

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, Laboratoire de l'Informatique du Parallélisme (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 powerful paradigm, suggested recently 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, e.g., 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. Here, we deal with the basic question of whether and to what extent we can simultaneously optimize the proof size and the tightness of security reductions, allowing for important applications with tight security (which are typically to date quite lengthy) to be of shorter size. In this paper, we resolve this question by describing 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 constant-size 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 employ a number of techniques, and explicitly point at a technique – which may be of independent interest – of hiding the linear span of a structure-preserving homomorphic 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, chosen-ciphertext security, threshold cryptosystems, KDM-CCA2 security, UC commitments, bilinear groups, Decision Linear assumption.

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 [34,37], 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 has to rely on some hardness assumption. In order 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 security 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 [24,25,45,23,59,1,43] and in the standard model [59,38,13,6].

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 [37] and it was more a feasibility result than a practical solution. In the context of identity-based encryption (IBE), Chen and Wee [22] 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].

NIZK PROOFS AND SIMULATION-SOUNDNESS. Non-interactive zero-knowledge proofs [14] 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 [35] 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 [28] in the idealized random oracle model [9]. recently, Jutla and Roy [41] 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 [41] gave computationally sound QA-NIZK proofs of length $\Theta(n - t)$ where the

⁵ Using random oracles, Katz and Wang [45] previously gave a tightly secure variant of the Boneh-Franklin IBE [17].

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 [42], regardless of the number of equations or the number of variables. Libert *et al.* [48] independently obtained similar improvements using different techniques.

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 [34] 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 [44,50,40] 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 [48] gave very efficient unbounded simulation-sound quasi-adaptive NIZK proofs which do not require quadratic pairing product equations or IND-CCA2-secure encryption. Interestingly, 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 [48] do not imply tight chosen-ciphertext security [37] 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 [34,37] are the only ones to enable tight security in this setting and it is unclear how to render them as efficient as [48] for linear multi-exponentiation equations.

TIGHTNESS AND CHOSEN-CIPHERTEXT SECURITY. Bellare, Boldyreva and Micali [7] provided evidence 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 [27], the degradation factor only depends on the number of challenges per user. More recently, Hofheinz and Jager [37] used a tightly secure simulation-sound proof system to construct the first IND-CCA2 secure encryption system whose IND-CCA 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.* [4]. 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 [36] —where the number of input elements depends on the number of adversarial queries— or incurred long ciphertexts [37,4] comprised of hundreds of group elements (or both). One of the reasons is that solutions based on standard assumptions [37,4,49] build on simulation-sound proof systems relying on OR proofs. Recently, Libert *et al.* [49] 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 [49] 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 [42, Section 6] if we accept a loose reduction.

OUR CONTRIBUTIONS. We present short QA-NIZK proofs of linear subspace membership (motivated by those in [48,42]) where the unbounded simulation-soundness property can be *almost tightly*—in the terminology of Chen and Wee [22]—related to the standard Decision Linear (DLIN) assumption [15]. As in [22], 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 [48]. Our construction only lengthens the QA-NIZK proofs of Libert *et al.* [48] by a factor of 2 and thus retains the constant proof length of [48], 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 [55], 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 [37,4], 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 [37,4], 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 [37,4] as they only consist of 48 group elements under the DLIN assumption, whereas the most efficient construction based on the same assumption [49] entails 69 group elements per ciphertext.

Our constant-size proofs offer much more dramatic savings when it comes to encrypting long messages without affecting the compatibility with zero-knowledge proofs. Indeed, 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 [26,20] with adaptive security in the erasure model. In particular, using the same design principle as previous UC commitments [51,29,41] 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) [11,18]—in the sense of [19]—with shorter ciphertexts. Analogously to the Jutla-Roy construction [42, 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. Unlike [42], 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 [37] and Abe *et al.* [4], which incur rather long proofs. Our QA-NIZK proofs yield ciphertexts that are about 75% shorter, as we show in Appendix G.

OUR TECHNIQUES. Our simulation-sound QA-NIZK arguments (as the construction in [48]) build on linearly homomorphic structure-preserving signatures (LHSPS) [47]. In [48], each proof of sub-

space 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 [52], we let the prover generate a Groth-Sahai CRS $\mathbf{F} = (\mathbf{f}_1, \mathbf{f}_2, \mathbf{F})$ of its choice, 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., its last vector \mathbf{F} lives in $\text{span}(\mathbf{f}_1, \mathbf{f}_2)$). 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 [22], which was proved almost tightly secure under the DLIN assumption, recently proposed by Libert, Joye, Yung and Peters [49]. 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 [49]— 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 [49] 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 [49]. 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 [49] 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 [48] to show that the reduction is able to extract a non-trivial homomorphic signature, thus breaking the DLIN assumption.

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 ([15]). *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 [21].

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:

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

Quasi-Adaptive Soundness:

$$\Pr[\Gamma \leftarrow K_0(\lambda); \rho \leftarrow \mathcal{D}_\Gamma; \psi \leftarrow K_1(\Gamma, \rho); (x, \pi, \text{lbl}) \leftarrow \mathcal{A}_2(\Gamma, \psi, \rho) : \\ \forall (\psi, x, \pi, \text{lbl}) = 1 \wedge \neg(\exists w : R_\rho(x, w) = 1)] \in \text{negl}(\lambda) .$$

Quasi-Adaptive Zero-Knowledge:

$$\Pr[\Gamma \leftarrow K_0(\lambda); \rho \leftarrow \mathcal{D}_\Gamma; \psi \leftarrow K_1(\Gamma, \rho) : \mathcal{A}_3^{P(\psi, \dots)}(\Gamma, \psi, \rho) = 1] \\ \approx \Pr[\Gamma \leftarrow K_0(\lambda); \rho \leftarrow \mathcal{D}_\Gamma; (\psi, \tau_{sim}) \leftarrow S_1(\Gamma, \rho) : \mathcal{A}_3^{S(\psi, \tau_{sim}, \dots)}(\Gamma, \psi, \rho) = 1] ,$$

where

- $P(\psi, \cdot, \cdot, \cdot)$ 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 .
- $S(\psi, \tau_{sim}, \cdot, \cdot, \cdot)$ is an oracle that takes as input (x, w) and lbl . It outputs a simulated proof $S_2(\psi, \tau_{sim}, x, \text{lbl})$ 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 ,

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

where the adversary is allowed unbounded access to an oracle $\mathsf{S}_2(\psi, \tau, \cdot, \cdot, \cdot)$ 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

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

Here, \mathcal{A}_4 has access to an oracle $\mathsf{S}_2(\{\psi_\kappa\}_{\kappa=1}^\mu, \{\tau_{sim, \kappa}\}_{\kappa=1}^\mu, \cdot, \cdot, \cdot)$ 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 [3,2] 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.* [47] considered structure-preserving with linear homomorphic properties (see Appendix B for formal definitions). This section reviews the one-time linearly homomorphic structure-preserving signature (LHSPS) of [47].

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 $\mathbf{sk} = \{(\chi_i, \gamma_i, \delta_i)\}_{i=1}^n$ and the public key is $\mathbf{pk} = (g_z, g_r, h_z, h_u, \{(g_i, h_i)\}_{i=1}^n) \in \mathbb{G}^{2n+4}$.

Sign($\mathbf{sk}, (M_1, \dots, M_n)$): to sign a vector $(M_1, \dots, M_n) \in \mathbb{G}^n$ using $\mathbf{sk} = \{(\chi_i, \gamma_i, \delta_i)\}_{i=1}^n$, 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($\mathbf{pk}, \{(\omega_i, \sigma^{(i)})\}_{i=1}^\ell$): given \mathbf{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 [47].

2.4 Short Signatures Almost Tightly Related to the DLIN Assumption

Recently, Libert *et al.* [49] proposed the following signature scheme which was shown almost tightly related to the DLIN assumption. The scheme can be seen as an instantiation of a general construction suggested by Blazy *et al.* [13], where each signature consists of an algebraic message authentication code (MAC) and a NIZK proof that the MAC is valid w.r.t. a committed key included in the public key. In [49], the NIZK proof is a short QA-NIZK argument and the MAC is a triple of the form $(g^{\omega_1 + \omega_2} \cdot H(\mathbf{V}, M)^r \cdot H(\mathbf{W}, M)^s, f^r, h^s)$, for some hash function $H(\mathbf{V}, \cdot)$ and $H(\mathbf{W}, \cdot)$, where the public key contains $(\Omega_1, \Omega_2) = (u_1^{\omega_1}, u_2^{\omega_2})$. The description hereunder assumes symmetric pairings.

Keygen(λ): Choose bilinear groups $(\mathbb{G}, \mathbb{G}_T)$ of prime order p together with $f, g, h, u_1, u_2 \xleftarrow{R} \mathbb{G}$.

1. For $\ell = 1$ to L , choose $V_{\ell,0}, V_{\ell,1}, W_{\ell,0}, W_{\ell,1} \xleftarrow{R} \mathbb{G}$ to assemble row vectors

$$\mathbf{V} = (V_{1,0}, V_{1,1}, \dots, V_{L,0}, V_{L,1}) \in \mathbb{G}^{2L}, \quad \mathbf{W} = (W_{1,0}, W_{1,1}, \dots, W_{L,0}, W_{L,1}) \in \mathbb{G}^{2L} .$$

2. Define the matrix $\mathbf{M} = (M_{i,j})_{i,j}$ given by

$$\mathbf{M} = \begin{pmatrix} \mathbf{V}^\top & \mathbf{Id}_{f,2L} & \mathbf{1}^{2L \times 2L} & \mathbf{1}^{2L \times 1} & \mathbf{1}^{2L \times 1} \\ \mathbf{W}^\top & \mathbf{1}^{2L \times 2L} & \mathbf{Id}_{h,2L} & \mathbf{1}^{2L \times 1} & \mathbf{1}^{2L \times 1} \\ g & \mathbf{1}^{1 \times 2L} & \mathbf{1}^{1 \times 2L} & u_1 & 1 \\ g & \mathbf{1}^{1 \times 2L} & \mathbf{1}^{1 \times 2L} & 1 & u_2 \end{pmatrix} \in \mathbb{G}^{(4L+2) \times (4L+3)} \quad (2)$$

with $\mathbf{Id}_{f,2L} = f\mathbf{I}_{2L} \in \mathbb{G}^{2L \times 2L}$, $\mathbf{Id}_{h,2L} = h\mathbf{I}_{2L} \in \mathbb{G}^{2L \times 2L}$, where $\mathbf{I}_{2L} \in \mathbb{Z}_p^{2L \times 2L}$ is the identity.

3. Generate a key pair $(\mathbf{sk}_{hsp_s}, \mathbf{pk}_{hsp_s})$ for the one-time linearly homomorphic signature of Section 2.3 in order to sign vectors of dimension $n = 4L+3$. Let $\mathbf{sk}_{hsp_s} = \{(\chi_i, \gamma_i, \delta_i)\}_{i=1}^{4L+3}$ be the private key, of which the corresponding public key is $\mathbf{pk}_{hsp_s} = (g_z, g_r, h_z, h_u, \{(g_i, h_i)\}_{i=1}^{4L+3})$.
4. Using the LHSPS private key $\mathbf{sk}_{hsp_s} = \{\chi_i, \gamma_i, \delta_i\}_{i=1}^{4L+3}$, generate one-time homomorphic signatures $\{(Z_j, R_j, U_j)\}_{j=1}^{4L+2}$ on the rows $\mathbf{M}_j = (M_{j,1}, \dots, M_{j,4L+3}) \in \mathbb{G}^{4L+3}$ of \mathbf{M} . These are obtained as

$$(Z_j, R_j, U_j) = \left(\prod_{i=1}^{4L+3} M_{j,i}^{-\chi_i}, \prod_{i=1}^{4L+3} M_{j,i}^{-\gamma_i}, \prod_{i=1}^{4L+3} M_{j,i}^{-\delta_i} \right),$$

for each $j \in \{1, \dots, 4L+2\}$.

5. Choose $\omega_1, \omega_2 \xleftarrow{R} \mathbb{Z}_p$ and compute $\Omega_1 = u_1^{\omega_1} \in \mathbb{G}$, $\Omega_2 = u_2^{\omega_2} \in \mathbb{G}$.

The private key consists of $SK = (\omega_1, \omega_2)$ and the public key is

$$PK = (f, g, h, u_1, u_2, \Omega_1, \Omega_2, \mathbf{V}, \mathbf{W}, \mathbf{pk}_{hsp_s}, \{(Z_j, R_j, U_j)\}_{j=1}^{4L+2}).$$

Sign(SK, M): Given an L -bit message $M = M[1] \dots M[L] \in \{0, 1\}^L$ and $SK = (\omega_1, \omega_2)$:

1. Let $H(\mathbf{V}, M) = \prod_{\ell=1}^L V_{\ell, M[\ell]}$ and $H(\mathbf{W}, M) = \prod_{\ell=1}^L W_{\ell, M[\ell]}$. Choose $r, s \xleftarrow{R} \mathbb{Z}_p$ and compute

$$\sigma_1 = g^{\omega_1 + \omega_2} \cdot H(\mathbf{V}, M)^r \cdot H(\mathbf{W}, M)^s, \quad \sigma_2 = f^r, \quad \sigma_3 = h^s. \quad (3)$$

2. Using $\{(Z_j, R_j, U_j)\}_{j=1}^{4L+2}$, derive a one-time homomorphic signature (Z, R, U) which serves as a QA-NIZK argument showing that the vector

$$(\sigma_1, \sigma_2^{1-M[1]}, \sigma_2^{M[1]}, \dots, \sigma_2^{1-M[L]}, \sigma_2^{M[L]}, \sigma_3^{1-M[1]}, \sigma_3^{M[1]}, \dots, \sigma_3^{1-M[L]}, \sigma_3^{M[L]}, \Omega_1, \Omega_2)$$

belongs to the row space of \mathbf{M} , so that $(\sigma_1, \sigma_2, \sigma_3)$ is of the form (3). Namely, compute

$$Z = Z_{4L+1}^{\omega_1} \cdot Z_{4L+2}^{\omega_2} \cdot \prod_{i=1}^L (Z_{2i-\overline{M[i]}}^r \cdot Z_{2L+2i-\overline{M[i]}}^s) \text{ and}$$

$$R = R_{4L+1}^{\omega_1} \cdot R_{4L+2}^{\omega_2} \cdot \prod_{i=1}^L (R_{2i-\overline{M[i]}}^r \cdot R_{2L+2i-\overline{M[i]}}^s),$$

$$U = U_{4L+1}^{\omega_1} \cdot U_{4L+2}^{\omega_2} \cdot \prod_{i=1}^L (U_{2i-\overline{M[i]}}^r \cdot U_{2L+2i-\overline{M[i]}}^s).$$

Return the signature $\sigma = (\sigma_1, \sigma_2, \sigma_3, Z, R, U) \in \mathbb{G}^6$.

Verify(PK, M, σ): Parse σ as $(\sigma_1, \sigma_2, \sigma_3, Z, R, U) \in \mathbb{G}^6$ and return 1 if and only if

$$\begin{aligned}
e(g_z, Z) \cdot e(g_r, R) &= e(g_1, \sigma_1)^{-1} \cdot e\left(\prod_{i=1}^L g_{2i+M[i]}, \sigma_2\right)^{-1} \cdot \\
&\quad e\left(\prod_{i=1}^L g_{2L+2i+M[i]}, \sigma_3\right)^{-1} \cdot e(g_{4L+2}, \Omega_1)^{-1} \cdot e(g_{4L+3}, \Omega_2)^{-1} \\
e(h_z, Z) \cdot e(h_u, U) &= e(h_1, \sigma_1)^{-1} \cdot e\left(\prod_{i=1}^L h_{2i+M[i]}, \sigma_2\right)^{-1} \cdot \\
&\quad e\left(\prod_{i=1}^L h_{2L+2i+M[i]}, \sigma_3\right)^{-1} \cdot e(h_{4L+2}, \Omega_1)^{-1} \cdot e(h_{4L+3}, \Omega_2)^{-1} .
\end{aligned}$$

We will rely on a non-standard security property of this signature scheme. Specifically, we implicitly prove that no PPT adversary having access to a signing oracle can output a group element $X = g^x$, for some $x \in \mathbb{Z}_p$, along with a tuple $(\sigma_1, \sigma_2, \sigma_3, Z, R, U)$ which forms a valid signature with respect to the modified public key

$$PK = \left(f, g, h, u_1, u_2, \Omega_1^x = u_1^{\omega_1 \cdot x}, \Omega_2^x = u_2^{\omega_2 \cdot x}, \mathbf{V}, \mathbf{W}, \text{pk}_{h_{sps}}, \{(Z_j, R_j, U_j)\}_{j=1}^{4L+2} \right) . \quad (4)$$

Unfortunately, our QA-NIZK proof system of Section 3 cannot rely on this property in a modular way since, in the CRS, the matrix (2) will be mixed with other linear subspaces. We will thus implicitly prove this property “from scratch” in our proof of simulation-soundness. Under the DLIN assumption, we show that the adversary would be able to predict the value of a random function on a non-trivial input if it were able to contradict this property.

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.* [48] with several modifications allowing to tightly relate the simulation-soundness property to the DLIN assumption. The construction also uses the signature scheme of [49] which is recalled in Section 2.4.

3.1 Intuition

Like [48], 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

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 [34] is satisfiable requires a proof length proportional to the number of pairings (which is linear in the dimension n here) in pairing product equations.

would have been unable to compute and solve a DLIN instance. To implement this approach, the system of [48] uses Waters’ partitioning technique [61] in the fashion of [52], which inevitably [38] 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 recalled in Section 2.4.

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 vector $\mathbf{F} = \mathbf{f}_1^{\mu_1} \cdot \mathbf{f}_2^{\mu_2}$ 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, $\mathbf{F}^* \in \text{span}\langle \mathbf{f}_1, \mathbf{f}_2 \rangle$). 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 $\mathbf{F} \in \text{span}\langle \mathbf{f}_1, \mathbf{f}_2 \rangle$; (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 are able to 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 recalled in Section 2.4. 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 of the form (4).

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.* [49] 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 [49] requires quadratic equations and is thus relatively inefficient. 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, no matter how large the dimensions of the subspace are.

3.2 Construction

For simplicity, the description below assumes symmetric pairings $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$. However, as explained in Appendix H, instantiations in asymmetric pairings $e : \mathbb{G} \times \hat{\mathbb{G}} \rightarrow \mathbb{G}_T$ (with $\mathbb{G} \neq \hat{\mathbb{G}}$) are possible under a natural asymmetric analogue of the DLIN assumption. We leave it as an interesting open problem to build an even more efficient scheme under the Symmetric eXternal Diffie-Hellman assumption.

As in [41], 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 $\rho = g^{\mathbf{A}}$.

$\mathbf{K}_0(\lambda)$: 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 .

$\mathbf{K}_1(\Gamma, \rho)$: 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 linearly homomorphic signature of Section 2.3 in order to sign vectors of \mathbb{G}^n and \mathbb{G}^{4L+6} , respectively. Namely, choose generators $g_z, g_r, h_z, h_u \xleftarrow{R} \mathbb{G}$, $G_z, G_r, H_z, H_u \xleftarrow{R} \mathbb{G}$. 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

$$\mathbf{V} = (V_{1,0}, V_{1,1}, \dots, V_{L,0}, V_{L,1}) \in \mathbb{G}^{2L}, \quad \mathbf{W} = (W_{1,0}, W_{1,1}, \dots, W_{L,0}, W_{L,1}) \in \mathbb{G}^{2L}.$$

4. Choose n $\omega_1, \omega_2 \xleftarrow{R} \mathbb{Z}_p$ and $u_1, u_2 \xleftarrow{R} \mathbb{G}$, 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

$$(M_{i,j})_{i,j} = \begin{pmatrix} 1 & \mathbf{1}^{1 \times 2L} & \mathbf{1}^{1 \times 2L} & 1 & 1 & \mathbf{f}_1 \\ 1 & \mathbf{1}^{1 \times 2L} & \mathbf{1}^{1 \times 2L} & 1 & 1 & \mathbf{f}_2 \\ \mathbf{V}^\top & \mathbf{Id}_{f,2L} & \mathbf{1}^{2L \times 2L} & \mathbf{1}^{2L \times 1} & \mathbf{1}^{2L \times 1} & \mathbf{1}^{2L \times 3} \\ \mathbf{W}^\top & \mathbf{1}^{2L \times 2L} & \mathbf{Id}_{h,2L} & \mathbf{1}^{2L \times 1} & \mathbf{1}^{2L \times 1} & \mathbf{1}^{2L \times 3} \\ g & \mathbf{1}^{1 \times 2L} & \mathbf{1}^{1 \times 2L} & u_1 & 1 & \mathbf{1}^{1 \times 3} \\ g & \mathbf{1}^{1 \times 2L} & \mathbf{1}^{1 \times 2L} & 1 & u_2 & \mathbf{1}^{1 \times 3} \\ 1 & \mathbf{1}^{1 \times 2L} & \mathbf{1}^{1 \times 2L} & \Omega_1^{-1} & \Omega_2^{-1} & \mathbf{f}_0 \end{pmatrix} \quad (5)$$

with $\mathbf{Id}_{f,2L} = f^{\mathbf{I}_{2L}} \in \mathbb{G}^{2L \times 2L}$, $\mathbf{Id}_{h,2L} = h^{\mathbf{I}_{2L}} \in \mathbb{G}^{2L \times 2L}$, 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 the LHSPS private key 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 triples $(z_i, r_i, u_i) = (\prod_{j=1}^n G_{i,j}^{-\chi_j}, \prod_{j=1}^n G_{i,j}^{-\gamma_j}, \prod_{j=1}^n G_{i,j}^{-\delta_j})$ 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 (5) and obtain signatures

$$(Z_j, R_j, U_j) = \left(\prod_{k=1}^{4L+6} M_{j,k}^{-\varphi_k}, \prod_{k=1}^{4L+6} M_{j,k}^{-\phi_k}, \prod_{k=1}^{4L+6} M_{j,k}^{-\vartheta_k} \right) \quad j \in \{1, \dots, 4L+5\}.$$

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

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

$$\mathbf{CRS}_2 = \left(\mathbf{f}, \mathbf{f}_0, \text{pk}_0, \text{pk}_1, \Omega_1, \Omega_2, \mathbf{V}, \mathbf{W} \right),$$

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

$\text{P}(T, \psi, \mathbf{v}, x, \text{lbl})$: 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

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

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 (6). 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 $\pi_{\sigma_1}, \pi_{\sigma_2} \in \mathbb{G}^3$ that committed variables (σ_1, Z, R, U) satisfy

$$e(Z, G_z) \cdot e(R, G_r) \cdot e(\sigma_1, G_1) = t_G, \quad e(Z, H_z) \cdot e(U, H_u) \cdot e(\sigma_1, H_1) = t_H, \quad (7)$$

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 $(\pi_1, \pi_2) \in \mathbb{G}^6$ that \mathbf{v} and (z, r, u) satisfy (1). Let $(\mathbf{C}_z, \mathbf{C}_r, \mathbf{C}_u, \pi_1, \pi_2) \in \mathbb{G}^{15}$ be the resulting commitments and proofs.

7. Set $\sigma = \mathcal{S}(\text{SK}, (\mathbf{v}, \mathbf{F}, \mathbf{C}_{\sigma_1}, \sigma_2, \sigma_3, \mathbf{C}_Z, \mathbf{C}_R, \mathbf{C}_U, \mathbf{C}_z, \mathbf{C}_r, \mathbf{C}_u, \pi_{\sigma_1}, \pi_{\sigma_2}, \pi_1, \pi_2, \text{lbl}))$ and output

$$\pi = (\text{VK}, \mathbf{F}, \mathbf{C}_{\sigma_1}, \sigma_2, \sigma_3, \mathbf{C}_Z, \mathbf{C}_R, \mathbf{C}_U, \mathbf{C}_z, \mathbf{C}_r, \mathbf{C}_u, \pi_{\sigma_1}, \pi_{\sigma_2}, \pi_1, \pi_2, \sigma). \quad (8)$$

$\mathbf{V}(\Gamma, \psi, \mathbf{v}, \pi, \text{lbl})$: parse π as in (8) 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) $\mathcal{V}(\text{VK}, (\mathbf{v}, \mathbf{F}, \mathbf{C}_{\sigma_1}, \sigma_2, \sigma_3, \mathbf{C}_Z, \mathbf{C}_R, \mathbf{C}_U, \mathbf{C}_z, \mathbf{C}_r, \mathbf{C}_u, \pi_{\sigma_1}, \pi_{\sigma_2}, \pi_1, \pi_2, \text{lbl}), \sigma) = 1$;
- (ii) $\pi_{\sigma_1}, \pi_{\sigma_2}$ 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 (7).
- (iii) The tuple $(\mathbf{C}_z, \mathbf{C}_r, \mathbf{C}_u, \pi_1, \pi_2)$ forms a valid a valid NIWI proof for the Groth-Sahai CRS $\mathbf{F} = (\mathbf{f}_1, \mathbf{f}_2, \mathbf{F})$. Namely, $\pi_1 = (\pi_{1,1}, \pi_{1,2}, \pi_{1,3})$ and $\pi_2 = (\pi_{2,1}, \pi_{2,2}, \pi_{2,3})$ satisfy

$$\begin{aligned} \prod_{i=1}^n E(g_i, \iota(v_i))^{-1} &= E(g_z, \mathbf{C}_z) \cdot E(g_r, \mathbf{C}_r) \cdot E(\pi_{1,1}, \mathbf{f}_1) \cdot E(\pi_{1,2}, \mathbf{f}_2) \cdot E(\pi_{1,3}, \mathbf{F}) \\ \prod_{i=1}^n E(h_i, \iota(v_i))^{-1} &= E(h_z, \mathbf{C}_z) \cdot E(h_u, \mathbf{C}_u) \cdot E(\pi_{2,1}, \mathbf{f}_1) \cdot E(\pi_{2,2}, \mathbf{f}_2) \cdot E(\pi_{2,3}, \mathbf{F}). \end{aligned} \quad (9)$$

The proof only requires 38 elements of \mathbb{G} and a pair (VK, σ) . In instantiations using the one-time signature of [37], its total size amounts to 42 group elements, which only lengthens the QA-NIZK proofs of [48] 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 Appendix C for a detailed proof that the simulation is perfect. This simulator (S_1, S_2) is defined by having S_1 generate the CRS ψ as in the real 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 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 the first step of the simulation instead of using the actual witness $\mathbf{x} \in \mathbb{Z}_p^t$ as in the real proving algorithm P. At step 2, it defines the vector $(F_1, F_2, F_3) = \mathbf{f}_0 \cdot \mathbf{f}_1^{\mu_1} \cdot \mathbf{f}_2^{\mu_2}$ with $\mu_1, \mu_2 \xleftarrow{R} \mathbb{Z}_p$. At step 3, it randomly picks $r, s \xleftarrow{R} \mathbb{Z}_p$ to compute a triple $(\sigma_1, \sigma_2, \sigma_3) = (g^{\omega_1 + \omega_2} \cdot H(\mathbf{V}, \mathbf{VK})^r \cdot H(\mathbf{W}, \mathbf{VK})^s, f^r, h^s)$ 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 P.

In Appendix C, we prove that the simulation is perfect in that the simulated CRS ψ is distributed as a real CRS and, for all $\mathbf{v} \in \mathbb{G}^n$ for which there exists $\mathbf{x} \in \mathbb{Z}_p^t$ such that $\mathbf{v} = g^{\mathbf{x} \cdot \mathbf{A}}$, 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 in Theorem 2 since the reduction can be made 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 $\text{Adv}_{\mathcal{A}}^{\text{e-sound}}(\lambda) \leq 2 \cdot \text{Adv}_{\mathcal{B}}^{\text{DLIN}}(\lambda) + 2/p$. (The proof is in Appendix D.)*

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*

$$\text{Adv}_{\mathcal{A}}^{\text{e-uss}}(\lambda) \leq \text{Adv}_{\mathcal{B}'}^{q\text{-suf-ots}}(\lambda) + 3 \cdot (L + 2) \cdot \text{Adv}_{\mathcal{B}}^{\text{DLIN}}(\lambda) + 4/p, \quad (10)$$

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 $S_2(\psi, \tau_{sim}, \cdot, \cdot)$ which is described in the proof of Theorem 1. At each invocation, $S_2(\psi, \tau_{sim}, \cdot, \cdot)$ 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 $\psi = (\mathbf{CRS}_1, \mathbf{CRS}_2)$. 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 $|\Pr[S_3] - \Pr[S_2]| \leq 2 \cdot \text{Adv}^{\text{DLIN}}(\mathcal{B})$.

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 VK^* that is recycled from an output of the $S_2(\psi, \tau_{sim}, \dots)$ 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 [37]), the probability of this event can be bounded independently of the number q of queries to $S_2(\psi, \tau_{sim}, \dots)$. We have $|\Pr[S_4] - \Pr[S_3]| \leq \text{Adv}_{\mathcal{B}}^{q\text{-suf-ots}}(\lambda)$.

Game₅: This game is identical to **Game₄** but we raise a failure event E_5 . When \mathcal{A} outputs its fake proof $\pi^* = (\text{VK}^*, \mathbf{F}^*, \mathbf{C}_{\sigma_1}^*, \sigma_2^*, \sigma_3^*, \mathbf{C}_Z^*, \mathbf{C}_R^*, \mathbf{C}_U^*, \mathbf{C}_z^*, \mathbf{C}_r^*, \mathbf{C}_u^*, \pi_{\sigma,1}^*, \pi_{\sigma,2}^*, \pi_1^*, \pi_2^*, \sigma^*)$, \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 the equality $F_3^* \neq F_1^{*1/y_1} \cdot F_2^{*1/y_2}$ holds, 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 $\Pr[E_5] \leq (2 \cdot L + 1) \cdot \text{Adv}_{\mathcal{B}}^{\text{DLIN}}(\lambda) + 2/p$, 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.

⁷ This notion (see Definition 4 in [37]) is defined via a game where the adversary is given q verification keys $\{\text{VK}_i\}_{i=1}^q$ 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 VK_{i^*} . Hofheinz and Jager [37, Section 4.2] gave a discrete-log-based one-time signature with tight security in the multi-key setting.

In Game_5 , 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 the probability of event S_5 is at most $\Pr[S_5] \leq (L + 1) \cdot \mathbf{Adv}_{\mathcal{B}}^{\text{DLIN}}(\lambda) + \frac{1}{p} + \Pr[S_5 \wedge \neg E_5]$.

In Game_5 , 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 $\mathbf{S}_2(\psi, \tau_{sim}, \dots)$ 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 trapdoor $(y_1, y_2) = (\log_g(f_1), \log_g(f_2))$ to recover (z^*, r^*, u^*) from their commitments $(\mathbf{C}_z^*, \mathbf{C}_r^*, \mathbf{C}_u^*)$ by performing BBS decryptions. Indeed, $\mathbf{C}_z^* = \iota(z^*) \cdot \mathbf{f}_1^{\theta_{z,1}} \cdot \mathbf{f}_2^{\theta_{z,2}} \cdot \mathbf{F}^{*\theta_{z,3}}$ is of the form $\mathbf{C}_z^* = \iota(z^*) \cdot \mathbf{f}_1^{\theta_{z,1} + \mu_1^* \cdot \theta_{z,3}} \cdot \mathbf{f}_2^{\theta_{z,2} + \mu_2^* \cdot \theta_{z,3}}$, 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

$$e(g_z, z^*) \cdot e(g_r, r^*) \cdot \prod_{i=1} e(g_i, v_i^*) = e(h_z, z^*) \cdot e(h_u, u^*) \cdot \prod_{i=1} e(h_i, v_i^*) = 1_{\mathbb{G}_T} . \quad (11)$$

In addition, \mathcal{B} computes $(z^\dagger, r^\dagger, u^\dagger) = (\prod_{i=1}^n v_i^{*-x_i}, \prod_{i=1}^n v_i^{*-y_i}, \prod_{i=1}^n v_i^{*-\delta_i})$, which also satisfies the equations (11). Since $(z^\dagger, r^\dagger, u^\dagger)$ and (z^*, r^*, u^*) both satisfy (11), 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 $\mathbf{S}_2(\psi, \tau_{sim}, \dots)$. Indeed, in simulated proofs, the only components that depend on (χ_1, \dots, χ_n) are $(\mathbf{C}_z, \mathbf{C}_r, \mathbf{C}_u, \pi_1, \pi_2)$, which are generated for a perfectly NIWI Groth-Sahai CRS $(\mathbf{f}_1, \mathbf{f}_2, \mathbf{F})$. Consequently, the same arguments as in [47, 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 the inequality $\Pr[S_5 \wedge \neg E_5] \leq \mathbf{Adv}_{\mathcal{B}}^{\text{SDP}}(\lambda) + 1/p$, which yields the bound (10) 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 [31]) of the signature scheme recalled in Section 2.4, 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 [49, Appendix G]) that the signature scheme of Section 2.4 provides tight security in the multi-user setting.

Lemma 1. *In Game₅, there is a DLIN distinguisher \mathcal{B} such that the probability of event 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 in Appendix E.)*

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) Public-Key Encryption with Shorter Ciphertexts

Like [37,49], our scheme builds on the Naor-Yung paradigm [55] and the encryption scheme of Boneh, Boyen and Shacham (BBS) [15].

In short, the encryption phase computes BBS ciphertexts $(C_0, C_1, C_2) = (M \cdot g^{\theta_1 + \theta_2}, X_1^{\theta_1}, Y_1^{\theta_2})$ 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 [49], our reduction is not quite as tight as in [37,4] 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.* [49] incurs 69 group elements per ciphertext. Our technique thus improves upon [49] by 30% and also outperforms the most efficient perfectly tight solution [4], 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

$$\begin{aligned} \mathbf{CRS}_1 &= \left(\rho, \mathbf{f}, \mathbf{f}_0, \{u_i\}_{i=1}^2, \{\Omega_i\}_{i=1}^2, \mathbf{V}, \mathbf{W}, \{\text{pk}_i\}_{i=1}^2, \{(z_i, r_i, u_i)\}_{i=1}^4, \{(Z_j, R_j, U_j)\}_{j=1}^{4L+5} \right), \\ \mathbf{CRS}_2 &= \left(\mathbf{f}, \mathbf{f}_0, \{\text{pk}_i\}_{i=1}^2, \{\Omega_i\}_{i=1}^2, \mathbf{V}, \mathbf{W} \right). \end{aligned}$$

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

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

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 $(C_0, C_1, C_2) = (M \cdot g^{\theta_1 + \theta_2}, X_1^{\theta_1}, Y_1^{\theta_2})$ as well as $(D_0, D_1, D_2) = (M \cdot g^{\theta_3 + \theta_4}, X_2^{\theta_3}, Y_2^{\theta_4})$.
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 simulation-sound QA-NIZK proof is

$$\pi = (\text{VK}, \mathbf{F}, \mathbf{C}_{\sigma_1, \sigma_2, \sigma_3}, \mathbf{C}_Z, \mathbf{C}_R, \mathbf{C}_U, \mathbf{C}_z, \mathbf{C}_r, \mathbf{C}_u, \pi_{\sigma_1}, \pi_{\sigma_2}, \pi_1, \pi_2, \sigma) .$$

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 [37], the ciphertext size amounts to that of 48 group elements, instead of 69 in [49].

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, as explained in Appendix H.

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 in Appendix F.)

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 [37, 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., [30]), the public verifiability of ciphertexts makes our scheme amenable for non-interactive threshold decryption in a static corruption model. Like many other pairing-based CCA2-secure threshold cryptosystems (e.g., [16, 46, 5]), the resulting scheme can be made robust against malicious adversaries in a simple manner, by leveraging the verifiability properties enabled by bilinear groups.

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 Appendix G.

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

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

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 Appendix 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 [26,20] are commitment schemes that provably remain secure when composed with arbitrary other protocols. They are 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 [51], 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

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 $(f, g, h, f^\alpha, h^\beta, \eta)$, where $\eta = g^{a+b}$ or $\eta \in_R \mathbb{G}$, the challenge ciphertext is computed as $(C_1, C_2, C_3) = (f^\alpha, h^\beta, M_\beta \cdot (f^\alpha)^\alpha \cdot (h^\beta)^\beta \cdot \eta^\gamma)$.

static adversaries that have to choose whom to corrupt upfront⁹. Subsequently, Fischlin *et al.* [29] 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 [41] 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 [26,20] that allows committing to long messages with expansion rate 2. Note that, in constructions of UC commitments from IND-CCA2-secure encryption (e.g., [20,29,41]), 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 [29,41] can be optimized using the techniques of [48,42] 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 [41], 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 $\text{span}\langle \mathbf{X}_{i,1}, \mathbf{X}_{i,2}, \mathbf{Y}_{i,1}, \mathbf{Y}_{i,2} \rangle_{i=1}^N$. 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.* [33], it retains adaptive security assuming reliable erasures while [33] is only known to be secure against static adversaries.

References

1. 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.
2. 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.
3. M. Abe, K. Haralambiev, M. Ohkubo. Signing on elements in bilinear groups for modular protocol design. In Cryptology ePrint Archive: Report 2010/133, 2010.
4. 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.
5. S. Arita, K. Tsurudome. Construction of Threshold Public-Key Encryption Through Tag-Based Encryptions. In *ACNS '09*, LNCS 5536, pp. pp. 186–200, 2009.
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. on Information 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.

⁹ Lindell’s commitment can actually be made adaptively secure (modulo a patch [12]), but even its optimized variant [12] remains interactive with 3 rounds of communication during the commitment phase.

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. 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.
12. 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.
13. 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.
14. M. Blum, P. Feldman, S. Micali. Non-interactive zero-knowledge and its applications. In *STOC '88*, pp. 103–112, ACM Press, 1988.
15. D. Boneh, X. Boyen, H. Shacham. Short group signatures. In *Crypto '04*, LNCS 3152, pp. 41–55, Springer, 2004.
16. D. Boneh, X. Boyen, S. Halevi. Chosen-Ciphertext Secure Threshold Public-Key Encryption Without Random Oracles. In *CT-RSA '06*, LNCS 3860, pp. 226–243, Springer, 2006.
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. 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.
22. J. Chen, H. Wee. Fully, (almost) tightly secure IBE and dual system groups. In *Crypto '13*, LNCS 8043, pp. 435–460, Springer, 2013.
23. B. Chevallier-Mames. An efficient CDH-based signature scheme with a tight security reduction. In *Crypto '05*, LNCS 3621, pp. 511–526, Springer, 2005.
24. J.-S. Coron. On the exact security of full domain hash. In *Crypto '00*, LNCS 1880, pp. 229–235, Springer, 2000.
25. J.-S. Coron. Optimal security proofs for PSS and other signature schemes. In *Eurocrypt '02*, LNCS 2332, pp. 229–235, Springer, 2002.
26. R. Canetti. Universally Composable Security: A New Paradigm for Cryptographic Protocols. In *FOCS'01*, pp. 136–145, 2001.
27. 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.
28. 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.
29. 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.
30. P.-A. Fouque, D. Pointcheval. Threshold Cryptosystems Secure against Chosen-Ciphertext Attacks. In *Asiacrypt '01*, LNCS 2248, pp. 351–368, Springer, 2001.
31. 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.
32. S. Galbraith, K. Paterson, N. Smart. Pairings for cryptographers. In *Discrete Applied Mathematics*, 156(16), pp. 3113–3121, 2008.
33. J. Garay, Y. Ishai, R. Kumaresan, H. Wee. On the Complexity of UC Commitments. In *Eurocrypt '14*, LNCS 8441, pp. 677–694, Springer, 2014.
34. 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.
35. J. Groth, A. Sahai. Efficient non-interactive proof systems for bilinear groups. In *Eurocrypt '08*, LNCS 4965, pp. 415–432, Springer, 2008.
36. D. Hofheinz. All-but-many lossy trapdoor functions. In *Eurocrypt '12*, LNCS 7237, pp. 209–227, Springer, 2012.
37. D. Hofheinz, T. Jager. Tightly secure signatures and public-key encryption. In *Crypto '12*, LNCS 7417, pp. 590–607, Springer, 2012.
38. D. Hofheinz, T. Jager, E. Knapp. Waters signatures with optimal security reduction. In *PKC '12*, LNCS 7293, pp. 66–83, Springer, 2012.

39. D. Hofheinz. Circular chosen-ciphertext security with compact ciphertexts. In *Eurocrypt '13*, LNCS 7881, pp. 520–536, Springer, 2013.
40. C. Jutla, A. Roy. Relatively-sound NIZKs and password-based key-exchange. In *PKC '12*, LNCS 7293, pp. 485–503, Springer, 2012.
41. 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.
42. 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.
43. S. Kakvi, E. Kiltz. Optimal security proofs for full domain hash, revisited. In *Eurocrypt '12*, LNCS 7237, pp. 537–553, Springer, 2012.
44. J. Katz, V. Vaikuntanathan. Round-optimal password-based authenticated key exchange. In *TCC '11*, LNCS 6597, pp. 293–310, Springer, 2011.
45. J. Katz, N. Wang. Efficiency improvements for signature schemes with tight security reductions. In *ACM-CCS '03*, pp. 155–164, ACM Press, 2003.
46. E. Kiltz. Chosen-Ciphertext Security from Tag-Based Encryption. In *TCC '06*, Springer, 2006.
47. 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.
48. 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.
49. 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.
50. 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.
51. Y. Lindell. Highly-efficient universally-composable commitments based on the DDH assumption. In *Eurocrypt '11*, LNCS 6632, pp. 446–466, Springer, 2011.
52. 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.
53. M. Naor, O. Reingold. Number-theoretic constructions of efficient pseudo-random functions. In *FOCS '97*, pp. 458–467, IEEE Press, 1997.
54. M. Naor. On cryptographic assumptions and challenges. In *Crypto '03*, LNCS 2729, pp. 96–109, Springer, 2003.
55. M. Naor, M. Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. In *STOC '90*, ACM Press, 1990.
56. P. Paillier. Public-key cryptosystems based on composite degree residuosity classes. In *Eurocrypt '99*, LNCS 1592, pp. 223–238, Springer, 1999.
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.

A Groth-Sahai Non-Interactive Proof Systems

In the following, vectors are always considered as row vectors unless stated otherwise. For any generator $g \in \mathbb{G}$, we denote by $\iota(g) \in \mathbb{G}^3$ the vector $(1_{\mathbb{G}}, 1_{\mathbb{G}}, g)$.

In their DLIN-based assumption in symmetric pairings, the Groth-Sahai (GS) proof systems [35] use a common reference string (CRS) made of vectors $\mathbf{f}_1, \mathbf{f}_2, \mathbf{f}_3 \in \mathbb{G}^3$, where $\mathbf{f}_1 = (f_1, 1, g)$, $\mathbf{f}_2 = (1, f_2, g)$ for some $f_1, f_2 \in \mathbb{G}$. A commitment to a group element $X \in \mathbb{G}$ is obtained by computing

$\mathbf{C} = (1, 1, X) \cdot \mathbf{f}_1^r \cdot \mathbf{f}_2^s \cdot \mathbf{f}_3^t$ with $r, s, t \xleftarrow{R} \mathbb{Z}_p$. In order to have perfectly sound proofs, \mathbf{f}_3 is chosen as $\mathbf{f}_3 = \mathbf{f}_1^{\xi_1} \cdot \mathbf{f}_2^{\xi_2}$ with $\xi_1, \xi_2 \xleftarrow{R} \mathbb{Z}_p$. In this case, commitments $\mathbf{C} = (f_1^{r+\xi_1 t}, f_2^{s+\xi_2 t}, X \cdot g^{r+s+t(\xi_1+\xi_2)})$ are nothing but Boneh-Boyen-Shacham (BBS) ciphertexts as X can be recovered using the BBS private key $(\alpha_1, \alpha_2) = (\log_g(f_1), \log_g(f_2))$. In order to have perfectly witness indistinguishable (WI) proofs, $\mathbf{f}_1, \mathbf{f}_2$ and \mathbf{f}_3 are linearly independent vectors, so that \mathbf{C} is a perfectly hiding commitment to $X \in \mathbb{G}$: a typical choice is $\mathbf{f}_3 = \mathbf{f}_1^{\xi_1} \cdot \mathbf{f}_2^{\xi_2} \cdot \iota(g)^{-1}$. Under the DLIN assumption, the two distributions of CRS are computationally indistinguishable.

In the perfect and computational NIWI settings, efficient proofs are available for pairing-product equations, which are relations of the type

$$\prod_{i=1}^n e(\mathcal{A}_i, \mathcal{X}_i) \cdot \prod_{i=1}^n \prod_{j=1}^n e(\mathcal{X}_i, \mathcal{X}_j)^{a_{ij}} = t_T,$$

for variables $\mathcal{X}_1, \dots, \mathcal{X}_n \in \mathbb{G}$ and constants $t_T \in \mathbb{G}_T$, $\mathcal{A}_1, \dots, \mathcal{A}_n \in \mathbb{G}$, $a_{ij} \in \mathbb{Z}_p$, for $i, j \in \{1, \dots, n\}$. Non-interactive proofs for linear equations (with $a_{ij} = 0$ for all i, j) only require 3 group elements.

B Definitions for Linearly Homomorphic Structure-Preserving Signatures

Let $(\mathbb{G}, \mathbb{G}_T)$ be groups of prime order p such that a bilinear map $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$ can be efficiently computed.

A signature scheme is *structure-preserving* [3,2] if messages, signatures and public keys all live in the group \mathbb{G} . In linearly homomorphic structure-preserving signatures, the message space \mathcal{M} consists of pairs $\mathcal{M} := \mathcal{T} \times \mathbb{G}^n$, for some $n \in \mathbb{N}$, where \mathcal{T} is a tag space. Depending on the application, one may want the tags to be group elements or not. In this paper, they can be arbitrary strings.

Definition 3. A linearly homomorphic structure-preserving signature scheme over $(\mathbb{G}, \mathbb{G}_T)$ is a tuple of efficient algorithms $\Sigma = (\text{Keygen}, \text{Sign}, \text{SignDerive}, \text{Verify})$ for which the message space consists of $\mathcal{M} := \mathcal{T} \times \mathbb{G}^n$, for some integer $n \in \text{poly}(\lambda)$ and some set \mathcal{T} , and with the following specifications.

Keygen(λ, n) is a randomized algorithm that takes in a security parameter $\lambda \in \mathbb{N}$ and an integer $n \in \text{poly}(\lambda)$ denoting the dimension of vectors to be signed. It outputs a key pair (pk, sk) , where pk includes the description of a tag space \mathcal{T} , where each tag serves as a file identifier.

Sign($\text{sk}, \tau, \mathbf{M}$) takes as input a private key sk , a file identifier $\tau \in \mathcal{T}$ and a vector of group elements $\mathbf{M} = (M_1, \dots, M_n) \in \mathbb{G}^n$. It outputs a signature $\sigma \in \mathbb{G}^{n_s}$, for some $n_s \in \text{poly}(\lambda)$.

SignDerive($\text{pk}, \tau, \{(\omega_i, \sigma^{(i)})\}_{i=1}^{\ell}$) is a homomorphic signature derivation algorithm. It inputs a public key pk , a file identifier τ as well as ℓ pairs $(\omega_i, \sigma^{(i)})$, each of which consists of a coefficient $\omega_i \in \mathbb{Z}_p$ and a signature $\sigma^{(i)} \in \mathbb{G}^{n_s}$. It outputs a signature $\sigma \in \mathbb{G}^{n_s}$ on the vector $\mathbf{M} = \prod_{i=1}^{\ell} M_i^{\omega_i}$, where $\sigma^{(i)}$ is a signature on M_i .

Verify($\text{pk}, \tau, \mathbf{M}, \sigma$) is a deterministic verification algorithm that takes as input a public key pk , a file identifier $\tau \in \mathcal{T}$, a signature σ and a vector $\mathbf{M} = (M_1, \dots, M_n)$. It outputs 0 or 1 depending on whether σ is deemed valid or not.

In a *one-time* linearly homomorphic SPS, the tag τ can be omitted in the specification as a given key pair (pk, sk) only allows signing one linear subspace.

As in all linearly homomorphic signatures, the desired security notion mandates the adversary's inability to come up with a valid triple $(\tau^*, \mathbf{M}^*, \sigma^*)$ for a new file identifier τ^* or, if τ^* appeared in signatures generated by the signing oracle, for a vector \mathbf{M}^* outside the linear span of the vectors that have been legitimately signed for the tag τ^* .

C Proof of Theorem 1

Proof. To prove the quasi-adaptive zero-knowledge property, we consider a sequence of three games which begins with a game where the adversary has access to a real prover P on a real CRS ψ . In the last game, the adversary interacts with a simulator (S_1, S_2) .

Game₁: is a game where the adversary \mathcal{A} is given the description of the language \mathcal{L}_ρ and is granted access to a real CRS ψ and an actual prover $P(\psi, \cdot, \cdot)$ which takes as input a vector \mathbf{v} along with a witness $\mathbf{x} \in \mathbb{Z}_p^t$ such that $\mathbf{v} = g^{\mathbf{x} \cdot \mathbf{A}}$. At each invocation, the oracle outputs a genuine proof π by running the legal P algorithm. The adversary is allowed to query $P(\psi, \cdot, \cdot)$ a polynomial number of times and eventually outputs a bit $\beta \in \{0, 1\}$. We denote by S_1 the event that $\beta = 1$.

Game₂: This game like **Game₁** with the difference that, when the $P(\psi, \cdot, \cdot)$ oracle is queried on a pair (\mathbf{v}, \mathbf{x}) , the witness $\mathbf{x} \in \mathbb{Z}_p^t$ is no longer used at step 1 of the proving algorithm. Instead, $P(\psi, \cdot, \cdot)$ uses the homomorphic signature's private key $\{(\chi_i, \gamma_i, \delta_i)\}_{i=1}^n$ to compute a one-time homomorphic signature $(z, r, u) = (\prod_{j=1}^n v_j^{-\chi_j}, \prod_{j=1}^n v_j^{-\gamma_j}, \prod_{j=1}^n v_j^{-\delta_j})$ on the input vector $\mathbf{v} = (v_1, \dots, v_n) \in \mathbb{G}^n$. All other parts of π are generated as in the real $P(\psi, \cdot, \cdot)$ oracle in steps 2-7 of the proof generation algorithm. Although, the witness $\mathbf{x} \in \mathbb{Z}_p^t$ is not used at any time, it is easy to see that (z, r, u) has exactly the same distribution as in **Game₁** if $\mathbf{v} \in \mathcal{L}_\rho$ (i.e., as long as there exists $\mathbf{x} \in \mathbb{Z}_p^t$ such that $\mathbf{v} = g^{\mathbf{x} \cdot \mathbf{A}}$). We thus have $\Pr[S_2] = \Pr[S_1]$.

Game₃: In this game, we bring another modification to the $P(\psi, \cdot, \cdot)$ oracle. Namely, steps 2 and 3 are conducted as follows. At step 2, the modified oracle $P(\psi, \cdot, \cdot)$ chooses $\mu_1, \mu_2 \xleftarrow{R} \mathbb{Z}_p$ and computes the vector $\mathbf{F} \in \mathbb{G}^3$ as $\mathbf{F} = (F_1, F_2, F_3) = \mathbf{f}_0 \cdot \mathbf{f}_1^{\mu_1} \cdot \mathbf{f}_2^{\mu_2}$ instead of $(F_1, F_2, F_3) = \mathbf{f}_1^{\mu_1} \cdot \mathbf{f}_2^{\mu_2}$ as previously. At Step 3, $P(\psi, \cdot, \cdot)$ uses $\omega_1, \omega_2 \in \mathbb{Z}_p$ to compute

$$(\sigma_1, \sigma_2, \sigma_3) = (g^{\omega_1 + \omega_2} \cdot H(\mathbf{V}, \mathbf{VK})^r \cdot H(\mathbf{W}, \mathbf{VK})^s, f^r, h^s).$$

Observe that the vector (6) is no longer confined in the row space of the first $4L + 2$ rows of $\mathbf{M} \in \mathbb{G}^{(4L+5) \times (4L+6)}$ as it now depends on all rows. At step 4 of the simulated P , however, the simulator can still compute a signature (Z, R, U) on the vector (6) by deriving it from $\{(Z_j, R_j, U_j)\}_{j=1}^{4L+5}$ via the coefficients $\mu_1, \mu_2, r, s, \omega_1, \omega_2, 1 \in \mathbb{Z}_p$. We claim that these changes do not affect the distribution of proofs π whatsoever. Indeed, since we have $\mathbf{f}_0 \in \text{span}\langle \mathbf{f}_1, \mathbf{f}_2 \rangle$, \mathbf{F} remains uniform in $\text{span}\langle \mathbf{f}_1, \mathbf{f}_2 \rangle$. Although the joint distribution of $(\sigma_1, \sigma_2, \sigma_3)$ has changed, the adversary only gets to see perfectly hiding Groth-Sahai commitments $\mathbf{C}_{\sigma_1}, \mathbf{C}_Z, \mathbf{C}_R, \mathbf{C}_U$ to (σ_1, Z, R, U) and perfectly NIWI proofs $\pi_{\sigma_1}, \pi_{\sigma_2} \in \mathbb{G}^3$, which retain the same distribution as in **Game₃** (recall that these are generated for a perfectly witness indistinguishable CRS $(\mathbf{f}_1, \mathbf{f}_2, \mathbf{f}_3)$). Hence, even an unbounded adversary is unable to tell if the committed signature (Z, R, U) authenticates a partially committed vector $\boldsymbol{\sigma}$ (6) that belongs to the subspace of the first $4L + 2$ rows of \mathbf{M} or the entire row space. It comes that $\Pr[S_3] = \Pr[S_2]$.

The simulator (S_1, S_2) lets S_1 generate ψ as in **Game₃** (so that ψ is distributed as the real CRS). As for S_2 , it generates proofs without using $\mathbf{x} \in \mathbb{Z}_p^t$ as in **Game₃**. Specifically, S_2 uses $\{(\chi_i, \gamma_i, \delta_i)\}_{i=1}^n$ to

compute (z, r, u) at the first step of the simulation. 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 randomly picks $r, s \xleftarrow{R} \mathbb{Z}_p$ to compute a triple $(\sigma_1, \sigma_2, \sigma_3) = (g^{\omega_1 + \omega_2} \cdot H(\mathbf{V}, \mathbf{VK})^r \cdot H(\mathbf{W}, \mathbf{VK})^s, f^r, h^s)$ before using the coefficients $\mu_1, \mu_2, r, s, \omega_1, \omega_2, 1 \in \mathbb{Z}_p$ to derive (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 algorithm P .

The system is perfectly quasi-adaptive zero-knowledge since the simulated CRS ψ is distributed as a real CRS and, for all vectors $\mathbf{v} \in \mathbb{G}^n$ in the row space of $g^{\mathbf{A}}$, simulated proofs have exactly the same distribution as real proofs. \square

D Proof of Theorem 2

Proof. The result immediately follows from the unforgeability of the linearly homomorphic signature recalled in Section 2.3, which was proved in [47, Theorem 1], but we give it for completeness.

We first note that the challenger \mathcal{B} can efficiently detect when the adversary wins since it can generate $\rho = g^{\mathbf{A}}$ by first choosing $\mathbf{A} \in \mathbb{Z}_p^{t \times n}$ given that the distribution \mathcal{D}_Γ is efficiently witness-samplable. Knowing \mathbf{A} , \mathcal{B} can thus recognize when the adversary outputs a vector $\mathbf{v}^* \notin \mathcal{L}_\rho$. We first show that, as long as the DLIN assumption holds, if the adversary \mathcal{A} can break the enhanced quasi-adaptive soundness with non-negligible probability in the real attack game, called Game_0 hereunder, the same holds in a modified game called Game_1 .

Game₁: It proceeds identically to Game_0 except that the Groth-Sahai CRS $\mathbf{f} = (\mathbf{f}_1, \mathbf{f}_2, \mathbf{f}_3)$ is now chosen as a perfectly binding Groth-Sahai CRS, where $\mathbf{f}_3 = \mathbf{f}_1^{\xi_1} \cdot \mathbf{f}_2^{\xi_2}$ for randomly chosen $\xi_1, \xi_2 \xleftarrow{R} \mathbb{Z}_p$, which amounts to setting $\xi_3 = 0$ at step 2 of K_1 . If \mathcal{A} wins with significantly different probabilities in Game_0 and Game_1 , the challenger can be turned into a DLIN distinguisher. More precisely, we have the inequality $|\Pr[\mathcal{A} \text{ wins } \mathsf{Game}_0] - \Pr[\mathcal{A} \text{ wins } \mathsf{Game}_1]| \leq \mathbf{Adv}^{\text{DLIN}}(\lambda)$.

In Game_1 , we build an algorithm \mathcal{B} that uses the soundness adversary \mathcal{A} to break the security of the LHSPS system of [47]. At the very beginning of the soundness game, \mathcal{B} flips a fair coin $d \xleftarrow{R} \{0, 1\}$ which will define its strategy depending on whether \mathcal{A} 's fake proof π^* is expected to involve a perfectly hiding (for $d = 1$) or a perfectly binding (with $d = 0$) Groth-Sahai CRS $\mathbf{F}^* = (\mathbf{f}_1, \mathbf{f}_2, \mathbf{F}^*)$.

- If $d = 0$, \mathcal{B} decides to interact with a challenger for an instance of the linearly homomorphic signature allowing to sign vectors of dimension n . Algorithm \mathcal{B} thus receives a LHSPS public key $\mathbf{pk}_0 = (g_z, g_r, h_z, h_u, \{g_i, h_i\}_{i=1}^n)$ from its challenger and uses it to build a CRS $\psi = (\mathbf{CRS}_1, \mathbf{CRS}_2)$ for \mathcal{A} . To this end, \mathcal{B} first generates $\rho \in \mathbb{G}^{t \times n}$ by choosing a matrix $\mathbf{A} \in \mathbb{Z}_p^{t \times n}$ with the appropriate distribution (recall that \mathcal{D}_Γ is efficiently witness-samplable) and computes $\rho = g^{\mathbf{A}}$. Having defined ρ , \mathcal{B} queries its challenger in order to obtain homomorphic signatures $\{(z_i, r_i, u_i)\}_{i=1}^t$ on the rows of ρ . Then, \mathcal{B} generates the rest of $\psi = (\mathbf{CRS}_1, \mathbf{CRS}_2)$ as in the real K_1 algorithm. In particular, \mathcal{B} chooses the second LHSPS key pair $(\mathbf{sk}_1, \mathbf{pk}_1)$ itself, which allows it to compute homomorphic signatures $\{(Z_j, R_j, U_j)\}_{j=1}^{4L+5}$ on the rows of $\mathbf{M} \in \mathbb{G}^{(4L+5) \times (4L+6)}$. At step 2 of K_1 , \mathcal{B} also defines $(f_1, f_2) = (g^{y_1}, g^{y_2})$ and $\mathbf{f}_0 = \mathbf{f}_1^{\nu_1} \cdot \mathbf{f}_2^{\nu_2}$, for randomly chosen $y_1, y_2, \nu_1, \nu_2 \xleftarrow{R} \mathbb{Z}_p$, where $\mathbf{f}_1 = (f_1, 1, g)$ and $\mathbf{f}_2 = (1, f_2, g)$. Steps 3-6 are conducted as in the real algorithm K_1 and the adversary is given $\psi = (\mathbf{CRS}_1, \mathbf{CRS}_2)$ as well as the matrix $\mathbf{A} = \log_g(\rho) \in \mathbb{Z}_p^{t \times n}$ which serves as a trapdoor τ_m for testing membership of \mathcal{L}_ρ .

By hypothesis, \mathcal{A} must be able to create a fake proof π^* for a vector $\mathbf{v}^* \notin \mathcal{L}_\rho$. At this point, \mathcal{B} parses π^* as $(\mathbf{VK}^*, \mathbf{F}^*, \mathbf{C}_{\sigma_1}^*, \sigma_2^*, \sigma_3^*, \mathbf{C}_Z^*, \mathbf{C}_R^*, \mathbf{C}_U^*, \mathbf{C}_z^*, \mathbf{C}_r^*, \mathbf{C}_u^*, \pi_{\sigma_1}^*, \pi_{\sigma_2}^*, \pi_1^*, \pi_2^*, \sigma^*)$ and the underlying $\mathbf{F}^* \in \mathbb{G}^3$ as (F_1^*, F_2^*, F_3^*) . Then, \mathcal{B} aborts if $F_3^* \neq F_1^{*1/y_1} \cdot F_2^{*1/y_2}$. Otherwise, there exists $\nu_1^*, \nu_2^* \in \mathbb{Z}_p$ such that $(F_1^*, F_2^*, F_3^*) = \mathbf{f}_1^{\nu_1^*} \cdot \mathbf{f}_2^{\nu_2^*}$, which means that $(\mathbf{f}_1, \mathbf{f}_2, \mathbf{F}^*)$ is a perfectly binding Groth-Sahai CRS. In this case, \mathcal{B} can use $(y_1, y_2) \in \mathbb{Z}_p^*$ as an extraction trapdoor to extract (z^*, r^*, u^*) from the Groth-Sahai commitments $\mathbf{C}_z^*, \mathbf{C}_r^*, \mathbf{C}_u^*$ and the perfect soundness of π_1^*, π_2^* ensures that they satisfy

$$e(g_z, z^*) \cdot e(g_r, r^*) \cdot \prod_{i=1}^n e(g_i, v_i^*) = 1_{\mathbb{G}_T},$$

$$e(h_z, z^*) \cdot e(h_u, r^*) \cdot \prod_{i=1}^n e(h_i, v_i^*) = 1_{\mathbb{G}_T}.$$

Since the vector $\mathbf{v}^* = (v_1^*, \dots, v_n^*) \in \mathbb{G}^n$ is linearly independent of the rows of $\boldsymbol{\rho} = g^{\mathbf{A}}$ by hypothesis, \mathcal{B} wins the game against its LHSPS challenger by outputting \mathbf{v}^* and (z^*, r^*, u^*) .

- If $d = 1$, \mathcal{B} interacts with a challenger for an instance of the homomorphic signature allowing to sign vectors of \mathbb{G}^{4L+6} . It receives as input a public key $\mathbf{pk}_1 = (G_z, G_r, H_z, H_u, \{(G_i, H_i)\}_{i=1}^{4L+6})$ from its LHSPS challenger. Then, it faithfully conducts steps 2-5 of \mathbf{K}_1 and queries its challenger in order to obtain homomorphic signatures $\{(Z_j, R_j, U_j)\}_{j=1}^{4L+5}$ on the rows of the matrix $\mathbf{M} \in \mathbb{G}^{(4L+5) \times (4L+6)}$. It also generates $(\mathbf{sk}_0, \mathbf{pk}_0)$ on its own and uses \mathbf{sk}_0 to compute homomorphic signatures $\{(z_i, r_i, u_i)\}_{i=1}^t$ on the rows of $\boldsymbol{\rho} \in \mathbb{G}^{t \times n}$. The resulting $\psi = (\mathbf{CRS}_1, \mathbf{CRS}_2)$ is given as input to \mathcal{A} along with the matrix $\mathbf{A} = \log_g(\boldsymbol{\rho}) \in \mathbb{Z}_p^{t \times n}$.

When \mathcal{A} halts, it outputs a fake proof π^* for a vector $\mathbf{v}^* \in \mathbb{G}^n$ outside \mathcal{L}_ρ . Then, \mathcal{B} parses π^* as $(\mathbf{VK}^*, \mathbf{F}^*, \mathbf{C}_{\sigma_1}^*, \sigma_2^*, \sigma_3^*, \mathbf{C}_Z^*, \mathbf{C}_R^*, \mathbf{C}_U^*, \mathbf{C}_z^*, \mathbf{C}_r^*, \mathbf{C}_u^*, \pi_{\sigma_1}^*, \pi_{\sigma_2}^*, \pi_1^*, \pi_2^*, \sigma^*)$ and \mathbf{F}^* as $(F_1^*, F_2^*, F_3^*) \in \mathbb{G}^3$. At this point, \mathcal{B} aborts in the event that $F_3^* \neq F_1^{*1/y_1} \cdot F_2^{*1/y_2}$. Otherwise (namely, if (F_1^*, F_2^*, F_3^*) is not a BBS encryption of $1_{\mathbb{G}}$), we know that $\mathbf{F}^* = (F_1^*, F_2^*, F_3^*)$ is linearly independent of the vectors $(\mathbf{f}_1, \mathbf{f}_2, \mathbf{f}_0)$ since $\mathbf{f}_0 \in \text{span}\langle \mathbf{f}_1, \mathbf{f}_2 \rangle$ and any linear combination of these is a BBS encryption of $1_{\mathbb{G}}$. Hence, since $\mathbf{f} = (\mathbf{f}_1, \mathbf{f}_2, \mathbf{f}_3)$ is an extractable Groth-Sahai CRS, \mathcal{B} can use the discrete logarithms $(y_1, y_2) = (\log_g(f_1), \log_g(f_2))$ to extract $(\sigma_1^*, Z^*, R^*, U^*)$ from their Groth-Sahai commitments $\mathbf{C}_{\sigma_1}^*, \mathbf{C}_Z^*, \mathbf{C}_R^*, \mathbf{C}_U^*$. The perfect soundness of $\pi_{\sigma_1}^*, \pi_{\sigma_2}^*$ for the CRS $\mathbf{f} = (\mathbf{f}_1, \mathbf{f}_2, \mathbf{f}_3)$ guarantees that extracted elements $(\sigma_1^*, Z^*, R^*, U^*)$ satisfy the verification equations (7). This means that \mathcal{B} can win the LHSPS security game by outputting (Z^*, R^*, U^*) and the vector

$$\boldsymbol{\sigma}^* = (\sigma_1^*, \sigma_2^{*1-\mathbf{VK}^*[1]}, \sigma_2^{*\mathbf{VK}^*[1]}, \dots, \sigma_2^{*1-\mathbf{VK}^*[L]}, \sigma_2^{*\mathbf{VK}^*[L]},$$

$$\sigma_3^{*1-\mathbf{VK}^*[1]}, \sigma_3^{*\mathbf{VK}^*[1]}, \dots, \sigma_3^{*1-\mathbf{VK}^*[L]}, \sigma_3^{*\mathbf{VK}^*[L]}, 1_{\mathbb{G}}, 1_{\mathbb{G}}, F_1^*, F_2^*, F_3^*) \in \mathbb{G}^{4L+6},$$

which is necessarily outside the row space of \mathbf{M} (5) since \mathbf{F}^* is linearly independent of $(\mathbf{f}_1, \mathbf{f}_2, \mathbf{f}_0)$.

Since the bit $d \in_{\mathcal{R}} \{0, 1\}$ is chosen independently of \mathcal{A} 's view, \mathcal{B} breaks the security of the LHSPS scheme with probability $\epsilon/2$ if the soundness adversary has advantage ϵ . The security result of [47, Theorem 1] implies the upper bound $\mathbf{Adv}_{\mathcal{B}}^{\text{LHSPS}}(\lambda) \leq \frac{1}{2} \cdot \mathbf{Adv}^{\text{DLIN}}(\lambda) + \frac{1}{p}$ for the advantage of any LHSPS forger \mathcal{B} . It follows that $\Pr[\mathcal{A} \text{ wins Game}_1] \leq \mathbf{Adv}^{\text{DLIN}}(\lambda) + 2/p$, which yields the announced upper bound. \square

E Proof of Lemma 1

Proof. We proceed using a sequence of games where several kinds of simulated proofs and Groth-Sahai CRSes $(\mathbf{f}_1, \mathbf{f}_2, \mathbf{F})$ may be used by $\mathcal{S}_2(\psi, \tau_{sim}, \cdot, \cdot)$ and the adversary.

Type A proof: A proof $\pi = (\text{VK}, \mathbf{F}, \mathbf{C}_{\sigma_1, \sigma_2, \sigma_3}, \mathbf{C}_Z, \mathbf{C}_R, \mathbf{C}_U, \mathbf{C}_z, \mathbf{C}_r, \mathbf{C}_u, \pi_{\sigma,1}, \pi_{\sigma,2}, \pi_1, \pi_2, \sigma)$ of Type A involves a CRS $(\mathbf{f}_1, \mathbf{f}_2, \mathbf{F})$ and a committed signature $(\sigma_1, \sigma_2, \sigma_3, Z, R, U)$ which are identical to those used by the simulator $\mathcal{S}_2(\psi, \tau_{sim}, \cdot, \cdot)$ in Game_5 . They are obtained by using $\omega_1, \omega_2 \in \mathbb{Z}_p$ to construct a signature $(\sigma_1, \sigma_2, \sigma_3, Z, R, U)$ on the one-time verification key VK . Let $\mathbf{V} = f^{\mathbf{v}}$ and $\mathbf{W} = h^{\mathbf{w}}$ for vectors $\mathbf{v} = (v_{1,0}, v_{1,1}, \dots, v_{L,0}, v_{L,1}) \in \mathbb{Z}_p^{2L}$, $\mathbf{w} = (w_{1,0}, w_{1,1}, \dots, w_{L,0}, w_{L,1}) \in \mathbb{Z}_p^{2L}$ and $f_1 = g^{y_1}$, $f_2 = g^{y_2}$. In a Type A proof, if we write $\mathbf{F} = (F_1, F_2, F_3) = \mathbf{f}_1^{\mu_1} \cdot \mathbf{f}_2^{\mu_2} \cdot \mathbf{f}_0^x$ for some unique triple $(\mu_1, \mu_2, x) \in \mathbb{Z}_p^3$ (recall that $(\mathbf{f}_1, \mathbf{f}_2, \mathbf{f}_0)$ form a basis of \mathbb{G}^3 in Game_5), the hidden signature $(\sigma_1, \sigma_2, \sigma_3, Z, R, U)$ satisfies

$$\sigma_1 = g^{(\omega_1 + \omega_2) \cdot x} \cdot \sigma_2^{F(\mathbf{v}, \text{VK})} \cdot \sigma_3^{F(\mathbf{w}, \text{VK})}, \quad (12)$$

where $F(\mathbf{v}, \text{VK}) = \sum_{\ell=1}^L v_{\ell, \text{VK}[\ell]}$ and $F(\mathbf{w}, \text{VK}) = \sum_{\ell=1}^L w_{\ell, \text{VK}[\ell]}$, whereas (Z, R, U) is a valid linearly homomorphic signature on the vector (6). Since this vector is in the row space of \mathbf{M} , (Z, R, U) can be obtained as

$$\begin{cases} Z = Z_1^{\mu_1} \cdot Z_2^{\mu_2} \cdot Z_{4L+3}^{\omega_1 \cdot x} \cdot Z_{4L+4}^{\omega_2 \cdot x} \cdot Z_{4L+5}^x \cdot \\ \quad \prod_{i=1}^L (Z_{2+2i-\overline{\text{VK}[i]}}^r \cdot Z_{2+2L+2i-\overline{\text{VK}[i]}}^s) \\ R = R_1^{\mu_1} \cdot R_2^{\mu_2} \cdot R_{4L+3}^{\omega_1 \cdot x} \cdot R_{4L+4}^{\omega_2 \cdot x} \cdot R_{4L+5}^x \cdot \\ \quad \prod_{i=1}^L (R_{2+2i-\overline{\text{VK}[i]}}^r \cdot R_{2+2L+2i-\overline{\text{VK}[i]}}^s) \\ U = U_1^{\mu_1} \cdot U_2^{\mu_2} \cdot U_{4L+3}^{\omega_1 \cdot x} \cdot U_{4L+4}^{\omega_2 \cdot x} \cdot U_{4L+5}^x \cdot \\ \quad \prod_{i=1}^L (U_{2+2i-\overline{\text{VK}[i]}}^r \cdot U_{2+2L+2i-\overline{\text{VK}[i]}}^s) \end{cases} \quad (13)$$

with $r = \log_f(\sigma_2)$ and $s = \log_h(\sigma_3)$. Note that, since $\mathbf{f}_0 = \iota(g) \cdot \mathbf{f}_1^{\nu_1} \cdot \mathbf{f}_2^{\nu_2}$ for some $\nu_1, \nu_2 \in \mathbb{Z}_p$, $\mathbf{F} = (F_1, F_2, F_3)$ can be seen as a BBS encryption of g^x since $g^x = F_1^{-1/y_1} \cdot F_2^{-1/y_2} \cdot F_3$.

We further define **Type A' proof** as a broader class of proofs where only conditions (12) are required. We do not impose any condition on (Z, R, U) besides being valid homomorphic signature (or quasi-adaptive proof) on the vectors (6).

In Game_5 , the simulator always sets $x = 1$ in all simulated proofs. If the adversary outputs a Type A' proof, however, we may have $x \neq 1$. In honestly generated proofs, we always have $x = 0$. We also consider another class of proofs.

Type B proofs: A proof $\pi = (\text{VK}, \mathbf{F}, \mathbf{C}_{\sigma_1, \sigma_2, \sigma_3}, \mathbf{C}_Z, \mathbf{C}_R, \mathbf{C}_U, \mathbf{C}_z, \mathbf{C}_r, \mathbf{C}_u, \pi_{\sigma,1}, \pi_{\sigma,2}, \pi_1, \pi_2, \sigma)$ of Type B is a valid proof that is not of Type A. In these proofs, the commitments $\mathbf{C}_{\sigma_1}, \mathbf{C}_Z, \mathbf{C}_R, \mathbf{C}_U$ contain group elements (σ_1, Z, R, U) such that, if we write $\mathbf{F} = (F_1, F_2, F_3)$ as $\mathbf{F} = \mathbf{f}_1^{\mu_1} \cdot \mathbf{f}_2^{\mu_2} \cdot \mathbf{f}_0^x$ for some unique triple $(\mu_1, \mu_2, x) \in \mathbb{Z}_p^3$, then $(\sigma_1, \sigma_2, \sigma_3)$ are of the form

$$\sigma_1 = g^{(\omega_1 + \omega_2 + \tau) \cdot x} \cdot H(\mathbf{V}, \text{VK})^r \cdot H(\mathbf{W}, \text{VK})^s, \quad \sigma_2 = f^r, \quad \sigma_3 = h^s,$$

for some $r, s \in \mathbb{Z}_p$ and some non-zero $\tau \in \mathbb{Z}_p^*$. Since (Z, R, U) is a valid homomorphic signature on the vector (6), it must be of the form

$$\begin{cases} Z = g^{-\tau \cdot x \cdot \varphi_1} \cdot Z_1^{\mu_1} \cdot Z_2^{\mu_2} \cdot Z_{4L+3}^{\omega_1 \cdot x} \cdot Z_{4L+4}^{\omega_2 \cdot x} \cdot Z_{4L+5}^x \cdot \\ \quad \prod_{i=1}^L (Z_{2+2i-\overline{\text{VK}[i]}}^r \cdot Z_{2+2L+2i-\overline{\text{VK}[i]}}^s) \\ R = g^{-\tau \cdot x \cdot \phi_1} \cdot R_1^{\mu_1} \cdot R_2^{\mu_2} \cdot R_{4L+3}^{\omega_1 \cdot x} \cdot R_{4L+4}^{\omega_2 \cdot x} \cdot R_{4L+5}^x \cdot \\ \quad \prod_{i=1}^L (R_{2+2i-\overline{\text{VK}[i]}}^r \cdot R_{2+2L+2i-\overline{\text{VK}[i]}}^s) \\ U = g^{-\tau \cdot x \cdot \vartheta_1} \cdot U_1^{\mu_1} \cdot U_2^{\mu_2} \cdot U_{4L+3}^{\omega_1 \cdot x} \cdot U_{4L+4}^{\omega_2 \cdot x} \cdot U_{4L+5}^x \cdot \\ \quad \prod_{i=1}^L (U_{2+2i-\overline{\text{VK}[i]}}^r \cdot U_{2+2L+2i-\overline{\text{VK}[i]}}^s) \end{cases}$$

for certain exponents $(\varphi_1, \phi_1, \vartheta_1) \in \mathbb{Z}_p^3$ such that $(G_1, H_1) = (G_z^{\varphi_1} G_r^{\phi_1}, H_z^{\varphi_1} H_u^{\vartheta_1})$. We observe that the vector (6) is outside the row space of \mathbf{M} whenever $\tau \neq 0$.

Similarly to Type A proofs, we define **Type B' proofs** as a generalization of Type B proofs where no condition is imposed on the distribution of (Z, R, U) beyond the fact that it should be a valid linearly homomorphic signature on the vector (6).

It is easy to see that, if event E_5 occurs in Game_5 , the adversary's fake proof π^* is necessarily a Type A' or Type B' proof with $x \neq 0$ either way. We further consider several sub-classes of Type B proofs.

Type B- k proofs ($1 \leq k \leq L$): These are a special kind of Type B proofs produced by the simulator at some steps of the sequence of games. They are generated by choosing $r, s, \mu_1, \mu_2, x \xleftarrow{R} \mathbb{Z}_p$ and setting $\mathbf{F} = (F_1, F_2, F_3) = \mathbf{f}_1^{\mu_1} \cdot \mathbf{f}_2^{\mu_2} \cdot \mathbf{f}_0^x$, as well as

$$\sigma_1 = g^{(\omega_1 + \omega_2) \cdot x} \cdot R_k(\text{VK}_{|k})^x \cdot H(\mathbf{V}, \text{VK})^r \cdot H(\mathbf{W}, \text{VK})^s, \quad \sigma_2 = f^r, \quad \sigma_3 = h^s,$$

where we define $H(\mathbf{V}, \text{VK}) = \prod_{\ell=1}^L V_{\ell, \text{VK}[\ell]}$, $H(\mathbf{W}, \text{VK}) = \prod_{\ell=1}^L W_{\ell, \text{VK}[\ell]}$ and

$$R_k: \{0, 1\}^k \rightarrow \mathbb{G}, \text{VK}_{|k} \mapsto R_k(\text{VK}_{|k})$$

is a random function that depends on the first k bits of its input $\text{VK} \in \{0, 1\}^L$. The (Z, R, U) components are simulated QA-NIZK proofs of linear subspace membership. They are obtained using $\text{sk}_1 = \{(\varphi_i, \phi_i, \vartheta_i)\}_{i=1}^{4L+6}$ to compute a homomorphic signature on the vector (6) by computing

$$\begin{cases} Z = \sigma_1^{-\varphi_1} \cdot \sigma_2^{-\sum_{i=1}^L \varphi_{2i+\text{VK}[i]}} \cdot \sigma_3^{-\sum_{i=1}^L \varphi_{2L+2i+\text{VK}[i]}} \cdot \\ \quad F_1^{-\varphi_{4L+4}} \cdot F_2^{-\varphi_{4L+5}} \cdot F_3^{-\varphi_{4L+6}} \\ R = \sigma_1^{-\phi_1} \cdot \sigma_2^{-\sum_{i=1}^L \phi_{2i+\text{VK}[i]}} \cdot \sigma_3^{-\sum_{i=1}^L \phi_{2L+2i+\text{VK}[i]}} \cdot \\ \quad F_1^{-\phi_{4L+4}} \cdot F_2^{-\phi_{4L+5}} \cdot F_3^{-\phi_{4L+6}} \\ U = \sigma_1^{-\vartheta_1} \cdot \sigma_2^{-\sum_{i=1}^L \vartheta_{2i+\text{VK}[i]}} \cdot \sigma_3^{-\sum_{i=1}^L \vartheta_{2L+2i+\text{VK}[i]}} \cdot \\ \quad F_1^{-\vartheta_{4L+4}} \cdot F_2^{-\vartheta_{4L+5}} \cdot F_3^{-\vartheta_{4L+6}} \end{cases},$$

or, equivalently,

$$\begin{cases} Z = R_k(\text{VK}_{|k})^{-x \cdot \varphi_1} \cdot Z_1^{\mu_1} \cdot Z_2^{\mu_2} \cdot Z_{4L+3}^{\omega_1 \cdot x} \cdot Z_{4L+4}^{\omega_2 \cdot x} \cdot Z_{4L+5}^x \cdot \\ \quad \prod_{i=1}^L (Z_{2+2i-\overline{\text{VK}[i]}}^r \cdot Z_{2+2L+2i-\overline{\text{VK}[i]}}^s) \\ R = R_k(\text{VK}_{|k})^{-x \cdot \phi_1} \cdot R_1^{\mu_1} \cdot R_2^{\mu_2} \cdot R_{4L+3}^{\omega_1 \cdot x} \cdot R_{4L+4}^{\omega_2 \cdot x} \cdot R_{4L+5}^x \cdot \\ \quad \prod_{i=1}^L (R_{2+2i-\overline{\text{VK}[i]}}^r \cdot R_{2+2L+2i-\overline{\text{VK}[i]}}^s) \\ U = R_k(\text{VK}_{|k})^{-x \cdot \vartheta_1} \cdot U_1^{\mu_1} \cdot U_2^{\mu_2} \cdot U_{4L+3}^{\omega_1 \cdot x} \cdot U_{4L+4}^{\omega_2 \cdot x} \cdot U_{4L+5}^x \cdot \\ \quad \prod_{i=1}^L (U_{2+2i-\overline{\text{VK}[i]}}^r \cdot U_{2+2L+2i-\overline{\text{VK}[i]}}^s) \end{cases} .$$

Finally, we also define **Type B'-k proofs** as a generalization of Type B-k proofs where (Z, R, U) is only constrained to be a valid homomorphic signature on the vector (6).

To prove the result, we consider a sequence of games $\text{Game}_{5,0}, \text{Game}_{5,1}, \text{Game}_{5,2,1}, \dots, \text{Game}_{5,2,L}$, where $\text{Game}_{5,0}$ coincides with Game_5 . For each $i \in \{0, \dots, L\}$, we call $E_{5,i}$ the counterpart of event E_5 in $\text{Game}_{5,i}$ (namely, the event that \mathcal{A} outputs a fake proof π^* containing a vector $\mathbf{F}^* = (F_1^*, F_2^*, F_3^*)$ which is not a BBS encryption of $1_{\mathbb{G}}$). We also define Match_i to be the event that, in $\text{Game}_{5,i}$, the adversary's fake proof π^* has the extended type as the simulated proofs it observes. Namely, if $\mathcal{S}_2(\psi, \tau_{sim}, \dots)$ generates a Type A (resp. Type B-k) proof at each query in $\text{Game}_{5,i}$, Match_i denotes the event that \mathcal{A} outputs a Type A' (resp. Type B'-k) fake proof π^* .

Game_{5,0}: This game is exactly Game_5 . Namely, the adversary obtains Type A simulated proofs at each query to $\mathcal{S}_2(\psi, \tau_{sim}, \dots)$. At the end of the game, the challenger \mathcal{B} checks if \mathcal{A} 's fake proof $\pi^* = (\text{VK}^*, \mathbf{F}^*, \mathbf{C}_{\sigma_1}^*, \sigma_2^*, \sigma_3^*, \mathbf{C}_Z^*, \mathbf{C}_R^*, \mathbf{C}_U^*, \mathbf{C}_z^*, \mathbf{C}_r^*, \mathbf{C}_u^*, \pi_{\sigma,1}^*, \pi_{\sigma,2}^*, \pi_1^*, \pi_2^*, \sigma^*)$ is a Type A' proof and we call Match_0 this event. We have $\Pr[E_{5,0}] = \Pr[E_{5,0} \wedge \text{Match}_0] + \Pr[E_{5,0} \wedge \neg \text{Match}_0]$. Lemma 2 shows that, if the DLIN assumption holds, π^* can only be a Type B' proof with negligible probability. Concretely, we prove that $\Pr[E_{5,0} \wedge \neg \text{Match}_0] \leq \text{Adv}_{\mathbb{G}}^{\text{DLIN}}(\lambda) + 1/p$. Instead of trying to bound $\Pr[E_{5,0} \wedge \text{Match}_0]$ right now, we will wait for a later game where it will be easier.

Game_{5,1}: This game is like $\text{Game}_{5,0}$ except that, at each query to the simulator $\mathcal{S}_2(\psi, \tau_{sim}, \dots)$, the signature components (Z, R, U) are obtained as simulated QA-NIZK proofs of linear subspace membership. Namely, instead of computing (Z, R, U) using $\mu_1, \mu_2, \omega_1, \omega_2, r, s \in \mathbb{Z}_p$ as in (13), the challenger uses $\{\varphi_i, \phi_i, \vartheta_i\}_{i=1}^{4L+6}$ to compute (Z, R, U) as a one-time linearly homomorphic signature on the vector σ (6). Clearly, (Z, R, U) has the same distribution as in $\text{Game}_{5,0}$ since $\sigma \in \mathbb{G}^{4L+6}$ remains in the row space of the matrix \mathbf{M} . Hence, \mathcal{A} 's view is the same as in $\text{Game}_{5,0}$ and we have $\Pr[E_{5,1} \wedge \text{Match}_1] = \Pr[E_{5,0} \wedge \text{Match}_0]$, where Match_1 is the equivalent of event Match_0 in $\text{Game}_{5,1}$.

Game_{5,2,k} ($1 \leq k \leq L$): In $\text{Game}_{5,2,k}$, all queries to $\mathcal{S}_2(\psi, \tau_{sim}, \dots)$ are answered by returning Type B-k proofs with $x = 1$ (this choice of x suffices to guarantee that $(\mathbf{f}_1, \mathbf{f}_2, \mathbf{F})$ will be a perfectly NIWI CRS). For each k , we define $\text{Match}_{2,k}$ to be the event that \mathcal{A} outputs a Type B'-k fake proof π^* in $\text{Game}_{5,2,k}$. Lemma 3 shows that $\text{Game}_{5,2,1}$ is computationally indistinguishable from $\text{Game}_{5,1}$ under the DLIN assumption in \mathbb{G} . In particular, \mathcal{A} 's fake proof π^* is not significantly more likely to depart from the output distribution of $\mathcal{S}_2(\psi, \tau_{sim}, \dots)$ than in $\text{Game}_{5,1}$: concretely, we have $|\Pr[E_{5,2,1} \wedge \text{Match}_{2,1}] - \Pr[E_{5,1} \wedge \text{Match}_1]| \leq 2 \cdot \text{Adv}_{\mathbb{G}}^{\text{DLIN}}(\lambda)$. Lemma 4 further shows that, under the DLIN assumption, the probability of \mathcal{A} 's fake proof π^* to be of the same type

as the outputs of $S_2(\psi, \tau_{sim}, \cdot, \cdot)$ is nearly the same in $\text{Game}_{5.2,k}$ and in $\text{Game}_{5.2,(k-1)}$. We have

$$|\Pr[E_{5.2,k} \wedge \text{Match}_{2,k}] - \Pr[E_{5.2,(k-1)} \wedge \text{Match}_{2,(k-1)}]| \leq 2 \cdot \mathbf{Adv}_{\mathbb{G}}^{\text{DLIN}}(\lambda).$$

In $\text{Game}_{5.2,L}$, we obtain $|\Pr[E_{5.2,L} \wedge \text{Match}_{2,L}] - \Pr[E_{5.0} \wedge \text{Match}_0]| \leq 2 \cdot L \cdot \mathbf{Adv}_{\mathbb{G}}^{\text{DLIN}}(\lambda)$ by the triangle inequality. In $\text{Game}_{5.2,L}$, although \mathcal{A} only obtains Type B-L proofs during the game, it remains information theoretically unable to output a Type B'-L proof π^* with a different verification key \mathbf{VK}^* . Indeed, a Type B'-L fake proof $\pi^* = (\mathbf{VK}^*, \mathbf{F}^*, \mathbf{C}_{\sigma_1}^*, \sigma_2^*, \sigma_3^*, \mathbf{C}_Z^*, \mathbf{C}_R^*, \mathbf{C}_U^*, \mathbf{C}_z^*, \mathbf{C}_r^*, \mathbf{C}_u^*, \pi_{\sigma_1,1}^*, \pi_{\sigma_2,2}^*, \pi_1^*, \pi_2^*, \sigma^*)$ would uniquely determine the function evaluation $R_L(\mathbf{VK}^*)$ from the values $(\sigma_1^*, \sigma_2^*, \sigma_3^*)$ and $\mathbf{F}^* = (F_1^*, F_2^*, F_3^*)$. To see this, recall that $E_{5.2,L}$ implies that \mathbf{F}^* is a BBS encryption of $g^{x^*} = F_1^{*-1/y_1} \cdot F_2^{*-1/y_2} \cdot F_3^* \neq 1_{\mathbb{G}}$, which in turn determines $R_L(\mathbf{VK}^*)$ from $(\sigma_1^*, \sigma_2^*, \sigma_3^*)$. The unpredictability of the random function $R_L(\cdot)$ ensures that this event can only occur by pure chance, so that we have $\Pr[E_{5.2,L} \wedge \text{Match}_{2,L}] \leq 1/p$.

When combining the above, we find

$$\Pr[E_{5.0} \wedge \text{Match}_0] \leq (2L + 1) \cdot \mathbf{Adv}_{\mathbb{G}}^{\text{DLIN}}(\lambda) + 2/p,$$

which yields the announced upper bound for $\Pr[E_5]$. \square

The proof of Lemma 1 is completed by the following lemmas.

Lemma 2. *In $\text{Game}_{5.0}$, any PPT adversary \mathcal{A} outputting a Type B' fake proof with noticeable probability ϵ implies an algorithm \mathcal{B} with comparable running time and breaking the DLIN assumption in \mathbb{G} with advantage at least $\epsilon - 1/p$.*

Proof. Let \mathcal{A} be a PPT adversary that outputs a Type B' fake proof π^* with probability ϵ in $\text{Game}_{5.0}$. We construct an algorithm \mathcal{B} that takes as input an SDP instance $(g_z, g_r, h_z, h_u) \in \mathbb{G}^4$ and finds a non-trivial $(Z, R, U) \in \mathbb{G}^3$ such that $e(g_z, Z) \cdot e(g_r, R) = 1_{\mathbb{G}_T}$ and $e(h_z, Z) \cdot e(h_u, U) = 1_{\mathbb{G}_T}$ with probability $\epsilon \cdot (1 - 1/p)$. Since the SDP assumption is implied by the DLIN assumption under a linear time reduction, \mathcal{B} immediately implies a DLIN distinguisher with the same advantage in \mathbb{G} .

We use \mathcal{A} to build a forger \mathcal{B} for the one-time linearly homomorphic signature of Section 2.3. Our LHSPS forger \mathcal{B} receives as input a public key pk_{hspS} for an instance of the LHSPS scheme that allows signing vectors of dimension $n = 4L + 6$ and defines $\text{pk}_1 = \text{pk}_{\text{hspS}}$ at step 1 of the K_1 algorithm. It generates the second LHSPS key pair $(\text{sk}_0, \text{pk}_0)$ on its own (note that \mathcal{B} can do this because the QA-NIZK simulator does it as well) and faithfully runs steps 2 to 5 of the K_1 algorithm. In particular, it chooses $\mathbf{V}, \mathbf{W} \xleftarrow{R} \mathbb{G}^{2L}$, $u_1, u_2 \xleftarrow{R} \mathbb{G}$, $\omega_1, \omega_2 \xleftarrow{R} \mathbb{Z}_p$ itself, just like the extraction trapdoor $(y_1, y_2) \xleftarrow{R} \mathbb{Z}_p^2$ of the Groth-Sahai CRS $\mathbf{f} = (\mathbf{f}_1, \mathbf{f}_2, \mathbf{f}_3)$. At step 6 of K_1 , it generates $\{(z_i, r_i, u_i)\}_{i=1}^t$ using sk_0 and invokes its LHSPS challenger $4L + 5$ times to obtain signatures $\{(Z_j, R_j, U_j)\}_{j=1}^{4L+5}$ on the on the rows of the matrix $\mathbf{M} \in \mathbb{G}^{(4L+5) \times (4L+6)}$ in (5). The adversary \mathcal{A} is run on input of

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

and \mathcal{B} retains the information $\tau_{sim} = (\omega_1, \omega_2, \{\chi_i, \gamma_i, \delta_i\}_{i=1}^n)$ which will make it possible to emulate $S_2(\psi, \tau_{sim}, \cdot, \cdot)$ exactly as in Game_5 . In particular, since \mathcal{B} knows $\omega_1, \omega_2 \in \mathbb{Z}_p$, it can simulate proofs without knowing $\text{sk}_1 = \{(\varphi_i, \phi_i, \vartheta_i)\}_{i=1}^{4L+6}$. When \mathcal{A} terminates, it outputs a triple $(\mathbf{v}^*, \pi^*, \text{lbl}^*)$ and a verifying proof $\pi^* = (\text{VK}^*, \mathbf{F}^*, \mathbf{C}_{\sigma_1}^*, \sigma_2^*, \sigma_3^*, \mathbf{C}_Z^*, \mathbf{C}_R^*, \mathbf{C}_U^*, \mathbf{C}_z^*, \mathbf{C}_r^*, \mathbf{C}_u^*, \pi_{\sigma_1}^*, \pi_{\sigma_2}^*, \pi_1^*, \pi_2^*, \sigma^*)$. At this point, \mathcal{B} uses the extraction trapdoor $(y_1, y_2) \in \mathbb{Z}_p^2$ of the perfectly binding CRS $(\mathbf{f}_1, \mathbf{f}_2, \mathbf{f}_3)$ to extract $(\sigma_1^*, Z^*, R^*, U^*)$ from $\mathbf{C}_{\sigma_1}^*, \mathbf{C}_Z^*, \mathbf{C}_R^*, \mathbf{C}_U^*$. The perfect soundness of $\pi_{\sigma_1}^*, \pi_{\sigma_2}^*$ ensures that (Z^*, R^*, U^*) is a valid homomorphic signature on the vector

$$\begin{aligned} \sigma^* = & (\sigma_1^*, \sigma_2^{*1-\text{VK}^*[1]}, \sigma_2^{*\text{VK}^*[1]}, \dots, \sigma_2^{*1-\text{VK}^*[L]}, \sigma_2^{*\text{VK}^*[L]}, \sigma_3^{*1-\text{VK}^*[1]}, \sigma_3^{*\text{VK}^*[1]}, \\ & \dots, \sigma_3^{*1-\text{VK}^*[L]}, \sigma_3^{*\text{VK}^*[L]}, 1_{\mathbb{G}}, 1_{\mathbb{G}}, F_1^*, F_2^*, F_3^*) . \end{aligned}$$

Moreover, the latter evades the row space of \mathbf{M} since π^* is a Type B' proof. Therefore \mathcal{B} can win against its LHSPS challenger by outputting (Z^*, R^*, U^*) and the above vector as a valid forgery. The result of [47, Theorem 1] implies that \mathcal{B} can in turn be used to solve the SDP problem —and *a fortiori* break the DLIN assumption in \mathbb{G} — with probability $\epsilon \cdot (1 - 1/p) \geq \epsilon - 1/p$. \square

In the multi-CRS setting, the proof of Lemma 2 readily extends and relies on the tight security of the LHSPS system of Section 2.3 in the multi-user setting, which was proved in [49, Appendix G]. Indeed, each CRS $\psi^{(\kappa)} = (\mathbf{CRS}_1^{(\kappa)}, \mathbf{CRS}_2^{(\kappa)})$ contains an independent LHSPS public key $\text{pk}_1^{(\kappa)}$ and the adversary is simply turned into a successful forger for one of these.

Lemma 3. *If the DLIN assumption holds in \mathbb{G} , \mathcal{A} 's probability to output a Type B'-1 proof in $\text{Game}_{5,2.1}$ and its probability of outputting a Type A' proof in Game_1 are about the same. Specifically, there is a DLIN distinguisher \mathcal{B} such that $|\Pr[E_{5,2.1} \wedge \text{Match}_{2.1}] - \Pr[E_{5,1} \wedge \text{Match}_1]| \leq 2 \cdot \text{Adv}_{\mathbb{G}}^{\text{DLIN}}(\lambda)$ and running in about the same time as \mathcal{A} .*

Proof. For the sake of contradiction, let us assume that events $E_{5,2.1} \wedge \text{Match}_{2.1}$ and $E_1 \wedge \text{Match}_1$ occur with substantially different probabilities in $\text{Game}_{5,2.1}$ and $\text{Game}_{5,1}$, respectively. We build a DLIN distinguisher \mathcal{B} which takes as input (f, g, h, f^a, h^b, T) and aims at deciding if $T = g^{a+b}$ or $T \in_R \mathbb{G}$. Analogously to [53] and [22, Lemma 6], \mathcal{B} leverages the random self-reducibility of DLIN so as to build q tuples $(K_j = f^{a_j}, L_j = h^{b_j}, T_j)$ such that, for each $j \in \{1, \dots, q\}$, we have

$$T_j = \begin{cases} g^{a_j+b_j} & \text{if } T = g^{a+b} \\ g^{a_j+b_j+\tau_0} & \text{if } T \in_R \mathbb{G} \end{cases}$$

for some $\tau_0 \in_R \mathbb{Z}_p$. This is done by picking $\rho_0 \xleftarrow{R} \mathbb{Z}_p$ and $\rho_{a_j}, \rho_{b_j} \xleftarrow{R} \mathbb{Z}_p$, for $j \in \{1, \dots, q\}$, and defining

$$(K_j, L_j, T_j) = ((f^a)^{\rho_0} \cdot f^{\rho_{a_j}}, (h^b)^{\rho_0} \cdot h^{\rho_{b_j}}, T^{\rho_0} \cdot g^{\rho_{a_j} + \rho_{b_j}}), \quad \forall j \in \{1, \dots, q\} .$$

Also, \mathcal{B} generates $(u_1, u_2, \Omega_1, \Omega_2) \in \mathbb{G}^4$ by drawing $\alpha_{u,1}, \alpha_{u,2} \xleftarrow{R} \mathbb{Z}_p$ and setting

$$u_1 = f^{\alpha_{u,1}}, \quad u_2 = h^{\alpha_{u,2}}, \quad \Omega_1 = (f^a)^{\alpha_{u,1}}, \quad \Omega_2 = (h^b)^{\alpha_{u,2}} .$$

Before generating the rest of \mathbf{CRS}_1 and \mathbf{CRS}_2 , \mathcal{B} flips a fair coin $b^\dagger \xleftarrow{R} \{0, 1\}$ as a guess for the first bit of the one-time verification key $\text{VK}^* = \text{VK}[1]^* \dots \text{VK}[L]^* \in \{0, 1\}^L$ that will

appear in the adversary's fake proof π^* . To build \mathbf{CRS}_1 and \mathbf{CRS}_2 , \mathcal{B} chooses vectors $\boldsymbol{\alpha} = (\alpha_{1,0}, \alpha_{1,1}, \dots, \alpha_{L,0}, \alpha_{L,1}) \xleftarrow{R} \mathbb{Z}_p^{2L}$, $\boldsymbol{\beta} = (\beta_{1,0}, \beta_{1,1}, \dots, \beta_{L,0}, \beta_{L,1}) \xleftarrow{R} \mathbb{Z}_p^{2L}$ and $\zeta, y_1, y_2 \xleftarrow{R} \mathbb{Z}_p$. It sets $f_1 = g^{y_1}$ and $f_2 = g^{y_2}$. It also defines the vectors $\mathbf{V} = (V_{1,0}, V_{1,1}, \dots, V_{L,0}, V_{L,1}) \in \mathbb{G}^{2L}$, $\mathbf{W} = (W_{1,0}, W_{1,1}, \dots, W_{L,0}, W_{L,1}) \in \mathbb{G}^{2L}$ by setting

$$\begin{aligned} (V_{\ell,0}, V_{\ell,1}) &= (f^{\alpha_{\ell,0}}, f^{\alpha_{\ell,1}}), \quad (W_{\ell,0}, W_{\ell,1}) = (h^{\beta_{\ell,0}}, h^{\beta_{\ell,1}}), \quad \text{if } \ell \neq 1 \\ (V_{1,1-b^\dagger}, V_{1,b^\dagger}) &= (f^{\alpha_{1,1-b^\dagger}} \cdot g^\zeta, f^{\alpha_{1,b^\dagger}}), \quad (W_{1,1-b^\dagger}, W_{1,b^\dagger}) = (h^{\beta_{1,1-b^\dagger}} \cdot g^\zeta, h^{\beta_{1,b^\dagger}}). \end{aligned}$$

The remaining components of \mathbf{CRS}_1 and \mathbf{CRS}_2 , including $(\text{sk}_1, \text{pk}_1)$ and $\{(Z_j, R_j, U_j)\}_{j=1}^{4L+5}$, are generated as in the real setup algorithm \mathbf{K}_1 . The adversary \mathcal{A} is run on input of

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

and the challenger \mathcal{B} keeps $(y_1, y_2, \{\chi_i, \gamma_i, \delta_i\}_{i=1}^n, \{\varphi_i, \phi_i, \vartheta_i\}_{i=1}^{4L+6})$ to itself. Note that the simulation trapdoor τ_{sim} is not entirely known to \mathcal{B} since $(\omega_1, \omega_2) = (a, b) \in \mathbb{Z}_p^2$, which are part of the original DLIN instance, are not available. Fortunately, \mathcal{B} will be able to simulate proofs using $\{\varphi_i, \phi_i, \vartheta_i\}_{i=1}^{4L+6}$ and its challenge value T which is either g^{a+b} or a random element of \mathbb{G} .

During the game, the adversary's queries to the simulator $\mathbf{S}_2(\psi, \tau_{sim}, \dots)$ are handled as follows. If $\mathbf{VK}^j = \mathbf{VK}[1]^j \dots \mathbf{VK}[L]^j \in \{0, 1\}^L$ denotes the verification key involved in the j -th query, \mathcal{B} 's answer depends on the first bit $\mathbf{VK}[1]^j$ of \mathbf{VK}^j . Specifically, \mathcal{B} considers two cases.

- If $\mathbf{VK}[1]^j = b^\dagger$, \mathcal{B} randomly chooses $r, s, \mu_1, \mu_2 \xleftarrow{R} \mathbb{Z}_p$ and defines the vector $\mathbf{F} = (F_1, F_2, F_3)$ as $\mathbf{F} = \mathbf{f}_0 \cdot \mathbf{f}_1^{\mu_1} \cdot \mathbf{f}_2^{\mu_2}$,

$$\sigma_1 = T \cdot H(\mathbf{V}, \mathbf{VK}^j)^r \cdot H(\mathbf{W}, \mathbf{VK}^j)^s, \quad \sigma_2 = f^r, \quad \sigma_3 = h^s,$$

where $H(\mathbf{V}, \mathbf{VK}^j) = \prod_{\ell=1}^L V_{\ell, \mathbf{VK}[\ell]^j}$ and $H(\mathbf{W}, \mathbf{VK}^j) = \prod_{\ell=1}^L W_{\ell, \mathbf{VK}[\ell]^j}$. The (Z, R, U) components of the proof are computed by generating a homomorphic signature on the vector

$$\begin{aligned} (\sigma_1, \sigma_2^{1-\mathbf{VK}[1]^j}, \sigma_2^{\mathbf{VK}[1]^j}, \dots, \sigma_2^{1-\mathbf{VK}[L]^j}, \sigma_2^{\mathbf{VK}[L]^j}, \sigma_3^{1-\mathbf{VK}[1]^j}, \sigma_3^{\mathbf{VK}[1]^j}, \\ \dots, \sigma_3^{1-\mathbf{VK}[L]^j}, \sigma_3^{\mathbf{VK}[L]^j}, 1_{\mathbb{G}}, 1_{\mathbb{G}}, F_1, F_2, F_3), \end{aligned}$$

which is done by computing

$$\begin{cases} Z = \sigma_1^{-\varphi_1} \cdot \sigma_2^{-\sum_{i=1}^L \varphi_{2i+\mathbf{VK}[i]^j}} \cdot \sigma_3^{-\sum_{i=1}^L \varphi_{2L+2i+\mathbf{VK}[i]^j}} \cdot \\ \quad F_1^{-\varphi_{4L+4}} \cdot F_2^{-\varphi_{4L+5}} \cdot F_3^{-\varphi_{4L+6}} \\ R = \sigma_1^{-\phi_1} \cdot \sigma_2^{-\sum_{i=1}^L \phi_{2i+\mathbf{VK}[i]^j}} \cdot \sigma_3^{-\sum_{i=1}^L \phi_{2L+2i+\mathbf{VK}[i]^j}} \cdot \\ \quad F_1^{-\phi_{4L+4}} \cdot F_2^{-\phi_{4L+5}} \cdot F_3^{-\phi_{4L+6}} \\ U = \sigma_1^{-\vartheta_1} \cdot \sigma_2^{-\sum_{i=1}^L \vartheta_{2i+\mathbf{VK}[i]^j}} \cdot \sigma_3^{-\sum_{i=1}^L \vartheta_{2L+2i+\mathbf{VK}[i]^j}} \cdot \\ \quad F_1^{-\vartheta_{4L+4}} \cdot F_2^{-\vartheta_{4L+5}} \cdot F_3^{-\vartheta_{4L+6}} \end{cases} \quad (14)$$

Note that, if $T = g^{a+b+\tau}$ for some $\tau \in_R \mathbb{Z}_p$, the obtained triple (Z, R, U) can be written

$$\begin{cases} Z = g^{-\tau \cdot \varphi_1} \cdot Z_1^{\mu_1} \cdot Z_2^{\mu_2} \cdot Z_{4L+3}^a \cdot Z_{4L+4}^b \cdot Z_{4L+5} \cdot \\ \quad \prod_{i=1}^L (Z_{2+2i-\overline{\text{VK}[i]^j}}^r \cdot Z_{2+2L+2i-\overline{\text{VK}[i]^j}}^s) \\ R = g^{-\tau \cdot \phi_1} \cdot R_1^{\mu_1} \cdot R_2^{\mu_2} \cdot R_{4L+3}^a \cdot R_{4L+4}^b \cdot R_{4L+5} \cdot \\ \quad \prod_{i=1}^L (R_{2+2i-\overline{\text{VK}[i]^j}}^r \cdot R_{2+2L+2i-\overline{\text{VK}[i]^j}}^s) \\ U = g^{-\tau \cdot \vartheta_1} \cdot U_1^{\mu_1} \cdot U_2^{\mu_2} \cdot U_{4L+3}^a \cdot U_{4L+4}^b \cdot U_{4L+5} \cdot \\ \quad \prod_{i=1}^L (U_{2+2i-\overline{\text{VK}[i]^j}}^r \cdot U_{2+2L+2i-\overline{\text{VK}[i]^j}}^s) \end{cases} .$$

We observe that $(F_1, F_2, F_3, \sigma_1, \sigma_2, \sigma_3, Z, R, U)$ matches the distribution of π in $\text{Game}_{5.2.1}$ if $\tau \neq 0$ and $\text{Game}_{5.1}$ if $\tau = 0$. Indeed, in the former case, we implicitly define the constant function $R_0(\varepsilon) = g^\tau$ and define the function R_1 so that $R_1(b^\dagger || M') = R_0(\varepsilon)$ for any string $M' \in \{0, 1\}^{L-1}$.

– If $\text{VK}[1]^j = 1 - b^\dagger$, \mathcal{B} implicitly defines

$$R_1(\text{VK}_1^j) = R_1(1 - b^\dagger) = \begin{cases} R_0(\varepsilon) \cdot g^{\zeta \cdot \tau_0} & \text{if } T \in_R \mathbb{G} \\ 1 & \text{if } T = g^{a+b} \end{cases} .$$

Namely, for randomly chosen $\mu_1, \mu_2 \xleftarrow{R} \mathbb{Z}_p$, \mathcal{B} defines $\mathbf{F} = \mathbf{f}_0 \cdot \mathbf{f}_1^{\mu_1} \cdot \mathbf{f}_2^{\mu_2}$ and uses the j -th tuple (K_j, L_j, T_j) to set

$$\begin{aligned} \sigma_1 &= T \cdot K_j^{\sum_{\ell=1}^L \alpha_{\ell, \text{VK}[\ell]^j}} \cdot L_j^{\sum_{\ell=1}^L \beta_{\ell, \text{VK}[\ell]^j}} \cdot T_j^\zeta, \\ \sigma_2 &= K_j = f^{a_j}, \quad \sigma_3 = L_j = h^{b_j}. \end{aligned}$$

If $T = g^{a+b}$ (and thus $T_j = g^{a_j+b_j}$), this implicitly defines $\sigma_1 = g^{(a+b)} \cdot H(\mathbf{V}, \text{VK}^j)^{a_j} \cdot H(\mathbf{W}, \text{VK}^j)^{b_j}$, so that $(\sigma_1, \sigma_2, \sigma_3)$ has the same distribution as in $\text{Game}_{5.1}$. If $T = g^{a+b+\tau}$ (so that $T_j = g^{a_j+b_j+\tau_0}$), we can write

$$\sigma_1 = g^{(a+b)} \cdot R_1(\text{VK}_1^j) \cdot H(\mathbf{V}, \text{VK}^j)^{a_j} \cdot H(\mathbf{W}, \text{VK}^j)^{b_j},$$

since $R_1(\text{VK}_1^j) = R_0(\varepsilon) \cdot g^{\zeta \cdot \tau_0}$, which is distributed as in $\text{Game}_{5.2.1}$. In either case, (Z, R, U) are computed using $\text{sk}_1 = \{(\varphi_i, \phi_i, \vartheta_i)\}_{i=1}^{4L+6}$ as in the previous case (i.e., as per (14)).

At the end of the game, the adversary \mathcal{A} halts and outputs a triple $(\mathbf{v}^*, \pi^*, \text{bl}^*)$ and a valid fake proof $\pi^* = (\text{VK}^*, \mathbf{F}^*, \mathbf{C}_{\sigma_1}^*, \sigma_2^*, \sigma_3^*, \mathbf{C}_Z^*, \mathbf{C}_R^*, \mathbf{C}_U^*, \mathbf{C}_Z^*, \mathbf{C}_R^*, \mathbf{C}_U^*, \pi_{\sigma_1}^*, \pi_{\sigma_2}^*, \pi_1^*, \pi_2^*, \sigma^*)$. At this point, our distinguisher \mathcal{B} aborts and outputs a random bit in the event that $\text{VK}[1]^* \neq b^\dagger$. Otherwise, it must determine if the fake proof π^* has the same type as the outputs of the simulator. To this end, \mathcal{B} uses the extraction trapdoor $(y_1, y_2) = (\log_g(f_1), \log_g(f_2))$ to extract $(\sigma_1^*, Z^*, R^*, U^*)$ from $\mathbf{C}_{\sigma_1}^*, \mathbf{C}_Z^*, \mathbf{C}_R^*, \mathbf{C}_U^*$. Since $\text{VK}[1]^* = b^\dagger$, \mathcal{B} can parse \mathbf{F}^* as (F_1^*, F_2^*, F_3^*) and compute $F(\mathbf{v}, \text{VK}^*) = \sum_{\ell=1}^L \alpha_{\ell, \text{VK}[\ell]^*}$ and $F(\mathbf{w}, \text{VK}^*) = \sum_{\ell=1}^L \beta_{\ell, \text{VK}[\ell]^*}$, which in turn yield

$$\eta^* = \sigma_1^* \cdot \sigma_2^{*-F(\mathbf{v}, \text{VK}^*)} \cdot \sigma_3^{*-F(\mathbf{w}, \text{VK}^*)}, \quad F_0^* = F_1^{*-1/y_1} \cdot F_2^{*-1/y_2} \cdot F_3^* .$$

If the equality

$$e(\eta^*, g) = (T, F_0^*) \quad (15)$$

holds, \mathcal{B} knows that $\sigma_1^* = T^x \cdot H(\mathbf{V}, \mathbf{VK}^*)^{r^*} \cdot H(\mathbf{W}, \mathbf{VK}^*)^{s^*}$, where $r^* = \log_f(\sigma_2^*)$, $s^* = \log_h(\sigma_3^*)$ and $x = \log_g(F_0^*)$. It thus considers π^* as a fake proof of the same extended type as the outputs of $\mathbf{S}_2(\psi, \tau_{sim}, \dots)$ and returns 1. Recall that $\mathbf{R}_0(\varepsilon) = T/g^{a+b}$, so that $(\sigma_1^*, \sigma_2^*, \sigma_3^*, F_1^*, F_2^*, F_3^*)$ corresponds to a Type A' proof (resp. Type B'-1) if $T = g^{a+b}$ (resp. $T = g^{a+b+\tau}$ with $\tau \neq 0$). Otherwise, \mathcal{B} concludes that π^* has a different distribution than proofs produced by the simulator and outputs 0. If the difference between the adversary's probability to output the same kind of proofs as $\mathbf{S}_2(\psi, \tau_{sim}, \dots)$ in $\text{Game}_{5.2.1}$ and $\text{Game}_{5.1}$ is ϵ , then \mathcal{B} 's distinguishing advantage is at least $\epsilon/2$ since $b^\dagger \in \{0, 1\}$ was chosen independently of \mathcal{A} 's view. \square

In the multi-CRS setting, the proof of Lemma 3 easily goes through by having the reduction \mathcal{B} generate a set μ CRSes $\{\psi^{(\kappa)} = (\mathbf{CRS}_1^{(\kappa)}, \mathbf{CRS}_2^{(\kappa)})\}_{\kappa=1}^\mu$, where each $\mathbf{CRS}_1^{(\kappa)}$ contains independent values of $u_1^{(\kappa)}, u_2^{(\kappa)}, \Omega_1^{(\kappa)}, \Omega_2^{(\kappa)}$, for unknown $(a^{(\kappa)}, b^{(\kappa)}) = (\log_{u_1^{(\kappa)}}(\Omega_1^{(\kappa)}), \log_{u_2^{(\kappa)}}(\Omega_2^{(\kappa)}))$, which are generated using the random self-reducibility of a given DLIN instance. At the beginning of the game, \mathcal{B} generates $q\mu$ tuples $(K_j^{(\kappa)}, L_j^{(\kappa)}, T_j^{(\kappa)}) = (f^{a_j^{(\kappa)}}, h^{b_j^{(\kappa)}}, g^{a_j^{(\kappa)}+b_j^{(\kappa)}+\tau_0^{(\kappa)}})$, for each $j \in \{1, \dots, q\}$ and $\kappa \in \{1, \dots, \mu\}$, where $\{\tau_0^{(\kappa)}\}_{\kappa=1}^\mu$ are either all zeroes or uniformly random in \mathbb{Z}_p . These instances will be used to implicitly define an independent random function $\mathbf{R}_1^{(\kappa)}$ for each CRS $\psi^{(\kappa)}$.

Lemma 4. *If the DLIN assumption holds in \mathbb{G} , \mathcal{A} 's probability to output the same (extended) type of proof as the simulator is about the same in $\text{Game}_{5.2.k}$ and $\text{Game}_{5.2.(k-1)}$ for any $k \in \{2, \dots, L\}$. Namely, there exists a DLIN distinguisher \mathcal{B} such that $|\Pr[E_{5.2.k} \wedge \text{Match}_{2.k}] - \Pr[E_{5.2.(k-1)} \wedge \text{Match}_{2.(k-1)}]| \leq 2 \cdot \text{Adv}_{\mathbb{G}}^{\text{DLIN}}(\lambda)$ and running in about the same time as \mathcal{A} .*

Proof. Let us assume that there exist an adversary \mathcal{A} and an index $k \in \{2, \dots, L\}$ such that events $E_{5.2.k} \wedge \text{Match}_{2.k}$ and $E_{5.2.(k-1)} \wedge \text{Match}_{2.(k-1)}$ have significantly different occurrence probabilities in $\text{Game}_{5.2.k}$ and $\text{Game}_{5.2.(k-1)}$, respectively. Out of \mathcal{A} , we construct a DLIN distinguisher \mathcal{B} as follows. Algorithm \mathcal{B} inputs (f, g, h, f^a, h^b, T) and has to decide if $T = g^{a+b}$ or $T \in_R \mathbb{G}$. As in [22, Lemma 6], \mathcal{B} can take advantage of the random self-reducibility of DLIN in order to build q tuples

$$(K_j = f^{a_j}, L_j = h^{b_j}, T_j)$$

such that, for each $j \in \{1, \dots, q\}$, we have

$$T_j = \begin{cases} g^{a_j+b_j} & \text{if } T = g^{a+b} \\ g^{a_j+b_j+\tau_j} & \text{if } T \in_R \mathbb{G} \end{cases}$$

for $\tau_1, \dots, \tau_q \in_R \mathbb{Z}_p$. This is achieved by picking $\rho_j, \rho_{a_j}, \rho_{b_j} \xleftarrow{R} \mathbb{Z}_p$ and setting

$$(K_j, L_j, T_j) = ((f^a)^{\rho_j} \cdot f^{\rho_{a_j}}, (h^b)^{\rho_j} \cdot h^{\rho_{b_j}}, T^{\rho_j} \cdot g^{\rho_{a_j}+\rho_{b_j}}), \quad \forall j \in \{1, \dots, q\}.$$

Before generating \mathbf{CRS}_1 and \mathbf{CRS}_2 , \mathcal{B} flips a fair binary coin $b^\dagger \xleftarrow{R} \{0, 1\}$ as a guess for the k -th bit $\mathbf{VK}[k]^*$ of the one-time verification key $\mathbf{VK}^* = \mathbf{VK}[1]^* \dots \mathbf{VK}[L]^* \in \{0, 1\}^L$ contained in \mathcal{A} 's fake proof π^* . To prepare \mathbf{CRS}_1 and \mathbf{CRS}_2 , \mathcal{B} picks $u_1, u_2 \xleftarrow{R} \mathbb{G}$, $\omega_1, \omega_2 \xleftarrow{R} \mathbb{Z}_p$, $y_1, y_2 \xleftarrow{R} \mathbb{Z}_p$,

$\zeta \xleftarrow{R} \mathbb{Z}_p$ and vectors $\boldsymbol{\alpha} = (\alpha_{1,0}, \alpha_{1,1}, \dots, \alpha_{L,0}, \alpha_{L,1}) \xleftarrow{R} \mathbb{Z}_p^{2L}$, $\boldsymbol{\beta} = (\beta_{1,0}, \beta_{1,1}, \dots, \beta_{L,0}, \beta_{L,1}) \xleftarrow{R} \mathbb{Z}_p^{2L}$. It sets $\Omega_1 = u_1^{\omega_1}$, $\Omega_2 = u_2^{\omega_2}$, $f_1 = g^{y_1}$, $f_2 = g^{y_2}$ and defines the vectors $\mathbf{V} = (V_{1,0}, V_{1,1}, \dots, V_{L,0}, V_{L,1})$, $\mathbf{W} = (W_{1,0}, W_{1,1}, \dots, W_{L,0}, W_{L,1})$ as

$$(V_{\ell,0}, V_{\ell,1}) = (f^{\alpha_{\ell,0}}, f^{\alpha_{\ell,1}}), \quad (W_{\ell,0}, W_{\ell,1}) = (h^{\beta_{\ell,0}}, h^{\beta_{\ell,1}}), \quad \text{if } \ell \neq k,$$

$$(V_{k,1-b^\dagger}, V_{k,b^\dagger}) = (f^{\alpha_{k,1-b^\dagger}} \cdot g^\zeta, f^{\alpha_{k,b^\dagger}}), \quad (W_{k,1-b^\dagger}, W_{k,b^\dagger}) = (h^{\beta_{k,1-b^\dagger}} \cdot g^\zeta, h^{\beta_{k,b^\dagger}}).$$

Other components of $(\mathbf{CRS}_1, \mathbf{CRS}_2)$, including $(\text{sk}_1, \text{pk}_1)$ and $\{(Z_j, R_j, U_j)\}_{j=1}^{4L+5}$, are faithfully generated as in the real algorithm \mathbf{K}_1 .

The adversary \mathcal{A} is run on input of

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

$$\mathbf{CRS}_2 = \left(\mathbf{f}, \mathbf{f}_0, \text{pk}_0, \text{pk}_1, \Omega_1, \Omega_2, \mathbf{V}, \mathbf{W} \right),$$

and the challenger \mathcal{B} retains $(y_1, y_2, \omega_1, \omega_2, \{\chi_i, \gamma_i, \delta_i\}_{i=1}^n, \{\varphi_i, \phi_i, \vartheta_i\}_{i=1}^{4L+6})$ for later use.

At the outset of the game, \mathcal{B} picks a random function $\mathbf{R}_{k-1} : \{0, 1\}^{k-1} \rightarrow \mathbb{G}$ which will be used to construct another random function $\mathbf{R}_k : \{0, 1\}^k \rightarrow \mathbb{G}$ such that, for any string $M \in \{0, 1\}^{k-1}$, we have $\mathbf{R}_k(M||b^\dagger) = \mathbf{R}_{k-1}(M)$ while $\mathbf{R}_k(M||1-b^\dagger)$ takes an independent value.

Then, \mathcal{B} starts answering simulation queries. Let $\mathbf{VK}^j = \mathbf{VK}[1]^j \dots \mathbf{VK}[L]^j$ be the one-time verification key involved in the j -th query to $\mathbf{S}_2(\psi, \tau_{sim}, \dots)$. The response of \mathcal{B} will depend on the k -th bit $\mathbf{VK}[k]^j$ of \mathbf{VK}^j . Namely, \mathcal{B} considers the following three cases.

- If $\mathbf{VK}[k]^j = b^\dagger$, \mathcal{B} exploits the property that $\mathbf{R}_k(\mathbf{VK}_{|k}^j) = \mathbf{R}_{k-1}(\mathbf{VK}_{|k-1}^j)$. It picks $r, s \xleftarrow{R} \mathbb{Z}_p$ and defines

$$\sigma_1 = g^{\omega_1 + \omega_2} \cdot \mathbf{R}_{k-1}(\mathbf{VK}_{|k-1}^j) \cdot H(\mathbf{V}, \mathbf{VK}^j)^r \cdot H(\mathbf{W}, \mathbf{VK}^j)^s,$$

$$\sigma_2 = f^r, \quad \sigma_3 = h^s,$$

where $H(\mathbf{V}, \mathbf{VK}^j) = \prod_{\ell=1}^L V_{\ell, \mathbf{VK}[\ell]^j}$ and $H(\mathbf{W}, \mathbf{VK}^j) = \prod_{\ell=1}^L W_{\ell, \mathbf{VK}[\ell]^j}$. It also computes the vector \mathbf{F} as $\mathbf{F} = (F_1, F_2, F_3) = \mathbf{f}_0 \cdot \mathbf{f}_1^{\mu_1} \cdot \mathbf{f}_2^{\mu_2}$ for random $\mu_1, \mu_2 \xleftarrow{R} \mathbb{Z}_p$. The (Z, R, U) components of the proof are obtained by using sk_1 to generate a homomorphic structure-preserving signature on the vector

$$(\sigma_1, \sigma_2^{1-\mathbf{VK}[1]^j}, \sigma_2^{\mathbf{VK}[1]^j}, \dots, \sigma_2^{1-\mathbf{VK}[L]^j}, \sigma_2^{\mathbf{VK}[L]^j}, \sigma_3^{1-\mathbf{VK}[1]^j}, \sigma_3^{\mathbf{VK}[1]^j}, \\ \dots, \sigma_3^{1-\mathbf{VK}[L]^j}, \sigma_3^{\mathbf{VK}[L]^j}, 1_{\mathbb{G}}, 1_{\mathbb{G}}, F_1, F_2, F_3)$$

by computing

$$\begin{cases} Z = \sigma_1^{-\varphi_1} \cdot \sigma_2^{-\sum_{i=1}^L \varphi_{2i+\mathbf{VK}[i]^j}} \cdot \sigma_3^{-\sum_{i=1}^L \varphi_{2L+2i+\mathbf{VK}[i]^j}} \cdot \\ \quad F_1^{-\varphi_{4L+4}} \cdot F_2^{-\varphi_{4L+5}} \cdot F_3^{-\varphi_{4L+6}} \\ R = \sigma_1^{-\phi_1} \cdot \sigma_2^{-\sum_{i=1}^L \phi_{2i+\mathbf{VK}[i]^j}} \cdot \sigma_3^{-\sum_{i=1}^L \phi_{2L+2i+\mathbf{VK}[i]^j}} \cdot \\ \quad F_1^{-\phi_{4L+4}} \cdot F_2^{-\phi_{4L+5}} \cdot F_3^{-\phi_{4L+6}} \\ U = \sigma_1^{-\vartheta_1} \cdot \sigma_2^{-\sum_{i=1}^L \vartheta_{2i+\mathbf{VK}[i]^j}} \cdot \sigma_3^{-\sum_{i=1}^L \vartheta_{2L+2i+\mathbf{VK}[i]^j}} \cdot \\ \quad F_1^{-\vartheta_{4L+4}} \cdot F_2^{-\vartheta_{4L+5}} \cdot F_3^{-\vartheta_{4L+6}} \end{cases} \quad (16)$$

Note that (Z, R, U) can be written

$$\begin{cases} Z = R_{k-1}(\text{VK}_{|k-1}^j)^{-\varphi_1} \cdot Z_1^{\mu_1} \cdot Z_2^{\mu_2} \cdot Z_{4L+3}^{\omega_1} \cdot Z_{4L+4}^{\omega_2} \cdot Z_{4L+5} \cdot \\ \quad \prod_{i=1}^L (Z_{2+2i-\overline{\text{VK}[i]^j}}^r \cdot Z_{2+2L+2i-\overline{\text{VK}[i]^j}}^s) \\ R = R_{k-1}(\text{VK}_{|k-1}^j)^{-\phi_1} \cdot R_1^{\mu_1} \cdot R_2^{\mu_2} \cdot R_{4L+3}^{\omega_1} \cdot R_{4L+4}^{\omega_2} \cdot R_{4L+5} \cdot \\ \quad \prod_{i=1}^L (R_{2+2i-\overline{\text{VK}[i]^j}}^r \cdot R_{2+2L+2i-\overline{\text{VK}[i]^j}}^s) \\ U = R_{k-1}(\text{VK}_{|k-1}^j)^{-\vartheta_1} \cdot U_1^{\mu_1} \cdot U_2^{\mu_2} \cdot U_{4L+3}^{\omega_1} \cdot U_{4L+4}^{\omega_2} \cdot U_{4L+5} \cdot \\ \quad \prod_{i=1}^L (U_{2+2i-\overline{\text{VK}[i]^j}}^r \cdot U_{2+2L+2i-\overline{\text{VK}[i]^j}}^s) \end{cases} .$$

We remark that $(F_1, F_2, F_3, \sigma_1, \sigma_2, \sigma_3, Z, R, U)$ have the appropriate distribution in both $\text{Game}_{5.2.(k-1)}$ and $\text{Game}_{5.2.k}$ since $R_{k-1}(\text{VK}_{|k-1}^j) = R_k(\text{VK}_{|k}^j)$.

– If $\text{VK}[k]^j = 1 - b^\dagger$ and $R_k(\text{VK}_{|k}^j)$ has not been defined yet, \mathcal{B} will implicitly define

$$R_k(\text{VK}_{|k}^j) = R_k(\text{VK}_{|k-1}^j || 1 - b^\dagger) = \begin{cases} R_{k-1}(\text{VK}_{|k-1}^j) \cdot g^{\zeta \cdot \tau_j} & \text{if } T \in_R \mathbb{G} \\ R_{k-1}(\text{VK}_{|k-1}^j) & \text{if } T = g^{a+b} \end{cases} .$$

Namely, \mathcal{B} computes $\mathbf{F} = (F_1, F_2, F_3) = \mathbf{f}_0 \cdot \mathbf{f}_1^{\mu_1} \cdot \mathbf{f}_2^{\mu_2}$ for $\mu_1, \mu_2 \xleftarrow{R} \mathbb{Z}_p$ and uses the j -th tuple (K_j, L_j, T_j) to set

$$\begin{aligned} \sigma_1 &= g^{\omega_1 + \omega_2} \cdot R_{k-1}(\text{VK}_{|k-1}^j) \cdot K_j^{\sum_{\ell=1}^L \alpha_{\ell, \text{VK}[\ell]^j}} \cdot L_j^{\sum_{\ell=1}^L \beta_{\ell, \text{VK}[\ell]^j}} \cdot T_j^\zeta, \\ \sigma_2 &= K_j = f^{a_j}, \quad \sigma_3 = L_j = h^{b_j} . \end{aligned}$$

If $T_j = g^{a_j + b_j}$, the above implicitly defines

$$\sigma_1 = g^{\omega_1 + \omega_2} \cdot R_{k-1}(\text{VK}_{|k-1}^j) \cdot H(\mathbf{V}, \text{VK}^j)^{a_j} \cdot H(\mathbf{W}, \text{VK}^j)^{b_j},$$

so that $(\sigma_1, \sigma_2, \sigma_3)$ has the same distribution as in $\text{Game}_{5.2.(k-1)}$. If $T_j = g^{a_j + b_j + \tau_j}$, we can write

$$\sigma_1 = g^{\omega_1 + \omega_2} \cdot R_k(\text{VK}_{|k}^j) \cdot H(\mathbf{V}, \text{VK}^j)^{a_j} \cdot H(\mathbf{W}, \text{VK}^j)^{b_j},$$

since $R_k(\text{VK}_{|k}^j) = R_{k-1}(\text{VK}_{|k-1}^j) \cdot g^{\zeta \cdot \tau_j}$, which is distributed as in $\text{Game}_{5.2.k}$. In either case, (Z, R, U) is computed using $\text{sk}_1 = \{(\varphi_i, \phi_i, \vartheta_i)\}_{i=1}^{4L+6}$ as per (16).

– If $\text{VK}[k]^j = 1 - b^\dagger$ and $R_k(\text{VK}_{|k}^j)$ was defined, \mathcal{B} recalls the index $j' < j$ of the query where this value was defined. It sets $\mathbf{f}_0 = (F_1, F_2, F_3) = \mathbf{f}_0 \cdot \mathbf{f}_1^{\mu_1} \cdot \mathbf{f}_2^{\mu_2}$ for randomly chosen $\mu_1, \mu_2 \xleftarrow{R} \mathbb{Z}_p$. It also picks $r, s \xleftarrow{R} \mathbb{Z}_p$ and re-uses the j' -th tuple $(K_{j'}, H_{j'}, T_{j'})$ to set

$$\begin{aligned} \sigma_1 &= g^{\omega_1 + \omega_2} \cdot R_{k-1}(\text{VK}_{|k-1}^j) \cdot K_{j'}^{\sum_{\ell=1}^L \alpha_{\ell, \text{VK}[\ell]^j}} \cdot H_{j'}^{\sum_{\ell=1}^L \beta_{\ell, \text{VK}[\ell]^j}} \cdot T_{j'}^\zeta \cdot \\ &\quad H(\mathbf{V}, \text{VK}^j)^r \cdot H(\mathbf{W}, \text{VK}^j)^s, \\ \sigma_2 &= K_{j'} \cdot f^r = f^{a_{j'} + r}, \quad \sigma_3 = H_{j'}^x \cdot h^s = h^{b_{j'} + s}, \end{aligned}$$

and generates (Z, R, U) using $\text{sk}_1 = \{(\varphi_i, \phi_i, \vartheta_i)\}_{i=1}^{4L+6}$ as in the previous cases.

At the end of the game, the adversary \mathcal{A} outputs a triple $(\mathbf{v}^*, \pi^*, \text{lbl}^*)$ together with a convincing proof $\pi^* = (\text{VK}^*, \mathbf{F}^*, \mathbf{C}_{\sigma_1}^*, \sigma_2^*, \sigma_3^*, \mathbf{C}_Z^*, \mathbf{C}_R^*, \mathbf{C}_U^*, \mathbf{C}_z^*, \mathbf{C}_r^*, \mathbf{C}_u^*, \pi_{\sigma_1,1}^*, \pi_{\sigma_2,2}^*, \pi_1^*, \pi_2^*, \sigma^*)$. At this point, our distinguisher \mathcal{B} halts and outputs a random bit if it happens that $\text{VK}[k]^* \neq b^\dagger$. Otherwise, \mathcal{B} can figure out if π^* is of the same (extended) type as the outputs of $\mathcal{S}_2(\psi, \tau_{sim}, \cdot, \cdot)$. Since $\text{VK}[k]^* = b^\dagger$, \mathcal{B} is able to compute $F(\mathbf{v}, \text{VK}^*) = \sum_{\ell=1}^L \alpha_{\ell, \text{VK}[\ell]^*}$ and $F(\mathbf{w}, \text{VK}^*) = \sum_{\ell=1}^L \beta_{\ell, \text{VK}[\ell]^*}$ and can use the extraction trapdoor $(y_1, y_2) = (\log_g(f_1), \log_g(f_2))$ to extract $(\sigma_1^*, Z^*, R^*, U^*)$ from their commitments $\mathbf{C}_{\sigma_1}^*, \mathbf{C}_Z^*, \mathbf{C}_R^*, \mathbf{C}_U^*$. This allows \mathcal{B} to compute

$$\eta^* = \sigma_1^* \cdot \sigma_2^{*-F(\mathbf{v}, \text{VK}^*)} \cdot \sigma_3^{*-F(\mathbf{w}, \text{VK}^*)}, \quad F_0^* = F_1^{*-1/y_1} \cdot F_2^{*-1/y_2} \cdot F_3^* .$$

If $e(\eta^*, g) = e(g^{\omega_1 + \omega_2} \cdot \text{R}_{k-1}(\text{VK}_{|k-1}^*), F_0^*)$, \mathcal{B} knows that σ_1^* is of the form

$$\sigma_1^* = g^{(\omega_1 + \omega_2) \cdot x} \cdot \text{R}_{k-1}(\text{VK}_{|k-1}^*)^x \cdot H(\mathbf{V}, \text{VK}^*)^{r^*} \cdot H(\mathbf{W}, \text{VK}^*)^{s^*}, \quad (17)$$

where $r^* = \log_f(\sigma_2^*)$, $s^* = \log_h(\sigma_3^*)$ and $x = \log_g(F_0^*)$. Therefore \mathcal{B} can conclude that π^* is a fake proof of the same (extended) type as those generated by the simulated $\mathcal{S}_2(\psi, \tau_{sim}, \cdot, \cdot)$ and outputs 1. Indeed, given that $\text{R}_{k-1}(\text{VK}_{|k-1}^*) = \text{R}_k(\text{VK}_k^*)$, if the equality (17) holds, π^* has the same distribution as outputs of $\mathcal{S}_2(\psi, \tau_{sim}, \cdot, \cdot)$ in both $\text{Game}_{5.2.k}$ and $\text{Game}_{5.2.(k-1)}$. Otherwise, \mathcal{B} concludes that π^* deviates from the output distribution of $\mathcal{S}_2(\psi, \tau_{sim}, \cdot, \cdot)$ and outputs 0. If the difference between π^* 's probability to emulate the behavior of $\mathcal{S}_2(\psi, \tau_{sim}, \cdot, \cdot)$ in $\text{Game}_{5.2.k}$ and $\text{Game}_{5.2.(k-1)}$ is ϵ , then \mathcal{B} 's advantage as a DLIN distinguisher is at least $\epsilon/2$ given that the choice of $b^\dagger \in \{0, 1\}$ is independent of \mathcal{A} 's view. \square

In the multi-CRS setting, the proof of Lemma 4 can be adapted in the same way as the proof of Lemma 3. In short, \mathcal{B} generates a set μ independent CRSes $\{\psi^{(\kappa)} = (\mathbf{CRS}_1^{(\kappa)}, \mathbf{CRS}_2^{(\kappa)})\}_{\kappa=1}^\mu$, where each $(\mathbf{CRS}_1^{(\kappa)}, \mathbf{CRS}_2^{(\kappa)})$ contains fresh vectors $\mathbf{V}^{(\kappa)}, \mathbf{W}^{(\kappa)}$ in which \mathcal{B} embeds the generators (f, h, g) . At the beginning of the game, \mathcal{B} generates $q \cdot \mu$ tuples

$$(K_j^{(\kappa)}, L_j^{(\kappa)}, T_j^{(\kappa)}) = (f^{a_j^{(\kappa)}}, h^{b_j^{(\kappa)}}, g^{a_j^{(\kappa)} + b_j^{(\kappa)} + \tau_j^{(\kappa)}}) \quad j \in \{1, \dots, q\}, \kappa \in \{1, \dots, \mu\},$$

where $\tau_j^{(\kappa)}$ may be zero or not. These randomized DLIN instances—which all have the same answer determined by that of the original DLIN instance—will be used to define μ independent random functions $\text{R}_{k-1}^{(\kappa)} : \{0, 1\}^k \rightarrow \mathbb{G}$.

F Proof of Theorem 4

Before describing the scheme and giving its security proof, let us first recall the definition of chosen-ciphertext security in the multi-user setting in the sense of Bellare, Boldyreva and Micali [7].

F.1 Public-Key Encryption in the Multi-User Setting

In the multi-user setting [7], a public-key encryption scheme consists of algorithms (Par-Gen, Keygen, Encrypt, Decrypt), where Par-Gen takes as input a security parameter λ and generates common public parameters Γ shared by all users, Keygen takes as input Γ and outputs a key pair (SK, PK) , and algorithms Encrypt and Decrypt that proceed in the usual way.

Definition 4 ([7]). A public-key encryption scheme is (μ, q_e) -IND-CCA secure, for integers $\mu, q_e \in \text{poly}(\lambda)$, if no PPT adversary has noticeable advantage in this game:

1. The challenger first generates $\Gamma \leftarrow \text{Par-Gen}(\lambda)$ and runs $(SK^{(i)}, PK^{(i)}) \leftarrow \text{Keygen}(\Gamma)$ for $i = 1$ to μ . It gives $\{PK^{(i)}\}_{i=1}^\mu$ to the adversary \mathcal{A} and retains $\{SK^{(i)}\}_{i=1}^\mu$. In addition, the challenger initializes a set $\mathcal{D} \leftarrow \emptyset$ and a counter $j_q \leftarrow 0$. Finally, it chooses a random bit $d \xleftarrow{R} \{0, 1\}$.
2. The adversary \mathcal{A} adaptively makes queries to the following oracles on multiple occasions:
 - Encryption query: \mathcal{A} chooses an index $i \in \{1, \dots, \mu\}$ and a pair (M_0, M_1) of equal-length messages. If $j_q = q_e$, the oracle returns \perp . Otherwise, it computes $C \leftarrow \text{Encrypt}(PK^{(i)}, M_d)$ and returns C . In addition, it sets $\mathcal{D} \leftarrow \mathcal{D} \cup \{(i, C)\}$ and $j_q \leftarrow j_q + 1$.
 - Decryption query: \mathcal{A} can also invoke the decryption oracle on arbitrary ciphertexts C and indexes $i \in \{1, \dots, \mu\}$. If $(i, C) \in \mathcal{D}$, the oracle returns \perp . Otherwise, the oracle returns $M \leftarrow \text{Decrypt}(SK^{(i)}, C)$, which may be \perp if C is an invalid ciphertext.
3. The adversary \mathcal{A} outputs a bit d' and is deemed successful if $d' = d$. As usual, \mathcal{A} 's advantage is measured as the distance $\mathbf{Adv}(\mathcal{A}) = |2 \cdot \Pr[d' = d] - 1|$.

The proof that the scheme of Section 5.1 provides $(1, q_e)$ -IND-CCA security applies standard techniques [55,58] and proceeds as follows.

Proof. The proof uses of a sequence of games starting with a game where the challenger's hidden bit is $d = 0$ and ending with a game where $d = 1$. For each i , S_i is the event that \mathcal{A} wins in Game_i .

Game₁: is the real attack game where the challenger's bit is $d = 0$. In details, the adversary is given the public key PK while the challenger keeps the private keys SK to itself. At each decryption query, \mathcal{B} faithfully runs the real decryption algorithm using the private key $SK = (x_1, y_1)$. At the j -th encryption query, for $j \in \{1, \dots, q_e\}$, the adversary \mathcal{A} chooses messages $M_0^{(j)}, M_1^{(j)} \in \mathbb{G}$ and obtains a challenge ciphertext $C_j^* = (C_{j,0}^*, C_{j,1}^*, C_{j,2}^*, D_{j,0}^*, D_{j,1}^*, D_{j,2}^*, \pi_j^*)$ which is an encryption of $M_0^{(j)}$. Decryption queries are disallowed for ciphertexts C returned by the encryption oracle. Eventually, \mathcal{A} halts and outputs a bit $d' \in \{0, 1\}$. We denote by S_1 the event that $d' = 0$.

Game₂: We change the decryption oracle. Instead of faithfully using the private key $SK = (x_1, y_1)$ to compute $M = C_0 \cdot C_1^{-1/x_1} \cdot C_2^{-1/y_1}$ at each valid decryption query $C = (C_0, C_1, C_2, D_0, D_1, D_2, \pi)$, \mathcal{B} recalls the values $(x_2, y_2) \in \mathbb{Z}_p^2$ and computes the plaintext as $M = D_0 \cdot D_1^{1/x_2} \cdot D_2^{1/y_2}$. Clearly, \mathcal{A} 's view will not be affected by this change unless it is able to invoke the decryption oracle on a valid-looking ciphertext although (C_0, C_1, C_2) and (D_0, D_1, D_2) are BBS encryptions of distinct messages. If we call the latter event E_2 , we have the inequality $|\Pr[S_1] - \Pr[S_2]| \leq \Pr[E_2]$. Moreover, event E_2 would contradict the enhanced soundness of the proof system, which is tightly related to the DLIN assumption as shown by Theorem 2. Concretely, the probability of E_2 is at most $\Pr[E_2] \leq 2 \cdot \mathbf{Adv}^{\text{DLIN}}(\lambda) + 2/p$.

Game₃: This game is like **Game₂** except that, at each encryption query $(M_0^{(j)}, M_1^{(j)})$, the returned ciphertext $C_j^* = (C_{j,0}^*, C_{j,1}^*, C_{j,2}^*, D_{j,0}^*, D_{j,1}^*, D_{j,2}^*, \pi^*)$ is obtained by computing π^* as a simulated proof using the simulation trapdoor τ_{sim} associated with the language ρ defined by PK . The quasi-adaptive zero-knowledge property of the simulation-sound proof system guarantees that \mathcal{A} 's view will not be affected by this change. We have $\Pr[S_3] = \Pr[S_2]$.

Game₄: We modify the treatment of encryption queries $\{(M_0^{(j)}, M_1^{(j)})\}_{j=1}^{q_e}$. When \mathcal{B} computes the j -th challenge ciphertext $C_j^* = (C_{j,0}^*, C_{j,1}^*, C_{j,2}^*, D_{j,0}^*, D_{j,1}^*, D_{j,2}^*, \pi^*)$, it computes a hybrid ciphertext where $(C_{j,0}^*, C_{j,1}^*, C_{j,2}^*)$ is a BBS encryption of $M_1^{(j)}$ and $(D_{j,0}^*, D_{j,1}^*, D_{j,2}^*)$ is a BBS

encryption of $M_0^{(j)}$. It is easy to prove that any PPT adversary \mathcal{A} having noticeably different behaviors in Game_3 and Game_4 would imply an adversary against the semantic security of the BBS cryptosystem in the multi-challenge setting, which would contradict the DLIN assumption. Indeed, Hofheinz and Jager proved [37, Theorem 6] that the multi-challenge (and multi-user) semantic security of BBS is tightly related to the DLIN assumption. The result of [37, Theorem 6] implies that $|\Pr[S_4] - \Pr[S_3]| \leq \mathbf{Adv}^{\text{DLIN}}(\lambda) + 1/p$.

Game₅: We modify again the decryption oracle. This time, instead of using the secondary private key (x_2, y_2) to recover the plaintext $M = D_0 \cdot D_1^{-1/x_2} \cdot D_2^{-1/y_2}$ at each valid decryption query $C = (C_0, C_1, C_2, D_0, D_1, D_2, \pi)$, the challenger \mathcal{B} switches back to using the actual private key $SK = (x_1, y_1)$ to compute $M = C_0 \cdot C_1^{-1/x_1} \cdot C_2^{-1/y_1}$. It is easy to see that \mathcal{A} 's view will be the same as in Game_4 until \mathcal{A} manages to query the decryption oracle on a valid-looking ciphertext C for which (C_0, C_1, C_2) and (D_0, D_1, D_2) encrypt distinct messages. If we denote by E_5 the latter event, we claim that it contradicts the enhanced unbounded simulation-soundness of the underlying proof system.

Indeed, \mathcal{B} can be turned into an adversary that breaks the latter property with advantage $\Pr[E_5]$ for the language \mathcal{L}_ρ , where

$$\rho = \begin{pmatrix} X_1 & 1 & g & X_1 & 1 \\ 1 & Y_1 & g & Y_1 & 1 \\ X_2 & 1 & g & 1 & X_2 \\ 1 & Y_2 & g & 1 & Y_2 \end{pmatrix} \in \mathbb{G}^{4 \times 5},$$

for uniformly random $X_1, X_2, Y_1, Y_2, g \xleftarrow{R} \mathbb{G}$. To this end, \mathcal{B} interacts with a simulation-soundness challenger which supplies it with a common reference string $(\Gamma, \psi = (\mathbf{CRS}_1, \mathbf{CRS}_2))$, where \mathbf{CRS}_1 contains a description of the language \mathcal{L}_ρ . In addition, \mathcal{B} receives the matrix of discrete logarithms $\mathbf{A} = \log(\rho) \in \mathbb{Z}_p^{4 \times 5}$ which allows deciding membership in \mathcal{L}_ρ . Using \mathbf{A} and ψ , \mathcal{B} can construct a properly distributed public key PK of which it knows the private key $SK = (x_1, y_1)$ and its twin (x_2, y_2) that are part of $\mathbf{A} = \log(\rho)$. Hence, \mathcal{B} can interact with the CCA2 adversary \mathcal{A} in the same way as the challenger of Game_5 does. Also, \mathcal{B} can recognize the first fatal decryption query $C = (C_0^\dagger, C_1^\dagger, C_2^\dagger, D_0^\dagger, D_1^\dagger, D_2^\dagger, \pi^\dagger)$ (i.e., the first valid query where $(C_0^\dagger, C_1^\dagger, C_2^\dagger)$ and $(D_0^\dagger, D_1^\dagger, D_2^\dagger)$ are BBS encryptions of distinct messages) since it knows (x_1, y_1, x_2, y_2) . At this point, \mathcal{B} halts and outputs $(\mathbf{v}^\dagger, \pi^\dagger)$ where

$$\mathbf{v}^\dagger = (C_1^\dagger/D_1^\dagger, C_2^\dagger/D_2^\dagger, C_0^\dagger/D_0^\dagger, C_1^\dagger \cdot C_2^\dagger, D_1^{\dagger-1} \cdot D_2^{\dagger-1}),$$

which breaks the enhanced unbounded simulation-soundness of the QA-NIZK proof system. The result of Theorem 3 implies that

$$|\Pr[S_5] - \Pr[S_4]| \leq \Pr[E_5] \leq \mathbf{Adv}_B^{q_e\text{-suf-ots}}(\lambda) + 3 \cdot (L + 2) \cdot \mathbf{Adv}_B^{\text{DLIN}}(\lambda) + 4/p.$$

Game₆: In this game, we bring yet another modification to the generation of challenge ciphertexts $C_j^* = (C_{j,0}^*, C_{j,1}^*, C_{j,2}^*, D_{j,0}^*, D_{j,1}^*, D_{j,2}^*, \pi^*)$. Namely, in all encryption queries $\{(M_0^{(j)}, M_1^{(j)})\}_{j=1}^{q_e}$, instead of generating a hybrid ciphertext where the built-in BBS ciphertexts encrypt distinct messages, $(C_{j,0}^*, C_{j,1}^*, C_{j,2}^*)$, and $(D_{j,0}^*, D_{j,1}^*, D_{j,2}^*)$ are both calculated by encrypting $M_1^{(j)}$. It is easy to prove that any noticeable change in \mathcal{A} 's behavior between Game_5 and Game_6 would imply an IND-CPA adversary against the BBS cryptosystem in the multi-challenge setting.

The result of [37, Theorem 6] thus implies $|\Pr[S_6] - \Pr[S_5]| \leq \mathbf{Adv}^{\text{DLIN}}(\lambda) + 1/p$. Note that, for each ciphertext C_j^* , the vectors $(C_{j,1}^*/D_{j,1}^*, C_{j,2}^*/D_{j,2}^*, C_{j,0}^*/D_{j,0}^*, C_{j,1}^* \cdot C_{j,2}^*, (D_{j,1}^* \cdot D_{j,2}^*)^{-1})$ are now back in the linear span of $\mathbf{X}_1, \mathbf{Y}_1, \mathbf{X}_2, \mathbf{Y}_2$.

Game₇: We bring one last change to the generation of the challenge ciphertexts $\{C_j^*\}_{j=1}^{q_e}$. For each ciphertext C_j^* , instead of computing π_j^* using the simulation trapdoor τ_{sim} , we compute it using the witnesses $(\theta_{j,1}, \theta_{j,2}, -\theta_{j,3}, -\theta_{j,4}) \in \mathbb{Z}_p^4$. This change is only conceptual since the obtained proofs have the same distribution as in **Game₆**. We have $\Pr[S_7] = \Pr[S_6]$.

We observe that **Game₇** corresponds to the actual game where the challenger's bit is $d = 1$. If we combine the above, we thus find the announced upper bound for the distance $|\Pr[S_1] - \Pr[S_7]|$. \square

G More Efficient (Almost) Tight Chosen-Ciphertext Security in the Key-Dependent Message Setting

In this section, we apply our unbounded simulation-sound QA-NIZK argument system in the context of key-dependent-message (KDM) security. Specifically, we describe an almost tightly secure variant of the KDM-CCA2 secure cryptosystem suggested by Camenisch, Chandran and Shoup [19] with substantially shorter ciphertexts than in previous tightly secure solutions.

Key-dependent message security [11] refers to encryption schemes that remain secure when the adversary obtains encryptions of functions of the secret key. The Camenisch *et al.* scheme can be seen as a chosen-ciphertext-secure variant of the construction initially given by Boneh, Halevi, Hamburg and Ostrovsky [18] (BHHO). The schemes of [18,19] both need $O(\ell)$ group elements in the ciphertext, where ℓ is proportional to the security parameter.

Camenisch, Chandran and Shoup [19] combined the BHHO construction with the Naor-Yung paradigm [55] in order to attain chosen-ciphertext security. In their constructions, they need NIZK proofs of plaintext equalities consisting of $O(\ell)$ group elements. Recent works by Libert *et al.* [48] and Jutla and Roy [42] independently showed how to achieve KDM-CCA2 security using proofs of plaintext equalities made of $O(1)$ group elements. Unfortunately, their security reduction is affected by the number q_e of challenge ciphertexts obtained by the adversary. Using our simulation-sound proofs, we can eliminate the latter disadvantage and obtain nearly tight KDM-CCA2 without sacrificing the constant-size proofs of plaintext equalities. In comparison with [19], the $O(\ell)$ -size Groth-Sahai-based proof of [19] is traded for a tightly simulation-sound argument comprised of 42 group elements.

Before outlining the construction, we first recall the definition of chosen-ciphertext security in the KDM setting. Let \mathcal{S} be the space of secret keys produced by the key generation algorithm of a public-key encryption scheme. Let $N > 0$ be an integer and let $\mathcal{C} = \{f : \mathcal{S}^N \rightarrow \mathcal{M}\}$ be a family of functions, where \mathcal{M} is the message space.

Definition 5 ([19]). *A public-key encryption scheme (Keygen, Encrypt, Decrypt) is KDM-CCA2 secure for the function family \mathcal{C} if no PPT adversary has non-negligible advantage in the game below.*

1. *The challenger runs the Keygen algorithm N times to generate pairs $(pk_1, sk_1), \dots, (pk_N, sk_N)$ and sends (pk_1, \dots, pk_N) to the adversary. The challenger also flips a random binary coin $d \xleftarrow{R} \{0, 1\}$.*
2. *The adversary interleaves the following kinds of queries.*

Encryption queries: \mathcal{A} specifies a pair (i, f) , where $i \in \{1, \dots, N\}$ and $f \in \mathcal{C}$. The challenger sets $M = f(sk_1, \dots, sk_N) \in \mathcal{M}$. If $d = 0$, the challenger returns $C = \text{Encrypt}(pk_i, M)$. If $d = 1$, it returns $C = \text{Encrypt}(pk_i, 0^{|M|})$ and stores the pair (i, C) in the list L of target ciphertexts, which is initially empty.

Decryption queries: \mathcal{A} submits a pair (i, C) . If $(i, C) \in L$, the challenger returns \perp . Otherwise, it returns $M \leftarrow \text{Decrypt}(sk_i, C)$.

3. The adversary \mathcal{A} outputs a bit $d' \in \{0, 1\}$ and wins if $d' = d$. As usual, \mathcal{A} 's advantage is defined to be $\text{Adv}^{\text{kdm-cca2}}(\mathcal{A}) = |\Pr[d' = d] - 1/2|$.

So far, the most efficient KDM-CCA2-secure construction is a scheme due to Hofheinz [39], which relies on the DLIN and Composite Residuosity [56] assumptions. While much more efficient than [19], his construction relies on several number theoretic assumptions and it is not known to provide tight security. The results of [37,4] do imply tight KDM-CCA2 security but, for typical choices of parameters, their simulation-sound proofs introduce several hundreds of group elements in the ciphertext, as discussed below. Our unbounded simulation-sound QA-NIZK proofs improve upon those tightly KDM-CCA2 constructions in that the KDM-CCA2 system is only longer than its underlying KDM-CPA variant (due to [18] and described in [19, Section 4.2]) by 42 group elements. Also, while our scheme does not compete with [39] from an efficiency point of view, its security proof only requires the DLIN assumption. Like [19], it applies the Naor-Yung paradigm to show that a BHHO ciphertext [18] encrypts the same message as a BBS ciphertext [15]. It goes as follows.

Keygen(λ):

1. Run the K_0 algorithm of Section 3 to obtain $\Gamma = ((\mathbb{G}, \mathbb{G}_T), f, g, h, \Sigma)$ and set $\ell = \lceil 4 \log p \rceil$. Then, choose generators $g \xleftarrow{R} \mathbb{G}, g_1, \dots, g_\ell \xleftarrow{R} \mathbb{G}, h_1, \dots, h_\ell \xleftarrow{R} \mathbb{G}$ as well as $x, y \xleftarrow{R} \mathbb{Z}_p$ a ℓ -bit string $s = s_1 \dots s_\ell \xleftarrow{R} \{0, 1\}^\ell$. Then, define $X = g^x, Y = g^y, g_0 = \prod_{i=1}^\ell g_i^{-s_i}, h_0 = \prod_{i=1}^\ell h_i^{-s_i}$. Next, construct the linearly independent vectors

$$\begin{aligned} \mathbf{g} &= (g_0, g_1, \dots, g_\ell, 1, 1) \in \mathbb{G}^{\ell+3}, & \mathbf{h} &= (h_0, h_1, \dots, h_\ell, 1, 1) \in \mathbb{G}^{\ell+3} \\ \mathbf{X} &= (g, 1, \dots, 1, X, 1) \in \mathbb{G}^{\ell+3}, & \mathbf{Y} &= (g, 1, \dots, 1, 1, Y) \in \mathbb{G}^{\ell+3} \end{aligned}$$

and erase the exponents $x, y \in \mathbb{Z}_p$.

2. Run algorithm $K_1(\Gamma, \rho)$ of Section 3 to generate a CRS for a QA-NIZK proof system, where the language parameter is the matrix $\rho \in \mathbb{G}^{4 \times (\ell+3)}$ whose rows consist of the vectors $\mathbf{g}, \mathbf{h}, \mathbf{X}$ and \mathbf{Y} . Let $\psi = (\text{CRS}_1, \text{CRS}_2)$ be the obtained CRS, where

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

3. Define the private key as $SK = (g^{s_1}, \dots, g^{s_\ell}) \in \mathbb{G}^\ell$. The public key is defined to be

$$PK = \left(g, X, Y, \{g_i\}_{i=0}^\ell, \{h_i\}_{i=0}^\ell, \psi = (\text{CRS}_1, \text{CRS}_2) \right).$$

Encrypt(M, PK): to encrypt $M \in \mathbb{G}$ under the public key PK , do the following.

1. Choose $\theta_1, \theta_2 \xleftarrow{R} \mathbb{Z}_p$ and compute

$$C_0 = M \cdot g_0^{\theta_1} \cdot h_0^{\theta_2}, \quad C_1 = g_1^{\theta_1} h_1^{\theta_2}, \quad C_2 = g_2^{\theta_1} h_2^{\theta_2}, \quad \dots, \quad C_\ell = g_\ell^{\theta_1} h_\ell^{\theta_2}$$

2. Choose $\theta_3, \theta_4 \xleftarrow{R} \mathbb{Z}_p$ and compute

$$D_0 = M \cdot g^{\theta_3 + \theta_4}, \quad D_1 = X^{\theta_3}, \quad D_2 = Y^{\theta_4} .$$

3. Define the label $\text{lbl} = (C_0, C_1, \dots, C_\ell, D_0, D_1, D_2)$ and, for this label, generate a simulation-sound QA-NIZK proof π that the vector

$$\begin{aligned} \mathbf{v} &= (C_0/D_0, C_1, \dots, C_\ell, D_1^{-1}, D_2^{-1}) \\ &= (g_0^{\theta_1} \cdot h_0^{\theta_2} \cdot g^{-\theta_3 - \theta_4}, g_1^{\theta_1} \cdot h_1^{\theta_2}, \dots, g_\ell^{\theta_1} \cdot h_\ell^{\theta_2}, X^{-\theta_3}, Y^{-\theta_4}) \end{aligned} \quad (18)$$

belongs to $\text{span}\langle \mathbf{g}, \mathbf{h}, \mathbf{X}, \mathbf{Y} \rangle$. This is done by running Steps 1-7 of algorithm P in Section 3.2. Then, output the ciphertext

$$C = (C_0, C_1, \dots, C_\ell, D_0, D_1, D_2, \pi), \quad (19)$$

where the proof π is comprised of 38 group elements and a one-time key pair (VK, σ) .

Decrypt(SK, C): parse the private key as $SK = (v_1, \dots, v_\ell) \in \mathbb{G}^\ell$ and the ciphertext as in (19).

Return \perp if C does not parse properly or if π is not an accepting QA-NIZK proof that the vector $(C_0/D_0, C_1, \dots, C_\ell, D_1^{-1}, D_2^{-1})$ lives in $\text{span}\langle \mathbf{g}, \mathbf{h}, \mathbf{X}, \mathbf{Y} \rangle$. Otherwise, conduct the following steps.

1. For $i = 1$ to ℓ , set $s_i = 1$ if $v_i \neq 1_{\mathbb{G}}$ and $s_i = 0$ otherwise.
2. Output $M = C_0 \cdot \prod_{i=1}^{\ell} C_i^{s_i}$.

In the above system, π only consists of 42 group elements in an instantiation with the one-time signature of [37]. In comparison, all previous simulation-sound or simulation-extractable proof systems [37,4,49] enabling tight KDM-CCA2 security would require to prove $\Theta(\ell)$ linear pairing product equations, each of which takes 3 group elements. For example, the simulation-extractable proof of [4] – which allows eliminating (D_0, D_1, D_2) from the ciphertext – requires Groth-Sahai commitments to $(M, W_1, W_2) = (M, g^{\theta_1}, g^{\theta_2})$ and NIWI proofs for the equations $e(C_0/M, g) = e(g_0, W_1) \cdot e(h_0, W_1)$ and $e(C_i, g) = e(g_i, W_1) \cdot e(h_i, W_2)$ for each $i \in \{1, \dots, \ell\}$, which demands $\lceil 12 \log p \rceil + 69$ group elements in an instantiation with the signature scheme of [49]. This incurs a total of $\lceil 16 \log p \rceil + 70$ group elements per ciphertext.

Our ciphertexts only take $\lceil 4 \log p \rceil + 45$ group elements, which is nearly 75% shorter than in the best previous tightly secure KDM-CCA2 system for any realistic security parameter. Our scheme is – up to an additive overhead of 42 group elements – essentially as efficient as its KDM-CPA variant. Although it remains significantly less efficient than Hofheinz’s KDM-CCA2-secure construction [39], it turns out to be the most efficient scheme with (nearly) tight KDM-CCA2 security to date.

The security in the sense of Definition 5 is proved using standard arguments and we omit the details of the proof, which proceeds exactly like the one of Camenisch *et al.* [19, Appendix A.1]. It naturally reduces the KDM-CCA2 security of the system to the KDM-CPA security of the BHHO construction [19, Section 4.2] (in particular, all encryption queries are relayed to the KDM-CPA challenger) and the enhanced unbounded simulation-soundness of our QA-NIZK proof system of Section 3.2 in the multi-CRS setting. The adversary’s advantage is bounded as in Theorem 4 but, in the multi-user setting, we need an additional term N/p .

Theorem 5. *The scheme is KDM-CCA2 secure assuming that: (i) Σ is a strongly unforgeable one-time signature; (ii) The DLIN assumption holds in \mathbb{G} . For any adversary \mathcal{A} , there is 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 $\text{Adv}_{\mathcal{A}}^{\text{kdm-cca}}(\lambda) \leq \text{Adv}_{\mathcal{B}'}^{q_e\text{-suf-ots}}(\lambda) + (3L + 10) \cdot \text{Adv}_{\mathcal{B}}^{\text{DLIN}}(\lambda) + (8 + N)/p$, where L is the length of one-time verification keys, N is the number of public keys and q_e is the number of encryption queries.*

H Extension to Asymmetric Pairings

In this section, we explain how to adapt our QA-NIZK proof system in the context of asymmetric pairings $e : \mathbb{G} \times \hat{\mathbb{G}} \rightarrow \mathbb{G}_T$, with $\mathbb{G} \neq \hat{\mathbb{G}}$. We consider both Type II pairings, where an isomorphism $\psi : \hat{\mathbb{G}} \rightarrow \mathbb{G}$ is efficiently computable, and Type III pairings, as defined in [32].

The only changes are that certain proof elements have to be in $\hat{\mathbb{G}}$ and, in the case of Type III pairings, the security proof has to rely on a slightly stronger version of the Decision Linear assumption where the challenge value is in $\mathbb{G} \times \hat{\mathbb{G}}$.

$\mathsf{K}_0(\lambda)$: choose asymmetric bilinear groups $(\mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_T)$ 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 $\Gamma = (\mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_T, f, g, h, \Sigma)$.

The dimensions (t, n) of the matrix $\mathbf{A} \in \mathbb{Z}_p^{t \times n}$ such that $\boldsymbol{\rho} = g^{\mathbf{A}}$ can be either fixed or part of the language, so that t, n can be given as input to the CRS generation algorithm K_1 .

$\mathsf{K}_1(\Gamma, \boldsymbol{\rho})$: parse Γ as $(\mathbb{G}, \mathbb{G}_T, f, g, h, \Sigma)$ and $\boldsymbol{\rho}$ as a matrix $\boldsymbol{\rho} = (G_{i,j})_{\substack{1 \leq i \leq t, \\ 1 \leq j \leq n}} \in \mathbb{G}^{t \times n}$.

1. Generate key pairs $\{(\text{sk}_b, \text{pk}_b)\}_{b=0}^1$ for the one-time linearly homomorphic signature of Section 2.3 in order to sign vectors of \mathbb{G}^n and \mathbb{G}^{4L+6} , respectively. Namely, pick generators $\hat{g}_z, \hat{g}_r, \hat{h}_z, \hat{h}_u \xleftarrow{R} \hat{\mathbb{G}}$ and $\hat{G}_z, \hat{G}_r, \hat{H}_z, \hat{H}_u \xleftarrow{R} \hat{\mathbb{G}}$. Then, for $i = 1$ to n , choose $\chi_i, \gamma_i, \delta_i \xleftarrow{R} \mathbb{Z}_p$ and compute $\hat{g}_i = \hat{g}_z^{\chi_i} \hat{g}_r^{\gamma_i}$ and $\hat{h}_i = \hat{h}_z^{\chi_i} \hat{h}_u^{\delta_i}$. Let $\text{sk}_0 = \{\chi_i, \gamma_i, \delta_i\}_{i=1}^n$ be the private key and let $\text{pk}_0 = (\hat{g}_z, \hat{g}_r, \hat{h}_z, \hat{h}_u, \{\hat{g}_i, \hat{h}_i\}_{i=1}^n)$ be the matching public key. The second 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

$$\text{pk}_1 = \left(\hat{G}_z, \hat{G}_r, \hat{H}_z, \hat{H}_u, \{\hat{G}_i = \hat{G}_z^{\varphi_i} \hat{G}_r^{\phi_i}, \hat{H}_i = \hat{H}_z^{\varphi_i} \hat{H}_u^{\vartheta_i}\}_{i=1}^{4L+6} \right).$$

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

$$\mathbf{V} = (V_{1,0}, V_{1,1}, \dots, V_{L,0}, V_{L,1}) \in \mathbb{G}^{2L}, \quad \mathbf{W} = (W_{1,0}, W_{1,1}, \dots, W_{L,0}, W_{L,1}) \in \mathbb{G}^{2L}.$$

4. Choose random $\omega_1, \omega_2 \xleftarrow{R} \mathbb{Z}_p$, $u_1, u_2 \xleftarrow{R} \mathbb{G}$, 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

$$(M_{i,j})_{i,j} = \begin{pmatrix} 1 & \mathbf{1}^{1 \times 2L} & \mathbf{1}^{1 \times 2L} & 1 & 1 & \mathbf{f}_1 \\ 1 & \mathbf{1}^{1 \times 2L} & \mathbf{1}^{1 \times 2L} & 1 & 1 & \mathbf{f}_2 \\ \mathbf{V}^\top & \mathbf{Id}_{f,2L} & \mathbf{1}^{2L \times 2L} & \mathbf{1}^{2L \times 1} & \mathbf{1}^{2L \times 1} & \mathbf{1}^{2L \times 3} \\ \mathbf{W}^\top & \mathbf{1}^{2L \times 2L} & \mathbf{Id}_{h,2L} & \mathbf{1}^{2L \times 1} & \mathbf{1}^{2L \times 1} & \mathbf{1}^{2L \times 3} \\ g & \mathbf{1}^{1 \times 2L} & \mathbf{1}^{1 \times 2L} & u_1 & 1 & \mathbf{1}^{1 \times 3} \\ g & \mathbf{1}^{1 \times 2L} & \mathbf{1}^{1 \times 2L} & 1 & u_2 & \mathbf{1}^{1 \times 3} \\ 1 & \mathbf{1}^{1 \times 2L} & \mathbf{1}^{1 \times 2L} & \Omega_1^{-1} & \Omega_2^{-1} & \mathbf{f}_0 \end{pmatrix} \quad (20)$$

with $\mathbf{Id}_{f,2L} = f \mathbf{I}_{2L} \in \mathbb{G}^{2L \times 2L}$, $\mathbf{Id}_{h,2L} = h \mathbf{I}_{2L} \in \mathbb{G}^{2L \times 2L}$, where $\mathbf{I}_{2L} \in \mathbb{Z}_p^{2L \times 2L}$ is 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 linearly 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 signatures are obtained as

$$(z_i, r_i, u_i) = \left(\prod_{j=1}^n G_{i,j}^{-\chi_j}, \prod_{j=1}^n G_{i,j}^{-\gamma_j}, \prod_{j=1}^n G_{i,j}^{-\delta_j} \right)$$

for each $i \in \{1, \dots, t\}$. Then, use the second LHSPS private key sk_1 to sign the rows $\mathbf{M}_j = (M_{j,1}, \dots, M_{j,4L+6}) \in \mathbb{G}^{4L+6}$ of the matrix (20) 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 $\tau_{sim} = (\omega_1, \omega_2, \{\chi_i, \gamma_i, \delta_i\}_{i=1}^n)$.

$\mathbf{P}(T, \psi, \mathbf{v}, x, \text{lbl})$: 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 $(\mathbf{VK}, \mathbf{SK}) \leftarrow \mathcal{G}(\lambda)$ and conduct the following steps.

1. Using $\{(z_j, r_j, u_j)\}_{j=1}^t$ from \mathbf{CRS}_1 , derive a one-time homomorphic signature $(z, r, u) \in \mathbb{G}^3$ on the vector \mathbf{v} by computing $z = \prod_{i=1}^t z_i^{x_i}$, $r = \prod_{i=1}^t r_i^{x_i}$ and $u = \prod_{i=1}^t u_i^{x_i}$.
2. Define a vector $\mathbf{F} = (F_1, F_2, F_3) = \mathbf{f}_1^{\mu_1} \cdot \mathbf{f}_2^{\mu_2}$, for randomly chosen $\mu_1, \mu_2 \xleftarrow{R} \mathbb{Z}_p$.
3. Pick $r, s \xleftarrow{R} \mathbb{Z}_p$ and compute a pseudo-signature $\sigma = (\sigma_1, \sigma_2, \sigma_3)$ on the one-time verification key $\mathbf{VK} = \mathbf{VK}[1] \dots \mathbf{VK}[L]$, which is obtained as

$$\sigma_1 = H(\mathbf{V}, \mathbf{VK})^r \cdot H(\mathbf{W}, \mathbf{VK})^s, \quad \sigma_2 = f^r, \quad \sigma_3 = h^s,$$

with $H(\mathbf{V}, \mathbf{VK}) = \prod_{\ell=1}^L V_{\ell, \mathbf{VK}[\ell]}$ and $H(\mathbf{W}, \mathbf{VK}) = \prod_{\ell=1}^L W_{\ell, \mathbf{VK}[\ell]}$.

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

$$\boldsymbol{\sigma} = (\sigma_1, \sigma_2^{1-\text{VK}[1]}, \sigma_2^{\text{VK}[1]}, \dots, \sigma_2^{1-\text{VK}[L]}, \sigma_2^{\text{VK}[L]}, \sigma_3^{1-\text{VK}[1]}, \sigma_3^{\text{VK}[1]}, \dots, \sigma_3^{1-\text{VK}[L]}, \sigma_3^{\text{VK}[L]}, 1_{\mathbb{G}}, 1_{\mathbb{G}}, F_1, F_2, F_3) \in \mathbb{G}^{4L+6} \quad (21)$$

which lives in the subspace spanned by the first $4L + 2$ rows of $\mathbf{M} \in \mathbb{G}^{(4L+5) \times (4L+6)}$. The coefficients $r, s, \mu_1, \mu_2 \in \mathbb{Z}_p$ thus allow deriving a homomorphic signature (Z, R, U) on the vector $\boldsymbol{\sigma} \in \mathbb{G}^{4L+6}$.

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 $\hat{\boldsymbol{\pi}}_{\sigma,1}, \hat{\boldsymbol{\pi}}_{\sigma,2} \in \hat{\mathbb{G}}^3$ that committed variables (σ_1, Z, R, U) satisfy the pairing product equations

$$\begin{aligned} e(Z, \hat{G}_z) \cdot e(R, \hat{G}_r) \cdot e(\sigma_1, \hat{G}_1) &= e(\sigma_2, \prod_{i=1}^L \hat{G}_{2i+\text{VK}[i]})^{-1} \cdot \\ &e(\sigma_3, \prod_{i=1}^L \hat{G}_{2L+2i+\text{VK}[i]})^{-1} \cdot \prod_{i=1}^3 e(F_i, \hat{G}_{4L+3+i})^{-1}, \\ e(Z, \hat{H}_z) \cdot e(U, \hat{H}_u) \cdot e(\sigma_1, \hat{H}_1) &= e(\sigma_2, \prod_{i=1}^L H_{2i+\text{VK}[i]})^{-1} \cdot \\ &e(\sigma_3, \prod_{i=1}^L \hat{H}_{2L+2i+\text{VK}[i]})^{-1} \cdot \prod_{i=1}^3 e(F_i, \hat{H}_{4L+3+i})^{-1}. \end{aligned} \quad (22)$$

6. Using the vector $\mathbf{F} \in \mathbb{G}^3$ defined at step 2, define a new Groth-Sahai CRS $\mathbf{F} = (\mathbf{f}_1, \mathbf{f}_2, \mathbf{F})$ and use it to compute Groth-Sahai 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 $(z, r, u) \in \mathbb{G}^3$ along with NIWI proofs $(\hat{\boldsymbol{\pi}}_1, \hat{\boldsymbol{\pi}}_2) \in \hat{\mathbb{G}}^6$ that \mathbf{v} and (z, r, u) satisfy

$$e(z, \hat{g}_z) \cdot e(r, \hat{g}_r) = \prod_{i=1}^n e(v_i, \hat{g}_i)^{-1}, \quad e(z, \hat{h}_z) \cdot e(u, \hat{h}_u) = \prod_{i=1}^n e(v_i, \hat{h}_i)^{-1}.$$

Let $(\mathbf{C}_z, \mathbf{C}_r, \mathbf{C}_u, \hat{\boldsymbol{\pi}}_1, \hat{\boldsymbol{\pi}}_2) \in \mathbb{G}^9 \times \hat{\mathbb{G}}^6$ be the resulting commitments and proofs.

7. Set $\sigma = \mathcal{S}(\text{SK}, (\mathbf{v}, \mathbf{F}, \mathbf{C}_{\sigma_1}, \sigma_2, \sigma_3, \mathbf{C}_Z, \mathbf{C}_R, \mathbf{C}_U, \mathbf{C}_z, \mathbf{C}_r, \mathbf{C}_u, \hat{\boldsymbol{\pi}}_{\sigma,1}, \hat{\boldsymbol{\pi}}_{\sigma,2}, \hat{\boldsymbol{\pi}}_1, \hat{\boldsymbol{\pi}}_2, \text{lbl}))$ and output

$$\pi = (\text{VK}, \mathbf{F}, \mathbf{C}_{\sigma_1}, \sigma_2, \sigma_3, \mathbf{C}_Z, \mathbf{C}_R, \mathbf{C}_U, \mathbf{C}_z, \mathbf{C}_r, \mathbf{C}_u, \hat{\boldsymbol{\pi}}_{\sigma,1}, \hat{\boldsymbol{\pi}}_{\sigma,2}, \hat{\boldsymbol{\pi}}_1, \hat{\boldsymbol{\pi}}_2, \sigma) \quad (23)$$

$\mathcal{V}(\Gamma, \psi, \mathbf{v}, \pi, \text{lbl})$: parse π as in (23) and \mathbf{v} as $(v_1, \dots, v_n) \in \mathbb{G}^n$. Return 1 if the conditions below are all satisfied. Otherwise, return 0.

- (i) $\mathcal{V}(\text{VK}, (\mathbf{v}, \mathbf{F}, \mathbf{C}_{\sigma_1}, \sigma_2, \sigma_3, \mathbf{C}_Z, \mathbf{C}_R, \mathbf{C}_U, \mathbf{C}_z, \mathbf{C}_r, \mathbf{C}_u, \hat{\boldsymbol{\pi}}_{\sigma,1}, \hat{\boldsymbol{\pi}}_{\sigma,2}, \hat{\boldsymbol{\pi}}_1, \hat{\boldsymbol{\pi}}_2, \text{lbl}), \sigma) = 1$;
- (ii) $\hat{\boldsymbol{\pi}}_{\sigma,1}, \hat{\boldsymbol{\pi}}_{\sigma,2}$ 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 the pairing product equations (7).

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

$$\prod_{i=1}^n E(\iota(v_i), \hat{g}_i)^{-1} = E(\mathbf{C}_z, \hat{g}_z) \cdot E(\mathbf{C}_r, \hat{g}_r) \cdot E(\mathbf{f}_1, \hat{\pi}_{1,1}) \cdot E(\mathbf{f}_2, \hat{\pi}_{1,2}) \cdot E(\mathbf{F}, \hat{\pi}_{1,3})$$

$$\prod_{i=1}^n E(\iota(v_i), \hat{h}_i)^{-1} = E(\mathbf{C}_z, \hat{h}_z) \cdot E(\mathbf{C}_u, \hat{h}_u) \cdot E(\mathbf{f}_1, \hat{\pi}_{2,1}) \cdot E(\mathbf{f}_2, \hat{\pi}_{2,2}) \cdot E(\mathbf{F}, \hat{\pi}_{2,3}).$$

The proof is comprised of 26 elements of \mathbb{G} , 12 elements of $\hat{\mathbb{G}}$ and a pair (VK, σ) . At the 128-bit security level, if elements of \mathbb{G} and $\hat{\mathbb{G}}$ can be represented using 256 bits and 512 bits, respectively, the proof fits within 1.68 kB if the system is instantiated using the one-time signature of [37].

The simulation-soundness property can be proved under a variant of the DLIN assumption. This assumption posits that the two distributions

$$D_0 = \{(f, g, h, \hat{g}, f^a, h^b, g^{a+b}, \hat{g}^{a+b}) \mid f, g, h \xleftarrow{R} \mathbb{G}, \hat{g} \xleftarrow{R} \hat{\mathbb{G}}, a, b \xleftarrow{R} \mathbb{Z}_p\},$$

$$D_1 = \{(f, g, h, \hat{g}, f^a, h^b, g^c, \hat{g}^c) \mid f, g, h \xleftarrow{R} \mathbb{G}, \hat{g} \xleftarrow{R} \hat{\mathbb{G}}, a, b, c \xleftarrow{R} \mathbb{Z}_p\}$$

are indistinguishable. Note that, in Type II pairings where an isomorphism $\psi : \hat{\mathbb{G}} \rightarrow \mathbb{G}$ is efficiently computable, this assumption is implied by the standard DLIN assumption in $\hat{\mathbb{G}}$.

In the security proof, the main changes occur in the proofs of Lemmas 3 and 4. In the former, the reduction takes as input $(f, g, h, \hat{g}, f^a, h^b)$ and twinned challenge values $(T, \hat{T}) = (g^z, \hat{g}^z) \in \mathbb{G} \times \hat{\mathbb{G}}$, where either $z = a + b$ or $z \in_R \mathbb{Z}_p$. The reason why (T, \hat{T}) are both necessary is that, while T is used to compute σ_1 in simulated proofs, implementing the test (15) requires \hat{T} to test the equality $e(\eta^*, \hat{g}) = e(F_0^*, \hat{T})$ at the end of the game.

In the proof of Lemma 4, the main change is that, in order to perform an asymmetric analogue of the test $e(\eta^*, g) = e(g^{\omega_1 + \omega_2} \cdot \mathbf{R}_{k-1}(\text{VK}_{|k-1}^*), F_0^*)$ at the end of the game, the reduction \mathcal{B} needs two correlated random functions

$$\mathbf{R}_{k-1} : \{0, 1\}^{k-1} \rightarrow \mathbb{G}, \quad \hat{\mathbf{R}}_{k-1} : \{0, 1\}^{k-1} \rightarrow \hat{\mathbb{G}}$$

such that $e(\mathbf{R}_{k-1}(\text{VK}_{|k-1}^*), \hat{g}) = e(g, \hat{\mathbf{R}}_{k-1}(\text{VK}_{|k-1}^*))$ for any $\text{VK}^* \in \{0, 1\}^L$ in order to test whether the equality $e(\eta^*, \hat{g}) = e(F_0^*, \hat{g}^{\omega_1 + \omega_2} \cdot \hat{\mathbf{R}}_{k-1}(\text{VK}_{|k-1}^*))$ holds. This can be simply achieved by defining

$$\mathbf{R}_{k-1}(\text{VK}_{|k-1}^*) = g^{\mathbf{R}'_{k-1}(\text{VK}_{|k-1}^*)}, \quad \hat{\mathbf{R}}_{k-1}(\text{VK}_{|k-1}^*) = \hat{g}^{\mathbf{R}'_{k-1}(\text{VK}_{|k-1}^*)}$$

using any random function $\mathbf{R}'_{k-1}(\text{VK}_{|k-1}^*) : \{0, 1\}^{k-1} \rightarrow \mathbb{Z}_p$.