

Tightly CCA-Secure Encryption Without Pairings

Romain Gay¹(✉), Dennis Hofheinz³, Eike Kiltz², and Hoeteck Wee¹

¹ ENS, Paris, France

{rgay, wee}@di.ens.fr

² Ruhr-Universität Bochum, Bochum, Germany

eike.kiltz@rub.de

³ Karlsruhe Institute of Technology, Karlsruhe, Germany

Dennis.Hofheinz@kit.edu

Abstract. We present the first CCA-secure public-key encryption scheme based on DDH where the security loss is independent of the number of challenge ciphertexts and the number of decryption queries. Our construction extends also to the standard k -Lin assumption in pairing-free groups, whereas all prior constructions starting with Hofheinz and Jager (Crypto '12) rely on the use of pairings. Moreover, our construction improves upon the concrete efficiency of existing schemes, reducing the ciphertext overhead by about half (to only 3 group elements under DDH), in addition to eliminating the use of pairings.

We also show how to use our techniques in the NIZK setting. Specifically, we construct the first tightly simulation-sound designated-verifier NIZK for linear languages without pairings. Using pairings, we can turn our construction into a highly optimized publicly verifiable NIZK with tight simulation-soundness.

1 Introduction

The most basic security guarantee we require of a public key encryption scheme is that of semantic security against chosen-plaintext attacks (CPA) [14]: it is infeasible to learn anything about the plaintext from the ciphertext. On the other hand, there is a general consensus within the cryptographic research community that in virtually every practical application, we require semantic security against adaptive chosen-ciphertext attacks (CCA) [12, 30], wherein an adversary is given access to decryptions of ciphertexts of her choice.

Romain Gay—CNRS, INRIA. Supported by ERC Project aSCEND (639554).

Dennis Hofheinz—Supported by DFG grants HO 4534/2-2, HO 4534/4-1.

Eike Kiltz—Partially supported by DFG grant KI 795/4-1 and ERC Project ERCC (FP7/615074).

Hoeteck Wee—CNRS (UMR 8548), INRIA and Columbia University. Partially supported by the Alexander von Humboldt Foundation, NSF Award CNS-1445424 and ERC Project aSCEND (639554).

In this work, we focus on the issue of security reduction and security loss in the construction of CPA and CCA-secure public-key encryption from the DDH assumption. Suppose we have such a scheme along with a security reduction showing that attacking the scheme in time t with success probability ϵ implies breaking the DDH assumption in time roughly t with success probability ϵ/L ; we refer to L as the security loss. In general, L would depend on the security parameter λ as well as the number of challenge ciphertexts Q_{enc} and the number decryption queries Q_{dec} , and we say that we have a *tight security reduction* if L depends only on the security parameter and is independent of both Q_{enc} and Q_{dec} . Note that for typical settings of parameters (e.g., $\lambda = 80$ and $Q_{\text{enc}}, Q_{\text{dec}} \approx 2^{20}$, or even $Q_{\text{enc}}, Q_{\text{dec}} \approx 2^{30}$ in truly large settings), λ is much smaller than Q_{enc} and Q_{dec} .

In the simpler setting of CPA-secure encryption, the ElGamal encryption scheme already has a tight security reduction to the DDH assumption [6, 27], thanks to random self-reducibility of DDH with a tight security reduction. In the case of CCA-secure encryption, the best result is still the seminal Cramer-Shoup encryption scheme [11], which achieves security loss Q_{enc} .¹ This raises the following open problem:

Does there exist a CCA-secure encryption scheme with a tight security reduction to the DDH assumption?

Hofheinz and Jager [16] gave an affirmative answer to this problem under stronger (and pairing-related) assumptions, notably the 2-Lin assumptions in bilinear groups, albeit with large ciphertexts and secret keys; a series of follow-up works [5, 15, 22, 24] leveraged techniques introduced in the context of tightly-secure IBE [7, 10, 18] to reduce the size of ciphertext and secret keys to a relatively small constant. However, all of these works rely crucially on the use of pairings, and seem to shed little insight on constructions under the standard DDH assumption; in fact, a pessimist may interpret the recent works as strong indication that the use of pairings is likely to be necessary for tightly CCA-secure encryption.

We may then restate the open problem as eliminating the use of pairings in these prior CCA-secure encryption schemes while still preserving a tight security reduction. From a theoretical stand-point, this is important because an affirmative answer would yield tightly CCA-secure encryption under qualitatively weaker assumptions, and in addition, shed insight into the broader question of whether tight security comes at the cost of qualitative stronger assumptions.

Eliminating the use of pairings is also important in practice as it allows us to instantiate the underlying assumption over a much larger class of groups that admit more efficient group operations and more compact representations, and also avoid the use of expensive pairing operations. Similarly, tight reductions matter in practice because as L increases, we should increase the size of the underlying groups in order to compensate for the security loss, which in turn

¹ We ignore contributions to the security loss that depend only on a statistical security parameter.

increases the running time of the implementation. Note that the impact on performance is quite substantial, as exponentiation in a r -bit group takes time roughly $\mathcal{O}(r^3)$.

1.1 Our Results

We settle the main open problem affirmatively: we construct a tightly CCA-secure encryption scheme from the DDH assumption without pairings. Moreover, our construction improves upon the concrete efficiency of existing schemes, reducing the ciphertext overhead by about half, in addition to eliminating the use of pairings. We refer to Fig. 2 for a comparison with prior works.

Overview of Our Construction. Fix an additively written group \mathbb{G} of order q . We rely on implicit representation notation [13] for group elements: for a fixed generator P of \mathbb{G} and for a matrix $\mathbf{M} \in \mathbb{Z}_q^{n \times t}$, we define $[\mathbf{M}] := \mathbf{M}P \in \mathbb{G}^{n \times t}$ where multiplication is done component-wise. We rely on the \mathcal{D}_k -MDDH Assumption [13], which stipulates that given $[\mathbf{M}]$ drawn from a matrix distribution \mathcal{D}_k over $\mathbb{Z}_q^{(k+1) \times k}$, $[\mathbf{M}\mathbf{x}]$ is computationally indistinguishable from a uniform vector in \mathbb{G}^k ; this is a generalization of the k -Lin Assumption.

We outline the construction under the k -Lin assumption over \mathbb{G} , of which the DDH assumption is a special case corresponding to $k = 1$.

In this overview, we will consider a weaker notion of security, namely tag-based KEM security against plaintext check attacks (PCA) [29]. In the PCA security experiment, the adversary gets no decryption oracle (as with CCA security), but a PCA oracle that takes as input a tag and a ciphertext/plaintext pair and checks whether the ciphertext decrypts to the plaintext. Furthermore, we restrict the adversary to only query the PCA oracle on tags different from those used in the challenge ciphertexts. PCA security is strictly weaker than the CCA security we actually strive for, but allows us to present our solution in a clean and simple way. (We show how to obtain full CCA security separately.)

The starting point of our construction is the Cramer-Shoup KEM, in which $\text{Enc}_{\text{KEM}}(\text{pk}, \tau)$ outputs the ciphertext/plaintext pair

$$([\mathbf{y}], [\mathbf{z}]) = ([\mathbf{x}^\top \mathbf{M}^\top], [\mathbf{x}^\top \mathbf{M}^\top \mathbf{k}_\tau]), \quad (1)$$

where $\mathbf{k}_\tau = \mathbf{k}_0 + \tau \mathbf{k}_1$ and $\text{pk} := ([\mathbf{M}], [\mathbf{M}^\top \mathbf{k}_0], [\mathbf{M}^\top \mathbf{k}_1])$ for $\mathbf{M} \leftarrow_{\text{R}} \mathbb{Z}_q^{(k+1) \times k}$. The KEM is PCA-secure under k -Lin, with a security loss that depends on the number of ciphertexts Q (via a hybrid argument) but independently of the number of PCA queries [1, 11].

Following the “randomized Naor-Reingold” paradigm introduced by Chen and Wee on tightly secure IBE [10], our starting point is (1), where we replace $\mathbf{k}_\tau = \mathbf{k}_0 + \tau \mathbf{k}_1$ with

$$\mathbf{k}_\tau = \sum_{j=1}^{\lambda} \mathbf{k}_{j, \tau_j}$$

and $\mathbf{pk} := ([\mathbf{M}], [\mathbf{M}^\top \mathbf{k}_{j,b}]_{j=1,\dots,\lambda, b=0,1})$, where $(\tau_1, \dots, \tau_\lambda)$ denotes the binary representation of the tag $\tau \in \{0, 1\}^\lambda$.

Following [10], we want to analyze this construction by a sequence of games in which we first replace $[\mathbf{y}]$ in the challenge ciphertexts by uniformly random group elements via random self-reducibility of MDDH (k -Lin), and then incrementally replace \mathbf{k}_τ in both the challenge ciphertexts and in the PCA oracle by $\mathbf{k}_\tau + \mathbf{m}^\perp \text{RF}(\tau)$, where RF is a truly random function and \mathbf{m}^\perp is a random element from the kernel of \mathbf{M} , i.e., $\mathbf{M}^\top \mathbf{m}^\perp = 0$. Concretely, in Game i , we will replace \mathbf{k}_τ with $\mathbf{k}_\tau + \mathbf{m}^\perp \text{RF}_i(\tau)$ where RF_i is a random function on $\{0, 1\}^i$ applied to the i -bit prefix of τ . We proceed to outline the two main ideas needed to carry out this transition. Looking ahead, note that once we reach Game λ , we would have replaced \mathbf{k}_τ with $\mathbf{k}_\tau + \mathbf{m}^\perp \text{RF}(\tau)$, upon which security follows from a straight-forward information-theoretic argument (and the fact that ciphertexts and decryption queries carry pairwise different τ).

First Idea. First, we show how to transition from Game i to Game $i+1$, under the restriction that the adversary is only allowed to query the encryption oracle on tags whose $i+1$ -st bit is 0; we show how to remove this unreasonable restriction later. Here, we rely on an *information-theoretic* argument similar to that of Cramer and Shoup to increase the entropy from RF_i to RF_{i+1} . This is in contrast to prior works which rely on a computational argument; note that the latter requires encoding secret keys as group elements and thus a pairing to carry out decryption.

More precisely, we pick a random function RF'_i on $\{0, 1\}^i$, and implicitly define RF_{i+1} as follows:

$$\text{RF}_{i+1}(\tau) = \begin{cases} \text{RF}_i(\tau) & \text{if } \tau_{i+1} = 0 \\ \text{RF}'_i(\tau) & \text{if } \tau_{i+1} = 1 \end{cases}$$

Observe all of the challenge ciphertexts leak no information about RF'_i or $\mathbf{k}_{i+1,1}$ since they all correspond to tags whose $i+1$ -st bit is 0. To handle a PCA query $(\tau, [\mathbf{y}], [z])$, we proceed via a case analysis:

- if $\tau_{i+1} = 0$, then $\mathbf{k}_\tau + \text{RF}_{i+1}(\tau) = \mathbf{k}_\tau + \text{RF}_i(\tau)$ and the PCA oracle returns the same value in both Games i and $i+1$.
- if $\tau_{i+1} = 1$ and \mathbf{y} lies in the span of \mathbf{M} , we have

$$\mathbf{y}^\top \mathbf{m}^\perp = 0 \implies \mathbf{y}^\top (\mathbf{k}_\tau + \mathbf{m}^\perp \text{RF}_i(\tau)) = \mathbf{y}^\top (\mathbf{k}_\tau + \mathbf{m}^\perp \text{RF}_{i+1}(\tau)),$$

and again the PCA oracle returns the same value in both Games i and $i+1$.

- if $\tau_{i+1} = 1$ and \mathbf{y} lies outside the span of \mathbf{M} , then $\mathbf{y}^\top \mathbf{k}_{i+1,1}$ is uniformly random given $\mathbf{M}, \mathbf{M}^\top \mathbf{k}_{i+1,1}$. (Here, we crucially use that the adversary does not query encryptions with $\tau_{i+1} = 1$, which ensures that the challenge ciphertexts do not leak additional information about $\mathbf{k}_{i+1,1}$.) This means that $\mathbf{y}^\top \mathbf{k}_\tau$ is uniformly random from the adversary's view-point, and therefore the PCA oracle will reject with high probability in both Games i and $i+1$. (At this point, we crucially rely on the fact that the PCA oracle only outputs a *single* check bit and not all of $\mathbf{k}_\tau + \text{RF}(\tau)$.)

Via a hybrid argument, we may deduce that the distinguishing advantage between Games i and $i + 1$ is at most Q/q where Q is the number of PCA queries.

Second Idea. Next, we remove the restriction on the encryption queries using an idea of Hofheinz et al. [18] for tightly-secure IBE in the multi-ciphertext setting, and its instantiation in prime-order groups [15]. The idea is to create two “independent copies” of $(\mathbf{M}^\perp, \mathbf{RF}_i)$; we use one to handle encryption queries on tags whose $i+1$ -st bit is 0, and the other to handle those whose $i+1$ -st bit is 1. We call these two copies $(\mathbf{M}_0^*, \mathbf{RF}_i^{(0)})$ and $(\mathbf{M}_1^*, \mathbf{RF}_i^{(1)})$, where $\mathbf{M}^\top \mathbf{M}_0^* = \mathbf{M}^\top \mathbf{M}_1^* = \mathbf{0}$.

Concretely, we replace $\mathbf{M} \leftarrow_{\mathbb{R}} \mathbb{Z}_q^{(k+1) \times k}$ with $\mathbf{M} \leftarrow_{\mathbb{R}} \mathbb{Z}_q^{3k \times k}$. We decompose \mathbb{Z}_q^{3k} into the span of the respective matrices $\mathbf{M}, \mathbf{M}_0, \mathbf{M}_1$, and we will also decompose the span of $\mathbf{M}^\perp \in \mathbb{Z}_q^{3k \times 2k}$ into that of $\mathbf{M}_0^*, \mathbf{M}_1^*$. Similarly, we decompose $\mathbf{M}^\perp \mathbf{RF}_i(\tau)$ into $\mathbf{M}_0^* \mathbf{RF}_i^{(0)}(\tau) + \mathbf{M}_1^* \mathbf{RF}_i^{(1)}(\tau)$. We then refine the prior transition from Games i to $i + 1$ as follows:

- Game $i.0$ (= Game i): pick $\mathbf{y} \leftarrow \mathbb{Z}_q^{3k}$ for ciphertexts, and replace \mathbf{k}_τ with $\mathbf{k}_\tau + \mathbf{M}_0^* \mathbf{RF}_i^{(0)}(\tau) + \mathbf{M}_1^* \mathbf{RF}_i^{(1)}(\tau)$;
- Game $i.1$: replace $\mathbf{y} \leftarrow_{\mathbb{R}} \mathbb{Z}_q^{3k}$ with $\mathbf{y} \leftarrow_{\mathbb{R}} \text{span}(\mathbf{M}, \mathbf{M}_{\tau_{i+1}})$;
- Game $i.2$: replace $\mathbf{RF}_i^{(0)}(\tau)$ with $\mathbf{RF}_{i+1}^{(0)}(\tau)$;
- Game $i.3$: replace $\mathbf{RF}_i^{(1)}(\tau)$ with $\mathbf{RF}_{i+1}^{(1)}(\tau)$;
- Game $i.4$ (= Game $i + 1$): replace $\mathbf{y} \leftarrow_{\mathbb{R}} \text{span}(\mathbf{M}, \mathbf{M}_{\tau_{i+1}})$ with $\mathbf{y} \leftarrow_{\mathbb{R}} \mathbb{Z}_q^{3k}$.

For the transition from Game $i.0$ to Game $i.1$, we rely on the fact that the uniform distributions over \mathbb{Z}_q^{3k} and $\text{span}(\mathbf{M}, \mathbf{M}_{\tau_{i+1}})$ encoded in the group are computationally indistinguishable, even given a random basis for $\text{span}(\mathbf{M}^\perp)$ (in the clear). This extends to the setting with multiple samples, with a tight reduction to the \mathcal{D}_k -MDDH Assumption independent of the number of samples.




Fig. 1. Solid lines mean orthogonal, that is: $\mathbf{M}^\top \mathbf{M}_0^* = \mathbf{M}_1^\top \mathbf{M}_0^* = \mathbf{0} = \mathbf{M}^\top \mathbf{M}_1^* = \mathbf{M}_0^\top \mathbf{M}_1^*$.

For the transition from Game $i.1$ to $i.2$, we rely on an information-theoretic argument like the one we just outlined, replacing $\text{span}(\mathbf{M})$ with $\text{span}(\mathbf{M}, \mathbf{M}_1)$ and \mathbf{M}^\perp with \mathbf{M}_0^* in the case analysis. In particular, we will exploit the fact that if \mathbf{y} lies outside $\text{span}(\mathbf{M}, \mathbf{M}_1)$, then $\mathbf{y}^\top \mathbf{k}_{i+1,1}$ is uniformly random even given $\mathbf{M}, \mathbf{M} \mathbf{k}_{i+1,1}, \mathbf{M}_1, \mathbf{M}_1 \mathbf{k}_{i+1,1}$. The transition from Game $i.2$ to $i.3$ is completely analogous.

From PCA to CCA. Using standard techniques from [4, 8, 11, 19, 21], we could transform our basic tag-based PCA-secure scheme into a “full-fledged” CCA-secure encryption scheme by adding another hash proof system (or an authenticated symmetric encryption scheme) and a one-time signature scheme. However, this would incur an additional overhead of several group elements in the ciphertext. Instead, we show how to directly modify our tag-based PCA-secure scheme to obtain a more efficient CCA-secure scheme with the minimal additional overhead of a single symmetric-key authenticated encryption. In particular, the overall ciphertext overhead in our tightly CCA-secure encryption scheme is merely *one* group element more than that for the best known non-tight schemes [17, 21].

To encrypt a message M in the CCA-secure encryption scheme, we will (i) pick a random \mathbf{y} as in the tag-based PCA scheme, (ii) derive a tag τ from \mathbf{y} , (iii) encrypt M using a one-time authenticated encryption under the KEM key $[\mathbf{y}^\top \mathbf{k}_\tau]$. The naive approach is to derive the tag τ by hashing $[\mathbf{y}] \in \mathbb{G}^{3k}$, as in [21]. However, this creates a circularity in Game $i.1$ where the distribution of $[\mathbf{y}]$ depends on the tag. Instead, we will derive the tag τ by hashing $[\bar{\mathbf{y}}] \in \mathbb{G}^k$, where $\bar{\mathbf{y}} \in \mathbb{Z}_q^k$ are the top k entries of $\mathbf{y} \in \mathbb{Z}_q^{3k}$. We then modify $\mathbf{M}_0, \mathbf{M}_1$ so that the top k rows of both matrices are zero, which avoids the circularity issue. In the proof of security, we will also rely on the fact that for any $\mathbf{y}_0, \mathbf{y}_1 \in \mathbb{Z}_q^{3k}$, if $\bar{\mathbf{y}}_0 = \bar{\mathbf{y}}_1$ and $\mathbf{y}_0 \in \text{span}(\mathbf{M})$, then either $\mathbf{y}_0 = \mathbf{y}_1$ or $\mathbf{y}_1 \notin \text{span}(\mathbf{M})$. This allows us to deduce that if the adversary queries the CCA oracle on a ciphertext which shares the same tag as some challenge ciphertext, then the CCA oracle will reject with overwhelming probability.

Alternative View-Point. Our construction can also be viewed as applying the BCHK IBE \rightarrow PKE transform [8] to the scheme from [18], and then writing the exponents of the secret keys in the clear, thereby avoiding the pairing. This means that we can no longer apply a computational assumption and the randomized Naor-Reingold argument to the secret key space. Indeed, we replace this with an information-theoretic Cramer-Shoup-like argument as outlined above.

Prior Approaches. Several approaches to construct tightly CCA-secure PKE schemes exist: first, the schemes of [2, 3, 16, 22–24] construct a tightly secure NIZK scheme from a tightly secure signature scheme, and then use the tightly secure NIZK in a CCA-secure PKE scheme following the Naor-Yung double encryption paradigm [12, 28]. Since these approaches build on the public verifiability of the used NIZK scheme (in order to faithfully simulate a decryption oracle), their reliance on a pairing seems inherent.

Next, the works of [5, 7, 10, 15, 18] used a (Naor-Reingold-based) MAC instead of a signature scheme to design tightly secure IBE schemes. Those IBE schemes can then be converted (using the BCHK transformation [8]) into tightly CCA-secure PKE schemes. However, the derived PKE schemes still rely on pairings, since the original IBE schemes do (and the BCHK does not remove the reliance on pairings).

In contrast, our approach directly fuses a Naor-Reingold-like randomization argument with the encryption process. We are able to do so since we substitute

a computational randomization argument (as used in the latter line of works) with an information-theoretic one, as described above. Hence, we can apply that argument to *exponents* rather than group elements. This enables us to trade pairing operations for exponentiations in our scheme.

Efficiency Comparison with Non-tightly Secure Schemes. We finally mention that our DDH-based scheme compares favorably even with the most efficient (non-tightly) CCA-secure DDH-based encryption schemes [17,21]. To make things concrete, assume $\lambda = 80$ and a setting with $Q_{\text{enc}} = Q_{\text{dec}} = 2^{30}$. The best known reductions for the schemes of [17,21] lose a factor of $Q_{\text{enc}} = 2^{30}$, whereas our scheme loses a factor of about $4\lambda \leq 2^9$. Hence, the group size for [17,21] should be at least $2^{2 \cdot (80+30)} = 2^{220}$ compared to $2^{2 \cdot (80+9)} = 2^{178}$ in our case. Thus, the ciphertext overhead (ignoring the symmetric encryption part) in our scheme is $3 \cdot 178 = 534$ bits, which is close to $2 \cdot 220 = 440$ bits with [17,21].²

Perhaps even more interestingly, we can compare computational efficiency of encryption in this scenario. For simplicity, we only count exponentiations and assume a naive square-and-multiply-based exponentiation with no further multi-exponentiation optimizations.³ Encryption in [17,21] takes about 3.5 exponentiations (where we count an exponentiation with a $(\lambda + \log_2(Q_{\text{enc}} + Q_{\text{dec}}))$ -bit hash value⁴ as 0.5 exponentiations). In our scheme, we have about 4.67 exponentiations, where we count the computation of $[\mathbf{M}^\top \mathbf{k}_r]$ – which consists of 2λ multiplications – as 0.67 exponentiations.) Since exponentiation (under our assumptions) takes time cubic in the bitlength, we get that encryption with our scheme is actually about 29% *less expensive* than with [17,21].

However, of course we should also note that public and secret key in our scheme are significantly larger (e.g., $4\lambda + 3 = 323$ group elements in pk) than with [17,21] (4 group elements in pk).

Extension: NIZK Arguments. We also obtain tightly simulation-sound non-interactive zero-knowledge (NIZK) arguments from our encryption scheme in a semi-generic way.

Let us start with any designated-verifier quasi-adaptive NIZK (short: DVQANIZK) argument system Π for a given language. Recall that in a designated-verifier NIZK, proofs can only be verified with a secret verification key, and soundness only holds against adversaries who do not know that key. Furthermore, quasi-adaptivity means that the language has to be fixed at setup time of the scheme. Let Π_{PKE} be the variant of Π in which proofs are encrypted using

² In this calculation, we do not consider the symmetric authenticated encryption of the actual plaintext (and a corresponding MAC value), which is the same with [17,21] and our scheme.

³ Here, optimizations would improve the schemes of [17,21] and ours similarly, since the schemes are very similar.

⁴ It is possible to prove the security of [17,21] using a *target*-collision-resistant hash function, such that $|\tau| = \lambda$. However, in the multi-user setting, a hybrid argument is required, such that the output size of the hash function will have to be increased to at least $|\tau| = \lambda + \log_2(Q_{\text{enc}} + Q_{\text{dec}})$.

Reference	$ \text{pk} $	$ \text{ct} - m $	security loss	assumption	pairing
CS98 [11]	$\mathcal{O}(1)$	3	$\mathcal{O}(Q)$	DDH	no
KD04, HK07 [23, 18]	$\mathcal{O}(1)$	2	$\mathcal{O}(Q)$	DDH	no
HJ12 [17]	$\mathcal{O}(1)$	$\mathcal{O}(\lambda)$	$\mathcal{O}(1)$	2-Lin	yes
LPJY15 [24, 26]	$\mathcal{O}(\lambda)$	47	$\mathcal{O}(\lambda)$	2-Lin	yes
AHY15 [5]	$\mathcal{O}(\lambda)$	12	$\mathcal{O}(\lambda)$	2-Lin	yes
GCDCT15 [16]	$\mathcal{O}(\lambda)$	10 (resp. $6k + 4$)	$\mathcal{O}(\lambda)$	SXDH (resp. k -Lin)	yes
Ours §4	$\mathcal{O}(\lambda)$	3 (resp. $3k$)	$\mathcal{O}(\lambda)$	DDH (resp. k -Lin)	no

Fig. 2. Comparison amongst CCA-secure encryption schemes, where Q is the number of ciphertexts, $|\text{pk}|$ denotes the size (i.e. the number of groups elements, or exponent of group elements) of the public key, and $|\text{ct}| - |m|$ denotes the ciphertext overhead, ignoring smaller contributions from symmetric-key encryption. We omit [18] from this table since we only focus on prime-order groups here.

a CCA-secure PKE scheme PKE. Public and secret key of PKE are of course made part of CRS and verification key, respectively. Observe that Π_{PKE} enjoys simulation-soundness, assuming that simulated proofs are simply encryptions of random plaintexts. Indeed, the CCA security of PKE guarantees that authentic Π_{PKE} -proofs can be substituted with simulated ones, while being able to verify (using a decryption oracle) a purported Π_{PKE} -proof generated by an adversary. Furthermore, if PKE is tightly secure, then so is Π_{PKE} .

When using a hash proof system for Π and our encryption scheme for PKE, this immediately yields a tightly simulation-sound DVQANIZK for linear languages (i.e., languages of the form $\{\mathbf{M}\mathbf{x} \mid \mathbf{x} \in \mathbb{Z}_q^t\}$ for some matrix $\mathbf{M} \in \mathbb{Z}_q^{n \times t}$ with $t < n$) that does not require pairings. We stress that our DVQANIZK is tightly secure in a setting with many simulated proofs and many adversarial verification queries.

Using the semi-generic transformation of [20], we can then derive a tightly simulation-sound QANIZK proof system (with public verification), that however relies on pairings. We note that the transformation of [20] only requires a DVQANIZK that is secure against a single adversarial verification query, since the pairing enables the public verifiability of proofs. Hence, we can first optimize and trim down our DVQANIZK (such that only a single adversarial verification query is supported), and then apply the transformation. This yields a QANIZK with particularly compact proofs. See Fig. 3 for a comparison with relevant existing proof systems.

Roadmap. We recall some notation and basic definitions (including those concerning our algebraic setting and for tightly secure encryption) in Sect. 2. Section 3 presents our basic PCA-secure encryption scheme and represents the core of our results. In Sect. 4, we present our optimized CCA-secure PKE scheme. Our NIZK-related applications are presented in the full version of this paper.

Reference	type	$ \text{crs} $	$ \pi $	sec. loss	assumption	pairing
CCS09 [9]	NIZK	$\mathcal{O}(1)$	$2n + 6t + 52$	$\mathcal{O}(Q_{\text{sim}})$	2-Lin	yes
HJ12 [17]	NIZK	$\mathcal{O}(1)$	$\gg 500$	$\mathcal{O}(1)$	2-Lin	yes
LPJY14 [25]	QANIZK	$\mathcal{O}(n + \lambda)$	20	$\mathcal{O}(Q_{\text{sim}})$	2-Lin	yes
KW15 [22]	QANIZK	$\mathcal{O}(kn)$	$2k + 2$	$\mathcal{O}(Q_{\text{sim}})$	k -Lin	yes
LPJY15 [27]	QANIZK	$\mathcal{O}(n + \lambda)$	42	$\mathcal{O}(\lambda)$	2-Lin	yes
Ours (full version)	DVQANIZK	$\mathcal{O}(t + k\lambda)$	$3k + 1$	$\mathcal{O}(\lambda)$	k -Lin	no
Ours (full version)	QANIZK	$\mathcal{O}(k^2\lambda + kn)$	$2k + 1$	$\mathcal{O}(\lambda)$	k -Lin	yes

Fig. 3. (DV)QANIZK schemes for subspaces of \mathbb{G}^n of dimension $t < n$. $|\text{crs}|$ and $|\pi|$ denote the size (in group elements) of the CRS and of proofs. Q_{sim} is the number of simulated proofs in the simulation-soundness experiment. The scheme from [20] (as well as our own schemes) can also be generalized to matrix assumptions [13], at the cost of a larger CRS.

2 Preliminaries

2.1 Notations

If $\mathbf{x} \in \mathcal{B}^n$, then $|\mathbf{x}|$ denotes the length n of the vector. Further, $x \leftarrow_{\text{R}} \mathcal{B}$ denotes the process of sampling an element x from set \mathcal{B} uniformly at random. For any bit string $\tau \in \{0, 1\}^*$, we denote by τ_i the i 'th bit of τ . We denote by λ the security parameter, and by $\text{negl}(\cdot)$ any negligible function of λ . For all matrix $\mathbf{A} \in \mathbb{Z}_q^{\ell \times k}$ with $\ell > k$, $\overline{\mathbf{A}} \in \mathbb{Z}_q^{k \times k}$ denotes the upper square matrix of \mathbf{A} and $\underline{\mathbf{A}} \in \mathbb{Z}_q^{\ell - k \times k}$ denotes the lower $\ell - k$ rows of \mathbf{A} . With $\text{span}(\mathbf{A}) := \{\mathbf{A}\mathbf{r} \mid \mathbf{r} \in \mathbb{Z}_q^k\} \subset \mathbb{Z}_q^\ell$, we denote the span of \mathbf{A} .

2.2 Collision Resistant Hashing

A hash function generator is a PPT algorithm \mathcal{H} that, on input 1^λ , outputs an efficiently computable function $\text{H} : \{0, 1\}^* \rightarrow \{0, 1\}^\lambda$.

Definition 1 (Collision Resistance). *We say that a hash function generator \mathcal{H} outputs collision-resistant functions H if for all PPT adversaries \mathcal{A} ,*

$$\text{Adv}_{\mathcal{H}}^{\text{cr}}(\mathcal{A}) := \Pr \left[x \neq x' \wedge \text{H}(x) = \text{H}(x') \mid \begin{array}{l} \text{H} \leftarrow_{\text{R}} \mathcal{H}(1^\lambda), \\ (x, x') \leftarrow \mathcal{A}(1^\lambda, \text{H}) \end{array} \right] = \text{negl}(\lambda).$$

2.3 Prime-Order Groups

Let GGen be a probabilistic polynomial time (PPT) algorithm that on input 1^λ returns a description $\mathcal{G} = (\mathbb{G}, q, P)$ of an additive cyclic group \mathbb{G} of order q for a λ -bit prime q , whose generator is P .

We use implicit representation of group elements as introduced in [13]. For $a \in \mathbb{Z}_q$, define $[a] = aP \in \mathbb{G}$ as the *implicit representation* of a in \mathbb{G} . More generally, for a matrix $\mathbf{A} = (a_{ij}) \in \mathbb{Z}_q^{n \times m}$ we define $[\mathbf{A}]$ as the implicit representation of \mathbf{A} in \mathbb{G} :

$$[\mathbf{A}] := \begin{pmatrix} a_{11}P & \dots & a_{1m}P \\ \vdots & & \vdots \\ a_{n1}P & \dots & a_{nm}P \end{pmatrix} \in \mathbb{G}^{n \times m}$$

We will always use this implicit notation of elements in \mathbb{G} , i.e., we let $[a] \in \mathbb{G}$ be an element in \mathbb{G} . Note that from $[a] \in \mathbb{G}$ it is generally hard to compute the value a (discrete logarithm problem in \mathbb{G}). Obviously, given $[a], [b] \in \mathbb{G}$ and a scalar $x \in \mathbb{Z}_q$, one can efficiently compute $[ax] \in \mathbb{G}$ and $[a + b] \in \mathbb{G}$.

2.4 Matrix Diffie-Hellman Assumption

We recall the definitions of the Matrix Decision Diffie-Hellman (MDDH) Assumption [13].

Definition 2 (Matrix Distribution). *Let $k, \ell \in \mathbb{N}$, with $\ell > k$. We call $\mathcal{D}_{\ell, k}$ a matrix distribution if it outputs matrices in $\mathbb{Z}_q^{\ell \times k}$ of full rank k in polynomial time. We write $\mathcal{D}_k := \mathcal{D}_{k+1, k}$.*

Without loss of generality, we assume the first k rows of $\mathbf{A} \leftarrow_{\mathbb{R}} \mathcal{D}_{\ell, k}$ form an invertible matrix. The $\mathcal{D}_{\ell, k}$ -Matrix Diffie-Hellman problem is to distinguish the two distributions $([\mathbf{A}], [\mathbf{Aw}])$ and $([\mathbf{A}], [\mathbf{u}])$ where $\mathbf{A} \leftarrow_{\mathbb{R}} \mathcal{D}_{\ell, k}$, $\mathbf{w} \leftarrow_{\mathbb{R}} \mathbb{Z}_q^k$ and $\mathbf{u} \leftarrow_{\mathbb{R}} \mathbb{Z}_q^\ell$.

Definition 3 ($\mathcal{D}_{\ell, k}$ -Matrix Diffie-Hellman Assumption $\mathcal{D}_{\ell, k}$ -MDDH). *Let $\mathcal{D}_{\ell, k}$ be a matrix distribution. We say that the $\mathcal{D}_{\ell, k}$ -Matrix Diffie-Hellman ($\mathcal{D}_{\ell, k}$ -MDDH) Assumption holds relative to GGen if for all PPT adversaries \mathcal{A} ,*

$$\begin{aligned} \text{Adv}_{\mathcal{D}_{\ell, k}, \text{GGen}}^{\text{mddh}}(\mathcal{A}) := \\ |\Pr[\mathcal{A}(\mathcal{G}, [\mathbf{A}], [\mathbf{Aw}]) = 1] - \Pr[\mathcal{A}(\mathcal{G}, [\mathbf{A}], [\mathbf{u}]) = 1]| = \text{negl}(\lambda), \end{aligned}$$

where the probability is over $\mathcal{G} \leftarrow_{\mathbb{R}} \text{GGen}(1^\lambda)$, $\mathbf{A} \leftarrow_{\mathbb{R}} \mathcal{D}_{\ell, k}$, $\mathbf{w} \leftarrow_{\mathbb{R}} \mathbb{Z}_q^k$, $\mathbf{u} \leftarrow_{\mathbb{R}} \mathbb{Z}_q^\ell$.

For each $k \geq 1$, [13] specifies distributions $\mathcal{L}_k, \mathcal{SC}_k, \mathcal{C}_k$ (and others) over $\mathbb{Z}_q^{(k+1) \times k}$ such that the corresponding \mathcal{D}_k -MDDH assumptions are generically secure in bilinear groups and form a hierarchy of increasingly weaker assumptions. \mathcal{L}_k -MDDH is the well known k -Linear Assumption k -Lin with 1-Lin = DDH. In this work we are mostly interested in the uniform matrix distribution $\mathcal{U}_{\ell, k}$.

Definition 4 (Uniform Distribution). *Let $\ell, k \in \mathbb{N}$, with $\ell > k$. We denote by $\mathcal{U}_{\ell, k}$ the uniform distribution over all full-rank $\ell \times k$ matrices over \mathbb{Z}_q . Let $\mathcal{U}_k := \mathcal{U}_{k+1, k}$.*

Lemma 1 (\mathcal{U}_k -MDDH $\Leftrightarrow \mathcal{U}_{\ell, k}$ -MDDH). *Let $\ell, k \in \mathbb{N}$, with $\ell > k$. For any PPT adversary \mathcal{A} , there exists an adversary \mathcal{B} (and vice versa) such that $\mathbf{T}(\mathcal{B}) \approx \mathbf{T}(\mathcal{A})$ and $\text{Adv}_{\mathcal{U}_{\ell, k}, \text{GGen}}^{\text{mddh}}(\mathcal{A}) = \text{Adv}_{\mathcal{U}_k, \text{GGen}}^{\text{mddh}}(\mathcal{B})$.*

Proof. This follows from the simple fact that a $\mathcal{U}_{\ell,k}$ -MDDH instance $([\mathbf{A}], [\mathbf{z}])$ can be transformed into an \mathcal{U}_k -MDDH instance $([\mathbf{A}'] = [\mathbf{TA}], [\mathbf{z}'] = [\mathbf{Tz}])$ for a random $(k+1) \times \ell$ matrix \mathbf{T} . If $\mathbf{z} = \mathbf{Aw}$, then $\mathbf{z}' = \mathbf{TAw} = \mathbf{A}'\mathbf{w}$; if \mathbf{z} is uniform, so is \mathbf{z}' . Similarly, a \mathcal{U}_k -MDDH instance $([\mathbf{A}'], [\mathbf{z}'])$ can be transformed into an $\mathcal{U}_{\ell,k}$ -MDDH instance $([\mathbf{A}] = [\mathbf{T}'\mathbf{A}'], [\mathbf{z}] = [\mathbf{T}'\mathbf{z}'])$ for a random $\ell \times (k+1)$ matrix \mathbf{T}' . \square

Among all possible matrix distributions $\mathcal{D}_{\ell,k}$, the uniform matrix distribution \mathcal{U}_k is the hardest possible instance, so in particular $k\text{-Lin} \Rightarrow \mathcal{U}_k\text{-MDDH}$.

Lemma 2 ($\mathcal{D}_{\ell,k}$ -MDDH $\Rightarrow \mathcal{U}_k$ -MDDH, [13]). *Let $\mathcal{D}_{\ell,k}$ be a matrix distribution. For any PPT adversary \mathcal{A} , there exists an adversary \mathcal{B} such that $\mathbf{T}(\mathcal{B}) \approx \mathbf{T}(\mathcal{A})$ and $\text{Adv}_{\mathcal{D}_{\ell,k}, \text{GGen}}^{\text{mddh}}(\mathcal{A}) = \text{Adv}_{\mathcal{U}_k, \text{GGen}}^{\text{mddh}}(\mathcal{B})$.*

Let $Q \geq 1$. For $\mathbf{W} \leftarrow_{\mathbb{R}} \mathbb{Z}_q^{k \times Q}$, $\mathbf{U} \leftarrow_{\mathbb{R}} \mathbb{Z}_q^{\ell \times Q}$, we consider the Q -fold $\mathcal{D}_{\ell,k}$ -MDDH Assumption which consists in distinguishing the distributions $([\mathbf{A}], [\mathbf{AW}])$ from $([\mathbf{A}], [\mathbf{U}])$. That is, a challenge for the Q -fold $\mathcal{D}_{\ell,k}$ -MDDH Assumption consists of Q independent challenges of the $\mathcal{D}_{\ell,k}$ -MDDH Assumption (with the same \mathbf{A} but different randomness \mathbf{w}). In [13] it is shown that the two problems are equivalent, where (for $Q \geq \ell - k$) the reduction loses a factor $\ell - k$. In combination with Lemma 1 we obtain the following tighter version for the special case of $\mathcal{D}_{\ell,k} = \mathcal{U}_{\ell,k}$.

Lemma 3 (Random Self-reducibility of $\mathcal{U}_{\ell,k}$ -MDDH, [13]). *Let $\ell, k, Q \in \mathbb{N}$ with $\ell > k$. For any PPT adversary \mathcal{A} , there exists an adversary \mathcal{B} such that $\mathbf{T}(\mathcal{B}) \approx \mathbf{T}(\mathcal{A}) + Q \cdot \text{poly}(\lambda)$ with $\text{poly}(\lambda)$ independent of $\mathbf{T}(\mathcal{A})$, and*

$$\text{Adv}_{\mathcal{U}_{\ell,k}, \text{GGen}}^{Q\text{-mddh}}(\mathcal{A}) \leq \text{Adv}_{\mathcal{U}_{\ell,k}, \text{GGen}}^{\text{mddh}}(\mathcal{B}) + \frac{1}{q-1}$$

where $\text{Adv}_{\mathcal{U}_{\ell,k}, \text{GGen}}^{Q\text{-mddh}}(\mathcal{B}) := |\Pr[\mathcal{B}(\mathcal{G}, [\mathbf{A}], [\mathbf{AW}]) = 1] - \Pr[\mathcal{B}(\mathcal{G}, [\mathbf{A}], [\mathbf{U}]) = 1]|$ and the probability is over $\mathcal{G} \leftarrow_{\mathbb{R}} \text{GGen}(1^\lambda)$, $\mathbf{A} \leftarrow_{\mathbb{R}} \mathcal{U}_{\ell,k}$, $\mathbf{W} \leftarrow_{\mathbb{R}} \mathbb{Z}_q^{k \times Q}$, $\mathbf{U} \leftarrow_{\mathbb{R}} \mathbb{Z}_q^{\ell \times Q}$.

2.5 Public-Key Encryption

Definition 5 (PKE). *A Public-Key Encryption (PKE) consists of three PPT algorithms $\text{PKE} = (\text{Param}_{\text{PKE}}, \text{Gen}_{\text{PKE}}, \text{Enc}_{\text{PKE}}, \text{Dec}_{\text{PKE}})$:*

- The probabilistic key generation algorithm $\text{Gen}_{\text{PKE}}(1^\lambda)$ generates a pair of public and secret keys (pk, sk) .
- The probabilistic encryption algorithm $\text{Enc}_{\text{PKE}}(\text{pk}, M)$ returns a ciphertext ct .
- The deterministic decryption algorithm $\text{Dec}_{\text{PKE}}(\text{pk}, \text{sk}, \text{ct})$ returns a message M or \perp , where \perp is a special rejection symbol.

We define the following properties:

Perfect Correctness. *For all λ , we have*

$$\Pr \left[\text{Dec}_{\text{PKE}}(\text{pk}, \text{sk}, \text{ct}) = M \mid \begin{array}{l} (\text{pk}, \text{sk}) \leftarrow_{\mathbb{R}} \text{Gen}_{\text{PKE}}(1^\lambda); \\ \text{ct} \leftarrow_{\mathbb{R}} \text{Enc}_{\text{PKE}}(\text{pk}, M) \end{array} \right] = 1.$$

Multi-ciphertext CCA Security [6]. For any adversary \mathcal{A} , we define

$$\text{Adv}_{\text{PKE}}^{\text{ind-cca}}(\mathcal{A}) := \left| \Pr \left[b = b' \mid b' \leftarrow \mathcal{A}^{\text{Setup}, \text{DecO}(\cdot, \cdot), \text{EncO}(\cdot, \cdot)}(1^\lambda) \right] - 1/2 \right|$$

where:

- Setup sets $\mathcal{C}_{\text{enc}} := \emptyset$, samples $(\text{pk}, \text{sk}) \leftarrow_{\text{R}} \text{Gen}_{\text{KEM}}(1^\lambda)$ and $b \leftarrow_{\text{R}} \{0, 1\}$, and returns pk . Setup must be called once at the beginning of the game.
- DecO(ct) returns $\text{Dec}_{\text{PKE}}(\text{pk}, \text{sk}, \text{ct})$ if $\text{ct} \notin \mathcal{C}_{\text{enc}}$, \perp otherwise.
- If M_0 and M_1 are two messages of equal length, $\text{EncO}(M_0, M_1)$ returns $\text{Enc}_{\text{PKE}}(\text{pk}, M_b)$ and sets $\mathcal{C}_{\text{enc}} := \mathcal{C}_{\text{enc}} \cup \{\text{ct}\}$.

We say PKE is IND-CCA secure if for all PPT adversaries \mathcal{A} , the advantage $\text{Adv}_{\text{PKE}}^{\text{ind-cca}}(\mathcal{A})$ is a negligible function of λ .

2.6 Key-Encapsulation Mechanism

Definition 6 (Tag-based KEM). A tag-based Key-Encapsulation Mechanism (KEM) consists of three PPT algorithms $\text{KEM} = (\text{Gen}_{\text{KEM}}, \text{Enc}_{\text{KEM}}, \text{Dec}_{\text{KEM}})$:

- The probabilistic key generation algorithm $\text{Gen}_{\text{KEM}}(1^\lambda)$ generates a pair of public and secret keys (pk, sk) .
- The probabilistic encryption algorithm $\text{Enc}_{\text{KEM}}(\text{pk}, \tau)$ returns a pair (K, C) where K is a uniformly distributed symmetric key in \mathcal{K} and C is a ciphertext, with respect to the tag $\tau \in \mathcal{T}$.
- The deterministic decryption algorithm $\text{Dec}_{\text{KEM}}(\text{pk}, \text{sk}, \tau, C)$ returns a key $K \in \mathcal{K}$.

We define the following properties:

Perfect Correctness. For all λ , for all tags $\tau \in \mathcal{T}$, we have

$$\Pr \left[\text{Dec}_{\text{KEM}}(\text{pk}, \text{sk}, \tau, C) = K \mid \begin{array}{l} (\text{pk}, \text{sk}) \leftarrow_{\text{R}} \text{Gen}_{\text{KEM}}(1^\lambda); \\ (K, C) \leftarrow_{\text{R}} \text{Enc}_{\text{KEM}}(\text{pk}, \tau) \end{array} \right] = 1.$$

Multi-ciphertext PCA Security [29]. For any adversary \mathcal{A} , we define

$$\text{Adv}_{\text{KEM}}^{\text{ind-pca}}(\mathcal{A}) := \left| \Pr \left[b = b' \mid b' \leftarrow \mathcal{A}^{\text{Setup}, \text{DecO}(\cdot, \cdot), \text{EncO}(\cdot)}(1^\lambda) \right] - 1/2 \right|$$

where:

- Setup sets $\mathcal{T}_{\text{enc}} = \mathcal{T}_{\text{dec}} := \emptyset$, samples $(\text{pk}, \text{sk}) \leftarrow_{\text{R}} \text{Gen}_{\text{KEM}}(1^\lambda)$, picks $b \leftarrow_{\text{R}} \{0, 1\}$, and returns pk . Setup is called once at the beginning of the game.
- The decryption oracle $\text{DecO}(\tau, C, \hat{K})$ computes $K := \text{Dec}_{\text{KEM}}(\text{pk}, \text{sk}, \tau, C)$. It returns 1 if $\hat{K} = K \wedge \tau \notin \mathcal{T}_{\text{enc}}$, 0 otherwise. Then it sets $\mathcal{T}_{\text{dec}} := \mathcal{T}_{\text{dec}} \cup \{\tau\}$.
- EncO(τ) computes $(K, C) \leftarrow_{\text{R}} \text{Enc}_{\text{KEM}}(\text{pk}, \tau)$, sets $K_0 := K$ and $K_1 \leftarrow_{\text{R}} \mathcal{K}$. If $\tau \notin \mathcal{T}_{\text{dec}} \cup \mathcal{T}_{\text{enc}}$, it returns (C, K_b) , and sets $\mathcal{T}_{\text{enc}} := \mathcal{T}_{\text{enc}} \cup \{\tau\}$; otherwise it returns \perp .

We say KEM is IND-PCA secure if for all PPT adversaries \mathcal{A} , the advantage $\text{Adv}_{\text{KEM}}^{\text{ind-pca}}(\mathcal{A})$ is a negligible function of λ .

2.7 Authenticated Encryption

Definition 7 (AE [17]). An authenticated symmetric encryption (AE) with message-space \mathcal{M} and key-space \mathcal{K} consists of two polynomial-time deterministic algorithms $(\text{Enc}_{\text{AE}}, \text{Dec}_{\text{AE}})$:

- The encryption algorithm $\text{Enc}_{\text{AE}}(K, M)$ generates C , encryption of the message M with the secret key K .
- The decryption algorithm $\text{Dec}_{\text{AE}}(K, C)$, returns a message M or \perp .

We require that the algorithms satisfy the following properties:

Perfect Correctness. For all λ , for all $K \in \mathcal{K}$ and $M \in \mathcal{M}$, we have

$$\text{Dec}_{\text{AE}}(K, \text{Enc}_{\text{AE}}(K, M)) = M.$$

One-Time Privacy and Authenticity. For any PPT adversary \mathcal{A} ,

$$\begin{aligned} & \text{Adv}_{\text{AE}}^{\text{ae-ot}}(\mathcal{A}) \\ & := \left| \Pr \left[b' = b \mid \begin{array}{l} K \leftarrow_{\text{R}} \mathcal{K}; b \leftarrow_{\text{R}} \{0, 1\} \\ b' \leftarrow_{\text{R}} \mathcal{A}^{\text{ot-EncO}(\cdot, \cdot), \text{ot-DecO}(\cdot)}(1^\lambda, \mathcal{K}) \end{array} \right] - 1/2 \right| \end{aligned}$$

is negligible, where $\text{ot-EncO}(M_0, M_1)$, on input two messages M_0 and M_1 of the same length, $\text{Enc}_{\text{AE}}(K, M_b)$, and $\text{ot-DecO}(\phi)$ returns $\text{Dec}_{\text{AE}}(K, \phi)$ if $b = 0, \perp$ otherwise. \mathcal{A} is allowed at most one call to each oracle ot-EncO and ot-DecO , and the query to ot-DecO must be different from the output of ot-EncO . \mathcal{A} is also given the description of the key-space \mathcal{K} as input.

3 Multi-ciphertext PCA-secure KEM

In this section we describe a tag-based Key Encapsulation Mechanism KEM_{PCA} that is IND-PCA-secure (see Definition 6).

For simplicity, we use the matrix distribution $\mathcal{U}_{3k,k}$ in our scheme in Fig. 4, and prove it secure under the \mathcal{U}_k -MDDH Assumption ($\Leftrightarrow \mathcal{U}_{3k,k}$ -MDDH Assumption, by Lemma 1), which in turn admits a tight reduction to the standard k -Lin Assumption. However, using a matrix distribution $\mathcal{D}_{3k,k}$ with more compact representation yields a more efficient scheme, secure under the $\mathcal{D}_{3k,k}$ -MDDH Assumption (see Remark 1).

3.1 Our Construction

Remark 1 (On the use of the \mathcal{U}_k -MDDH Assumption). In our scheme, we use a matrix distribution $\mathcal{U}_{3k,k}$ for the matrix \mathbf{M} , therefore proving security under the $\mathcal{U}_{3k,k}$ -MDDH Assumption $\Leftrightarrow \mathcal{U}_k$ -MDDH Assumption (see Lemma 2). This is for simplicity of presentation. However, for efficiency, one may want to use an assumption with a more compact representation, such as the $\mathcal{CI}_{3k,k}$ -MDDH Assumption [26] with representation size $2k$ instead of $3k^2$ for $\mathcal{U}_{3k,k}$.

$\text{Gen}_{\text{KEM}}(1^\lambda):$ $\mathcal{G} \leftarrow_{\text{R}} \text{GGen}(1^\lambda); \mathbf{M} \leftarrow_{\text{R}} \mathcal{U}_{3k,k}$ $\mathbf{k}_{1,0}, \dots, \mathbf{k}_{\lambda,1} \leftarrow_{\text{R}} \mathbb{Z}_q^{3k}$ $\text{pk} := \left(\mathcal{G}, [\mathbf{M}], ([\mathbf{M}^\top \mathbf{k}_{j,\beta}])_{1 \leq j \leq \lambda, 0 \leq \beta \leq 1} \right)$ $\text{sk} := (\mathbf{k}_{j,\beta})_{1 \leq j \leq \lambda, 0 \leq \beta \leq 1}$ $\text{Return } (\text{pk}, \text{sk})$	$\text{Enc}_{\text{KEM}}(\text{pk}, \tau):$ $\mathbf{r} \leftarrow_{\text{R}} \mathbb{Z}_q^k; C := [\mathbf{r}^\top \mathbf{M}^\top]$ $\mathbf{k}_\tau := \sum_{j=1}^{\lambda} \mathbf{k}_{j,\tau_j}$ $K := [\mathbf{r}^\top \cdot \mathbf{M}^\top \mathbf{k}_\tau]$ $\text{Return } (C, K) \in \mathbb{G}^{1 \times 3k} \times \mathbb{G}$ $\text{Dec}_{\text{KEM}}(\text{pk}, \text{sk}, \tau, C):$ $\mathbf{k}_\tau := \sum_{j=1}^{\lambda} \mathbf{k}_{j,\tau_j}$ $\text{Return } K := C \cdot \mathbf{k}_\tau$
---	--

Fig. 4. KEM_{PCA} , an IND-PCA-secure KEM under the \mathcal{U}_k -MDDH Assumption, with tag-space $\mathcal{T} = \{0, 1\}^\lambda$. Here, GGen is a prime-order group generator (see Sect. 2.3).

3.2 Security Proof

Theorem 1. *The tag-based Key Encapsulation Mechanism KEM_{PCA} defined in Fig. 4 has perfect correctness. Moreover, if the \mathcal{U}_k -MDDH Assumption holds in \mathbb{G} , KEM_{PCA} is IND-PCA secure. Namely, for any adversary \mathcal{A} , there exists an adversary \mathcal{B} such that $\mathbf{T}(\mathcal{B}) \approx \mathbf{T}(\mathcal{A}) + (Q_{\text{dec}} + Q_{\text{enc}}) \cdot \text{poly}(\lambda)$ and*

$$\text{Adv}_{\text{KEM}_{\text{PCA}}}^{\text{ind-pca}}(\mathcal{A}) \leq (4\lambda + 1) \cdot \text{Adv}_{\mathcal{U}_k, \text{GGen}}^{\text{mddh}}(\mathcal{B}) + (Q_{\text{dec}} + Q_{\text{enc}}) \cdot 2^{-\Omega(\lambda)},$$

where $Q_{\text{enc}}, Q_{\text{dec}}$ are the number of times \mathcal{A} queries EncO, DecO , respectively, and $\text{poly}(\lambda)$ is independent of $\mathbf{T}(\mathcal{A})$.

Proof of Theorem 1. Perfect correctness follows readily from the fact that for all $\mathbf{r} \in \mathbb{Z}_q^k$ and $C = \mathbf{r}^\top \mathbf{M}^\top$, for all $\mathbf{k} \in \mathbb{Z}_q^{3k}$:

$$\mathbf{r}^\top (\mathbf{M}^\top \mathbf{k}) = C \cdot \mathbf{k}.$$

We now prove the IND-PCA security of KEM_{PCA} . We proceed via a series of games described in Figs. 6 and 7 and we use Adv_i to denote the advantage of \mathcal{A} in game G_i . We also give a high-level picture of the proof in Fig. 5, summarizing the sequence of games.

Lemma 4 (G_0 to G_1). *There exists an adversary \mathcal{B}_0 such that $\mathbf{T}(\mathcal{B}_0) \approx \mathbf{T}(\mathcal{A}) + (Q_{\text{enc}} + Q_{\text{dec}}) \cdot \text{poly}(\lambda)$ and*

$$|\text{Adv}_0 - \text{Adv}_1| \leq \text{Adv}_{\mathcal{U}_k, \text{GGen}}^{\text{mddh}}(\mathcal{B}_0) + \frac{1}{q-1},$$

where $Q_{\text{enc}}, Q_{\text{dec}}$ are the number of times \mathcal{A} queries EncO, DecO , respectively, and $\text{poly}(\lambda)$ is independent of $\mathbf{T}(\mathcal{A})$.

Here, we use the MDDH assumption to “tightly” switch the distribution of all the challenge ciphertexts.

game	\mathbf{y} uniform in:	\mathbf{k}'_τ used by EncO and DecO	justification/remark
G_0	$\text{span}(\mathbf{M})$	\mathbf{k}_τ	actual scheme
G_1	\mathbb{Z}^{3k}_q	\mathbf{k}_τ	$\mathcal{U}_{3k,k}$ -MDDH on $[\mathbf{M}]$
$G_{2,i}$	\mathbb{Z}^{3k}_q	$\mathbf{k}_\tau + \mathbf{M}^\perp \text{RF}_i(\tau_i)$	$G_1 \equiv G_{2,0}$
$G_{2,i.1}$	$\tau_{i+1} = 0$:	$\text{span}(\mathbf{M}, \mathbf{M}_0)$	$\mathcal{U}_{3k,k}$ -MDDH on $[\mathbf{M}_0]$
	$\tau_{i+1} = 1$:	$\text{span}(\mathbf{M}, \mathbf{M}_1)$	
$G_{2,i.2}$	$\tau_{i+1} = 0$:	$\mathbf{k}_\tau + \mathbf{M}_0^* \text{RF}_{i+1}^{(0)}(\tau_{i+1}) + \mathbf{M}_1^* \text{RF}_i^{(1)}(\tau_i)$	Cramer-Shoup argument
	$\tau_{i+1} = 1$:	$\text{span}(\mathbf{M}, \mathbf{M}_1)$	
$G_{2,i.3}$	$\tau_{i+1} = 0$:	$\mathbf{k}_\tau + \mathbf{M}_0^* \text{RF}_{i+1}^{(0)}(\tau_{i+1}) + \mathbf{M}_1^* \text{RF}_{i+1}^{(1)}(\tau_{i+1})$	Cramer-Shoup argument
	$\tau_{i+1} = 1$:	$\text{span}(\mathbf{M}, \mathbf{M}_1)$	
$G_{2,i+1}$	\mathbb{Z}^{3k}_q	$\mathbf{k}_\tau + \mathbf{M}^\perp \text{RF}_{i+1}(\tau_{i+1})$	$\mathcal{U}_{3k,k}$ -MDDH on $[\mathbf{M}_0]$ and $[\mathbf{M}_1]$

Fig. 5. Sequence of games for the proof of Theorem 1. Throughout, we have (i) $\mathbf{k}_\tau := \sum_{j=1}^{\lambda} \mathbf{k}_{j,\tau_j}$; (ii) $\text{EncO}(\tau) = ([\mathbf{y}], K_b)$ where $K_0 = [\mathbf{y}^\top \mathbf{k}'_\tau]$ and $K_1 \leftarrow_{\mathbb{R}} \mathbb{G}$; (iii) $\text{DecO}(\tau, [\mathbf{y}], \hat{K})$ computes the encapsulation key $K := [\mathbf{y}^\top \cdot \mathbf{k}'_\tau]$. Here, $(\mathbf{M}_0^*, \mathbf{M}_1^*)$ is a basis for $\text{span}(\mathbf{M}^\perp)$, so that $\mathbf{M}_0^\perp \mathbf{M}_0^* = \mathbf{M}_1^\perp \mathbf{M}_1^* = \mathbf{0}$, and we write $\mathbf{M}^\perp \text{RF}_i(\tau_i) := \mathbf{M}_0^* \text{RF}_i^{(0)}(\tau_i) + \mathbf{M}_1^* \text{RF}_i^{(1)}(\tau_i)$. The second column shows which set \mathbf{y} is uniformly picked from by EncO, the third column shows the value of \mathbf{k}'_τ used by both EncO and DecO.

Setup:	$G_0, G_1, \boxed{G_{2.i}}$	EncO(τ):	$G_0, \boxed{G_1, G_{2.i}}$
$\mathcal{T}_{\text{enc}} = \mathcal{T}_{\text{dec}} := \emptyset; b \leftarrow_{\mathbb{R}} \{0, 1\}$	$\mathcal{G} \leftarrow_{\mathbb{R}} \text{GGen}(1^\lambda); \mathbf{M} \leftarrow_{\mathbb{R}} \mathcal{U}_{3k, k}$	$\mathbf{r} \leftarrow_{\mathbb{R}} \mathbb{Z}_q^k; \mathbf{y} := \mathbf{M}\mathbf{r}; \mathbf{y} \leftarrow_{\mathbb{R}} \mathbb{Z}_q^{3k}$	$K_0 := [\mathbf{y}^\top \cdot \mathbf{k}'_\tau]; K_1 \leftarrow_{\mathbb{R}} \mathbb{G}$
$\mathbf{M}^\perp \leftarrow_{\mathbb{R}} \mathcal{U}_{3k, 2k} \text{ s.t. } \mathbf{M}^\top \mathbf{M}^\perp = \mathbf{0}$	Pick random $\text{RF}_i : \{0, 1\}^i \rightarrow \mathbb{Z}_q^{2k}$	If $\tau \notin \mathcal{T}_{\text{dec}} \cup \mathcal{T}_{\text{enc}}$, return $(C := [\mathbf{y}], K_b)$, and set $\mathcal{T}_{\text{enc}} := \mathcal{T}_{\text{enc}} \cup \{\tau\}$.	Otherwise, return \perp .
$\mathbf{k}_{1,0}, \dots, \mathbf{k}_{\lambda,1} \leftarrow_{\mathbb{R}} \mathbb{Z}_q^{3k}$	For all $\tau \in \{0, 1\}^\lambda, \mathbf{k}_\tau := \sum_{j=1}^\lambda \mathbf{k}_{j, \tau_j}$	$\text{DecO}(\tau, C := [\mathbf{y}], \widehat{K}):$	$G_0, G_1, G_{2.i}$
$\mathbf{k}'_\tau := \mathbf{k}_\tau + \boxed{\mathbf{M}^\perp \text{RF}_i(\tau_i)}$	Return	$K := [\mathbf{y}^\top \cdot \mathbf{k}'_\tau]$	Return $\begin{cases} 1 & \text{if } \widehat{K} = K \wedge \tau \notin \mathcal{T}_{\text{enc}} \\ 0 & \text{otherwise} \end{cases}$
$\text{pk} := (\mathcal{G}, [\mathbf{M}], ([\mathbf{M}^\top \mathbf{k}_{j,\beta}])_{1 \leq j \leq \lambda, 0 \leq \beta \leq 1})$		$\mathcal{T}_{\text{dec}} := \mathcal{T}_{\text{dec}} \cup \{\tau\}$	

Fig. 6. Games $G_0, G_1, G_{2.i}$ (for $1 \leq i \leq \lambda$) for the proof of multi-ciphertext PCA security of KEM_{PCA} in Fig. 4. For all $0 \leq i \leq \lambda$, $\text{RF}_i : \{0, 1\}^i \rightarrow \mathbb{Z}_q^{2k}$ is a random function, and for all $\tau \in \mathcal{T}$, τ_i denotes the i -bit prefix of τ . In each procedure, the components inside a solid (dotted) frame are only present in the games marked by a solid (dotted) frame.

Proof of Lemma 4. To go from G_0 to G_1 , we switch the distribution of the vectors $[\mathbf{y}]$ sampled by EncO, using the Q_{enc} -fold $\mathcal{U}_{3k, k}$ -MDDH Assumption on $[\mathbf{M}]$ (see Definition 4 and Lemma 3).

We build an adversary \mathcal{B}'_0 against the Q_{enc} -fold $\mathcal{U}_{3k, k}$ -MDDH Assumption, such that $\mathbf{T}(\mathcal{B}'_0) \approx \mathbf{T}(\mathcal{A}) + (Q_{\text{enc}} + Q_{\text{dec}}) \cdot \text{poly}(\lambda)$ with $\text{poly}(\lambda)$ independent of $\mathbf{T}(\mathcal{A})$, and

$$|\mathbf{Adv}_0 - \mathbf{Adv}_1| \leq \mathbf{Adv}_{\mathcal{U}_{3k, k}, \text{GGen}}^{Q_{\text{enc}}\text{-mddh}}(\mathcal{B}'_0).$$

This implies the lemma by Lemma 3 (self-reducibility of $\mathcal{U}_{3k, k}$ -MDDH), and Lemma 1 ($\mathcal{U}_{3k, k}$ -MDDH $\Leftrightarrow \mathcal{U}_k$ -MDDH).

Upon receiving a challenge $(\mathcal{G}, [\mathbf{M}] \in \mathbb{G}^{3k \times k}, [\mathbf{H}] := [\mathbf{h}_1 | \dots | \mathbf{h}_{Q_{\text{enc}}}] \in \mathbb{G}^{3k \times Q_{\text{enc}}})$ for the Q_{enc} -fold $\mathcal{U}_{3k, k}$ -MDDH Assumption, \mathcal{B}'_0 picks $b \leftarrow_{\mathbb{R}} \{0, 1\}$, $\mathbf{k}_{1,0}, \dots, \mathbf{k}_{\lambda,1} \leftarrow_{\mathbb{R}} \mathbb{Z}_q^{3k}$, and simulates Setup, DecO as described in Fig. 6. To simulate EncO on its j 'th query, for $j = 1, \dots, Q_{\text{enc}}$, \mathcal{B}'_0 sets $[\mathbf{y}] := [\mathbf{h}_j]$, and computes K_b as described in Fig. 6. \square

Lemma 5 (G_1 to $G_{2.0}$). $|\mathbf{Adv}_1 - \mathbf{Adv}_{2.0}| = 0$.

Proof of Lemma 5. We show that the two games are statistically equivalent. To go from G_1 to $G_{2.0}$, we change the distribution of $\mathbf{k}_{1,\beta} \leftarrow_{\mathbb{R}} \mathbb{Z}_q^{3k}$ for $\beta = 0, 1$, to $\mathbf{k}_{1,\beta} + \mathbf{M}^\perp \text{RF}_0(\varepsilon)$, where $\mathbf{k}_{1,\beta} \leftarrow_{\mathbb{R}} \mathbb{Z}_q^{3k}$, $\text{RF}_0(\varepsilon) \leftarrow_{\mathbb{R}} \mathbb{Z}_q^{2k}$, and $\mathbf{M}^\perp \leftarrow_{\mathbb{R}} \mathcal{U}_{3k, 2k}$ such that $\mathbf{M}^\top \mathbf{M}^\perp = \mathbf{0}$. Note that the extra term $\mathbf{M}^\perp \text{RF}_0(\varepsilon)$ does not appear in pk , since $\mathbf{M}^\top (\mathbf{k}_{1,\beta} + \mathbf{M}^\perp \text{RF}_0(\varepsilon)) = \mathbf{M}^\top \mathbf{k}_{1,\beta}$. \square

Lemma 6 ($G_{2.i}$ to $G_{2.i+1}$). For all $0 \leq i \leq \lambda - 1$, there exists an adversary $\mathcal{B}_{2.i}$ such that $\mathbf{T}(\mathcal{B}_{2.i}) \approx \mathbf{T}(\mathcal{A}) + (Q_{\text{enc}} + Q_{\text{dec}}) \cdot \text{poly}(\lambda)$ and

$$|\text{Adv}_{2.i} - \text{Adv}_{2.i+1}| \leq 4 \cdot \text{Adv}_{\mathcal{U}_{k, \text{GGen}}}^{\text{mddh}}(\mathcal{B}_{2.i}) + \frac{4Q_{\text{dec}} + 2k}{q} + \frac{4}{q-1},$$

where $Q_{\text{enc}}, Q_{\text{dec}}$ are the number of times \mathcal{A} queries EncO, DecO , respectively, and $\text{poly}(\lambda)$ is independent of $\mathbf{T}(\mathcal{A})$.

Proof of Lemma 6. To go from $G_{2.i}$ to $G_{2.i+1}$, we introduce intermediate games $G_{2.i.1}, G_{2.i.2}$ and $G_{2.i.3}$, defined in Fig. 7. We prove that these games are indistinguishable in Lemmas 7, 8, 9, and 10.

Setup:	$G_{2.i}, G_{2.i.1}, G_{2.i.2}, G_{2.i.3}$	EncO(τ):	$G_{2.i}, G_{2.i.1}, G_{2.i.2}, G_{2.i.3}$
$\mathcal{T}_{\text{enc}} = \mathcal{T}_{\text{dec}} := \emptyset; b \leftarrow_{\mathcal{R}} \{0, 1\}$ $\mathcal{G} \leftarrow_{\mathcal{R}} \text{GGen}(1^\lambda); \mathbf{M} \leftarrow_{\mathcal{R}} \mathcal{U}_{3k, k}$ $\mathbf{M}^\perp \leftarrow_{\mathcal{R}} \mathcal{U}_{3k, 2k}$ s.t. $\mathbf{M}^\top \mathbf{M}^\perp = \mathbf{0}$ $\mathbf{M}_0, \mathbf{M}_1 \leftarrow_{\mathcal{R}} \mathcal{U}_{3k, k}$ $\mathbf{M}_0^*, \mathbf{M}_1^* \leftarrow_{\mathcal{R}} \mathcal{U}_{3k, k}$ s.t. $\text{span}(\mathbf{M}^\perp) = \text{span}(\mathbf{M}_0^*, \mathbf{M}_1^*)$ $\mathbf{M}^\top \mathbf{M}_0^* = \mathbf{M}_1^* \mathbf{M}_0^* = \mathbf{0}$ $\mathbf{M}^\top \mathbf{M}_1^* = \mathbf{M}_0^* \mathbf{M}_1^* = \mathbf{0}$	Pick random $\text{RF}_i : \{0, 1\}^i \rightarrow \mathbb{Z}_q^{2k}$. Pick $\text{RF}_{i+1}^{(0)} : \{0, 1\}^{i+1} \rightarrow \mathbb{Z}_q^k$ and $\text{RF}_i^{(1)} : \{0, 1\}^i \rightarrow \mathbb{Z}_q^k$ Pick $\text{RF}_{i+1}^{(0)}, \text{RF}_{i+1}^{(1)} : \{0, 1\}^{i+1} \rightarrow \mathbb{Z}_q^k$. $\mathbf{k}_{1,0}, \dots, \mathbf{k}_{\lambda,1} \leftarrow_{\mathcal{R}} \mathbb{Z}_q^{3k}$ For all $\tau \in \{0, 1\}^\lambda$, $\mathbf{k}_\tau := \sum_{j=1}^\lambda \mathbf{k}_{j, \tau_j}$ $\mathbf{k}'_\tau := \mathbf{k}_\tau + \mathbf{M}^\perp \text{RF}_i(\tau_i)$ $\mathbf{k}'_\tau := \mathbf{k}_\tau + \mathbf{M}_0^* \text{RF}_{i+1}^{(0)}(\tau_{i+1})$ $+ \mathbf{M}_1^* \text{RF}_i^{(1)}(\tau_i)$ $\mathbf{k}'_\tau := \mathbf{k}_\tau + \mathbf{M}_0^* \text{RF}_{i+1}^{(0)}(\tau_{i+1})$ $+ \mathbf{M}_1^* \text{RF}_{i+1}^{(1)}(\tau_{i+1})$	$\mathbf{y} \leftarrow_{\mathcal{R}} \mathbb{Z}_q^{3k}$ If $\tau_{i+1} = 0$: $\mathbf{r} \leftarrow_{\mathcal{R}} \mathbb{Z}_q^k; \mathbf{r}_0 \leftarrow_{\mathcal{R}} \mathbb{Z}_q^k; \mathbf{y} := \mathbf{M}\mathbf{r} + \mathbf{M}_0\mathbf{r}_0$ If $\tau_{i+1} = 1$: $\mathbf{r} \leftarrow_{\mathcal{R}} \mathbb{Z}_q^k; \mathbf{r}_1 \leftarrow_{\mathcal{R}} \mathbb{Z}_q^k; \mathbf{y} := \mathbf{M}\mathbf{r} + \mathbf{M}_1\mathbf{r}_1$ $K_0 := [\mathbf{y}^\top \cdot \mathbf{k}'_\tau];$ $K_1 \leftarrow_{\mathcal{R}} \mathbb{G}$ If $\tau \notin \mathcal{T}_{\text{dec}} \cup \mathcal{T}_{\text{enc}}$, return $([\mathbf{y}], K_b)$ and set $\mathcal{T}_{\text{enc}} := \mathcal{T}_{\text{enc}} \cup \{\tau\}$. Otherwise, return \perp .	$\text{DecO}(\tau, [\mathbf{y}], \widehat{K}) : G_{2.i}, G_{2.i.1}, G_{2.i.2}, G_{2.i.3}$ $K := [\mathbf{y}^\top \mathbf{k}'_\tau]$ Return $\begin{cases} 1 & \text{if } \widehat{K} = K \wedge \tau \notin \mathcal{T}_{\text{enc}} \\ 0 & \text{otherwise} \end{cases}$ $\mathcal{T}_{\text{dec}} := \mathcal{T}_{\text{dec}} \cup \{\tau\}$.
Return $\text{pk} := (\mathcal{G}, [\mathbf{M}], ([\mathbf{M}^\top \mathbf{k}_{j,\beta}])_{1 \leq j \leq \lambda, 0 \leq \beta \leq 1})$			

Fig. 7. Games $G_{2.i}$ (for $0 \leq i \leq \lambda$), $G_{2.i.1}, G_{2.i.2}$ and $G_{2.i.3}$ (for $0 \leq i \leq \lambda - 1$) for the proof of Lemma 6. For all $0 \leq i \leq \lambda$, $\text{RF}_i : \{0, 1\}^i \rightarrow \mathbb{Z}_q^{2k}$, $\text{RF}_i^{(0)}, \text{RF}_i^{(1)} : \{0, 1\}^i \rightarrow \mathbb{Z}_q^k$ are random functions, and for all $\tau \in \mathcal{T}$, we denote by τ_i the i -bit prefix of τ . In each procedure, the components inside a solid (dotted, gray) frame are only present in the games marked by a solid (dotted, gray) frame.

Lemma 7 ($G_{2,i}$ to $G_{2,i.1}$). *For all $0 \leq i \leq \lambda - 1$, there exists an adversary $\mathcal{B}_{2,i.0}$ such that $\mathbf{T}(\mathcal{B}_{2,i.0}) \approx \mathbf{T}(\mathcal{A}) + (Q_{\text{enc}} + Q_{\text{dec}}) \cdot \text{poly}(\lambda)$ and*

$$|\mathbf{Adv}_{2,i} - \mathbf{Adv}_{2,i.1}| \leq 2 \cdot \mathbf{Adv}_{\mathcal{U}_{3k,k}, \text{GGen}}^{\text{mddh}}(\mathcal{B}_{2,i.0}) + \frac{2}{q-1},$$

where $Q_{\text{enc}}, Q_{\text{dec}}$ are the number of times \mathcal{A} queries EncO, DecO, respectively, and $\text{poly}(\lambda)$ is independent of $\mathbf{T}(\mathcal{A})$.

Here, we use the MDDH Assumption to “tightly” switch the distribution of all the challenge ciphertexts. We proceed in two steps, first, by changing the distribution of all the ciphertexts with a tag τ such that $\tau_{i+1} = 0$, and then, for those with a tag τ such that $\tau_{i+1} = 1$. We use the MDDH Assumption with respect to an independent matrix for each step.

Proof of Lemma 7. To go from $G_{2,i}$ to $G_{2,i.1}$, we switch the distribution of the vectors $[\mathbf{y}]$ sampled by EncO, using the Q_{enc} -fold $\mathcal{U}_{3k,k}$ -MDDH Assumption.

We introduce an intermediate game $G_{2,i.0}$ where EncO(τ) is computed as in $G_{2,i.1}$ if $\tau_{i+1} = 0$, and as in $G_{2,i}$ if $\tau_{i+1} = 1$. Setup, DecO are as in $G_{2,i.1}$. We build adversaries $\mathcal{B}'_{2,i.0}$ and $\mathcal{B}''_{2,i.0}$ such that $\mathbf{T}(\mathcal{B}'_{2,i.0}) \approx \mathbf{T}(\mathcal{B}''_{2,i.0}) \approx \mathbf{T}(\mathcal{A}) + (Q_{\text{enc}} + Q_{\text{dec}}) \cdot \text{poly}(\lambda)$ with $\text{poly}(\lambda)$ independent of $\mathbf{T}(\mathcal{A})$, and

Claim 1: $|\mathbf{Adv}_{2,i} - \mathbf{Adv}_{2,i.0}| \leq \mathbf{Adv}_{\mathcal{U}_{3,k}, \text{GGen}}^{\text{Q}_{\text{enc}}\text{-mddh}}(\mathcal{B}'_{2,i.0})$.

Claim 2: $|\mathbf{Adv}_{2,i.0} - \mathbf{Adv}_{2,i.1}| \leq \mathbf{Adv}_{\mathcal{U}_{3k,k}, \text{GGen}}^{\text{Q}_{\text{enc}}\text{-mddh}}(\mathcal{B}''_{2,i.0})$.

This implies the lemma by Lemma 3 (self-reducibility of $\mathcal{U}_{3k,k}$ -MDDH), and Lemma 1 ($\mathcal{U}_{3k,k}$ -MDDH $\Leftrightarrow \mathcal{U}_k$ -MDDH).

Let us prove Claim 1. Upon receiving a challenge $(\mathcal{G}, [\mathbf{M}_0] \in \mathbb{G}^{3k \times k}, [\mathbf{H}] := [\mathbf{h}_1 | \dots | \mathbf{h}_{Q_{\text{enc}}}] \in \mathbb{G}^{3k \times Q_{\text{enc}}})$ for the Q_{enc} -fold $\mathcal{U}_{3k,k}$ -MDDH Assumption with respect to $\mathbf{M}_0 \leftarrow_{\mathbb{R}} \mathcal{U}_{3k,k}$, $\mathcal{B}'_{2,i.0}$ does as follows:

Setup: $\mathcal{B}'_{2,i.0}$ picks $\mathbf{M} \leftarrow_{\mathbb{R}} \mathcal{U}_{3k,k}$, $\mathbf{k}_{1,0}, \dots, \mathbf{k}_{\lambda,1} \leftarrow_{\mathbb{R}} \mathbb{Z}_q^{3k}$, and computes pk as described in Fig. 7. For each τ queried to EncO or DecO, it computes on the fly $\text{RF}_i(\tau_i)$ and $\mathbf{k}'_{\tau} := \mathbf{k}_{\tau} + \mathbf{M}^{\perp} \text{RF}_i(\tau_i)$, where $\mathbf{k}_{\tau} := \sum_{j=1}^{\lambda} \mathbf{k}_{j,\tau_j}$, $\text{RF}_i : \{0,1\}^i \rightarrow \mathbb{Z}_q^{3k}$ is a random function, and τ_i denotes the i -bit prefix of τ (see Fig. 7). Note that $\mathcal{B}'_{2,i.0}$ can compute efficiently \mathbf{M}^{\perp} from \mathbf{M} .

EncO: To simulate the oracle EncO(τ) on its j 'th query, for $j = 1, \dots, Q_{\text{enc}}$, $\mathcal{B}'_{2,i.0}$ computes $[\mathbf{y}]$ as follows:

$$\begin{aligned} \text{if } \tau_{i+1} = 0 : \mathbf{r} &\leftarrow_{\mathbb{R}} \mathbb{Z}_q^k; [\mathbf{y}] := [\mathbf{M}\mathbf{r} + \mathbf{h}_j] \\ \text{if } \tau_{i+1} = 1 : [\mathbf{y}] &\leftarrow_{\mathbb{R}} \mathbb{G}^{3k} \end{aligned}$$

This way, $\mathcal{B}'_{2,i.0}$ simulates EncO as in $G_{2,i.0}$ when $[\mathbf{h}_j] := [\mathbf{M}_0 \mathbf{r}_0]$ with $\mathbf{r}_0 \leftarrow_{\mathbb{R}} \mathbb{Z}_q^k$, and as in $G_{2,i}$ when $[\mathbf{h}_j] \leftarrow_{\mathbb{R}} \mathbb{G}^{3k}$.

DecO: Finally, $\mathcal{B}'_{2,i.0}$ simulates DecO as described in Fig. 7.

Therefore, $|\mathbf{Adv}_{2,i} - \mathbf{Adv}_{2,i.0}| \leq \mathbf{Adv}_{\mathcal{U}_{3k,k}, \text{GGen}}^{\text{Q}_{\text{enc}}\text{-mddh}}(\mathcal{B}'_{2,i.0})$.

To prove Claim 2, we build an adversary $\mathcal{B}''_{2,i.0}$ against the Q_{enc} -fold $\mathcal{U}_{3k,k}$ -MDDH Assumption with respect to a matrix $\mathbf{M}_1 \leftarrow_{\mathbb{R}} \mathcal{U}_{3k,k}$, independent from \mathbf{M}_0 , similarly than $\mathcal{B}'_{2,i.0}$. \square

Lemma 8 ($G_{2.i.1}$ to $G_{2.i.2}$). *For all $0 \leq i \leq \lambda - 1$,*

$$|\mathbf{Adv}_{2.i.1} - \mathbf{Adv}_{2.i.2}| \leq \frac{2Q_{\text{dec}} + 2k}{q},$$

where Q_{dec} is the number of times \mathcal{A} queries DecO.

Here, we use a variant of the Cramer-Shoup information-theoretic argument to move from RF_i to RF_{i+1} , thereby increasing the entropy of \mathbf{k}'_τ computed by Setup. For the sake of readability, we proceed in two steps: in Lemma 8, we move from RF_i to an hybrid between RF_i and RF_{i+1} , and in Lemma 9, we move to RF_{i+1} .

Proof of Lemma 8. In $G_{2.i.2}$, we decompose $\text{span}(\mathbf{M}^\perp)$ into two subspaces $\text{span}(\mathbf{M}_0^*)$ and $\text{span}(\mathbf{M}_1^*)$, and we increase the entropy of the components of \mathbf{k}'_τ which lie in $\text{span}(\mathbf{M}_0^*)$. To argue that $G_{2.i.1}$ and $G_{2.i.2}$ are statistically close, we use a Cramer-Shoup argument [11].

Let us first explain how the matrices \mathbf{M}_0^* and \mathbf{M}_1^* are sampled. Note that with probability at least $1 - \frac{2k}{q}$ over the random coins of Setup, $(\mathbf{M} \parallel \mathbf{M}_0 \parallel \mathbf{M}_1)$ forms a basis of \mathbb{Z}_q^{3k} . Therefore, we have

$$\text{span}(\mathbf{M}^\perp) = \text{Ker}(\mathbf{M}^\top) = \text{Ker}((\mathbf{M} \parallel \mathbf{M}_1)^\top) \oplus \text{Ker}((\mathbf{M} \parallel \mathbf{M}_0)^\top).$$

We pick uniformly \mathbf{M}_0^* and \mathbf{M}_1^* in $\mathbb{Z}_q^{3k \times k}$ that generate $\text{Ker}((\mathbf{M} \parallel \mathbf{M}_1)^\top)$ and $\text{Ker}((\mathbf{M} \parallel \mathbf{M}_0)^\top)$, respectively (see Fig. 1). This way, for all $\tau \in \{0, 1\}^\lambda$, we can write

$$\mathbf{M}^\perp \text{RF}_i(\tau_i) := \mathbf{M}_0^* \text{RF}_i^{(0)}(\tau_i) + \mathbf{M}_1^* \text{RF}_i^{(1)}(\tau_i),$$

where $\text{RF}_i^{(0)}, \text{RF}_i^{(1)} : \{0, 1\}^i \rightarrow \mathbb{Z}_q^k$ are independent random functions.

We define $\text{RF}_{i+1}^{(0)} : \{0, 1\}^{i+1} \rightarrow \mathbb{Z}_q^k$ as follows:

$$\text{RF}_{i+1}^{(0)}(\tau_{i+1}) := \begin{cases} \text{RF}_i^{(0)}(\tau_i) & \text{if } \tau_{i+1} = 0 \\ \text{RF}_i^{(0)}(\tau_i) + \text{RF}'_i(\tau_i) & \text{if } \tau_{i+1} = 1 \end{cases}$$

where $\text{RF}'_i(\tau_i) : \{0, 1\}^i \rightarrow \mathbb{Z}_q^k$ is a random function independent from $\text{RF}_i^{(0)}$. This way, $\text{RF}_{i+1}^{(0)}$ is a random function.

We show that the outputs of EncO and DecO are statistically close in $G_{2.i.1}$ and $G_{2.i.2}$. We decompose the proof in two cases (delimited with \blacksquare): the queries with a tag $\tau \in \{0, 1\}^\lambda$ such that $\tau_{i+1} = 0$, and the queries with a tag τ such that $\tau_{i+1} = 1$.

Queries with $\tau_{i+1} = 0$:

The only difference between $G_{2.i.1}$ and $G_{2.i.2}$ is that Setup computes \mathbf{k}'_τ using the random function $\text{RF}_i^{(0)}$ in $G_{2.i.1}$, whereas it uses the random function $\text{RF}_{i+1}^{(0)}$ in $G_{2.i.2}$ (see Fig. 7). Therefore, by definition of $\text{RF}_{i+1}^{(0)}$, for all $\tau \in \{0, 1\}^\lambda$ such that $\tau_{i+1} = 0$, \mathbf{k}'_τ is the same in $G_{2.i.1}$ and $G_{2.i.2}$, and the outputs of EncO and DecO are identically distributed. \blacksquare

Queries with $\tau_{i+1} = 1$:

Observe that for all $\mathbf{y} \in \text{span}(\mathbf{M}, \mathbf{M}_1)$ and all $\tau \in \{0, 1\}^\lambda$ such that $\tau_{i+1} = 1$,

$$\begin{aligned}
& \overbrace{\mathbf{y}^\top \left(\mathbf{k}_\tau + \mathbf{M}_0^* \text{RF}_i^{(0)}(\tau_{|i}) + \mathbf{M}_1^* \text{RF}_i^{(1)}(\tau_{|i}) + \boxed{\mathbf{M}_0^* \text{RF}'_i^{(0)}(\tau_{|i})} \right)}^{\text{G}_{2.i.2}} \\
&= \mathbf{y}^\top \left(\mathbf{k}_\tau + \mathbf{M}_0^* \text{RF}_i^{(0)}(\tau_{|i}) + \mathbf{M}_1^* \text{RF}_i^{(1)}(\tau_{|i}) \right) + \underbrace{\mathbf{y}^\top \mathbf{M}_0^* \text{RF}'_i^{(0)}(\tau_{|i})}_{=0} \\
&= \mathbf{y}^\top \cdot \overbrace{\left(\mathbf{k}_\tau + \mathbf{M}_0^* \text{RF}_i^{(0)}(\tau_{|i}) + \mathbf{M}_1^* \text{RF}_i^{(1)}(\tau_{|i}) \right)}^{\text{G}_{2.i.1}}
\end{aligned}$$

where the second equality uses the fact that $\mathbf{M}^\top \mathbf{M}_0^* = \mathbf{M}_1^\top \mathbf{M}_0^* = \mathbf{0}$ and thus $\mathbf{y}^\top \mathbf{M}_0^* = \mathbf{0}$.

This means that:

- the output of EncO on any input τ such that $\tau_{i+1} = 1$ is identically distributed in $\text{G}_{2.i.1}$ and $\text{G}_{2.i.2}$;
- the output of DecO on any input $(\tau, [\mathbf{y}], \hat{K})$ where $\tau_{i+1} = 1$, and $\mathbf{y} \in \text{span}(\mathbf{M}, \mathbf{M}_1)$ is the same in $\text{G}_{2.i.1}$ and $\text{G}_{2.i.2}$.

Henceforth, we focus on the *ill-formed* queries to DecO, namely those corresponding to $\tau_{i+1} = 1$, and $\mathbf{y} \notin \text{span}(\mathbf{M}, \mathbf{M}_1)$. We introduce intermediate games $\text{G}_{2.i.1.j}$, and $\text{G}'_{2.i.1.j}$ for $j = 0, \dots, Q_{\text{dec}}$, defined as follows:

- $\text{G}_{2.i.1.j}$: DecO is as in $\text{G}_{2.i.1}$ except that for the first j times it is queried, it outputs 0 to any ill-formed query. EncO is as in $\text{G}_{2.i.2}$.
- $\text{G}'_{2.i.1.j}$: DecO as in $\text{G}_{2.i.2}$ except that for the first j times it is queried, it outputs 0 to any ill-formed query. EncO is as in $\text{G}_{2.i.2}$.

We show that:

$$\begin{aligned}
\text{G}_{2.i.1} &\equiv \text{G}_{2.i.1.0} \approx_s \text{G}_{2.i.1.1} \approx_s \dots \approx_s \text{G}_{2.i.1.Q_{\text{dec}}} \equiv \text{G}'_{2.i.1.Q_{\text{dec}}} \\
\text{G}'_{2.i.1.Q_{\text{dec}}} &\approx_s \text{G}'_{2.i.1.Q_{\text{dec}}-1} \approx_s \dots \approx_s \text{G}'_{2.i.1.0} \equiv \text{G}_{2.i.2}
\end{aligned}$$

where we denote statistical closeness with \approx_s and statistical equality with \equiv .

It suffices to show that for all $j = 0, \dots, Q_{\text{dec}} - 1$:

Claim 1: in $\text{G}_{2.i.1.j}$, if the $j + 1$ -st query is ill-formed, then DecO outputs 0 with overwhelming probability $1 - 1/q$ (this implies $\text{G}_{2.i.1.j} \approx_s \text{G}_{2.i.1.j+1}$, with statistical difference $1/q$);

Claim 2: in $\text{G}'_{2.i.1.j}$, if the $j + 1$ -st query is ill-formed, then DecO outputs 0 with overwhelming probability $1 - 1/q$ (this implies $\text{G}'_{2.i.1.j} \approx_s \text{G}'_{2.i.1.j+1}$, with statistical difference $1/q$);

where the probabilities are taken over the random coins of Setup.

Let us prove Claim 1. Recall that in $G_{2.i.1.j}$, on its $j + 1$ -st query, $\text{DecO}(\tau, [\mathbf{y}], \widehat{K})$ computes $K := [\mathbf{y}^\top \mathbf{k}'_\tau]$, where $\mathbf{k}'_\tau := (\mathbf{k}_\tau + \mathbf{M}_0^* \text{RF}_i^{(0)}(\tau_{i+1}) + \mathbf{M}_1^* \text{RF}_i^{(1)}(\tau_i))$ (see Fig. 7). We prove that if $(\tau, [\mathbf{y}], \widehat{K})$ is ill-formed, then K is completely hidden from \mathcal{A} , up to its $j + 1$ -st query to DecO . The reason is that the vector $\mathbf{k}_{i+1,1}$ in sk contains some entropy that is hidden from \mathcal{A} . This entropy is “released” on the $j + 1$ -st query to DecO if it is ill-formed. More formally, we use the fact that the vector $\mathbf{k}_{i+1,1} \leftarrow_{\mathbb{R}} \mathbb{Z}_q^{3k}$ is identically distributed as $\mathbf{k}_{i+1,1} + \mathbf{M}_0^* \mathbf{w}$, where $\mathbf{k}_{i+1,1} \leftarrow_{\mathbb{R}} \mathbb{Z}_q^{3k}$, and $\mathbf{w} \leftarrow_{\mathbb{R}} \mathbb{Z}_q^k$. We show that \mathbf{w} is completely hidden from \mathcal{A} , up to its $j + 1$ -st query to DecO .

– The public key pk does not leak any information about \mathbf{w} , since

$$\mathbf{M}^\top (\mathbf{k}_{i+1,1} + \boxed{\mathbf{M}_0^* \mathbf{w}}) = \mathbf{M}^\top \mathbf{k}_{i+1,1}.$$

This is because $\mathbf{M}^\top \mathbf{M}_0^* = \mathbf{0}$.

– The outputs of EncO also hide \mathbf{w} .

- For τ such that $\tau_{i+1} = 0$, \mathbf{k}'_τ is independent of $\mathbf{k}_{i+1,1}$, and therefore, so does $\text{EncO}(\tau)$.
- For τ such that $\tau_{i+1} = 1$, and for any $\mathbf{y} \in \text{span}(\mathbf{M}, \mathbf{M}_1)$, we have:

$$\mathbf{y}^\top (\mathbf{k}'_\tau + \boxed{\mathbf{M}_0^* \mathbf{w}}) = \mathbf{y}^\top \mathbf{k}'_\tau \quad (2)$$

since $\mathbf{M}^\top \mathbf{M}_0^* = \mathbf{M}_1^\top \mathbf{M}_0^* = \mathbf{0}$, which implies $\mathbf{y}^\top \mathbf{M}_0^* = \mathbf{0}$.

– The first j outputs of DecO also hide \mathbf{w} .

- For τ such that $\tau_{i+1} = 0$, \mathbf{k}'_τ is independent of $\mathbf{k}_{i+1,1}$, and therefore, so does $\text{DecO}([\mathbf{y}], \tau, \widehat{K})$.
- For τ such that $\tau_{i+1} = 1$ and $\mathbf{y} \in \text{span}(\mathbf{M}, \mathbf{M}_1)$, the fact that $\text{DecO}(\tau, [\mathbf{y}], \widehat{K})$ is independent of \mathbf{w} follows readily from Equation (2).
- For τ such that $\tau_{i+1} = 1$ and $\mathbf{y} \notin \text{span}(\mathbf{M}, \mathbf{M}_1)$, that is, for an ill-formed query, DecO outputs 0, independently of \mathbf{w} , by definition of $G_{2.i.1.j}$.

This proves that \mathbf{w} is uniformly random from \mathcal{A} 's viewpoint.

Finally, because the $j + 1$ -st query $(\tau, [\mathbf{y}], \widehat{K})$ is ill-formed, we have $\tau_{i+1} = 1$, and $\mathbf{y} \notin \text{span}(\mathbf{M}, \mathbf{M}_1)$, which implies that $\mathbf{y}^\top \mathbf{M}_0^* \neq \mathbf{0}$. Therefore, the value

$$K = [\mathbf{y}^\top (\mathbf{k}'_\tau + \mathbf{M}_0^* \mathbf{w})] = [\mathbf{y}^\top \mathbf{k}'_\tau + \underbrace{\mathbf{y}^\top \mathbf{M}_0^* \mathbf{w}}_{\neq 0}]$$

computed by DecO is uniformly random over \mathbb{G} from \mathcal{A} 's viewpoint. Thus, with probability $1 - 1/q$ over $K \leftarrow_{\mathbb{R}} \mathbb{G}$, we have $\widehat{K} \neq K$, and $\text{DecO}(\tau, [\mathbf{y}], \widehat{K}) = 0$.

We prove Claim 2 similarly, arguing than in $G'_{2.i.1.j}$, the value $K := [\mathbf{y}^\top \mathbf{k}'_\tau]$, where $\mathbf{k}'_\tau := (\mathbf{k}_\tau + \mathbf{M}_0^* \text{RF}_{i+1}^{(0)}(\tau_{i+1}) + \mathbf{M}_1^* \text{RF}_i^{(1)}(\tau_i))$, computed by $\text{DecO}(\tau, [\mathbf{y}], \widehat{K})$ on its $j + 1$ -st query, is completely hidden from \mathcal{A} , up to its $j + 1$ -st query to DecO , if $(\tau, [\mathbf{y}], \widehat{K})$ is ill-formed. The argument goes exactly as for Claim 1. ■

Lemma 9 ($G_{2.i.2}$ to $G_{2.i.3}$). *For all $0 \leq i \leq \lambda - 1$,*

$$|\mathbf{Adv}_{2.i.2} - \mathbf{Adv}_{2.i.3}| \leq \frac{2Q_{\text{dec}}}{q},$$

where Q_{dec} is the number of times \mathcal{A} queries DecO.

Proof of Lemma 9. In $G_{2.i.3}$, we use the same decomposition $\text{span}(\mathbf{M}^\perp) = \text{span}(\mathbf{M}_0^*, \mathbf{M}_1^*)$ as that in $G_{2.i.2}$. The entropy of the components of \mathbf{k}'_τ that lie in $\text{span}(\mathbf{M}_1^*)$ increases from $G_{2.i.2}$ to $G_{2.i.3}$. To argue that these two games are statistically close, we use a Cramer-Shoup argument [11], exactly as for Lemma 8.

We define $\text{RF}_{i+1}^{(1)}\{0, 1\}^{i+1} \rightarrow \mathbb{Z}_q^k$ as follows:

$$\text{RF}_{i+1}^{(1)}(\tau_{i+1}) := \begin{cases} \text{RF}_i^{(1)}(\tau_i) + \text{RF}'_i^{(1)}(\tau_i) & \text{if } \tau_{i+1} = 0 \\ \text{RF}_i^{(1)}(\tau_i) & \text{if } \tau_{i+1} = 1 \end{cases}$$

where $\text{RF}'_i^{(1)} : \{0, 1\}^i \rightarrow \mathbb{Z}_q^k$ is a random function independent from $\text{RF}_i^{(1)}$. This way, $\text{RF}_{i+1}^{(1)}$ is a random function.

We show that the outputs of EncO and DecO are statistically close in $G_{2.i.1}$ and $G_{2.i.2}$. We decompose the proof in two cases (delimited with \blacksquare): the queries with a tag $\tau \in \{0, 1\}^\lambda$ such that $\tau_{i+1} = 0$, and the queries with tag τ such that $\tau_{i+1} = 1$.

Queries with $\tau_{i+1} = 1$:

The only difference between $G_{2.i.2}$ and $G_{2.i.3}$ is that Setup computes \mathbf{k}'_τ using the random function $\text{RF}_i^{(1)}$ in $G_{2.i.2}$, whereas it uses the random function $\text{RF}_{i+1}^{(1)}$ in $G_{2.i.3}$ (see Fig. 7). Therefore, by definition of $\text{RF}_{i+1}^{(1)}$, for all $\tau \in \{0, 1\}^\lambda$ such that $\tau_{i+1} = 1$, \mathbf{k}'_τ is the same in $G_{2.i.2}$ and $G_{2.i.3}$, and the outputs of EncO and DecO are identically distributed. \blacksquare

Queries with $\tau_{i+1} = 0$:

Observe that for all $\mathbf{y} \in \text{span}(\mathbf{M}, \mathbf{M}_0)$ and all $\tau \in \{0, 1\}^\lambda$ such that $\tau_{i+1} = 0$,

$$\begin{aligned} & \overbrace{\mathbf{y}^\top \left(\mathbf{k}_\tau + \mathbf{M}_0^* \text{RF}_{i+1}^{(0)}(\tau_{i+1}) + \mathbf{M}_1^* \text{RF}_i^{(1)}(\tau_i) + \boxed{\mathbf{M}_1^* \text{RF}'_i^{(1)}(\tau_i)} \right)}^{G_{2.i.3}} \\ &= \mathbf{y}^\top \left(\mathbf{k}_\tau + \mathbf{M}_0^* \text{RF}_{i+1}^{(0)}(\tau_{i+1}) + \mathbf{M}_1^* \text{RF}_i^{(1)}(\tau_i) \right) + \underbrace{\mathbf{y}^\top \mathbf{M}_1^* \text{RF}'_i^{(1)}(\tau_i)}_{=0} \\ &= \mathbf{y}^\top \cdot \overbrace{\left(\mathbf{k}_\tau + \mathbf{M}_0^* \text{RF}_{i+1}^{(0)}(\tau_{i+1}) + \mathbf{M}_1^* \text{RF}_i^{(1)}(\tau_i) \right)}^{G_{2.i.2}} \end{aligned}$$

where the second equality uses the fact $\mathbf{M}^\top \mathbf{M}_1^* = \mathbf{M}_0^\top \mathbf{M}_1^* = \mathbf{0}$, which implies $\mathbf{y}^\top \mathbf{M}_1^* = \mathbf{0}$.

This means that:

- the output of EncO on any input τ such that $\tau_{i+1} = 0$ is identically distributed in $G_{2.i.2}$ and $G_{2.i.3}$;
- the output of DecO on any input $(\tau, [\mathbf{y}], \widehat{K})$ where $\tau_{i+1} = 0$, and $\mathbf{y} \in \text{span}(\mathbf{M}, \mathbf{M}_0)$ is the same in $G_{2.i.2}$ and $G_{2.i.3}$.

Henceforth, we focus on the *ill-formed* queries to DecO, namely those corresponding to $\tau_{i+1} = 0$, and $\mathbf{y} \notin \text{span}(\mathbf{M}, \mathbf{M}_0)$. The rest of the proof goes similarly than the proof of Lemma 8. See the latter for further details. \square

Lemma 10 ($G_{2.i.3}$ to $G_{2.i+1}$). *For all $0 \leq i \leq \lambda - 1$, there exists an adversary $\mathcal{B}_{2.i.3}$ such that $\mathbf{T}(\mathcal{B}_{2.i.3}) \approx \mathbf{T}(\mathcal{A}) + (Q_{\text{enc}} + Q_{\text{dec}}) \cdot \text{poly}(\lambda)$ and*

$$|\mathbf{Adv}_{2.i.3} - \mathbf{Adv}_{2.i+1}| \leq 2 \cdot \mathbf{Adv}_{\mathcal{U}_{k,k}, \text{GGen}}^{\text{mddh}}(\mathcal{B}_{2.i.3}) + \frac{2}{q-1}$$

where $Q_{\text{enc}}, Q_{\text{dec}}$ are the number of times \mathcal{A} queries EncO, DecO, respectively, and $\text{poly}(\lambda)$ is independent of $\mathbf{T}(\mathcal{A})$.

Here, we use the MDDH Assumption to “tightly” switch the distribution of all the challenge ciphertexts, as for Lemma 7. We proceed in two steps, first, by changing the distribution of all the ciphertexts with a tag τ such that $\tau_{i+1} = 0$, and then, the distribution of those with a tag τ such that $\tau_{i+1} = 1$, using the MDDH Assumption with respect to an independent matrix for each step.

Proof of Lemma 10. To go from $G_{2.i.3}$ to $G_{2.i+1}$, we switch the distribution of the vectors $[\mathbf{y}]$ sampled by EncO, using the Q_{enc} -fold $\mathcal{U}_{3k,k}$ -MDDH Assumption. This transition is symmetric to the transition between $G_{2.i}$ and $G_{2.i.1}$ (see the proof of Lemma 7 for further details). Finally, we use the fact that for all $\tau \in \{0, 1\}^\lambda$, $\mathbf{M}_0^* \text{RF}_{i+1}^{(0)}(\tau_i) + \mathbf{M}_1^* \text{RF}_{i+1}^{(1)}(\tau_{i+1})$ is identically distributed to $\mathbf{M}^\perp \text{RF}_{i+1}(\tau_{i+1})$, where $\text{RF}_{i+1} : \{0, 1\}^{i+1} \rightarrow \mathbb{Z}_q^{2k}$ is a random function. This is because $(\mathbf{M}_0^*, \mathbf{M}_1^*)$ is a basis of $\text{span}(\mathbf{M}^\perp)$. \square

The proof of Lemma 6 follows readily from Lemmas 7, 8, 9, and 10. \square

Lemma 11 ($G_{2.\lambda}$). $\mathbf{Adv}_{2.\lambda} \leq \frac{Q_{\text{enc}}}{q}$.

Proof of Lemma 11. We show that the joint distribution of all the values K_0 computed by EncO is statistically close to uniform over $\mathbb{G}^{Q_{\text{enc}}}$. Recall that on input τ , EncO(τ) computes

$$K_0 := [\mathbf{y}^\top (\mathbf{k}_\tau + \mathbf{M}^\perp \text{RF}_\lambda(\tau))],$$

where $\text{RF}_\lambda : \{0, 1\}^\lambda \rightarrow \mathbb{Z}_q^{2k}$ is a random function, and $\mathbf{y} \leftarrow_{\mathbb{R}} \mathbb{Z}_q^{3k}$ (see Fig. 6).

We make use of the following properties:

Property 1: all the tags τ queried to EncO, such that $\text{EncO}(\tau) \neq \perp$, are distinct.

Property 2: the outputs of DecO are independent of $\{\text{RF}(\tau) : \tau \in \mathcal{T}_{\text{enc}}\}$.

This is because for all queries $(\tau, [\mathbf{y}], \widehat{K})$ to DecO such that $\tau \in \mathcal{T}_{\text{enc}}$, $\text{DecO}(\tau, [\mathbf{y}], \widehat{K}) = 0$, independently of $\text{RF}_\lambda(\tau)$, by definition of $G_{2.\lambda}$.

Property 3: with probability at least $1 - \frac{Q_{\text{enc}}}{q}$ over the random coins of EncO, all the vectors \mathbf{y} sampled by EncO are such that $\mathbf{y}^\top \mathbf{M}^\perp \neq \mathbf{0}$.

We deduce that the joint distribution of all the values $\text{RF}_\lambda(\tau)$ computed by EncO is uniformly random over $(\mathbb{Z}_q^{2k})^{Q_{\text{enc}}}$ (from Property 1), independent of the outputs of DecO (from Property 2). Finally, from Property 3, we get that the joint distribution of all the values K_0 computed by EncO is statistically close to uniform over $\mathbb{G}^{Q_{\text{enc}}}$, since:

$$K_0 := [\mathbf{y}^\top (\mathbf{k}_\tau + \mathbf{M}^\perp \text{RF}_\lambda(\tau))] = [\mathbf{y}^\top \mathbf{k}_\tau + \underbrace{\mathbf{y}^\top \mathbf{M}^\perp}_{\neq \mathbf{0} \text{ w.h.p.}} \text{RF}_\lambda(\tau)].$$

This means that the values K_0 and K_1 are statistically close, and therefore, $\text{Adv}_3 \leq \frac{Q_{\text{enc}}}{q}$. \square

Finally, Theorem 1 follows readily from Lemmas 4, 5, 6, and 11. \square

4 Multi-ciphertext CCA-secure Public Key Encryption Scheme

4.1 Our Construction

We now describe the optimized IND-CCA-secure PKE scheme. Compared to the PCA-secure KEM from Sect. 3, we add an authenticated (symmetric) encryption scheme $(\text{Enc}_{\text{AE}}, \text{Dec}_{\text{AE}})$, and set the KEM tag τ as the hash value of a suitable part of the KEM ciphertext (as explained in the introduction). A formal definition with highlighted differences to our PCA-secure KEM appears in Fig. 8.

We prove the security under the \mathcal{U}_k -MDDH Assumption, which admits a tight reduction to the standard k -Lin Assumption.

Theorem 2. *The Public Key Encryption scheme PKE_{CCA} defined in Fig. 8 has perfect correctness, if the underlying Authenticated Encryption scheme AE has perfect correctness. Moreover, if the \mathcal{U}_k -MDDH Assumption holds in \mathbb{G} , AE has one-time privacy and authenticity, and \mathcal{H} generates collision resistant hash functions, then PKE_{CCA} is IND-CCA secure. Namely, for any adversary \mathcal{A} , there exist adversaries \mathcal{B} , \mathcal{B}' , \mathcal{B}'' such that $\mathbf{T}(\mathcal{B}) \approx \mathbf{T}(\mathcal{B}') \approx \mathbf{T}(\mathcal{B}'') \approx \mathbf{T}(\mathcal{A}) + (Q_{\text{dec}} + Q_{\text{enc}}) \cdot \text{poly}(\lambda)$ and*

$$\begin{aligned} \text{Adv}_{\text{PKE}_{\text{CCA}}}^{\text{ind-cca}}(\mathcal{A}) &\leq (4\lambda + 1) \cdot \text{Adv}_{\mathcal{U}_k, \mathbb{G}\text{Gen}}^{\text{mddh}}(\mathcal{B}) \\ &\quad + ((4\lambda + 2)Q_{\text{dec}} + Q_{\text{enc}} + Q_{\text{enc}}Q_{\text{dec}}) \cdot \text{Adv}_{\text{AE}}^{\text{ae-ot}}(\mathcal{B}') \\ &\quad + \text{Adv}_{\mathcal{H}}^{\text{cr}}(\mathcal{B}'') + Q_{\text{enc}}(Q_{\text{enc}} + Q_{\text{dec}}) \cdot 2^{-\Omega(\lambda)}, \end{aligned} \quad (3)$$

where Q_{enc} , Q_{dec} are the number of times \mathcal{A} queries EncO, DecO, respectively, and $\text{poly}(\lambda)$ is independent of $\mathbf{T}(\mathcal{A})$.

$\text{Gen}_{\text{PKE}}(1^\lambda):$ $\mathcal{G} \leftarrow_{\text{r}} \text{GGen}(1^\lambda); \mathbf{H} \leftarrow_{\text{r}} \mathcal{H}(1^\lambda); \mathbf{M} \leftarrow_{\text{r}} \mathcal{U}_{3k,k}$ $\mathbf{k}_{1,0}, \dots, \mathbf{k}_{\lambda,1} \leftarrow_{\text{r}} \mathbb{Z}_q^{3k}$ $\text{pk} := \left(\mathcal{G}, [\mathbf{M}], \mathbf{H}, ([\mathbf{M}^\top \mathbf{k}_{j,\beta}])_{1 \leq j \leq \lambda, 0 \leq \beta \leq 1} \right)$ $\text{sk} := (\mathbf{k}_{j,\beta})_{1 \leq j \leq \lambda, 0 \leq \beta \leq 1}$ $\text{Return } (\text{pk}, \text{sk})$	$\text{Enc}_{\text{PKE}}(\text{pk}, M):$ $\mathbf{r} \leftarrow_{\text{r}} \mathbb{Z}_q^k; \mathbf{y} := \mathbf{M}\mathbf{r}$ $\tau := \text{H}([\bar{\mathbf{y}}])$ $\mathbf{k}_\tau := \sum_{j=1}^\lambda \mathbf{k}_{j,\tau_j}$ $K := [\mathbf{r}^\top \cdot \mathbf{M}^\top \mathbf{k}_\tau]$ $\phi := \text{Enc}_{\text{CAE}}(K, M)$ $\text{Return } ([\mathbf{y}], \phi)$ $\text{Dec}_{\text{PKE}}(\text{pk}, \text{sk}, ([\mathbf{y}], \phi)):$ $\tau := \text{H}([\bar{\mathbf{y}}]); \mathbf{k}_\tau := \sum_{j=1}^\lambda \mathbf{k}_{j,\tau_j}$ $K := [\mathbf{y}^\top \mathbf{k}_\tau]$ $\text{Return } \text{Dec}_{\text{CAE}}(K, \phi).$
--	--

Fig. 8. PKE_{CCA} , an IND-CCA-secure PKE. We color in blue the differences with KEM_{PCA} , the IND-PCA-secure KEM in Fig. 4. Here, GGen is a prime-order group generator (see Sect. 2.3), and $\text{AE} := (\text{Enc}_{\text{CAE}}, \text{Dec}_{\text{CAE}})$ is an Authenticated Encryption scheme with key-space $\mathcal{K} := \mathcal{G}$ (see Definition 7).

We note that the Q_{enc} and Q_{dec} factors in (3) are only related to AE. Hence, when using a statistically secure (one-time) authenticated encryption scheme, the corresponding terms in (3) become exponentially small.

Remark 2 (Extension to the Multi-user CCA Security). We only provide an analysis in the multi-ciphertext (but single-user) setting. However, we remark (without proof) that our analysis generalizes to the multi-user, multi-ciphertext scenario, similar to [6, 16, 18]. Indeed, all computational steps (not counting the steps related to the AE scheme) modify all ciphertexts simultaneously, relying for this on the re-randomizability of the \mathcal{U}_k -MDDH Assumption relative to a fixed matrix \mathbf{M} . The same modifications can be made to many PKE_{CCA} simultaneously by using that the \mathcal{U}_k -MDDH Assumption is also re-randomizable across many matrices \mathbf{M}_i . (A similar property for the DDH, DLIN, and bilinear DDH assumptions is used in [6], [16], and [18], respectively.)

We defer the proof of Theorem 2 to the full version of this paper.

Acknowledgments. We would like to thank Jie Chen for insightful and inspiring discussions, and the reviewers for helpful comments. This work was done in part while the first and last authors were visiting the Simons Institute for the Theory of Computing, supported by the Simons Foundation and NSF grant CNS-1523467.

References

1. Abdalla, M., Benhamouda, F., Pointcheval, D.: Public-key encryption indistinguishable under plaintext-checkable attacks. In: Katz, J. (ed.) PKC 2015. LNCS, vol. 9020, pp. 332–352. Springer, Heidelberg (2015)

2. Abe, M., Chase, M., David, B., Kohlweiss, M., Nishimaki, R., Ohkubo, M.: Constant-size structure-preserving signatures: generic constructions and simple assumptions. In: Wang, X., Sako, K. (eds.) ASIACRYPT 2012. LNCS, vol. 7658, pp. 4–24. Springer, Heidelberg (2012)
3. Abe, M., David, B., Kohlweiss, M., Nishimaki, R., Ohkubo, M.: Tagged one-time signatures: tight security and optimal tag size. In: Kurosawa, K., Hanaoka, G. (eds.) PKC 2013. LNCS, vol. 7778, pp. 312–331. Springer, Heidelberg (2013)
4. Abe, M., Gennaro, R., Kurosawa, K.: Tag-KEM/DEM: a new framework for hybrid encryption. *J. Cryptology* **21**(1), 97–130 (2008)
5. Attrapadung, N., Hanaoka, G., Yamada, S.: A framework for identity-based encryption with almost tight security. Cryptology ePrint Archive, Report 2015/566 (2015). <http://eprint.iacr.org/2015/566>
6. Bellare, M., Boldyreva, A., Micali, S.: Public-key encryption in a multi-user setting: security proofs and improvements. In: Preneel, B. (ed.) EUROCRYPT 2000. LNCS, vol. 1807, pp. 259–274. Springer, Heidelberg (2000)
7. Blazy, O., Kiltz, E., Pan, J.: (Hierarchical) Identity-based encryption from affine message authentication. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014, Part I. LNCS, vol. 8616, pp. 408–425. Springer, Heidelberg (2014)
8. Boneh, D., Canetti, R., Halevi, S., Katz, J.: Chosen-ciphertext security from identity-based encryption. *SIAM J. Comput.* **36**(5), 1301–1328 (2007)
9. Camenisch, J., Chandran, N., Shoup, V.: A public key encryption scheme secure against key dependent chosen plaintext and adaptive chosen ciphertext attacks. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 351–368. Springer, Heidelberg (2009)
10. Chen, J., Wee, H.: Fully, (almost) tightly secure IBE and dual system groups. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part II. LNCS, vol. 8043, pp. 435–460. Springer, Heidelberg (2013)
11. Cramer, R., Shoup, V.: Design and analysis of practical public-key encryption schemes secure against adaptive chosen ciphertext attack. *SIAM J. Comput.* **33**(1), 167–226 (2003)
12. Dolev, D., Dwork, C., Naor, M.: Nonmalleable cryptography. *SIAM J. Comput.* **30**(2), 391–437 (2000)
13. Escala, A., Herold, G., Kiltz, E., Ràfols, C., Villar, J.: An algebraic framework for Diffie-Hellman assumptions. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part II. LNCS, vol. 8043, pp. 129–147. Springer, Heidelberg (2013)
14. Goldwasser, S., Micali, S.: Probabilistic encryption. *J. Comput. Syst. Sci.* **28**(2), 270–299 (1984)
15. Gong, J., Chen, J., Dong, X., Cao, Z., Tang, S.: Extended nested dual system groups, revisited. Cryptology ePrint Archive, Report 2015/820 (2015). <http://eprint.iacr.org/>
16. Hofheinz, D., Jager, T.: Tightly secure signatures and public-key encryption. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 590–607. Springer, Heidelberg (2012)
17. Hofheinz, D., Kiltz, E.: Secure hybrid encryption from weakened key encapsulation. In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 553–571. Springer, Heidelberg (2007)
18. Hofheinz, D., Koch, J., Striecks, C.: Identity-based encryption with (almost) tight security in the multi-instance, multi-ciphertext setting. In: Katz, J. (ed.) PKC 2015. LNCS, vol. 9020, pp. 799–822. Springer, Heidelberg (2015)

19. Kiltz, E.: Chosen-ciphertext security from tag-based encryption. In: Halevi, S., Rabin, T. (eds.) TCC 2006. LNCS, vol. 3876, pp. 581–600. Springer, Heidelberg (2006)
20. Kiltz, E., Wee, H.: Quasi-adaptive NIZK for linear subspaces revisited. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015. LNCS, vol. 9057, pp. 101–128. Springer, Heidelberg (2015)
21. Kurosawa, K., Desmedt, Y.G.: A new paradigm of hybrid encryption scheme. In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 426–442. Springer, Heidelberg (2004)
22. Libert, B., Joye, M., Yung, M., Peters, T.: Concise multi-challenge CCA-secure encryption and signatures with almost tight security. In: Sarkar, P., Iwata, T. (eds.) ASIACRYPT 2014, Part II. LNCS, vol. 8874, pp. 1–21. Springer, Heidelberg (2014)
23. Libert, B., Peters, T., Joye, M., Yung, M.: Non-malleability from malleability: simulation-sound quasi-adaptive NIZK proofs and CCA2-secure encryption from homomorphic signatures. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 514–532. Springer, Heidelberg (2014)
24. Libert, B., Peters, T., Joye, M., Yung, M.: Compactly hiding linear spans: Tightly secure constant-size simulation-sound QA-NIZK proofs and applications. Cryptology ePrint Archive, Report 2015/242 (2015). <http://eprint.iacr.org/2015/242>
25. Libert, B., Peters, T., Joye, M., Yung, M.: Compactly hiding linear spans: tightly secure constant-size simulation-sound QA-NIZK proofs and applications. Cryptology ePrint Archive, Report 2015/242 (2015). <http://eprint.iacr.org/>
26. Morillo, P., Ràfols, C., Villar, J.L.: Matrix computational assumptions in multilinear groups. IACR Cryptology ePrint Archive, 2015:353 (2015)
27. Naor, M., Reingold, O.: Number-theoretic constructions of efficient pseudo-random functions. *J. ACM* **51**(2), 231–262 (2004)
28. Naor, M., Yung, M.: Public-key cryptosystems provably secure against chosen ciphertext attacks. In: 22nd ACM STOC, pp. 427–437. ACM Press, May 1990
29. Okamoto, T., Pointcheval, D.: REACT: rapid enhanced-security asymmetric cryptosystem transform. In: Naccache, D. (ed.) CT-RSA 2001. LNCS, vol. 2020, pp. 159–175. Springer, Heidelberg (2001)
30. Rackoff, C., Simon, D.R.: Non-interactive zero-knowledge proof of knowledge and chosen ciphertext attack. In: Feigenbaum, J. (ed.) CRYPTO 1991. LNCS, vol. 576, pp. 433–444. Springer, Heidelberg (1992)