Simulation-Sound Arguments for LWE and Applications to KDM-CCA2 Security

Benoît Libert[1,2], Khoa Nguyen[3], Alain Passelègue[4,2], and Radu Titiu[5,2]

[1] CNRS, Laboratoire LIP, France
[2] ENS de Lyon, Laboratoire LIP (U. Lyon, CNRS, ENSL, Inria, UCBL), France
[3] Nanyang Technological University, SPMS, Singapore
[4] Inria, France
[5] Bitdefender, Bucharest, Romania

Abstract. The Naor-Yung paradigm is a well-known technique that constructs IND-CCA2-secure encryption schemes by means of non-interactive zero-knowledge proofs satisfying a notion of simulation-soundness. Until recently, it was an open problem to instantiate it under the sole Learning-With-Errors (LWE) assumption without relying on random oracles. While the recent results of Canetti *et al.* (STOC'19) and Peikert-Shiehian (Crypto'19) provide a solution to this problem by applying the Fiat-Shamir transform in the standard model, the resulting constructions are extremely inefficient as they proceed via a reduction to an NP-complete problem. In this paper, we give a direct, non-generic method for instantiating Naor-Yung under the LWE assumption outside the random oracle model. Specifically, we give a direct construction of an unbounded simulation-sound NIZK argument system which, for carefully chosen parameters, makes it possible to express the equality of plaintexts encrypted under different keys in Regev's cryptosystem. We also give a variant of our argument that provides tight security. As an application, we obtain an LWE-based public-key encryption scheme for which we can prove (tight) key-dependent message security under chosen-ciphertext attacks in the standard model.

Keywords. LWE, standard model, Naor-Yung, NIZK arguments, simulation-soundness, KDM-CCA2 security, tight security.

1 Introduction

The Fiat-Shamir transformation [43] is a well-known technique that turns any 3-move honest-verifier zero-knowledge proof system (a.k.a. Σ-protocol [36]) into a non-interactive zero-knowledge proof (NIZK) by replacing the verifier's challenge by a hash value of the transcript so far. Bellare and Rogaway [11] showed that this approach is secure if the underlying hash function is modeled as a random oracle. Since then, the Fiat-Shamir heuristic has been used in the design of countless cryptographic schemes, including digital signatures [78] and chosen-ciphertext-secure public-key encryption schemes [44]. In the standard model, however, counter-examples [49] showed that it may fail to guarantee soundness.

Until recently, it was not known to be securely instantiable without random oracles under any standard assumption. This situation drastically changed with the works of Canetti *et al.* [26] and Peikert and Shiehian [76], which imply the existence of Fiat-Shamir-based NIZK proofs for all NP languages under the sole Learning-With-Errors (LWE) assumption [79]. Their results followed a line of research [82,27,25] showing that Fiat-Shamir can provide soundness in the standard model if the underlying hash function is *correlation intractable* (CI). In short, correlation intractability for a relation R captures the infeasibility of finding an x such that $(x, H_k(x)) \in R$ given a random hashing key k. Intuitively, the reason why this property provides soundness is that a cheating prover's first message cannot be hashed into a verifier message admitting an accepting transcript, except with negligible probability.

While [26,76] resolve the challenging problem of realizing NIZK proofs for all NP under standard lattice assumptions, they leave open the question of building more efficient instantiations of Fiat-Shamir for specific languages, such as those arising in the context of chosen-ciphertext security [75,80,44].

In order to instantiate the Naor-Yung paradigm of CCA2-secure encryption [75] in the lattice setting, the only known solution is to proceed via a general NP reduction to graph Hamiltonicity and apply the Σ-protocol of Feige, Lapidot and Shamir [42] with the modifications suggested by Canetti *et al.* [26,30]. In addition, a direct application of [26,30,76] to CCA2 security requires to apply the generic compiler of [39] that turns any NIZK proof system into simulation-sound [80] proofs. Here, we consider the problem of more efficiently instantiating Naor-Yung in the standard model under lattice assumptions. Using correlation intractable hash functions, our goal is to directly construct simulation-sound arguments of plaintext equality without using generic techniques.

1.1 Our Contributions

We describe the most efficient post-quantum realization of the Naor-Yung paradigm so far and its first non-trivial instantiation under lattice assumptions. As an application, we obtain the most efficient public-key encryption scheme providing key-dependent message security under chosen-ciphertext attacks (or KDM-CCA2 security for short) under the standard Learning-With-Errors (LWE) assumption [79]. Our scheme is *not* the result of merely combining generic NIZK techniques [80,39] with the results [26,30,76] on NIZK proofs based on correlation intractable hash functions. In particular, we bypass the use of a Karp reduction to the graph Hamiltonicity language [42,26,30]. Instead, as a key building block, we directly build a simulation-sound NIZK proof system showing that two dual Regev ciphertexts [46] are encryptions of the same plaintext.

As a result of independent interest, we also obtain a multi-theorem NIZK argument system without using the Feige-Lapidot-Shamir (FLS) transformation [42]. Recall that the FLS compiler constructs a multi-theorem NIZK proof system for an NP language from a single-theorem NIZK proof system by using the latter to prove OR statements of the form "either element x is in the language OR some CRS component is in the range of a pseudorandom generator". Unlike FLS, our

multi-theorem NIZK argument avoids the non-black-box use of a PRG. Another advantage is that it provides multi-theorem statistical NIZK in the common *random* string model while proving soundness under the LWE assumption. In contrast, achieving statistical multi-theorem NIZK by applying FLS to [76,30] requires a common reference string sampled from a non-uniform distribution.

We further show that our argument system provides *unbounded* (as opposed to one-time [80]) simulation-soundness (USS) [39], meaning that the adversary remains unable to prove a false statement, even after having seen simulated arguments for polynomially many (possibly false) statements. This makes our argument system suitable to prove KDM-CCA2 security by applying the Naor-Yung technique to the KDM-CPA system of Applebaum, Cash, Peikert, and Sahai (ACPS) [6], which is known to provide key-dependent message security for affine functions. In addition, we provide a variant of our USS argument that can be proved tightly secure, meaning that the reduction's advantage is not affected by the number of simulated proofs obtained by the adversary. The simulation-soundness property is indeed tightly related to the security of the underlying pseudorandom function. By exploiting a result of Lai *et al.* [64], it can be combined with a tightly secure lattice-based PRF so as to instantiate our scheme with a polynomial modulus.

Our first simulation-sound NIZK argument implies a public-key encryption (PKE) scheme providing KDM-CCA2 security under the LWE assumption with polynomial approximation factors. Our second NIZK argument yields an instantiation that enjoys *tight* KDM-CCA2 security. Until recently, this was only possible under an LWE assumption with large approximation factors for lack of a tightly secure low-depth lattice-based PRF based on an LWE assumption with polynomial inverse-error rate. Lai *et al.* [64] recently showed that many tightly secure LWE-based schemes (e.g., [17,67,18]) can actually be obtained using a PRF outside NC1 without going through Barrington's theorem [9]. Their technique [64] applies to our setting and ensure that any (possibly sequential) PRF with a tight security reduction from LWE with polynomial modulus and inverse-error rate allows instantiating the scheme under a similarly standard assumption.

Recall that KDM security is formalized by an experiment where the adversary obtains N public keys. On polynomially many occasions, it sends encryption queries (i, f), for functions $f \in \mathcal{F}$ belonging to some family, and expects to receive an encryption of $f(SK_1, \ldots, SK_N)$ under PK_i. Security requires the adversary to be unable to distinguish the real encryption oracle from an oracle that always returns an encryption of 0. Our KDM-CCA2 construction supports the same function family (namely, affine functions) as the KDM-CPA system it builds on. However, like previous LWE-based realizations [6,4], it can be bootstrapped using Applebaum's technique [5] so as to retain KDM security for arbitrary functions that are computable in a priori bounded polynomial time.

We believe our LWE-based instantiation of Naor-Yung to be of interest beyond KDM security. For example, it makes possible to publicly recognize ciphertexts that correctly decrypt, which is a rare feature among LWE-based schemes and comes in handy in the threshold decryption setting (see, e.g., [44]). It can also

be used to obtain chosen-ciphertext security in settings – such as inner product functional encryption [1,3] or receiver selective-opening security [54] – for which we do not know how to apply the Canetti-Halevi-Katz technique [29].

1.2 Technical Overview

Our starting point is a trapdoor Σ-protocol [26,30] allowing to prove the well-formed of ciphertexts in the KDM-CPA system of Applebaum *et al.* [6]. Namely, it allows proving that a given vector $\mathbf{c} = (\mathbf{u}, u) \in \mathbb{Z}_q^{n+1}$ is of the form $(\mathbf{u}, \mathbf{u}^\top \mathbf{s} + \mu\lfloor q/p \rfloor + \mathsf{noise})$, where $\mu \in \mathbb{Z}_p$ is the message, $\mathbf{s} \in \mathbb{Z}^n$ is the secret key and the public key is $(\mathbf{A}, \mathbf{b} = \mathbf{A}^\top \mathbf{s} + \mathsf{noise}) \in \mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^m$ for some $m = \Omega(n \cdot \log q)$. Recall that a standard Σ-protocol [36,35] is a 3-move protocol with transcripts of the form $(\mathbf{a}, \mathsf{Chall}, \mathbf{z})$ where Chall is the verifier's challenge and messages $\mathbf{a}$ and $\mathbf{z}$ are sent by the prover. In the common reference string model, a trapdoor Σ-protocol [26,30] has the property that, for any statement x outside the language $\mathcal{L}$ and any first message $\mathbf{a}$ sent by the prover, a trapdoor makes it possible to determine the unique challenge Chall for which a valid response $\mathbf{z}$ exists. There is an efficiently computable function $\mathsf{BadChallenge}$ that takes as input a trapdoor τ, a false statement $x \notin \mathcal{L}$, and a first prover message $\mathbf{a}$, and computes the unique Chall such that there exists an accepting transcript $(\mathbf{a}, \mathsf{Chall}, \mathbf{z})$ (that is, there is no accepting transcript of the form $(\mathbf{a}, \mathsf{Chall}', \mathbf{z})$ for any $\mathsf{Chall}' \neq \mathsf{Chall}$).

Our first observation is that, in order to preserve the soundness of Fiat-Shamir, it suffices for a trapdoor Σ-protocol to have a $\mathsf{BadChallenge}$ function that outputs "if there is a bad challenge at all for $\mathbf{a}$, it can only be Chall". Indeed, false positives do not hurt soundness as we only need the CI hash function to sidestep the bad challenge whenever it exists. Based on this observation, we can build a trapdoor Σ-protocol showing that a Regev ciphertext $\mathbf{c} \in \mathbb{Z}_q^{n+1}$ encrypts 0. Letting $\bar{\mathbf{A}} = [\mathbf{A}^\top \mid \mathbf{b}]^\top \in \mathbb{Z}_q^{(n+1) \times m}$, this can be done using by showing knowledge of a short $\mathbf{r} \in \mathbb{Z}^m$ such that $\mathbf{c} = \bar{\mathbf{A}} \cdot \mathbf{r}$. In Σ-protocols like [70,71], the verifier accepts transcripts $(\mathbf{a}, \mathsf{Chall}, \mathbf{z})$ such that $\mathbf{a} + \mathsf{Chall} \cdot \mathbf{c} = \bar{\mathbf{A}} \cdot \mathbf{z}$ if $\mathbf{z} \in \mathbb{Z}^m$ is short enough. Since the right-hand side member of the verification equation is an encryption of 0, the $\mathsf{BadChallenge}$ function can use the decryption key $\mathbf{s}$ to infer that no valid response exists for the challenge $\mathsf{Chall} = b$ when $\mathbf{a} + b \cdot \mathbf{c}$ does not decrypt to 0.

The next step is to argue that $\mathbf{c}$ encrypts an arbitrary $\mu \in \mathbb{Z}_p$. To this end, we exploit the fact the KDM-CPA scheme of [6] uses a square modulus $q = p^2$ when we compute part of the response $z_\mu = r_u + \mathsf{Chall} \cdot \mu \bmod p$ over $\mathbb{Z}_p$, while using a uniform mask $r_u \in \mathbb{Z}_p$ to hide $\mu \in \mathbb{Z}_p$ as in standard Schnorr-like protocols [81]. Now, the $\mathsf{BadChallenge}$ function can output $\mathsf{Chall} = 1 - b$ if it detects that $\mathbf{a} + b \cdot \mathbf{c}$ is not of the form $(\mathbf{u}, \mathbf{u}^\top \mathbf{s} + z_\mu \cdot p + \mathsf{noise})$, for some $z_\mu \in \mathbb{Z}_p$. Indeed, this rules out the existence of a short enough $\mathbf{z} \in \mathbb{Z}^m$ such that $\mathbf{a} + b \cdot \mathbf{c} = \bar{\mathbf{A}} \cdot \mathbf{z} + z_\mu \cdot [\mathbf{0}^{n\top} | p]^\top$ with $z_\mu \in \mathbb{Z}_p$. The above technique extends into a trapdoor Σ-protocol for proving plaintext equalities in the ACPS cryptosystem [6]. Our instantiation of Naor-Yung thus requires to work with LWE over a composite modulus q and we leave it as an open problem to extend it to prime moduli.

The main difficulty, however, is to turn the aforementioned trapdoor Σ-protocol into a non-interactive proof system with unbounded simulation-soundness.

This problem is non-trivial since the Canetti *et al.* protocol [26,30] is not known to satisfy this security notion.[6] The NIZK simulator of [26,30] generates simulated proofs by "programming" the CI hash function from which the verifier's challenge is derived. In the context of unbounded simulation-soundness [80,39], we cannot proceed in the same way since the simulator would have to program the hash function for each simulated proof (and thus for each challenge ciphertext in the proof of KDM-CCA2 security). Since the number of simulated proofs is not a priori bounded, it is not clear how to do that using a hashing key of length independent of the number of adversarial queries.

Our solution to this problem is inspired by the modification introduced by Canetti *et al.* [30,26] in the Feige-Lapidot-Shamir protocol [42]. In [30, Section 5.2], the first prover message $\mathbf{a}$ is computed using a lossy encryption scheme [10] instead of an ordinary commitment. Recall that, depending on the distribution of the public key PK, a lossy encryption scheme behaves either as an extractable non-interactive commitment or a statistically-hiding commitment. The extractable mode is used to prove the soundness property (by using the secret key SK corresponding to PK to compute the $\mathsf{BadChallenge}$ function) while the statistically hiding mode allows proving zero-knowledge. Our unbounded simulation-sound proof system exploits the observation made by Bellare *et al.* [10] that specific lossy encryption schemes admit an efficient opening algorithm. Namely, ciphertexts encrypted under a lossy public key can be equivocated in the same way as a trapdoor commitment using the lossy secret key SK. This suggests that, if the protocol of Canetti *et al.* [30,26] is instantiated using a lossy encryption scheme with efficient opening, we can use a strategy introduced by Damgård [38] to simulate NIZK proofs without programming the CI hash function. Namely, the simulator can generate the first prover message as a lossy encryption of 0. When receiving the verifier's challenge Chall, it can run the HVZK simulator to obtain $(\mathbf{a}, \mathbf{z})$ before using the lossy secret key SK to explain the lossy ciphertext as an encryption of the simulated $\mathbf{a}$. By doing this, we also obtain a multi-theorem NIZK argument without using the FLS transformation [42] and without using any primitive in a non-black-box way. The language of the underlying trapdoor Σ-protocol is exactly the same as that of the multi-theorem NIZK argument, so that, if the former is efficient, so is the latter.

However, standard lossy encryption schemes with efficient opening do not suffice to prove unbounded simulation-soundness: We do not only need to equivocate lossy ciphertexts in all simulated proofs, but we should also make sure that the adversary's fake proof is generated for a statistically binding (and even extractable) commitment. For this reason, we rely on a lossy encryption flavor, called $\mathcal{R}$-lossy encryption by Boyle *et al.* [19], where a tag determines whether a ciphertext is lossy or injective. The public key is generated for a computationally hidden initialization value $K \in \mathcal{K}$ and ciphertexts are encrypted under a tag $t \in \mathcal{T}$. If $\mathcal{R} \subset \mathcal{K} \times \mathcal{T}$ is a binary relation, the syntax of $\mathcal{R}$-lossy encryption [19] is

[6] It can be generically achieved using NIZK for general NP relations [39] but our goal is to obtain a more efficient solution than generic NIZK techniques. In fact, even one-time simulation-soundness is not proven in [26,30]

that a ciphertext encrypted for a tag $t \in \mathcal{T}$ is injective if $\mathcal{R}(K,t) = 1$ and lossy otherwise. For our purposes, we need to enrich the syntax of $\mathcal{R}$-lossy encryption in two aspects. First, we require lossy ciphertexts to be efficiently equivocable (i.e., the secret key SK should make it possible to find random coins that explain a lossy ciphertext as an encryption of any target plaintext). Second, in order to simplify the description of our NIZK simulator, we need the syntax to support lossy/injective tags *and* lossy/injective keys. When the public key PK is lossy, all ciphertexts are lossy, no matter which tag is used to encrypt. In contrast, injective public keys lead to injective ciphertexts whenever $\mathcal{R}(K,t) = 1$. Our NIZK simulator actually uses lossy public keys while injective keys only show up in the proof of simulation-soundness.

We then construct an $\mathcal{R}$-lossy encryption scheme for the bit-matching relation (i.e., $\mathcal{R}_{\mathsf{BM}}(K,t) = 1$ if and only if K and t agree in all positions where K does not contain a "don't care entry") under the LWE assumption. The scheme can be viewed as a combination of the primal Regev cryptosystem [79] – which is known [77] to be a lossy PKE scheme and is easily seen to support efficient openings as defined in [10] – with the lattice trapdoors of Micciancio and Peikert [74]. An injective public key consists of a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ with short vectors in its row space. In order to encrypt $\boldsymbol{\mu} \in \{0,1\}^{n_0}$ under a tag t, we sample a short Gaussian $\mathbf{r} \in \mathbb{Z}^{2m}$ and compute $\mathbf{c} = [\mathbf{A} \mid \mathbf{A} \cdot \mathbf{R}_t + (1 - \mathcal{R}(K,t)) \cdot \mathbf{G}] \cdot \mathbf{r} + [\mathbf{0} \mid \boldsymbol{\mu} \cdot \lfloor q/2 \rfloor]^\top$, for some small-norm $\mathbf{R}_t \in \mathbb{Z}^{m \times m}$, where $\mathbf{G} \in \mathbb{Z}_q^{n \times m}$ is the gadget matrix of [74]. In each lossy tag, we have $\mathcal{R}(K,t) = 0$, in which case the matrix $\mathbf{R}_t$ can be used as a trapdoor (using the techniques of [2,74]) to sample a Gaussian $\mathbf{r} \in \mathbb{Z}^{2m}$ that explains $\mathbf{c}$ as an encryption of any arbitrary $\boldsymbol{\mu} \in \{0,1\}^{n_0}$. In injective tags, we have $\mathcal{R}(K,t) = 1$, so that the gadget matrix vanishes from the matrix $\mathbf{A}_t = [\mathbf{A} \mid \mathbf{A} \cdot \mathbf{R}_t + (1 - \mathcal{R}(K,t)) \cdot \mathbf{G}]$. Since $\mathbf{A}$ has short vectors in its row space, so does $\mathbf{A}_t$ and we can thus use these short vectors to recover $\boldsymbol{\mu}$ from $\mathbf{c}$ exactly as in the primal Regev cryptosystem. When the public key PK is lossy, the matrix $\mathbf{A}$ is replaced by a statistically uniform matrix over $\mathbb{Z}_q^{n \times m}$. We can then use a trapdoor for $\Lambda^\perp(\mathbf{A})$ to equivocate lossy ciphertexts for any arbitrary tag.

Our USS argument system uses our $\mathcal{R}$-lossy encryption scheme – with the standard trick of using the verification key of a one-time signature as a tag – to compute the first prover message $\mathbf{a}$ by encrypting the first message $\mathbf{a}'$ of a basic trapdoor Σ-protocol. In the security proof, we have a noticeable probability that: (i) For all adversarially-chosen statements, proofs can be simulated by equivocating lossy ciphertexts; (ii) When the adversary comes up with a proof of its own, the underlying commitment is an injective ciphertext. If these conditions are fulfilled, we can annihilate the adversary's chance of proving a false statement by using a hash function which is statistically CI for the relation that evaluates the $\mathsf{BadChallenge}$ function on input of the decryption of an $\mathcal{R}$-lossy ciphertext.

At a high-level, our simulation-sound proof system bears similarities with interactive proof systems described by MacKenzie and Yang [72]. Our extension of $\mathcal{R}$-lossy encryption resembles their notion of simulation-sound trapdoor commitments. The difference is that, while [72] only requires commitments to be computationally binding for tags that have never been equivocated, we need

adversarially-chosen tags to be extractable.

Our first USS argument system does not provide tight security because it relies on admissible hash functions [14] to partition the tag space of the $\mathcal{R}$-lossy PKE scheme into two disjoint subspaces (which contain equivocable and extractable tags, respectively). In order to obtain tight simulation-soundness, our second USS argument partitions the tag space of an $\mathcal{R}$-lossy PKE scheme using a pseudorandom function instead of an admissible hash function. For this purpose, we build an $\mathcal{R}$-lossy PKE scheme for a relation $\mathcal{R}_{\mathsf{PRF}}$ induced by a PRF family. Analogously to [55], we consider tags $t = (t_c, t_a)$ consisting of an auxiliary component t_a (which can be an arbitrary string) and core component t_c. The PRF-induced relation $\mathcal{R}_{\mathsf{PRF}}$ is then defined as $\mathcal{R}_{\mathsf{PRF}}(K, (t_c, t_a)) = 1$ if and only if $t_c \neq \mathsf{PRF}_K(t_a)$, where K is the PRF secret key. Our $\mathcal{R}_{\mathsf{PRF}}$-lossy PKE then proceeds as in [67] and uses a public key containing Gentry-Sahai-Waters encryptions [47] $\mathbf{A}_i = \mathbf{A} \cdot \mathbf{R}_i + k_i \cdot \mathbf{G}$ of the bits of K. To encrypt $\boldsymbol{\mu} \in \{0,1\}^{n_0}$ under a tag $t = (t_c, t_a)$, the encryptor first homomorphically computes $\mathbf{A}_{F,t} = \mathbf{A} \cdot \mathbf{R}_t + (1 - \mathcal{R}_{\mathsf{PRF}}(K, t)) \cdot \mathbf{G}$ before sampling a short Gaussian $\mathbf{r} \in \mathbb{Z}^{2m}$ and computing $\mathbf{c} = [\mathbf{A} \mid \mathbf{A}_{F,t}] \cdot \mathbf{r} + [\mathbf{0} \mid \boldsymbol{\mu} \cdot \lfloor q/2 \rfloor]^\top$. In the proof of simulation-soundness, the reduction simulates all arguments by "adaptively programming" all tags $t = (\mathsf{PRF}_K(t_a), t_a)$ to ensure equivocability. At the same time, the adversary can only output an argument on an extractable tag $t^\star = (t_c^\star, t_a^\star)$, where $\mathcal{R}_{\mathsf{PRF}}(K, t^\star) = 1$, unless it can predict $t_c^\star = \mathsf{PRF}_K(t_a^\star)$.

1.3 Related Work

Fiat-Shamir in the Standard Model. The Fiat-Shamir methodology was shown [49] not to be sound in the standard model in general. Known negative results (see [49,12] and references therein) nevertheless left open the existence of secure instantiations of the paradigm when specific protocols are transformed using concrete hash functions. Of particular interest is the notion of *correlation intractable* hash function [28], which rules out specific relations between an input and its hash value. It was actually shown [52] that correlation intractability for all sparse relations[7] suffices to ensure soundness as long as the underlying protocol is statistically sound. A recent line of work [82,27,58,25] focused on the design of correlation intractable hash functions leading to sound instantiation of Fiat–Shamir in the standard model. Canetti *et al.* [26] showed that it is actually sufficient to obtain correlation intractable hash families for *efficiently searchable* relations (i.e., where each x has at most one corresponding y, which is computable within some polynomial time bound). This opened the way to CI hash candidates based on more established assumptions like the circular security of fully homomorphic encryption (FHE) schemes [30]. Peikert and Shiehian [76] recently gave an elegant FHE-based solution relying on the hardness of the LWE problem [79] with polynomial approximation factors. While specific to the Gentry-Sahai-Waters (GSW) FHE [47], their construction does not require any non-standard circular security assumption. Together with the techniques of

[7] A relation $R \subset \mathcal{X} \times \mathcal{Y}$ is sparse if, for a given $x \in \mathcal{X}$, the fraction of $y \in \mathcal{Y}$ for which $(x, y) \in R$ is negligible.

[30,26], it implies NIZK for all NP languages.

In [30,26], Canetti *et al.* showed that, besides the language of Hamiltonian graphs considered in [42], trapdoor Σ-protocols also exist for other languages like that of quadratic residues modulo a composite integer [48]. Using the CI hash function of [76], they thus obtained a NIZK proof for the Quadratic Residuosity language under the LWE assumption. Choudhuri *et al.* [32] showed that the hash families of [26] make the transformation sound for the sumcheck protocol.

Multi-Theorem NIZK. Several multi-theorem NIZK constructions are available in the literature (see, e.g., [42,40,31,50]). Under the LWE assumption, all solutions so far either rely on the FLS transformation [34,76] – thus incurring proofs of OR statements via non-black-box techniques – or restrict themselves to the designated verifier setting [34,68]. While the meta-proof approach of De Santis and Yung [40] provides an alternative to FLS, it makes non-black-box use of a single-theorem proof system for an NP-complete language. Our construction uses a single-theorem argument for the same language as the one for which we need a multi-theorem argument. Hence, if the former is efficient, so is the latter.

KDM Security. This security notion was first formalized by Black, Rogaway and Shrimpton [13] and motivated by applications in anonymous credentials [23] or in disk encryption (e.g., in the BitLocker encryption utility [16]), where the key may be stored on the disk being encrypted. The first examples of KDM-secure secret-key encryption were given by Black *et al.* [13] in the random oracle model.

In the standard model, Boneh *et al.* [16] designed the first public-key scheme with provable KDM-CPA security w.r.t. all affine functions under the decisional Diffie-Hellman (DDH) assumption. Applebaum *et al.* [6] showed that a variant of Regev's system [79] is KDM-secure for all affine functions under the LWE assumption. They also gave a secret-key construction based on the hardness of the Learning Parity with Noise (LPN) problem for which Döttling gave a public-key variant [41]. Under the Quadratic Residuosity (QR) and Decisional Composite Residuosity (DCR) assumptions, Brakerski and Goldwasser [20] gave alternative constructions that additionally provide security under key leakage. Alperin-Sheriff and Peikert [4] showed that a variant of the identity-based encryption scheme of Agrawal *et al.* [2] provides KDM security for a bounded number of challenge ciphertexts.

Brakerski *et al.* [21] and Barak *et al.* [8] came up with different techniques to prove KDM security for richer function families. Malkin *et al.* [73] suggested a much more efficient scheme with ciphertexts of $O(d)$ group elements for function families containing degree d polynomials. Applebaum [5] put forth a generic technique that turns any PKE scheme with KDM security for projection functions – where each output bit only depends on a single input bit – into a scheme providing KDM security for any circuit of a priori bounded polynomial size.

KDM-CCA Security. The first PKE scheme with KDM-CCA2 security in the standard model appeared in the work of Camenisch, Chandran, and Shoup [24]. They gave a generic construction based on the Naor-Yung paradigm that combines a KDM-CPA system, a standard CPA-secure encryption scheme, and a simulation-sound NIZK proof system. For their purposes, they crucially need

unbounded simulation-soundness since the KDM setting inherently involves many challenge ciphertexts and single-challenge security is not known to imply multi-challenge security. They instantiated their construction using the DDH-based KDM-CPA system of Boneh *et al* [16] and Groth-Sahai proofs [51]. Our scheme is an instantiation of the generic construction of [24] in the lattice setting, where we cannot simply use Groth-Sahai proofs. Hofheinz [56] subsequently obtained chosen-ciphertext circular security (i.e., for selection functions where $f(SK_1, \ldots, SK_N) = SK_i$ for some $i \in [N]$) with shorter ciphertexts.

A first attempt of KDM-CCA security without pairings was made by Lu *et al.* [69]. Han *et al.* [53] identified a bug in [69] and gave a patch using the same methodology. They obtained KDM-CCA security for bounded-degree polynomial functions under the DDH and DCR assumptions. Kitigawa and Tanaka [63] described a framework for the design of KDM-CCA systems under a single number theoretic assumption (i.e., DDH, QR, or DCR). Their results were extended by Kitigawa *et al.* [62] so as to prove tight KDM-CCA2 security under the DCR assumption. Since the framework of [63] relies on hash proof systems [37], it is not known to provide LWE-based realizations (indeed, hash proof systems do not readily enable chosen-ciphertext security from LWE), let alone with tight security. To our knowledge, our scheme is thus the first explicit solution with tight KDM-CCA2 security under the LWE assumption. Before [62], the only pathway to tight KDM-CCA security was to instantiate the construction of Camenisch *et al.* [24] using a tightly secure USS proof/argument (e.g., [57]), which tends to incur very large ciphertexts. Our system also follows this approach with the difference that ciphertexts are not much longer than in its non-tight variant.

Kitigawa and Matsuda [61] generically obtained KDM-CCA security for bounded-size circuits from any system providing KDM-CPA security for projection functions. While their result shows the equivalence between KDM-CPA and KDM-CCA security, our scheme is conceptually simpler and significantly more efficient than an LWE-based instantiation of the construction in [61]. In particular, such an instantiation requires both garbling schemes and $\Omega(\lambda)$ designated-verifier proofs of plaintext equalities with negligible soundness error. While these proofs seem realizable by applying the techniques of [68] to specific Σ-protocols, each of them would cost $\Omega(\lambda^2)$ public-key encryptions. Our scheme is much simpler and only requires one argument of plaintext equality, thus compressing ciphertexts by a factor at least $\Omega(\lambda)$.

1.4 Organization

Section 2 first recalls the the building blocks of our constructions. Our first simulation-sound argument is presented in Section 3 together with the underlying $\mathcal{R}$-lossy PKE scheme. Its tightly secure variant is described in Section 4. In Section 5, we give a trapdoor Σ-protocol allowing to apply the Naor-Yung transformation to the ACPS cryptosystem. The resulting (tightly secure) KDM-CCA2 system is then detailed in the full version of the paper [65]. As written, our security proof only shows tightness in the number of challenge ciphertexts, but not in the

number of users. In the full version of the paper, we also explain how to also obtain tightness w.r.t. the number of users.

2 Background

We recall the main tools involved in our constructions. Additional standard tools, such as NIZK proofs, are defined in the full version of the paper.

2.1 Lattices

For any $q \geq 2$, $\mathbb{Z}_q$ denotes the ring of integers with addition and multiplication modulo q. If $\mathbf{x} \in \mathbb{R}^n$ is a vector, $\|\mathbf{x}\| = \sqrt{\sum_{i=1}^{n} x_i^2}$ denotes its Euclidean norm and $\|\mathbf{x}\|_\infty = \max_i |x_i|$ its infinity norm. If $\mathbf{M}$ is a matrix over $\mathbb{R}$, then $\|\mathbf{M}\| := \sup_{\mathbf{x}\neq 0} \frac{\|\mathbf{M}\mathbf{x}\|}{\|\mathbf{x}\|}$ and $\|\mathbf{M}\|_\infty := \sup_{\mathbf{x}\neq 0} \frac{\|\mathbf{M}\mathbf{x}\|_\infty}{\|\mathbf{x}\|_\infty}$ denote its induced norms. For a finite set S, $U(S)$ stands for the uniform distribution over S. If X and Y are distributions over the same domain, $\Delta(X, Y)$ denotes their statistical distance.

Let $\mathbf{\Sigma} \in \mathbb{R}^{n\times n}$ be a symmetric positive-definite matrix, and $\mathbf{c} \in \mathbb{R}^n$. We define the Gaussian function on $\mathbb{R}^n$ by $\rho_{\mathbf{\Sigma},\mathbf{c}}(\mathbf{x}) = \exp(-\pi(\mathbf{x}-\mathbf{c})^\top \mathbf{\Sigma}^{-1}(\mathbf{x}-\mathbf{c}))$ and if $\mathbf{\Sigma} = \sigma^2 \cdot \mathbf{I}_n$ and $\mathbf{c} = \mathbf{0}$ we denote it by ρ_σ. For an n dimensional lattice $\Lambda \subset \mathbb{R}^n$ and for any lattice vector $\mathbf{x} \in \Lambda$ the discrete Gaussian is defined by $\rho_{\Lambda,\mathbf{\Sigma},\mathbf{c}}(\mathbf{x}) = \frac{\rho_{\mathbf{\Sigma},\mathbf{c}}}{\rho_{\mathbf{\Sigma},\mathbf{c}}(\Lambda)}$.

For an n-dimensional lattice Λ, we define $\eta_\varepsilon(\Lambda)$ as the smallest $r > 0$ such that $\rho_{1/r}(\widehat{\Lambda} \setminus \mathbf{0}) \leq \varepsilon$ with $\widehat{\Lambda}$ denoting the dual of Λ, for any $\varepsilon \in (0, 1)$.

For a matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$, we define $\Lambda^\perp(\mathbf{A}) = \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{A}\cdot\mathbf{x} = \mathbf{0} \bmod q\}$ and $\Lambda(\mathbf{A}) = \mathbf{A}^\top \cdot \mathbb{Z}^n + q\mathbb{Z}^m$. For an arbitrary vector $\mathbf{u} \in \mathbb{Z}_q^n$, we also define the shifted lattice $\Lambda^{\mathbf{u}}(\mathbf{A}) = \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{A}\cdot\mathbf{x} = \mathbf{u} \bmod q\}$.

Definition 2.1 (LWE). *Let $m \geq n \geq 1$, $q \geq 2$ and $\alpha \in (0, 1)$ be functions of a security parameter λ. The* LWE *problem consists in distinguishing between the distributions $(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e})$ and $U(\mathbb{Z}_q^{m\times n} \times \mathbb{Z}_q^m)$, where $\mathbf{A} \sim U(\mathbb{Z}_q^{m\times n})$, $\mathbf{s} \sim U(\mathbb{Z}_q^n)$ and $\mathbf{e} \sim D_{\mathbb{Z}^m,\alpha q}$. For an algorithm $\mathcal{A} : \mathbb{Z}_q^{m\times n} \times \mathbb{Z}_q^m \to \{0, 1\}$, we define:*

$$\mathbf{Adv}^{\mathsf{LWE}}_{q,m,n,\alpha}(\lambda) = |\Pr[\mathcal{A}(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}) = 1] - \Pr[\mathcal{A}(\mathbf{A}, \mathbf{u}) = 1| \ ,$$

where the probabilities are over $\mathbf{A} \sim U(\mathbb{Z}_q^{m\times n})$, $\mathbf{s} \sim U(\mathbb{Z}_q^n)$, $\mathbf{u} \sim U(\mathbb{Z}_q^m)$ and $\mathbf{e} \sim D_{\mathbb{Z}^m,\alpha q}$ and the internal randomness of $\mathcal{A}$. We say that $\mathsf{LWE}_{q,m,n,\alpha}$ is hard if, for any PPT algorithm $\mathcal{A}$, the advantage $\mathbf{Adv}^{\mathsf{LWE}}_{q,m,n,\alpha}(\mathcal{A})$ is negligible.

Micciancio and Peikert [74] described a trapdoor mechanism for LWE. Their technique uses a "gadget" matrix $\mathbf{G} \in \mathbb{Z}_q^{n\times w}$, with $w = n \log q$, for which anyone can publicly sample short vectors $\mathbf{x} \in \mathbb{Z}^w$ such that $\mathbf{G}\cdot\mathbf{x} = \mathbf{0}$.

Lemma 2.2 (**[74, Section 5]**). *Assume that $\bar{m} \geq n \log q + O(\lambda)$ and $m = \bar{m} + n\lceil \log q \rceil$. There exists a probabilistic polynomial time (PPT) algorithm* GenTrap *that takes as inputs matrices $\bar{\mathbf{A}} \in \mathbb{Z}_q^{n\times\bar{m}}$, $\mathbf{H} \in \mathbb{Z}_q^{n\times n}$ and outputs*

matrices $\mathbf{R} \in \{-1,1\}^{\bar{m} \times n \cdot \lceil \log q \rceil}$ *and* $\mathbf{A} = \left[\bar{\mathbf{A}} \mid \bar{\mathbf{A}}\mathbf{R} + \mathbf{H} \cdot \mathbf{G}\right] \in \mathbb{Z}_q^{n \times m}$ *such that if* $\mathbf{H} \in \mathbb{Z}_q^{n \times n}$ *is invertible, then* $\mathbf{R}$ *is a* $\mathbf{G}$*-trapdoor for* $\mathbf{A}$ *with tag* $\mathbf{H}$*; and if* $\mathbf{H} = \mathbf{0}$*, then* $\mathbf{R}$ *is a punctured trapdoor.*

Further, in case of a $\mathbf{G}$*-trapdoor, one can efficiently compute from* $\mathbf{A}, \mathbf{R}$ *and* $\mathbf{H}$ *a basis* $(\mathbf{t}_i)_{i \le m}$ *of* $\Lambda^{\perp}(\mathbf{A})$ *such that* $\max_i \|\mathbf{t}_i\| \le O(m^{3/2})$.

Lemma 2.3 ([46, Theorem 4.1]). *There is a PPT algorithm that, given a basis* $\mathbf{B}$ *of an* n*-dimensional* $\Lambda = \Lambda(\mathbf{B})$*, a parameter* $s > \|\tilde{\mathbf{B}}\| \cdot \omega(\sqrt{\log n})$*, and a center* $\mathbf{c} \in \mathbb{R}^n$*, outputs a sample from a distribution statistically close to* $D_{\Lambda,s,\mathbf{c}}$.

2.2 Correlation Intractable Hash Functions

We consider unique-output searchable binary relations [26]. These are binary relations such that, for every x, there is at most one y such that $R(x,y) = 1$ and y is efficiently computable from x.

Definition 2.4. *A relation* $R \subseteq \mathcal{X} \times \mathcal{Y}$ *is* **searchable** *in time* T *if there exists a function* $f : \mathcal{X} \to \mathcal{Y}$ *which is computable in time* T *and such that, if there exists* y *such that* $(x,y) \in R$*, then* $f(x) = y$.

Letting $\lambda \in \mathbb{N}$ denote a security parameter, a hash family with input length $n(\lambda)$ and output length $m(\lambda)$ is a collection $\mathcal{H} = \{h_\lambda : \{0,1\}^{s(\lambda)} \times \{0,1\}^{n(\lambda)} \to \{0,1\}^{m(\lambda)}\}$ of keyed hash functions implemented by efficient algorithms $(\mathsf{Gen}, \mathsf{Hash})$, where $\mathsf{Gen}(1^\lambda)$ outputs a key $k \in \{0,1\}^{s(\lambda)}$ and $\mathsf{Hash}(k,x)$ computes a hash value $h_\lambda(k,x) \in \{0,1\}^{m(\lambda)}$.

Definition 2.5. *For a relation ensemble* $\{R_\lambda \subseteq \{0,1\}^{n(\lambda)} \times \{0,1\}^{m(\lambda)}\}$*, a hash function family* $\mathcal{H} = \{h_\lambda : \{0,1\}^{s(\lambda)} \times \{0,1\}^{n(\lambda)} \to \{0,1\}^{m(\lambda)}\}$ *is* R**-correlation intractable** *if, for any probabilistic polynomial time (PPT) adversary* $\mathcal{A}$*, we have* $\Pr\left[k \leftarrow \mathsf{Gen}(1^\lambda)),\ x \leftarrow \mathcal{A}(k) : (x, h_\lambda(k,x)) \in R\right] = \mathsf{negl}(\lambda)$.

Peikert and Shiehian [76] described a correlation-intractable hash family for any searchable relation (in the sense of Definition 2.4) defined by functions f of bounded depth. Their construction relies on the standard Short Integer Solution assumption (which is implied by LWE) with polynomial approximation factors.

2.3 Admissible Hash Functions

Admissible hash functions were introduced in [14] as a combinatorial tool for partitioning-based security proofs. A simplified definition was given in [45].

Definition 2.6 ([14,45]). *Let* $\ell(\lambda), L(\lambda) \in \mathbb{N}$ *be functions of* $\lambda \in \mathbb{N}$*. Let an efficiently computable function* $\mathsf{AHF} : \{0,1\}^\ell \to \{0,1\}^L$*. For each* $K \in \{0,1,\perp\}^L$*, let the partitioning function* $F_{\mathsf{ADH}}(K,\cdot) : \{0,1\}^\ell \to \{0,1\}$ *such that*

$$F_{\mathsf{ADH}}(K,X) := \begin{cases} 0 & \text{if} \quad \forall i \in [L] \quad (\mathsf{AHF}(X)_i = K_i) \ \vee \ (K_i = \perp) \\ 1 & \text{otherwise} \end{cases}$$

We say that AHF *is an* **admissible hash function** *if there exists an efficient algorithm* $\mathsf{AdmSmp}(1^\lambda, Q, \delta)$ *that takes as input* $Q \in \mathsf{poly}(\lambda)$ *and a non-negligible* $\delta(\lambda) \in (0,1]$ *and outputs a key* $K \in \{0,1,\bot\}^L$ *such that, for all* $X^{(1)}, \ldots, X^{(Q)}, X^\star \in \{0,1\}^\ell$ *such that* $X^\star \notin \{X^{(1)}, \ldots, X^{(Q)}\}$*, we have*

$$\Pr_K \left[F_{\mathsf{ADH}}(K, X^{(1)}) = \cdots = F_{\mathsf{ADH}}(K, X^{(Q)}) = 1 \ \wedge \ F_{\mathsf{ADH}}(K, X^\star) = 0 \right] \geq \delta(Q(\lambda)) \ .$$

It is known that admissible hash functions exist for $\ell, L = \Theta(\lambda)$.

Theorem 2.7 ([59, Theorem 1]). *Let* $(C_\ell)_{\ell \in \mathbb{N}}$ *be a family of codes* $C_\ell : \{0,1\}^\ell \to \{0,1\}^L$ *with minimal distance* $c \cdot L$ *for some constant* $c \in (0, 1/2)$*. Then,* $(C_\ell)_{\ell \in \mathbb{N}}$ *is a family of admissible hash functions. Furthermore,* $\mathsf{AdmSmp}(1^\lambda, Q, \delta)$ *outputs a key* $K \in \{0,1,\bot\}^L$ *for which* $\eta = O(\log \lambda)$ *components are not* $\bot$ *and* $\delta(Q(\lambda))$ *is a non-negligible function of* λ*.*

Jager proved [59] Theorem 2.7 for *balanced* admissible hash functions, which provide both a lower bound and a close upper bound for the probability in Definition 2.6. Here, we only need the standard definition of admissible hash functions since we use them in a game where the adversary aims at outputting a hard-to-compute result (instead of breaking an indistinguishability property). However, the result of Theorem 2.7 applies to standard admissible hash functions.

2.4 Trapdoor Σ-protocols

Canetti *et al.* [30] considered a definition of Σ-protocols that slightly differs from the usual formulation [36,35].

Definition 2.8 (Adapted from [30,7]). *Let a language* $\mathcal{L} = (\mathcal{L}_{\mathsf{zk}}, \mathcal{L}_{\mathsf{sound}})$ *associated with two NP relations* $R_{\mathsf{zk}}, R_{\mathsf{sound}}$*. A 3-move interactive proof system* $\Pi = (\mathsf{Gen}_{\mathsf{par}}, \mathsf{Gen}_{\mathcal{L}}, \mathsf{P}, \mathsf{V})$ *in the common reference string model is a Gap* Σ*-protocol for* $\mathcal{L}$ *if it satisfies the following conditions:*

- **3-Move Form:** *The prover and the verifier both take as input* $\mathsf{crs} = (\mathsf{par}, \mathsf{crs}_{\mathcal{L}})$*, with* $\mathsf{par} \leftarrow \mathsf{Gen}_{\mathsf{par}}(1^\lambda)$ *and* $\mathsf{crs}_{\mathcal{L}} \leftarrow \mathsf{Gen}_{\mathcal{L}}(\mathsf{par}, \mathcal{L})$*, and a statement* x *and proceed as follows: (i)* P *takes in* $w \in R_{\mathsf{zk}}(x)$*, computes* $(\mathbf{a}, st) \leftarrow \mathsf{P}(\mathsf{crs}, x, w)$ *and sends* $\mathbf{a}$ *to the verifier; (ii)* V *sends back a random challenge* Chall *from the challenge space* $\mathcal{C}$*; (iii)* P *finally sends a response* $\mathbf{z} = \mathsf{P}(\mathsf{crs}, x, w, \mathbf{a}, \mathsf{Chall}, st)$ *to* V*; (iv) On input of* $(\mathbf{a}, \mathsf{Chall}, \mathbf{z})$*,* V *outputs* 1 *or* 0*.*
- **Completeness:** *If* $(x, w) \in R_{\mathsf{zk}}$ *and* P *honestly computes* $(\mathbf{a}, \mathbf{z})$ *for a challenge* Chall*,* $\mathsf{V}(\mathsf{crs}, x, (\mathbf{a}, \mathsf{Chall}, \mathbf{z}))$ *outputs* 1 *with probability* $1 - \mathsf{negl}(\lambda)$*.*
- **Special zero-knowledge:** *There is a PPT simulator* ZKSim *that, on input of* crs*,* $x \in \mathcal{L}_{\mathsf{zk}}$ *and a challenge* $\mathsf{Chall} \in \mathcal{C}$*, outputs* $(\mathbf{a}, \mathbf{z}) \leftarrow \mathsf{ZKSim}(\mathsf{crs}, x, \mathsf{Chall})$ *such that* $(\mathbf{a}, \mathsf{Chall}, \mathbf{z})$ *is computationally indistinguishable from a real transcript with challenge* Chall *(for* $w \in R_{zk}(x)$*).*
- **Special soundness:** *For any CRS* $\mathsf{crs} = (\mathsf{par}, \mathsf{crs}_{\mathcal{L}})$ *obtained as* $\mathsf{par} \leftarrow \mathsf{Gen}_{\mathsf{par}}(1^\lambda)$*,* $\mathsf{crs}_{\mathcal{L}} \leftarrow \mathsf{Gen}_{\mathcal{L}}(\mathsf{par}, \mathcal{L})$*, any* $x \notin \mathcal{L}_{\mathsf{sound}}$*, and any first message* $\mathbf{a}$ *sent by* P*, there is at most one challenge* $\mathsf{Chall} = f(\mathsf{crs}, x, \mathbf{a})$ *for which an*

accepting transcript $(\mathsf{crs}, x, \mathbf{a}, \mathsf{Chall}, \mathbf{z})$ *exists for some third message* $\mathbf{z}$*. The function* f *is called the "bad challenge function" of* Π*. That is, if* $x \notin \mathcal{L}_{\mathsf{sound}}$ *and the challenge differs from the bad challenge, the verifier never accepts.*

Definition 2.8 is taken from [30,7] and relaxes the standard special soundness property in that extractability is not required. Instead, it considers a bad challenge function f, which may not be efficiently computable. Canetti *et al.* [30] define *trapdoor* Σ-protocols as Σ-protocols where the bad challenge function is efficiently computable using a trapdoor. They also define instance-dependent trapdoor Σ-protocol where the trapdoor τ_Σ should be generated as a function of some instance $x \notin \mathcal{L}_{\mathsf{sound}}$. Here, we use a definition where x need not be known in advance (which is not possible in applications to chosen-ciphertext security, where x is determined by a decryption query) and the trapdoor does not depend on a specific x. However, the common reference string and the trapdoor may depend on the language (which is determined by the public key in our application).

The common reference string $\mathsf{crs} = (\mathsf{par}, \mathsf{crs}_{\mathcal{L}})$ consists of a fixed part par and a language-dependent part $\mathsf{crs}_{\mathcal{L}}$ which is generated as a function of par and a language parameter $\mathcal{L} = (\mathcal{L}_{\mathsf{zk}}, \mathcal{L}_{\mathsf{sound}})$.

Definition 2.9 (Adapted from [30]). *A* Σ*-protocol* $\Pi = (\mathsf{Gen}_{\mathsf{par}}, \mathsf{Gen}_{\mathcal{L}}, \mathsf{P}, \mathsf{V})$ *with bad challenge function* f *for a trapdoor language* $\mathcal{L} = (\mathcal{L}_{\mathsf{zk}}, \mathcal{L}_{\mathsf{sound}})$ *is a* **trapdoor** Σ**-protocol** *if it satisfies the properties of Definition 2.8 and there exist PPT algorithms* $(\mathsf{TrapGen}, \mathsf{BadChallenge})$ *with the following properties.*

- $\mathsf{Gen}_{\mathsf{par}}$ *inputs* $\lambda \in \mathbb{N}$ *and outputs public parameters* $\mathsf{par} \leftarrow \mathsf{Gen}_{\mathsf{par}}(1^\lambda)$.
- $\mathsf{Gen}_{\mathcal{L}}$ *is a randomized algorithm that, on input of public parameters* par*, outputs the language-dependent part* $\mathsf{crs}_{\mathcal{L}} \leftarrow \mathsf{Gen}_{\mathcal{L}}(\mathsf{par}, \mathcal{L})$ *of* $\mathsf{crs} = (\mathsf{par}, \mathsf{crs}_{\mathcal{L}})$.
- $\mathsf{TrapGen}(\mathsf{par}, \mathcal{L}, \tau_{\mathcal{L}})$ *takes as input public parameters* par *and a membership-testing trapdoor* $\tau_{\mathcal{L}}$ *for the language* $\mathcal{L}_{\mathsf{sound}}$*. It outputs a common reference string* $\mathsf{crs}_{\mathcal{L}}$ *and a trapdoor* $\tau_\Sigma \in \{0,1\}^{\ell_\tau}$*, for some* $\ell_\tau(\lambda)$.
- $\mathsf{BadChallenge}(\tau_\Sigma, \mathsf{crs}, x, \mathbf{a})$ *takes in a trapdoor* τ_Σ*, a CRS* $\mathsf{crs} = (\mathsf{par}, \mathsf{crs}_{\mathcal{L}})$*, an instance* x*, and a first prover message* $\mathbf{a}$*. It outputs a challenge* Chall.

In addition, the following properties are required.

- **CRS indistinguishability:** *For any* $\mathsf{par} \leftarrow \mathsf{Gen}_{\mathsf{par}}(1^\lambda)$*, and any trapdoor* $\tau_{\mathcal{L}}$ *for the language* $\mathcal{L}$*, an honestly generated* $\mathsf{crs}_{\mathcal{L}}$ *is computationally indistinguishable from a CRS produced by* $\mathsf{TrapGen}(\mathsf{par}, \mathcal{L}, \tau_{\mathcal{L}})$*. Namely, for any* aux *and any PPT distinguisher* $\mathcal{A}$*, we have*

$$\begin{aligned}\mathbf{Adv}^{\text{indist-}\Sigma}_{\mathcal{A}}(\lambda) := |\Pr[\mathsf{crs}_{\mathcal{L}} \leftarrow \mathsf{Gen}_{\mathcal{L}}(\mathsf{par}, \mathcal{L}) : \mathcal{A}(\mathsf{par}, \mathsf{crs}_{\mathcal{L}}) = 1] \\ - \Pr[(\mathsf{crs}_{\mathcal{L}}, \tau_\Sigma) \leftarrow \mathsf{TrapGen}(\mathsf{par}, \mathcal{L}, \tau_{\mathcal{L}}) : \mathcal{A}(\mathsf{par}, \mathsf{crs}_{\mathcal{L}}) = 1]| \leq \mathsf{negl}(\lambda).\end{aligned}$$

- **Correctness:** *There exists a language-specific trapdoor* $\tau_{\mathcal{L}}$ *such that, for any instance* $x \notin \mathcal{L}_{\mathsf{sound}}$ *and all pairs* $(\mathsf{crs}_{\mathcal{L}}, \tau_\Sigma) \leftarrow \mathsf{TrapGen}(\mathsf{par}, \mathcal{L}, \tau_{\mathcal{L}})$*, we have* $\mathsf{BadChallenge}(\tau_\Sigma, \mathsf{crs}, x, \mathbf{a}) = f(\mathsf{crs}, x, \mathbf{a})$.

Note that the $\mathsf{TrapGen}$ algorithm does not take a specific statement x as input, but only a trapdoor $\tau_{\mathcal{L}}$ allowing to recognize elements of $\mathcal{L}_{\mathsf{sound}}$.

2.5 $\mathcal{R}$-Lossy Public-Key Encryption With Efficient Opening

We generalize the notion of $\mathcal{R}$-lossy public-key encryption introduced by Boyle *et al.* [19]. As defined in [19], it is a tag-based encryption scheme [60] where the tag space $\mathcal{T}$ is partitioned into a set of *injective* tags and a set of *lossy* tags. When ciphertexts are generated for an injective tag, the decryption algorithm correctly recovers the underlying plaintext. When messages are encrypted under lossy tags, the ciphertext is statistically independent of the plaintext. In $\mathcal{R}$-lossy PKE schemes, the tag space is partitioned according to a binary relation $\mathcal{R} \subseteq \mathcal{K} \times \mathcal{T}$. The key generation algorithm takes as input an initialization value $K \in \mathcal{K}$ and partitions $\mathcal{T}$ in such a way that injective tags $t \in \mathcal{T}$ are exactly those for which $(K, t) \in \mathcal{R}$ (i.e., all tags t for which $(K, t) \notin \mathcal{R}$ are lossy).

From a security standpoint, the definitions of [19] require the initialization value K to be computationally hidden by the public key. For our purposes, we need to introduce additional requirements.

First, we require the existence of a lossy key generation algorithm $\mathsf{LKeygen}$ which outputs public keys with respect to which all tags t are lossy (in contrast with injective keys where the only lossy tags are those for which $(K, t) \notin \mathcal{R}$). Second, we also ask that the secret key makes it possible to equivocate lossy ciphertexts (a property called *efficient opening* by Bellare *et al.* [10]) using an algorithm called Opener. Finally, we use two distinct opening algorithms Opener and $\mathsf{LOpener}$. The former operates over (lossy and injective) public keys for lossy tags while the latter can equivocate ciphertexts encrypted under lossy keys for any tag.

Definition 2.10. *Let $\mathcal{R} \subseteq \mathcal{K}_\lambda \times \mathcal{T}_\lambda$ be an efficiently computable binary relation. An $\mathcal{R}$-lossy PKE scheme with efficient opening is a 7-uple of PPT algorithms* $(\mathsf{Par\text{-}Gen}, \mathsf{Keygen}, \mathsf{LKeygen}, \mathsf{Encrypt}, \mathsf{Decrypt}, \mathsf{Opener}, \mathsf{LOpener})$ *such that:*

Parameter generation: *On input a security parameter λ, $\mathsf{Par\text{-}Gen}(1^\lambda)$ outputs public parameters Γ.*

Key generation: *For an initialization value $K \in \mathcal{K}_\lambda$ and public parameters Γ, algorithm $\mathsf{Keygen}(\Gamma, K)$ outputs an injective public key $pk \in \mathcal{PK}$, a decryption key $sk \in \mathcal{SK}$ and a trapdoor key $tk \in \mathcal{TK}$. The public key specifies a ciphertext space CtSp and a randomness space R^{LPKE}.*

Lossy Key generation: *Given an initialization value $K \in \mathcal{K}_\lambda$ and public parameters Γ, the lossy key generation algorithm $\mathsf{LKeygen}(\Gamma, K)$ outputs a lossy public key $pk \in \mathcal{PK}$, a lossy secret key $sk \in \mathcal{SK}$ and a trapdoor key $tk \in \mathcal{TK}$.*

Decryption under injective tags: *For any initialization value $K \in \mathcal{K}$, any tag $t \in \mathcal{T}$ such that $(K, t) \in \mathcal{R}$, and any message $\mathsf{Msg} \in \mathsf{MsgSp}$, we have*

$$\Pr\left[\exists r \in R^{\mathsf{LPKE}} : \mathsf{Decrypt}\big(sk, t, \mathsf{Encrypt}(pk, t, \mathsf{Msg}; r)\big) \neq \mathsf{Msg}\right] < \nu(\lambda) \ ,$$

for some negligible function $\nu(\lambda)$, where $(pk, sk, tk) \leftarrow \mathsf{Keygen}(\Gamma, K)$ and the probability is taken over the randomness of Keygen.

Indistinguishability: *Algorithms* $\mathsf{LKeygen}$ *and* Keygen *satisfy the following:*

(i) For any $K \in \mathcal{K}_\lambda$, the distributions $D_{\text{inj}} = \{(pk, tk) \mid (pk, sk, tk) \leftarrow \mathsf{Keygen}(\Gamma, K)\}$ and $D_{\text{loss}} = \{(pk, tk) \mid (pk, sk, tk) \leftarrow \mathsf{LKeygen}(\Gamma, K)\}$ are computationally indistinguishable. Namely, for any PPT adversary $\mathcal{A}$, we have $\mathbf{Adv}^{\mathsf{indist\text{-}LPKE\text{-}1}}_{\mathcal{A}}(\lambda) \leq \mathsf{negl}(\lambda)$, where

$$\mathbf{Adv}^{\mathsf{indist\text{-}LPKE\text{-}1}}_{\mathcal{A}}(\lambda) := |\Pr[(pk, tk) \hookleftarrow D_{\text{inj}} : \mathcal{A}(pk, tk) = 1] - \Pr[(pk, tk) \hookleftarrow D_{\text{loss}} : \mathcal{A}(pk, tk) = 1]| \ .$$

(ii) For any distinct initialization values $K, K' \in \mathcal{K}_\lambda$, the two distributions $\{pk \mid (pk, sk, tk) \leftarrow \mathsf{LKeygen}(\Gamma, K)\}$ and $\{pk \mid (pk, sk, tk) \leftarrow \mathsf{LKeygen}(\Gamma, K')\}$ are statistically indistinguishable. We require them to be $2^{-\Omega(\lambda)}$-close in terms of statistical distance.

Lossiness: *For any initialization value $K \in \mathcal{K}_\lambda$ and tag $t \in \mathcal{T}_\lambda$ such that $(K, t) \notin \mathcal{R}$, any $(pk, sk, tk) \leftarrow \mathsf{Keygen}(\Gamma, K)$, and any $\mathsf{Msg}_0, \mathsf{Msg}_1 \in \mathsf{MsgSp}$, the following distributions are statistically close:*

$$\{C \mid C \leftarrow \mathsf{Encrypt}(pk, t, \mathsf{Msg}_0)\} \quad \approx_s \quad \{C \mid C \leftarrow \mathsf{Encrypt}(pk, t, \mathsf{Msg}_1)\}.$$

For any $(pk, sk, tk) \leftarrow \mathsf{LKeygen}(\Gamma, K)$, the above holds for any tag t (and not only those for which $(K, t) \notin \mathcal{R}$).

Efficient opening under lossy tags: *Let D_R denote the distribution, defined over the randomness space R^{LPKE}, from which the random coins used by $\mathsf{Encrypt}$ are sampled. For any message $\mathsf{Msg} \in \mathsf{MsgSp}$ and ciphertext C, let $D_{PK,\mathsf{Msg},C,t}$ denote the probability distribution on R^{LPKE} with support*

$$S_{PK,\mathsf{Msg},C,t} = \{\overline{r} \in R^{\mathsf{LPKE}} \mid \mathsf{Encrypt}(pk, t, \mathsf{Msg}, \overline{r}) = C\} \ ,$$

and such that, for each $\overline{r} \in S_{PK,\mathsf{Msg},C,t}$, we have

$$D_{PK,\mathsf{Msg},C,t}(\overline{r}) = \Pr_{r' \hookleftarrow D_R}[r' = \overline{r} \mid \mathsf{Encrypt}(pk, t, \mathsf{Msg}, r') = C] \ .$$

There exists a PPT algorithm Opener such that, for any $K \in \mathcal{K}_\lambda$, any keys $(pk, sk, tk) \leftarrow \mathsf{Keygen}(\Gamma, K)$ and $(pk, sk, tk) \leftarrow \mathsf{LKeygen}(\Gamma, K)$, any random coins $r \hookleftarrow D_R$, any tag $t \in \mathcal{T}_\lambda$ such that $(K, t) \notin \mathcal{R}$, and any messages $\mathsf{Msg}_0, \mathsf{Msg}_1 \in \mathsf{MsgSp}$, takes as inputs $pk, C = \mathsf{Encrypt}(pk, t, \mathsf{Msg}_0, r)$, t, and tk. It outputs a sample $\overline{r}$ from a distribution statistically close to $D_{PK,\mathsf{Msg}_1,C,t}$.

Efficient opening under lossy keys: *There exists a PPT sampling algorithm $\mathsf{LOpener}$ such that, for any $K \in \mathcal{K}_\lambda$, any keys $(pk, sk, tk) \leftarrow \mathsf{LKeygen}(\Gamma, K)$, any random coins $r \hookleftarrow D_R$, any tag $t \in \mathcal{T}_\lambda$, and any distinct messages $\mathsf{Msg}_0, \mathsf{Msg}_1 \in \mathsf{MsgSp}$, takes as input $C = \mathsf{Encrypt}(pk, t, \mathsf{Msg}_0, r)$, t and sk. It outputs a sample $\overline{r}$ from a distribution statistically close to $D_{PK,\mathsf{Msg}_1,C,t}$.*

In Definition 2.10, some of the first four properties were defined in [19, Definition 4.1]. The last two properties are a natural extension of the definition of efficient opening introduced by Bellare *et al.* [10]. We note that property of

decryption under injective tags does not assume that random coins are honestly sampled, but only that they belong to some pre-defined set R^{LPKE}.

For our applications to simulation-sound proofs, it would be sufficient to have algorithms $(\mathsf{Opener}, \mathsf{LOpener})$ that have access to the initial messages Msg_0 *and* the random coins r_0 of the ciphertext to be equivocated (as was the case in the opening algorithms of [10]). In our LWE-based construction, however, the initial messages and random coins are not needed.

3 Direct Construction of Unbounded Simulation-Sound NIZK Arguments

We provide a method that directly compiles any trapdoor Σ-protocol into an unbounded simulation-sound NIZK argument using an $\mathcal{R}$-lossy encryption scheme for the bit-matching relation $\mathcal{R}_{\mathsf{BM}}$ and a correlation intractable hash function.

Definition 3.1. *Let $\mathcal{K} = \{0,1,\perp\}^L$ and $\mathcal{T} = \{0,1\}^\ell$, for some $\ell, L \in \mathsf{poly}(\lambda)$ such that $\ell < L$. Let F_{ADH} the partitioning function defined by $\mathsf{AHF} : \{0,1\}^\ell \to \{0,1\}^L$ in Definition 2.6. The* **bit-matching relation** *$\mathcal{R}_{\mathsf{BM}} : \mathcal{K} \times \mathcal{T} \to \{0,1\}$ for AHF is the relation where $\mathcal{R}_{\mathsf{BM}}(K,t) = 1$ if and only if $K = K_1 \dots K_L$ and $t = t_1 \dots t_\ell$ satisfy $F_{\mathsf{ADH}}(K,t) = 0$ (namely, $\bigwedge_{i=1}^{L}(K_i = \perp) \vee (K_i = \mathsf{AHF}(t)_i)$).*

3.1 An $\mathcal{R}_{\mathsf{BM}}$-Lossy PKE Scheme from LWE

We describe an $\mathcal{R}_{\mathsf{BM}}$-lossy PKE scheme below. Our scheme builds on a variant of the primal Regev cryptosystem [79] suggested in [46].

Let $\mathsf{AHF} : \{0,1\}^\ell \to \{0,1\}^L$ an admissible hash function with key space $\mathcal{K} = \{0,1,\perp\}^L$ and let $\mathcal{R}_{\mathsf{BM}} \subset \mathcal{K} \times \{0,1\}^\ell$ the corresponding bit-matching relation. We construct an $\mathcal{R}_{\mathsf{BM}}$-lossy PKE scheme in the following way.

Par-Gen(1^λ)**:** Given a security parameter $\lambda \in \mathbb{N}$, let $n_0 = \mathsf{poly}(\lambda)$ the length of messages. Choose a prime modulus $q = \mathsf{poly}(\lambda)$; dimensions $n = n_0 + \Omega(\lambda)$ and $m = 2n\lceil \log q \rceil + O(\lambda)$. Define the tag space as $\mathcal{T} = \{0,1\}^\ell$ where $\ell = \Theta(\lambda)$. Define the initialization value space $\mathcal{K} = \{0,1,\perp\}^L$ and Gaussian parameters $\sigma = O(m) \cdot L$ and $\alpha \in (0,1)$ such that $m\alpha q \cdot (L+1) \cdot \sigma\sqrt{2m} < q/4$. Define public parameters as $\Gamma = (\ell, L, n_0, q, n, m, \alpha, \sigma)$.

Keygen(Γ, K)**:** On input of public parameters Γ and an initialization value $K \in \{0,1,\perp\}^L$, generate a key pair as follows.

1. Sample random matrices $\bar{\mathbf{B}} \hookleftarrow U(\mathbb{Z}_q^{(n-n_0)\times m})$, $\mathbf{S} \hookleftarrow U(\mathbb{Z}_q^{(n-n_0)\times n_0})$ and a small-norm $\mathbf{E} \hookleftarrow \chi^{m\times n_0}$ to compute

$$\mathbf{A} = \left[\frac{\bar{\mathbf{B}}}{\mathbf{S}^\top \cdot \bar{\mathbf{B}} + \mathbf{E}^\top}\right] \in \mathbb{Z}_q^{n\times m}.$$

2. Parse K as $K_1 \dots K_L \in \{0,1,\perp\}^L$. Letting $\mathbf{G} \in \mathbb{Z}_q^{n\times m}$ denote the gadget matrix, for each $i \in [L]$ and $b \in \{0,1\}$, compute matrices $\mathbf{A}_{i,b} \in \mathbb{Z}_q^{n\times m}$ as

$$\mathbf{A}_{i,b} = \begin{cases} \mathbf{A}\cdot\mathbf{R}_{i,b} + \mathbf{G} & \text{if } (K_i \neq \perp) \wedge (b = 1 - K_i) \\ \mathbf{A}\cdot\mathbf{R}_{i,b} & \text{if } (K_i = \perp) \vee (b = K_i). \end{cases} \tag{1}$$

where $\mathbf{R}_{i,b} \hookleftarrow U(\{-1,1\}^{m\times m})$ for all $i \in [L]$ and $b \in \{0,1\}$.

Define $R^{\mathsf{LPKE}} = \{\mathbf{r} \in \mathbb{Z}^{2m} \mid \|\mathbf{r}\| \leq \sigma\sqrt{2m}\}$ and output $sk = (K, \mathbf{S})$ as well as

$$pk := \Big(\mathbf{A},\ \{\mathbf{A}_{i,b}\}_{(i,b)\in[L]\times\{0,1\}}\Big), \qquad tk = (K, \{\mathbf{R}_{i,b}\}_{(i,b)\in[L]\times\{0,1\}}).$$

LKeygen(Γ, K)**:** This algorithm proceeds identically to Keygen except that steps 1 and 2 are modified in the following way.

1. Run $(\mathbf{A}, \mathbf{T}_{\mathbf{A}}) \leftarrow \mathsf{GenTrap}(1^\lambda, 1^n, 1^m, q)$ so as to obtain a statistically uniform matrix $\mathbf{A} \sim U(\mathbb{Z}_q^{n\times m})$ with a trapdoor for the lattice $\Lambda^\perp(\mathbf{A})$.
2. Define matrices $\{\mathbf{A}_{i,b} \in \mathbb{Z}_q^{n\times m}\}_{(i,b)\in[L]\times\{0,1\}}$ as in (1).

Define R^{LPKE} as in Keygen and output

$$pk := \Big(\mathbf{A},\ \{\mathbf{A}_{i,b}\}_{(i,b)\in[L]\times\{0,1\}}\Big), \quad sk = \mathbf{T}_{\mathbf{A}}, \quad tk = (K, \{\mathbf{R}_{i,b}\}_{(i,b)\in[L]\times\{0,1\}}).$$

Encrypt(pk, t, Msg)**:** To encrypt $\mathsf{Msg} \in \{0,1\}^{n_0}$ for the tag $t = t_1 \dots t_\ell \in \{0,1\}^\ell$, conduct the following steps.

1. Encode the tag t as $t' = t'_1 \dots t'_L = \mathsf{AHF}(t) \in \{0,1\}^L$ and compute $\mathbf{A}_{F,t} = \sum_{i=1}^L \mathbf{A}_{i,t'_i} \in \mathbb{Z}_q^{n\times m}$. Note that $\mathbf{A}_{F,t} = \mathbf{A}\cdot\mathbf{R}_{F,t} + d_t\cdot\mathbf{G}$ for some $\mathbf{R}_{F,t} \in \mathbb{Z}^{m\times m}$ of norm $\|\mathbf{R}_{F,t}\|_\infty \leq L$ and where $d_t \in \{0,\dots,L\}$ is the number of non-$\perp$ entries of K for which $K_i \neq t'_i$.
2. Choose $\mathbf{r} \hookleftarrow D_{\mathbb{Z}^{2m},\sigma}$ and output $\perp$ if $\mathbf{r} \notin R^{\mathsf{LPKE}}$. Otherwise, output

$$\mathbf{c} = [\mathbf{A} \mid \mathbf{A}_{F,t}]\cdot\mathbf{r} + \begin{bmatrix} \mathbf{0}^{n-n_0} \\ \mathsf{Msg}\cdot\lfloor q/2\rfloor \end{bmatrix} \in \mathbb{Z}_q^n \ . \tag{2}$$

Decrypt$(sk, t, \mathbf{c})$**:** Given $sk = (K, \mathbf{S})$ and the tag $t \in \{0,1\}^\ell$, compute $t' = t'_1 \dots t'_L = \mathsf{AHF}(t) \in \{0,1\}^L$ and return $\perp$ if $R_{\mathsf{BM}}(K, t') = 0$. Otherwise, compute $\mathbf{w} = [-\mathbf{S}^\top \mid \mathbf{I}_{n_0}]\cdot\mathbf{c} \in \mathbb{Z}_q^{n_0}$. For each $i \in [n_0]$, do the following:

1. If neither $\mathbf{w}[i]$ nor $|\mathbf{w}[i] - \lfloor q/2\rfloor|$ is close to 0, halt and return $\perp$.
2. Otherwise, set $\mathsf{Msg}[i] \in \{0,1\}$ so as to minimize $|\mathbf{w}[i] - \mathsf{Msg}[i]\cdot\lfloor q/2\rfloor|$.

Return $\mathsf{Msg} = \mathsf{Msg}[1]\dots\mathsf{Msg}[n_0]$.

Opener$(pk, tk, t, \mathbf{c}, \mathsf{Msg}_1)$**:** Given $tk = (K, \{\mathbf{R}_{i,b}\}_{i,b})$ and $t \in \{0,1\}^\ell$, compute $t' = t'_1 \dots t'_L = \mathsf{AHF}(t) \in \{0,1\}^L$ and return $\perp$ if $R_{\mathsf{BM}}(K, t') = 1$. Otherwise,

1. Compute the small-norm matrix $\mathbf{R}_{F,t} = \sum_{i=1}^L \mathbf{R}_{i,t'_i} \in \mathbb{Z}^{m\times m}$ such that $\mathbf{A}_{F,t} = \mathbf{A}\cdot\mathbf{R}_{F,t} + d_t\cdot\mathbf{G}$ and $\|\mathbf{R}_{F,t}\|_\infty \leq L$ with $d_t \in [L]$.

2. Use $\mathbf{R}_{F,t} \in \mathbb{Z}^{m\times m}$ as a trapdoor for the matrix

$$\bar{\mathbf{A}}_{F,t} = [\mathbf{A} \mid \mathbf{A}_{F,t}] = [\mathbf{A} \mid \mathbf{A}\cdot\mathbf{R}_{F,t} + d_t \cdot \mathbf{G}] \ \in \mathbb{Z}_q^{n\times 2m}$$

to sample a Gaussian vector $\bar{\mathbf{r}} \in \mathbb{Z}^{2m}$ such that

$$\bar{\mathbf{A}}_{F,t}\cdot\bar{\mathbf{r}} = \mathbf{c} - \left[\frac{\mathbf{0}^{n-n_0}}{\mathsf{Msg}_1 \cdot \lfloor q/2 \rfloor}\right] \ . \tag{3}$$

Namely, defining $\mathbf{c}_{\mathsf{Msg}_1} = \mathbf{c} - [(\mathbf{0}^{n-n_0})^\top \mid \mathsf{Msg}_1^\top \cdot \lfloor q/2 \rfloor]^\top$, sample and output fake random coins $\bar{\mathbf{r}} \hookleftarrow D_{\Lambda_q^{\mathbf{c}_{\mathsf{Msg}_1}}(\bar{\mathbf{A}}_{F,t}),\sigma}$.

LOpener$(sk, t, \mathbf{c}, \mathsf{Msg}_1)$**:** Given $sk = \mathbf{T}_{\mathbf{A}}$ and $t \in \{0,1\}^\ell$, use $\mathbf{T}_{\mathbf{A}}$ to derive a trapdoor $\mathbf{T}_{\mathbf{A},t}$ for the lattice $\Lambda_q^\perp(\bar{\mathbf{A}}_{F,t})$ and use $\mathbf{T}_{\mathbf{A},t}$ to sample a Gaussian vector $\bar{\mathbf{r}} \hookleftarrow D_{\Lambda_q^{\mathbf{c}_{\mathsf{Msg}_1}}(\bar{\mathbf{A}}_{F,t}),\sigma}$ satisfying (3).

The above construction requires $2L = \Theta(\lambda)$ matrices in the public key but allows for a relatively small modulus $q = \Theta(m^{5/2}n^{1/2}L^2)$. A technique suggested by Yamada [83] can be used to reduce the number of public matrices to $O(\log^2 \lambda)$ at the expense of a larger (but still polynomial) modulus. Since our application to Naor-Yung requires a public key containing a large correlation-intractable hashing key anyway, we chose to minimize the modulus size.

Theorem 3.2 states that the construction has the required properties under the LWE assumption. The proof is given in the full version of the paper [65].

Theorem 3.2. *The above construction is an $\mathcal{R}_{\mathsf{BM}}$-lossy public-key encryption scheme with efficient opening under the* LWE *assumption.*

3.2 A Generic Construction from Trapdoor Σ-Protocols and $\mathcal{R}_{\mathsf{BM}}$-lossy PKE

We construct unbounded simulation-sound NIZK proofs by combining trapdoor Σ-protocols and $\mathcal{R}$-lossy public-key encryption schemes. Our proof system is inspired by ideas from [72] and relies on the following ingredients:

- A trapdoor Σ-protocol $\Pi' = (\mathsf{Gen}'_{\mathsf{par}}, \mathsf{Gen}'_{\mathcal{L}}, \mathsf{P}', \mathsf{V}')$ with challenge space $\mathcal{C}$, for the same language $\mathcal{L} = (\mathcal{L}_{\mathsf{zk}}, \mathcal{L}_{\mathsf{sound}})$ and which satisfies the properties of Definition 2.9. In addition, $\mathsf{BadChallenge}(\tau_\Sigma, \mathsf{crs}, x, \mathbf{a})$ should be computable within time $T \in \mathsf{poly}(\lambda)$ for any input $(\tau, \mathsf{crs}, x, \mathbf{a})$.
- A strongly unforgeable one-time signature scheme $\mathsf{OTS} = (\mathcal{G}, \mathcal{S}, \mathcal{V})$ with verification keys of length $\ell \in \mathsf{poly}(\lambda)$.
- An admissible hash function $\mathsf{AHF} : \{0,1\}^\ell \to \{0,1\}^L$, for some $L \in \mathsf{poly}(\lambda)$ with $L > \ell$, which induces the relation $\mathcal{R}_{\mathsf{BM}} : \{0,1,\perp\}^L \times \{0,1\}^\ell \to \{0,1\}$.
- An $\mathcal{R}$-lossy PKE scheme $\mathcal{R}$-$\mathsf{LPKE} = (\mathsf{Par\text{-}Gen}, \mathsf{Keygen}, \mathsf{LKeygen}, \mathsf{Encrypt}, \mathsf{Decrypt}, \mathsf{Opener}, \mathsf{LOpener})$ for the relation $\mathcal{R}_{\mathsf{BM}} : \{0,1,\perp\}^L \times \{0,1\}^\ell \to \{0,1\}$ with public (resp. secret) key space $\mathcal{PK}$ (resp. $\mathcal{SK}$). We assume that $\mathsf{Decrypt}$ is computable within time T. We denote the message (resp. ciphertext) space by MsgSp (resp. CtSp) and the randomness space by R^{LPKE}. Let also D_R^{LPKE} denote the distribution from which the random coins of $\mathsf{Encrypt}$ are sampled.

- A correlation intractable hash family $\mathcal{H} = (\mathsf{Gen}, \mathsf{Hash})$ for the class $\mathcal{R}_{\mathsf{CI}}$ of relations that are efficiently searchable within time T.

We also assume that these ingredients are compatible in the sense that P' outputs a first prover message $\mathbf{a}$ that fits in the message space MsgSp of $\mathcal{R}$-LPKE.

Our argument system $\Pi^{\mathrm{uss}} = (\mathsf{Gen}_{\mathsf{par}}, \mathsf{Gen}_{\mathcal{L}}, \mathsf{P}, \mathsf{V})$ allows P and V to input a label lbl consisting of public data. While this label will be the empty string in our KDM-CCA scheme of Section, it may be useful when several non-interactive arguments have to be bound together. The construction goes as follows.

$\mathbf{Gen}_{\mathsf{par}}(1^\lambda)$**:** Run $\mathsf{par} \leftarrow \mathsf{Gen}'_{\mathsf{par}}(1^\lambda)$ and output par.

$\mathbf{Gen}_{\mathcal{L}}(\mathsf{par}, \mathcal{L})$**:** Given public parameters par and a language $\mathcal{L} \subset \{0,1\}^N$, let $\mathcal{K} = \{0, 1, \perp\}^L$ and $\mathcal{T} = \{0,1\}^\ell$. The CRS is generated as follows.

1. Generate a CRS $\mathsf{crs}'_{\mathcal{L}} \leftarrow \mathsf{Gen}'_{\mathcal{L}}(\mathsf{par}, \mathcal{L})$ for the trapdoor Σ-protocol Π'.
2. Generate public parameters $\Gamma \hookleftarrow \mathsf{Par\text{-}Gen}(1^\lambda)$ for the $\mathcal{R}_{\mathsf{BM}}$-lossy PKE scheme where the relation $\mathcal{R}_{\mathsf{BM}} : \mathcal{K} \times \mathcal{T} \to \{0,1\}$ is defined by an admissible hash function $\mathsf{AHF} : \{0,1\}^\ell \to \{0,1\}^L$. Choose a random initialization value $K \leftarrow \mathcal{K}$ and generate lossy keys $(pk, sk, tk) \leftarrow \mathsf{LKeygen}(\Gamma, K)$.
3. Generate a key $k \leftarrow \mathsf{Gen}(1^\lambda)$ for a correlation intractable hash function with output length $\kappa = \Theta(\lambda)$.

Output the language-dependent $\mathsf{crs}_{\mathcal{L}} := \big(\mathsf{crs}'_{\mathcal{L}}, k\big)$ and the simulation trapdoor $\tau_{\mathrm{zk}} := sk$, which is the lossy secret key of $\mathcal{R}$-LPKE. The global common reference string consists of $\mathsf{crs} = (\mathsf{par}, \mathsf{crs}_{\mathcal{L}}, pk, \mathsf{AHF}, \mathsf{OTS})$.

$\mathbf{P}(\mathsf{crs}, x, w, \mathsf{lbl})$: To prove a statement x for a label $\mathsf{lbl} \in \{0,1\}^*$ using $w \in R_{\mathsf{zk}}(x)$, generate a one-time signature key pair $(\mathsf{VK}, \mathsf{SK}) \leftarrow \mathcal{G}(1^\lambda)$. Then,

1. Compute $\big(\mathbf{a}' = (\mathbf{a}'_1, \ldots, \mathbf{a}'_\kappa), st'\big) \leftarrow \mathsf{P}'(\mathsf{crs}'_{\mathcal{L}}, x, w)$ via κ invocations of the prover for Π'. Then, for each $i \in [\kappa]$, compute $\mathbf{a}_i \leftarrow \mathsf{Encrypt}(pk, \mathsf{VK}, \mathbf{a}'_i; \mathbf{r}_i)$ using random coins $\mathbf{r}_i \hookleftarrow D_R^{\mathsf{LPKE}}$. Let $\mathbf{a} = (\mathbf{a}_1, \ldots, \mathbf{a}_\kappa)$ and $\mathbf{r} = (\mathbf{r}_1, \ldots, \mathbf{r}_\kappa)$.
2. Compute $\mathsf{Chall} = \mathsf{Hash}(k, (x, \mathbf{a}, \mathsf{VK})) \in \{0,1\}^\kappa$.
3. Compute $\mathbf{z}' = (\mathbf{z}'_1, \ldots, \mathbf{z}'_\kappa) = \mathsf{P}'(\mathsf{crs}'_{\mathcal{L}}, x, w, \mathbf{a}', \mathsf{Chall}, st')$ via κ executions of the prover of Π'. Define $\mathbf{z} = (\mathbf{z}', \mathbf{a}', \mathbf{r})$.
4. Generate $sig \leftarrow \mathcal{S}(\mathsf{SK}, (x, \mathbf{a}, \mathbf{z}, \mathsf{lbl}))$ and output $\boldsymbol{\pi} = \big(\mathsf{VK}, (\mathbf{a}, \mathbf{z}), sig\big)$.

$\mathbf{V}(\mathsf{crs}, x, \boldsymbol{\pi}, \mathsf{lbl})$: Given a statement x, a label lbl as well as a purported proof $\boldsymbol{\pi} = \big(\mathsf{VK}, (\mathbf{a}, \mathbf{z}), sig\big)$, return 0 if $\mathcal{V}(\mathsf{VK}, (x, \mathbf{a}, \mathbf{z}, \mathsf{lbl}), sig) = 0$. Otherwise,

1. Write $\mathbf{z}$ as $\mathbf{z} = \big((\mathbf{z}'_1, \ldots, \mathbf{z}'_\kappa), (\mathbf{a}'_1, \ldots, \mathbf{a}'_\kappa), (\mathbf{r}_1, \ldots, \mathbf{r}_\kappa)\big)$ and return 0 if it does not parse properly. Return 0 if there exists $i \in [\kappa]$ such that $\mathbf{a}_i \neq \mathsf{Encrypt}(pk, \mathsf{VK}, \mathbf{a}'_i; \mathbf{r}_i)$ or $\mathbf{r}_i \notin R^{\mathsf{LPKE}}$.
2. Let $\mathsf{Chall} = \mathsf{Hash}\big(k, (x, (\mathbf{a}_1, \ldots, \mathbf{a}_\kappa), \mathsf{VK})\big)$. If $\mathsf{V}'(\mathsf{crs}'_{\mathcal{L}}, x, (\mathbf{a}'_i, \mathsf{Chall}[i], \mathbf{z}'_i)) = 1$ for each $i \in [\kappa]$, return 1. Otherwise, return 0.

Our NIZK simulator uses a technique due to Damgård [38], which uses a trapdoor commitment scheme to achieve a straight-line simulation of 3-move zero-knowledge proofs in the common reference string model.

Theorem 3.3. *The above argument system is multi-theorem zero-knowledge assuming that the trapdoor Σ-protocol Π' is special zero-knowledge.*

Proof (Sketch). We describe a simulator $(\mathsf{Sim}_0, \mathsf{Sim}_1)$ which uses the lossy secret key $\tau_{\mathrm{zk}} = sk$ of $\mathcal{R}$-LPKE to simulate transcripts $(\mathbf{a}, \mathsf{Chall}, \mathbf{z})$ without using the witnesses. Namely, on input of $\mathsf{par} \leftarrow \mathsf{Gen}_{\mathsf{par}}(1^\lambda)$, Sim_0 generates $\mathsf{crs}_{\mathcal{L}}$ by proceeding identically to $\mathsf{Gen}_{\mathcal{L}}$ while Sim_1 is described hereunder.

$\mathbf{Sim}_1(\mathsf{crs}, \tau_{\mathrm{zk}}, x, \mathsf{lbl})$**:** On input a statement $x \in \{0,1\}^N$, a label lbl and the simulation trapdoor $\tau_{\mathrm{zk}} = sk$, algorithm Sim_1 proceeds as follows.

1. Generate a one-time signature key pair $(\mathsf{VK}, \mathsf{SK}) \leftarrow \mathcal{G}(1^\lambda)$. Let $\mathbf{0}^{|\mathbf{a}'|}$ the all-zeroes string of length $|\mathbf{a}'|$. Sample random coins $\mathbf{r}_0 \hookleftarrow D_R^{\mathsf{LPKE}}$ from the distribution D_R^{LPKE} and compute $\mathbf{a} \leftarrow \mathsf{Encrypt}(pk, \mathsf{VK}, \mathbf{0}^{|\mathbf{a}'|}; \mathbf{r}_0)$.
2. Compute $\mathsf{Chall} = \mathsf{Hash}(k, (x, \mathbf{a}, \mathsf{VK}))$.
3. Run the special ZK simulator $(\mathbf{a}', \mathbf{z}') \leftarrow \mathsf{ZKSim}(\mathsf{crs}'_{\mathcal{L}}, x, \mathsf{Chall})$ of Π' to obtain a simulated transcript $(\mathbf{a}', \mathsf{Chall}, \mathbf{z}')$ of Π' for the challenge Chall.
4. Using the lossy secret key sk of $\mathcal{R}$-LPKE, compute random coins $\mathbf{r} \leftarrow \mathsf{LOpener}(sk, \mathsf{VK}, \mathbf{a}, \mathbf{a}')$ which explain $\mathbf{a}$ as an encryption of $(x, \mathbf{a}')$ under the tag VK. Then, define $\mathbf{z} = (\mathbf{z}', \mathbf{a}', \mathbf{r})$
5. Compute $sig \leftarrow \mathcal{S}(\mathsf{SK}, (x, \mathbf{a}, \mathbf{z}, \mathsf{lbl}))$ and output $\boldsymbol{\pi} = \big(\mathsf{VK}, (\mathbf{a}, \mathbf{z}), sig\big)$.

In the full version of the paper, we show that the simulation is statistically indistinguishable from proofs generated by the real prover. □

If we just target multi-theorem NIZK without simulation-soundness, the construction can be simplified as shown in the full version of the paper, where we explain how it can provide statistical zero-knowledge in the common random string (instead of the common reference string) model.

Going back to simulation-soundness, our proof builds on techniques used in [38,72]. The interactive proof systems of [72] rely on commitment schemes where the adversary cannot break the computational binding property of the commitment for some tag after having seen equivocations of commitments for different tags. Here, in order to use a correlation-intractable hash function, we need a commitment scheme which is equivocable on some tags but (with noticeable probability) becomes statistically binding on an adversarially-chosen tag. For this purpose, we exploit the observation that an $\mathcal{R}$-lossy PKE scheme can be used as a commitment scheme with these properties. Namely, it can serve as a trapdoor commitment to equivocate lossy encryptions of the first prover message in Π' while forcing the adversary to create a fake proof on a statistically binding (and even extractable) commitment.

At a high level, the proof also bears similarities with [66] in that they also use a commitment scheme that is statistically hiding in adversarial queries but becomes statistically binding in the adversary's output. The difference is that we need to equivocate the statistically-hiding commitment in simulated proofs here.

Theorem 3.4. *The above argument system provides unbounded simulation-soundness if: (i)* OTS *is a strongly unforgeable one-time signature; (ii)* $\mathcal{R}$-LPKE

is an $\mathcal{R}_{\mathsf{BM}}$-lossy PKE scheme; (iii) The hash family $\mathcal{H}$ is somewhere correlation-intractable for all relations that are searchable within time T, where T denotes the maximal running time of algorithms $\mathsf{BadChallenge}(\cdot,\cdot,\cdot,\cdot)$ *and* $\mathsf{Decrypt}(\cdot,\cdot,\cdot)$. (The proof is given in the full version of the paper.)

The work of Peikert and Shiehian [76] implies a correlation intractable hash function for the relation R_{bad} defined in the proof of Theorem 3.4. Their bootstrapping theorem actually implies the existence of such a hash family under the LWE assumption with polynomial approximation factors.

4 Tightly Secure Simulation-Sound Arguments

To achieve tight simulation-soundness, we describe an $\mathcal{R}$-lossy PKE scheme for a relation induced by a pseudorandom function family. In Definition 4.1, we assume that the tag space $\mathcal{T}$ has a special structure. Namely, each tag $t = (t_c, t_a) \in \mathcal{T}$ consists of a core component $t_c \in \{0,1\}^\lambda$ and an auxiliary component $t_a \in \{0,1\}^\ell$.

Definition 4.1. *Let a pseudorandom function* $\mathsf{PRF} : \mathcal{K} \times \{0,1\}^\ell \to \{0,1\}^\lambda$ *with key space* $\mathcal{K} = \{0,1\}^\lambda$ *and input space* $\{0,1\}^\ell$. *Let* $\mathcal{T} = \{0,1\}^\lambda \times \{0,1\}^\ell$, *for some* $\ell \in \mathsf{poly}(\lambda)$. *We define the PRF relation* $\mathcal{R}_{\mathsf{PRF}} : \mathcal{K} \times \mathcal{T} \to \{0,1\}$ *as* $\mathcal{R}_{\mathsf{PRF}}(K, (t_c, t_a)) = 1$ *if and only if* $t_c \neq \mathsf{PRF}(K, t_a)$.

We rely on the idea (previously used in [17,67]) of homomorphically evaluating the circuit of a PRF using the GSW FHE [47]. As observed in [22], when the circuit is in NC1, it is advantageous to convert it into a branching program using Barrington's theorem. This enables the use of a polynomial modulus q.

Lemma 4.2 (Adapted from [47,15]). *Let* $C : \{0,1\}^L \to \{0,1\}$ *be a NAND Boolean circuit of depth* d. *Let* $\mathbf{A}_i = \mathbf{A} \cdot \mathbf{R}_i + k_i \cdot \mathbf{G} \in \mathbb{Z}_q^{n \times m}$ *with* $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, $\mathbf{R}_i \in \{-1,1\}^{m \times m}$ *and* $k_i \in \{0,1\}$, *for* $i \leq L$. *There exist deterministic algorithms* $\mathsf{Eval}_{\mathsf{BP}}^{\mathsf{pub}}$ *and* $\mathsf{Eval}_{\mathsf{BP}}^{\mathsf{priv}}$ *with running time* $\mathsf{poly}(4^d, L, m, n, \log q)$ *that satisfy:* $\mathsf{Eval}_{\mathsf{BP}}^{\mathsf{pub}}(C, (\mathbf{A}_i)_i) = \mathbf{A} \cdot \mathsf{Eval}_{\mathsf{BP}}^{\mathsf{priv}}\big(C, ((\mathbf{R}_i, k_i))_i\big) + C(k_1, \ldots, k_L) \cdot \mathbf{G}$, *and* $\|\mathsf{Eval}_{\mathsf{BP}}^{\mathsf{priv}}(C, (\mathbf{R}_i, k_i)_i)\| \leq 4^d \cdot O(m^{3/2})$.

4.1 An $\mathcal{R}_{\mathsf{PRF}}$-Lossy PKE Scheme

We describe an $\mathcal{R}$-lossy PKE scheme for the relation $\mathcal{R}_{\mathsf{PRF}}$ of Definition 4.1.

Let $\mathsf{PRF} : \mathcal{K} \times \{0,1\}^\ell \to \{0,1\}^\lambda$ with key space $\mathcal{K} = \{0,1\}^\lambda$ and input space $\{0,1\}^\ell$ and let $\mathcal{R}_{\mathsf{PRF}} \subset \mathcal{K} \times \mathcal{T}$ the corresponding relation. We construct an $\mathcal{R}_{\mathsf{PRF}}$-lossy PKE scheme in the following way.

Par-Gen(1^λ)**:** Given a security parameter $\lambda \in \mathbb{N}$, let $n_0 = \mathsf{poly}(\lambda)$ the length of messages. Choose a prime modulus $q = \mathsf{poly}(\lambda)$; dimensions $n = n_0 + \Omega(\lambda)$ and $m = 2n\lceil \log q \rceil + O(\lambda)$. Define the tag space as $\mathcal{T} = \{0,1\}^\lambda \times \{0,1\}^\ell$ where $\ell = \Theta(\lambda)$. Define the initialization value space $\mathcal{K} = \{0,1\}^\lambda$ and Gaussian parameters $\sigma = 4^d \cdot O\left(m^2\right)$ and $\alpha \in (0,1)$ such that $4^d m^{3.5} \alpha q \cdot \sigma < q$. Define public parameters as $\Gamma = (\ell, L, n_0, q, n, m, u, \alpha, \sigma)$.

Keygen(Γ, K)**:** On input of public parameters Γ and an initialization value $K \in \{0,1\}^\lambda$, generate a key pair as follows.

1. Sample random matrices $\bar{\mathbf{B}} \hookleftarrow U(\mathbb{Z}_q^{(n-n_0)\times m})$, $\mathbf{S} \hookleftarrow U(\mathbb{Z}_q^{(n-n_0)\times n_0})$ and a small-norm $\mathbf{E} \hookleftarrow \chi^{m\times n_0}$ to compute $\mathbf{A} = [\bar{\mathbf{B}}^\top \mid \bar{\mathbf{B}}^\top\mathbf{S} + \mathbf{E}]^\top \in \mathbb{Z}_q^{n\times m}$.
2. Parse K as $k_1 \ldots k_\lambda \in \{0,1\}^\lambda$. For each $i \in [L]$, compute matrices $\mathbf{A}_i = \mathbf{A}\cdot\mathbf{R}_i + k_i \cdot \mathbf{G}$, where $\mathbf{R}_i \hookleftarrow U(\{-1,1\}^{m\times m})$, for all $i \in [\lambda]$.

Define $R^{\mathsf{LPKE}} = \{\mathbf{r} \in \mathbb{Z}^{2m} \mid \|\mathbf{r}\| \leq \sigma\sqrt{2m}\}$ and output $sk = (K, \mathbf{S})$ as well as

$$pk := \Big(\mathbf{A},\ \{\mathbf{A}_i\}_{i\in[\lambda]}\Big), \qquad tk = (K, \{\mathbf{R}_i\}_{i\in[\lambda]}).$$

LKeygen(Γ, K)**:** This algorithm proceeds identically to Keygen except that steps 1 and 2 are modified in the following way.

1. Run $(\mathbf{A}, \mathbf{T}_\mathbf{A}) \leftarrow \mathsf{GenTrap}(1^\lambda, 1^n, 1^m, q)$ to obtain a statistically uniform $\mathbf{A} \sim U(\mathbb{Z}_q^{n\times m})$ with a trapdoor for $\Lambda^\perp(\mathbf{A})$.
2. Define matrices $\{\mathbf{A}_i \in \mathbb{Z}_q^{n\times m}\}_{i\in[\lambda]}$ as in Keygen.

Output $pk := \Big(\mathbf{A},\ \{\mathbf{A}_i\}_{i\in[\lambda]}\Big)$, $sk = \mathbf{T}_\mathbf{A}$, and $tk = (K, \{\mathbf{R}_i\}_{i\in[\lambda]})$.

Encrypt(pk, t, Msg)**:** To encrypt a message $\mathsf{Msg} \in \{0,1\}^{n_0}$ for the structured tag $t = (t_c, t_a) \in \mathcal{T} = \{0,1\}^\lambda \times \{0,1\}^\ell$, conduct the following steps.

1. Let $C_{\mathsf{PRF},t} : \{0,1\}^\lambda \to \{0,1\}$ the circuit, where $t = (t_c, t_a)$ is hard-wired, which inputs a λ-bit key $K = k_1 \ldots k_\lambda \in \{0,1\}^\lambda$ and outputs $C_{\mathsf{PRF},t}(K)$ such that $C_{\mathsf{PRF},t}(K) = 1 \ \Leftrightarrow\ t_c = \mathsf{PRF}_K(t_a) \ \Leftrightarrow\ \mathcal{R}_{\mathsf{PRF}}(K,t) = 0$. Compute $\mathbf{A}_{F,t} \leftarrow \mathsf{Eval}_{\mathsf{BP}}^{\mathsf{pub}}(C_{\mathsf{PRF}}, (\mathbf{A}_i)_i) \in \mathbb{Z}_q^{n\times m}$ such that
$$\mathbf{A}_{F,t} = \mathbf{A}\cdot\mathbf{R}_t + C_{\mathsf{PRF},t}(K)\cdot\mathbf{G},$$
where $\mathbf{R}_t = \mathsf{Eval}_{\mathsf{BP}}^{\mathsf{priv}}(C_{\mathsf{PRF},t}, (\mathbf{R}_i, k_i)_i) \in \mathbb{Z}^{m\times m}$ s.t. $\|\mathbf{R}_t\| \leq 4^d \cdot O(m^{3/2})$.
2. Choose $\mathbf{r} \hookleftarrow D_{\mathbb{Z}^{2m},\sigma}$ and output $\perp$ if $\mathbf{r} \notin R^{\mathsf{LPKE}}$. Otherwise, output
$$\mathbf{c} = [\mathbf{A} \mid \mathbf{A}_{F,t}]\cdot\mathbf{r} + [\mathbf{0}^{n-n_0\top} \mid \mathsf{Msg}\cdot\lfloor q/2\rfloor]^\top \ \in \mathbb{Z}_q^n \ .$$

Decrypt$(sk, t, \mathbf{c})$**:** Given the secret key $sk = (K, \mathbf{S})$ and the tag $t = (t_c, t_a) \in \mathcal{T}$, compute $C_{\mathsf{PRF},t}(K) \in \{0,1\}$ and return $\perp$ if $C_{\mathsf{PRF},t}(K) = 1$. Otherwise, Compute and return $\mathsf{Msg} = \mathsf{Msg}[1]\ldots\mathsf{Msg}[n_0]$ exactly as in Section 3.1.

Opener$(pk, tk, t, \mathbf{c}, \mathsf{Msg}_1)$**:** Given $tk = (K, \{\mathbf{R}_i\}_i)$ and $t = (t_c, t_a) \in \mathcal{T}$, compute $C_{\mathsf{PRF},t}(K) \in \{0,1\}$ and return $\perp$ if $C_{\mathsf{PRF},t}(K) = 0$. Otherwise,

1. Compute the matrix $\mathbf{R}_t = \mathsf{Eval}_{\mathsf{BP}}^{\mathsf{priv}}(C_{\mathsf{PRF},t}, (\mathbf{R}_i, k_i)_i) \in \mathbb{Z}^{m\times m}$ such that $\mathbf{A}_{F,t} = \mathbf{A}\cdot\mathbf{R}_t + \mathbf{G}$ and $\|\mathbf{R}_t\| \leq 4^d \cdot O(m^{3/2})$.
2. Use $\mathbf{R}_t \in \mathbb{Z}^{m\times m}$ as a trapdoor for $\bar{\mathbf{A}}_{F,t} = [\mathbf{A} \mid \mathbf{A}_{F,t}] = [\mathbf{A} \mid \mathbf{A}\cdot\mathbf{R}_t + \mathbf{G}]$ to sample $\bar{\mathbf{r}} \in \mathbb{Z}^{2m}$ such that $\bar{\mathbf{A}}_{F,t}\cdot\bar{\mathbf{r}} = \mathbf{c} - [\mathbf{0}^{n-n_0\top} \mid \mathsf{Msg}_1\cdot\lfloor q/2\rfloor]^\top$. Namely, defining $\mathbf{c}_{\mathsf{Msg}_1} = \mathbf{c} - [(\mathbf{0}^{n-n_0})^\top \mid \mathsf{Msg}_1^\top\cdot\lfloor q/2\rfloor]^\top$, sample and output fake random coins $\bar{\mathbf{r}} \hookleftarrow D_{\Lambda_q^{\mathbf{c}_{\mathsf{Msg}_1}}(\bar{\mathbf{A}}_{F,t}),\sigma}$.

LOpener$(sk, t, \mathbf{c}, \mathsf{Msg}_1)$**:** Given $sk = \mathbf{T}_{\mathbf{A}}$ and $t = (t_c, t_a) \in \mathcal{T}$, use $\mathbf{T}_{\mathbf{A}}$ to derive a trapdoor $\mathbf{T}_{\mathbf{A},t}$ for the lattice $\Lambda_q^{\perp}(\bar{\mathbf{A}}_{F,t})$ and use $\mathbf{T}_{\mathbf{A},t}$ to sample a Gaussian vector $\bar{\mathbf{r}} \hookleftarrow D_{\Lambda_q^{\mathbf{c}_{\mathsf{Msg}_1}}(\bar{\mathbf{A}}_{F,t}),\sigma}$ in the same coset of $\Lambda_q^{\perp}(\bar{\mathbf{A}}_{F,t})$ as in Opener.

The proof of Theorem 4.3 is identical to that of Theorem 3.2 and omitted.

Theorem 4.3. *The above construction is an $\mathcal{R}_{\mathsf{PRF}}$-lossy public-key encryption scheme with efficient opening under the* LWE *assumption.*

4.2 Unbounded Simulation-Sound Argument

We construct a tightly secure USS argument from the following ingredients:

- A pseudorandom function family $\mathsf{PRF} : \mathcal{K} \times \{0,1\}^{\ell} \to \{0,1\}^{\lambda}$ with key space $\mathcal{K} = \{0,1\}^{\lambda}$ and input space $\{0,1\}^{\ell}$, which induces the relation $\mathcal{R}_{\mathsf{PRF}} : \mathcal{K} \times \mathcal{T} \to \{0,1\}$ of Definition 4.1.
- An $\mathcal{R}_{\mathsf{PRF}}$-lossy PKE scheme $\mathcal{R}$-$\mathsf{LPKE} = (\mathsf{Par\text{-}Gen}, \mathsf{Keygen}, \mathsf{LKeygen}, \mathsf{Encrypt}, \mathsf{Decrypt}, \mathsf{Opener}, \mathsf{LOpener})$ for the relation $\mathcal{R}_{\mathsf{PRF}} : \mathcal{K} \times \mathcal{T} \to \{0,1\}$ with public (resp. secret) key space $\mathcal{PK}$ (resp. $\mathcal{SK}$). We assume that $\mathsf{Decrypt}$ is computable within time T. We denote the message (resp. ciphertext) space by MsgSp (resp. CtSp) and the randomness space by R^{LPKE}. Let also D_R^{LPKE} denote the distribution of the random coins of $\mathsf{Encrypt}$.
- A trapdoor Σ-protocol $\Pi' = (\mathsf{Gen}'_{\mathsf{par}}, \mathsf{Gen}'_{\mathcal{L}}, \mathsf{P}', \mathsf{V}')$, a one-time signature scheme $\mathsf{OTS} = (\mathcal{G}, \mathcal{S}, \mathcal{V})$ and a correlation intractable hash family $\mathcal{H} = (\mathsf{Gen}, \mathsf{Hash})$ that satisfy the same conditions as in Section 3.2.

Our construction $\Pi^{\mathrm{uss}} = (\mathsf{Gen}_{\mathsf{par}}, \mathsf{Gen}_{\mathcal{L}}, \mathsf{P}, \mathsf{V})$ goes as follows.

Gen$_{\mathsf{par}}(1^{\lambda})$**:** Run $\mathsf{par} \leftarrow \mathsf{Gen}'_{\mathsf{par}}(1^{\lambda})$ and output par.

Gen$_{\mathcal{L}}(\mathsf{par}, \mathcal{L})$**:** Given public parameters par and a language $\mathcal{L} \subset \{0,1\}^N$, let $\mathcal{K} = \{0,1\}^{\lambda}$ and $\mathcal{T} = \{0,1\}^{\ell}$. The CRS is generated as follows.

1. Generate a CRS $\mathsf{crs}'_{\mathcal{L}} \leftarrow \mathsf{Gen}'_{\mathcal{L}}(\mathsf{par}, \mathcal{L})$ for the trapdoor Σ-protocol Π'.
2. Generate public parameters $\Gamma \hookleftarrow \mathsf{Par\text{-}Gen}(1^{\lambda})$ for the $\mathcal{R}_{\mathsf{PRF}}$-lossy PKE scheme where the relation $\mathcal{R}_{\mathsf{PRF}} : \mathcal{K} \times \mathcal{T} \to \{0,1\}$ is defined by a PRF family $\mathsf{PRF} : \mathcal{K} \times \{0,1\}^{\ell} \to \{0,1\}^{\lambda}$. Generate lossy keys $(pk, sk, tk) \leftarrow \mathsf{LKeygen}(\Gamma, \mathbf{0}^{\lambda})$, where the initialization value is the all-zeroes string $\mathbf{0}^{\lambda}$.
3. Generate a key $k \leftarrow \mathsf{Gen}(1^{\lambda})$ for a correlation intractable hash function with output length $\kappa = \Theta(\lambda)$.

Output the language-dependent $\mathsf{crs}_{\mathcal{L}} := (\mathsf{crs}'_{\mathcal{L}}, k)$ and the simulation trapdoor $\tau_{\mathrm{zk}} := sk$. The global CRS consists of $\mathsf{crs} = (\mathsf{par}, \mathsf{crs}_{\mathcal{L}}, pk, \mathsf{PRF}, \mathsf{OTS})$.

P$(\mathsf{crs}, x, w, \mathsf{lbl})$: To prove x with respect to a label lbl using $w \in R_{\mathsf{zk}}(x)$, generate a one-time signature key pair $(\mathsf{VK}, \mathsf{SK}) \leftarrow \mathcal{G}(1^{\lambda})$. Then, choose a random core tag component $t_c \hookleftarrow U(\{0,1\}^{\lambda})$ and do the following.

1. Compute $\big(\mathbf{a}' = (\mathbf{a}'_1, \ldots, \mathbf{a}'_\kappa), st'\big) \leftarrow \mathsf{P}'(\mathsf{crs}'_{\mathcal{L}}, x, w)$ via κ invocations of the prover for Π'. For each $i \in [\kappa]$, compute $\mathbf{a}_i \leftarrow \mathsf{Encrypt}(pk, (t_c, \mathsf{VK}), \mathbf{a}'_i; \mathbf{r}_i)$ using random coins $\mathbf{r}_i \hookleftarrow D_R^{\mathsf{LPKE}}$. Let $\mathbf{a} = (\mathbf{a}_1, \ldots, \mathbf{a}_\kappa)$ and $\mathbf{r} = (\mathbf{r}_1, \ldots, \mathbf{r}_\kappa)$.
2. Compute $\mathsf{Chall} = \mathsf{Hash}(k, (x, \mathbf{a}, t_c, \mathsf{VK})) \in \{0,1\}^\kappa$.
3. Compute $\mathbf{z}' = (\mathbf{z}'_1, \ldots, \mathbf{z}'_\kappa) = \mathsf{P}'(\mathsf{crs}'_{\mathcal{L}}, x, w, \mathbf{a}', \mathsf{Chall}, st')$ via κ executions of the prover of Π'. Define $\mathbf{z} = (\mathbf{z}', \mathbf{a}', \mathbf{r})$.
4. Generate a one-time signature $sig \leftarrow \mathcal{S}(\mathsf{SK}, (x, t_c, \mathbf{a}, \mathbf{z}, \mathsf{lbl}))$ and output the proof $\boldsymbol{\pi} = \big((t_c, \mathsf{VK}), (\mathbf{a}, \mathbf{z}), sig\big)$.

$\mathbf{V}(\mathsf{crs}, x, \boldsymbol{\pi}, \mathsf{lbl})$: Given a statement x, a label lbl and a candidate proof $\boldsymbol{\pi} = \big((t_c, \mathsf{VK}), (\mathbf{a}, \mathbf{z}), sig\big)$, return 0 if $\mathcal{V}(\mathsf{VK}, (x, t_c, \mathbf{a}, \mathbf{z}, \mathsf{lbl}), sig) = 0$. Otherwise,

1. Write $\mathbf{z}$ as $\mathbf{z} = \big((\mathbf{z}'_1, \ldots, \mathbf{z}'_\kappa), (\mathbf{a}'_1, \ldots, \mathbf{a}'_\kappa), (\mathbf{r}_1, \ldots, \mathbf{r}_\kappa)\big)$. Return 0 if there exists $i \in [\kappa]$ such that $\mathbf{a}_i \neq \mathsf{Encrypt}(pk, (t_c, \mathsf{VK}), \mathbf{a}'_i; \mathbf{r}_i)$ or $\mathbf{r}_i \notin R^{\mathsf{LPKE}}$.
2. Let $\mathsf{Chall} = \mathsf{Hash}\big(k, (x, (\mathbf{a}_1, \ldots, \mathbf{a}_\kappa), t_c, \mathsf{VK})\big)$. If there exists $i \in [\kappa]$ such that $\mathsf{V}'(\mathsf{crs}'_{\mathcal{L}}, x, (\mathbf{a}'_i, \mathsf{Chall}[i], \mathbf{z}'_i)) = 0$, return 0. Otherwise, return 1.

In the full version of the paper, we show that the unbounded simulation-soundness of the above argument system is tightly related to the security of its underlying building blocks, which are all instantiable (with tight security reductions) from LWE.

5 Trapdoor Σ-Protocols for ACPS Ciphertexts

The KDM-CPA system of Applebaum *et al.* [6] uses a modulus $q = p^2$, for some prime p. Its public key $(\mathbf{A}, \mathbf{b}) \in \mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^m$ contains a random matrix $\mathbf{A} \sim U(\mathbb{Z}_q^{n \times m})$ and a vector $\mathbf{b} = \mathbf{A}^\top \cdot \mathbf{s} + \mathbf{e}$, for some $\mathbf{s} \sim D_{\mathbb{Z}^n, \alpha q}$, $\mathbf{e} \sim D_{\mathbb{Z}^m, \alpha q}$. Its encryption algorithm proceeds analogously to the primal Regev cryptosystem [79] and computes $\mathbf{c} = (\bar{\mathbf{c}}, c) = (\mathbf{A} \cdot \mathbf{r}, \mathbf{b}^\top \mathbf{r} + \mu \cdot p + \chi) \in \mathbb{Z}_q^{n+1}$, where $\mathbf{r} \sim D_{\mathbb{Z}^m, r}$ is a Gaussian vector and $\chi \sim D_{\mathbb{Z}, r'}$ is sampled from a Gaussian with a slightly larger standard deviation. Decryption proceeds by rounding $c - \mathbf{s}^\top \cdot \bar{\mathbf{c}} \mod q$ to the nearest multiple of p.

In this section, we describe a trapdoor Σ-protocol allowing to prove that two ACPS ciphertexts $\mathbf{c}_0 = (\bar{\mathbf{c}}_0, c_0)$, $\mathbf{c}_1 = (\bar{\mathbf{c}}_1, c_1)$ are both encryptions of the same $\mu \in \mathbb{Z}_p$. This protocol is obtained by extending a simpler protocol (described in the full version of the paper), which argues that a given vector $\mathbf{c} \in \mathbb{Z}_q^{n+1}$ is an ACPS encryption of some plaintext $\mu \in \mathbb{Z}_p$.

We note that Ciampi *et al.* [33] recently gave a construction of trapdoor Σ-protocol from any Σ-protocol. The Σ-protocol described hereunder is natively trapdoor without applying the transformation of [33].

Proving Plaintext Equalities in ACPS Ciphertexts. Let $q = p^2$, for some prime p, and a matrix $\mathbf{A}$ which is used to set up two Regev public keys $(\mathbf{A}, \mathbf{b}_0) \in \mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^m$ and $(\mathbf{A}, \mathbf{b}_1) \in \mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^m$, where $\mathbf{b}_0 = \mathbf{A}^\top \cdot \mathbf{s}_0 + \mathbf{e}_0$ and $\mathbf{b}_1 = \mathbf{A}^\top \cdot \mathbf{s}_1 + \mathbf{e}_1$ for some $\mathbf{s}_0, \mathbf{s}_1 \sim D_{\mathbb{Z}^n, \alpha q}$, $\mathbf{e}_0, \mathbf{e}_1 \sim D_{\mathbb{Z}^m, \alpha q}$. Let also the matrix

$$\mathbf{A}_{\mathrm{eq}} = \left[\begin{array}{c|c|c|c} \mathbf{A} & & & \\ \hline \mathbf{b}_0^\top & 1 & & \\ \hline & & \mathbf{A} & \\ \hline & & \mathbf{b}_1^\top & 1 \end{array}\right] \in \mathbb{Z}_q^{2(n+1) \times 2(m+1)}, \tag{4}$$

We give a trapdoor Σ-protocol for the language $\mathcal{L}^{\text{eq}} = (\mathcal{L}^{\text{eq}}_{\mathsf{zk}}, \mathcal{L}^{\text{eq}}_{\mathsf{sound}})$, where

$$\mathcal{L}^{\text{eq}}_{\mathsf{zk}} := \Big\{(\mathbf{c}_0, \mathbf{c}_1) \in (\mathbb{Z}_q^{n+1})^2 \;\; | \; \exists \mathbf{r}_0, \mathbf{r}_1 \in \mathbb{Z}^m,\; \chi_0, \chi_1 \in \mathbb{Z},\; \mu \in \mathbb{Z}_p \;\; : \\ \|\mathbf{r}_b\| \le B_r,\; |\chi_b| \le B_\chi \quad \forall b \in \{0,1\} \\ \wedge \quad \mathbf{c}_b = \bar{\mathbf{A}}_b \cdot [\mathbf{r}_b^\top \mid \chi_b]^\top + \mu \cdot \big[\mathbf{0}^{n\,\top} \mid p\big]^\top \bmod q \;\Big\},$$

$$\mathcal{L}^{\text{eq}}_{\mathsf{sound}} := \Big\{(\mathbf{c}_0, \mathbf{c}_1) \in (\mathbb{Z}_q^{n+1})^2 \;\; | \;\; \exists\; \bar{\mathbf{c}}_0, \bar{\mathbf{c}}_1 \in \mathbb{Z}_q^n,\; v_0, v_1 \in [-B^*, B^*],\; \mu \in \mathbb{Z}_p \\ \wedge \;\; \mathbf{c}_b = \left\lceil \frac{\bar{\mathbf{c}}_b}{\mathbf{s}_b^\top \cdot \bar{\mathbf{c}}_b + p \cdot \mu + v_b} \right\rceil \quad \forall b \in \{0,1\} \Big\},$$

where

$$\bar{\mathbf{A}}_b = \left[\begin{array}{c|c} \mathbf{A} & \\ \hline \mathbf{b}_b^\top & 1 \end{array}\right] \in \mathbb{Z}_q^{(n+1)\times(m+1)} \qquad\qquad \forall b \in \{0,1\}.$$

We note that $\mathcal{L}^{\text{eq}}_{\mathsf{zk}} \subseteq \mathcal{L}^{\text{eq}}_{\mathsf{sound}}$ when $B_r \alpha q \sqrt{m} + B_\chi < B^* \ll p$. Also, $\mathcal{L}^{\text{eq}}_{\mathsf{sound}}$ is equivalently defined as the language of pairs $(\mathbf{c}_0, \mathbf{c}_1)$ such that such that

$$\left[\begin{array}{c|c|c|c} -\mathbf{s}_0^\top & 1 & & \\ \hline & & -\mathbf{s}_1^\top & 1 \end{array}\right] \cdot \left[\begin{array}{c} \mathbf{c}_0 \\ \hline \mathbf{c}_1 \end{array}\right] \bmod q \;=\; \left[\begin{array}{c} v_0 \\ \hline v_1 \end{array}\right] + \mu \cdot \left[\begin{array}{c} p \\ \hline p \end{array}\right]$$

for some $\mu \in \mathbb{Z}_p$, $v_0, v_1 \in [-B^*, B^*]$.

$\mathsf{Gen}_{\mathsf{par}}(1^\lambda)$: On input of a security parameter $\lambda \in \mathbb{N}$, choose moduli q, p with $q = p^2$, dimensions n, m, and error rate $\alpha > 0$ and a Gaussian parameter $\sigma_{\text{eq}} \ge \log(2m+2) \cdot \sqrt{B_r^2 + B_\chi^2}$. Define public parameters $\mathsf{par} = \{\lambda, q, p, n, m, \alpha, \sigma_{\text{eq}}\}$.

$\mathsf{Gen}_{\mathcal{L}}(\mathsf{par}, \mathcal{L}^{\text{eq}})$: Takes in global parameters par and the description of a language $\mathcal{L}^{\text{eq}} = (\mathcal{L}^{\text{eq}}_{\mathsf{zk}}, \mathcal{L}^{\text{eq}}_{\mathsf{sound}})$ specifying real numbers $B^*, B_r, B_\chi > 0$ such that $B_r \alpha q \sqrt{m} + B_\chi < B^* \ll p$, and a matrix $\mathbf{A}_{\text{eq}}$ from the distribution (4). It defines the language-dependent $\mathsf{crs}_{\mathcal{L}} = \{\bar{\mathbf{A}}, B^*, B_r, B_\chi\}$. The global CRS is

$$\mathsf{crs} = \big(\{\lambda, q, p, n, m, \alpha, \sigma_{\text{eq}}\}, \{\mathbf{A}_{\text{eq}}, B^*, B_r, B_\chi\}\big).$$

$\mathsf{TrapGen}(\mathsf{par}, \mathcal{L}, \tau_{\mathcal{L}})$: Given par and a language description $\mathcal{L}^{\text{eq}}$ that specifies $B^*, B_r, B_\chi > 0$ satisfying the same constraints as in $\mathsf{Gen}_{\mathcal{L}}$, a matrix $\mathbf{A}_{\text{eq}}$ sampled from the distribution (4), as well as a membership-testing trapdoor $\tau_{\mathcal{L}} = (\mathbf{s}_0, \mathbf{s}_1) \sim (D_{\mathbb{Z}^n, \alpha q})^2$ for $\mathcal{L}^{\text{eq}}_{\mathsf{sound}}$, output $\mathsf{crs}_{\mathcal{L}} = \{\bar{\mathbf{A}}, B^*, B_r, B_\chi\}$. The global CRS is $\mathsf{crs} = \big(\{\lambda, q, p, n, m, \alpha, \sigma_{\text{eq}}\}, \{\mathbf{A}_{\text{eq}}, B^*, B_r, B_\chi\}\big)$ and the trapdoor $\tau_\Sigma = (\mathbf{s}_0, \mathbf{s}_1) \in \mathbb{Z}^n \times \mathbb{Z}^n$.

$\mathsf{P}\big(\mathsf{crs}, (\mathbf{c}_0, \mathbf{c}_1), (\mu, \mathbf{w})\big) \leftrightarrow \mathsf{V}(\mathsf{crs}, \mathbf{x})$: Given crs and a statement

$$\left[\begin{array}{c} \mathbf{c}_0 \\ \hline \mathbf{c}_1 \end{array}\right] = \mathbf{A}_{\text{eq}} \cdot [\mathbf{r}_0{}^\top \mid \chi_0 \mid \mathbf{r}_1{}^\top \mid \chi_1]^\top + \mu \cdot [\mathbf{0}^{n\,\top} \mid p \mid \mathbf{0}^{n\,\top} \mid p]^\top \in \mathbb{Z}_q^{2(n+1)},$$

the prover P (who has $\mathbf{w} = [\mathbf{r}_0^\top \mid \chi_0 \mid \mathbf{r}_1^\top \mid \chi_1]^\top \in \mathbb{Z}^{2(m+1)}$ and $\mu \in \mathbb{Z}_p$) and the verifier V interact as follows.

1. The prover P samples a uniform scalar $r_\mu \hookleftarrow U(\mathbb{Z}_p)$ and Gaussian vector $\mathbf{r}_w \hookleftarrow D_{\mathbb{Z}^{2(m+1)},\sigma_{\text{eq}}}$. It computes the following which is sent to V:
$$\mathbf{a} = \mathbf{A}_{\text{eq}} \cdot \mathbf{r}_w + r_\mu \cdot [\mathbf{0}^{n\top} \mid p \mid \mathbf{0}^{n\top} \mid p]^\top \in \mathbb{Z}_q^{2(n+1)}.$$
2. V sends a random challenge $\mathsf{Chall} \in \{0,1\}$ to P.
3. P computes $\mathbf{z} = \mathbf{r}_w + \mathsf{Chall} \cdot \mathbf{w} \in \mathbb{Z}^{2(m+1)}$, $z_\mu = r_\mu + \mathsf{Chall} \cdot \mu \mod p$. It sends $(\mathbf{z}, z_\mu)$ to V with probability $\theta = \min\left(\frac{D_{\mathbb{Z}^{2(m+1)},\sigma_{\text{eq}}}(\mathbf{z})}{M \cdot D_{\mathbb{Z}^{2(m+1)},\sigma_{\text{eq}},\mathsf{Chall}\cdot\mathbf{w}}(\mathbf{z})}, 1\right)$, where $M = e^{12/\log(2(m+1))+1/(2\log^2(2(m+1)))}$. With probability $1-\theta$, P aborts.
4. Given $(\mathbf{z}, z_\mu) \in \mathbb{Z}^{2(m+1)} \times \mathbb{Z}_p$, V checks if $\|\mathbf{z}\| \le \sigma_{\text{eq}}\sqrt{2(m+1)}$ and
$$\mathbf{a} + \mathsf{Chall} \cdot \begin{bmatrix}\mathbf{c}_0\\ \mathbf{c}_1\end{bmatrix} = \mathbf{A}_{\text{eq}} \cdot \mathbf{z} + z_\mu \cdot [\mathbf{0}^{n\top} \mid p \mid \mathbf{0}^{n\top} \mid p]^\top \bmod q. \tag{5}$$
If these conditions do not both hold, V halts and returns $\perp$.

$\mathsf{BadChallenge}\big(\mathsf{par}, \tau_\Sigma, \mathsf{crs}, (\mathbf{c}_0, \mathbf{c}_1), \mathbf{a}\big)$: Given $\tau_\Sigma = (\mathbf{s}_0, \mathbf{s}_1) \in \mathbb{Z}^n \times \mathbb{Z}^n$, parse the first prover message as $\mathbf{a} = (\mathbf{a}_0^\top \mid \mathbf{a}_1^\top)^\top \in \mathbb{Z}_q^{2(n+1)}$. If there exists $d \in \{0,1\}$ such that no pair $(\mu'_d, \mathbf{v}_d) \in [-(p-1)/2, (p-1)/2] \times [-B^*/2, B^*/2]^2$ satisfies
$$\left[\begin{array}{c|c|c|c} -\mathbf{s}_0^\top & 1 & & \\ \hline & & -\mathbf{s}_1^\top & 1 \end{array}\right] \cdot (\mathbf{a} + d \cdot \mathbf{c}) \bmod q \;=\; \mathbf{v}_d + \mu'_d \cdot \begin{bmatrix}p\\ p\end{bmatrix} \tag{6}$$
over $\mathbb{Z}$, then return $\mathsf{Chall} = 1 - d$. Otherwise, return $\mathsf{Chall} = \perp$.

The completeness of the protocol crucially uses the fact that p divides q to ensure that the response $z_\mu = r_\mu + \mathsf{Chall} \cdot \mu \bmod p$ satisfies (5).

The intuition of $\mathsf{BadChallenge}$ is that, for a false statement $(\mathbf{c}_0, \mathbf{c}_1) \notin \mathcal{L}^{\text{eq}}_{\mathsf{sound}}$, there exists $d \in \{0,1\}$ such that no pair $(\mu'_d, \mathbf{v}_d)$ satisfies (6) for a small enough $\mathbf{v}_d \in \mathbb{Z}^2$. Moreover, for this challenge $\mathsf{Chall} = d$, no valid response can exist, as shown in the proof of Lemma 5.1. We note that $\mathsf{BadChallenge}$ may output a bit even when there is no bad challenge at all for a given $\mathbf{a}$. These "false positives" are not a problem since, in order to soundly instantiate Fiat-Shamir, we only need the somewhere CI hash function to avoid the bad challenge when it exists.

Lemma 5.1. *The above construction is a trapdoor Σ-protocol for $\mathcal{L}^{\text{eq}}$ if we set $\sigma_{\text{eq}} \ge \log(2m+2) \cdot \sqrt{B_r^2 + B_\chi^2}$ and*
$$B^* > \max(2\sigma_{\text{eq}}\sqrt{2m+2} \cdot (\alpha q\sqrt{m} + 1), B_r \alpha q\sqrt{m} + B_\chi).$$
(The proof is given in the full version of the paper.)

Parallel Repetitions. To achieve negligible soundness error, the protocol is repeated $\kappa = \Theta(\lambda)$ times in parallel by first computing $(\mathbf{a}_1, \ldots, \mathbf{a}_\kappa)$ before obtaining $\mathsf{Chall} = \mathsf{Chall}[1] \ldots \mathsf{Chall}[\kappa]$ and computing the response $\bar{\mathbf{z}} = (\mathbf{z}_1, \ldots, \mathbf{z}_\kappa)$, $(z_{\mu,1}, \ldots, z_{\mu,\kappa})$. We then handle $\bar{\mathbf{z}}$ as an integer vector in $\mathbb{Z}^{\kappa\cdot(m+1)}$ and reject it with probability $\theta = \min\big(1, D_{\mathbb{Z}^{2\kappa\cdot(m+1)},\sigma_{\text{eq}}}(\mathbf{z})/M \cdot D_{\mathbb{Z}^{2\kappa\cdot(m+1)},\sigma_{\text{eq}},\mathsf{Chall}\cdot(\mathbf{1}^\kappa\otimes\mathbf{w})}(\mathbf{z})\big)$, where $M = e^{12/\log(2\kappa\cdot(m+1))+1/(2\log^2(2\kappa\cdot(m+1)))}$. Then, we need to slightly increase σ_{eq} and set $\sigma_{\text{eq}} \ge \log(2\kappa(m+1)) \cdot \sqrt{\kappa(B_r^2 + B_\chi^2)}$.

Acknowledgements

Part of this research was funded by the French ANR ALAMBIC project (ANR-16-CE39-0006). This work was also supported in part by the European Union PROMETHEUS project (Horizon 2020 Research and Innovation Program, grant 780701). Khoa Nguyen was supported in part by the Gopalakrishnan - NTU PPF 2018, by A*STAR, Singapore under research grant SERC A19E3b0099, and by Vietnam National University HoChiMinh City (VNU-HCM) under grant number NCM2019-18-01.

References

1. M. Abdalla, F. Bourse, A. De Caro, and D. Pointcheval. Simple functional encryption schemes for inner products. In *PKC*, 2015.
2. S. Agrawal, D. Boneh, and X. Boyen. Efficient lattice (H)IBE in the standard model. In *Eurocrypt*, 2010.
3. S. Agrawal, B. Libert, and D. Stehlé. Fully secure functional encryption for inner products from standard assumptions. In *Crypto*, 2016.
4. J. Alperin-Sheriff and C. Peikert. Circular and KDM security for identity-based encryption. In *PKC*, 2012.
5. B. Applebaum. Key-dependent message security: Generic amplification and completeness theorems. *J. of Cryptology*, 27(3), 2013.
6. B. Applebaum, D. Cash, C. Peikert, and A. Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In *Crypto*, 2009.
7. G. Asharov, A. Jain, and D. Wichs. Multiparty computation with low communication, computation and interaction via threshold FHE. Cryptology ePrint Archive: Report 2011/613, 2012.
8. B. Barak, I. Haitner, D. Hofheinz, and Y. Ishai. Bounded key-dependent message security. In *Eurocrypt*, 2010.
9. D. Barrington. Bounded-width polynomial-size branching programs recognize exactly those languages in nc1. In *STOC*, 1986.
10. M. Bellare, D. Hofheinz, and S. Yilek. Possibility and impossibility results for encryption and commitment secure under selective opening. In *Eurocrypt*, 2009.
11. M. Bellare and P. Rogaway. Random oracles are practical: a paradigm for designing efficient protocols. In *ACM-CCS*, 1993.
12. N. Bitansky, D. Dachman-Soled, S. Garg, A. Jain, T. Tauman Kalai, A. Lopez-Alt, and D. Wichs. Why "Fiat-Shamir for proofs" lacks a proof. In *TCC*, 2013.
13. J. Black, P. Rogaway, and T. Shrimpton. Encryption-scheme security in the presence of key-dependent messages. In *SAC*, 2002.
14. D. Boneh and X. Boyen. Secure identity based encryption without random oracles. In *Crypto*, 2004.
15. D. Boneh, C. Gentry, S. Gorbunov, S. Halevi, V. Nikolaenko, G. Segev, V. Vaikuntanathan, and D. Vinayagamurthy. Fully key-homomorphic encryption, arithmetic circuit ABE and compact garbled circuits. In *Eurocrypt*, 2014.
16. D. Boneh, S. Halevi, M. Hamburg, and R. Ostrovsky. Circular-secure encryption from Decision Diffie-Hellman. In *Crypto*, 2008.
17. X. Boyen and Q. Li. Towards tightly secure lattice short signature and ID-based encryption. In *Asiacrypt*, 2016.

18. X. Boyen and Q. Li. Almost tight multi-instance multi-ciphertext identity-based encryption on lattices. In *ACNS*, 2018.
19. E. Boyle, G. Segev, and D. Wichs. Fully leakage-resilient signatures. In *Eurocrypt*, 2011.
20. Z. Brakerski and S. Goldwasser. Circular and leakage resilient public-key encryption under subgroup indistinguishability (or: Quadratic residuosity strikes back). In *Crypto*, 2010.
21. Z. Brakerski, S. Goldwasser, and Y. Tauman Kalai. Black-box circular-secure encryption beyond affine functions. In *TCC*, 2011.
22. Z. Brakerski and V. Vaikuntanathan. Lattice-based FHE as secure as PKE. In *ITCS*, 2014.
23. J. Camenisch and A. Lysyanskaya. An efficient system for non-transferable anonymous credentials with optional anonymity revocation. In *Eurocrypt*, 2001.
24. J. Camensich, N. Chandran, and V. Shoup. A public key encryption scheme secure against key dependent chosen plaintext and adaptive chosen ciphertext attacks. In *Eurocrypt*, 2009.
25. R. Canetti, Y. Chen, J. Holmgren, A. Lombardi, G. Rothblum, and R. Rothblum. Fiat-Shamir from simpler assumptions. Cryptology ePrint Archive: Report 2018/1004.
26. R. Canetti, Y. Chen, J. Holmgren, A. Lombardi, G. Rothblum, R. Rothblum, and D. Wichs. Fiat-Shamir: From practice to theory. In *STOC*, 2019.
27. R. Canetti, Y. Chen, L. Reyzin, and R. Rothblum. Fiat-Shamir and correlation intractability from strong KDM-secure encryption. In *Eurocrypt*, 2018.
28. R. Canetti, O. Goldreich, and S. Halevi. The random oracle methodology, revisted. *J. of the ACM*, 51(4), 2004.
29. R. Canetti, S. Halevi, and J. Katz. Chosen-ciphertext security from identity-based encryption. In *Eurocrypt*, 2004.
30. R. Canetti, A. Lombardi, and D. Wichs. Fiat-Shamir: From Practice to Theory, Part II (NIZK and Correlation Intractability from Circular-Secure FHE). Cryptology ePrint Archive: Report 2018/1248.
31. M. Chase and A. Lysyanskaya. Simulatable VRFs with applications to multi-theorem NIZK. In *Crypto*, 2007.
32. A. Choudhuri, P. Hubacek, K. C., K. Pietrzak, A. Rosen, and G. Rothblum. Finding a Nash equilibrium is no easier than breaking Fiat-Shamir. In *STOC*, 2019.
33. M. Ciampi, R. Parisella, and D. Ventury. On adaptive security of delayed-input sigma protocols and Fiat-Shamir NIZKs. In *SCN*, 2020.
34. G. Couteau and D. Hofheinz. Designated-verifier pseudorandom generators, and their applications. In *Eurocrypt*, 2019.
35. R. Cramer. Modular design of secure, yet practical cryptographic protocols. PhD thesis, University of Amsterdam, 1996.
36. R. Cramer, I. Damgård, and B. Schoenmaekers. Proofs of partial knowledge and simplified design of witness hiding protocols. In *Crypto*, 1994.
37. R. Cramer and V. Shoup. Universal hash proofs and a paradigm for adaptive chosen ciphertext secure public-key encryption. In *Eurocrypt*, 2002.
38. I. Damgård. Efficient concurrent zero-knowledge in the auxiliary string model. In *Eurocrypt*, 2000.
39. A. De Santis, G. Di Crescenzo, R. Ostrovsky, G. Persiano, and A. Sahai. Robust non-interactive zero-knowledge. In *Crypto*, 2001.
40. A. De Santis and M. Yung. Cryptograpic applications of the non-interactive metaproof and many-prover systems. In *Crypto*, 1990.

41. N. Döttling. Low-noise LPN: KDM secure public key encryption and sample amplification. In *PKC*, 2015.
42. U. Feige, D. Lapidot, and A. Shamir. Multiple non-interactive zero-knowledge under general assumptions. *SIAM J. of Computing*, 29(1), 1999.
43. A. Fiat and A. Shamir. How to prove yourself: Practical solutions to identification and signature problems. In *Crypto*, 1986.
44. P.-A. Fouque and D. Pointcheval. Threshold Cryptosystems Secure against Chosen-Ciphertext Attacks. In *Asiacrypt*, 2001.
45. E. Freire, D. Hofheinz, K. Paterson, and C. Striecks. Programmable hash functions in the multilinear setting. In *Crypto*, 2013.
46. C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, 2008.
47. C. Gentry, A. Sahai, and B. Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In *Crypto*, 2013.
48. S. Goldwasser, S. Micali, and C. Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, 1989.
49. S. Goldwasser and Y. Tauman Kalai. On the (in)security of the Fiat-Shamir paradigm. In *FOCS*, 2003.
50. J. Groth, R. Ostrovsky, and A. Sahai. New techniques for noninteractive zero-knowledge. *J. ACM*, 2012.
51. J. Groth and A. Sahai. Efficient non-interactive proof systems for bilinear groups. In *Eurocrypt*, 2008.
52. S. Halevi, S. Myers, and C. Rackoff. On seed-incompressible functions. In *TCC*, 2008.
53. S. Han, S. Liu, and L. Lyu. Efficient KDM-CCA secure public-key encryption for polynomial functions. In *Asiacrypt*, 2016.
54. K. Hara, K. F., T. Matsuda, G. Hanaoka, and K. Tanaka. Simulation-based receiver selective opening cca secure pke from standard computational assumptions. In *SCN*, 2018.
55. D. Hofheinz. All-but-many lossy trapdoor functions. In *Eurocrypt*, 2012.
56. D. Hofheinz. Circular chosen-ciphertext security with compact ciphertexts. In *Eurocrypt*, 2013.
57. D. Hofheinz and T. Jager. Tightly secure signatures and public-key encryption. In *Crypto*, 2012.
58. J. Holmgren and A. Lombardi. Cryptographic hashing from strong one-way functions (or: One-way product functions and their applications). In *FOCS*, 2018.
59. T. Jager. Verifiable random functions from weaker assumptions. In *TCC*, 2015.
60. E. Kiltz. Chosen-ciphertext security from tag-based encryption. In *TCC*, 2006.
61. F. Kitagawa and T. Matsuda. CPA-to-CCA transformation for KDM security. In *TCC*, 2019.
62. F. Kitagawa, T. Matsuda, and K. Tanaka. Simple and efficient KDM-CCA secure public key encryption. In *Asiacrypt*, 2019.
63. F. Kitagawa and K. Tanaka. A framework for achieving KDM-CCA secure public-key encryption. In *Asiacrypt*, 2018.
64. Q. Lai, F. Liu, and Z. Wang. Almost tight security in lattices with polynomial moduli - PRF, IBE, all-but-many LTF, and more. In *PKC*, 2020.
65. B. Libert, K. Nguyen, A. Passelègue, and R. Titiu. Simulation-sound arguments for LWE and applications to KDM-CCA2 security. Cryptology ePrint Archive, Report 2019/908, 2020. https://eprint.iacr.org/2019/908.

66. B. Libert, T. Peters, M. Joye, and M. Yung. Non-malleability from malleability: Simulation-sound quasi-adaptive NIZK proofs and CCA2-secure encryption from homomorphic signatures. In *Eurocrypt*, 2014.
67. B. Libert, A. Sakzad, D. Stehlé, and R. Steinfeld. All-but-many lossy trapdoor functions and selective opening chosen-ciphertext security from LWE. In *Crypto*, 2017.
68. A. Lombardi, W. Quach, R. Rothblum, D. Wichs, and D. Wu. New constructions of reusable designated-verifier NIZKs. In *Crypto*, 2019.
69. X. Lu, B. Li, and D. Jia. KDM-CCA security from RKA secure authenticated encryption. In *Eurocrypt*, 2015.
70. V. Lyubashevsky. Lattice-Based Identification Schemes Secure Under Active Attacks. In *PKC*, 2008.
71. V. Lyubashevsky. Fiat-Shamir with aborts: Applications to lattice and factoring-based signatures. In *Asiacrypt*, 2009.
72. P. MacKenzie and K. Yang. On simulation-sound trapdoor commitments. In *Eurocrypt*, 2004.
73. T. Malkin, I. Teranishi, and M. Yung. Efficient circuit-size independent public key encryption with KDM security. In *Eurocrypt*, 2012.
74. D. Micciancio and C. Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *Eurocrypt*, 2012.
75. M. Naor and M. Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. In *STOC*, 1990.
76. C. Peikert and S. Shiehian. Non-interactive zero knowledge for NP from (plain) Learning With Errors. In *Crypto*, 2019.
77. C. Peikert, V. Vaikuntanathan, and B. Waters. A framework for efficient and composable oblivious transfer. In *Crypto*, 2008.
78. D. Pointcheval and J. Stern. Security arguments for digital signatures and blind signatures. *Jo. of Cryptology*, 13(3), 2000.
79. O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In *STOC*, 2005.
80. A. Sahai. Non-malleable non-interactive zero-knowledge and adaptive chosen-ciphertext security. In *FOCS*, 1999.
81. C.-P. Schnorr. Efficient identification and signatures for smart cards. In *Crypto*, 1989.
82. Y. Tauman Kalai, G. Rothblum, and R. Rothblum. From obfuscation to the security of Fiat-Shamir for proofs. In *Crypto*, 2017.
83. S. Yamada. Asymptotically compact adaptively secure lattice IBEs and verifiable random functions via generalized partitioning techniques. In *Crypto*, 2017.