Sahai vectors $\mathbf{f}_1, \mathbf{f}_2 \in \mathbb{G}^3$ with the one in the public key of the signature scheme of [51].

⁶ At first, tight simulation-soundness may seem achievable via an OR proof showing the knowledge of either a homomorphic signature on \mathbf{v} or a digital signature on the verification key of a one-time signature. However, proving that a disjunction of pairing product equations [36] is satisfiable requires a proof length proportional to the number of pairings (which is linear in the dimension n here) in pairing product equations.

The NIWI proof of knowledge is thus generated for a Groth-Sahai CRS $\mathbf{F} = (\mathbf{f}_1, \mathbf{f}_2, \mathbf{F})$ where \mathbf{f}_1 and \mathbf{f}_2 are part of the global CRS but $\mathbf{F} \in \mathbb{G}^3$ is chosen by the prover and included in the proof. To prove that \mathbf{F} is a perfectly sound CRS, honest provers derive a homomorphic signature (Z, R, U) from the first $4L + 2$ rows of a matrix $\mathbf{M} \in \mathbb{G}^{(4L+5) \times (4L+6)}$ defined by the public key of the signature scheme and fixed vectors $\mathbf{f}_1, \mathbf{f}_2, \mathbf{f}_0 \in \mathbb{G}^3$. The first two rows allow deriving a signature on the honestly generated $\mathbf{F} = \mathbf{f}_1^{\mu_1} \cdot \mathbf{f}_2^{\mu_2}$ from publicly available homomorphic signatures on \mathbf{f}_1 and \mathbf{f}_2 . The next $4L$ rows are used to demonstrate the validity of a pseudo-signature $(\sigma_1, \sigma_2, \sigma_3) = (H(\mathbf{V}, \text{VK})^r \cdot H(\mathbf{W}, \text{VK})^s, f^r, h^s)$ on the verification key VK of a one-time signature. This allows the prover to derive a homomorphic signature (Z, R, U) that authenticates a specific vector $\boldsymbol{\sigma} \in \mathbb{G}^{(4L+6)}$ determined by \mathbf{F} and the pseudo-signature $(\sigma_1, \sigma_2, \sigma_3)$.

The proof of simulation-soundness uses a strategy where, with high probability, all simulated proofs will take place on a perfectly NIWI CRS $\mathbf{F} = (\mathbf{f}_1, \mathbf{f}_2, \mathbf{F})$ —where $\mathbf{F} \in \mathbb{G}^3$ is linearly independent of $(\mathbf{f}_1, \mathbf{f}_2)$ —whereas the adversary’s fake proof π^* will contain a vector $\mathbf{F}^* \in \mathbb{G}^3$ such that $\mathbf{F} = (\mathbf{f}_1, \mathbf{f}_2, \mathbf{F}^*)$ is an extractable CRS (namely, $\mathbf{F}^* \in \text{span}(\mathbf{f}_1, \mathbf{f}_2)$). In order to satisfy the above conditions, the key idea is to have each QA-NIZK proof demonstrate that either: (i) The vector \mathbf{F} contained in π satisfies $\mathbf{F} \in \text{span}(\mathbf{f}_1, \mathbf{f}_2)$; (ii) $(\sigma_1, \sigma_2, \sigma_3)$ is a real signature rather than a pseudo-signature. Since $\mathbf{F} \in \mathbb{G}^3$ is chosen by the simulator, we can prove this compound statement *without* resorting to quadratic equations, by appropriately mixing linear subspaces. In more details, using a perfectly NIWI CRS in all simulated proofs requires the reduction to introduce a dependency on the fixed $\mathbf{f}_0 \in \mathbb{G}^3$ in the vector \mathbf{F} which is included in the proof π . In turn, in order to obtain a valid homomorphic signature on the vector $\boldsymbol{\sigma} \in \mathbb{G}^{(4L+6)}$ determined by \mathbf{F} and $(\sigma_1, \sigma_2, \sigma_3)$, this forces the simulator to use the last row of the matrix \mathbf{M} which contains the vector $\mathbf{f}_0 \in \mathbb{G}^3$ and the public key components Ω_1, Ω_2 of the signature scheme in [51]. To satisfy the verification algorithm, the vector $\boldsymbol{\sigma}$ must contain $1_{\mathbb{G}}$ in the coordinates where Ω_1, Ω_2 are located in the last row of \mathbf{M} . In order to retain these $1_{\mathbb{G}}$ ’s at these places, the simulator must use two other rows of \mathbf{M} to cancel out the introduction of Ω_1, Ω_2 in $\boldsymbol{\sigma}$. Applying such a “correction” implies the capability of replacing the pseudo-signature $(\sigma_1, \sigma_2, \sigma_3, Z, R, U)$ by a pair $(\sigma, X = g^x)$, where $\sigma = (\sigma_1, \sigma_2, \sigma_3, Z, R, U)$ is a real signature for a possibly modified key.

In order to obtain a perfectly NIZK proof system, we need to unconditionally hide the actual subspace where $\boldsymbol{\sigma} \in \mathbb{G}^{(4L+6)}$ lives as well as the fact that $(\sigma_1, \sigma_2, \sigma_3)$ is a real signature in simulated proofs. To this end, we refrain from letting (σ_1, Z, R, U) appear in the clear and replace them by perfectly hiding commitments $\mathbf{C}_{\sigma_1}, \mathbf{C}_Z, \mathbf{C}_R, \mathbf{C}_U$ to the same values and a NIWI proof that (Z, R, U) is a valid homomorphic signature on the partially committed vector $\boldsymbol{\sigma}$. Using our technique, we only need to prove *linear* pairing product equations.

In a construction of nearly tightly CCA2-secure cryptosystem, Libert *et al.* [51] used a somewhat similar approach based on pseudo-signatures and consisting of hiding the subspace where a partially committed vector is chosen. However, besides falling short of providing constant-size QA-NIZK proofs of

subspace membership, the approach of [51] requires quadratic equations. In contrast, while we also relying on pseudo-signatures, our technique for compactly hiding the underlying linear span completely avoids quadratic equations. It further yields simulation-sound QA-NIZK arguments that is constant size fitting within 42 group elements, regardless of the dimensions of the subspace.

3.2 Construction

For simplicity, the description below assumes symmetric pairings $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$ but instantiations in asymmetric pairings $e : \mathbb{G} \times \hat{\mathbb{G}} \rightarrow \mathbb{G}_T$, with $\mathbb{G} \neq \hat{\mathbb{G}}$, are possible, as explained in the full version of the paper.

As in [43], we assume that the language parameter ρ is a matrix in $\mathbb{G}^{t \times n}$, for some integers $t, n \in \text{poly}(\lambda)$ such that $t < n$, with an underlying witness relation R_{par} such that, for any $\mathbf{A} \in \mathbb{Z}_p^{t \times n}$ and $\rho \in \mathbb{G}^{t \times n}$, $R_{\text{par}}(\mathbf{A}, \rho) = 1$ if and only if $\rho = g^{\mathbf{A}}$. We consider distributions $\mathcal{D}_\Gamma \subset \mathbb{G}^{t \times n}$ that are efficiently witness-samplable: namely, there is a PPT algorithm which outputs a pair (ρ, \mathbf{A}) such that $R_{\text{par}}(\mathbf{A}, \rho) = 1$ and describing a relation R_ρ with its associated language \mathcal{L}_ρ according to \mathcal{D}_Γ . For example, the sampling algorithm could pick a random matrix $\mathbf{A} \xleftarrow{R} \mathbb{Z}_p^{t \times n}$ and define $\rho = g^{\mathbf{A}}$.

K₀(λ): choose symmetric bilinear groups $(\mathbb{G}, \mathbb{G}_T)$ of prime order $p > 2^\lambda$ with $f, g, h \xleftarrow{R} \mathbb{G}$. Choose a strongly unforgeable one-time signature $\Sigma = (\mathcal{G}, \mathcal{S}, \mathcal{V})$ with verification keys consisting of L -bit strings, for a suitable $L \in \text{poly}(\lambda)$. Then, output $\Gamma = (\mathbb{G}, \mathbb{G}_T, f, g, h, \Sigma)$.

The dimensions (t, n) of the matrix $\mathbf{A} \in \mathbb{Z}_p^{t \times n}$ such that $\rho = g^{\mathbf{A}}$ can be part of the language, so that t, n can be given as input to algorithm \mathbf{K}_1 .

K₁(Γ, ρ): parse Γ as $(\mathbb{G}, \mathbb{G}_T, f, g, h, \Sigma)$ and ρ as $\rho = (G_{i,j})_{1 \leq i \leq t, 1 \leq j \leq n} \in \mathbb{G}^{t \times n}$.

1. Generate key pairs $\{(\text{sk}_b, \text{pk}_b)\}_{b=0}^1$ for the one-time homomorphic signature of Section 2.3 in order to sign vectors of \mathbb{G}^n and \mathbb{G}^{4L+6} , respectively. Namely, choose $g_z, g_r, h_z, h_u \xleftarrow{R} \mathbb{G}$, $G_z, G_r, H_z, H_u \xleftarrow{R} \mathbb{G}$. Then, for $i = 1$ to n , pick $\chi_i, \gamma_i, \delta_i \xleftarrow{R} \mathbb{Z}_p$ and compute $g_i = g_z^{\chi_i} g_r^{\gamma_i}$ and $h_i = h_z^{\chi_i} h_u^{\delta_i}$. Let $\text{sk}_0 = \{\chi_i, \gamma_i, \delta_i\}_{i=1}^n$ be the private key and let $\text{pk}_0 = (g_z, g_r, h_z, h_u, \{g_i, h_i\}_{i=1}^n)$ be the public key. The second LHSPS key pair $(\text{sk}_1, \text{pk}_1)$ is generated analogously as $\text{sk}_1 = \{\varphi_i, \phi_i, \vartheta_i\}_{i=1}^{4L+6}$ and

$$\text{pk}_1 = \left(G_z, G_r, H_z, H_u, \{G_i = G_z^{\varphi_i} G_r^{\phi_i}, H_i = H_z^{\varphi_i} H_u^{\vartheta_i}\}_{i=1}^{4L+6} \right).$$

2. Choose $y_1, y_2, \xi_1, \xi_2, \xi_3 \xleftarrow{R} \mathbb{Z}_p$ and compute $f_1 = g^{y_1}$, $f_2 = g^{y_2}$. Define vectors $\mathbf{f}_1 = (f_1, 1_{\mathbb{G}}, g)$, $\mathbf{f}_2 = (1_{\mathbb{G}}, f_2, g)$ and $\mathbf{f}_3 = \mathbf{f}_1^{\xi_1} \cdot \mathbf{f}_2^{\xi_2} \cdot \iota(g)^{\xi_3}$, where $\iota(g) = (1_{\mathbb{G}}, 1_{\mathbb{G}}, g)$. Define the Groth-Sahai CRS $\mathbf{f} = (\mathbf{f}_1, \mathbf{f}_2, \mathbf{f}_3)$. Then, define yet another vector $\mathbf{f}_0 = \mathbf{f}_1^{\nu_1} \cdot \mathbf{f}_2^{\nu_2}$, with $\nu_1, \nu_2 \xleftarrow{R} \mathbb{Z}_p$.
3. For $\ell = 1$ to L , choose $V_{\ell,0}, V_{\ell,1}, W_{\ell,0}, W_{\ell,1} \xleftarrow{R} \mathbb{G}$ and define row vectors $\mathbf{V} = (V_{1,0}, V_{1,1}, \dots, V_{L,0}, V_{L,1})$, $\mathbf{W} = (W_{1,0}, W_{1,1}, \dots, W_{L,0}, W_{L,1})$.

4. Choose random exponents $\omega_1, \omega_2 \xleftarrow{R} \mathbb{Z}_p$ and group elements $u_1, u_2 \xleftarrow{R} \mathbb{G}$, and compute $\Omega_1 = u_1^{\omega_1} \in \mathbb{G}$, $\Omega_2 = u_2^{\omega_2} \in \mathbb{G}$.
5. Define the matrix $\mathbf{M} = (M_{i,j})_{i,j} \in \mathbb{G}^{(4L+5) \times (4L+6)}$ as

$$(M_{i,j})_{i,j} = \begin{pmatrix} 1 & \mathbf{1}^{1 \times 2L} & \mathbf{1}^{1 \times 2L} & 1 & 1 & \mathbf{f}_1 \\ 1 & \mathbf{1}^{1 \times 2L} & \mathbf{1}^{1 \times 2L} & 1 & 1 & \mathbf{f}_2 \\ \mathbf{V}^\top & \mathbf{Id}_{f,2L} & \mathbf{1}^{2L \times 2L} & \mathbf{1}^{2L \times 1} & \mathbf{1}^{2L \times 1} & \mathbf{1}^{2L \times 3} \\ \mathbf{W}^\top & \mathbf{1}^{2L \times 2L} & \mathbf{Id}_{h,2L} & \mathbf{1}^{2L \times 1} & \mathbf{1}^{2L \times 1} & \mathbf{1}^{2L \times 3} \\ g & \mathbf{1}^{1 \times 2L} & \mathbf{1}^{1 \times 2L} & u_1 & 1 & \mathbf{1}^{1 \times 3} \\ g & \mathbf{1}^{1 \times 2L} & \mathbf{1}^{1 \times 2L} & 1 & u_2 & \mathbf{1}^{1 \times 3} \\ 1 & \mathbf{1}^{1 \times 2L} & \mathbf{1}^{1 \times 2L} & \Omega_1^{-1} & \Omega_2^{-1} & \mathbf{f}_0 \end{pmatrix} \quad (2)$$

with $\mathbf{Id}_{f,2L} = f^{\mathbf{1}^{2L}} \in \mathbb{G}^{2L \times 2L}$, $\mathbf{Id}_{h,2L} = h^{\mathbf{1}^{2L}} \in \mathbb{G}^{2L \times 2L}$, and where $\mathbf{I}_{2L} \in \mathbb{Z}_p^{2L \times 2L}$ stands for the identity matrix. Note that the last row allows linking \mathbf{f}_0 and Ω_1, Ω_2 .

6. Use sk_0 to generate one-time homomorphic signatures $\{(z_i, r_i, u_i)\}_{i=1}^t$ on the vectors $(G_{i1}, \dots, G_{in}) \in \mathbb{G}^n$ that form the rows of $\boldsymbol{\rho} \in \mathbb{G}^{t \times n}$. These are given by $(z_i, r_i, u_i) = (\prod_{j=1}^n G_{i,j}^{-\chi_j}, \prod_{j=1}^n G_{i,j}^{-\gamma_j}, \prod_{j=1}^n G_{i,j}^{-\delta_j})$ for each $i \in \{1, \dots, t\}$. Likewise, use sk_1 to sign the rows $\mathbf{M}_j = (M_{j,1}, \dots, M_{j,4L+6})$ of the matrix (2) and obtain signatures

$$(Z_j, R_j, U_j) = \left(\prod_{k=1}^{4L+6} M_{j,k}^{-\varphi_k}, \prod_{k=1}^{4L+6} M_{j,k}^{-\phi_k}, \prod_{k=1}^{4L+6} M_{j,k}^{-\vartheta_k} \right)$$

for each $j \in \{1, \dots, 4L+5\}$.

7. The CRS $\psi = (\mathbf{CRS}_1, \mathbf{CRS}_2)$ consists of two parts which are defined as

$$\begin{aligned} \mathbf{CRS}_1 &= \left(\boldsymbol{\rho}, \mathbf{f}, \mathbf{f}_0, u_1, u_2, \Omega_1, \Omega_2, \mathbf{V}, \mathbf{W}, \text{pk}_0, \text{pk}_1, \right. \\ &\quad \left. \{(z_i, r_i, u_i)\}_{i=1}^t, \{(Z_j, R_j, U_j)\}_{j=1}^{4L+5} \right), \\ \mathbf{CRS}_2 &= \left(\mathbf{f}, \mathbf{f}_0, \text{pk}_0, \text{pk}_1, \Omega_1, \Omega_2, \mathbf{V}, \mathbf{W} \right), \end{aligned}$$

while the simulation trapdoor is $\tau_{sim} = (\omega_1, \omega_2, \{\chi_i, \gamma_i, \delta_i\}_{i=1}^n)$.

P($\Gamma, \psi, \mathbf{v}, \mathbf{x}, \text{lbl}$): given $\mathbf{v} \in \mathbb{G}^n$ and a witness $\mathbf{x} = (x_1, \dots, x_t) \in \mathbb{Z}_p^t$ such that $\mathbf{v} = g^{\mathbf{x} \cdot \mathbf{A}}$, generate a one-time signature key pair $(\text{VK}, \text{SK}) \leftarrow \mathcal{G}(\lambda)$.

1. Using $\{(z_j, r_j, u_j)\}_{j=1}^t$, derive a one-time linearly homomorphic signature (z, r, u) on the vector \mathbf{v} with respect to pk_0 . Namely, compute $z = \prod_{i=1}^t z_i^{x_i}$, $r = \prod_{i=1}^t r_i^{x_i}$ and $u = \prod_{i=1}^t u_i^{x_i}$.
2. Choose a vector $\mathbf{F} = (F_1, F_2, F_3) = \mathbf{f}_1^{\mu_1} \cdot \mathbf{f}_2^{\mu_2}$, for random $\mu_1, \mu_2 \xleftarrow{R} \mathbb{Z}_p$.
3. Pick $r, s \xleftarrow{R} \mathbb{Z}_p$ and compute a pseudo-signature on $\text{VK} = \text{VK}[1] \dots \text{VK}[L]$, which is obtained as $(\sigma_1, \sigma_2, \sigma_3) = (H(\mathbf{V}, \text{VK})^r \cdot H(\mathbf{W}, \text{VK})^s, f^r, h^s)$, where $H(\mathbf{V}, \text{VK}) = \prod_{\ell=1}^L V_{\ell, \text{VK}[\ell]}$ and $H(\mathbf{W}, \text{VK}) = \prod_{\ell=1}^L W_{\ell, \text{VK}[\ell]}$.

4. Derive a one-time linearly homomorphic signature $(Z, R, U) \in \mathbb{G}^3$ for pk_1 on the vector

$$\boldsymbol{\sigma} = (\sigma_1, \sigma_2^{1-\text{VK}[1]}, \sigma_2^{\text{VK}[1]}, \dots, \sigma_2^{1-\text{VK}[L]}, \sigma_2^{\text{VK}[L]}, \sigma_3^{1-\text{VK}[1]}, \sigma_3^{\text{VK}[1]}, \dots, \sigma_3^{1-\text{VK}[L]}, \sigma_3^{\text{VK}[L]}, 1_{\mathbb{G}}, 1_{\mathbb{G}}, F_1, F_2, F_3) \in \mathbb{G}^{4L+6} \quad (3)$$

which belongs to subspace spanned by the first $4L + 2$ rows of the matrix $\mathbf{M} \in \mathbb{G}^{(4L+5) \times (4L+6)}$. Hence, the coefficients $r, s, \mu_1, \mu_2 \in \mathbb{Z}_p$ allow deriving a homomorphic signature (Z, R, U) on $\boldsymbol{\sigma}$ in (3). Note that the $(4L + 2)$ -th and the $(4L + 3)$ -th coordinates of $\boldsymbol{\sigma}$ must both equal $1_{\mathbb{G}}$.

5. Using the CRS $\mathbf{f} = (\mathbf{f}_1, \mathbf{f}_2, \mathbf{f}_3)$, generate Groth-Sahai commitments $\mathbf{C}_{\sigma_1}, \mathbf{C}_Z, \mathbf{C}_R, \mathbf{C}_U \in \mathbb{G}^3$. Then, compute NIWI proofs $\boldsymbol{\pi}_{\sigma_1}, \boldsymbol{\pi}_{\sigma_2} \in \mathbb{G}^3$ that committed variables (σ_1, Z, R, U) satisfy

$$\begin{aligned} e(Z, G_z) \cdot e(R, G_r) \cdot e(\sigma_1, G_1) &= t_G, \\ e(Z, H_z) \cdot e(U, H_u) \cdot e(\sigma_1, H_1) &= t_H, \end{aligned} \quad (4)$$

where

$$t_G = e(\sigma_2, \prod_{i=1}^L G_{2i+\text{VK}[i]})^{-1} \cdot e(\sigma_3, \prod_{i=1}^L G_{2L+2i+\text{VK}[i]})^{-1} \cdot \prod_{i=1}^3 e(F_i, G_{4L+3+i})^{-1}$$

and

$$t_H = e(\sigma_2, \prod_{i=1}^L H_{2i+\text{VK}[i]})^{-1} \cdot e(\sigma_3, \prod_{i=1}^L H_{2L+2i+\text{VK}[i]})^{-1} \cdot \prod_{i=1}^3 e(F_i, H_{4L+3+i})^{-1}.$$

6. Using the vector $\mathbf{F} = (F_1, F_2, F_3)$ of Step 2, define a new Groth-Sahai CRS $\mathbf{F} = (\mathbf{f}_1, \mathbf{f}_2, \mathbf{F})$ and use it to compute commitments

$$\begin{aligned} \mathbf{C}_z &= \iota(z) \cdot \mathbf{f}_1^{\theta_{z,1}} \cdot \mathbf{f}_2^{\theta_{z,2}} \cdot \mathbf{F}^{\theta_{z,3}}, & \mathbf{C}_r &= \iota(r) \cdot \mathbf{f}_1^{\theta_{r,1}} \cdot \mathbf{f}_2^{\theta_{r,2}} \cdot \mathbf{F}^{\theta_{r,3}}, \\ \mathbf{C}_u &= \iota(u) \cdot \mathbf{f}_1^{\theta_{u,1}} \cdot \mathbf{f}_2^{\theta_{u,2}} \cdot \mathbf{F}^{\theta_{u,3}} \end{aligned}$$

to the components of (z, r, u) along with NIWI proofs $(\boldsymbol{\pi}_1, \boldsymbol{\pi}_2) \in \mathbb{G}^6$ that \mathbf{v} and (z, r, u) satisfy (1). Let $(\mathbf{C}_z, \mathbf{C}_r, \mathbf{C}_u, \boldsymbol{\pi}_1, \boldsymbol{\pi}_2) \in \mathbb{G}^{15}$ be the resulting commitments and proofs.

7. Set $\sigma = \mathcal{S}(\text{SK}, (\mathbf{v}, \mathbf{F}, \mathbf{C}_{\sigma_1}, \sigma_2, \sigma_3, \mathbf{C}_Z, \mathbf{C}_R, \mathbf{C}_U, \mathbf{C}_z, \mathbf{C}_r, \mathbf{C}_u, \boldsymbol{\pi}_{\sigma_1}, \boldsymbol{\pi}_{\sigma_2}, \boldsymbol{\pi}_1, \boldsymbol{\pi}_2, \text{lbl}))$ and output

$$\pi = (\text{VK}, \mathbf{F}, \mathbf{C}_{\sigma_1}, \sigma_2, \sigma_3, \mathbf{C}_Z, \mathbf{C}_R, \mathbf{C}_U, \mathbf{C}_z, \mathbf{C}_r, \mathbf{C}_u, \boldsymbol{\pi}_{\sigma_1}, \boldsymbol{\pi}_{\sigma_2}, \boldsymbol{\pi}_1, \boldsymbol{\pi}_2, \sigma) \quad (5)$$

V($\Gamma, \psi, \mathbf{v}, \boldsymbol{\pi}, \text{lbl}$): parse π as in (5) and \mathbf{v} as $(v_1, \dots, v_n) \in \mathbb{G}^n$. Return 1 if the conditions hereunder all hold. Otherwise, return 0.

- (i) $\mathcal{V}(\text{VK}, (\mathbf{v}, \mathbf{F}, \mathbf{C}_{\sigma_1}, \sigma_2, \sigma_3, \mathbf{C}_Z, \mathbf{C}_R, \mathbf{C}_U, \mathbf{C}_z, \mathbf{C}_r, \mathbf{C}_u, \boldsymbol{\pi}_{\sigma_1}, \boldsymbol{\pi}_{\sigma_2}, \boldsymbol{\pi}_1, \boldsymbol{\pi}_2, \text{lbl}), \sigma) = 1$;
- (ii) $\boldsymbol{\pi}_{\sigma_1}, \boldsymbol{\pi}_{\sigma_2}$ are valid proofs that the variables (σ_1, Z, R, U) , which are contained in commitments $\mathbf{C}_{\sigma_1}, \mathbf{C}_Z, \mathbf{C}_R, \mathbf{C}_U$, satisfy equations (4).

- (iii) The tuple $(\mathbf{C}_z, \mathbf{C}_r, \mathbf{C}_u, \boldsymbol{\pi}_1, \boldsymbol{\pi}_2)$ forms a valid a valid NIWI proof for the Groth-Sahai CRS $\mathbf{F} = (\mathbf{f}_1, \mathbf{f}_2, \mathbf{F})$. Namely, $\boldsymbol{\pi}_1 = (\pi_{1,1}, \pi_{1,2}, \pi_{1,3})$ and $\boldsymbol{\pi}_2 = (\pi_{2,1}, \pi_{2,2}, \pi_{2,3})$ satisfy

$$\begin{aligned} \prod_{i=1}^n E(g_i, \iota(v_i))^{-1} &= E(g_z, \mathbf{C}_z) \cdot E(g_r, \mathbf{C}_r) \cdot E(\pi_{1,1}, \mathbf{f}_1) \cdot \\ &\quad E(\pi_{1,2}, \mathbf{f}_2) \cdot E(\pi_{1,3}, \mathbf{F}) \\ \prod_{i=1}^n E(h_i, \iota(v_i))^{-1} &= E(h_z, \mathbf{C}_z) \cdot E(h_u, \mathbf{C}_u) \cdot E(\pi_{2,1}, \mathbf{f}_1) \cdot \\ &\quad E(\pi_{2,2}, \mathbf{f}_2) \cdot E(\pi_{2,3}, \mathbf{F}) . \end{aligned} \tag{6}$$

The proof only requires 38 elements of \mathbb{G} and a pair (\mathbf{VK}, σ) . In instantiations using the one-time signature of [39], its total size amounts to 42 group elements, which only lengthens the QA-NIZK proofs of [50] by a factor of 2.

4 Security

To avoid unnecessarily overloading notations, we will prove our results in the single CRS setting. At the main steps, we will explain how the proof can be adapted to the multi-CRS setting without affecting the tightness of reductions.

Theorem 1. *The above proof system is perfectly quasi-adaptive zero-knowledge.*

Proof (sketch). We describe the QA-NIZK simulator here but we refer to the full paper for a detailed proof that the simulation is perfect. This simulator $(\mathbf{S}_1, \mathbf{S}_2)$ is defined by having \mathbf{S}_1 generate the CRS ψ as in the real \mathbf{K}_0 algorithm but retain the simulation trapdoor $\tau_{sim} = (\omega_1, \omega_2, \{\chi_i, \gamma_i, \delta_i\}_{i=1}^n)$ for later use. As for \mathbf{S}_2 , it generates a simulated proof for $\mathbf{v} = (v_1, \dots, v_n) \in \mathbb{G}^n$ by using $\{\{\chi_i, \gamma_i, \delta_i\}_{i=1}^n\}$ to compute $(z, r, u) = (\prod_{j=1}^n v_j^{-\chi_j}, \prod_{j=1}^n v_j^{-\gamma_j}, \prod_{j=1}^n v_j^{-\delta_j})$ at step 1 of the simulation instead of using the witness $\mathbf{x} \in \mathbb{Z}_p^t$ as in the real proving algorithm \mathbf{P} . At step 2, it defines $(F_1, F_2, F_3) = \mathbf{f}_0 \cdot \mathbf{f}_1^{\mu_1} \cdot \mathbf{f}_2^{\mu_2}$ with $\mu_1, \mu_2 \xleftarrow{R} \mathbb{Z}_p$. At step 3, it picks $r, s \xleftarrow{R} \mathbb{Z}_p$ to compute $(\sigma_1, \sigma_2, \sigma_3) = (g^{\omega_1 + \omega_2} \cdot H(\mathbf{V}, \mathbf{VK})^r \cdot H(\mathbf{W}, \mathbf{VK})^s, f^r, h^s)$ before using the coefficients $\mu_1, \mu_2, r, s, \omega_1, \omega_2, 1 \in \mathbb{Z}_p$ to derive a homomorphic signature (Z, R, U) from $\{(Z_j, R_j, U_j)\}_{j=1}^{4L+5}$ at step 4. Steps 5 to 7 are conducted as in the real \mathbf{P} .

In the full paper, we prove that the simulation is perfect in that the simulated CRS ψ is distributed as a real CRS and, for all $\mathbf{v} \in \mathbb{G}^n$ such that $\mathbf{v} = g^{\mathbf{x} \cdot \mathbf{A}}$ for some $\mathbf{x} \in \mathbb{Z}_p^t$, simulated proofs are distributed as real proofs. \square

We now prove that the system remains computationally sound and simulation-sound, even when the adversary is given the matrix $\mathbf{A} = \log_g(\boldsymbol{\rho}) \in \mathbb{Z}_p^{t \times n}$, which allows recognizing elements of \mathcal{L}_ρ . Although the enhanced soundness property is implied by that of enhanced simulation-soundness, we prove it separately (see the full paper for the proof) in Theorem 2 since the reduction is optimal.

Theorem 2. *The system provides quasi-adaptive soundness under the DLIN assumption. Any enhanced soundness adversary \mathcal{A} with running time $t_{\mathcal{A}}$ implies a DLIN distinguisher \mathcal{B} with running time $t_{\mathcal{B}} \leq t_{\mathcal{A}} + q \cdot \text{poly}(\lambda, L, t, n)$ and such that $\text{Adv}_{\mathcal{A}}^{\text{e-sound}}(\lambda) \leq 2 \cdot \text{Adv}_{\mathcal{B}}^{\text{DLIN}}(\lambda) + 2/p$.*

Theorem 3. *The above system provides quasi-adaptive unbounded simulation-soundness if: (i) Σ is a strongly unforgeable one-time signature; (ii) The DLIN assumption holds. For any enhanced unbounded simulation-soundness adversary \mathcal{A} , there exist a one-time signature forger \mathcal{B}' in the multi-key setting and a DLIN distinguisher \mathcal{B} with running times $t_{\mathcal{B}}, t_{\mathcal{B}'} \leq t_{\mathcal{A}} + q \cdot \text{poly}(\lambda, L, t, n)$ such that*

$$\text{Adv}_{\mathcal{A}}^{\text{e-uss}}(\lambda) \leq \text{Adv}_{\mathcal{B}'}^{q\text{-suf-ots}}(\lambda) + 3 \cdot (L + 2) \cdot \text{Adv}_{\mathcal{B}}^{\text{DLIN}}(\lambda) + 4/p, \quad (7)$$

where L is the verification key length of Σ and q is the number of simulations.

Proof. To prove the result, we consider a sequence of games. In Game_i , we denote by S_i the event that the challenger outputs 1.

Game₁: This game is the actual attack. Namely, the adversary \mathcal{A} receives as input the description of the language \mathcal{L}_{ρ} and has access to a simulated CRS ψ and the simulated prover $S_2(\psi, \tau_{sim}, \dots)$ which is described in the proof of Theorem 1. At each invocation, $S_2(\psi, \tau_{sim}, \dots)$ inputs a vector-label pair (\mathbf{v}, lbl) and outputs a simulated proof π that $\mathbf{v} \in \mathcal{L}_{\rho}$. In order to generate the matrix $\rho \in \mathbb{G}^{t \times n}$ with the appropriate distribution D_{Γ} , the challenger chooses a matrix $\mathbf{A} \in \mathbb{Z}_p^{t \times n}$ with the suitable distribution (which is possible since D_{Γ} is efficiently witness-samplable) and computes $\rho = g^{\mathbf{A}}$. Also, the challenger \mathcal{B} computes a basis $\mathbf{W} \in \mathbb{Z}_p^{n \times (n-t)}$ of the nullspace of \mathbf{A} . The adversary receives as input the simulated CRS ψ and the matrix $\mathbf{A} \in \mathbb{Z}_p^{t \times n}$, which serves as a membership testing trapdoor τ_m , and queries the simulator $S_2(\psi, \tau_{sim}, \dots)$ on a polynomial number of occasions. When the adversary \mathcal{A} halts, it outputs an element \mathbf{v}^* , a proof π^* and a label lbl^* . The adversary is declared successful and the challenger outputs 1 if and only if (π^*, lbl^*) is a verifying proof but $\mathbf{v}^* \notin \mathcal{L}_{\rho}$ (i.e., \mathbf{v}^* is linearly independent of the rows of $\rho \in \mathbb{G}^{t \times n}$) and (π^*, lbl^*) was not trivially obtained from the simulator. We call S_1 the latter event, which is easily recognizable by the challenger \mathcal{B} since the latter knows a basis $\mathbf{W} \in \mathbb{Z}_p^{n \times (n-t)}$ of the right kernel of \mathbf{A} . Indeed, \mathbf{W} allows testing if $\mathbf{v} = (v_1, \dots, v_n) \in \mathbb{G}^n$ satisfies $\prod_{j=1}^n v_j^{w_{ji}} = 1_{\mathbb{G}}$ for each column $\mathbf{w}_i^{\top} = (w_{1i}, \dots, w_{ni})^{\top}$ of \mathbf{W} . By definition, the adversary's advantage is $\text{Adv}(\mathcal{A}) := \Pr[S_1]$.

Game₂: We modify the generation of the CRS $\psi = (\text{CRS}_1, \text{CRS}_2)$. Instead of choosing $\mathbf{f}_3 \in_R \mathbb{G}^3$ as a uniformly random vector, S_1 sets $\mathbf{f}_3 = \mathbf{f}_1^{\xi_1} \cdot \mathbf{f}_2^{\xi_2}$, for random $\xi_1, \xi_2 \xleftarrow{R} \mathbb{Z}_p$. Hence, $\mathbf{f}_1, \mathbf{f}_2$ and \mathbf{f}_3 now underlie a subspace of dimension 2 and $\mathbf{f} = (\mathbf{f}_1, \mathbf{f}_2, \mathbf{f}_3)$ thus becomes a perfectly binding CRS. Under the DLIN assumption, this modification should have no noticeable impact on \mathcal{A} 's probability of success. We have $|\Pr[S_2] - \Pr[S_1]| \leq \text{Adv}^{\text{DLIN}}(\mathcal{B})$.

Game₃: We modify again the generation of ψ . Now, instead of choosing \mathbf{f}_0 in $\text{span}\langle \mathbf{f}_1, \mathbf{f}_2 \rangle$, S_1 sets $\mathbf{f}_0 = \mathbf{f}_1^{\nu_1} \cdot \mathbf{f}_2^{\nu_2} \cdot \iota(g)$, for random $\nu_1, \nu_2 \xleftarrow{R} \mathbb{Z}_p^*$. The vector

\mathbf{f}_0 is now linearly independent of $(\mathbf{f}_1, \mathbf{f}_2)$. Under the DLIN assumption, this modification will remain unnoticed to the adversary. In particular, \mathcal{A} 's winning probability should only change by a negligible amount. A two-step reduction from DLIN shows that $|\Pr[S_3] - \Pr[S_2]| \leq 2 \cdot \mathbf{Adv}^{\text{DLIN}}(\mathcal{B})$.

Game₄: This game is like **Game₃** but \mathcal{B} halts and outputs a random bit if \mathcal{A} outputs a proof π^* containing a one-time verification key \mathbf{VK}^* that is recycled from an output of the $\mathcal{S}_2(\psi, \tau_{sim}, \cdot, \cdot)$ oracle. **Game₄** and **Game₃** proceed identically until the latter event occurs. This event further contradicts the strong unforgeability of Σ . If Σ has tight multi-key security⁷ (in the sense of [39]), the probability of this event can be bounded independently of the number q of queries to $\mathcal{S}_2(\psi, \tau_{sim}, \cdot, \cdot)$. We have $|\Pr[S_4] - \Pr[S_3]| \leq \mathbf{Adv}_B^{q\text{-suf-ots}}(\lambda)$.

Game₅: This game is identical to **Game₄** but we raise a failure event E_5 . When \mathcal{A} outputs its fake proof $\pi^* = (\mathbf{VK}^*, \mathbf{F}^*, \mathbf{C}_{\sigma_1}^*, \sigma_2^*, \sigma_3^*, \mathbf{C}_Z^*, \mathbf{C}_R^*, \mathbf{C}_U^*, \mathbf{C}_z^*, \mathbf{C}_r^*, \mathbf{C}_u^*, \pi_{\sigma,1}^*, \pi_{\sigma,2}^*, \pi_1^*, \pi_2^*, \sigma^*)$, \mathcal{B} parses the vector \mathbf{F}^* as $(F_1^*, F_2^*, F_3^*) \in \mathbb{G}^3$ and uses the extraction trapdoor $(y_1, y_2) = (\log_g(f_1), \log_g(f_2))$ of the Groth-Sahai CRS $\mathbf{f} = (\mathbf{f}_1, \mathbf{f}_2, \mathbf{f}_3)$ to test if $F_3^* \neq F_1^{*1/y_1} \cdot F_2^{*1/y_2}$, meaning that $\mathbf{F}^* = (\mathbf{f}_1, \mathbf{f}_2, \mathbf{F}^*)$ is not a perfectly binding Groth-Sahai CRS. We denote by E_5 the latter event, which causes \mathcal{B} to abort and output a random bit if it occurs. Clearly, **Game₅** is identical to **Game₄** unless E_5 occurs, so that $|\Pr[S_5] - \Pr[S_4]| \leq \Pr[E_5]$. Lemma 1 demonstrates that event E_5 occurs with negligible probability if the DLIN assumption holds. More precisely, the probability $\Pr[E_5]$ is at most $\Pr[E_5] \leq (2 \cdot L + 1) \cdot \mathbf{Adv}_B^{\text{DLIN}}(\lambda) + 2/p$, where \mathcal{B} is a DLIN distinguisher whose computational complexity only exceeds that of \mathcal{A} by the cost of a polynomial number of exponentiations in \mathbb{G} and a constant number of pairing evaluations.

In **Game₅**, we have $\Pr[S_5] = \Pr[S_5 \wedge E_5] + \Pr[S_5 \wedge \neg E_5] = \frac{1}{2} \cdot \Pr[E_5] + \Pr[S_5 \wedge \neg E_5]$, so that $\Pr[S_5] \leq (L + 1) \cdot \mathbf{Adv}_B^{\text{DLIN}}(\lambda) + \frac{1}{p} + \Pr[S_5 \wedge \neg E_5]$.

In **Game₅**, we show that event $S_5 \wedge \neg E_5$ implies an algorithm \mathcal{B} solving a given SDP instance (g_z, g_r, h_z, h_u) , which also contradicts the DLIN assumption.

Assuming that event $S_5 \wedge \neg E_5$ indeed occurs, we know that the adversary \mathcal{A} manages to output a correct proof $\pi^* = (\mathbf{VK}^*, \mathbf{F}^*, \mathbf{C}_{\sigma_1}^*, \sigma_2^*, \sigma_3^*, \mathbf{C}_Z^*, \mathbf{C}_R^*, \mathbf{C}_U^*, \mathbf{C}_z^*, \mathbf{C}_r^*, \mathbf{C}_u^*, \pi_{\sigma,1}^*, \pi_{\sigma,2}^*, \pi_1^*, \pi_2^*, \sigma^*)$ for a vector $\mathbf{v}^* = (v_1^*, \dots, v_n^*)$ outside the row space of $\boldsymbol{\rho} = g^{\mathbf{A}}$ and such that $\mathbf{F}^* = (F_1^*, F_2^*, F_3^*)$ is a BBS encryption of $1_{\mathbb{G}}$ (namely, $F_3^* = F_1^{*1/y_1} \cdot F_2^{*1/y_1}$). This means that, although the simulated proofs produced by $\mathcal{S}_2(\psi, \tau_{sim}, \cdot, \cdot)$ were all generated for a perfectly NIWI Groth-Sahai CRS $\mathbf{F} = (\mathbf{f}_1, \mathbf{f}_2, \mathbf{F})$, the last part $(\mathbf{C}_z^*, \mathbf{C}_r^*, \mathbf{C}_u^*, \pi_1^*, \pi_2^*)$ of \mathcal{A} 's proof π^* takes place on a perfectly binding CRS $\mathbf{F}^* = (\mathbf{f}_1, \mathbf{f}_2, \mathbf{F}^*)$. Moreover, although \mathcal{B} does not know $\mu_1^*, \mu_2^* \in \mathbb{Z}_p$ such that $\mathbf{F}^* = \mathbf{f}_1^{\mu_1^*} \cdot \mathbf{f}_2^{\mu_2^*}$, \mathcal{B} can still use the extraction

⁷ This notion (see Definition 4 in [39]) is defined via a game where the adversary is given q verification keys $\{\mathbf{VK}_i\}_{i=1}^q$ and an oracle that returns exactly one signature for each key. The adversary's task is to output a triple (i^*, M^*, σ^*) , where $i^* \in \{1, \dots, q\}$ and (M^*, σ^*) was not produced by the signing oracle for \mathbf{VK}_{i^*} . Hofheinz and Jager [39, Section 4.2] gave a discrete-log-based one-time signature with tight security in the multi-key setting.

trapdoor $(y_1, y_2) = (\log_g(f_1), \log_g(f_2))$ to recover (z^*, r^*, u^*) from (C_z^*, C_r^*, C_u^*) by performing BBS decryptions. Indeed, $C_z^* = \iota(z^*) \cdot \mathbf{f}_1^{\theta_{z,1}} \cdot \mathbf{f}_2^{\theta_{z,2}} \cdot \mathbf{F}^{\theta_{z,3}}$ is of the form $C_z^* = \iota(z^*) \cdot \mathbf{f}_1^{\theta_{z,1} + \mu_1^* \cdot \theta_{z,3}} \cdot \mathbf{f}_2^{\theta_{z,2} + \mu_2^* \cdot \theta_{z,3}}$, which decrypts to z^* .

The perfect soundness of the Groth-Sahai CRS $\mathbf{F}^* = (\mathbf{f}_1, \mathbf{f}_2, \mathbf{F}^*)$ ensures that extracted group elements (z^*, r^*, u^*) satisfy the pairing product equations

$$e(g_z, z^*) \cdot e(g_r, r^*) \cdot \prod_{i=1} e(g_i, v_i^*) = e(h_z, z^*) \cdot e(h_u, u^*) \cdot \prod_{i=1} e(h_i, v_i^*) = 1_{\mathbb{G}_T} \quad (8)$$

In addition, \mathcal{B} computes $(z^\dagger, r^\dagger, u^\dagger) = (\prod_{i=1}^n v_i^{*\chi_i}, \prod_{i=1}^n v_i^{*\gamma_i}, \prod_{i=1}^n v_i^{*\delta_i})$, which also satisfies the equations (8). Since $(z^\dagger, r^\dagger, u^\dagger)$ and (z^*, r^*, u^*) both satisfy (8), the triple $(z^\dagger, r^\dagger, u^\dagger) = (\frac{z^*}{z^\dagger}, \frac{r^*}{r^\dagger}, \frac{u^*}{u^\dagger})$ necessarily satisfies the equalities $e(g_z, z^\dagger) \cdot e(g_r, r^\dagger) = e(h_z, z^\dagger) \cdot e(h_u, u^\dagger) = 1_{\mathbb{G}_T}$. We argue that $z^\dagger \neq 1_{\mathbb{G}}$ with probability $1 - 1/p$, so that $(z^\dagger, r^\dagger, u^\dagger)$ breaks the SDP assumption.

To see this, we remark that, if event $S_5 \wedge \neg E_5$ actually happens, \mathcal{B} never reveals any information about (χ_1, \dots, χ_n) when it emulates $S_2(\psi, \tau_{sim}, \dots)$. Indeed, in simulated proofs, the only components that depend on (χ_1, \dots, χ_n) are $(C_z, C_r, C_u, \pi_1, \pi_2)$, which are generated for a perfectly NIWI Groth-Sahai CRS $(\mathbf{f}_1, \mathbf{f}_2, \mathbf{F})$. Consequently, the same arguments as in [49, Theorem 1] show that $z^\dagger \neq z^*$ with probability $1 - 1/p$. In the CRS, $\{(g_i, h_i)\}_{i=1}^n$ and $\{(z_i, r_i, u_i)\}_{i=1}^t$ provide \mathcal{A} with a linear system of $2n + t < 3n$ equations in $3n$ unknowns $\{(\chi_i, \gamma_i, \delta_i)\}_{i=1}^n$, which leaves z^\dagger completely undetermined in \mathcal{A} 's view if \mathbf{v}^* is linearly independent of the rows of $\boldsymbol{\rho} = (G_{i,j})_{i,j}$. We thus find $\Pr[S_5 \wedge \neg E_5] \leq \mathbf{Adv}_{\mathcal{B}}^{\text{SDP}}(\lambda) + 1/p$, which yields the bound (7) since $\mathbf{Adv}_{\mathcal{B}}^{\text{SDP}}(\lambda) \leq \frac{1}{2} \cdot \mathbf{Adv}_{\mathcal{B}}^{\text{DLIN}}(\lambda)$ if we translate the SDP solver \mathcal{B} into a DLIN distinguisher. \square

The result easily extends to the multi-CRS setting via the following changes. In the transitions from Game_1 to Game_2 and Game_2 to Game_3 , we can simultaneously modify all CRSes $\{\psi^{(\kappa)}\}_{\kappa=1}^\mu$ by using the random self-reducibility of DLIN to build μ instances of the DLIN assumption from a given instance. In Game_5 , the probability $\Pr[E_5]$ can be bounded by implicitly relying on the multi-user security (in the sense of [33]) of the signature scheme of [51], which remains almost tight in the multi-key setting. In the proof of the following lemma, we will explain at each step how the proof can be adapted to the multi-CRS setting. Finally, the probability of event $S_5 \wedge \neg E_5$ in Game_5 can be proved by applying the same arguments as in the proof (see [51, Appendix G]) that the signature of [51] provides tight security in the multi-user setting.

Lemma 1. *In Game_5 , there is a DLIN distinguisher \mathcal{B} such that the probability $\Pr[E_5]$ is at most $\Pr[E_5] \leq (2 \cdot L + 1) \cdot \mathbf{Adv}_{\mathcal{B}}^{\text{DLIN}}(\lambda) + 2/p$. Moreover, \mathcal{B} 's complexity only exceeds that of \mathcal{A} by a polynomial number of exponentiations and a constant number of pairing computations. (The proof is given in the full version).*

5 Applications to Tightly Secure Primitives

As an application of our QA-NIZK proof system, we present a new encryption scheme whose IND-CCA2 security in the multi-challenge-multi-user setting

(almost) tightly relates to the DLIN assumption. We show that the resulting construction allows improving the expansion rate of non-interactive universally composable commitments based on IND-CCA2-secure public-key encryption.

5.1 CCA2-Secure (Threshold) Encryption with Shorter Ciphertexts

Like [39,51], our scheme builds on the Naor-Yung paradigm [56] and the encryption scheme of Boneh, Boyen and Shacham (BBS) [16].

The encryption phase computes $(C_0, C_1, C_2) = (M \cdot g^{\theta_1 + \theta_2}, X_1^{\theta_1}, Y_1^{\theta_2})$ and $(D_0, D_1, D_3) = (M \cdot g^{\theta_3 + \theta_4}, X_2^{\theta_3}, Y_2^{\theta_4})$, where (X_1, Y_1, X_2, Y_2) are part of the public key, and generates a QA-NIZK proof π that the vector

$$\begin{aligned} \mathbf{v} &= (C_1/D_1, C_2/D_2, C_0/D_0, C_1 \cdot C_2, D_1^{-1} \cdot D_2^{-1}) \in \mathbb{G}^5 \\ &= (X_1^{\theta_1} \cdot X_2^{-\theta_3}, Y_1^{\theta_2} \cdot Y_2^{-\theta_4}, g^{(\theta_1 + \theta_2) - (\theta_3 + \theta_4)}, X_1^{\theta_1} \cdot Y_1^{\theta_2}, X_2^{-\theta_3} \cdot Y_2^{-\theta_4}) \end{aligned}$$

is in the subspace spanned by $\mathbf{X}_1 = (X_1, 1, g, X_1, 1)$, $\mathbf{Y}_1 = (1, Y_1, g, Y_1, 1)$, $\mathbf{X}_2 = (X_2, 1, g, 1, X_2)$ and $\mathbf{Y}_2 = (1, X_2, g, 1, X_2)$. As in [51], our reduction is not quite as tight as in [39,5] since a factor $\Theta(\lambda)$ is lost. On the other hand, our scheme becomes nearly practical as the ciphertext overhead now decreases to 48 group elements. In comparison, the solution of Libert *et al.* [51] incurs 69 group elements per ciphertext. Our technique thus improves upon [51] by 30% and also outperforms the most efficient perfectly tight solution [5], which entails over 300 group elements per ciphertext.

The CRS of the proof system is included in the user's public key rather than in the common public parameters since, in the QA-NIZK setting, it depends on the considered language which is defined by certain public key components.

Par-Gen(λ): Run the K_0 algorithm of Section 3 in order to obtain common public parameters $\Gamma = ((\mathbb{G}, \mathbb{G}_T), f, g, h, \Sigma)$.

Keygen(Γ): Parse Γ as $((\mathbb{G}, \mathbb{G}_T), f, g, h, \Sigma)$ and conduct the following steps.

1. Choose random exponents $x_1, x_2, y_1, y_2 \xleftarrow{R} \mathbb{Z}_p$ and define $X_1 = g^{x_1}$, $X_2 = g^{x_2}$, $Y_1 = g^{y_1}$, $Y_2 = g^{y_2}$. Then, define the independent vectors $\mathbf{X}_1 = (X_1, 1, g, X_1, 1)$, $\mathbf{Y}_1 = (1, Y_1, g, Y_1, 1)$, $\mathbf{X}_2 = (X_2, 1, g, 1, X_2)$ and $\mathbf{Y}_2 = (1, X_2, g, 1, X_2)$.
2. Run algorithm $K_1(\Gamma, \rho)$ of Section 3 to generate the language-dependent part of the CRS for the proof system, where the rows of the matrix $\rho \in \mathbb{G}^{4 \times 5}$ consist of \mathbf{X}_1 , \mathbf{Y}_1 , \mathbf{X}_2 and \mathbf{Y}_2 . Let $\psi = (\mathbf{CRS}_1, \mathbf{CRS}_2)$ be the obtained CRS, where

$$\begin{aligned} \mathbf{CRS}_1 &= \left(\rho, \mathbf{f}, \mathbf{f}_0, \{u_i\}_{i=1}^2, \{\Omega_i\}_{i=1}^2, \mathbf{V}, \mathbf{W}, \right. \\ &\quad \left. \{\mathbf{pk}_i\}_{i=1}^2, \{(z_i, r_i, u_i)\}_{i=1}^4, \{(Z_j, R_j, U_j)\}_{j=1}^{4L+5} \right), \end{aligned}$$

$$\mathbf{CRS}_2 = \left(\mathbf{f}, \mathbf{f}_0, \{\mathbf{pk}_i\}_{i=1}^2, \{\Omega_i\}_{i=1}^2, \mathbf{V}, \mathbf{W} \right).$$

3. Define the private key as the pair $SK = (x_1, y_1) \in \mathbb{Z}_p^4$. The public key is

$$PK = (g, \mathbf{X}_1, \mathbf{Y}_1, \mathbf{X}_2, \mathbf{Y}_2, \psi = (\mathbf{CRS}_1, \mathbf{CRS}_2)).$$

Encrypt(M, PK): to encrypt $M \in \mathbb{G}$, conduct the following steps.

1. Pick random exponents $\theta_1, \theta_2, \theta_3, \theta_4 \xleftarrow{R} \mathbb{Z}_p$ and compute

$$\begin{aligned} (C_0, C_1, C_2) &= (M \cdot g^{\theta_1 + \theta_2}, X_1^{\theta_1}, Y_1^{\theta_2}) \\ (D_0, D_1, D_3) &= (M \cdot g^{\theta_3 + \theta_4}, X_2^{\theta_3}, Y_2^{\theta_4}) . \end{aligned}$$

2. Define $\text{lbl} = (C_0, C_1, C_2, D_0, D_1, D_2)$. Using the witness $\mathbf{x} = (\theta_1, \theta_2, -\theta_3, -\theta_4) \in \mathbb{Z}_p^4$ and the label lbl , run Steps 1-7 of Algorithm P in Section 3 to generate a proof π that the vector

$$\begin{aligned} \mathbf{v} &= (C_1/D_1, C_2/D_2, C_0/D_0, C_1 \cdot C_2, D_1^{-1} \cdot D_2^{-1}) \in \mathbb{G}^5 \\ &= (X_1^{\theta_1} \cdot X_2^{-\theta_3}, Y_1^{\theta_2} \cdot Y_2^{-\theta_4}, g^{(\theta_1 + \theta_2) - (\theta_3 + \theta_4)}, X_1^{\theta_1} \cdot Y_1^{\theta_2}, X_2^{-\theta_3} \cdot Y_2^{-\theta_4}) \end{aligned}$$

belongs to $\text{span}\langle \mathbf{X}_1, \mathbf{Y}_1, \mathbf{X}_2, \mathbf{Y}_2 \rangle$. The QA-NIZK proof is

$$\pi = (\text{VK}, \mathbf{F}, \mathbf{C}_{\sigma_1}, \sigma_2, \sigma_3, \mathbf{C}_Z, \mathbf{C}_R, \mathbf{C}_U, \mathbf{C}_z, \mathbf{C}_r, \mathbf{C}_u, \boldsymbol{\pi}_{\sigma_1}, \boldsymbol{\pi}_{\sigma_2}, \boldsymbol{\pi}_1, \boldsymbol{\pi}_2, \sigma) .$$

3. Output the ciphertext $C = (C_0, C_1, C_2, D_0, D_1, D_2, \pi)$.

Decrypt(SK, C): given $C = (C_0, C_1, C_2, D_0, D_1, D_2, \pi)$, do the following.

1. Run the verification algorithm V of Section 3 on input of $\text{lbl} = (C_0, C_1, C_2, D_0, D_1, D_2)$, the vector $\mathbf{v} = (C_1/D_1, C_2/D_2, C_0/D_0, C_1 \cdot C_2, D_1^{-1} \cdot D_2^{-1})$ and π . Return \perp if π is not a valid proof for the label lbl that \mathbf{v} is in $\text{span}\langle \mathbf{X}_1, \mathbf{Y}_1, \mathbf{X}_2, \mathbf{Y}_2 \rangle$.
2. Using $SK = (x_1, y_1) \in \mathbb{Z}_p^2$, compute and return $M = C_0 \cdot C_1^{-1/x_1} \cdot C_2^{-1/y_1}$.

Using our proof system of Section 3 and the one-time signature of [39], the ciphertext size amounts to that of 48 group elements, instead of 69 in [51].

While our construction is described in terms of symmetric pairings in order to lighten notations as much as possible, it readily extends to asymmetric pairings.

Theorem 4. *The scheme is $(1, q_e)$ -IND-CCA secure provided: (i) Σ is a strongly unforgeable one-time signature; (ii) The DLIN assumption holds in \mathbb{G} . For any adversary \mathcal{A} , there exist a one-time signature forger \mathcal{B}' and a DLIN distinguisher \mathcal{B} with running times $t_{\mathcal{B}}, t_{\mathcal{B}'} \leq t_{\mathcal{A}} + q_e \cdot \text{poly}(\lambda, L)$ such that*

$$\text{Adv}_{\mathcal{A}}^{(1, q_e)\text{-cca}}(\lambda) \leq \text{Adv}_{\mathcal{B}'}^{q_e\text{-suf-ots}}(\lambda) + (3L + 10) \cdot \text{Adv}_{\mathcal{B}}^{\text{DLIN}}(\lambda) + 8/p ,$$

where L is the length of one-time verification keys and q_e is the number of encryption queries. (The proof is given in the full version of the paper.)

The result of Theorem 4 carries over to a scenario involving $\mu > 1$ public keys modulo an additional negligible term μ/p in the bound which is inherited from [39, Theorem 6]. This is achieved by relying on the enhanced USS property of the QA-NIZK proof system in the multi-CRS setting.

Similarly to previous IND-CCA2-secure encryption schemes based on the Naor-Yung paradigm (e.g., [32]), the public verifiability of ciphertexts makes our scheme amenable for non-interactive threshold decryption in a static corruption model.

By instantiating the construction of Camenisch *et al.* [19] with our QA-NIZK proofs, we similarly obtain more efficient KDM-CCA2-secure systems with tight security, as explained in the full version of the paper.

5.2 Encrypting Long Messages

In some applications, it is useful to encrypt long messages while preserving the feasibility of efficiently proving statements about encrypted values using Groth-Sahai proofs. In this case, the amortized efficiency of our system can be significantly improved. Suppose that we want to encrypt messages $(M_1, \dots, M_N) \in \mathbb{G}^N$. The technique of Bellare *et al.* [8] allows doing so while making optimal use of encryption exponents. In more details, the public key consists of group elements $(g, h, \{(X_{i,1}, Y_{i,1}, X_{i,2}, Y_{i,2})\}_{i=1}^N)$, with $(X_{i,1}, Y_{i,1}, X_{i,2}, Y_{i,2}) = (g^{x_{i,1}}, h^{y_{i,1}}, g^{x_{i,2}}, h^{y_{i,2}})$ and the secret key is $\{(x_{i,1}, y_{i,1})\}_{i=1}^N$. The vector is encrypted by choosing $\theta_1, \theta_2, \theta_3, \theta_4 \xleftarrow{R} \mathbb{Z}_p$ and computing

$$\begin{aligned} C_0 &= f^{\theta_1}, & C'_0 &= h^{\theta_2}, & \{C_i &= M_i \cdot X_{i,1}^{\theta_1} \cdot Y_{i,1}^{\theta_2}\}_{i=1}^N, \\ D_0 &= f^{\theta_3}, & D'_0 &= h^{\theta_4}, & \{D_i &= M_i \cdot X_{i,2}^{\theta_3} \cdot Y_{i,2}^{\theta_4}\}_{i=1}^N, \end{aligned}$$

while appending a simulation-sound QA-NIZK argument that the vector

$$\begin{aligned} & (C_1/D_1, \dots, C_N/D_N, \overbrace{C_0, \dots, C_0}^{N \text{ times}}, \\ & \overbrace{D_0^{-1}, \dots, D_0^{-1}}^{N \text{ times}}, \overbrace{C'_0, \dots, C'_0}^{N \text{ times}}, \overbrace{D'_0{}^{-1}, \dots, D'_0{}^{-1}}^{N \text{ times}}) \in \mathbb{G}^{5N} \end{aligned}$$

lives in the $4N$ -dimensional linear subspace $\text{span}\langle \mathbf{X}_{i,1}, \mathbf{X}_{i,2}, \mathbf{Y}_{i,1}, \mathbf{Y}_{i,2} \rangle_{i=1}^N$, with

$$\begin{aligned} \mathbf{X}_{i,1} &= (\mathbf{1}^{i-1}, X_{i,1}, \mathbf{1}^{N-i}, \mathbf{1}^{i-1}, f, \mathbf{1}^{N-i}, \mathbf{1}^{3N}), \\ \mathbf{X}_{i,2} &= (\mathbf{1}^{i-1}, X_{i,2}, \mathbf{1}^{N-i}, \mathbf{1}^N, \mathbf{1}^{i-1}, f, \mathbf{1}^{N-i}, \mathbf{1}^{2N}), \\ \mathbf{Y}_{i,1} &= (\mathbf{1}^{i-1}, Y_{i,1}, \mathbf{1}^{N-i}, \mathbf{1}^{2N}, \mathbf{1}^{i-1}, h, \mathbf{1}^{N-i}, \mathbf{1}^N), \\ \mathbf{Y}_{i,2} &= (\mathbf{1}^{i-1}, Y_{i,2}, \mathbf{1}^{N-i}, \mathbf{1}^{3N}, \mathbf{1}^{i-1}, h, \mathbf{1}^{N-i}), \end{aligned}$$

where, for each $i \in \mathbb{N}$, $\mathbf{1}^i$ stands for the i -dimensional vector $(1_{\mathbb{G}}, \dots, 1_{\mathbb{G}}) \in \mathbb{G}^i$. The entire ciphertext fits within $2N + 46$ group elements, of which only 42 elements are consumed by the QA-NIZK proof.

The tight IND-CCA2 security can be proved in the same way as in Theorem 4. In particular, we rely on the tight IND-CPA security in the multi-challenge setting of a variant of the BBS encryption scheme where messages M are encrypted⁸ as $(f^{\theta_1}, h^{\theta_2}, M \cdot X^{\theta_1} \cdot Y^{\theta_2})$.

In Section 5.3, we explain how the compatibility of this construction with zero-knowledge proofs comes in handy to build non-interactive and adaptively secure universally composable commitments based on CCA2-secure encryption.

5.3 Application to UC Commitments

Universally composable commitments [27,20] are commitment schemes that provably remain secure when composed with arbitrary other protocols. They are

⁸ The reduction from the DLIN assumption is straightforward and sets up $X = f^\alpha \cdot g^\gamma$, $Y = h^\beta \cdot g^\gamma$. From a given DLIN instance $(f, g, h, f^a, h^b, \eta)$, where $\eta = g^{a+b}$ or $\eta \in_R \mathbb{G}$, the challenge ciphertext is computed as $(C_1, C_2, C_3) = (f^a, h^b, M_\beta \cdot (f^a)^\alpha \cdot (h^b)^\beta \cdot \eta^\gamma)$.

known [20] to require some setup assumption like a common reference string. In some constructions, the CRS can only be used in a single commitment. Back in 2001, Canetti and Fischlin [20] gave re-usable bit commitments based on chosen-ciphertext-secure public-key encryption. In [53], Lindell described a simple and practical re-usable construction which allows committing to strings rather than individual bits. In short, each commitment consists of an IND-CCA2-secure encryption. In order to open a commitment later on, the sender generates an interactive zero-knowledge proof that the ciphertext encrypts the underlying plaintext. In its basic variant, Lindell’s commitment only provides security against static adversaries that have to choose whom to corrupt upfront⁹. Subsequently, Fischlin *et al.* [31] showed that Lindell’s commitment can be made adaptively secure in the erasure model by the simple expedient of opening commitments via a NIZK proof (rather than an interactive one) which the sender generates at commitment time before erasing his encryption coins. Jutla and Roy [43] gave an optimization of the latter approach where the use of QA-NIZK proofs allows reducing the size of commitments and openings.

Using our CCA2-secure encryption scheme for long messages, we can build a tightly secure non-interactive universally composable commitment [27,20] that allows committing to long messages with expansion rate 2. In constructions of UC commitments from IND-CCA2-secure encryption (e.g., [20,31,43]), a multi-challenge definition of IND-CCA2 security is usually considered in proofs of UC security. In the erasure model, the non-interactive and adaptively secure variants of Lindell’s commitment [31,43] can be optimized using the techniques of [50,44] to achieve a two-fold expansion rate. However, these solutions are not known to provide tight security. At the cost of a CRS of size $\Theta(N)$, the labeled version of our encryption scheme for long messages (where the label L of the ciphertext is simply included in lbl) allows eliminating this limitation. As in [43], the sender can encrypt the message (M_1, \dots, M_N) he wants to commit to and open the commitment via a QA-NIZK proof that

$$(C_1/M_1, \dots, C_N/M_N, \overbrace{C_0, \dots, C_0}^{N \text{ times}}, \overbrace{1, \dots, 1}^{N \text{ times}}, \overbrace{C'_0, \dots, C'_0}^{N \text{ times}}, \overbrace{1, \dots, 1}^{N \text{ times}}) \in \mathbb{G}^{5N}$$

is in $\text{span}(\mathbf{X}_{i,1}, \mathbf{X}_{i,2}, \mathbf{Y}_{i,1}, \mathbf{Y}_{i,2})_{i=1}^N$. For long messages, this construction thus achieves a two-fold expansion rate. While not as efficient as the recent rate-1 commitments of Garay *et al.* [35], it retains adaptive security assuming reliable erasures while [35] is only known to be secure against static adversaries.

Acknowledgments

The first author’s work was supported by the “Programme Avenir Lyon Saint-Etienne de l’Université de Lyon” in the framework of the programme “Investissements d’Avenir” (ANR-11-IDEX-0007). The second author was supported by

⁹ Lindell’s commitment can actually be made adaptively secure (modulo a patch [13]), but even its optimized variant [13] remains interactive with 3 rounds of communication during the commitment phase.

the European Research Council (FP7/2007-2013 Grant Agreement no. 339563 CryptoCloud). Part of this work of the fourth author was done while visiting the Simons Institute for Theory of Computing, U.C. Berkeley.

References

1. M. Abdalla, F. Ben Hamouda, D. Pointcheval. Disjunctions for hash proof systems: New constructions and applications. In *Eurocrypt '15*, LNCS 9057, pp. 69–100, Springer, 2015.
2. M. Abdalla, P.-A. Fouque, V. Lyubashevsky, M. Tibouchi. Tightly-secure signatures from lossy identification schemes. In *Eurocrypt '12*, LNCS 7237, pp. 572–590, Springer, 2012.
3. M. Abe, G. Fuchsbauer, J. Groth, K. Haralambiev, M. Ohkubo. Structure-preserving signatures and commitments to group elements. In *Crypto '10*, LNCS 6223, pp. 209–236, Springer, 2010.
4. M. Abe, K. Haralambiev, M. Ohkubo. Signing on elements in bilinear groups for modular protocol design. In Cryptology ePrint Archive: Report 2010/133, 2010.
5. M. Abe, B. David, M. Kohlweiss, R. Nishimaki, M. Ohkubo. Tagged one-time signatures: Tight security and optimal tag size. In *PKC '13*, LNCS 7778, pp. 312–331, Springer, 2013.
6. C. Bader, D. Hofheinz, T. Jager, E. Kiltz, Y. Li. Tightly-secure authenticated key exchange. In *TCC 2015*, LNCS series, Springer, 2015.
7. M. Bellare, A. Boldyreva, S. Micali. Public-key encryption in a multi-user setting: Security proofs and improvements. In *Eurocrypt '00*, LNCS 1807, pp. 259–274, Springer, 2000.
8. M. Bellare, A. Boldyreva, K. Kurosawa, J. Staddon. Multi-recipient encryption schemes: How to save on bandwidth and computation without sacrificing security. In *IEEE Trans. Inf. Theory*, 53 (11), pp. 3927–3943, 2007.
9. M. Bellare, P. Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *ACM CCS '93*, pp. 62–73, ACM Press, 1993.
10. M. Bellare, P. Rogaway. The exact security of digital signatures - How to sign with RSA and Rabin. In *Eurocrypt '96*, LNCS 1070, pp. 399–416, Springer, 1996.
11. D. Bernstein. Proving tight security for Rabin-Williams signatures. In *Eurocrypt '08*, LNCS 4965, pp. 70–87, Springer, 2008.
12. J. Black, P. Rogaway, T. Shrimpton. Encryption scheme security in the presence of key-dependent messages. In *Selected Areas in Cryptography (SAC '02)*, LNCS 2595, pp. 62–75, pp. 62–75, Springer, 2002.
13. O. Blazy, C. Chevalier, D. Pointcheval, D. Vergnaud. Analysis and improvement of Lindell's UC-secure commitment schemes. In *ACNS '13*, LNCS 7954, pp. 70–87, pp. 534–551, Springer, 2014.
14. O. Blazy, E. Kiltz, J. Pan. (Hierarchical) identity-based encryption from affine message authentication. In *Crypto '14*, LNCS 8616, pp. 70–87, pp. 408–425, Springer, 2014.
15. M. Blum, P. Feldman, S. Micali. Non-interactive zero-knowledge and its applications. In *STOC '88*, pp. 103–112, ACM Press, 1988.
16. D. Boneh, X. Boyen, H. Shacham. Short group signatures. In *Crypto '04*, LNCS 3152, pp. 41–55, Springer, 2004.
17. D. Boneh, M. Franklin. Identity-based encryption from the Weil pairing. In *SIAM J. of Computing* 32(3), pp. 586–615, 2003. Earlier version in *Crypto '01*.

18. D. Boneh, S. Halevi, M. Hamburg, R. Ostrovsky. Circular-secure encryption from Decision Diffie-Hellman. In *Crypto '08*, LNCS 5157, pp. 108–125, Springer, 2008.
19. J. Camenisch, N. Chandran, and V. Shoup. A public key encryption scheme secure against key dependent chosen plaintext and adaptive chosen ciphertext attacks. In *Eurocrypt '09*, LNCS 5479, pp. 351–368, Springer, 2009.
20. R. Canetti, M. Fischlin. Universally composable commitments. In *Crypto '01*, LNCS 2139, pp. 19–40, Springer, 2001.
21. R. Canetti, S. Halevi, J. Katz. Chosen-ciphertext security from identity-based encryption. In *Eurocrypt '04*, LNCS 3027, pp. 207–222, 2004.
22. J. Cathalo, B. Libert, M. Yung. Group encryption: Non-interactive realization in the standard model. In *Asiacrypt '09*, LNCS 5912, pp. 179–196, Springer, 2009.
23. J. Chen, H. Wee. Fully, (almost) tightly secure IBE and dual system groups. In *Crypto '13*, LNCS 8043, pp. 435–460, Springer, 2013.
24. B. Chevallier-Mames. An efficient CDH-based signature scheme with a tight security reduction. In *Crypto '05*, LNCS 3621, pp. 511–526, Springer, 2005.
25. J.-S. Coron. On the exact security of full domain hash. In *Crypto '00*, LNCS 1880, pp. 229–235, Springer, 2000.
26. J.-S. Coron. Optimal security proofs for PSS and other signature schemes. In *Eurocrypt '02*, LNCS 2332, pp. 229–235, Springer, 2002.
27. R. Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *FOCS '01*, pp. 136–145, 2001.
28. R. Cramer, V. Shoup. A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In *Crypto '98*, LNCS 1462, pp. 13–25, Springer, 1998.
29. A. Escala, J. Groth. Fine-tuning Groth-Sahai proofs. In *PKC '14*, LNCS 8383, pp. 630–649, Springer, 2014.
30. A. Fiat, A. Shamir. How to prove yourself: Practical solutions to identification and signature problems. In *Crypto '86*, LNCS 263, pp. 186–194, Springer, 1986.
31. M. Fischlin, B. Libert, M. Manulis. Non-interactive and re-usable universally composable string commitments with adaptive security. In *Asiacrypt '11*, LNCS 7073, pp. 468–485, Springer, 2014.
32. P.-A. Fouque, D. Pointcheval. Threshold cryptosystems secure against chosen-ciphertext attacks. In *Asiacrypt '01*, LNCS 2248, pp. 351–368, Springer, 2001.
33. S. Galbraith, J. Malone-Lee, N. Smart. Public-key signatures in the multi-user setting. In *Information Processing Letters*, vol. 83(5), pp. 263–266, 2002.
34. S. Galbraith, K. Paterson, N. Smart. Pairings for cryptographers. In *Discrete Applied Mathematics*, 156(16), pp. 3113–3121, 2008.
35. J. Garay, Y. Ishai, R. Kumaresan, H. Wee. On the complexity of UC commitments. In *Eurocrypt '14*, LNCS 8441, pp. 677–694, Springer, 2014.
36. J. Groth. Simulation-sound NIZK proofs for a practical language and constant size group signatures. In *Asiacrypt '06*, LNCS 4284, pp. 444–459, Springer, 2006.
37. J. Groth, A. Sahai. Efficient non-interactive proof systems for bilinear groups. In *Eurocrypt '08*, LNCS 4965, pp. 415–432, Springer, 2008.
38. D. Hofheinz. All-but-many lossy trapdoor functions. In *Eurocrypt '12*, LNCS 7237, pp. 209–227, Springer, 2012.
39. D. Hofheinz, T. Jager. Tightly secure signatures and public-key encryption. In *Crypto '12*, LNCS 7417, pp. 590–607, Springer, 2012.
40. D. Hofheinz, T. Jager, E. Knapp. Waters signatures with optimal security reduction. In *PKC '12*, LNCS 7293, pp. 66–83, Springer, 2012.

41. D. Hofheinz, J. Koch, C. Striecks. Identity-based encryption with (almost) tight security in the multi-instance, multi-ciphertext setting. In *PKC '15*, LNCS 9020, pp. 799–822, Springer, 2015.
42. C. Jutla, A. Roy. Relatively-sound NIZKs and password-based key-exchange. In *PKC '12*, LNCS 7293, pp. 485–503, Springer, 2012.
43. C. Jutla, A. Roy. Shorter quasi-adaptive NIZK proofs for linear subspaces. In *Asiacrypt '13*, LNCS 8269, pp. 1–20, Springer, 2013. Cryptology ePrint Archive: Report 2013/109, 2013.
44. C. Jutla, A. Roy. Switching lemma for bilinear tests and constant-size NIZK proofs for linear subspaces. In *Crypto '14*, LNCS 8617, pp. 295–312, Springer, 2014.
45. S. Kakvi, E. Kiltz. Optimal security proofs for full domain hash, revisited. In *Eurocrypt '12*, LNCS 7237, pp. 537–553, Springer, 2012.
46. J. Katz, V. Vaikuntanathan. Round-optimal password-based authenticated key exchange. In *TCC '11*, LNCS 6597, pp. 293–310, Springer, 2011.
47. J. Katz, N. Wang. Efficiency improvements for signature schemes with tight security reductions. In *ACM-CCS '03*, pp. 155–164, ACM Press, 2003.
48. E. Kiltz, H. Wee. Quasi-adaptive NIZK for linear subspaces revisited. In *Eurocrypt '15*, LNCS 9057, pp. 101–128, Springer, 2015.
49. B. Libert, T. Peters, M. Joye, M. Yung. Linearly homomorphic structure-preserving signatures and their applications. In *Crypto '13*, LNCS 8043, pp. 289–307, Springer, 2013.
50. B. Libert, T. Peters, M. Joye, M. Yung. Non-malleability from malleability: Simulation-sound quasi-adaptive NIZK proofs and CCA2-secure encryption from homomorphic signatures. In *Eurocrypt '14*, LNCS 8441, pp. 514–532, Springer, 2014.
51. B. Libert, M. Joye, M. Yung, T. Peters. Concise multi-challenge CCA-secure encryption and signatures with almost tight security. In *Asiacrypt '14*, LNCS 8874, pp. 1–21, Springer, 2014.
52. B. Libert, M. Yung. Non-interactive CCA2-secure threshold cryptosystems with adaptive security: New framework and constructions. In *TCC '12*, LNCS 7194, pp. 75–93, Springer, 2012.
53. Y. Lindell. Highly-efficient universally-composable commitments based on the DDH assumption. In *Eurocrypt '11*, LNCS 6632, pp. 446–466, Springer, 2011.
54. T. Malkin, I. Teranishi, Y. Vahlis, M. Yung. Signatures resilient to continual leakage on memory and computation. In *TCC '11*, LNCS 6597, pp. 89–106, Springer, 2011.
55. M. Naor, O. Reingold. Number-theoretic constructions of efficient pseudo-random functions. In *FOCS '97*, pp. 458–467, IEEE Press, 1997.
56. M. Naor, M. Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. In *STOC '90*, ACM Press, 1990.
57. C. Rackoff, D. Simon. Non-interactive zero-knowledge proof of knowledge and chosen ciphertext attack. In *Crypto, '91*, LNCS 576, pp. 433–444, Springer, 1991.
58. A. Sahai. Non-malleable non-interactive zero-knowledge and adaptive chosen-ciphertext security. In *FOCS '99*, pp. 543–553, IEEE Press, 1999.
59. S. Schäge. Tight proofs for signature schemes without random oracles. In *Eurocrypt '11*, LNCS 6632, pp. 189–206, Springer, 2011.
60. A. Shamir. Identity-based cryptosystems and signature schemes. In *Crypto '84*, LNCS 196, pp. 47–53, Springer, 1984.
61. B. Waters. Efficient identity-based encryption without random oracles. In *Eurocrypt '05*, LNCS 3494, Springer, 2005.