

Efficient Updatable Public-Key Encryption from Lattices

Calvin Abou Haidar^{1,2}, Alain Passelègue^{1,2,3}, and Damien Stehlé^{1,3}

¹ ENS de Lyon

calvin.abou-haidar@ens-lyon.fr

² INRIA

alain.passelegue@cryptolab.co.kr

³ CryptoLab Inc., Lyon, France

damien.stehle@cryptolab.co.kr

Abstract. Updatable public key encryption has recently been introduced as a solution to achieve forward-security in the context of secure group messaging without hurting efficiency, but so far, no efficient lattice-based instantiation of this primitive is known.

In this work, we construct the first LWE-based UPKE scheme with polynomial modulus-to-noise rate, which is CPA-secure in the standard model. At the core of our security analysis is a generalized reduction from the standard LWE problem to (a stronger version of) the Extended LWE problem. We further extend our construction to achieve stronger security notions by proposing two generic transforms. Our first transform allows to obtain CCA security in the random oracle model and adapts the Fujisaki-Okamoto transform to the UPKE setting. Our second transform allows to achieve security against malicious updates by adding a NIZK argument in the update mechanism. In the process, we also introduce the notion of Updatable Key Encapsulation Mechanism (UKEM), as the updatable variant of KEMs. Overall, we obtain a CCA-secure UKEM in the random oracle model whose ciphertext sizes are of the same order of magnitude as that of CRYSTALS-Kyber.

1 Introduction

Secure group messaging aims to allow secure, long-lasting, communication for a large group of users. The larger the group and the longer the communication, the likelier one of the group member gets compromised. When the latter happens, ideally, one would like to guarantee that messages sent before the attack occurred remain hidden to the attacker. This corresponds to the notion of forward security [5,9,30,35,19,18,13] and can be achieved by relying on forward-secure public key encryption (FS-PKE), but it vastly hurts efficiency compared to relying on standard PKE. FS-PKE generates an initial key pair (pk_0, sk_0) which allows to derive a chain of key pairs $(pk_1, sk_1), (pk_2, sk_2), \dots$ where each pk_{t+1} can be derived publicly from pk_t (and sk_{t+1} from sk_t). Hence, the first epoch key pair of an FS-PKE scheme implicitly defines all the subsequent epoch key pairs.

Forward security further requires that it should be hard to go back in the secret key chain, as compromising sk_j should not hurt the confidentiality of messages encrypted under pk_t for $t < j$. Therefore, FS-PKE can be seen as a simple form of hierarchical identity-based encryption (HIBE) [28,18], and tight connections between the notions have been observed [26]. As of today, FS-PKE schemes have similar performances as HIBE constructions, and therefore relying on FS-PKE for building secure group messaging seems inherently inefficient. The extreme alternative is to rely on standard PKE and to require every user to refresh their key pair on a regular basis. This assumes users to be active and online, which is an undesirable assumption in the context of group messaging. Moreover, a user refreshing its own key only guarantees confidentiality of messages it receives (and therefore messages sent by other users) but does not provide any guarantee about messages it sent. For the latter, users have to rely on the willingness of receivers to update their keys.

Updatable public-key encryption (UPKE), recently introduced in [31,4], offers a compromise between the above two approaches by relaxing the update mechanism of FS-PKE: in a UPKE scheme, any user can update any other user's key pair by running an update algorithm with (high-entropy) private coins. As a consequence, a key pair does not need to contain any information about the next epoch key pair as this information can be provided by the external user who proceeds in the update. This change allows to hope for UPKE constructions with similar efficiency as standard PKE, but a sender can now protect the messages it sent by updating the receiver's key.

To be formal, a UPKE scheme consists in a standard PKE scheme (KeyGen , Enc , Dec) augmented with two additional algorithms (UpdatePk , UpdateSk). The UpdatePk algorithm can be run by any user on inputs a target public key pk_t^U of a user U used at epoch t and fresh private coins r . It produces a public key pk_{t+1}^U for user U for epoch $t+1$ as well as an update ciphertext up (encrypted under pk_t^U). The UpdateSk algorithm then allows user U , given an update ciphertext up and its secret key sk_t^U to update the latter to obtain secret key sk_{t+1}^U corresponding to pk_{t+1}^U . Security of UPKE guarantees that ciphertexts encrypted under U 's public key pk_t^U at any epoch t remain secret to an attacker which compromises sk_j^U for $j > t$, as long as any of the updates which occurred between epoch t and epoch j was performed by an honest user (i.e., using private coins unknown to the attacker). This is formalized by the notion of IND-CR-CPA security, in which the adversary can impose updates of the target user's public key with Chosen Randomness (CR) (i.e., providing the private coins used by the update mechanism to the challenger). This has been the main security notion studied so far [4,31,22]. For practical applications, stronger notions are desirable, and were introduced in [22]: first, the adversary could have access to a decryption oracle (using the secret key of the current epoch), which corresponds to CCA security. Second, the adversary could generate malicious updates. The latter notion corresponds to IND-CU-CPA/CCA security, where the adversary provides Chosen Updates (CU) by providing (possibly malicious) updates to the challenger rather

than providing private coins (used to honestly generate updates in the chosen randomness setting).

UPKE has been constructed from various assumptions over the past years. An efficient IND-CR-CPA construction based on the Computational Diffie-Hellman (CDH) assumption and in the random oracle model (ROM) was proposed in [31,4]. Constructions in the standard model were first proposed in [22] from the Learning with Errors (LWE) assumption and from the Decisional Diffie-Hellman (DDH) assumption, but the latter two constructions are mainly of theoretical interest as they rely on bit-by-bit encryption, and circular-secure and leakage-resilient PKE. The LWE-based construction notably relies on super-polynomial modulus-to-noise rate due to the use of the noise flooding technique. Generic transforms from IND-CR-CPA security to IND-CU-CCA security are described in [22] but rely on heavy tools, namely one-time, strong, true-simulation f -extractable Non-Interactive Zero-Knowledge (NIZK) arguments [21]. In [1], an efficient construction based on the Decisional Composite Residuosity (DCR) assumption was proposed. The authors show that a variant of the ElGamal Paillier encryption scheme [17] can be turned into a (standard model) IND-CR-CPA UPKE scheme, and achieve IND-CR-CCA and IND-CU-CCA UPKE by further adding NIZK proofs using the Naor-Yung paradigm [38]. Concrete instantiations of the latter NIZKs are proposed in the random oracle model, resulting in the first efficient IND-CR-CCA and IND-CU-CCA instantiations from the DCR assumption and the strong RSA assumption [8], in the ROM.

1.1 Contributions

We provide the first efficient UPKE instantiation based on the LWE assumption with polynomial modulus-to-noise rate.

First, we construct a UPKE encryption scheme which follows the lines of the PKE scheme from [34], which underlies CRYSTALS-Kyber [11]. The main technicalities lie in its security analysis: (i) we prove our construction to achieve IND-CR-CPA security in the standard model, based on a new assumption which generalizes the extended-LWE assumption defined in [39], and (ii) we show that the latter assumption reduces to the standard LWE assumption by extending the results from [16].

Second, we provide two simple generic transforms which allow to convert any IND-CR-CPA UPKE construction into IND-CR-CCA and IND-CU-CCA UPKE schemes in the ROM. As we aim for practical constructions, we focus on constructing updatable key encapsulation mechanism (UKEM), which we introduce as the updatable variant of KEM. Our first transformation is an adaption of the Fujisaki-Okamoto transform [24] to the context of UPKE and allows to generically transform an IND-CR-CPA UPKE into an IND-CR-CCA UKEM with a minimal cost, in the ROM. The second transformation relies on the existence of a NIZK argument for a specific language. As an important remark, the underlying NIZK only plays a role in the update mechanism and should only satisfy basic properties (namely, a single-theorem NIZK with computational soundness and computational zero-knowledge is sufficient) while prior constructions [22,1]

relied on strong NIZK notions (e.g., statistical-simulation-sound NIZKs for instantiating Naor-Yung). The latter NIZK argument can be efficiently instantiated from [33].

1.2 Technical Overview

We now present our contributions in more details, starting with our IND-CR-CPA UPKE construction.

IND-CR-CPA UPKE from lattices. Our IND-CR-CPA construction follows the LWE-variant of [34]: a public key is an LWE instance (\mathbf{A}, \mathbf{b}) with $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e}$ for $\mathbf{A} \in \mathbb{Z}_q^{n \times n}$ and $\mathbf{s}, \mathbf{e} \leftarrow \mathcal{D}_{\mathbb{Z}^n, \sigma}$, the LWE secret \mathbf{s} being the corresponding secret key. An encryption of a message $\boldsymbol{\mu} \in \mathbb{Z}_p^n$ is a pair $(\mathbf{ct}_0, \mathbf{ct}_1)$ of the form $(\mathbf{X}\mathbf{A} + \mathbf{E}, \mathbf{X}\mathbf{b} + \mathbf{f} + \lfloor q/p \rfloor \cdot \boldsymbol{\mu} \bmod q)$ with $\mathbf{X}, \mathbf{E} \leftarrow \mathcal{D}_{\mathbb{Z}^{n \times n}, \sigma}$, $\mathbf{f} \leftarrow \mathcal{D}_{\mathbb{Z}^n, \sigma}$. Such a ciphertext can be decrypted by rounding $\mathbf{ct}_1 - \mathbf{ct}_0\mathbf{s}$ since:

$$\mathbf{ct}_1 - \mathbf{ct}_0\mathbf{s} = \mathbf{X}\mathbf{b} + \mathbf{f} + \lfloor q/p \rfloor \cdot \boldsymbol{\mu} - (\mathbf{X}\mathbf{A} + \mathbf{E})\mathbf{s} = \mathbf{X}\mathbf{e} + \mathbf{f} - \mathbf{E}\mathbf{s} + \lfloor q/p \rfloor \cdot \boldsymbol{\mu}$$

where the term $\mathbf{X}\mathbf{e} + \mathbf{f} - \mathbf{E}\mathbf{s}$ is small. Updating a key pair is done by sampling small vectors $\mathbf{r}, \boldsymbol{\eta} \leftarrow \mathcal{D}_{\mathbb{Z}^n, \sigma}$. The public key is updated to $(\mathbf{A}, \mathbf{b} + \mathbf{A}\mathbf{r} + \boldsymbol{\eta}) = (\mathbf{A}, \mathbf{A}(\mathbf{s} + \mathbf{r}) + \mathbf{e} + \boldsymbol{\eta})$. The corresponding update ciphertext up is an encryption of \mathbf{r} (which fits in the plaintext space) under the original public key (\mathbf{A}, \mathbf{b}) . The updated secret key is then $\mathbf{s} + \mathbf{r}$. Correctness follows from the correctness of the PKE scheme. We emphasize that the secret key and noise term might have increased in norm, which can hurt correctness of decryption. We provide more details about how we handle this issue later, when we mention concrete instantiations.

Let us now focus on the security analysis. An IND-CR-CPA attacker first sees a public key $(\mathbf{A}, \mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e})$ and can make a first sequence of updates with private coins of its choice $(\mathbf{r}_1, \boldsymbol{\eta}_1), \dots, (\mathbf{r}_{chall}, \boldsymbol{\eta}_{chall})$ before asking for a challenge ciphertext for a pair of plaintexts $(\boldsymbol{\mu}_0, \boldsymbol{\mu}_1)$ at epoch *chall*. The challenge ciphertext is then encrypted under public key $pk_{chall} = (\mathbf{A}, \mathbf{b} + \mathbf{A}\Delta_{chall}^{\mathbf{r}} + \Delta_{chall}^{\boldsymbol{\eta}})$, where $\Delta_{chall}^{\mathbf{r}} = \sum_{i=1}^{chall} \mathbf{r}_i$ and $\Delta_{chall}^{\boldsymbol{\eta}} = \sum_{i=1}^{chall} \boldsymbol{\eta}_i$. It can then ask for an additional sequence of updates $(\mathbf{r}_{chall+1}, \boldsymbol{\eta}_{chall+1}), \dots, (\mathbf{r}_{last}, \boldsymbol{\eta}_{last})$ until it decides to compromise the secret key. When the latter happens, an honest update is performed by the challenger using randomness $\mathbf{r}^*, \boldsymbol{\eta}^*$. Let $\Delta_{last}^{\mathbf{r}}$ and $\Delta_{last}^{\boldsymbol{\eta}}$ denote respectively $\sum_{i=1}^{last} \mathbf{r}_i$ and $\sum_{i=1}^{last} \boldsymbol{\eta}_i$. Then, the adversary's goal is to guess which plaintext was encrypted, given the compromised secret key $\mathbf{s} + \Delta_{last}^{\mathbf{r}} + \mathbf{r}^*$ and the last update ciphertext which encrypts \mathbf{r}^* under public key $(\mathbf{A}, \mathbf{b} + \mathbf{A}\Delta_{last}^{\mathbf{r}} + \Delta_{last}^{\boldsymbol{\eta}})$.

The prior lattice-based construction from [22] has a similar structure (though it is based on the Dual-Regev PKE scheme [25]) and the authors argue about security by using the following observation, which we adapt to our construction for the exposition. First, notice that the final update ciphertext, which encrypts \mathbf{r}^* , can be transformed into an encryption of $-\mathbf{s}$ as we are given $\mathbf{s} + \Delta_{last}^{\mathbf{r}} + \mathbf{r}^*$ and $\Delta_{last}^{\mathbf{r}}$ is known. It then suffices to argue that the scheme is circular-secure, given the compromised secret key (which is additional leakage about \mathbf{s}). To do

so, observe that, for a ciphertext $(ct_0, ct_1) = (\mathbf{XA} + \mathbf{E}, \mathbf{Xb} + \mathbf{f} + \lfloor q/p \rfloor \cdot \boldsymbol{\mu})$, the second term can be re-written as $(\mathbf{XA} + \mathbf{E})\mathbf{s} + \mathbf{Xe} + \mathbf{f} + \lfloor q/p \rfloor \cdot \boldsymbol{\mu} - \mathbf{Es}$, where $\mathbf{XA} + \mathbf{E}$ is the first part ct_0 of the ciphertext. That is, we have:

$$ct_1 = ct_0\mathbf{s} + \mathbf{f} + \mathbf{Xe} - \mathbf{Es} + \lfloor q/p \rfloor \cdot \boldsymbol{\mu} .$$

Therefore, assuming \mathbf{f} is much larger than $\mathbf{Xe} - \mathbf{Es}$, the ciphertext distribution is statistically close to $(ct_0, ct_0\mathbf{s} + \mathbf{f} + \lfloor q/p \rfloor \cdot \boldsymbol{\mu})$. Under the LWE assumption, ct_0 is pseudorandom, and then any (linear) information about the secret \mathbf{s} contained in $\boldsymbol{\mu}$ can be absorbed by the term $ct_0\mathbf{s}$. The Leftover Hash Lemma allows to complete the security analysis by proving that the latter term is statistically close to uniform, as long as \mathbf{s} retains enough min-entropy (in this case, conditioned on the leaked key $\mathbf{s} + \Delta_{last}^r + \mathbf{r}^*$). Hence, using noise-flooding and assuming LWE, the scheme is proven secure. The proof additionally relies on the (key)-homomorphism of Dual-Regev to incorporate updates required by the adversary in the challenge/update ciphertext and keys.

Our analysis deviates from the above and avoids the noise-flooding step. Instead, we directly prove pseudorandomness of the above $\mathbf{XA} + \mathbf{E}$ term. It seems that the LWE assumption for the instance $(\mathbf{A}, \mathbf{XA} + \mathbf{E})$ would suffice, but the issue is that the second term $ct_1 = (\mathbf{XA} + \mathbf{E})\mathbf{s} + \mathbf{Xe} + \mathbf{f} + \lfloor q/p \rfloor \cdot \boldsymbol{\mu} - \mathbf{Es}$ of the above tuple contains information about \mathbf{X} and \mathbf{E} , namely the terms \mathbf{Xe} and $-\mathbf{Es}$. This is similar to the Extended-LWE assumption [39], which claims that pseudorandomness of an LWE instance $(\mathbf{A}, \mathbf{As} + \mathbf{e})$ still holds when the adversary is given an additional hint h computed as $\langle \mathbf{z}, \mathbf{e} \rangle \bmod q$ for a small \mathbf{z} chosen by the adversary independently of \mathbf{A} . However, the latter assumption is not sufficient: in our case, the sample contains a hint about both the error and the secret and, additionally, as we are interested in updatable encryption, the adversary can make update queries before asking for the challenge. To answer such queries, one needs to know \mathbf{A} , which is part of the public key, in advance. We introduce the Hermite Normal Form Adaptive Extended LWE assumption HNF-AextLWE, which precisely states that pseudorandomness of $(\mathbf{A}, \mathbf{As} + \mathbf{e})$ still holds, provided an additional hint of the form $\langle \mathbf{z}_0, \mathbf{s} \rangle + \langle \mathbf{z}_1, \mathbf{e} \rangle + g \bmod q$, with $\mathbf{z}_0, \mathbf{z}_1$ being small vectors arbitrarily chosen by the adversary after it sees \mathbf{A} and g being a small Gaussian noise. Equipped with this assumption, the rest of the proof can be adapted and we are able to prove the IND-CR-CPA security of our UPKE scheme under the HNF-AextLWE assumption. It remains to show that the latter assumption is implied by the standard LWE assumption.

Reduction from LWE. We first make a reduction from the adaptive extended-LWE (AextLWE) problem to the HNF-AextLWE problem. AextLWE generalizes the Extended-LWE problem by allowing the adversary to choose a small vector \mathbf{z} arbitrarily given \mathbf{A} . The reduction adapts the one from LWE to HNF-LWE given in [6] to our setting. It relies on the observation made in [16] that, if $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ for $m \geq 16n + 4 \log \log q$, one can extract an invertible matrix \mathbf{A}_0 from \mathbf{A} together with another matrix $\mathbf{A}_1 \in \mathbb{Z}_q^{m' \times n}$ with $m' = m - 16n - 4 \log \log q$ such that the matrix $\mathbf{A}_1 \cdot \mathbf{A}_0^{-1}$ is uniformly distributed. Importantly, a hint

$\langle \mathbf{z}_0, \mathbf{s}^* \rangle + \langle \mathbf{z}_1, \mathbf{e}^* \rangle + g \bmod q$ for an HNF-AextLWE instance using \mathbf{A}^* can be computed as a hint $\langle \mathbf{z}, \mathbf{e} \rangle + g \bmod q$ for an AextLWE instance using \mathbf{A} .

We then show that LWE reduces to this new adaptive version by showing that taking a larger standard deviation allows the additional Gaussian noise g in the hint $h = \langle \mathbf{z}, \mathbf{e} \rangle + g \bmod q$ to hide the information given by it. Precisely, the standard deviation must be taken larger by a factor $\|\mathbf{z}\|_2$ (which has to be small by definition). The proof technique is similar to that of [?,?], except that we need to show that the adaptive nature of our assumption still allows for a reduction.

The reduction goes as follows: Given an LWE instance (\mathbf{A}, \mathbf{b}) , first send \mathbf{A} to the AextLWE adversary to receive its choice of small hint vector \mathbf{z} . In response, sample an additional error \mathbf{e}' and Gaussian term g' from a well-chosen distribution that depends on the small vector \mathbf{z} chosen by the adversary, and return $\mathbf{b}' = \mathbf{b} + \mathbf{e}'$ and a hint $h = \langle \mathbf{z}, \mathbf{e}' \rangle + g'$.

One can rewrite the hint as $h = \langle \mathbf{z}, \mathbf{e} + \mathbf{e}' \rangle - \langle \mathbf{z}, \mathbf{e} \rangle + g' = \langle \mathbf{z}, \mathbf{e} + \mathbf{e}' \rangle + g$ for $g = -\langle \mathbf{z}, \mathbf{e} \rangle + g'$. If the vector \mathbf{b} is equal to $\mathbf{A}\mathbf{s} + \mathbf{e}$, as we have $\mathbf{b}' = \mathbf{A}\mathbf{s} + (\mathbf{e} + \mathbf{e}')$ and $h = \langle \mathbf{z}, \mathbf{e} + \mathbf{e}' \rangle + g$, it suffices to show that the joint distribution of $\mathbf{e} + \mathbf{e}'$ and g is a spherical Gaussian. This is achieved by applying a convolution lemma to the sum

$$\begin{pmatrix} \mathbf{e} + \mathbf{e}' \\ -\mathbf{z}^T \mathbf{e} + g' \end{pmatrix} = \begin{pmatrix} \mathbf{Id} \\ -\mathbf{z}^T \end{pmatrix} \mathbf{e} + \begin{pmatrix} \mathbf{e}' \\ g' \end{pmatrix},$$

which is possible if the standard deviation is larger by a factor of $\|\mathbf{z}\|_2$. The analysis for the case where of uniform \mathbf{b} is identical.

Combining the above two results, we then obtain an IND-CR-CPA UPKE construction based on the standard LWE assumption, leading to the first lattice-based UPKE with polynomial modulus-to-noise ratio. We now explain how we transform this construction in order to achieve IND-CU-CCA security.

A Fujisaki-Okamoto transform for UPKE. Prior works [22,1] have relied on the Naor-Yung paradigm [38] to achieve CCA-security, which requires simulation-sound NIZK proofs. While this allows to remain in the standard model, efficient instantiations of NIZKs rely on random oracles, which motivates us to consider a ROM-based transform following the Fujisaki-Okamoto transform [24]. As we aim for practical efficiency, we focus on constructing IND-CR-CCA updatable key encapsulation mechanism (UKEM), a notion we introduce in this work. Our transform allows to construct IND-CR-CCA UKEM in the ROM with similar efficiency as that of the underlying IND-CR-CPA UPKE. To encapsulate a key for a target user with public key pk_t (at epoch t), one produces a ciphertext ct as an encryption of a uniform message m with randomness extracted from applying a hash function \mathbf{G} (modeled as a random oracle) to the public key pk_t and the message m . The encapsulated key is defined as $\mathbf{H}(\text{ct}, m)$ for another hash function (also modeled as a random oracle). Decapsulation recovers m by decrypting ct and re-encrypts it to check that ct was properly generated, in which case one computes the key $\mathbf{H}(\text{ct}, m)$. The update mechanism $\text{UpdatePk}, \text{UpdateSk}$ are exactly the same as that of the underlying IND-CR-CPA UPKE scheme. Overall, this

is the same transform as for PKE [27] except that pk_t is fed as input to G . The security analysis follows the standard route for FO analyses: we modify oracles to allow the challenger to simulate the decapsulation oracle without knowledge of the secret key sk . The main change is that we rely on the additional pk_t which is fed as an additional input to the hash function G in order to keep track of possibly valid ciphertexts known by the adversary for each epoch t .

In a concurrent work, Asano et al. [7] define a similar FO transform to build IND-CR-CCA secure UPKEs. The authors point out a weakness in the generic CCA transform from [22]: the latter work does not consider the possibility of updates of the public key that would allow the adversary to come back to the challenge public key and then trivially break security by querying the CCA decryption oracle on the ciphertext. This is allowed as in [22], this query is forbidden only at the challenge epoch. This is solved in [7] by generalizing the technique of [1], which adjoins a counter to the public key that is incremented at each update. The construction of [7] relies on using this counter in the derandomization step of their FO transform, which then makes any ciphertext generated in a previous epoch invalid for decryption queries. Our security model for IND-CR-CCA UKEM deals with this problem by adding another sanity check in the decapsulation oracle: we require that the adversary is not allowed to make a decapsulation query of the challenge ciphertext only if its current public key is the same as the challenge one.

Adding security against malicious updates. Next, we extend our IND-CR-CCA construction to achieve IND-CU-CCA security. This is achieved via the standard Naor-Yung “double-encrypt + NIZK” paradigm [38] applied (only) to the update mechanism: a user’s public key is now a pair of public keys (pk_0^L, pk^R) . The first one is an evolving key, for which the user keeps the corresponding secret key sk_0^L , while the second one is never updated and its corresponding secret key is discarded after generation. To update a target public key (pk_t^L, pk^R) used at epoch t , one updates the first key as before by revealing the next epoch public key pk_{t+1}^L and encrypting the private coins r used for the update. However, rather than encrypting r under pk_t^L only, one also encrypts it under pk^R . Additionally, one produces a NIZK argument that the private coins underlying each ciphertext and used for updating the public key match. The encapsulation and decapsulation mechanisms are unchanged (and only use pk_t^L).

These changes allow us to argue about IND-CU-CCA security using standard techniques. Let $up^* = (ct_L^*, ct_R^*, \pi^*)$ denote the honest update generated by the challenger before leaking the secret key, and r^* denote the underlying private coins. In the IND-CU-CCA security reduction, one can then replace π^* by a simulated proof and ct_R^* by an encryption of 0 using the zero-knowledge property and the IND-CPA security of the underlying PKE, since no information about sk^R is revealed to the adversary. The soundness of the NIZK argument guarantees that the adversary cannot produce an accepting argument for invalid updates. Hence, security can be reduced to that of the underlying IND-CR-CCA UKEM: the IND-CR-CCA attacker can use the additional key sk_R to decrypt the private coins r used by the IND-CU-CCA adversary in its valid updates queries,

and forward r to its IND-CR-CCA challenger for producing the same update. A crucial remark is that the adversary gets to see an update (and then a NIZK argument) generated by the challenger only at the very end of the game, when it compromises the key. In particular, it can no longer query oracles from this point and therefore cannot use this proof as part of oracle queries. This allows us to rely on a NIZK argument which is only computational zero-knowledge.

Concrete parameters. We provide concrete parameters for our (IND-CR-CPA / IND-CR-CCA) scheme, following design choices of CRYSTALS-Kyber [11]: we instantiate our construction in the module lattices setting, using binomial distributions. In particular, we assume that our scheme is secure in the module setting though our security analysis does not immediately carries over to the Module Learning With Errors (MLWE) setting [15,32]. To extend it, one would need a similar reduction from decision entropic-MLWE to MLWE, which is currently lacking though a recent work from [12] shows a reduction for the search variants, providing a first step in this direction.

Notice that, as our modulus is small and the key can keep growing with (adversarially generated) updates, we can only guarantee correctness for a bounded number of updates as the decryption error might become too large at some point. We introduce a parameter k which is the maximal number of updates for which correctness is guaranteed with probability extremely close to 1. This parameter affects the size of the modulus q and forces us to use a larger modulus compared to Kyber (which uses $q = 3329$ and achieves a ciphertext size of 0.8KB for 128 bit CCA security). Note that in practice, if randomness is honestly sampled from centered distribution (e.g., $\mathbf{r} \leftarrow U(\{-1, 0, 1\}^n)$), the expected number of supported updates is $O(k^2)$. In Table 1, we provide parameters for our IND-CR-CPA/CCA UKEM schemes, for $k \in \{2^5, 2^{10}, 2^{15}, 2^{20}\}$, and for a security of λ close to 128 bits.

	λ	q	k	$ ct $	$ up $
DCR-based construction [1]	128	-	∞	8.3KB	1.5KB
Estimate for [22]	120	$\approx 2^{85}$	2^5	33KB	360KB
This work	128	$\approx 2^{21}$	2^5	1.8KB	5.4KB
	128	$\approx 2^{26}$	2^{10}	3.0KB	12KB
	116	$\approx 2^{31}$	2^{15}	5.8KB	12KB
	128	$\approx 2^{36}$	2^{20}	9.1KB	27KB

Table 1. Concrete parameters for our IND-CR-CCA UKEM.

We provide a brief comparison with the DCR-based (IND-CR-CPA) construction of [1], whose ciphertext/update size is about 1.5KB. Note that in the latter work, the authors achieve CCA-security by adding NIZKs, which hurts their ciphertext size for the CCA setting (about 8.3KB for 128 bits of security), while using our FO transform leaves us with the same numbers for our IND-CR-CCA

construction. In order to give an insight on the efficiency gain compared to the construction of [22] (which was not meant to be efficient), we provide estimates of practical parameters for their scheme. As it requires flooding, we first make the assumption that flooding by 64 bits suffices (see [37]). In order to give optimistic parameters, we relax their statistical leftover hash lemma to a computational one, i.e., we use an adaptation of the scheme from [34] rather than dual Regev encryption. This leads to considering parameters for our scheme but with flooding. Also, to achieve IND-CR-CCA security, we apply our efficient FO transform and not their generic one.

2 Preliminaries

We start by giving out the mathematical background and some useful lemmas needed in this paper.

Throughout this paper, we use bold upper case letters to denote matrices (\mathbf{A}), bold lower case letters for vectors (\mathbf{a}) and italic letters for scalars (a). For any vector $\mathbf{x} = (x_1, \dots, x_n)$, we use the ℓ_2 -norm $\|\mathbf{x}\|_2 = \sqrt{\sum x_i^2}$, the ℓ_1 -norm $\|\mathbf{x}\|_1 = \sum |x_i|$ and the ℓ_∞ -norm $\|\mathbf{x}\|_\infty = \max |x_i|$. For any matrix $\mathbf{A} = (\mathbf{a}_1 \parallel \dots \parallel \mathbf{a}_n)$, we define $\|\mathbf{F}\|_2 = \max \|\mathbf{a}_i\|_2$, $\|\mathbf{F}\|_1 = \max \|\mathbf{a}_i\|_1$ and $\|\mathbf{F}\|_\infty = \max \|\mathbf{a}_i\|_\infty$. We let $\lfloor \cdot \rfloor$ denote the floor function and $\lceil \cdot \rceil$ denote the rounding to the closest integer with ties being rounded up, which are extended to vectors by considering their coefficient-wise application. For $\mathbf{x} \in \mathbb{Q}^n$ and $q > p > 0$, we write $\lfloor \mathbf{x} \rfloor_{p,q}$ for $\lfloor p/q \cdot \mathbf{x} \bmod q \rfloor$. In this work, the modulus q will always be implicit and omitted.

For a distribution \mathcal{S} , we note $s \leftarrow \mathcal{S}$ the fact that s is sampled using distribution \mathcal{S} . For a random variable X , we write $X \sim \mathcal{S}$ if X follows the distribution \mathcal{S} . We let $\mathcal{B}(p)$ denote the Bernoulli distribution of parameter p . We write $a \approx_\delta b$ for $a, b, \delta > 0$ if there exists $\varepsilon < \delta$ such that $|a - b| = \varepsilon$.

We say an algorithm is PPT if it is probabilistic, polynomial-time. We use \log to denote the logarithm in base 2 and \ln to denote the logarithm in base e .

We use the convolution product to express the distribution of a sum of random variables, which we remind below as well as some additional basic operations and properties of probability distributions and discrete Gaussian distributions.

Definition 1 (Convolution). *Let $m \in \mathbb{N}$. Let $\mathcal{S}_1, \mathcal{S}_2$ be two probability distribution on \mathbb{Z}^m . We define the convolution product $\mathcal{S}_1 * \mathcal{S}_2$ as:*

$$\mathcal{S}_1 * \mathcal{S}_2(x) = \sum_{y \in \mathbb{Z}^m} \mathcal{S}_1(x - y) \mathcal{S}_2(y).$$

*If $X \sim \mathcal{S}_1$ and $Y \sim \mathcal{S}_2$ are independent random variables, then $X + Y \sim \mathcal{S}_1 * \mathcal{S}_2$.*

We recall the definition of min-entropy.

Definition 2 (Min-entropy). *Let X, Y be random variables. We define the min-entropy*

$$H_\infty(X) = -\log \left(\max_x \mathbb{P}[X = x] \right)$$

and the average conditional min-entropy:

$$H_\infty(X|Y) = -\log\left(\mathbb{E}_y[\max_x \mathbb{P}[X = x|Y = y]]\right).$$

Definition 3 (Statistical distance). Let $\mathcal{S}_1, \mathcal{S}_2$ be two distributions on \mathbb{Z}^n . We define the statistical $\Delta(\mathcal{S}_1, \mathcal{S}_2)$ as:

$$\Delta(\mathcal{S}_1, \mathcal{S}_2) = \frac{1}{2} \sum_{x \in \mathbb{Z}^n} |\mathcal{S}_1(x) - \mathcal{S}_2(x)|.$$

2.1 Gaussian distributions

We give the definition of Gaussian distribution and several useful lemmas that are used afterwards.

Definition 4 (Gaussian distribution). Let $m \in \mathbb{N}$. For any symmetric positive-definite matrix $\Sigma \in \mathbb{R}^{m \times m}$, define the function $g_\Sigma : \mathbb{R}^m \rightarrow \mathbb{R}$ as

$$\rho_\Sigma(\mathbf{x}) = \exp\left(-\pi \frac{\mathbf{x}^T \Sigma^{-1} \mathbf{x}}{2}\right).$$

We define the Gaussian distribution on \mathbb{Z}^m with center parameter \mathbf{c} and covariance matrix parameter Σ as $\mathcal{D}_{\mathbb{Z}^m, \Sigma, \mathbf{c}}(\mathbf{x}) = \rho_\Sigma(\mathbf{x} - \mathbf{c}) / \rho_\Sigma(\mathbb{Z}^m - \mathbf{c})$. We will also use, for $\sigma > 0$, the notation $\mathcal{D}_{\mathbb{Z}^m, \sigma}$ to denote $\mathcal{D}_{\mathbb{Z}^m, \sigma^2 \mathbf{Id}, \mathbf{0}}$. Additionally, we will let $\mathcal{D}_{\mathbb{Z}^{m \times n}, \sigma}$ denote the distribution obtained by sampling n vectors from $\mathcal{D}_{\mathbb{Z}^m, \sigma}$ and viewing them as the columns of a matrix in $\mathbb{Z}^{m \times n}$.

Lemma 1 (Gaussian tail-bound, [20, Lemma 2.13]). Let $\mathbf{x} \sim \mathcal{D}_{\mathbb{Z}^m, \sigma}$, then for all $t > 1$, we have

$$\mathbb{P}\left[\|\mathbf{x}\|_2 \geq t\sigma \sqrt{\frac{m}{2\pi}}\right] \leq e^{-\frac{m}{2}(1-t)^2}.$$

Lemma 2 (Gaussian convolution, [10, Lemma 4.12]). Let $\mathbf{c}_1, \mathbf{c}_2 \in \mathbb{Z}^n$. Let $X \sim \mathcal{D}_{\mathbb{Z}^n, \sigma, \mathbf{c}_1}$, $Y \sim \mathcal{D}_{\mathbb{Z}^n, \sigma', \mathbf{c}_2}$ and let \mathcal{S} be the distribution followed by $X + Y$. Then, if

$$\left(\frac{1}{\sigma^2} + \frac{1}{\sigma'^2}\right)^{-1/2} > \sqrt{\frac{\ln(2n(1 + \frac{1}{\epsilon}))}{\pi}},$$

then we have the following inequality

$$\Delta\left(\mathcal{S}, \mathcal{D}_{\mathbb{Z}^n, \sqrt{\sigma^2 + \sigma'^2}, \mathbf{c}_1 + \mathbf{c}_2}\right) < \frac{2\epsilon}{1 - \epsilon}.$$

We now state a discrete Gaussian decomposition result.

Lemma 3 (Gaussian decomposition, instantiated from [36, Lemma 1]).

For $m \geq n$, let $\mathbf{F} \in \mathbb{Z}^{m \times n}$ be a matrix and let $s_1(\mathbf{F})$ be the largest singular value of \mathbf{F} . Take $\sigma, \sigma_1 > 0$. Let $\mathbf{e}_1 \sim \mathcal{D}_{\mathbb{Z}^n, \sigma_1}$ and $\mathbf{e}_2 \sim \mathcal{D}_{\mathbb{Z}^m, \Sigma}$ for

$$\Sigma = \sigma^2 \mathbf{Id} - \sigma_1^2 \mathbf{F}^T \mathbf{F} .$$

Then, if $\sigma > \sqrt{2} \sigma_1 s_1(\mathbf{F})$ and $\sigma_1 > \sqrt{2 \ln(2n(1 + 1/\varepsilon)) / \pi}$, we have:

$$\Delta(\mathcal{S}, \mathcal{D}_{\mathbb{Z}^m, \sigma}) < \frac{2\varepsilon}{1 - \varepsilon} ,$$

where \mathcal{S} is the distribution of $\mathbf{F}\mathbf{e}_1 + \mathbf{e}_2$.

In order to apply Lemma 3, one needs to control the ratio $s_1(\mathbf{F})$. This is the purpose of the following result.

Lemma 4 (Adapted from [2, Lemma 8]). *There exists a constant $K > 1$ such that the following holds. For $m \geq 2n$, $\sigma > K\sqrt{n}$ and $\mathbf{F} \leftarrow \mathcal{D}_{\mathbb{Z}^{m \times n}, \sigma}$*

$$\mathbb{P}[s_1(\mathbf{F}) > K\sigma\sqrt{m}] < e^{-m/K} ,$$

where $s_1(\mathbf{F})$ denotes the largest singular value of \mathbf{F}

2.2 Updatable Public Key Encryption

We recall the syntax of Updatable Public Key Encryption (UPKE) and adapt the underlying IND-CR-CPA security notion defined in [22], with a minor modification: we define correctness and security with a bound on the number of updates. This is motivated by the fact that, in our LWE-based scheme, updates make the key slightly larger and then after a (large but polynomial) number of updates, correctness of decryption is no longer guaranteed. This results from the fact that we are able to work over a (small) polynomial modulus.

Definition 5. *(Updatable Public Key Encryption) An updatable public key encryption scheme is a tuple $\text{UPKE} = (\text{KeyGen}, \text{Enc}, \text{Dec}, \text{UpdatePk}, \text{UpdateSk})$ of PPT algorithms with the following syntax:*

- $\text{KeyGen}(1^\lambda)$ takes as input a security parameter 1^λ and outputs a pair (pk, sk) .
- $\text{Enc}(pk, m)$ takes as input a public key pk and a message m and outputs a ciphertext ct .
- $\text{Dec}(sk, ct)$ takes as input a secret key sk and a ciphertext ct and outputs a message m' .
- $\text{UpdatePk}(pk)$ takes as input a public key pk and outputs an update up and a new public key pk' .
- $\text{UpdateSk}(sk, up)$ takes as input a secret key sk and an update up and outputs a new secret key sk' .

(k, δ) -Correctness: Let $(pk_0, sk_0) \leftarrow \text{KeyGen}(1^\lambda)$ be a key pair and $k > 0$ be an integer. For $t < k$, define

$$(up_{t+1}, pk_{t+1}) \leftarrow \text{UpdatePk}(pk_t) \text{ and } sk_{t+1} \leftarrow \text{UpdateSk}(sk_t, up_{t+1}).$$

The UPKE scheme is said to be (k, δ) -correct, for $\delta > 0$, if for all messages m and $t \leq k$

$$\mathbb{P}[\text{Dec}(sk_t, \text{Enc}(pk_t, m)) \neq m] < \delta,$$

where the probability is over the coins of the underlying algorithms.

We give the definition from [22] which we adapt to the bounded number of updates setting by adding a parameter k for the number of updates.

Definition 6 (k -IND-CR-CPA security). Let $k > 0$ be an integer and $(\text{KeyGen}, \text{Enc}, \text{Dec}, \text{UpdatePk}, \text{UpdateSk})$ be a UPKE scheme. Let \mathcal{R} be the randomness space of

UpdatePk . We give the k -IND-CR-CPA security game in Figure 1.

The advantage of \mathcal{A} in winning the above game is

$$\text{Adv}_{\text{UPKE}}^{\text{IND-CR-CPA}}(\mathcal{A}) = \left| \Pr[\beta = \beta'] - \frac{1}{2} \right|.$$

A UPKE scheme is k -IND-CR-CPA-secure if for all PPT attackers \mathcal{A} , the advantage $\text{Adv}_{\text{UPKE}}^{\text{IND-CR-CPA}}(\mathcal{A})$ is negligible.

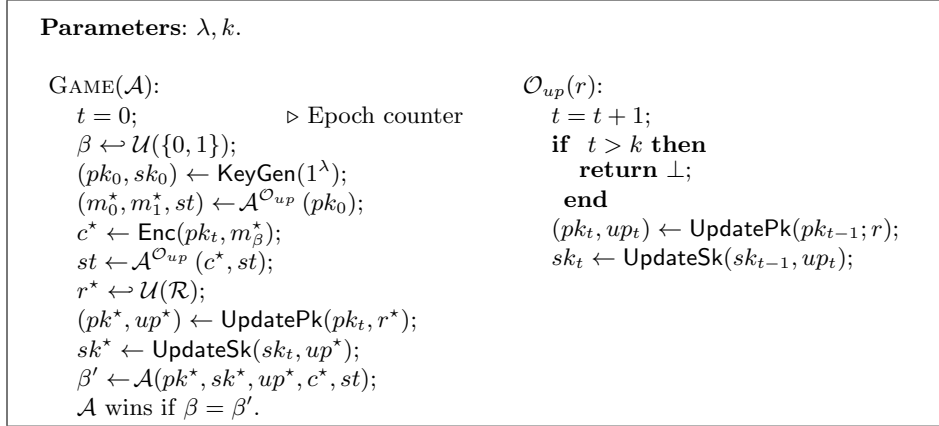


Fig. 1: k -IND-CR-CPA security game.

We also recall the definition of γ -spreadness, which allows to bound the probability that a specific randomness r was used to produce a valid encryption. It is used in Section 5 for our FO transform.

Definition 7 (γ -spreadness, adapted from [23, Section 2.1]). Let $\gamma > 0$. We say that a UPKE $(\text{KeyGen}, \text{Enc}, \text{Dec}, \text{UpdatePk}, \text{UpdateSk})$ is γ -spread if for all m, c and $(pk, sk) \leftarrow \text{KeyGen}(1^\lambda)$, we have

$$\mathbb{P}[\text{Enc}(pk, m) = c] \leq \gamma.$$

2.3 Updatable Key Encapsulation Mechanism

We introduce the KEM variant of UPKE, which we term Updatable KEM or UKEM. Defining the KEM equivalent of UPKE seems particularly relevant considering that UPKE was introduced as a group messaging primitive, hence requiring real-world efficiency.

We adapt the definitions of IND-CR-CCA and IND-CU-CCA security notions defined by [22] for UPKEs.

Definition 8 (Updatable KEM (UKEM)). An updatable KEM is a tuple $(\text{KeyGen}, \text{Encaps}, \text{Decaps}, \text{UpdatePk}, \text{UpdateSk})$ of algorithms with the following syntax:

- $\text{KeyGen}(1^\lambda)$ takes as input a security parameter 1^λ and outputs a pair (pk, sk) .
- $\text{Encaps}(pk)$ takes as input a public key pk and outputs an encapsulation c and a key K .
- $\text{Decaps}(sk, c)$ takes as input a secret key sk and an encapsulation c and outputs a key K' .
- $\text{UpdatePk}(pk)$ takes as input a public key pk and outputs an update up and a new public key pk' .
- $\text{UpdateSk}(sk, up)$ takes as input a secret key sk and an update up and outputs a new secret key sk' .

(k, δ) -Correctness: Let $(pk_0, sk_0) \leftarrow \text{KeyGen}(1^\lambda)$ be a key pair and $k > 0$ be an integer. For $t < k$, define

$$(up_{t+1}, pk_{t+1}) \leftarrow \text{UpdatePk}(pk_t) \text{ and } sk_{t+1} \leftarrow \text{UpdateSk}(sk_t, up_{t+1}).$$

The UKEM scheme is said to be (k, δ) -correct, for $\delta > 0$, if for all $t \leq k$

$$\mathbb{P}[\text{Decaps}(sk_t, c_t) \neq K_t \mid (c_t, K_t) \leftarrow \text{Encaps}(pk_t)] < \delta,$$

where the probability is over the coins of the underlying algorithms.

The k -IND-CR-CCA security corresponds to a variant of k -IND-CR-CPA where the adversary is given access to a decapsulation oracle. We define k -IND-CR-CCA in the Random Oracle Model (ROM), as we make use of the Fujisaki-Okamoto transform in Section 5 in order to build our IND-CR-CCA UKEM.

Definition 9 (k -IND-CR-CCA KEM security in the ROM). Let $(\text{KeyGen}, \text{Encaps}, \text{Decaps}, \text{UpdatePk}, \text{UpdateSk})$ be a UKEM with key space \mathcal{K} . Let \mathcal{R} denote the randomness space of UpdatePk . We give the game for k -IND-CR-CCA security for an adversary that has access to a random oracle H in Figure 2.

Parameters: λ, k .	
GAME (\mathcal{A}): $t = 0;$ \triangleright Epoch counter $\beta \leftarrow \mathcal{U}(\{0, 1\});$ $(pk_0, sk_0) \leftarrow \text{KeyGen}(1^\lambda);$ $st \leftarrow \mathcal{A}^{\mathcal{O}_{up}, \mathcal{O}_{dec}, \mathcal{H}}(pk_0);$ $(c^*, K^*) \leftarrow \text{Encaps}(pk_t);$ if $\beta = 1$ then $\quad K^* = \mathcal{U}(\mathcal{K});$ end $pk^{chall} = pk_t;$ $st \leftarrow \mathcal{A}^{\mathcal{O}_{up}, \mathcal{O}_{dec}, \mathcal{H}}(c^*, st);$ $r^* \leftarrow \mathcal{U}(\mathcal{R});$ $(up^*, pk^*) \leftarrow \text{UpdatePk}(pk_t, r^*);$ $sk^* \leftarrow \text{UpdateSk}(sk_t, up^*);$ $\beta' \leftarrow \mathcal{A}^{\mathcal{H}}(pk^*, sk^*, up^*, c^*, st);$ \mathcal{A} wins if $\beta = \beta'$.	$\mathcal{O}_{up}(r):$ $t = t + 1;$ if $t > k$ then $\quad \text{return } \perp;$ end $(pk_t, up_t) \leftarrow \text{UpdatePk}(pk_{t-1}; r);$ $sk_t \leftarrow \text{UpdateSk}(sk_{t-1}, up_t);$ $\mathcal{O}_{dec}(c):$ if $pk_t = pk^{chall} \wedge c = c^*$ then $\quad \text{return } \perp;$ end return $\text{Decaps}(sk_t, c)$.

Fig. 2: k -IND-CR-CCA security game in the ROM. Note that if $\beta = 0$, then the value of the key K^* is the output of Encaps.

The advantage of \mathcal{A} in winning the above game is

$$\text{Adv}_{\text{UKEM}}^{\text{IND-CR-CCA}}(\mathcal{A}) = \left| \Pr[\beta = \beta'] - \frac{1}{2} \right|.$$

A UKEM scheme is k -IND-CR-CCA-secure if for all PPT attackers \mathcal{A} , the advantage $\text{Adv}_{\text{UKEM}}^{\text{IND-CR-CCA}}(\mathcal{A})$ is negligible.

Notice that compared to the IND-CR-CCA definition for UPKE given in [22], we add a check in the \mathcal{O}_{dec} oracle that the current public key pk_t is different from the challenge public key pk^{chall} . This disallows trivial attacks in which an adversary might make carefully chosen updates that would cancel out in order to get back to the challenge public key and issue a decryption query on the challenge. Another approach to solve this is given in [7], which generalizes the one considered in [1].

In order to define the stronger k -IND-CU-CCA security notions for UKEM, we add an algorithm `VerifyUpdate` to the UKEM syntax that allows a user to check the validity of an update. Specifically, `VerifyUpdate`($pk, (pk', up)$) takes as input the current epoch public key pk and a proposed update (pk', up) and returns a Boolean value. k -IND-CU-CCA security aims to guarantee security against adversaries who makes malicious updates.

Definition 10 (*k -IND-CU-CCA KEM security in the ROM*). Let (KeyGen, Encaps, Decaps, UpdatePk, UpdateSk, VerifyUpdate) be a UKEM. The security game for IND-CU-CCA is identical to the IND-CR-CCA game, except for the modified $\mathcal{O}_{up}(\cdot)$ oracle. We present the modified \mathcal{O}_{up} oracle in Figure 3.

```

 $\mathcal{O}_{up}(pk', up):$ 
  if VerifyUpdate( $pk_t, (pk', up)$ ) =  $\perp$  then
    return  $\perp$ ;
  end
   $pk_{t+1} = pk'$ ;
   $sk_{t+1} \leftarrow \text{UpdateSk}(sk_t, up_{t+1})$ ;
   $t = t + 1$ .

```

Fig. 3: k -IND-CU-CCA security game in the ROM.

A UKEM scheme is k -IND-CU-CCA-secure if for all PPT attackers \mathcal{A} , its advantage $\text{Adv}_{\text{UKEM}}^{\text{IND-CU-CCA}}(\mathcal{A})$ is negligible.

In the rest of the paper, we omit the k in k -IND-CR-CPA/ k -IND-CR-CCA/ k -IND-CU-CCA when it is implicit.

3 Extended LWE

We start by recalling the Learning With Errors (LWE) assumption.

Definition 11. (*Learning With Errors - LWE*) Let $\lambda \geq 0$ be a security parameter. Let $q = q(\lambda), n = n(\lambda), m = m(\lambda) \geq 0$, \mathcal{S} be a distribution on \mathbb{Z}_q^n and χ be an error distribution on \mathbb{Z}^m . The goal of $\text{LWE}_{q,n,m,\chi}(\mathcal{S})$ for an adversary \mathcal{A} is to distinguish between $(\mathbf{A}, \mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e})$ and (\mathbf{A}, \mathbf{u}) , for $\mathbf{A} \leftarrow \mathcal{U}(\mathbb{Z}_q^{m \times n})$, $\mathbf{s} \leftarrow \mathcal{S}$, $\mathbf{e} \leftarrow \chi^m$ and $\mathbf{u} \leftarrow \mathcal{U}(\mathbb{Z}_q^m)$. We define the advantage of \mathcal{A} in the LWE game as

$$\text{Adv}^{\text{LWE}}(\mathcal{A}) := |\mathbb{P}[\mathcal{A}(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}) \rightarrow 1] - \mathbb{P}[\mathcal{A}(\mathbf{A}, \mathbf{u}) \rightarrow 1]| .$$

To keep the notations simple, we write $\text{LWE}_{q,n,m,\sigma}$ for $\sigma > 0$, to denote $\text{LWE}_{q,n,m,\mathcal{D}_{\mathbb{Z}^m,\sigma}}(\mathcal{U}(\mathbb{Z}_q^n))$.

The extended-LWE assumption claims that pseudorandomness of an LWE instance $(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e})$ still holds when the adversary is given an additional hint h computed as $\langle \mathbf{z}, \mathbf{e} \rangle \bmod q$ for a small \mathbf{z} chosen by the adversary independently of \mathbf{A} . We define Adaptive extended-LWE, an adaptive version of this assumption. As the name suggests, it allows the adversary to choose the hint vector \mathbf{z} adaptively, i.e. after having seen the matrix \mathbf{A} , which is not allowed in the definition of the extended-LWE from [39]. In Theorem 1, we prove that LWE reduces to this adaptive version.

Definition 12 (Adaptive extended-LWE - AextLWE). Let $\lambda \geq 0$ be a security parameter. Let $q = q(\lambda), n = n(\lambda), m = m(\lambda), B = B(\lambda) \in \mathbb{N}$ and χ be an error distribution on \mathbb{Z}^m . The goal of $\text{AextLWE}_{q,n,m,\chi,B}$ for an adversary \mathcal{A} is to distinguish between the case where $\beta = 0$ and $\beta = 1$ in the interactive game depicted in Figure 4. We define the advantage of \mathcal{A} in the AextLWE game as

$$\text{Adv}^{\text{AextLWE}}(\mathcal{A}) := |\mathbb{P}[\mathcal{A}(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}, \mathbf{z}, h) \rightarrow 1] - \mathbb{P}[\mathcal{A}(\mathbf{A}, \mathbf{u}, \mathbf{z}, h) \rightarrow 1]| ,$$

where the elements are distributed as shown in Figure 4.

To keep the notations simple, we write $\text{AextLWE}_{q,n,m,\sigma,B}$, for $\sigma > 0$, to denote $\text{AextLWE}_{q,n,m,\mathcal{D}_{\mathbb{Z}^m},\sigma,B}$.




Fig. 4: The decision game for $\text{AextLWE}_{q,n,m,\chi}$.

We define the Hermite Normal Form (HNF) variant of Adaptive extended-LWE, based on the normal form reduction from [6, Lemma 2]. Lemma 5 shows that the HNF variant reduces to the standard Adaptive extended-LWE.

Definition 13 (HNF Adaptive extended-LWE - HNF-AextLWE). Let $\lambda \in \mathbb{N}$ be a security parameter. Let $q = q(\lambda), n = n(\lambda), m = m(\lambda), B = B(\lambda) \in \mathbb{N}$ and χ be an error distribution on \mathbb{R}^m . The goal of $\text{HNF-AextLWE}_{q,n,m,\chi,B}$ for an adversary \mathcal{A} is to distinguish between the case where $\beta = 0$ and $\beta = 1$ in the interactive game depicted in Figure 5. We define the advantage of \mathcal{A} in the HNF-AextLWE game as

$$\text{Adv}^{\text{HNF-AextLWE}}(\mathcal{A}) = |\mathbb{P}[\mathcal{A}(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}, \mathbf{z}_0, \mathbf{z}_1, h) \rightarrow 1] - \mathbb{P}[\mathcal{A}(\mathbf{A}, \mathbf{u}, \mathbf{z}_0, \mathbf{z}_1, h) \rightarrow 1]|$$

where the elements are distributed as shown in Figure 5.

To keep the notations simple, we write $\text{HNF-AextLWE}_{q,n,m,\sigma,B}$, for $\sigma > 0$, to denote $\text{HNF-AextLWE}_{q,n,m,\mathcal{D}_{\mathbb{Z}^m},\sigma,B}$.

Multiple-secret variants. We consider the multiple-secret variants of all our assumptions $\text{Asp} \in \{\text{LWE}, \text{AextLWE}, \text{HNF-AextLWE}\}$ which consist in considering k distinct secrets for the same public matrix \mathbf{A} , thus replacing the secret vector $\mathbf{s} \in \mathbb{Z}_q^n$ by a secret matrix $\mathbf{S} \in \mathbb{Z}_q^{n \times k}$ and the error vector \mathbf{e} by an error matrix $\mathbf{E} \in \mathbb{Z}_q^{m \times k}$. Note that for AextLWE and HNF-AextLWE , the hint $h \in \mathbb{Z}_q$ also becomes a vector $\mathbf{h} \in \mathbb{Z}_q^k$. Also, the multiple-secret variants for AextLWE and HNF-AextLWE could allow for a different \mathbf{z} for each secret, but we restrict




Fig. 5: The decision game for $\text{HNF-AextLWE}_{q,n,m,\chi}$.

ourselves to the case where the \mathbf{z} is the same for all secrets, as it is all we need for our proofs.

Using a hybrid argument, one can show that for every adversary \mathcal{A} for the multiple-secret variant of Asp with k secrets, there exists an adversary \mathcal{B} with a similar run-time against the single-secret problem Asp such that \mathcal{A} 's advantage is bounded by $k \cdot \text{Adv}^{\text{Asp}}(\mathcal{B})$.

Lemma 5. *Let $q \geq 25, n \geq 1, m \geq 16n + 4 \log \log q$, then any adversary \mathcal{A} for $\text{HNF-AextLWE}_{q,n,m',\sigma,B}$, where $m' = m - 16n - 4 \log \log q$, running in time T can be used to build an adversary \mathcal{B} for $\text{AextLWE}_{q,n,m,\sigma,B}$ running in time $\approx T$, with advantage*

$$\text{Adv}^{\text{HNF-AextLWE}}(\mathcal{A}) \leq 4 \cdot \text{Adv}^{\text{AextLWE}}(\mathcal{B}) .$$

Proof. Assume \mathcal{A} is an adversary against HNF-AextLWE . We construct an adversary \mathcal{B} against AextLWE with the claimed advantage as follows.

Adversary \mathcal{B} receives a matrix $\mathbf{A} = (\mathbf{A}_0^T \parallel \mathbf{A}_1^T)^T \in \mathbb{Z}_q^{m \times n}$ from the AextLWE challenger, with $\mathbf{A}_0 \in \mathbb{Z}_q^{n \times n}$ and $\mathbf{A}_1 \in \mathbb{Z}_q^{m-n \times n}$. According to [16, Claim 2.13], with probability at least $1 - 2e^{-1} \geq 1/4$, there exist n linearly independent rows within the first $16n + 4 \log \log q$ rows of \mathbf{A} and an efficient way to find them, so that \mathcal{B} can reorder the matrix so that \mathbf{A}_0 is invertible. If it cannot find such n rows, adversary \mathcal{B} aborts. To avoid keeping track of the indices for the reordering, assume that \mathbf{A} is such that \mathbf{A}_0 is invertible and denote by \mathbf{A}_d the last $15n + 4 \log \log q$ rows of \mathbf{A}_1 so that $\mathbf{A}_1 = (\tilde{\mathbf{A}}_1^T \parallel \mathbf{A}_d^T)^T$ with $\tilde{\mathbf{A}}_1 \in \mathbb{Z}_q^{m' \times n}$.

It then computes $\mathbf{A}^* = -\tilde{\mathbf{A}}_1 \mathbf{A}_0^{-1} \in \mathbb{Z}_q^{m' \times n}$ and sends \mathbf{A}^* to adversary \mathcal{A} . Adversary \mathcal{A} responds with the hint vectors $\mathbf{z}_0 \in \mathbb{Z}_q^n, \mathbf{z}_1 \in \mathbb{Z}_q^{m'}$. Then, adversary \mathcal{B} forwards $\mathbf{z} = (\mathbf{z}_0^T \parallel \mathbf{z}_1^T \parallel 0^{m-m'-n})^T \in \mathbb{Z}_q^m$ to its challenger and receives a vector $\mathbf{b} = (\mathbf{b}_0^T \parallel \mathbf{b}_1^T \parallel \mathbf{d}^T)^T$ and a hint $h = \langle \mathbf{z}, \mathbf{e} \rangle + g \pmod q$ from the AextLWE

challenger, with $\mathbf{b}_0 \in \mathbb{Z}_q^n$, $\mathbf{b}_1 \in \mathbb{Z}_q^{m'}$, $\mathbf{d} \in \mathbb{Z}_q^{m-m'}$ and $g \leftarrow \mathcal{D}_{\mathbb{Z},\sigma}$. It then computes $\mathbf{b}^* = \mathbf{b}_1 + \mathbf{A}^* \mathbf{b}_0$ and sends (\mathbf{b}^*, h) to \mathcal{A} . Finally, it receives a response bit β from \mathcal{A} , which it forwards to its challenger.

In the case where \mathbf{b} was a uniform vector, as \mathbf{A}_0 is an invertible matrix, matrix \mathbf{A}^* is uniform and so is $\mathbf{b}^* = \mathbf{b}_1 + \mathbf{A}^* \mathbf{b}_0$.

If we are in the case where

$$\begin{pmatrix} \mathbf{b}_0 \\ \mathbf{b}_1 \\ \mathbf{d} \end{pmatrix} = \begin{pmatrix} \mathbf{A}_0 \\ \tilde{\mathbf{A}}_1 \\ \mathbf{A}_d \end{pmatrix} \mathbf{s} + \begin{pmatrix} \mathbf{e}_0 \\ \mathbf{e}_1 \\ \mathbf{e}_d \end{pmatrix}$$

for $\mathbf{s} \leftarrow \mathcal{D}_{\mathbb{Z}^n,\sigma}$, $\mathbf{e}_0 \leftarrow \mathcal{D}_{\mathbb{Z}^n,\sigma}$, $\mathbf{e}_1 \leftarrow \mathcal{D}_{\mathbb{Z}^{m'},\sigma}$ and $\mathbf{e}_d \leftarrow \mathcal{D}_{\mathbb{Z}^{m-m'},\sigma}$, then

$$\mathbf{b}^* = \tilde{\mathbf{A}}_1 \mathbf{s} + \mathbf{e}_1 - \tilde{\mathbf{A}}_1 \mathbf{A}_0^{-1} \mathbf{A}_0 \mathbf{s} + \mathbf{A}^* \mathbf{e}_0 = \mathbf{A}^* \mathbf{e}_0 + \mathbf{e}_1.$$

Furthermore, the hint is exactly

$$\begin{aligned} \langle \mathbf{z}, \mathbf{e} \rangle + g &= \left\langle \mathbf{z}_0^T \|\mathbf{z}_1^T\| 0^{m-m'}, \mathbf{e}_0^T \|\mathbf{e}_1^T\| \mathbf{e}_d^T \right\rangle + g \\ &= \langle \mathbf{z}_0, \mathbf{e}_0 \rangle + \langle \mathbf{z}_1, \mathbf{e}_1 \rangle + g \pmod{q}. \end{aligned}$$

Consequently, adversary \mathcal{A} receives a valid HNF Adaptive extended-LWE instance.

Adversary \mathcal{B} runs \mathcal{A} only once and has to compute the reordering which is feasible in time $\text{poly}(\lambda)$. It has advantage at least $\text{Adv}^{\text{HNF-AextLWE}}(\mathcal{A})/4$, completing the proof of the lemma. \square

We now show that LWE reduces to Adaptive extended-LWE .

Theorem 1. *Let q be a prime, $\varepsilon > 0$ and $n, m, B, \gamma, \sigma \geq 0$. Assume that $\sigma > \sqrt{2 \ln(2(n+1)(1+1/\varepsilon))}/\pi$ and $\gamma > \sigma \sqrt{2(1+nB^2)}$. Then for any adversary \mathcal{A} for $\text{AextLWE}_{q,n,m,\sigma,B}$ running in time T , there exists an adversary \mathcal{B} for $\text{LWE}_{q,n,m,\gamma,B}$ running in time $\text{poly}(m, \log q) \cdot T$ such that:*

$$\text{Adv}^{\text{AextLWE}}(\mathcal{A}) \leq \text{Adv}^{\text{LWE}}(\mathcal{B}) + 2 \frac{2\varepsilon}{1-\varepsilon} .$$

Proof. Let \mathcal{A} be an adversary against $\text{AextLWE}_{q,n,m,\sigma}$. We build an adversary \mathcal{B} against $\text{LWE}_{q,n,m,\gamma}$ as follows. Adversary \mathcal{B} receives from its LWE challenger a tuple (\mathbf{A}, \mathbf{b}) . It forwards \mathbf{A} to \mathcal{A} and receives a small hint vector \mathbf{z} such that $\|\mathbf{z}\|_\infty \leq B$.

It then samples $[\mathbf{e}'^T \|g'\|^T] \leftarrow \mathcal{D}_{\mathbb{Z}^{n+1}, \Sigma}$, for some Σ defined later on. It sets $\mathbf{b}' = \mathbf{b} + \mathbf{e}'$ and $h = \langle \mathbf{z}, \mathbf{e}' \rangle + g'$ and sends (\mathbf{b}', h) to \mathcal{A} . Adversary \mathcal{B} receives a final bit from \mathcal{A} which it forwards to its challenger.

Assume we are in the $\beta = 0$ case of the LWE game. Then $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e}$ and $\mathbf{b}' = \mathbf{A}\mathbf{s} + (\mathbf{e} + \mathbf{e}')$. The hint $h = \langle \mathbf{z}, \mathbf{e}' \rangle + g'$ can be rewritten as $h = \langle \mathbf{z}, (\mathbf{e} + \mathbf{e}') \rangle - \langle \mathbf{z}, \mathbf{e} \rangle + g'$. Notice that if

$$\begin{pmatrix} \mathbf{e} + \mathbf{e}' \\ -\langle \mathbf{z}, \mathbf{e} \rangle + g' \end{pmatrix} = \begin{pmatrix} \mathbf{e} \\ -\mathbf{z}^T \mathbf{e} \end{pmatrix} + \begin{pmatrix} \mathbf{e}' \\ g' \end{pmatrix} \sim \mathcal{D}_{\mathbb{Z}^{n+1}, \gamma} \quad (1)$$

this corresponds to the $\beta = 0$ case of the AextLWE game.

It thus suffices to set Σ accordingly. Notice that as $\mathbf{e} \sim \mathcal{D}_{\mathbb{Z}^n, \sigma}$, we have

$$\begin{pmatrix} \mathbf{e} \\ -\mathbf{z}^T \mathbf{e} \end{pmatrix} = \begin{pmatrix} \mathbf{Id} \\ -\mathbf{z}^T \end{pmatrix} \mathbf{e} \sim \mathcal{D}_{\mathbb{Z}^{n+1}, \sigma^2 \mathbf{F} \mathbf{F}^T}$$

for $\mathbf{F} = [\mathbf{Id} \parallel -\mathbf{z}]^T$. Let us then take $\Sigma = \gamma^2 \mathbf{Id} - \sigma^2 \mathbf{F} \mathbf{F}^T$. Note that $s_1(\mathbf{F})^2 = 1 + \|\mathbf{z}\|_2^2 \leq 1 + nB^2$. By assumption, we have $\gamma > \sqrt{2}\sigma s_1(\mathbf{F})$. By applying Lemma 3, we get

$$\begin{pmatrix} \mathbf{e} \\ -\mathbf{z}^T \mathbf{e} \end{pmatrix} + \begin{pmatrix} \mathbf{e}' \\ g' \end{pmatrix} \approx_{\delta} \mathcal{D}_{\mathbb{Z}^{n+1}, \gamma} \quad (2)$$

for $\delta = 2\varepsilon/(1 - \varepsilon)$.

In the $\beta = 1$ case of the LWE game, the vector \mathbf{b} is uniform and so is \mathbf{b}' . The same analysis holds for the distribution of the hint h , so this case matches with the $\beta = 1$ case of the AextLWE game for \mathcal{A} . \square

4 IND-CR-CPA UPKE from LWE

We now describe a UPKE scheme with security based on the HNF-AextLWE assumption. As already shown, it is implied by the standard LWE assumption. Our scheme, detailed in Figure 6, avoid noise flooding by taking advantage of the HNF-AextLWE assumption defined in Section 3. We then provide the first efficient UPKE scheme based on lattices. Our construction follows the lines of [34] which underlies Kyber [11].

In contrast, the only prior lattice-based construction, proposed in [22] and based on the Dual-Regev PKE from [25], is highly inefficient: (i) it supports only binary plaintexts, (ii) updates are done via bit-by-bit encryption of the private coins, and (iii) the security analysis relies on noise flooding, which requires a super-polynomial modulus.

Theorem 2. *Let $\varepsilon, \delta \in (0, 1), k > 0$. Let q, p be primes and $n, m > 0$ and $\sigma, \sigma_c > 0$ such that $\sigma \geq \sqrt{2 \ln(2n(1 + 1/\varepsilon))}/\pi$ and $\sigma_c > 2\sigma \sqrt{1 + n((k + 1)y\sigma)^2}$, where $y = \sqrt{-2 \log(\delta/(4n))}$.*

Assuming the hardness of HNF-AextLWE, the scheme presented in Figure 6 is k -IND-CR-CPA secure. More precisely, for any adversary \mathcal{A} for the k -IND-CR-CPA game, there exists an adversary \mathcal{B} for HNF-AextLWE running in similar time as \mathcal{A} such that:

$$\text{Adv}_{\text{UPKE}}^{\text{IND-CR-CPA}}(\mathcal{A}) \leq (2n + 8) \cdot \frac{2\varepsilon}{1 - \varepsilon} + (2n + 1) \cdot \text{Adv}^{\text{HNF-AextLWE}}(\mathcal{B}).$$

Furthermore, assuming $q > 2p\sigma_c \cdot (2y^2\sigma nk + y)$ and $p > 2y\sigma$, the scheme is (k, δ) -correct.

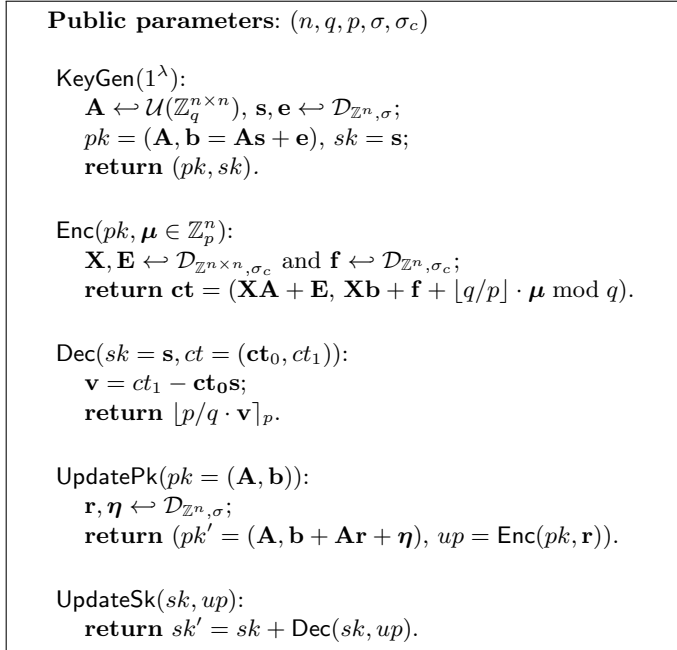


Fig. 6: LWE-based IND-CR-CPA UPKE construction.

Proof. The proof of correctness is detailed in the appendix (Appendix B). We also provide a proof of γ -spreadness there, which is relevant for the next section.

We show the IND-CR-CPA security of the scheme. Let us start by defining all the security games.

Game G_0 : This is the original IND-CR-CPA game. Adversary \mathcal{A} receives $pk_0 = (\mathbf{A}, \mathbf{b}_0 = \mathbf{A}\mathbf{s} + \mathbf{e})$ and queries the $\mathcal{O}_{up}(\cdot)$ oracle with randomness $(\mathbf{r}_1, \boldsymbol{\eta}_1), \dots, (\mathbf{r}_{chall}, \boldsymbol{\eta}_{chall})$ until it asks for a challenge at epoch *chall* for a pair of plaintexts $(\boldsymbol{\mu}_0, \boldsymbol{\mu}_1)$. At this epoch, the secret key is $sk_{chall} = \mathbf{s} + \Delta_{chall}^{\mathbf{r}}$ where $\Delta_{chall}^{\mathbf{r}} = \sum_{i=1}^{chall} \mathbf{r}_i$ and the public key is

$$pk_{chall} = \left(\mathbf{A}, \mathbf{b}_{chall} = \mathbf{A}(\mathbf{s} + \Delta_{chall}^{\mathbf{r}}) + \mathbf{e} + \Delta_{chall}^{\boldsymbol{\eta}} \right),$$

with $\Delta_{chall}^{\boldsymbol{\eta}} = \sum_{i=1}^{chall} \boldsymbol{\eta}_i$. It receives a challenge

$$\mathbf{c}^* = \left(\mathbf{T}_{chall} = \mathbf{X}_{chall}\mathbf{A} + \mathbf{E}_{chall}, \mathbf{pad}_{chall} = \mathbf{X}_{chall}\mathbf{b}_{chall} + \mathbf{f}_{chall} + \lfloor q/p \rfloor \cdot \boldsymbol{\mu}_\beta \right),$$

for $\beta \in \{0, 1\}$ uniform.

Then the adversary queries the $\mathcal{O}_{up}(\cdot)$ oracle until the last epoch *last*. At this epoch, the secret key is $sk_{last} = \mathbf{s} + \Delta_{last}^{\mathbf{r}}$, where $\Delta_{last}^{\mathbf{r}} = \sum_{i=1}^{last} \mathbf{r}_i$ and the public key is $pk_{last} = (\mathbf{A}, \mathbf{b}_{last} = \mathbf{A}(\mathbf{s} + \Delta_{last}^{\mathbf{r}}) + \mathbf{e} + \Delta_{last}^{\boldsymbol{\eta}})$, where $\Delta_{last}^{\boldsymbol{\eta}} = \sum_{i=1}^{last} \boldsymbol{\eta}_i$.

The challenger samples the final update $\mathbf{r}^*, \boldsymbol{\eta}^* \leftarrow \mathcal{D}_{\mathbb{Z}^n, \sigma}$ and sends

$$\begin{aligned} up^* &= \text{Enc}(pk_{last}, \mathbf{r}^*) \\ &= (\mathbf{T}_{last} = \mathbf{X}_{last}\mathbf{A} + \mathbf{E}_{last}, \mathbf{pad}_{last} = \mathbf{X}_{last}\mathbf{b}_{last} + \mathbf{f}_{last} + \lfloor q/p \rfloor \cdot \mathbf{r}^*) \end{aligned}$$

together with $pk^* = (\mathbf{A}, \mathbf{b}_{last} + \mathbf{A}\mathbf{r}^* + \boldsymbol{\eta}^*)$ and $sk^* = \mathbf{s} + \Delta_{last}^{\mathbf{r}} + \mathbf{r}^*$ to the adversary.

Game G_1 : In this game we modify the update up^* . Instead of computing it as

$$up^* = (\mathbf{T}_{last} = \mathbf{X}_{last}\mathbf{A} + \mathbf{E}_{last}, \mathbf{pad}_{last} = \mathbf{X}_{last}\mathbf{b}_{last} + \mathbf{f}_{last} + \lfloor q/p \rfloor \cdot \mathbf{r}^*),$$

the challenger sets

$$up^* = (\mathbf{T}_{last} = \mathbf{X}_{last}\mathbf{A} + \mathbf{E}_{last}, \mathbf{pad}_{last} = \mathbf{X}_{last}\mathbf{b}_{last} + \mathbf{f}_{last} + \lfloor q/p \rfloor \cdot (-\mathbf{s})).$$

This modification results in a computationally equivalent game. Indeed adversary receives up^* together with $sk^* = \mathbf{s} + \Delta_{last}^{\mathbf{r}} + \mathbf{r}^*$ with $\Delta_{last}^{\mathbf{r}}$ known to the adversary. This modification is just a subtraction of $\lfloor q/p \rfloor \cdot (\mathbf{s} + \mathbf{r}^*)$ in \mathbf{pad}_{last} .

Game G_2 : In this game, we again modify the update. This time the challenger computes the update up^* as

$$\begin{aligned} \mathbf{T}_{last} &= \mathbf{X}_{last}\mathbf{A} + \mathbf{E}_{last} - \lfloor q/p \rfloor \cdot \mathbf{Id}, \\ \mathbf{pad}_{last} &= \mathbf{T}_{last}(\mathbf{s} + \Delta_{last}^{\mathbf{r}}) - \mathbf{E}_{last}(\mathbf{s} + \Delta_{last}^{\mathbf{r}}) + \mathbf{X}_{last}(\mathbf{e} + \Delta_{last}^{\boldsymbol{\eta}}) \\ &\quad + \mathbf{f}_{last} + \lfloor q/p \rfloor \cdot \Delta_{last}^{\mathbf{r}}. \end{aligned}$$

Notice that

$$\mathbf{pad}_{last} = \mathbf{X}_{last}\mathbf{b}_{last} + \mathbf{f}_{last} + \lfloor q/p \rfloor \cdot (-\mathbf{s}).$$

Therefore, the only difference with the previous game is that we subtract a publicly computable element $\lfloor q/p \rfloor \cdot \mathbf{Id}$ in \mathbf{T}_{last} , which implies that this game is computationally equivalent to the last one.

Game G_3 : In this game, instead of computing \mathbf{T}_{last} as $\mathbf{T}_{last} = \mathbf{X}_{last}\mathbf{A} + \mathbf{E}_{last} - \lfloor q/p \rfloor \cdot \mathbf{Id}$ the challenger sets \mathbf{T}_{last} uniformly, i.e., $\mathbf{T}_{last} \leftarrow \mathcal{U}(\mathbb{Z}_q^{n \times n})$.

Lemma 6 below states that games G_2 and G_3 are computationally indistinguishable. The proof relies on the hardness of HNF-AextLWE. In particular, any adversary \mathcal{B} has advantage at most $\text{Adv}(\mathcal{B}) \leq n \cdot \text{Adv}^{\text{HNF-AextLWE}}$ at distinguishing games G_2 and G_3 .

Game G_4 : Here, instead of having the challenger sample $\mathbf{s}, \mathbf{e} \leftarrow \mathcal{D}_{\mathbb{Z}^n, \sigma}$ at the start of the game, and $\mathbf{r}^*, \boldsymbol{\eta}^* \leftarrow \mathcal{D}_{\mathbb{Z}^n, \sigma}$ at the end and setting $sk^* = \mathbf{s} + \mathbf{r}^* + \Delta_{last}^{\mathbf{r}}$ and $pk^* = (\mathbf{A}, \mathbf{A}(\mathbf{s} + \Delta_{last}^{\mathbf{r}} + \mathbf{r}^*) + \mathbf{e} + \Delta_{last}^{\boldsymbol{\eta}} + \boldsymbol{\eta}^*)$, we do the following.

Let us define distributions $\mathcal{S}, \mathcal{S}_{\mathbf{t}}$ and $\mathcal{S}_{\bar{\mathbf{e}}}$ as:

$$\mathcal{S} = \mathcal{D}_{\mathbb{Z}^n, \sigma\sqrt{2}}, \quad \mathcal{S}_{\mathbf{t}} = \mathcal{D}_{\mathbb{Z}^n, \frac{\sigma}{\sqrt{2}}, \frac{\mathbf{t}}{2}}, \quad \text{and} \quad \mathcal{S}_{\bar{\mathbf{e}}} = \mathcal{D}_{\mathbb{Z}^n, \frac{\sigma}{\sqrt{2}}, \frac{\bar{\mathbf{e}}}{2}}.$$

Then, in game G_4 , the challenger samples $\mathbf{t}, \tilde{\mathbf{e}} \leftarrow \mathcal{S}$ at the beginning of the game, then samples $\mathbf{s} \leftarrow \mathcal{S}_{\mathbf{t}}, \mathbf{e} \leftarrow \mathcal{S}_{\tilde{\mathbf{e}}}$ and finally sets $sk^* = \mathbf{t} + \Delta_{last}^{\mathbf{r}}$ and $pk^* = (\mathbf{A}, \mathbf{A}\mathbf{t} + \tilde{\mathbf{e}} + \mathbf{A}\Delta_{last}^{\mathbf{r}} + \Delta_{last}^{\eta})$.

Let $\delta = 2\varepsilon/(1 - \varepsilon)$. Lemma 2 shows that this change only induces a statistically negligible bias. Specifically, assuming $\sigma \geq \sqrt{2 \ln(2n(1 + 1/\varepsilon))}/\pi$, \mathbf{t} is within statistical distance at most δ from the distribution of $\mathbf{s} + \mathbf{r}^*$ in game G_3 , and the marginal distribution of \mathbf{s} in game G_4 with respect to the adversary's view is:

$$\begin{aligned} \mathbb{P}[\mathbf{s} = \mathbf{x}] &= \sum_{\mathbf{y} \in \mathbb{Z}^n} \mathbb{P}[\mathbf{s} = \mathbf{x} | \mathbf{t} = \mathbf{y}] \mathbb{P}[\mathbf{t} = \mathbf{y}] \\ &= \sum_{\mathbf{y} \in \mathbb{Z}^n} \mathcal{D}_{\mathbb{Z}^n, \frac{\sigma}{\sqrt{2}}} \left(\mathbf{x} - \frac{\mathbf{y}}{2} \right) \mathcal{D}_{\mathbb{Z}^n, \sigma\sqrt{2}}(\mathbf{y}) \\ &= \sum_{\mathbf{y} \in \mathbb{Z}^n} \mathcal{D}_{\mathbb{Z}^n, \sigma\sqrt{2}}(2\mathbf{x} - \mathbf{y}) \mathcal{D}_{\mathbb{Z}^n, \sigma\sqrt{2}}(\mathbf{y}) \\ &\approx_{\delta} \mathcal{D}_{\mathbb{Z}^n, 2\sigma}(2\mathbf{x}) = \mathcal{D}_{\mathbb{Z}^n, \sigma}(\mathbf{x}). \end{aligned}$$

The fourth equality comes from applying Lemma 2 for the convolution of two Gaussian distributions with the same standard deviation. The same argument applies for $\tilde{\mathbf{e}}$ and \mathbf{e} . Hence any adversary \mathcal{B} has advantage at most $4\delta = 8\varepsilon/(1 - \varepsilon)$ in distinguishing games G_3 and G_4 .

Game G_5 : In this game, we replace \mathbf{b}_0 and $up^* = (\mathbf{T}_{last}, \mathbf{pad}_{last})$ by uniform elements. Note that \mathbf{T}_{last} is already uniform since game G_3 . Hence, the challenger samples $\mathbf{b}_0, \mathbf{pad}_{last} \leftarrow \mathcal{U}(\mathbb{Z}_q^n)$, and sets $pk_0 = (\mathbf{A}, \mathbf{b}_0)$ at the start of the game, and returns $up^* = (\mathbf{T}_{last}, \mathbf{pad}_{last})$ as the last update message.

Lemma 7 below states that this game and the previous one are computationally indistinguishable under the LWE assumption.

Game G_6 : This is the final game. Here, the challenger replaces the challenge \mathbf{c}^* to make it uniform: it samples $\mathbf{T}_{chall} \leftarrow \mathcal{U}(\mathbb{Z}_q^{n \times n})$ and $\mathbf{pad}_{chall} \leftarrow \mathcal{U}(\mathbb{Z}_q^n)$, and then sets $\mathbf{c}^* = (\mathbf{T}_{chall}, \mathbf{pad}_{chall})$.

Remember that in game G_5 , we have $\mathbf{c}^* = (\mathbf{X}_{chall}\mathbf{A} + \mathbf{E}_{chall}, \mathbf{X}_{chall}\mathbf{b}_{chall} + \mathbf{f}_{chall} + \lfloor q/p \rfloor \cdot \boldsymbol{\mu}_{\beta})$. We can rewrite \mathbf{c}^* in a matrix form as:

$$\mathbf{X}_{chall} (\mathbf{A} \parallel \mathbf{b}_{chall}) + (\mathbf{E}_{chall} \parallel \mathbf{f}_{chall}) + \lfloor q/p \rfloor \cdot (\mathbf{0} \parallel \boldsymbol{\mu}_{\beta}) \quad (3)$$

with $\mathbf{A} \leftarrow \mathcal{U}(\mathbb{Z}_q^{n \times n})$ and $\mathbf{b}_{chall} = \mathbf{b}_0 + \mathbf{A}\Delta_{chall}^{\mathbf{r}} + \Delta_{chall}^{\eta}$. Recall that we have $\mathbf{b}_0 \leftarrow \mathcal{U}(\mathbb{Z}_q^n)$ since game G_5 . The last column of Equation (3) is

$$(\mathbf{X}_{chall}\mathbf{b}_0 + \mathbf{f}_{chall}) + (\mathbf{X}_{chall}(\mathbf{A}\Delta_{chall}^{\mathbf{r}} + \Delta_{chall}^{\eta})) + \lfloor q/p \rfloor \cdot \boldsymbol{\mu}_{\beta}.$$

and can be rewritten as

$$\begin{aligned} &(\mathbf{X}_{chall}\mathbf{b}_0 + \mathbf{f}_{chall,0}) \\ &+ \mathbf{T}_{chall}\Delta_{chall}^{\mathbf{r}} - \mathbf{E}_{chall}\Delta_{chall}^{\mathbf{r}} + \mathbf{X}_{chall}\Delta_{chall}^{\eta} + \mathbf{f}_{chall,1} \\ &+ \lfloor q/p \rfloor \cdot \boldsymbol{\mu}_{\beta}, \end{aligned}$$

where $\mathbf{f}_{chall} = \mathbf{f}_{chall,0} + \mathbf{f}_{chall,1}$ for $\mathbf{f}_{chall,0}, \mathbf{f}_{chall,1} \leftarrow \mathcal{D}_{\mathbb{Z}^n, \sigma_c/\sqrt{2}}$. Consider we are working with standard deviation $\sigma_c/\sqrt{2}$ instead of σ_c . The first term is a multiple-secret LWE sample that is independent of any adversarially chosen value. The second one can be computed from \mathbf{T}_{chall} . The next three can be viewed as an HNF-AextLWE hint on the secret \mathbf{X}_{chall} and the error \mathbf{E}_{chall} with hint vector $\mathbf{z}_0 = \Delta_{chall}^\eta$ and $\mathbf{z}_1 = \Delta_{chall}^r$, which are small vectors. The difference in standard deviation can be handled as in Lemma 7 by adding terms sampled from $\mathcal{D}_{\mathbb{Z}^n, \sigma_c/\sqrt{2}}$ to the matrices \mathbf{X}_{chall} and \mathbf{E}_{chall} . Applying Lemma 2 proves this change to be statistically unnoticeable.

The above indicates that the modification between this game and game G_5 can be analyzed by using the multiple-secret variant of HNF-AextLWE $_{q,n,n+1,\sigma_c/\sqrt{2},ky\sigma}$ with n secrets and hint vector $\mathbf{z} = [(\Delta_{chall}^\eta)^T \| (\Delta_{chall}^r)^T]^T$. Consequently, any adversary \mathcal{A} has advantage at most $n \cdot \text{Adv}^{\text{HNF-AextLWE}} + (2n+1) \cdot 2\varepsilon/(1-\varepsilon)$ in distinguishing between games G_5 and G_6 .

Note that in game G_6 , the adversary has no information on the challenge μ_β . Hence $\text{Adv}^{G_6}(\mathcal{A}) = 0$. We obtain

$$\text{Adv}_{\text{UPKE}}^{\text{IND-CR-CPA}}(\mathcal{A}) \leq (2n+8) \cdot \frac{2\varepsilon}{1-\varepsilon} + (2n+1) \cdot \text{Adv}^{\text{HNF-AextLWE}}.$$

This completes the proof, up to Lemmas 6 and 7 below. \square

Lemma 6. *For any adversary \mathcal{A} that distinguishes between games G_2 and G_3 , there exists an efficient algorithm \mathcal{B} for HNF-AextLWE $_{q,n,n,\sigma_c,B}$ (for $B = (k+1)y\sigma$), calling \mathcal{A} once, such that $\text{Adv}_{G_2,G_3}^{\text{dist}}(\mathcal{A}) \leq n \cdot \text{Adv}^{\text{HNF-AextLWE}}(\mathcal{B})$.*

Proof. This proof constructs an algorithm \mathcal{B} for the multiple-secret variant of the HNF-AextLWE assumption with n secrets, using a distinguisher \mathcal{A} for games G_2 and G_3 .

Algorithm \mathcal{B} receives a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times n}$ from the HNF-AextLWE challenger. Then it samples $\mathbf{s}, \mathbf{e} \leftarrow \mathcal{D}_{\mathbb{Z}^n, \sigma}$ and sets $pk_0 = (\mathbf{A}, \mathbf{b}_0 = \mathbf{A}\mathbf{s} + \mathbf{e})$, forwards pk_0 to \mathcal{A} and acts as \mathcal{A} 's challenger until the last update phase where it has to send up^* and sk^* to \mathcal{A} . At this stage, algorithm \mathcal{B} knows the sum of all the updates Δ_{last}^r and the sum of all the noises used for each updates Δ_{last}^η as \mathcal{A} has finished querying the \mathcal{O}_{up} oracle.

The HNF-AextLWE challenger expects small vectors $\mathbf{z}_0, \mathbf{z}_1$ for which to send a hint \mathbf{h} . Let $\mathbf{X}_{last} \leftarrow \mathcal{D}_{\mathbb{Z}^n \times n, \sigma_c}$ be the secret matrix and $\mathbf{E}_{last} \leftarrow \mathcal{D}_{\mathbb{Z}^n \times n, \sigma_c}$ be the error matrix sampled by the challenger in the multiple-secret variant of HNF-AextLWE. Algorithm \mathcal{B} sets $\mathbf{z}_0 = \mathbf{e} + \Delta_{last}^\eta$ and $\mathbf{z}_1 = -(\mathbf{s} + \Delta_{last}^r)$.

It then receives from the challenger a matrix $\mathbf{B} \in \mathbb{Z}_q^{n \times n}$ and a hint

$$\mathbf{h} = \mathbf{X}_{last}\mathbf{z}_0 + \mathbf{E}_{last}\mathbf{z}_1 = (\mathbf{X}_{last} \| \mathbf{E}_{last})\mathbf{z} + \mathbf{f}_{last},$$

where $\mathbf{z} = (\mathbf{z}_0^T \| \mathbf{z}_1^T)^T$ and $\mathbf{f}_{last} \leftarrow \mathcal{D}_{\mathbb{Z}^n, \sigma_c}$. The matrix \mathbf{B} is either uniform or of the form $\mathbf{X}_{last}\mathbf{A} + \mathbf{E}_{last}$.

Adversary \mathcal{B} sets

$$\begin{aligned} up^* &= (\mathbf{T}_{last} = \mathbf{B} - \lfloor q/p \rfloor \cdot \mathbf{Id}, \mathbf{T}_{last}(\mathbf{s} + \Delta_{last}^{\mathbf{r}}) + \mathbf{h} + \lfloor q/p \rfloor \Delta_{last}^{\mathbf{r}}) \\ &= \left(\mathbf{T}_{last}, \mathbf{T}_{last}(\mathbf{s} + \Delta_{last}^{\mathbf{r}}) + \mathbf{X}_{last}(\mathbf{e} + \Delta_{last}^{\boldsymbol{\eta}}) - \mathbf{E}_{last}(\mathbf{s} + \Delta_{last}^{\mathbf{r}}) \right. \\ &\quad \left. + \mathbf{f}_{last} + \lfloor q/p \rfloor \Delta_{last}^{\mathbf{r}} \right). \end{aligned}$$

It also sets $pk^* = (\mathbf{A}, \mathbf{b}_0 + \mathbf{A}(\Delta_{last}^{\mathbf{r}} + \mathbf{r}^*) + \Delta_{last}^{\boldsymbol{\eta}} + \boldsymbol{\eta}^*)$ and $sk^* = \mathbf{s} + \Delta_{last}^{\mathbf{r}} + \mathbf{r}^*$, where $\mathbf{r}^*, \boldsymbol{\eta}^* \leftarrow \mathcal{D}_{\mathbb{Z}^n, \sigma}$.

The case where \mathbf{B} is uniform corresponds to adversary \mathcal{A} playing game G_3 and the case where $\mathbf{B} = \mathbf{X}_{last}\mathbf{A} + \mathbf{E}_{last}$ corresponds to \mathcal{A} playing game G_2 . Hence \mathcal{B} has the same advantage as \mathcal{A} .

By a hybrid argument, there exists an adversary \mathcal{B}' for $\text{HNF-AextLWE}_{q,n,n,\sigma,B}$ such that the advantage of \mathcal{B} in the multiple-secret variant of HNF-AextLWE with n secrets can be bounded by $n \cdot \text{Adv}^{\text{HNF-AextLWE}}(\mathcal{B}')$, completing the proof. \square

Lemma 7. *For any adversary \mathcal{A} that distinguishes between games G_4 and G_5 , there exists an adversary \mathcal{B} for $\text{LWE}_{q,n,2n,\sigma/2}$ calling \mathcal{A} once, such that:*

$$\text{Adv}_{G_4, G_5}^{\text{dist}}(\mathcal{A}) \leq \text{Adv}^{\text{LWE}}(\mathcal{B}) + \frac{6\varepsilon}{1 - \varepsilon}.$$

Proof. Let us build an adversary \mathcal{B} for $\text{LWE}_{q,n,2n,\sigma/2}$ that uses any distinguisher \mathcal{A} between games G_4 and G_5 .

Adversary \mathcal{B} receives a uniform $\mathbf{B} \in \mathbb{Z}_q^{2n \times n}$ and a vector $\mathbf{c} \in \mathbb{Z}_q^{2n}$ from the LWE challenger. The vector \mathbf{c} is either uniform or computed as an LWE sample with secret $\mathbf{s} \leftarrow \mathcal{D}_{\mathbb{Z}^n, \sigma/2}$. Now adversary \mathcal{B} samples $\mathbf{E}_{last}, \mathbf{X}_{last} \leftarrow \mathcal{D}_{\mathbb{Z}^n \times n, \sigma_e}$. It then computes

$$\mathbf{B}' = \mathbf{M}\mathbf{B} + \begin{pmatrix} \mathbf{0} \\ \mathbf{E}_{last} \end{pmatrix}, \quad \text{with } \mathbf{M} = \begin{pmatrix} \mathbf{Id} & \mathbf{0} \\ \mathbf{X}_{last} & \mathbf{Id} \end{pmatrix} \in \mathbb{Z}_q^{2n \times 2n}$$

and parses \mathbf{B}' as $(\mathbf{A}^T \parallel \mathbf{T}_{last}^T)^T$. Let $\mathbf{t}, \tilde{\mathbf{e}} \leftarrow \mathcal{S} = \mathcal{D}_{\mathbb{Z}^n, \sigma\sqrt{2}}$. After that, it samples elements $\mathbf{s}' \leftarrow \mathcal{D}_{\mathbb{Z}^n, \sigma/2, \mathbf{t}/2}$, $\boldsymbol{\eta} \leftarrow \mathcal{D}_{\mathbb{Z}^n, \sigma/2, \tilde{\mathbf{e}}/2}$ and $\mathbf{f}' \leftarrow \mathcal{D}_{\mathbb{Z}^n, (\sigma_e^2 - \sigma^2/4)} \mathbf{Id}$ that are used to adjust the standard deviations of the discrete Gaussian distributions involved in the proof. Then it sets $\mathbf{e}' = (\boldsymbol{\eta}^T \parallel \mathbf{f}'^T)^T$ and $\mathbf{c}' = \mathbf{M}(\mathbf{c} + \mathbf{e}') + \mathbf{M}\mathbf{B}\mathbf{s}'$ and parses \mathbf{c}' as $(\mathbf{b}_0^T \parallel \mathbf{u}_1^T)^T$.

From there, adversary \mathcal{B} runs as \mathcal{A} 's challenger and sets $pk_0 = (\mathbf{A}, \mathbf{b}_0)$. At epoch $last$, it computes

$$up^* = (\mathbf{T}_{last}, \mathbf{u}_1 + (\mathbf{T}_{last} - \mathbf{E}_{last} + \lfloor q/p \rfloor \cdot \mathbf{Id})\Delta_{last}^{\mathbf{r}}) + \mathbf{X}_{last}\Delta_{last}^{\boldsymbol{\eta}}.$$

If \mathcal{A} returns G_4 then \mathcal{B} guesses that \mathbf{c} is an LWE sample and if \mathcal{A} returns G_5 it guesses that it is uniform.

If \mathbf{c} is uniform, as \mathbf{M} is invertible, \mathbf{B}' and \mathbf{c}' are also uniformly distributed and adversary \mathcal{A} is playing game G_5 .

If $\mathbf{c} = \mathbf{B}\mathbf{s} + (\mathbf{e}^T \|\mathbf{f}^T)^T$, for $\mathbf{s} \leftarrow \mathcal{D}_{\mathbb{Z}^n, \sigma/2}$ and $\mathbf{e}, \mathbf{f} \leftarrow \mathcal{D}_{\mathbb{Z}^n, \sigma/2}$, then

$$\begin{aligned} \mathbf{c}' &= \mathbf{M} \left(\mathbf{B}\mathbf{s} + \begin{pmatrix} \mathbf{e} + \boldsymbol{\eta} \\ \mathbf{f} + \mathbf{f}' \end{pmatrix} \right) + \mathbf{M}\mathbf{B}\mathbf{s}' \\ &= \begin{pmatrix} \mathbf{A} \\ \mathbf{T}_{last} - \mathbf{E}_{last} \end{pmatrix} (\mathbf{s} + \mathbf{s}') + \begin{pmatrix} \mathbf{e} + \boldsymbol{\eta} \\ \mathbf{X}_{last}(\mathbf{e} + \boldsymbol{\eta}) + \mathbf{f} + \mathbf{f}' \end{pmatrix} = \begin{pmatrix} \mathbf{b}_0 \\ \mathbf{u}_1 \end{pmatrix}. \end{aligned}$$

Let us set $\bar{\mathbf{s}} = \mathbf{s} + \mathbf{s}'$, $\bar{\mathbf{e}} = \mathbf{e} + \boldsymbol{\eta}$, $\bar{\mathbf{f}} = \mathbf{f} + \mathbf{f}'$. Then, using the equation above, we have the following:

$$\begin{aligned} up^* &= (\mathbf{T}_{last}, \mathbf{u}_1 + (\mathbf{T}_{last} - \mathbf{E}_{last} + \lfloor q/p \rfloor \cdot \mathbf{Id}) \Delta_{last}^r + \mathbf{X}_{last} \Delta_{last}^\eta) \\ &= (\mathbf{T}_{last}, (\mathbf{T}_{last} - \mathbf{E}_{last}) \bar{\mathbf{s}} + \mathbf{X}_{last} (\bar{\mathbf{e}} + \Delta_{last}^\eta) + \bar{\mathbf{f}} \\ &\quad + (\mathbf{T}_{last} - \mathbf{E}_{last} + \lfloor q/p \rfloor \cdot \mathbf{Id}) \Delta_{last}^r) \\ &= (\mathbf{T}_{last}, (\mathbf{T}_{last} - \mathbf{E}_{last}) (\bar{\mathbf{s}} + \Delta_{last}^r) + \mathbf{X}_{last} (\bar{\mathbf{e}} + \Delta_{last}^\eta) + \bar{\mathbf{f}} + \lfloor q/p \rfloor \cdot \Delta_{last}^r). \end{aligned} \tag{4}$$

Let $\delta = 2\varepsilon/(1 - \varepsilon)$, for $\varepsilon \in (0, 1)$. As $\mathbf{s} \leftarrow \mathcal{D}_{\mathbb{Z}^n, \sigma/2}$ and $\mathbf{s}' \leftarrow \mathcal{D}_{\mathbb{Z}^n, \sigma/2, \mathbf{t}/2}$, Lemma 2 gives that the distribution of $\bar{\mathbf{s}}$ has statistical distance at most δ from $\mathcal{D}_{\mathbb{Z}^n, \sigma/\sqrt{2}, \mathbf{t}/2}$. Similarly, errors $\boldsymbol{\eta}$ and \mathbf{f}' were chosen such that $\bar{\mathbf{e}}$ and $\bar{\mathbf{f}}$ are within statistical distance at most δ from $\mathcal{D}_{\mathbb{Z}^n, \sigma/\sqrt{2}, \bar{\mathbf{e}}}$ and $\mathcal{D}_{\mathbb{Z}^n, \sigma_c}$. The equation above shows that up^* is statistically close (at distance at most 3δ) from its value in game G_4 , thus \mathcal{A} can be viewed as playing game G_4 .

Overall, algorithm \mathcal{B} has advantage at least $\text{Adv}_{G_4, G_5}^{\text{dist}}(\mathcal{A}) - 3\delta$, completing the proof. \square

5 A UPKE Fujisaki-Okamoto Transform

In this section, we describe a transform from an IND-CR-CPA UPKE into an IND-CR-CCA UKEM following the Fujisaki-Okamoto [24] technique.

Definition 14 (FO-transform for UPKEs). *Let UPKE be a UPKE, and G and H be two functions modeled as random oracles. We define the transform $\text{FO}(\text{UPKE}, G, H)$ in Figure 7.*

<p>KeyGen = UPKE.KeyGen.</p> <p>Encaps(pk):</p> <p style="padding-left: 20px;">$m \leftarrow \mathcal{U}(\mathcal{M});$</p> <p style="padding-left: 20px;">$c \leftarrow \text{UPKE.Enc}(pk, m; G(pk, m));$</p> <p style="padding-left: 20px;">$K = H(m, c);$</p> <p style="padding-left: 20px;">return (c, K).</p>	<p>Decaps(sk, c):</p> <p style="padding-left: 20px;">$m' \leftarrow \text{UPKE.Dec}(sk, c);$</p> <p style="padding-left: 20px;">if $c \neq \text{UPKE.Enc}(pk, m'; G(pk, m'))$</p> <p style="padding-left: 40px;">then return \perp;</p> <p style="padding-left: 20px;">return $K' = H(m', c)$.</p> <p>UpdatePk = UPKE.UpdatePk.</p> <p>UpdateSk = UPKE.UpdateSk.</p>
--	--

Fig. 7: Transform $\text{FO}(\text{UPKE}, G, H)$ for a UPKE using random oracles G, H .

Our FO transform is essentially the KEM⁺ construction from [27]. We add pk to the inputs of the hash function used to determinize the Enc algorithm in order to prevent trivial attacks, given the ability of the adversary to update the key pair.

Theorem 3 (FO transform for UPKEs). *Let $\gamma, \delta \in (0, 1), k > 0$. Let $\text{UPKE} = (\text{Enc}, \text{Dec}, \text{UpdatePk}, \text{UpdateSk})$ denote a γ -spread and (k, δ) -correct k -IND-CR-CPA UPKE scheme. Then the UPKE FO(UPKE, G, H) is a (k, δ) -correct k -IND-CR-CCA UKEM in the ROM.*

More precisely, for any adversary \mathcal{A} for the k -IND-CR-CCA UKEM game in the ROM making at most q_G queries to oracle G, q_H queries to oracle H and q_D queries to oracle \mathcal{O}_{dec} , there exists an adversary \mathcal{B} for the k -IND-CR-CPA game of UPKE with a similar running time such that:

$$\text{Adv}^{\text{IND-CR-CCA}}(\mathcal{A}) \leq q_G \cdot \delta + q_D \cdot \gamma + 2 \left(\text{Adv}^{\text{IND-CR-CPA}}(\mathcal{B}) + \frac{q_G + q_H}{|\mathcal{M}|} \right).$$

The proof of the above theorem follows standard techniques for FO analysis (e.g., [27]), and we postpone it to the appendix (Appendix C).

Note that we rely on the γ -spreadness of the underlying UPKE scheme. We prove this property for the scheme from Section 4 in the appendix (Appendix B.3).

6 Obtaining IND-CU-CCA Security

In this section, we further boost security in order to achieve IND-CU-CCA-security. As in [1], we use a NIZK argument that two keys encrypt the same message in order to make a reduction from IND-CU-CCA to IND-CR-CCA. This technique allows to extract the randomness used by the adversary for the oracle queries to $\mathcal{O}_{\text{up}}(\cdot)$, to forward it to the update oracle of the IND-CR-CCA challenger. We give the definitions about Non Interactive Zero Knowledge (NIZK) argument in the ROM in the appendix (Appendix A).

Let $\text{UPKE} = (\text{KeyGen}, \text{Enc}, \text{Dec}, \text{UpdatePk}, \text{UpdateSk})$ be a k -IND-CR-CPA UPKE, for some $k > 0$. Define UKEM = (KeyGen, Encaps, Decaps, UpdatePk, UpdateSk) as the k -IND-CR-CCA UKEM scheme obtained by applying our FO transform from Section 5 to UPKE, using G, H modeled as random oracles. Let F be a third function, also modeled as a random oracle. We assume that UpdatePk proceeds in two parts (this is the case for all known constructions, including the one from Section 4): $\text{UpdatePk}(pk) = (\text{Enc}(pk, r), \text{NewPk}(pk, r))$, i.e., a first part which encrypts the randomness of the update using the UKEM encryption algorithm, and a second one which returns the updated public key. Let us define the language

$$\begin{aligned} \mathcal{L}_{\text{up}}^{\text{UKEM}} &= \{(pk_0, pk_1, pk', ct_0, ct_1) \mid \exists r_0, r_1, r, \\ &ct_0 = \text{Enc}(pk_0, r; r_0) \wedge ct_1 = \text{Enc}(pk_1, r; r_1) \wedge (pk', ct_0) = \text{UpdatePk}(pk_0; r)\}. \end{aligned}$$

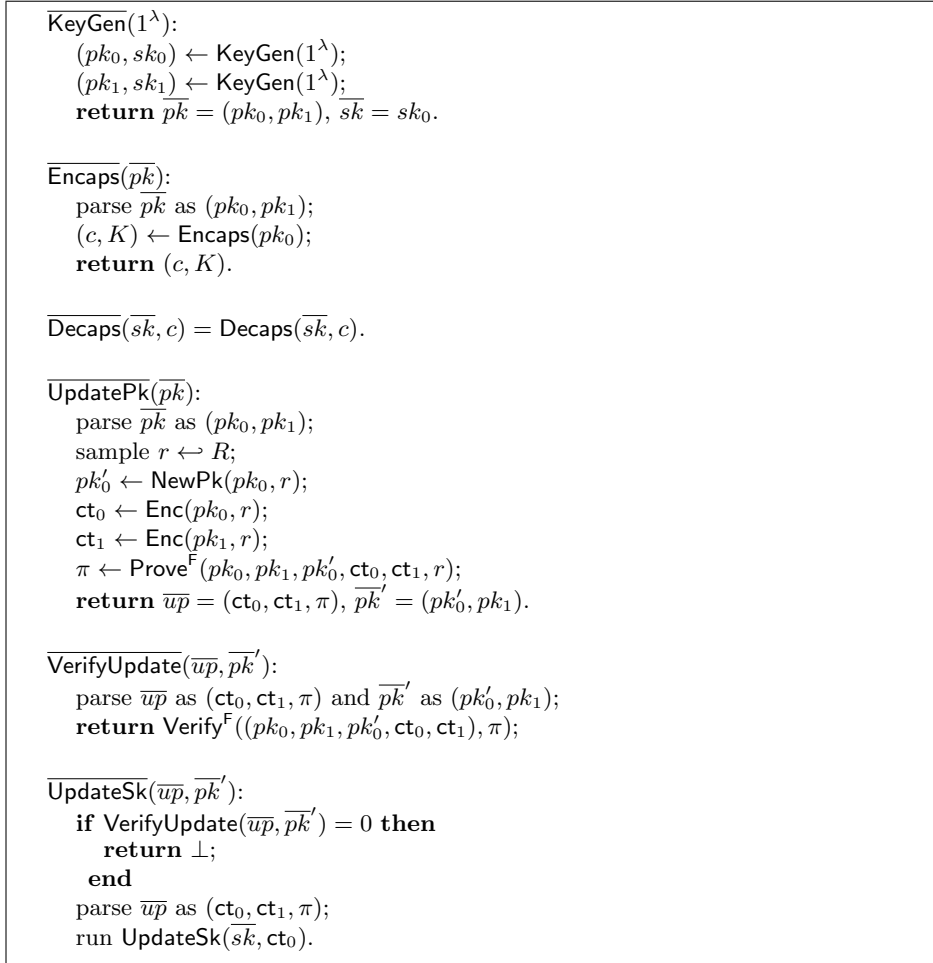


Fig. 8: Construction of a IND-CU-CCA UKEM.

Let $\Pi = (\text{Prove}^F, \text{Verify}^F)$ a NIZK argument in the random oracle for $\mathcal{L}_{up}^{\text{UKEM}}$. We construct an k -IND-CU-CCA UKEM as described in Figure 8.

Theorem 4. *Let UPKE, UKEM, Π be defined as above. Then, the construction $\overline{\text{UKEM}}$ described in Figure 8 is an k -IND-CU-CCA UKEM. Specifically, for any adversary \mathcal{A} against the k -IND-CU-CCA security of $\overline{\text{UKEM}}$, there exist adversaries $\mathcal{B}, \mathcal{C}, \mathcal{D}, \mathcal{E}$ with running times similar to \mathcal{A} 's such that:*

$$\text{Adv}^{\text{IND-CU-CCA}}(\mathcal{A}) \leq \text{Adv}_{\text{UKEM}}^{\text{IND-CR-CCA}}(\mathcal{B}) + \text{Adv}_{\text{UPKE}}^{\text{IND-CR-CPA}}(\mathcal{C}) + \text{Adv}_{\Pi}^{\text{zk}}(\mathcal{D}) + \text{Adv}_{\Pi}^{\text{sound}}(\mathcal{E}).$$

The proof closely follows the one of IND-CU-CCA security of the construction from [1] and is detailed in the appendix (Appendix D).

7 Concrete Parameters

In this section, we give some concrete parameters for the scheme presented in Section 4, which can directly be transformed into an IND-CR-CCA UKEM by applying the FO transform from Section 5. We focus on the latter. We conjecture that security holds in the module setting and use the lattice-estimator SAGE module (commit fd4a460) from [3] to estimate the security of the given parameter sets. For our UPKE/UKEM, we consider the module variant of the scheme presented in Section 4, i.e., we define $\mathcal{R} = \mathbb{Z}[X]/(X^d + 1)$ and $\mathcal{R}_q = \mathcal{R}/q\mathcal{R}$ and we consider the base ring to be \mathcal{R} instead of \mathbb{Z} .

Note that, for $p > 0$ a prime, the message space of Enc for the module variant is $\mathcal{M} = \mathcal{R}_p^n$ which is of size p^{dn} . For optimization purposes, we drop the last $n - 1$ rows of the whole ciphertext computed by Enc in our encapsulation mechanism, so that an encapsulation is just:

$$c = (\mathbf{x}^T \mathbf{A} + \mathbf{e}^T, \mathbf{x}\mathbf{b} + f + \lfloor q/p \rfloor m)$$

for $\mathbf{x}, \mathbf{e} \in \mathcal{R}_q^n$, $f \in \mathcal{R}_q$ and $m \in \mathcal{R}_p$. The message space is now $\mathcal{M} = \mathcal{R}_p$, of size p^d . This optimization is made possible by considering the UKEM, which only require a message space with at least λ bits of entropy, which is the case when setting $d = 256$. The whole message space \mathcal{R}_p^n is only used to encrypt updates, as an update changes all components of the secret key.

Also, as done in [11], we replace Gaussian distributions by the centered binomial distributions B_η , which for $\eta > 0$, samples elements $(a_i, b_i)_{i \leq \eta} \leftarrow \mathcal{U}(\{0, 1\}^2)$ and returns $\sum_{i=1}^{\eta} (a_i - b_i)$. Samples from B_η are contained in $[-\eta, \eta]$, and we choose the modulus q such that perfect correctness ($\delta = 0$) is guaranteed up to a bounded number of (possibly malicious) updates. We let k denote this bound, and provide parameters for $k \in \{2^5, 2^{10}, 2^{15}, 2^{20}\}$. We are assuming worst-case updates and then make q scale linearly with k . It could be tempting to make it scale with \sqrt{k} as updates are symmetric and centered in 0 though we should not, as they are chosen by the attacker. Due to this requirement, our UPKE/UKEM suffers from a loss compared to Kyber, which can take q as small as 3329 and then have ciphertexts of size 0.8KB.

As we are working in the UPKE setting, we consider that the adversary gets a leakage $\mathbf{s} + \mathbf{r}$ on the initial secret key \mathbf{s} , which roughly halves the variance of the distribution of \mathbf{s} in the adversary's view (as shown in the proof of Theorem 2). We use a script to compute the average variance left on \mathbf{s} conditioned on the value of $\mathbf{s} + \mathbf{r}$. We obtain that for $\mathbf{s} \leftarrow B_{2\eta}^n$, we are left on average as if \mathbf{s} was sampled from B_η^n . This is taken into account for the security estimates.

Our parameters are given in Table 2. Note that as done in Kyber, in order to have fast multiplication using the Number Theoretic Transform in the ring, we take modulus $q = 1 \pmod{2d}$. This is the first practical lattice-based construction of UPKE/UKEM, hence there are no equivalent constructions to compare our results to. We achieve similar efficiency as the IND-CR-CPA construction of [1], which is based on the DCR assumption achieves a ciphertext and update size of 1.5KB (for the CPA case only, although our FO transform applies to their

scheme). Note that by increasing d , the matrices involved become smaller. Hence, a tradeoff can be made to reduce the sizes of the updates at the cost of increasing ciphertext size. For small number of updates, we also apply the bit-dropping technique from Kyber to improve parameters. This optimization drops parts of the least significant bits of the ciphertexts to reduce their size. We use the script provided at <https://github.com/pq-crystals/security-estimates> to estimate the correctness loss implied by using this technique.

	λ	q	n	d	p	η	δ	k	$ ct $	$ up $
DCR-based construction [1]	128	-	-	-	-	-	0	∞	8.3KB	1.5KB
Estimate for [22]	120	$\approx 2^{85}$	11	256	21	10	0	2^5	33KB	360KB
This work	128	$\approx 2^{21}$	3	256	5	2	2^{-136}	2^5	1.8KB	5.4KB
	128	$\approx 2^{26}$	4	256	5	2	0	2^{10}	3.0KB	12KB
	116	$\approx 2^{31}$	2	512	5	2	0	2^{15}	5.8KB	12KB
	128	$\approx 2^{36}$	3	512	5	2	0	2^{20}	9.1KB	27KB

Table 2. Parameter sets for the module variant of our IND-CR-CCA UKEM.

IND-CU-CCA instantiation. In order to add security against chosen updates via our transform from Section 6, we can further add a computationally sound NIZK argument for $\mathcal{L}_{up}^{\text{UKEM}}$ in the updates. In the module setting, the language $\mathcal{L}_{up}^{\text{UKEM}}$ can be defined as:

$$\begin{aligned} \mathcal{L}_{up}^{\text{UKEM}} = \{ & (pk_0, pk_1, pk', ct_0, ct_1) \mid \exists \mathbf{X}_0, \mathbf{X}_1, \mathbf{E}_0, \mathbf{E}_1 \in \mathcal{R}^{n \times n}, \mathbf{f}_0, \mathbf{f}_1, \mathbf{r} \in \mathcal{R}^n \\ & ct_0 = (\mathbf{X}_0 \mathbf{A} + \mathbf{E}_0, \mathbf{X}_0 \mathbf{b} + \mathbf{f}_0 + \lfloor q/p \rfloor \cdot \mathbf{r}) \bmod q \wedge \|\mathbf{X}_0\|_2, \|\mathbf{E}_0\|_2, \|\mathbf{f}_0\|_2 < B_0 \\ & \wedge ct_1 = (\mathbf{X}_1 \tilde{\mathbf{A}} + \mathbf{E}_1, \mathbf{X}_1 \tilde{\mathbf{b}} + \mathbf{f}_1 + \lfloor q/p \rfloor \cdot \mathbf{r}) \bmod q \wedge \|\mathbf{X}_1\|_2, \|\mathbf{E}_1\|_2, \|\mathbf{f}_1\|_2 < B_0 \\ & \wedge \|\mathbf{b}' - (\mathbf{b} + \mathbf{A}\mathbf{r})\|_2 \leq B_1 \wedge \|\mathbf{r}\|_2 < B_1 \}. \end{aligned}$$

where $pk_0 = (\mathbf{A}, \mathbf{b})$, $pk_1 = (\tilde{\mathbf{A}}, \tilde{\mathbf{b}})$, $pk' = (\mathbf{A}, \mathbf{b}')$ and B_0, B_1 are bounds for correctness.

Proving membership in $\mathcal{L}_{up}^{\text{UKEM}}$ then corresponds to proving 4 norm bounds for matrices, 4 norm bounds for vectors and $2n^2 + 2n$ linear equations over \mathcal{R}_q . This can be achieved by applying [33], which allows to prove exact norm bounds and linear relations using a commit-and-prove protocol. This only affects the size of the updates, since the ciphertext remains the same as in the IND-CR-CCA setting.

Acknowledgments. We thank Michael Rosenberg for pointing out a flaw in a prior version of this work. The first author is funded by the Direction Générale de l'Armement (Pôle de Recherche CYBER). This work was supported by the France 2030 ANR Project ANR-22-PECY-003 SecureCompute and the France 2030 ANR Project ANR-22-PETQ-0008 PQ-TLS.

References

1. Abou Haidar, C., Libert, B., Passelègue, A.: Updatable public key encryption from DCR: efficient constructions with stronger security. In: CCS (2022)
2. Agrawal, S., Gentry, C., Halevi, S., Sahai, A.: Sampling discrete gaussians efficiently and obliviously. In: ASIACRYPT (2013)
3. Albrecht, M.R., Player, R., Scott, S.: On the concrete hardness of learning with errors. ePrint 2015/046 (2015)
4. Alwen, J., Coretti, S., Dodis, Y., Tselekounis, Y.: Security analysis and improvements for the IETF MLS standard for group messaging. In: CRYPTO (2020)
5. Anderson, R.: Two remarks on public-key cryptology. In: CCS (1997), invited talk
6. Applebaum, B., Cash, D., Peikert, C., Sahai, A.: Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In: CRYPTO (2009)
7. Asano, K., Watanabe, Y.: Updatable public key encryption with strong CCA security: Security analysis and efficient generic construction. Cryptology ePrint Archive 2023/976 (2023)
8. Barić, N., Pfitzmann, B.: Collision-free accumulators and fail-stop signature schemes without trees. In: EUROCRYPT (1997)
9. Bellare, M., Miner, S.: A forward-secure digital signature scheme. In: CRYPTO (1999)
10. Boneh, D., Freeman, D.M.: Linearly homomorphic signatures over binary fields and new tools for lattice-based signatures. In: PKC (2011)
11. Bos, J.W., Ducas, L., Kiltz, E., Lepoint, T., Lyubashevsky, V., Schanck, J.M., Schwabe, P., Seiler, G., Stehlé, D.: CRYSTALS - Kyber: A CCA-secure module-lattice-based KEM. In: EuroS&P (2018)
12. Boudgoust, K., Jeudy, C., Roux-Langlois, A., Wen, W.: Entropic hardness of Module-LWE from Module-NTRU. In: Indocrypt (2022)
13. Boyen, X., Shacham, H., Shen, E., Waters, B.: Forward-secure signatures with untrusted update. In: CCS (2006)
14. Brakerski, Z., Döttling, N.: Hardness of LWE on general entropic distributions. In: EUROCRYPT (2020)
15. Brakerski, Z., Gentry, C., Vaikuntanathan, V.: (Leveled) fully homomorphic encryption without bootstrapping. In: ITCS (2012)
16. Brakerski, Z., Langlois, A., Peikert, C., Regev, O., Stehle, D.: Classical hardness of learning with errors. In: STOC (2013)
17. Camenisch, J., Shoup, V.: Practical verifiable encryption and decryption of discrete logarithms. In: CRYPTO (2003)
18. Canetti, R., Halevi, S., Katz, J.: A forward-secure public-key encryption scheme. In: EUROCRYPT (2003)
19. Cronin, E., Jamin, S., Malkin, T., McDaniel, P.: On the performance, feasibility, and use of forward-secure signatures. In: CCS (2003)
20. Dadush, D., Regev, O., Stephens-Davidowitz, N.: On the closest vector problem with a distance guarantee. In: CCC (2014)
21. Dodis, Y., Haralambiev, K., Lopez-Alt, A., Wichs, D.: Efficient public-key cryptography in the presence of key leakage. In: ASIACRYPT (2010)
22. Dodis, Y., Karthikeyan, H., Wichs, D.: Updatable public key encryption in the standard model. In: TCC (2021)
23. Duman, J., Hövelmanns, K., Kiltz, E., Lyubashevsky, V., Seiler, G.: Faster lattice-based KEMs via a generic Fujisaki-Okamoto transform using prefix hashing. In: CCS (2021)

24. Fujisaki, E., Okamoto, T.: Statistical zero knowledge protocols to prove modular polynomial relations. In: CRYPTO (1997)
25. Gentry, C., Peikert, C., Vaikunthanathan, V.: Trapdoors for hard lattices and new cryptographic constructions. In: STOC (2008)
26. Gentry, C., Silverberg, A.: Hierarchical ID-based cryptography. In: ASIACRYPT (2002)
27. Hofheinz, D., Hövelmanns, K., Kiltz, E.: A modular analysis of the Fujisaki-Okamoto transformation. In: TCC (2017)
28. Horwitz, J., Lynn, B.: Toward hierarchical identity-based encryption. In: EUROCRYPT (2002)
29. Hövelmanns, K., Hülsing, A., Majenz, C.: Failing gracefully: Decryption failures and the fujisaki-okamoto transform. In: ASIACRYPT (2022)
30. Itkis, G., Reyzin, L.: Forward-secure signatures with optimal signing and verifying. In: CRYPTO (2001)
31. Jost, D., Maurer, U., Mularczyk, M.: Efficient ratcheting: Almost-optimal guarantees for secure messaging. In: EUROCRYPT (2019)
32. Langlois, A., Stehlé, D.: Worst-case to average-case reductions for module lattices. *Des. Codes Cryptogr.* (2015)
33. Lyubashevsky, V., Nguyen, N.K., Plançon, M.: Lattice-based zero-knowledge proofs and applications: Shorter, simpler, and more general. In: CRYPTO (2022)
34. Lyubashevsky, V., Palacio, A., Segev, G.: Public-key cryptographic primitives provably as secure as subset sum. In: TCC (2010)
35. Malkin, T., Micciancio, D., Miner, S.: Efficient generic forward-secure signatures with an unbounded number of time periods. In: EUROCRYPT (2002)
36. Mera, J.M.B., Karmakar, A., Marc, T., Soleimani, A.: Efficient lattice-based inner-product functional encryption. In: PKC (2022)
37. Micciancio, D., Walter, M.: On the bit security of cryptographic primitives. In: EUROCRYPT (2018)
38. Naor, M., Yung, M.: Public-key cryptosystems provably secure against chosen ciphertext attacks. In: STOC (1990)
39. O’Neill, A., Peikert, C., Waters, B.: Bi-deniable public-key encryption. In: CRYPTO (2011)

Appendix

A Non-Interactive Zero-Knowledge Argument

In this section, we recall the definition of NIZK arguments. We focus on definitions in the ROM, without a trusted setup, as these are the ones we aim to use for our instantiations.

Definition 15 (NIZK in the ROM). Let $\mathcal{L} \subseteq \{0, 1\}^*$ an NP language, defined by an efficient relation \mathcal{R} , such that $x \in \mathcal{L} \Leftrightarrow \exists w \in \{0, 1\}^{q(|x|)} : R(x, w) = 1$, for some polynomial q . Let H denote a hash function, modeled as a random oracle. A non-interactive zero-knowledge argument in the ROM for \mathcal{L} is a pair of PPT algorithms $(\text{Prove}^H, \text{Verify}^H)$, such that:

- $\text{Prove}^H(x, w)$: on input a statement x and a witness w , returns a proof π ;
- $\text{Verify}^H(x, \pi)$: on input a statement x , and a proof π , returns 0 or 1.

We further require the following properties.

Completeness. For all x, w such that $R(x, w) = 1$, we have:

$$\text{Verify}^H(x, \text{Prove}^H(x, w)) = 1 .$$

Computational soundness. For all $x \notin \mathcal{L}$, for all PPT adversary \mathcal{A} :

$$\Pr[\text{Verify}^H(x, \pi^*) = 1 \mid \pi^* \leftarrow \mathcal{A}^H(x)] \leq \text{negl}(\lambda) ,$$

where the probability is over the choice of the random oracle. We let the above probability be denoted by $\text{Adv}_{\Pi}^{\text{sound}}(\mathcal{A})$.

Computational zero-knowledge. There exists a PPT simulator Sim which can program the random oracle H values such that, for all x, w such that $R(x, w) = 1$, for all PPT adversary \mathcal{A} , we have:

$$\text{Adv}_{\Pi}^{\text{zk}}(\mathcal{A}) := \left| \Pr[\mathcal{A}^H(x, \text{Prove}^H(x, w)) = 1] - \Pr[\mathcal{A}^H(x, \text{Sim}(x)) = 1] \right| \leq \text{negl}(\lambda) .$$

B Complements on the IND-CR-CPA UPKE Scheme

B.1 Proof of Lemma ??

Proof. By [14, Lemma 5.1], we have

$$H_{\infty}(\mathbf{s} \mid \mathbf{s} + \mathbf{e}) \geq H_{\infty}(\mathbf{s}) - \log \left(\sum_{\mathbf{y} \in \mathbb{Z}_q^n} \max_{\mathbf{x} \in \mathbb{Z}_q^n} \mathbb{P}[\mathbf{e} = \mathbf{y} - \mathbf{x}] \right) .$$

Hence it suffices to show $\sum_{\mathbf{y} \in \mathbb{Z}_q^n} \max_{\mathbf{x} \in \mathbb{Z}_q^n} \mathbb{P}[\mathbf{e} = \mathbf{y} - \mathbf{x}] \leq 2q^n/\sigma^n$. It holds that,

$$\begin{aligned} \sum_{\mathbf{y} \in \mathbb{Z}_q^n} \max_{\mathbf{x} \in \mathbb{Z}_q^n} \mathbb{P}[\mathbf{e} = \mathbf{y} - \mathbf{x}] &= \frac{1}{\rho_\sigma(\mathbb{Z}^n)} \sum_{\mathbf{y} \in \mathbb{Z}_q^n} \max_{\mathbf{x} \in \mathbb{Z}_q^n} \rho_\sigma(\mathbf{y} - \mathbf{x} + q\mathbb{Z}^n) \\ &\leq \frac{1}{\rho_\sigma(\mathbb{Z}^n)} \sum_{\mathbf{y} \in \mathbb{Z}_q^n} 2 \\ &= 2 \frac{q^n}{\rho_\sigma(\mathbb{Z}^n)} \\ &\leq 2 \frac{q^n}{\sigma^n}. \end{aligned}$$

The second inequality comes from applying [14, Lemma 2.4], which requires $q/\sigma \geq \sqrt{\ln(4n)}/\pi$. The last equality comes from the Poisson summation formula which gives that $\rho_\sigma(\mathbb{Z}^n) = \sigma^n \rho_{1/\sigma}(\mathbb{Z}^n) \geq \sigma^n$. \square

B.2 Correctness of the scheme

Proof. Let us prove the correctness of our UPKE. Let $(pk = (\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}), sk = \mathbf{s}) \leftarrow \text{KeyGen}(1^\lambda)$ be an honestly generated key pair. In order to consider the worst case scenario where k updates to the key have been performed, assume that \mathbf{s} and \mathbf{e} satisfy $\|\mathbf{s}\|_\infty, \|\mathbf{e}\|_\infty \leq ky\sigma$, for y a parameter that we set afterwards.

Let $\mu \in \mathbb{Z}_p^n$ and

$$\text{Enc}(pk, \mu) = (\mathbf{X}\mathbf{A} + \mathbf{E}, \mathbf{X}\mathbf{b} + \mathbf{f} + \lfloor q/p \rfloor \cdot \mu \bmod q).$$

Then, we have

$$\begin{aligned} \text{Dec}(\mathbf{s}, ct) &= \lfloor ct_1 - \mathbf{c}\mathbf{t}_0 \cdot \mathbf{s} \rfloor_p \\ &= \lfloor \mathbf{X}\mathbf{b} + \mathbf{f} + \lfloor q/p \rfloor \cdot \mu - (\mathbf{X}\mathbf{A} + \mathbf{E})\mathbf{s} \rfloor_p \\ &= \lfloor \mathbf{X}\mathbf{e} - \mathbf{E}\mathbf{s} + \mathbf{f} + \lfloor q/p \rfloor \cdot \mu \rfloor_p. \end{aligned}$$

We obtain that $\text{Dec}(\mathbf{s}, ct) = \mu$ if $\|\mathbf{X}\mathbf{e} - \mathbf{E}\mathbf{s} + \mathbf{f}\|_\infty < q/(2p)$. By the triangular inequality, it suffices to have $\|\mathbf{X}\mathbf{e}\|_\infty + \|\mathbf{E}\mathbf{s}\|_\infty + \|\mathbf{f}\|_\infty < q/(2p)$. By using that $\|\mathbf{M}\mathbf{v}\|_\infty \leq \|\mathbf{M}^T\|_1 \|\mathbf{v}\|_\infty \leq \sqrt{n} \|\mathbf{M}^T\|_2 \|\mathbf{v}\|_\infty$ for any matrix $\mathbf{M} \in \mathbb{Z}^{n \times n}$ and any vector $\mathbf{v} \in \mathbb{Z}^n$, we obtain another sufficient condition:

$$\sqrt{n} \|\mathbf{X}^T\|_2 \|\mathbf{e}\|_\infty + \sqrt{n} \|\mathbf{E}^T\|_2 \|\mathbf{s}\|_\infty + \|\mathbf{f}\|_\infty < q/(2p). \quad (5)$$

If we assume that $\|\mathbf{X}^T\|_2, \|\mathbf{E}^T\|_2 < y\sqrt{n}\sigma_c$ and $\|\mathbf{f}\|_\infty < y\sigma_c$, for some $y > 0$, then (5) is verified if

$$q > 2p \cdot (2y^2\sigma_c\sigma n k + y\sigma_c).$$

We bound the ℓ_2 -norms using Lemma 1 and a union-bound, and the ℓ_∞ -norms with Lemma 1 in dimension 1 and a union-bound. Using the independence of

the random variables, the assumption we made on the norms are verified with probability at least

$$\begin{aligned}
& \mathbb{P} [\|\mathbf{X}^T\|_2, \|\mathbf{E}^T\|_2 < y\sqrt{n}\sigma_c \wedge \|\mathbf{f}\|_\infty < y\sigma_c] \\
& > \left(1 - ny^n e^{\frac{n}{2}(1-y^2)}\right)^2 \left(1 - 2ne^{-\frac{y^2}{2}}\right) \\
& > \left(1 - 2ny^n e^{\frac{n}{2}(1-y^2)}\right) \left(1 - 2ne^{-\frac{y^2}{2}}\right) \\
& > 1 - 4ne^{-\frac{y^2}{2}}.
\end{aligned}$$

In order to achieve (k, δ) -correctness, it suffices to set $y = \sqrt{-2\log(\delta/(4n))}$.

Notice that we implicitly assumed that the norm of the updates were bounded by $y\sigma$. As the plaintext space is \mathbb{Z}_p^n , in order to fit a secret key into an encryption, it suffices that $p > 2y\sigma$. \square

B.3 γ -Spreadness of the Scheme

We adapt the proof of [29, Lemma 6] for FrodoKEM to prove the following result.

Lemma 8. *The UPKE construction given in Section 4 is γ -spread.*

Proof. Let $\text{ct} = (\text{ct}_0, \text{ct}_1)$ be an element of the ciphertext space, μ be a message and $pk = (\mathbf{A}, \mathbf{b})$ be a public key. We have:

$$\begin{aligned}
\mathbb{P} [\text{ct} = \text{Enc}(pk, \mu)] & \leq \mathbb{P}_{\mathbf{X}, \mathbf{E}} [\text{ct}_0 = \mathbf{X}\mathbf{A} + \mathbf{E}] \\
& = \sum \mathbb{P}_{\mathbf{X}, \mathbf{E}} [\text{ct}_0 = \tilde{\mathbf{X}}\mathbf{A} + \mathbf{E} \wedge \mathbf{X} = \tilde{\mathbf{X}}] \\
& = \sum \mathbb{P}_{\mathbf{E}} [\mathbf{E} = \tilde{\mathbf{X}}\mathbf{A} + \text{ct}_0] \cdot \mathbb{P} [\mathbf{X} = \tilde{\mathbf{X}}] \\
& \leq \sum \mathbb{P} [\mathbf{E} = \mathbf{0}] \cdot \mathbb{P} [\mathbf{X} = \tilde{\mathbf{X}}] \\
& = \mathbb{P} [\mathbf{E} = \mathbf{0}] \\
& = (\mathcal{D}_{\mathbb{Z}, \sigma}(0))^{n^2}
\end{aligned}$$

where the fourth inequality stems from the fact that the distribution $\mathcal{D}_{\mathbb{Z}, \sigma}(x)$ is maximal at $x = 0$. \square

C Analysis of the FO Transform (Theorem 3)

The (k, δ) -correctness of FO(UPKE, G, H) in the ROM follows from the (k, δ) -correctness of the underlying UPKE scheme, since Encaps runs the Enc algorithm, Decaps runs the Dec algorithm and the underlying KeyGen, UpdatePk algorithms are unchanged.

In Figure 1, we present the random oracles and decapsulation oracles as they are in the original IND-CR-CCA UKEM game. The idea of the proof is the same as for usual proofs of FO: we modify oracles to allow the challenger to simulate the decapsulation oracle without knowledge of the secret key sk . The additional pk in the inputs of the oracle \mathbf{G} allows the challenger to keep track of the ciphertexts known by the adversary for any public key pk , through epochs.

Algorithm 1: Oracles \mathbf{G} , \mathbf{H} and \mathcal{O}_{dec} for Game 0.

<pre> 1 $\mathbf{G}_0(pk, m)$: 2 if $\exists r : (pk, m, r) \in \mathcal{L}_G$ then 3 return r 4 end 5 $r \leftarrow \mathcal{U}(\mathcal{R})$; 6 $\mathcal{L}_G = \mathcal{L}_G \cup \{(pk, m, r)\}$; 7 return r </pre>	<pre> 8 $\mathbf{H}_0(m, c)$: 9 if $\exists K : (m, c, K) \in \mathcal{L}_H$ then 10 return K 11 end 12 $K \leftarrow \mathcal{U}(\mathcal{K})$; 13 $\mathcal{L}_H = \mathcal{L}_H \cup \{(m, c, K)\}$; 14 return K </pre>
	<pre> 15 $\mathcal{O}_{dec,0}(c)$: 16 if $c = c^* \wedge pk_t = pk_{chall}$ then 17 Abort 18 end 19 return $\text{Decaps}(sk_t, c)$ </pre>

We add a subscript i to the oracle names to refer to the implementation of this oracle in Game i . For instance, oracle \mathbf{G}_0 refers to the oracle \mathbf{G} in Game 0. When the context is clear, we omit the subscript. We let \mathcal{K} denote the key space and \mathcal{R} the space of the randomness used by algorithm Enc .

Let us define the following sequence of games. Note that, in each game, the challenger initializes all relevant lists \mathcal{L}_H , \mathcal{L}_G , or \mathcal{L}_E to \emptyset at the start of the game.

- Game 0: This is the original IND-CR-CCA UKEM game, using oracles as they are described in Figure 1.
- Game 1: In this game, we modify both the random oracles and the decapsulation oracle. We replace the oracles of Figure 1 by those in Figure 2. The main difference is that oracle \mathbf{G} on input (pk, m) keeps track of $(pk, m, \text{Enc}(pk, m; r), r)$, where r is the output of $\mathbf{G}(pk, m)$. This allows for oracles \mathbf{H} and \mathcal{O}_{dec} to know, for every epoch t , if they are queried on valid encapsulations for pk_t .
- Game 2: In this game, the challenger additionally aborts if the adversary makes a query $\mathbf{G}(pk, m^*)$ or $\mathbf{H}(m^*, c)$ with m^* being the (uniformly random) message used to compute the challenge encapsulation c^* , where pk and c are arbitrary. As the adversary \mathcal{A} cannot learn $\mathbf{H}(m^*, c^*)$, no information about it is available to the adversary. Hence $\text{Adv}^{G_2}(\mathcal{A}) = 0$, and Games 1 and 2 are indistinguishable up to the adversary making a query using m^* .

Algorithm 2: Oracles G , H and \mathcal{O}_{dec} for Game 1. Here pk_t denotes the public key at the current epoch t .

<pre> 1 $\mathsf{G}_1(pk, m)$: 2 if $\exists r : (pk, m, r) \in \mathcal{L}_G$ then 3 return r 4 end 5 $r \leftarrow \mathcal{U}(\mathcal{R})$; 6 $c = \text{Enc}(pk, m; r)$; 7 $\mathcal{L}_E = \mathcal{L}_E \cup \{(pk, m, r, c)\}$; 8 $\mathcal{L}_G = \mathcal{L}_G \cup \{(pk, m, r)\}$; 9 return r 10 $\mathsf{H}_1(m, c)$: 11 if $\exists K : (m, c, K) \in \mathcal{L}_H$ then 12 return K 13 end 14 $K \leftarrow \mathcal{U}(\mathcal{K})$; 15 if $\exists pk, r : (pk, m, c, r) \in \mathcal{L}_E$ then 16 $\mathcal{L}_D = \mathcal{L}_D \cup \{(pk, m, c, K)\}$; 17 end 18 $\mathcal{L}_H = \mathcal{L}_H \cup \{(m, c, K)\}$; 19 return K </pre>	<pre> 20 $\mathcal{O}_{dec,1}(c)$: 21 if $c = c^* \wedge pk_t = pk_{chall}$ then 22 abort 23 end 24 if $\exists m, K : (pk_t, m, c, K) \in \mathcal{L}_D$ then 25 return K 26 end 27 if $\exists m, r : (pk_t, m, c, r) \in \mathcal{L}_E$ then 28 $K \leftarrow \mathcal{U}(\mathcal{K})$; 29 $\mathcal{L}_H = \mathcal{L}_H \cup \{(m, c, K)\}$; 30 $\mathcal{L}_D = \mathcal{L}_D \cup \{(pk_t, m, c, K)\}$; 31 return K 32 end 33 return \perp </pre>
--	--

Let us now prove that the above games are indistinguishable in the adversary's view.

Indistinguishability of Games 0 and 1. Compared to G_0 , oracle G_1 only performs additional bookkeeping operations. Hence there is no difference between G_0 and G_1 for the adversary. Oracle H_1 might behave differently than H_0 only if a decapsulation query is made to \mathcal{O}_{dec} for a c such that $(pk_t, m, c, r) \in \mathcal{L}_E$ for some (m, r) , where t is the current epoch. Consider the case where the adversary makes a query c to the decapsulation oracle \mathcal{O}_{dec} at epoch t :

1. Assume that $\mathcal{O}_{dec,0}(c) = \perp$ and $\mathcal{O}_{dec,1}(c) \neq \perp$: then by the definition of $\mathcal{O}_{dec,1}$, this implies that there exists⁴ $(pk_t, m, r, c) \in \mathcal{L}_E$ such that $c = \text{Enc}(pk_t, m; r)$, where $r = \mathsf{G}(pk_t, m)$. As we assumed $\mathcal{O}_{dec,0}(c) = \perp$, the original decapsulation function fails on c , hence r is such that we have $\text{Dec}(sk_t, \text{Enc}(pk_t, m; r)) \neq m$. By the (k, δ) -correctness, this happens with probability at most δ .
2. Assume that $\mathcal{O}_{dec,0}(c) \neq \perp$ and $\mathcal{O}_{dec,1}(c) = \perp$: by the definition of $\mathcal{O}_{dec,1}$, this implies that there is no $(pk_t, m, r, c) \in \mathcal{L}_E$, hence \mathcal{A} did not make any query $\mathsf{G}(pk_t, m)$ but was able to compute a valid ciphertext of m under pk_t . By γ -spreadness, this happens only with probability at most γ .
3. Assume that $\mathcal{O}_{dec,0}(c) = K$ and $\mathcal{O}_{dec,1}(c) = K' \neq \perp$: by the definition of $\mathcal{O}_{dec,1}$ we know that there exists $(pk_t, m, r, c) \in \mathcal{L}_E$ such that

⁴ Note that the only way $\mathcal{O}_{dec,1}(c)$ returns $K \neq \perp$ is that either $(pk_t, m, r, c) \in \mathcal{L}_E$ or that oracle H_1 added (pk_t, m, c, K) to \mathcal{L}_D . However, the latter only happens if $(pk_t, m, r, c) \in \mathcal{L}_E$. Thus $\mathcal{O}_{dec,1}(c)$ does not return \perp only if $(pk_t, m, r, c) \in \mathcal{L}_E$.

$c = \text{Enc}(pk_t, m; r)$, and as $\mathcal{O}_{dec,0}(c) \neq \perp$, this is a valid encryption. Hence $K = H_0(m, c)$. We consider the two following sub-cases:

- (a) Adversary \mathcal{A} first made the decryption query $\mathcal{O}_{dec}(c)$ without knowing $H(m, c)$. By definition of $\mathcal{O}_{dec,1}(c)$, the challenger samples $K' \leftarrow \mathcal{U}(\mathcal{K})$ and adds (m, c, K') to \mathcal{L}_H . By definition of H_1 , we have $H_1(m, c) = K'$. Thus K' has the same distribution as K and H_1 has the same behaviour as H_0 .
- (b) Adversary \mathcal{A} already knows $H(m, c)$ as it queried it before to the oracle H . It is then set to a uniformly random value $K' \leftarrow \mathcal{U}(\mathcal{K})$. Then, when the adversary makes the decryption query $\mathcal{O}_{dec}(c)$, the definition of $H_1(m, c)$ guarantees that $\mathcal{O}_{dec,1}(c)$ returns K' , which has the same distribution as K and H_1 behaves identically to H_0 .

We just showed that except with probability at most $q_G \cdot \delta + q_D \cdot \gamma$, Games 1 and 2 behave identically. Further note that in Game 1, for any epoch t , oracle queries to \mathcal{O}_{dec} can be simulated without the knowledge of the secret key sk_t .

Indistinguishability of Games 1 and 2. Let us call FIND the event that an adversary \mathcal{A} makes a query $G(pk, m^*)$ or $H(m^*, c)$ with m^* being the (uniformly random) message used to compute the challenge encapsulation c^* , where pk and c are arbitrary. As already detailed, adversary \mathcal{A} has advantage at most $\mathbb{P}[\text{FIND}]$ in distinguishing between Games 1 and 2. We now bound the probability $\mathbb{P}[\text{FIND}]$ by constructing an adversary \mathcal{B} for the IND-CR-CPA game such that

$$\mathbb{P}[\text{FIND}] \leq 2 \left(\text{Adv}^{\text{IND-CR-CPA}}(\mathcal{B}) + \frac{q_G + q_H}{|\mathcal{M}|} \right). \quad (6)$$

Adversary \mathcal{B} first receives pk_0 from its IND-CR-CPA challenger and forwards pk_0 to \mathcal{A} . Whenever \mathcal{A} makes an \mathcal{O}_{up} oracle query, adversary \mathcal{B} makes the same \mathcal{O}_{up} query to its challenger. Whenever \mathcal{A} makes a G, H or \mathcal{O}_{dec} query, adversary \mathcal{B} runs them as in Game 1, which is possible as it does not need to know the secret key, as observed above. When \mathcal{A} requests a challenge, \mathcal{B} samples two random messages $m_0, m_1 \leftarrow \mathcal{U}(\mathcal{M})$ and sends them to its challenger.

The challenger answers with the IND-CR-CPA challenge c^* . Adversary \mathcal{B} samples $K^* \leftarrow \mathcal{U}(\mathcal{K})$ and sends the challenge (K^*, c^*) to \mathcal{A} .

From now, adversary \mathcal{B} continues to simulate \mathcal{A} 's challenger. If \mathcal{A} makes a query $G(pk, m_{b'})$ or $H(m_{b'}, c)$ for any $b' \in \{0, 1\}$, adversary \mathcal{B} stops running \mathcal{A} and returns b' to its challenger. If \mathcal{A} makes no such request, then \mathcal{B} samples $b' \leftarrow \mathcal{U}(\{0, 1\})$ and returns b' .

Call WRG the event that \mathcal{A} makes an oracle query to G or H containing m_{1-b} , where b is the challenge bit. Since \mathcal{A} has absolutely no information about m_{1-b} ,

this happens with probability at most $\mathbb{P}[\text{WRG}] \leq (q_G + q_H)/|\mathcal{M}|$. Then:

$$\begin{aligned}
\text{Adv}^{\text{IND-CR-CPA}}(\mathcal{B}) &= \left| \mathbb{P}[b = b'] - \frac{1}{2} \right| \\
&= \left| \mathbb{P}[\text{FIND} \wedge \neg\text{WRG}] + \frac{1}{2}\mathbb{P}[\neg\text{FIND}] - \frac{1}{2} \right| \\
&= \left| \mathbb{P}[\text{FIND}] - \mathbb{P}[\text{FIND} \wedge \text{WRG}] + \frac{1}{2}\mathbb{P}[\neg\text{FIND}] - \frac{1}{2} \right| \\
&\geq \frac{1}{2}\mathbb{P}[\text{FIND}] - \mathbb{P}[\text{FIND} \wedge \text{WRG}] \\
&\geq \frac{1}{2}\mathbb{P}[\text{FIND}] - \mathbb{P}[\text{WRG}].
\end{aligned}$$

The second equality holds as \mathcal{B} finds b' if and only if \mathcal{A} makes an oracle query containing m_b (i.e., both FIND and $\neg\text{WRG}$ occur) or if no such query occurs, by guessing randomly. For the third equality, we use that for any two events A, B , we have $\mathbb{P}[A \wedge B] = \mathbb{P}[A] - \mathbb{P}[A \wedge \neg B]$. Equation (6) then follows, which completes the proof of Theorem 3. \square

D Analysis of the IND-CU-CCA Transform (Theorem 4)

We proceed by a sequence of hybrid games.

Game 0: This is the original IND-CU-CCA game where the challenger's bit is set to $b = 0$.

Game 1: We replace the proof π^* in the final update $up^* = (\text{ct}_0^*, \text{ct}_1^*, \pi^*)$ by a simulated NIZK proof. As the adversary only sees this simulated proof at the very end of the game and cannot submit any additional update or decryption queries, the two games are indistinguishable thanks to the computational zero-knowledge property of the underlying proof system.

Game 2: We now change the plaintext underlying ct_1^* to an encryption of 0 rather than r . This change remains undetected thanks to the IND-CPA security of the underlying encryption scheme. As an important remark, note that IND-CPA security (which is implied by IND-CR-CCA security) suffices here as no information about sk_1 is provided to the adversary, since neither the decapsulation oracle nor the final secret contain information about sk_1 .

Game 3: In this game, when the adversary makes an update query which passes `VerifyUpdate`, the challenger does the following. Let $((\text{ct}_0, \text{ct}_1, \pi), (pk'_0, pk_1))$ denote such a query. Then, the challenger uses both secret keys sk_0 and sk_1 to decrypt ct_0 and ct_1 and verify that the underlying plaintexts are indeed equal and that the new public key pk'_0 is computed honestly. It halts if it is not the case. Unless Game 3 aborts, the two games are identical. The computational soundness of H guarantees that any PPT adversary cannot trigger an abort, except with negligible probability. Here, we insist that standard (computational)

soundness suffices as the adversary does not receive any proof until it can no longer make queries.

Game 4: This final game is identical to the previous game except that the challenger's bit is 1. We show that these two games are indistinguishable under the IND-CR-CCA security of UKEM.

Assume there exists a PPT adversary \mathcal{A} that can distinguish Game 3 and Game 4. We construct a PPT adversary \mathcal{B} against the IND-CR-CCA security of UKEM as follows. Adversary \mathcal{B} gets pk_0 from its challenger and further samples an additional key pair $(pk_1, sk_1) \leftarrow \text{KeyGen}(1^\lambda)$. It also implements a random oracle F by storing a table. It forwards (pk_0, pk_1) to \mathcal{A} as the public key.

When \mathcal{A} makes a decapsulation query, adversary \mathcal{B} simply submits the same query to its decapsulation oracle and returns the result to \mathcal{A} . When \mathcal{A} makes an update query $((ct_0, ct_1, \pi), (pk'_0, pk_1))$, adversary \mathcal{B} verifies the validity of π and if it passes verification, uses sk_1 to decrypt ct_1 in order to recover the randomness r used by \mathcal{A} to generate its update. Adversary \mathcal{B} can then submit r to its own update oracle to produce the same update.

When \mathcal{A} asks for a challenge, so does \mathcal{B} , and the latter forwards its challenge encapsulation c^* to the former.

Finally, when \mathcal{A} stops making updates, so does \mathcal{B} . Its challenger then replies by (pk^*, sk^*, up^*) , where up^* is simply an encryption ct_0^* of the last (unknown) update under the last epoch public key pk_{last}^ℓ . It generates an encryption ct_1^* of 0 under pk_1 , as well as a simulated proof π^* that $(pk_{last}^\ell, pk_1, pk^*, ct_0^*, ct_1^*)$ is a valid update. It finally sends the tuple $((pk^*, pk_1), sk^*, ct_0^*, ct_1^*, \pi^*)$ to \mathcal{A} . When \mathcal{A} halts with output b' , so does \mathcal{B} .

This completes the proof of Theorem 4. □