

Unidirectional Updatable Encryption and Proxy Re-encryption from DDH or LWE*

Peihan Miao¹, Sikhar Patranabis², and Gaven Watson³

¹ University of Illinois Chicago, peihan@uic.edu

² IBM Research India, sikhar.patranabis@ibm.com

³ Visa Research, gawatson@visa.com

Abstract

Updatable Encryption (UE) and Proxy Re-encryption (PRE) allow *re-encrypting* a ciphertext from one key to another in the symmetric-key and public-key settings, respectively, without decryption. A longstanding open question has been the following: do *unidirectional* UE and PRE schemes (where ciphertext re-encryption is permitted in only one direction) necessarily require stronger/more structured assumptions as compared to their bidirectional counterparts? Known constructions of UE and PRE seem to exemplify this “gap” – while bidirectional schemes can be realized as relatively simple extensions of public-key encryption from standard assumptions such as DDH or LWE, unidirectional schemes typically rely on stronger assumptions such as FHE or indistinguishability obfuscation (iO), or highly structured cryptographic tools such as bilinear maps or lattice trapdoors.

In this paper, we bridge this gap by showing the first feasibility results for realizing unidirectional UE and PRE from a new generic primitive that we call Key and Plaintext Homomorphic Encryption (KPHE) – a public-key encryption scheme that supports additive homomorphisms on its plaintext and key spaces simultaneously. We show that KPHE can be instantiated from DDH or LWE. This yields, in particular, the first constructions of unidirectional UE and PRE from DDH, as well as the first constructions of unidirectional UE and PRE from LWE that do not resort to FHE or lattice trapdoors.

Our constructions achieve the strongest notions of *post-compromise security* in the standard model. Our UE schemes also achieve “backwards-leak directionality” of key updates (a notion we discuss is equivalent, from a security perspective, to that of unidirectionality with no-key updates). Our results establish (somewhat surprisingly) that unidirectional UE and PRE schemes satisfying such strong security notions *do not*, in fact, require stronger/more structured cryptographic assumptions as compared to bidirectional schemes.

1 Introduction

Cryptographic encryption is a powerful tool for ensuring data confidentiality. A common security guarantee offered by any encryption scheme (either symmetric-key or public-key) is the following:

*The authors grant IACR a non-exclusive and irrevocable license to distribute the article under the CC BY-NC (creative commons attribution-noncommercial) license. S. Patranabis did the work while at Visa Research. P. Miao did part of the work while at Visa Research. P. Miao is supported in part by the NSF CNS Award 2055358 and a 2020 DPI Science Team Seed Grant.

encrypted data can only be decrypted using a certain secret key. However, a limitation of traditional encryption schemes is that once data is encrypted, it is generally hard to allow a third party to transform the ciphertext so that it can be decrypted with different key, without sharing either the original or the new secret key with the third party.

Re-encryption schemes such as Proxy Re-encryption (PRE) [BBS98] and Updatable Encryption (UE) [BLMR13] circumvent this limitation by enabling a public transformation of ciphertexts from encryption under one key to that of another, while protecting the underlying secret keys. Classic applications for such schemes include key rotation for secure outsourced storage [BBB⁺12, Pay18], access control, delegation of email access, and many more.

Proxy Re-encryption (PRE). PRE is a public-key encryption scheme which enables a party Alice, with the help of a proxy, to re-encrypt her ciphertexts for decryption by an alternate party Bob. To facilitate re-encryption, Alice and Bob, with key pairs (pk_A, sk_A) and (pk_B, sk_B) respectively, will together compute a re-encryption key rk_{AB} and then provide this to the proxy. Whenever the proxy needs to perform re-encryption, it can use rk_{AB} to transform a ciphertext encrypted under pk_A into a ciphertext encrypted under pk_B . Security of the PRE scheme guarantees that the proxy learns nothing about the underlying plaintext during the re-encryption process.

Updatable Encryption (UE). UE was introduced by Boneh et al. [BLMR13] to address the problem of key rotation for secure outsourced storage. UE addresses re-encryption by using similar techniques to those of PRE, with two main differences: (1) UE is a symmetric-key encryption scheme, and (2) UE typically only allows sequential updates. More specifically, in UE we divide time into a series of epochs. In the first epoch a fresh symmetric key k_0 is chosen and used to encrypt all data. When we rotate a key from k_{e-1} to k_e , we transition to the next epoch by calculating an update token Δ_e . All new ciphertexts are encrypted under the new key k_e and all existing ciphertexts ct_{e-1} are re-encrypted using the update token Δ_e so that they can be decrypted by k_e . The benefit of this approach is that the storage server can perform the re-encryption of data using the update token without the risk of exposing any plaintext data.

There are two variants of UE schemes, *ciphertext-dependent* schemes [EPRS17, BEKS20] and *ciphertext-independent* schemes [LT18, KLR19, BDGJ20, Jia20]. In ciphertext-dependent UE schemes, the update token $\Delta_{e, ct_{e-1}}$ depends on the ciphertext ct_{e-1} to be updated, while in ciphertext-independent schemes, the update token Δ_e is generated independent of the updated ciphertext, hence a single token can be used to update all ciphertexts on the storage server. In the rest of the paper, when we refer to UE, we mean a ciphertext-independent scheme unless otherwise specified.

Directionality of PRE and UE. Re-encryption schemes are either *bidirectional* or *unidirectional*. A scheme is said to be bidirectional if a re-encryption key/update token can be used to re-encrypt a ciphertext to either the next party/epoch or the previous party/epoch. In contrast, the re-encryption key/update token of a unidirectional scheme can only be used to re-encrypt a ciphertext to the next party/epoch and *not* the previous. So far we have only discussed the directionality within the context of ciphertext updates. The (uni)directionality with regards to keys differs slightly between PRE and UE, as we discuss next.

Bidirectionality vs. Unidirectionality in PRE. In bidirectional PRE schemes, the re-encryption key rk_{AB} from Alice to Bob is generated from Alice’s key-pair (pk_A, sk_A) and Bob’s

key-pair (pk_B, sk_B) . Given rk_{AB} along with sk_A (resp. sk_B), it is usually possible to derive sk_B (resp. sk_A). In unidirectional PRE schemes, the re-encryption key rk_{AB} is derived from $((pk_A, sk_A), pk_B)$; Bob’s secret key sk_B is not used. In fact, given the re-encryption key rk_{AB} and Alice’s secret key sk_A , it should be impossible to derive any knowledge of Bob’s secret key sk_B .

Unidirectionality in UE. For UE schemes, there is an extra level of subtlety regarding the *directionality of keys* in addition to ciphertexts. A recent work of Jiang [Jia20] extensively studied the question: given an update token Δ_e along with either k_{e-1} or k_e , is it possible to derive the other key? A scheme has bidirectional key updates if Δ_e can be used to derive keys in both directions, and has unidirectional key updates if Δ_e can be used in one direction, to derive k_e from k_{e-1} . Jiang [Jia20] showed that UE with bidirectional key and ciphertext updates implies UE with unidirectional key and ciphertext updates.

In the same work, Jiang postulated that to capture the same security level as the unidirectional PRE schemes, one requires even stronger UE schemes with *no-directional key updates*, where k_e cannot be derived from k_{e-1} and Δ_e . In Jiang [Jia20], the definition of no-directional key updates intuitively requires that it is *also* impossible to derive k_{e-1} from Δ_e and k_e . The recent work of Nishimaki [Nis21] proposed a seemingly weaker notion called *backward-leak unidirectional key updates* where Δ_e can only be used in one direction to derive k_{e-1} from k_e . However, we observe that this new notion is essentially equivalent to no-directional key updates because derivation of k_{e-1} does not increase the adversary’s advantage in breaking the scheme. In particular, if the adversary obtains a ciphertext ct_{e-1} and corrupts Δ_e and k_e , then it can first update the ciphertext to ct_e and decrypt it using k_e . Jiang emphasized that UE with no-directional key updates is the ideal security model, which by our argument above, extends to backwards-leak key updates. Henceforth, when we refer to unidirectional UE, we mean unidirectional UE with backwards-leak directional key updates unless otherwise specified.

Gap between Unidirectionality and Bidirectionality. In general, unidirectional UE and PRE schemes are more ideally suited to real-world applications as compared to their bidirectional counterparts due to their superior security guarantees. For example, unlike bidirectional UE schemes, unidirectional UE schemes guarantee security of data as if “freshly encrypted” in epoch e (i.e., not re-encrypted from epoch $(e - 1)$) even if the adversary gains access to the secret key k_{e-1} and the update token Δ_e . Unidirectional PRE schemes also offer similarly superior security guarantees over their bidirectional counterparts.

Another natural point of comparison between unidirectional and bidirectional UE and PRE schemes is the nature of cryptographic assumptions from which such schemes can be realized. Known constructions of UE and PRE seemingly exemplify an apparent “gap” in terms of the assumptions required – unidirectional schemes have historically relied on stronger/more structured cryptographic assumptions as compared to their bidirectional counterparts.

Blaze et al. [BBS98] showed how to construct bidirectional PRE schemes from the Decisional Diffie-Hellman (DDH) assumption by suitably extending the well-known ElGamal encryption scheme [Gam85]. Similarly, a long line of works [BLMR13,LT18,KLR19,BDGJ20,Jia20] have shown how to realize bidirectional UE schemes as relatively simple extensions of public-key encryption from standard assumptions such as DDH and Learning With Errors (LWE).

On the other hand, unidirectional UE and PRE schemes typically rely on a stronger set of assumptions such as FHE [Gen09] and indistinguishability obfuscation (iO) [BGI⁺12], or highly

structured cryptographic tools such as bilinear maps [BF03] and “hard” lattice trapdoors [GPV08]. Examples of constructions of unidirectional PRE from FHE and/or structured lattice trapdoors can be found in [NX15, CCL⁺14, Kir14, NAL15, PWA⁺16, FL17, PRSV17]. Constructions of unidirectional PRE schemes have also been shown to exist from bilinear maps [AFGH06, LV11]; however, these constructions are restricted to the *single-hop* setting in the sense that they only permit a single re-encryption of a ciphertext. Known constructions of unidirectional UE include the construction in [SS21] (which relies on bilinear maps), and two constructions in [Nis21] (one which achieves backward-leak key updates from lattice-specific techniques, and one which achieves non-directional key updates from iO). Sehwat and Desmedt [SD19] show a construction of UE from bi-homomorphic lattice-based pseudorandom functions; however, their construction only achieves unidirectional ciphertext updates while still incurring bidirectional key updates (and is hence effectively bidirectional as per the recent findings in [Jia20]). To date, there exist no constructions of unidirectional PRE or UE from the plain DDH assumption (to our knowledge).

In this paper, we are motivated by the following longstanding open question in the study of UE and PRE:

Do unidirectional UE and PRE schemes necessarily require stronger/more structured assumptions as compared to their bidirectional counterparts?

More concretely, we ask the following questions:

Can we construct unidirectional UE/PRE schemes from DDH?

Can we construct unidirectional UE/PRE schemes from LWE without resorting to FHE or highly structured tools such as lattice trapdoors?

1.1 Our Results

In this paper, we bridge this gap between the assumptions for unidirectional and bidirectional UE/PRE. We establish (somewhat surprisingly) that unidirectional UE and PRE schemes *do not*, in fact, require stronger/more structured cryptographic assumptions as compared to their bidirectional counterparts.

More concretely, we present generic constructions of unidirectional UE and PRE from a new primitive that we call Key and Plaintext Homomorphic Encryption (KPHE). We also show that such a KPHE scheme can be instantiated from the BHHO encryption scheme [BHHO08] based on the DDH assumption, and Regev’s encryption scheme [Reg09] based on the LWE assumption. This yields, in particular, the first constructions of unidirectional UE and PRE from the plain DDH assumption, as well as the first constructions of unidirectional UE and PRE from LWE without resorting to FHE or lattice trapdoors.

Our main result is summarized by the following (informal) theorem:

Theorem 1.1 (Informal). *Assuming the existence of a Key and Plaintext Homomorphic Encryption (KPHE) scheme that satisfies certain special properties, there exist post-compromise secure unidirectional UE and PRE schemes.*

On KPHE. The KPHE scheme with special properties required in our constructions can be viewed as a generalization of the BHHO public-key encryption scheme due to Boneh et al. [BHHO08]. It is a public key encryption scheme where the secret key is a bit-string $\mathbf{sk} \in \{0, 1\}^\ell$ and the plaintext is also a bit-string $\mathbf{m} \in \{0, 1\}^{\ell'}$ (in our constructions we use $\ell = \ell' = 2n$). The specialized KPHE scheme satisfies the following three properties:

- **Distributional Semantic Security:** We require a KPHE scheme to achieve semantic security even when the secret keys are sampled from a specific distribution. In particular, we use KPHE schemes with $2n$ -bit secret keys where the secret key is uniformly random subject to the constraint that it has equally many 0 and 1 bits (i.e., n bits of 0 and n bits of 1).
- **Additive Key and Plaintext Homomorphisms:** We require a KPHE scheme to satisfy the following property: let T, T' be two arbitrary affine transformations that map 0-1 vectors to 0-1 vectors of the same length (in our constructions we use permutation maps over the bits of a $2n$ -bit string). Then, given a public key \mathbf{pk} corresponding to some secret key \mathbf{sk} and a ciphertext $\mathbf{ct} \stackrel{\$}{\leftarrow} \text{Enc}(\mathbf{pk}, \mathbf{m})$, one can generate a public key \mathbf{pk}' corresponding to the secret key $T(\mathbf{sk})$ and a ciphertext $\mathbf{ct}' \stackrel{\$}{\leftarrow} \text{Enc}(\mathbf{pk}', T'(\mathbf{m}))$, without the knowledge of the original secret key \mathbf{sk} or the original message \mathbf{m} .
- **Blinding:** We also require the KPHE scheme to satisfy an associated security property called “blinding”, that (informally) argues that the public key and ciphertext generated via the aforementioned homomorphic transformations are indistinguishable from freshly generated public keys and ciphertexts (we make this more formal in Section 2).

For our PRE constructions, we also require that the KPHE scheme satisfies a notion of distributional circular security (i.e., circular security when the secret keys are sampled from a specific distribution). This is not required for our UE constructions.

Instantiating KPHE. We show how to concretely instantiate a KPHE scheme satisfying all of the aforementioned properties from DDH (based on the BHHO scheme [BHHO08]) or LWE (based on Regev’s encryption scheme [Reg09]).

Lemma 1.2 (Informal). *Assuming decisional Diffie-Hellman (DDH) or learning with errors (LWE) holds, there exists a secure construction of KPHE that satisfies the aforementioned properties.*

Corollary 1.3 (Informal). *Assuming decisional Diffie-Hellman (DDH) or learning with errors (LWE) holds, there exist post-compromise secure unidirectional UE and PRE schemes.*

Security of Our Constructions. Our constructions of unidirectional UE and PRE achieve the strongest notions of *post-compromise security* in the standard model. Our construction of unidirectional UE achieves the state-of-the-art post-compromise security definition due to Boyd et al. [BDGJ20], while also ensuring backward-leak unidirectional key updates [Nis21]. Our unidirectional PRE construction achieves the post-compromise security definition recently proposed by Davidson et al. [DDL19], which is, to our knowledge, the only notion of post-compromise PRE security to be proposed to date. We present a more detailed discussion on post-compromise security (and other related security notions) of UE and PRE in the next subsection.

	scheme	ctx updates	key updates	security	assumption
UE	[BLMR13]	bidirectional	bidirectional	IND-ENC	DDH/LWE
UE	[LT18, KLR19, BDGJ20]	bidirectional	bidirectional	IND-UE	DDH
UE	[Jia20]	bidirectional	bidirectional	IND-UE	DLWE
UE	[SD19]	unidirectional	bidirectional	IND-UE	LWE
UE	[Nis21]	unidirectional	backward-leak unidirectional	IND-UE	LWE
UE	[Nis21]	unidirectional	no-directional	IND-UE	iO
UE	[SS21]	unidirectional	no-directional	IND-UE	SXDH
UE	Ours	unidirectional	backward-leak unidirectional	IND-UE	DDH/LWE
PRE	[BBS98]	bidirectional	bidirectional	IND-CPA	DDH
PRE	[PWA ⁺ 16]	unidirectional	unidirectional	IND-CCA	LWE
PRE	[PRSV17]	unidirectional	unidirectional	IND-CPA	RLWE
PRE	Ours	unidirectional	unidirectional	IND-HRA	DDH/LWE

Table 1: Summary of bi/unidirectional UE and PRE schemes. We focus on ciphertext-independent UE and multi-hop PRE.

1.2 Background and Related Work

There has been extensive research on both UE and PRE, including various settings, definitions, and constructions. Below we only mention works that are the most directly relevant. For both UE and PRE, we focus on the CPA-type definitions, which are by far the most well-studied notions.

When modeling the security of re-encryption schemes at the most basic level we must ensure confidentiality of the underlying message. In addition to this we also want to ensure that messages remain protected even in the face of key compromise – a motivating factor for key rotation. A scheme which provides *post-compromise security* (PCS) ensures that a re-encrypted/updated ciphertext remains secure even when one of the previous keys has been compromised.

Security Notions for UE. Since the introduction of UE in [BLMR13], several works have explored its security notions [EPRS17, LT18, AMP19, KLR19, BDGJ20, Jia20]. Most notable is the work of Lehmann and Tackmann [LT18], which improved the model and studied the notion of post-compromise security for UE. Their Indistinguishability of Update notion (IND-UPD) returns a challenge ciphertext ct^* which is either the re-encryption of a ciphertext ct_0 or ct_1 . A scheme is IND-UPD secure if an adversary is unable to determine which of the ciphertexts was re-encrypted.

In subsequent works a stronger combined notion of IND-UE security has been used, first defined by Boyd et al. [BDGJ20]. The IND-UE notion requires an adversary be unable to distinguish between a fresh encryption of a plaintext m and the re-encryption of a ciphertext ct . As a result this notion captures both CPA (specifically IND-ENC) and IND-UPD security.

Security Notions for PRE. In the context of PRE, the traditional notion of IND-CPA security [ID03, AFGH06] have been shown to be insufficient in practice. To address this, Cohen [Coh19] introduced the notion of Honest Re-Encryption Attack (HRA) security where an adversary is additionally permitted to re-encrypt (from honest to corrupt users) ciphertexts previously output by the encryption oracle. While only recently considered in the analysis of PRE, the essence of this notion is also fundamental in formalizing security for UE.

More recently, Davidson et al. [DDL19] have investigated achieving post-compromise secure

PRE schemes. They introduced a notion of IND-PCS security for PRE, which can be viewed as the analogue of IND-UPD security of UE in the context of PRE, albeit for more complex re-encryption graphs. To date this is the only paper which studies the PCS security of PRE schemes. Their work again demonstrates the challenges in constructing such schemes in the unidirectional setting. They discuss two PCS-secure constructions which are based on prior unidirectional PRE scheme, Construction 7b of Fuchsbauer et al. [FKKP19] and an extension of BV-PRE [PRSV17].

Updatable Public Key Encryption. In order to achieve forward security in public key encryption (PKE), a notion called *updatable PKE (UPKE)* has recently been proposed and studied [JMM19,ACDT20,DKW21], where any sender (encryptor) can initiate a key update by sending a special update ciphertext to the receiver (decryptor). This ciphertext updates the public key and also, once processed by the receiver, will update its secret key. These are PKE schemes that encrypt messages under different public keys and aim to achieve forward security. In contrast, UE and PRE schemes studied in this paper aim to update ciphertexts encrypted under an old key to a new key without leaking the message content. The notions of UE/PRE as well as our techniques are very different from UPKE despite the partial naming collision.

1.3 Technical Overview

In this section, we provide a high-level overview of our techniques for constructing unidirectional UE and PRE from any generic KPHE scheme satisfying the special properties described earlier.

IND-ENC Secure UE. Our first attempt is to build a unidirectional IND-ENC secure UE scheme, and we start with a naïve idea. Take an arbitrary symmetric-key encryption scheme and each epoch key is a freshly generated key of this encryption scheme. The update token Δ_e from k_{e-1} to k_e is an encryption of k_{e-1} under k_e , namely $\Delta_e = \text{Enc}_{k_e}(k_{e-1})$. When we update a ciphertext from epoch $(e-1)$ to epoch e , we just attach the update token Δ_e to the end of the ciphertext. For a message m first encrypted in epoch e and then updated through epoch e' , the resulting ciphertext is of the form:

$$ct_{e'} = (\text{Enc}_{k_e}(m), \text{Enc}_{k_{e+1}}(k_e), \dots, \text{Enc}_{k_{e'}}(k_{e'-1})).$$

Given $k_{e'}$, one can easily decrypt $ct_{e'}$ layer by layer to recover m .

This naïve approach does not achieve IND-ENC security. We show a concrete attack in the following. Let e^* be the challenge epoch and m^* be the challenge message queried by the adversary. Let $ct_{e^*} = \text{Enc}_{k_{e^*}}(m^*)$ be the challenge ciphertext. To extract the secret key k_{e^*} , the adversary proceeds as follows. It first queries for an encryption of an arbitrary message m in epoch 0 and then updates it to epoch e (for some $e > e^*$) via a sequence of update queries. This way the adversary obtains a ciphertext of m of the form:

$$ct_e = (\text{Enc}_{k_0}(m), \text{Enc}_{k_1}(k_0), \dots, \text{Enc}_{k_e}(k_{e-1})).$$

Now the adversary corrupts the secret key k_e . Then it can recover all the previous keys from k_0 to k_{e-1} (including k_{e^*}) during decryption of the ciphertext ct_e .

Nonetheless, this simple approach demonstrates some nice properties of unidirectionality. For key updates, it is impossible to derive k_e from k_{e-1} and Δ_e . For ciphertext updates, given a *fresh* ciphertext ct_e in epoch e and the previous update token Δ_e (from epoch $(e-1)$ to e), it is impossible to transition the ciphertext ct_e to the previous epoch ct_{e-1} (i.e. the epoch prior to its existence).

In fact, Cohen [Coh19] applied this idea to PRE and showed a CPA-secure but not HRA-secure PRE scheme (HRA security is inherently required in IND-ENC UE schemes).

Re-randomizing the Secret Keys. The problem with these chained ciphertexts is that during decryption of a single ciphertext, all the previous secret keys are also leaked. To resolve this problem, our hope is to somehow re-randomize all the previous secret keys in the chain, in a consistent and homomorphic manner. In particular, we want the ciphertext to be of the form

$$\text{ct}_e = \left(\text{Enc}_{\bar{k}_0}(m), \text{Enc}_{\bar{k}_1}(\bar{k}_0), \dots, \text{Enc}_{\bar{k}_{e-1}}(\bar{k}_{e-2}), \text{Enc}_{k_e}(\bar{k}_{e-1}) \right),$$

where $\bar{k}_0, \bar{k}_1, \dots, \bar{k}_{e-1}$ are all re-randomized secret keys that are different for each ciphertext. During the decryption of ct_e , only these re-randomized secret keys are leaked, which does not affect the security of other ciphertexts.

To enable such re-randomization, our idea is inspired by the re-randomizable Yao’s garbled circuits [GHV10]. We propose a new primitive called Key and Plaintext Homomorphic Encryption (KPHE), which can be seen as a generalization of the circular secure encryption scheme of Boneh et al. [BHHO08]. Instead of using an arbitrary symmetric-key encryption scheme, we use the KPHE scheme for encryption, where the UE secret key k_e is a key pair (pk_e, sk_e) of the KPHE scheme. The update token is a KPHE encryption of the previous epoch’s secret key under the current epoch’s public key, namely $\Delta_e = \text{KPHE.Enc}_{pk_e}(sk_{e-1})$.

To update a ciphertext we exploit the two homomorphism properties of the KPHE scheme, in both the message space and the key space. Given an update token Δ_e and a ciphertext of the form

$$\text{ct}_{e-1} = \left(\text{KPHE.Enc}_{\bar{pk}_0}(m), \text{KPHE.Enc}_{\bar{pk}_1}(\bar{sk}_0), \dots, \text{KPHE.Enc}_{pk_{e-1}}(\bar{sk}_{e-2}) \right),$$

we focus on the last component $\text{ctx} = \text{KPHE.Enc}_{pk_{e-1}}(\bar{sk}_{e-2})$ and the update token $\Delta_e = \text{KPHE.Enc}_{pk_e}(sk_{e-1})$. In our update operation we first generate a random permutation π and then perform two important steps:

- Use the KPHE key-space homomorphism to transform ctx from an encryption under sk_{e-1} to an encryption under $\pi(sk_{e-1})$.
- Use the KPHE message-space homomorphism to transform Δ_e from an encryption of sk_{e-1} to an encryption of $\pi(sk_{e-1})$.

The updated ciphertext becomes

$$\text{ct}_e = \left(\text{KPHE.Enc}_{\bar{pk}_0}(m), \text{KPHE.Enc}_{\bar{pk}_1}(\bar{sk}_0), \dots, \text{KPHE.Enc}_{\bar{pk}_{e-1}}(\bar{sk}_{e-2}), \text{KPHE.Enc}_{pk_e}(\bar{sk}_{e-1}) \right),$$

where $\bar{sk}_{e-1} = \pi(sk_{e-1})$ (with corresponding public key \bar{pk}_{e-1}). In our construction, the KPHE secret key is a $2n$ -bit string, which is randomly sampled with exactly n bits of 0 and n bits of 1. The affine transformation π is a random permutation on the $2n$ bits of the string. By transforming from sk_{e-1} to \bar{sk}_{e-1} we ensure that a fresh secret key is used for each update operation and hence there is appropriate isolation between all ciphertexts updated in a given epoch. The blinding property of KPHE ensures that re-randomization can be done without knowledge of the underlying secret keys, and that the re-randomized ciphertexts are computationally indistinguishable from freshly generated ciphertexts.

Achieving Post-Compromise Secure UE. We can extend the IND-ENC secure UE construction to achieve post-compromise security. To achieve IND-UPD security, we can modify the update operation to ensure that all the chained ciphertexts are updated (rather than just the last one). In effect what our enhanced construction does is again exploit properties of the KPHE scheme to re-randomize each of the ciphertext components. This ensures that two updated ciphertexts of the same length are computationally indistinguishable. To further achieve the combined IND-UE security, we need to additionally guarantee that a freshly generated ciphertext has the same length as an updated ciphertext in a certain epoch. More details on our UE constructions are given in Section 3.

Achieving Unidirectional PRE. We can use the same high-level approach to construct a unidirectional PRE scheme, where a ciphertext consists of a chain of KPHE ciphertexts, and re-encryption exploits the two KPHE homomorphisms to transform each new KPHE ciphertext to a fresh secret key. The crucial subtlety in the PRE case, which makes proving security slightly more involved, is that we no longer consider sequential ciphertext updates but must consider re-encryption between all possible key pairs. As a result we need to further exploit the circular security properties of the KPHE scheme to prove security. This is further detailed in Section 4.

Connections between UE and PRE. Generally speaking, unidirectional PRE can be viewed as a stronger primitive than unidirectional UE because UE only allows for sequential updates while PRE allows for re-encryption between every pair of keys. In fact, we observe that if we treat the public-secret key pair of PRE as a secret key for UE, and the PRE re-encryption key as an update token for UE, then IND-HRA secure PRE implies IND-ENC secure UE, and IND-PCS secure PRE implies IND-UPD secure UE. This is also why our constructions for unidirectional UE and PRE follow a very similar framework. On the other hand, since PRE supports re-encryption between (potentially) every pair of keys, our constructions of PRE require stronger security guarantees (in particular, circular security) from the underlying KPHE scheme.

Efficiency and Feasibility. We acknowledge that the ciphertext length in our UE/PRE constructions grows linearly with the number of epochs/re-encryption hops, unlike certain existing constructions (e.g. in [Nis21, PWA⁺16, PRSV17]) where the ciphertext size remains the same. In this context, we emphasize that we primarily aim to demonstrate feasibility of unidirectional-UE/PRE from standard assumptions e.g. DDH, which has been a longstanding open problem. We believe that our results should be viewed with emphasis on the new theoretical insights/understanding into unidirectional-UE/PRE that they enable as opposed to concrete efficiency. We also point out that while the size of ciphertexts in our constructions grows linearly, the secret keys and update tokens/re-encryption keys remain constant-sized. We also note that for the basic versions of our UE/PRE construction (IND-CPA unidirectional UE and IND-HRA secure unidirectional PRE), the work done per update/re-encryption operation is also constant (independent of the number of epochs/update hops).

We note here that a naïve approach to achieving unidirectional UE is the so called “download–decrypt–re-encrypt–upload” approach, where the client downloads the encrypted data (e.g. from the server storing the encrypted data), locally decrypts it, re-encrypts it using the new key, and re-uploads the newly encrypted data to the server. Our UE constructions are non-trivial in the sense that we achieve significantly better properties as compared to this naïve approach. In particular,

for applications of UE (e.g. key rotation) where the client outsources encrypted data to the server, this entails constant computational/ communication/storage overheads at the client during key rotation (the client simply generates and sends the update token to the server); the corresponding client-overheads are linear (in database size) in the naïve solution.

1.4 Paper Outline

The rest of the paper is organized as follows. Section 2 formally defines a KPHE scheme and its associated security properties. Section 3 presents our constructions of IND-CPA and IND-UPD secure UE from any KPHE scheme (parts of the corresponding proofs of security are detailed in Appendices A and B). Section 4 presents our construction of IND-HRA secure PRE from any KPHE scheme (the corresponding proof of security is detailed in Appendix C). Appendix D presents our final construction of IND-UE secure UE, along with the security proof. Finally, Appendix E describes how to instantiate KHPE from DDH or LWE (here, we mostly rely on known results from the literature).

For readers not familiar with the formal definitions of UE and PRE, we present relatively self-contained background material on UE and PRE in Sections 3.1 and 4.1, respectively.

1.5 Notations

We summarize here the notations used in the rest of the paper. We write $x \stackrel{\$}{\leftarrow} \mathcal{X}$ to represent that an element x is sampled randomly from a set/distribution \mathcal{X} . The output x of a deterministic (resp. randomized) algorithm \mathcal{A} is denoted by $x = \mathcal{A}$ (resp. $x \stackrel{\$}{\leftarrow} \mathcal{A}$). For $a \in \mathbb{N}$ such that $a \geq 1$, we denote by $[a]$ the set of integers lying between 1 and a (both inclusive). We refer to $\lambda \in \mathbb{N}$ as the security parameter, and denote by $\text{poly}(\lambda)$ and $\text{negl}(\lambda)$ any generic (unspecified) polynomial function and negligible function in λ , respectively ¹.

2 Key and Plaintext Homomorphic Encryption

In this section, we present the definitions for the core building block for our constructions, namely key and plaintext homomorphic encryption (KPHE). Informally, a KPHE scheme has the following features:

- **Keys and Plaintexts:** Each secret key sk is an ℓ -bit string for some $\ell = \text{poly}(\lambda)$ (λ being the security parameter). Additionally, each plaintext message \mathbf{m} is an ℓ' -bit string for some $\ell' = \text{poly}(\lambda)$.
- **Key Distribution:** Each secret key is sampled according to some distribution \mathcal{D} over $\{0, 1\}^\ell$. In particular, for our applications, we assume KPHE schemes where each secret key sk is a $2n$ -bit string with equally many 0 and 1 entries.
- **Key Homomorphism:** Let T be any linear transformation that maps ℓ -bit strings to ℓ -bit strings. Then, it is possible to efficiently evaluate the following:

¹Note that a function $f : \mathbb{N} \rightarrow \mathbb{N}$ is said to be negligible in λ if for every positive polynomial p , $f(\lambda) < 1/p(\lambda)$ when λ is sufficiently large.

- Given a public key \mathbf{pk} corresponding to some secret key $\mathbf{sk} \in \{0, 1\}^\ell$, it is possible to efficiently compute a valid public key \mathbf{pk}' corresponding to the transformed secret key $\mathbf{sk}' = T(\mathbf{sk})$, without the knowledge of \mathbf{sk} .
 - Given a ciphertext \mathbf{ct} encrypting a message \mathbf{m} under some secret key $\mathbf{sk} \in \{0, 1\}^\ell$, it is possible to efficiently compute a ciphertext \mathbf{ct}' encrypting the same message \mathbf{m} under the transformed secret key $\mathbf{sk}' = T(\mathbf{sk})$, without the knowledge of \mathbf{sk} .
- **Plaintext Homomorphism:** Let T' be any linear transformation that maps ℓ' -bit strings to ℓ' -bit strings. Then, given a ciphertext \mathbf{ct} encrypting a message $\mathbf{m} \in \{0, 1\}^{\ell'}$ under some secret key \mathbf{sk} , it is possible to efficiently compute a ciphertext \mathbf{ct}'' encrypting the transformed message $\mathbf{m}' = T'(\mathbf{m})$ under the same secret key \mathbf{sk} .

We now summarize these features of KPHE formally below.

Definition 2.1 (KPHE). *A KPHE scheme is a tuple of PPT algorithms of the form $\text{KPHE} = (\text{Setup}, \text{SKGen}, \text{PKGen}, \text{Enc}, \text{Dec}, \text{Eval})$ that are defined as follows:*

- $\mathbf{pp} \stackrel{\$}{\leftarrow} \text{Setup}(1^\lambda)$: *On input the security parameter λ , the setup algorithm outputs a public parameter \mathbf{pp} .*
- $\mathbf{sk} \stackrel{\$}{\leftarrow} \text{SKGen}(\mathbf{pp}, \mathcal{D})$: *On input the public parameter \mathbf{pp} and a distribution \mathcal{D} over $\{0, 1\}^\ell$ (for $\ell = \text{poly}(\lambda)$), the secret key generation algorithm outputs a secret key $\mathbf{sk} \stackrel{\$}{\leftarrow} \mathcal{D}$.*
- $\mathbf{pk} \stackrel{\$}{\leftarrow} \text{PKGen}(\mathbf{pp}, \mathbf{sk})$: *On input the public parameter \mathbf{pp} and a secret key $\mathbf{sk} \in \{0, 1\}^\ell$, the public key generation algorithm outputs a public key \mathbf{pk} .*
- $\mathbf{ct} \stackrel{\$}{\leftarrow} \text{Enc}(\mathbf{pk}, \mathbf{m})$: *On input a public key \mathbf{pk} and a message $\mathbf{m} \in \{0, 1\}^{\ell'}$ (for $\ell' = \text{poly}(\lambda)$), the encryption algorithm outputs a ciphertext \mathbf{ct} .*
- $\mathbf{m}/\perp \stackrel{\$}{\leftarrow} \text{Dec}(\mathbf{sk}, \mathbf{ct})$: *On input a secret key $\mathbf{sk} \in \{0, 1\}^\ell$ and a ciphertext \mathbf{ct} , the decryption algorithm outputs a plaintext message string \mathbf{m} or an error symbol \perp .*
- $(\mathbf{pk}', \mathbf{ct}') \stackrel{\$}{\leftarrow} \text{Eval}(\mathbf{pk}, \mathbf{ct}, T, T')$: *On input a public key \mathbf{pk} , a ciphertext \mathbf{ct} , and a pair of (linear) transformations $T : \{0, 1\}^\ell \rightarrow \{0, 1\}^{\ell'}$ and $T' : \{0, 1\}^{\ell'} \rightarrow \{0, 1\}^{\ell'}$, the homomorphic evaluation algorithm outputs a tuple consisting of a transformed public key and a transformed ciphertext $(\mathbf{pk}', \mathbf{ct}')$.*

Correctness. A KPHE scheme $(\text{Setup}, \text{SKGen}, \text{PKGen}, \text{Enc}, \text{Dec}, \text{Eval})$ is said to be correct with respect to a distribution \mathcal{D} over $\{0, 1\}^\ell$ if for any $\mathbf{pp} \stackrel{\$}{\leftarrow} \text{Setup}(1^\lambda)$, any $\mathbf{sk} \stackrel{\$}{\leftarrow} \text{SKGen}(\mathbf{pp}, \mathcal{D})$, any $\mathbf{pk} \stackrel{\$}{\leftarrow} \text{PKGen}(\mathbf{pp}, \mathbf{sk})$, any $\mathbf{m} \in \{0, 1\}^{\ell'}$, and any pair of (linear) transformations $T : \{0, 1\}^\ell \rightarrow \{0, 1\}^{\ell'}$ and $T' : \{0, 1\}^{\ell'} \rightarrow \{0, 1\}^{\ell'}$, letting $\mathbf{sk}' = T(\mathbf{sk})$, $\mathbf{m}' = T'(\mathbf{m})$ and

$$\mathbf{ct} \stackrel{\$}{\leftarrow} \text{Enc}(\mathbf{pk}, \mathbf{m}), \quad (\mathbf{pk}', \mathbf{ct}') \stackrel{\$}{\leftarrow} \text{Eval}(\mathbf{pk}, \mathbf{ct}, T, T'),$$

both of the following hold with overwhelmingly large probability:

Experiment $\text{Expt}_{\mathcal{D}}^{\text{DSS-KPHE}}(\lambda, \mathcal{A})$:

1. The challenger generates $\text{pp} \xleftarrow{\$} \text{Setup}(1^\lambda)$, $\text{sk} \xleftarrow{\$} \text{SKGen}(\text{pp}, \mathcal{D})$, and $\text{pk} \xleftarrow{\$} \text{PKGen}(\text{pp}, \text{sk})$, and provides the adversary \mathcal{A} with (pp, pk) .
2. The adversary \mathcal{A} issues a challenge encryption query for a pair of messages (m_0, m_1) .
3. The challenger samples $b \xleftarrow{\$} \{0, 1\}$, creates the challenge ciphertext

$$\text{ct}^* \xleftarrow{\$} \text{Enc}(\text{pk}, m_b),$$

and sends ct^* to the adversary \mathcal{A} .

4. The adversary \mathcal{A} outputs a bit $b' \in \{0, 1\}$.
5. Output 1 if $b = b'$ and 0 otherwise.

Figure 1: The \mathcal{D} -Semantic Security Experiment for KPHE

- pk' is a valid public key with respect to $\text{sk}' = T(\text{sk})$, i.e., for any $\bar{m} \in \{0, 1\}^\ell$, it holds that $\text{Dec}(\text{sk}', \text{Enc}(\text{pk}', \bar{m})) = \bar{m}$.
- ct' is a valid encryption of m' under (pk', sk') , i.e., $\text{Dec}(\text{sk}', \text{ct}') = m'$.

Distributional Semantic Security. We (informally) say that a KPHE satisfies distributional semantic security with respect to some distribution \mathcal{D} over $\{0, 1\}^\ell$ if it remains semantically secure even when the secret key sk is sampled according to the distribution \mathcal{D} . Formally, this is modeled using a semantic security game where the secret key is sampled by the challenger as per the distribution \mathcal{D} .

Definition 2.2 (*\mathcal{D} -Semantic Security*). *A KPHE scheme with ℓ -bit secret keys is said to be \mathcal{D} -semantically secure with respect to a distribution \mathcal{D} over $\{0, 1\}^\ell$ if for any security parameter $\lambda \in \mathbb{N}$ and any PPT adversary \mathcal{A} , the following holds with overwhelmingly large probability:*

$$|\Pr[\text{Expt}_{\mathcal{D}}^{\text{DSS-KPHE}}(\lambda, \mathcal{A}) = 1] - 1/2| < \text{negl}(\lambda),$$

where the experiment $\text{Expt}_{\mathcal{D}}^{\text{DSS-KPHE}}(\lambda, \mathcal{A})$ is as defined in Figure 1.

Distributional Circular Security. We (informally) say that a KPHE satisfies distributional circular security with respect to some distribution \mathcal{D} over $\{0, 1\}^\ell$ if it satisfies the standard notion of circular security [CL01, BRS02, BHHO08, ACPS09] even when each secret key is sampled from the distribution \mathcal{D} . Formally, this is modeled using a circular security game where the secret keys are sampled by the challenger as per the distribution \mathcal{D} .

Definition 2.3 (*\mathcal{D} -Circular Security*). *A KPHE scheme with ℓ -bit secret keys and ℓ -bit messages is said to be \mathcal{D} -circular secure with respect to a distribution \mathcal{D} over $\{0, 1\}^\ell$ if for any security parameter $\lambda \in \mathbb{N}$ and any PPT adversary \mathcal{A} , the following holds with overwhelmingly large probability:*

$$|\Pr[\text{Expt}_{\mathcal{D}}^{\text{DCC-KPHE}}(\lambda, \mathcal{A}) = 1] - 1/2| < \text{negl}(\lambda),$$

Experiment $\text{Expt}_{\mathcal{D}}^{\text{DCC-KPHE}}(\lambda, \mathcal{A})$:

1. The challenger generates $\text{pp} \xleftarrow{\$} \text{Setup}(1^\lambda)$ and provides it to the adversary.
2. The adversary \mathcal{A} outputs $n = \text{poly}(\lambda)$.
3. The challenger samples $\text{sk}_1, \dots, \text{sk}_n \xleftarrow{\$} \text{SKGen}(\text{pp}, \mathcal{D})$, sets

$$\text{pk}_1 \xleftarrow{\$} \text{PKGen}(\text{pp}, \text{sk}_1), \dots, \text{pk}_n \xleftarrow{\$} \text{PKGen}(\text{pp}, \text{sk}_n),$$

and provides $(\text{pk}_1, \dots, \text{pk}_n)$ to the adversary \mathcal{A} .

4. The challenger also sets the following for each $i, j \in [n]$:

$$\text{ct}_{i,j,0} \xleftarrow{\$} \text{Enc}(\text{pk}_i, \text{sk}_j), \quad \text{ct}_{i,j,1} \xleftarrow{\$} \text{Enc}(\text{pk}_i, 0^{|\text{sk}_j|}).$$

5. The challenger finally samples a bit $b \xleftarrow{\$} \{0, 1\}$ and provides the adversary \mathcal{A} with the ensemble $\{\text{ct}_{i,j,b}\}_{i,j \in [n]}$.
6. The adversary \mathcal{A} outputs a bit $b' \in \{0, 1\}$.
7. Output 1 if $b = b'$ and 0 otherwise.

Figure 2: The \mathcal{D} -Circular Security Experiment for KPHE

where the experiment $\text{Expt}_{\mathcal{D}}^{\text{DCC-KPHE}}(\lambda, \mathcal{A})$ is as defined in Figure 2.

Blinding. We (informally) say that a KPHE scheme satisfies public key and ciphertext blinding if the homomorphic evaluation algorithm outputs a public key-ciphertext pair (pk', ct') corresponding to the transformed secret key sk' and the transformed message m' such that:

- The transformed public key pk' is computationally indistinguishable from a public key sampled uniformly at random from the set of all valid public keys corresponding to the secret key sk' .
- The transformed ciphertext ct' is computationally indistinguishable from a ciphertext sampled uniformly at random from the set of all valid ciphertexts corresponding to the transformed message m' under pk' .

More formally, we define this blinding property as follows.

Definition 2.4 (Blinding). *A KPHE scheme with ℓ -bit secret keys and ℓ' -bit messages is said to satisfy blinding security with respect to a distribution \mathcal{D} over $\{0, 1\}^\ell$ if for any security parameter $\lambda \in \mathbb{N}$ and any PPT adversary \mathcal{A} , the following holds with overwhelmingly large probability:*

$$|\Pr[\text{Expt}_{\mathcal{D}}^{\text{Blind-KPHE}}(\lambda, \mathcal{A}) = 1] - 1/2| < \text{negl}(\lambda),$$

where the experiment $\text{Expt}_{\mathcal{D}}^{\text{Blind-KPHE}}(\lambda, \mathcal{A})$ is as defined in Figure 3.

Experiment $\text{Expt}_{\mathcal{D}}^{\text{Blind-KPHE}}(\lambda, \mathcal{A})$:

1. The challenger generates $\text{pp} \xleftarrow{\$} \text{Setup}(1^\lambda)$, $\text{sk} \xleftarrow{\$} \mathcal{D}$, and $\text{pk} \xleftarrow{\$} \text{PKGen}(\text{pp}, \text{sk})$, and provides the adversary \mathcal{A} with $(\text{pp}, \text{sk}, \text{pk})$.
2. The adversary \mathcal{A} sends a message $\mathbf{m} \in \{0, 1\}^{\ell'}$ to the challenger.
3. The challenger responds to \mathcal{A} with a ciphertext $\text{ct} \xleftarrow{\$} \text{Enc}(\text{pk}, \mathbf{m})$.
4. The adversary \mathcal{A} then sends a pair of (linear) transformations

$$T : \{0, 1\}^\ell \rightarrow \{0, 1\}^\ell, \quad T' : \{0, 1\}^{\ell'} \rightarrow \{0, 1\}^{\ell'}.$$

5. The challenger sets

$$\text{sk}' = T(\text{sk}), \quad \mathbf{m}' = T'(\mathbf{m}),$$

and computes the following:

$$(\text{pk}_0, \text{ct}_0) \xleftarrow{\$} \text{Eval}(\text{pk}, \text{ct}, T, T'), \quad \text{pk}_1 \xleftarrow{\$} \text{PKGen}(\text{pp}, \text{sk}'), \quad \text{ct}_1 \xleftarrow{\$} \text{Enc}(\text{pk}_1, \mathbf{m}'),$$

6. The challenger finally samples a bit $b \xleftarrow{\$} \{0, 1\}$ and provides the adversary \mathcal{A} with $(\text{pk}_b, \text{ct}_b)$.
7. The adversary \mathcal{A} outputs a bit $b' \in \{0, 1\}$.
8. Output 1 if $b = b'$ and 0 otherwise.

Figure 3: The Blinding Experiment for KPHE

KPHE from Concrete Assumptions. Looking ahead, in Appendix E, we show the following: assuming that either the decisional Diffie-Hellman (DDH) assumption holds *or* that the learning with errors (LWE) assumption holds, there exists a KPHE scheme that satisfies essentially all of the aforementioned properties. In particular, we (informally) have the following:

Theorem 2.5 (Informal). *Assuming DDH or LWE, there exists a KPHE scheme with $2n$ -bit secret keys that satisfies distributional semantic security with respect to the distribution \mathcal{U}_n , distributional circular security with respect to the distribution \mathcal{U}_n , and blinding, as defined above.*

In particular, we rely on known results from [BH08, NS12, GHV10] for the DDH-based instantiation of KPHE, and on known results from [Reg09, ACPS09, GKPV10] for proving certain properties of the LWE-based instantiation of KPHE.

3 Unidirectional UE from KPHE

In this section, we show how to construct unidirectional UE satisfying various security notions (IND-ENC, IND-UPD and IND-UE) from any KPHE scheme.

3.1 Definition

Definition 3.1. An updatable encryption (UE) scheme for message space \mathcal{M} is a tuple of PPT algorithms $\text{UE} = (\text{UE.setup}, \text{UE.next}, \text{UE.enc}, \text{UE.upd}, \text{UE.dec})$ with the following syntax:

- $\mathbf{k}_0 \xleftarrow{\$} \text{UE.setup}(1^\lambda)$: On input a security parameter 1^λ , it returns a secret key \mathbf{k}_e for epoch $e = 0$.
- $(\mathbf{k}_{e+1}, \Delta_{e+1}) \xleftarrow{\$} \text{UE.next}(\mathbf{k}_e)$: On input a secret key \mathbf{k}_e for epoch e , it outputs a new secret key \mathbf{k}_{e+1} and an update token Δ_{e+1} for epoch $e + 1$.
- $\text{ct}_e \xleftarrow{\$} \text{UE.enc}(\mathbf{k}_e, m)$: On input a secret key \mathbf{k}_e for epoch e and a message $m \in \mathcal{M}$, it outputs a ciphertext ct_e .
- $\text{ct}_{e+1} \xleftarrow{\$} \text{UE.upd}(\Delta_{e+1}, \text{ct}_e)$: On input a ciphertext ct_e from epoch e and the update token Δ_{e+1} , it returns the updated ciphertext ct_{e+1} .
- $m'/\perp \leftarrow \text{UE.dec}(\mathbf{k}_e, \text{ct}_e)$: On input a ciphertext ct_e and a secret key \mathbf{k}_e of some epoch e , it returns a message m' or \perp .

We stress that UE.next generates a new key along with an update token, which follows from the definition in the work of Lehmann and Tackmann [LT18]. In our constructions, the update token Δ_{e+1} can also be generated from \mathbf{k}_e and \mathbf{k}_{e+1} .

Definition 3.2 (Correctness). Let $\text{UE} = (\text{UE.setup}, \text{UE.next}, \text{UE.enc}, \text{UE.upd}, \text{UE.dec})$ be an updatable encryption scheme. We say UE is correct if for any $m \in \mathcal{M}$, any $\mathbf{k}_0 \xleftarrow{\$} \text{UE.setup}(1^\lambda)$, any sequence of $(\mathbf{k}_1, \Delta_1), \dots, (\mathbf{k}_e, \Delta_e)$ generated as $(\mathbf{k}_i, \Delta_i) \xleftarrow{\$} \text{UE.next}(\mathbf{k}_{i-1})$ for all $i \in [e]$, and for any $0 \leq \hat{e} \leq e$, let $\text{ct}_{\hat{e}} \xleftarrow{\$} \text{UE.enc}(\mathbf{k}_{\hat{e}}, m)$ and $\text{ct}_j \xleftarrow{\$} \text{UE.upd}(\Delta_j, \text{ct}_{j-1})$ for all $j = \hat{e} + 1, \dots, e$, then $\text{UE.dec}(\mathbf{k}_e, \text{ct}_e) = m$.

Confidentiality. The adversary \mathcal{A} has access to the oracles defined in Figure 4. We follow the bookkeeping techniques of [LT18, KLR19, BDGJ20, Jia20], using the following sets to keep track of the generated and updated ciphertexts, and the epochs in which the adversary corrupted a key or a token, or learned a version of the challenge ciphertext.

- \mathcal{L} : Set of non-challenge ciphertexts (e, ct_e) produced by calls to the $\mathcal{O}.enc$ or $\mathcal{O}.upd$ oracle. $\mathcal{O}.upd$ only updates ciphertexts obtained in \mathcal{L} .
- $\tilde{\mathcal{L}}$: Set of updated versions of the challenge ciphertexts $(e, \tilde{\text{ct}}_e)$. $\tilde{\mathcal{L}}$ is initiated with the challenge ciphertext $(\hat{e}, \tilde{\text{ct}}_{\hat{e}})$. Any call to the $\mathcal{O}.next$ oracle automatically updates the challenge ciphertext to the new epoch, which the adversary can fetch via a call to $\mathcal{O}.upd\tilde{\mathcal{C}}$.
- \mathcal{K} : Set of epochs e in which the adversary corrupted the secret key \mathbf{k}_e (from $\mathcal{O}.corr$).
- \mathcal{T} : Set of epochs e in which the adversary corrupted the update token Δ_e (from $\mathcal{O}.corr$).
- \mathcal{C} : Set of epochs e in which the adversary learned a version of the challenge ciphertext (from $\mathcal{O}.chall$ or $\mathcal{O}.upd\tilde{\mathcal{C}}$).

<p><u>Setup(1^λ):</u></p> $k_0 \xleftarrow{\$} \text{UE.setup}(1^\lambda)$ $e := 0; \text{phase} := 0$ $\mathcal{L}, \tilde{\mathcal{L}}, \mathcal{K}, \mathcal{T}, \mathcal{C} \xleftarrow{\$} \emptyset$ <p><u>$\mathcal{O}.\text{enc}(m)$:</u></p> $\text{ct} \xleftarrow{\$} \text{UE.enc}(k_e, m)$ $\mathcal{L} := \mathcal{L} \cup \{(e, \text{ct})\}$ <p>return ct</p> <p><u>$\mathcal{O}.\text{next}$:</u></p> $e := e + 1$ $(k_e, \Delta_e) \xleftarrow{\$} \text{UE.next}(k_{e-1})$ <p>if phase = 1 then</p> $\tilde{\text{ct}}_e \xleftarrow{\$} \text{UE.upd}(\Delta_e, \tilde{\text{ct}}_{e-1})$ $\tilde{\mathcal{L}} := \tilde{\mathcal{L}} \cup \{(e, \tilde{\text{ct}}_e)\}$ <p><u>$\mathcal{O}.\text{upd}(\text{ct}_{e-1})$:</u></p> <p>if $(e - 1, \text{ct}_{e-1}) \notin \mathcal{L}$ then</p> <p>return \perp</p> $\text{ct}_e \xleftarrow{\$} \text{UE.upd}(\Delta_e, \text{ct}_{e-1})$ $\mathcal{L} := \mathcal{L} \cup \{(e, \text{ct}_e)\}$ <p>return ct_e</p> <p><u>$\mathcal{O}.\text{corr}(\text{inp}, \hat{e})$:</u></p> <p>if $\hat{e} > e$ then</p> <p>return \perp</p> <p>if inp = key then</p> $\mathcal{K} := \mathcal{K} \cup \{\hat{e}\}$ <p>return $k_{\hat{e}}$</p> <p>if inp = token then</p> $\mathcal{T} := \mathcal{T} \cup \{\hat{e}\}$ <p>return $\Delta_{\hat{e}}$</p>	<p><u>$\mathcal{O}.\text{chall-IND-ENC}(\bar{m}_0, \bar{m}_1)$:</u></p> <p>if $\bar{m}_0 \neq \bar{m}_1$ then</p> <p>return \perp</p> <p>phase := 1; $\tilde{e} := e$</p> $\tilde{\text{ct}}_{\tilde{e}} \xleftarrow{\$} \text{UE.enc}(k_{\tilde{e}}, \bar{m}_b)$ $\mathcal{C} := \mathcal{C} \cup \{\tilde{e}\}$ $\tilde{\mathcal{L}} := \tilde{\mathcal{L}} \cup \{(\tilde{e}, \tilde{\text{ct}}_{\tilde{e}})\}$ <p>return $\tilde{\text{ct}}_{\tilde{e}}$</p> <p><u>$\mathcal{O}.\text{chall-IND-UPD}(\bar{\text{ct}}_0, \bar{\text{ct}}_1)$:</u></p> <p>if $(e - 1, \bar{\text{ct}}_0) \notin \mathcal{L}$ or $(e - 1, \bar{\text{ct}}_1) \notin \mathcal{L}$ or $\bar{\text{ct}}_0 \neq \bar{\text{ct}}_1$ then</p> <p>return \perp</p> <p>phase := 1; $\tilde{e} := e$</p> $\tilde{\text{ct}}_{\tilde{e}} \xleftarrow{\$} \text{UE.upd}(\Delta_{\tilde{e}}, \bar{\text{ct}}_b)$ $\mathcal{C} := \mathcal{C} \cup \{\tilde{e}\}$ $\tilde{\mathcal{L}} := \tilde{\mathcal{L}} \cup \{(\tilde{e}, \tilde{\text{ct}}_{\tilde{e}})\}$ <p>return $\tilde{\text{ct}}_{\tilde{e}}$</p> <p><u>$\mathcal{O}.\text{chall-IND-UE}(\bar{m}, \bar{\text{ct}})$:</u></p> <p>if $(e - 1, \bar{\text{ct}}) \notin \mathcal{L}$ then</p> <p>return \perp</p> <p>phase := 1; $\tilde{e} := e$</p> <p>if b = 0 then</p> $\tilde{\text{ct}}_{\tilde{e}} \xleftarrow{\$} \text{UE.enc}(k_{\tilde{e}}, \bar{m})$ <p>else</p> $\tilde{\text{ct}}_{\tilde{e}} \xleftarrow{\$} \text{UE.upd}(\Delta_{\tilde{e}}, \bar{\text{ct}})$ $\mathcal{C} := \mathcal{C} \cup \{\tilde{e}\}$ $\tilde{\mathcal{L}} := \tilde{\mathcal{L}} \cup \{(\tilde{e}, \tilde{\text{ct}}_{\tilde{e}})\}$ <p>return $\tilde{\text{ct}}_{\tilde{e}}$</p> <p><u>$\mathcal{O}.\text{upd}\tilde{\mathcal{C}}$:</u></p> <p>if phase = 0 then</p> <p>return \perp</p> $\mathcal{C} := \mathcal{C} \cup \{e\}$ <p>return $\tilde{\text{ct}}_e$</p>
---	--

Figure 4: Oracles in security games for updatable encryption.

We further define the epoch identification sets $\mathcal{C}^*, \mathcal{K}^*, \mathcal{T}^*$ as the extended sets of $\mathcal{C}, \mathcal{K}, \mathcal{T}$ in which the adversary learned or inferred information. We focus on *no-directional* key updates and *uni-directional* ciphertext updates.

$$\begin{aligned} \mathcal{K}^* &:= \mathcal{K} \\ \mathcal{T}^* &:= \{e \in \{0, \dots, e_{\text{end}}\} \mid (e \in \mathcal{T}) \vee (e - 1 \in \mathcal{K}^* \wedge e \in \mathcal{K}^*)\} \\ \mathcal{C}^* &:= \{e \in \{0, \dots, e_{\text{end}}\} \mid \text{ChallEq}(e) = \text{true}\} \\ &\text{where } \text{true} \leftarrow \text{ChallEq}(e) \iff (e \in \mathcal{C}) \vee (\text{ChallEq}(e - 1) \wedge e \in \mathcal{T}^*) \end{aligned}$$

Remark 3.3. *The constructions we present later will in fact permit backward-leak key updates. At first glance the backward-leak key updates notion proposed by Nishimaki [Nis21] is seemingly weaker than no-directionality key updates. However, as mentioned in the introduction, this notion is*

essentially equivalent to no-directional key updates because backward-leak derivation of \mathbf{k}_{e-1} does not increase the adversary's advantage in breaking the scheme. In particular, if the adversary obtains a challenge ciphertext $\tilde{\mathbf{ct}}_{e-1}$ and corrupts Δ_e and \mathbf{k}_e , then it does not matter if the adversary can derive \mathbf{k}_{e-1} or not, as it can always update the ciphertext to $\tilde{\mathbf{ct}}_e$ and decrypt it using \mathbf{k}_e .

Definition 3.4 (IND-ENC, IND-UPD, IND-UE security). *Let $\text{UE} = (\text{UE.setup}, \text{UE.next}, \text{UE.enc}, \text{UE.upd}, \text{UE.dec})$ be an updatable encryption scheme. We say UE is notion-secure for notion $\in \{\text{IND-ENC}, \text{IND-UPD}, \text{IND-UE}\}$ if for all PPT adversary \mathcal{A} it holds that*

$$\left| \Pr \left[\text{Exp}_{\mathcal{A}, \text{UE}}^{\text{notion}}(1^\lambda) = 1 \right] - \frac{1}{2} \right| \leq \text{negl}(\lambda)$$

for some negligible function $\text{negl}(\cdot)$.

Experiment $\text{Exp}_{\mathcal{A}, \text{UE}}^{\text{notion}}(1^\lambda)$:

Run $\text{Setup}(1^\lambda)$
 $(\text{state}, \text{Chall}_0, \text{Chall}_1) \xleftarrow{\$} \mathcal{A}^{\mathcal{O}.enc, \mathcal{O}.next, \mathcal{O}.upd, \mathcal{O}.corr}(1^\lambda)$
 $\mathbf{b} \xleftarrow{\$} \{0, 1\}$
 $\tilde{\mathbf{ct}} \xleftarrow{\$} \mathcal{O}.chall\text{-notion}(\text{Chall}_0, \text{Chall}_1)$
 Proceed only if $\tilde{\mathbf{ct}} \neq \perp$
 $\mathbf{b}' \xleftarrow{\$} \mathcal{A}^{\mathcal{O}.enc, \mathcal{O}.next, \mathcal{O}.upd, \mathcal{O}.corr, \mathcal{O}.upd\tilde{\mathbf{C}}}(\text{state}, \tilde{\mathbf{ct}})$
return 1 if $\mathbf{b} = \mathbf{b}'$ and $\mathcal{C}^* \cap \mathcal{K}^* = \emptyset$

3.2 IND-ENC Secure Unidirectional UE

We begin by showing that any KPHE scheme with $2n$ -bit secret keys that satisfies distributional semantic security with respect to the distribution \mathcal{U}_n , as well as public key and ciphertext blinding as described in Section 2 implies an IND-ENC secure unidirectional UE scheme.

Construction. Given a KHPE scheme of the form

$$\text{KPHE} = (\text{KPHE.Setup}, \text{KPHE.SKGen}, \text{KPHE.PKGen}, \text{KPHE.Enc}, \text{KPHE.Dec}, \text{KPHE.Eval}),$$

with $2n$ -bit secret keys, we construct a unidirectional UE scheme

$$\text{UE} = (\text{UE.setup}, \text{UE.next}, \text{UE.enc}, \text{UE.upd}, \text{UE.dec}),$$

with message space $\mathcal{M} = \{0, 1\}^{2n}$ as follows:

- $\text{UE.setup}(1^\lambda)$: Generate $\text{pp} \xleftarrow{\$} \text{KPHE.Setup}(1^\lambda)$, $\text{sk}_0 \xleftarrow{\$} \text{KPHE.SKGen}(\text{pp}, \mathcal{U}_n)$, and output

$$\mathbf{k}_0 = (\text{pp}, \text{sk}_0).$$

- $\text{UE.next}(\mathbf{k}_e)$: Parse $\mathbf{k}_e = (\text{pp}, \text{sk}_e)$. Generate $\text{sk}_{e+1} \xleftarrow{\$} \text{KPHE.SKGen}(\text{pp}, \mathcal{U}_n)$ and $\text{pk}_{e+1} \xleftarrow{\$} \text{KPHE.PKGen}(\text{pp}, \text{sk}_{e+1})$. Output

$$\mathbf{k}_{e+1} = (\text{pp}, \text{sk}_{e+1}), \quad \Delta_{e+1} = (\text{pk}_{e+1}, \text{KPHE.Enc}(\text{pk}_{e+1}, \text{sk}_e)).$$

- $\text{UE.enc}(\mathbf{k}_e, m)$: Parse $\mathbf{k}_e = (\text{pp}, \text{sk}_e)$. Generate $\text{pk}_e \stackrel{\$}{\leftarrow} \text{KPHE.PKGen}(\text{pp}, \text{sk}_e)$ and compute $\text{ctx}_e \stackrel{\$}{\leftarrow} \text{KPHE.Enc}(\text{pk}_e, m)$. Output

$$\text{ct}_e = (0, (\text{pk}_e, \text{ctx}_e)).$$

- $\text{UE.upd}(\Delta_{e+1}, \text{ct}_e)$: Parse the update token and the ciphertext as

$$\Delta_{e+1} = (\text{pk}_\Delta, \text{ctx}_\Delta), \quad \text{ct}_e = (t, (\overline{\text{pk}}_{e-t}, \overline{\text{ctx}}_{e-t}), \dots, (\overline{\text{pk}}_{e-1}, \overline{\text{ctx}}_{e-1}), (\text{pk}_e, \text{ctx}_e))$$

Sample a uniform random permutation $\pi : [2n] \rightarrow [2n]$. Also, let $\pi_{\text{id}} : [2n] \rightarrow [2n]$ denote the *identity* permutation. Compute

$$(\overline{\text{pk}}_e, \overline{\text{ctx}}_e) \stackrel{\$}{\leftarrow} \text{KPHE.Eval}(\text{pk}_e, \text{ctx}_e, \pi, \pi_{\text{id}}), \quad (\text{pk}_{e+1}, \text{ctx}_{e+1}) \stackrel{\$}{\leftarrow} \text{KPHE.Eval}(\text{pk}_\Delta, \text{ctx}_\Delta, \pi_{\text{id}}, \pi).$$

and output the updated ciphertext as:

$$\text{ct}_{e+1} = (t + 1, (\overline{\text{pk}}_{e-t}, \overline{\text{ctx}}_{e-t}), \dots, (\overline{\text{pk}}_e, \overline{\text{ctx}}_e), (\text{pk}_{e+1}, \text{ctx}_{e+1})).$$

- $\text{UE.dec}(\mathbf{k}_e, \text{ct}_e)$: Parse $\mathbf{k}_e = (\text{pp}, \text{sk}_e)$ and the ciphertext as

$$\text{ct}_e = (t, (\overline{\text{pk}}_{e-t}, \overline{\text{ctx}}_{e-t}), \dots, (\overline{\text{pk}}_{e-1}, \overline{\text{ctx}}_{e-1}), (\text{pk}_e, \text{ctx}_e)).$$

If $t = 0$, then output $m \leftarrow \text{KPHE.Dec}(\text{sk}_e, \text{ctx}_e)$.

Otherwise, compute $\overline{\text{sk}}_{e-1} \leftarrow \text{KPHE.Dec}(\text{sk}_e, \text{ctx}_e)$. Then for each j from $(e - 1)$ down to $(e - t + 1)$, compute

$$\overline{\text{sk}}_{j-1} \leftarrow \text{KPHE.Dec}(\overline{\text{sk}}_j, \overline{\text{ctx}}_j).$$

Finally, output the message $m \leftarrow \text{KPHE.Dec}(\overline{\text{sk}}_{e-t}, \overline{\text{ctx}}_{e-t})$.

Correctness. We first prove the correctness of the UE scheme. For any $m \in \mathcal{M}$, any $\mathbf{k}_0 \leftarrow \text{UE.setup}(1^\lambda)$, any sequence of $(\mathbf{k}_1, \Delta_1), \dots, (\mathbf{k}_e, \Delta_e)$ generated as $(\mathbf{k}_i, \Delta_i) \leftarrow \text{UE.next}(\mathbf{k}_{i-1})$ for all $i \in [e]$, let $\text{ct}_0 \leftarrow \text{UE.enc}(\mathbf{k}_0, m)$ and $\text{ct}_i \leftarrow \text{UE.upd}(\Delta_i, \text{ct}_{i-1})$ for all $j \in [e]$, then the final ciphertext is of the form $\text{ct}_e = (e, (\overline{\text{pk}}_0, \overline{\text{ctx}}_0), \dots, (\overline{\text{pk}}_{e-1}, \overline{\text{ctx}}_{e-1}), (\text{pk}_e, \text{ctx}_e))$. All the secret keys are of the form $\mathbf{k}_0 = (\text{pp}, \text{sk}_0), \dots, \mathbf{k}_e = (\text{pp}, \text{sk}_e)$. Let π_j be the random permutation sampled in $\text{UE.upd}(\Delta_{j+1}, \text{ct}_j)$ and let $\overline{\text{sk}}_j = \pi_j(\text{sk}_j)$ for all $j = 0, 1, \dots, e - 1$. We can prove by induction that $\text{KPHE.Dec}(\overline{\text{sk}}_0, \overline{\text{ctx}}_0) = m$, $\text{KPHE.Dec}(\overline{\text{sk}}_1, \overline{\text{ctx}}_1) = \overline{\text{sk}}_0, \dots, \text{KPHE.Dec}(\overline{\text{sk}}_{e-1}, \overline{\text{ctx}}_{e-1}) = \overline{\text{sk}}_{e-2}$, $\text{KPHE.Dec}(\text{sk}_e, \text{ctx}_e) = \overline{\text{sk}}_{e-1}$. Therefore, $\text{UE.dec}(\mathbf{k}_e, \text{ct}_e)$ outputs m . This argument is for any ciphertext starting from epoch 0. The same argument holds for any ciphertext starting from any epoch \hat{e} where $0 \leq \hat{e} \leq e$.

Confidentiality. Next we prove the IND-ENC security of our UE scheme. More formally, we state and prove the following theorem:

Theorem 3.5 (IND-ENC Security). *Assuming that KPHE is a KPHE scheme with $2n$ -bit secret keys that satisfies distributional security with respect to the distribution \mathcal{U}_n , as well as public key and ciphertext blinding as described in Section 2, the above UE construction is an IND-ENC secure unidirectional UE scheme.*

Proof. The proof proceeds via a hybrid argument.

Hyb₀ The challenger plays the real game with the adversary.

Hyb₁ Same as Hyb₀ but for $\text{UE.upd}(\Delta_e, \text{ct}_{e-1})$ in $\mathcal{O}.\text{upd}$ and $\text{UE.upd}(\Delta_e, \tilde{\text{ct}}_{e-1})$ in $\mathcal{O}.\text{next}$, do the following:

- Let $\mathbf{k}_{e-1} = (\text{pp}, \text{sk}_{e-1})$ and $\mathbf{k}_e = (\text{pp}, \text{sk}_e)$.
- Parse the ciphertext ct_{e-1} or $\tilde{\text{ct}}_{e-1}$ as

$$(t, (\overline{\text{pk}}_{e-1-t}, \overline{\text{ctx}}_{e-1-t}), \dots, (\overline{\text{pk}}_{e-2}, \overline{\text{ctx}}_{e-2}), (\text{pk}_{e-1}, \text{ctx}_{e-1})),$$

where $\text{ctx}_{e-1} = \text{KPHE.Enc}(\text{pk}_{e-1}, x)$. Note that if $t = 0$, then $x = \overline{\mathbf{m}}$ for some message, otherwise $x = \overline{\text{sk}}_{e-2}$ that is the KPHE secret key corresponding to $\overline{\text{pk}}_{e-2}$.

- Sample a uniform random permutation $\pi : [2n] \rightarrow [2n]$, let $\overline{\text{sk}}_{e-1} = \pi(\text{sk}_{e-1})$, and sample $\overline{\text{pk}}_{e-1} \stackrel{\$}{\leftarrow} \text{KPHE.PKGen}(\text{pp}, \overline{\text{sk}}_{e-1})$. Compute $\overline{\text{ctx}}_{e-1} \stackrel{\$}{\leftarrow} \text{KPHE.Enc}(\overline{\text{pk}}_{e-1}, x)$.
- Sample $\text{pk}_e \stackrel{\$}{\leftarrow} \text{KPHE.PKGen}(\text{pp}, \text{sk}_e)$ and compute $\text{ctx}_e \stackrel{\$}{\leftarrow} \text{KPHE.Enc}(\text{pk}_e, \overline{\text{sk}}_{e-1})$.
- Let ct_e or $\tilde{\text{ct}}_e$ be

$$(t+1, (\overline{\text{pk}}_{e-1-t}, \overline{\text{ctx}}_{e-1-t}), \dots, (\overline{\text{pk}}_{e-1}, \overline{\text{ctx}}_{e-1}), (\text{pk}_e, \text{ctx}_e)).$$

We prove in Lemma 3.6 that this hybrid is computationally indistinguishable from Hyb₀ to any PPT adversary by the blinding property of KPHE.

Hyb₂ Same as Hyb₁ but for $\text{UE.upd}(\Delta_e, \text{ct}_{e-1})$ in $\mathcal{O}.\text{upd}$ and $\text{UE.upd}(\Delta_e, \tilde{\text{ct}}_{e-1})$ in $\mathcal{O}.\text{next}$, instead of letting $\overline{\text{sk}}_{e-1} = \pi(\text{sk}_{e-1})$, sample $\overline{\text{sk}}_{e-1}$ from the distribution \mathcal{U}_n . This hybrid is statistically identical to Hyb₁.

Hyb₃ Let \tilde{e} be the challenge epoch, and let \bar{e} be the last epoch where the adversary corrupts continuous update tokens from \tilde{e} , namely the adversary corrupts $\Delta_{\tilde{e}+1}, \Delta_{\tilde{e}+2}, \dots, \Delta_{\bar{e}}$ but not $\Delta_{\bar{e}+1}$. This hybrid is the same as Hyb₂ except that the challenger guesses \tilde{e}^* and \bar{e}^* at the beginning of the game and aborts the game if guessing incorrectly. Let E be the upper bound on the number of epochs during the game. If the challenger does not abort, then this hybrid is identical to Hyb₂, which happens with probability at least $\frac{1}{E^2}$. In the remaining hybrids, we assume for simplicity that the challenger guesses \tilde{e} and \bar{e} correctly.

Hyb₄ Same as Hyb₃ except that for each $\mathbf{k}_e = (\text{pp}, \text{sk}_e)$, generate a single public key $\widehat{\text{pk}}_e \stackrel{\$}{\leftarrow} \text{KPHE.PKGen}(\text{pp}, \text{sk}_e)$. Then whenever $\text{KPHE.Enc}(\text{pk}_e, x)$ is computed for a freshly generated pk_e and some x , compute it as $\text{KPHE.Eval}(\widehat{\text{pk}}_e, \text{KPHE.Enc}(\widehat{\text{pk}}_e, x), \pi_{\text{id}}, \pi_{\text{id}})$. That is, instead of generating a fresh pk_e from sk_e every time, use the same $\widehat{\text{pk}}_e$ to encrypt x and use then KPHE.Eval to re-randomize it.

This hybrid is computationally indistinguishable from Hyb₃ by the blinding property of KPHE. We omit the detailed reduction here, but it is similar to the reduction in the proof of Lemma 3.6.

Hyb₅ Same as Hyb₄ except that for all $\tilde{e}+1 \leq e \leq \bar{e}$, $\text{UE.next}(\mathbf{k}_{e-1})$ is computed as follows. Generate $\text{sk}_e \stackrel{\$}{\leftarrow} \text{KPHE.SKGen}(\text{pp}, \mathcal{U}_n)$ and let $\widehat{\text{pk}}_e \stackrel{\$}{\leftarrow} \text{KPHE.PKGen}(\text{pp}, \text{sk}_e)$ be the single public key for \mathbf{k}_e (that will be used for every KPHE.Enc). Output

$$\mathbf{k}_e = (\text{pp}, \text{sk}_e), \quad \Delta_e = (\widehat{\text{pk}}_e, \text{KPHE.Enc}(\widehat{\text{pk}}_e, 0^{2n})).$$

We prove in Lemma 3.7 that this hybrid is computationally indistinguishable from Hyb_4 to any PPT adversary based on the distributional semantic security of KPHE.

Hyb_6 Same as Hyb_5 except that for each $\mathbf{k}_e = (\mathbf{pp}, \mathbf{sk}_e)$, generate a single public key $\widehat{\mathbf{pk}}_e \stackrel{\$}{\leftarrow} \text{KPHE.PKGen}(\mathbf{pp}, \mathbf{sk}_e)$ and use $\text{KPHE.Eval}(\widehat{\mathbf{pk}}_e, \text{KPHE.Enc}(\widehat{\mathbf{pk}}_e, \cdot), \pi_{\text{id}}, \pi_{\text{id}})$ for all the computation of $\text{KPHE.Enc}(\mathbf{k}_e, \cdot)$ (including the computation of Δ_e). The only exception is the challenge ciphertext $\widetilde{\text{ct}}_e$, which is computed using $\widehat{\mathbf{pk}}_e$ without re-randomization, namely

$$\widetilde{\text{ct}}_e = (0, (\widehat{\mathbf{pk}}_e, \text{KPHE.Enc}(\widehat{\mathbf{pk}}_e, \overline{\mathbf{m}}_b))).$$

This hybrid is computationally indistinguishable from Hyb_5 by the blinding property of KPHE. We omit the detailed reduction here, but it is similar to the reduction in the proof of Lemma 3.6.

Finally, we argue that in the final hybrid Hyb_6 , any PPT adversary cannot distinguish an encryption of $\overline{\mathbf{m}}_0$ or $\overline{\mathbf{m}}_1$ in the challenge epoch \widetilde{e} , which relies on the distributional semantic security of KPHE, which will conclude our proof.

Assume for the purpose of contradiction that there exists a PPT adversary \mathcal{A} that can distinguish an encryption of $\overline{\mathbf{m}}_0$ or $\overline{\mathbf{m}}_1$ in the challenge epoch. Then we construct a PPT adversary \mathcal{B} that breaks the distributional semantic security of KPHE. The adversary \mathcal{B} first receives $(\mathbf{pp}, \mathbf{pk})$ from the challenger in the semantic security game. Then \mathcal{B} plays the UE game with \mathcal{A} as a challenger in Hyb_6 . \mathcal{B} uses \mathbf{pp} to generate UE keys and update tokens as in Hyb_6 except that for epoch \widetilde{e} , the UE key $\mathbf{k}_{\widetilde{e}}$ is unknown. When \mathcal{B} receives the challenge messages $(\overline{\mathbf{m}}_0, \overline{\mathbf{m}}_1)$ from \mathcal{A} in the UE game, it forwards the two messages to the KPHE challenger and gets back ctx , and then responds to \mathcal{A} with $\text{ct} = (0, (\mathbf{pp}, \text{ctx}))$. Note that \mathcal{B} doesn't need to know $\mathbf{k}_{\widetilde{e}}$ because it is never used. In particular, \mathcal{B} can use \mathbf{pk} to compute all the $\text{UE.enc}(\mathbf{k}_{\widetilde{e}}, \cdot)$. Finally, \mathcal{B} outputs whatever \mathcal{A} outputs.

If \mathcal{A} can distinguish between encryptions of $\overline{\mathbf{m}}_0$ and $\overline{\mathbf{m}}_1$ with non-negligible probability, then \mathcal{B} can break the distributional semantic security of KPHE with non-negligible probability, which leads to contradiction.

Lemma 3.6. $\text{Hyb}_0 \stackrel{c}{\approx} \text{Hyb}_1$ in the proof of Theorem 3.5.

Lemma 3.7. $\text{Hyb}_4 \stackrel{c}{\approx} \text{Hyb}_5$ in the proof of Theorem 3.5.

We defer the formal proofs of Lemmas 3.6 and 3.7 to Appendix A. These proofs complete the overall proof of Theorem 3.5. □

Remark 3.8. In our construction one can derive \mathbf{k}_{e-1} from Δ_e and \mathbf{k}_e . It is for this reason that our construction permits backward-leak unidirectional key updates proposed by Nishimaki [Nis21] where secret keys can be derived in the backward direction but not forward direction. However, as discussed earlier, this notion is essentially equivalent to no-directional key updates (the optimal case) and has no bearing on our security analysis.

3.3 Post-Compromise Secure Unidirectional UE

In this section, we show that any KPHE scheme with $2n$ -bit secret keys that satisfies distributional security with respect to the distribution \mathcal{U}_n , as well as public key and ciphertext blinding as described in Section 2 implies a post-compromise secure unidirectional UE scheme.

3.3.1 IND-UPD Secure Unidirectional UE

We first show how to construct an UE scheme that satisfies the IND-UPD security definition as proposed in [LT18]. Given a KHPE scheme of the form

$$\text{KPHE} = (\text{KPHE.Setup}, \text{KPHE.KeyGen}, \text{KPHE.Enc}, \text{KPHE.Dec}, \text{KPHE.TransPK}, \text{KPHE.Eval}),$$

with $2n$ -bit secret keys, we construct a unidirectional UE scheme

$$\text{UE} = (\text{UE.setup}, \text{UE.next}, \text{UE.enc}, \text{UE.upd}, \text{UE.dec}),$$

that only differs from the IND-ENC construction in UE.upd :

- $\text{UE.upd}(\Delta_{e+1}, \text{ct}_e)$: Parse the update token and the ciphertext as

$$\Delta_{e+1} = (\text{pk}_\Delta, \text{ct}_{x_\Delta}), \quad \text{ct}_e = (t, (\overline{\text{pk}}_{e-t}, \overline{\text{ct}}_{x_{e-t}}), \dots, (\overline{\text{pk}}_{e-1}, \overline{\text{ct}}_{x_{e-1}}), (\text{pk}_e, \text{ct}_{x_e}))$$

Sample $(t+1)$ uniform random permutations $\pi_{e-t}, \dots, \pi_e : [2n] \rightarrow [2n]$. Also, let $\pi_{\text{id}} : [2n] \rightarrow [2n]$ denote the identity permutation. For each $i \in \{e-t+1, \dots, e-1\}$, compute

$$(\widetilde{\text{pk}}_i, \widetilde{\text{ct}}_{x_i}) \stackrel{\$}{\leftarrow} \text{KPHE.Eval}(\overline{\text{pk}}_i, \overline{\text{ct}}_{x_i}, \pi_i, \pi_{i-1}).$$

Additionally, compute

$$\begin{aligned} (\widetilde{\text{pk}}_{e-t}, \widetilde{\text{ct}}_{x_{e-t}}) &\stackrel{\$}{\leftarrow} \text{KPHE.Eval}(\overline{\text{pk}}_{e-t}, \overline{\text{ct}}_{x_{e-t}}, \pi_{e-t}, \pi_{\text{id}}), \\ (\overline{\text{pk}}_e, \overline{\text{ct}}_{x_e}) &\stackrel{\$}{\leftarrow} \text{KPHE.Eval}(\text{pk}_e, \text{ct}_{x_e}, \pi_e, \pi_{e-1}), \\ (\text{pk}_{e+1}, \text{ct}_{x_{e+1}}) &\stackrel{\$}{\leftarrow} \text{KPHE.Eval}(\text{pk}_\Delta, \text{ct}_{x_\Delta}, \pi_{\text{id}}, \pi_e). \end{aligned}$$

Output the updated ciphertext as:

$$\text{ct}_{e+1} = (t+1, (\widetilde{\text{pk}}_{e-t}, \widetilde{\text{ct}}_{x_{e-t}}), \dots, (\widetilde{\text{pk}}_{e-1}, \widetilde{\text{ct}}_{x_{e-1}}), (\overline{\text{pk}}_e, \overline{\text{ct}}_{x_e}), (\text{pk}_{e+1}, \text{ct}_{x_{e+1}})).$$

Correctness. We first prove the correctness of the UE scheme. For any $m \in \mathcal{M}$, any $\mathbf{k}_0 \leftarrow \text{UE.setup}(1^\lambda)$, any sequence of $(\mathbf{k}_1, \Delta_1), \dots, (\mathbf{k}_e, \Delta_e)$ generated as $(\mathbf{k}_i, \Delta_i) \leftarrow \text{UE.next}(\mathbf{k}_{i-1})$ for all $i \in [e]$, let $\text{ct}_0 \leftarrow \text{UE.enc}(\mathbf{k}_0, m)$ and $\text{ct}_i \leftarrow \text{UE.upd}(\Delta_i, \text{ct}_{i-1})$ for all $j \in [e]$, then the final ciphertext is of the form $\text{ct}_e = (e, (\overline{\text{pk}}_0, \overline{\text{ct}}_{x_0}), \dots, (\overline{\text{pk}}_{e-1}, \overline{\text{ct}}_{x_{e-1}}), (\text{pk}_e, \text{ct}_{x_e}))$. All the UE secret keys are of the form $\mathbf{k}_0 = (\text{pp}, \text{sk}_0), \dots, \mathbf{k}_e = (\text{pp}, \text{sk}_e)$. We can prove by induction that there exist permutations $\pi_0, \pi_1, \dots, \pi_{e-1} : [2n] \rightarrow [2n]$ such that $\overline{\text{sk}}_i = \pi_i(\text{sk}_i)$ for all $i = 0, 1, \dots, e-1$, and that $\text{KPHE.Dec}(\overline{\text{sk}}_0, \overline{\text{ct}}_{x_0}) = m, \text{KPHE.Dec}(\overline{\text{sk}}_1, \overline{\text{ct}}_{x_1}) = \overline{\text{sk}}_0, \dots, \text{KPHE.Dec}(\overline{\text{sk}}_{e-1}, \overline{\text{ct}}_{x_{e-1}}) = \overline{\text{sk}}_{e-2}, \text{KPHE.Dec}(\text{sk}_e, \text{ct}_{x_e}) = \overline{\text{sk}}_{e-1}$. Therefore, $\text{UE.dec}(\mathbf{k}_e, \text{ct}_e)$ outputs m . This argument is for any ciphertext starting from epoch 0. The same argument holds for any ciphertext starting from any epoch \hat{e} where $0 \leq \hat{e} \leq e$.

Confidentiality. Next we prove the IND-UPD security the UE scheme. More formally, we state and prove the following theorem (the proof is provided in Appendix B):

Theorem 3.9 (IND-UPD Security). *Assuming that KPHE is a KPHE scheme with $2n$ -bit secret keys that satisfies distributional security with respect to the distribution \mathcal{U}_n , as well as public key and ciphertext blinding as described in Section 2, the above UE construction is an IND-UPD secure unidirectional UE scheme.*

3.3.2 IND-UE Secure Unidirectional UE

The basic IND-UPD construction allows ciphertexts from the same epoch e to have different sizes. In particular, a freshly created ciphertext in epoch e can be trivially distinguished from a ciphertext that was created as an update of a ciphertext from epoch $(e - 1)$. So it cannot satisfy the combined security definition of post-compromise security for UE due to Boyd et al. [BDGJ20].

We showcase here a simple extension of the basic construction wherein we ensure that the size for any ciphertext in epoch e is the same, irrespective of whether it was freshly created, or created as an update of a ciphertext from epoch $(e - 1)$. The overall construction remains exactly the same; the key alteration is in how we generate fresh ciphertexts. At a high level, a freshly created ciphertext in epoch e is made to look exactly like a ciphertext that has undergone e update operations. We do this by having e “dummy wrapper” layers over and above the core ciphertext generated by the basic construction.

We describe the new encryption algorithm below (we assume here that information about the epoch e is available as part of \mathbf{k}_e):

- $\text{UE.enc}(\mathbf{k}_e, m)$: Parse $\mathbf{k}_e = (e, \text{pp}, \text{sk}_e)$ and generate $\text{pk}_e \xleftarrow{\$} \text{KPHE.PKGen}(\text{pp}, \text{sk}_e)$.

For each $i \in \{0, \dots, e - 1\}$, generate

$$\overline{\text{sk}}_i \xleftarrow{\$} \text{KPHE.SKGen}(\text{pp}, \mathcal{U}_n), \quad \overline{\text{pk}}_i \xleftarrow{\$} \text{KPHE.PKGen}(\text{pp}, \overline{\text{sk}}_i).$$

Compute $\overline{\text{ctx}}_0 \xleftarrow{\$} \text{KPHE.Enc}(\overline{\text{pk}}_0, m)$ and $\text{ctx}_e \xleftarrow{\$} \text{KPHE.Enc}(\text{pk}_e, \overline{\text{sk}}_{e-1})$.

For each $i \in [e - 1]$, compute

$$\overline{\text{ctx}}_i \xleftarrow{\$} \text{KPHE.Enc}(\overline{\text{pk}}_i, \overline{\text{sk}}_{i-1}).$$

Output the ciphertext $\text{ct}_e = (e, (\overline{\text{pk}}_0, \overline{\text{ctx}}_0), \dots, (\overline{\text{pk}}_{e-1}, \overline{\text{ctx}}_{e-1}), (\text{pk}_e, \text{ctx}_e))$.

The only difference between the above UE construction and the IND-UPD secure UE construction presented in Section 3.3.1 is that all the fresh and updated ciphertexts in epoch e have the same length. The correctness and confidentiality proofs of the new construction follow exactly the same way as the IND-UPD secure construction. In particular, for the security proof, we can use the same hybrid argument as in Theorem 3.9 and prove that the challenge ciphertext $\tilde{\text{ct}}_e$ is computationally indistinguishable from a ciphertext freshly generated in Hyb_7 of the form

$$(\tilde{e}, (\tilde{\text{pk}}_0, \tilde{\text{ctx}}_0), \dots, (\tilde{\text{pk}}_{e-1}, \tilde{\text{ctx}}_{e-1}), (\tilde{\text{pk}}_{\tilde{e}}, \text{ctx}_{\tilde{e}})),$$

where each ctx is a KPHE encryption of 0^{2n} . We state the theorem below and omit the detailed proof.

Theorem 3.10 (IND-UE Security). *Assuming that KPHE is a KPHE scheme with $2n$ -bit secret keys that satisfies distributional security with respect to the distribution \mathcal{U}_n , as well as public key and ciphertext blinding as described in Section 2, the above UE construction is an IND-UE secure unidirectional UE scheme.*

4 Unidirectional PRE from Circular-Secure KPHE

In this section, we show how to construct unidirectional PRE from any KPHE scheme that satisfies distributional circular security. We present the simpler construction of IND-HRA unidirectional PRE in Section 4.2. Subsequently, in Appendix D, we show how to augment it to achieve the stronger notion of strong post-compromise security (PCS) as introduced in a recent work by Davidson et al. [DDL19].

4.1 Definition

Definition 4.1 (Unidirectional Proxy Re-Encryption (PRE)). *A unidirectional PRE scheme is a tuple of PPT algorithms of the form*

$$\text{PRE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{ReKeyGen}, \text{ReEnc}, \text{Dec}),$$

described as follows:

- $\text{pp} \stackrel{\$}{\leftarrow} \text{Setup}(1^\lambda)$: *On input the security parameter λ , the setup algorithm outputs some public parameters pp (these parameters are implicit to all other algorithms).*
- $(\text{sk}, \text{pk}) \stackrel{\$}{\leftarrow} \text{KeyGen}(\text{pp})$: *On input the public parameters pp , the key-generation algorithm outputs a secret key-public key pair, (sk, pk) .*
- $\text{ct} \stackrel{\$}{\leftarrow} \text{Enc}(\text{pk}, \text{m})$: *On input a public key pk and a message m , the encryption algorithm outputs a ciphertext ct .*
- $\text{rk}_{i,j} \stackrel{\$}{\leftarrow} \text{ReKeyGen}((\text{sk}_i, \text{pk}_i), \text{pk}_j)$: *The re-key generation algorithm returns a re-encryption key $\text{rk}_{i,j}$ for translation of a ciphertext from a key-pair $(\text{sk}_i, \text{pk}_i)$ to a key-pair $(\text{sk}_j, \text{pk}_j)$. It takes as input $(\text{sk}_i, \text{pk}_i)$ and pk_j , and outputs the re-encryption key $\text{rk}_{i,j}$.²*
- $\text{ct}_j \stackrel{\$}{\leftarrow} \text{ReEnc}(\text{rk}_{i,j}, \text{ct}_i)$: *On input a re-encryption key $\text{rk}_{i,j}$ and a ciphertext ct_i , the re-encryption algorithm outputs an updated ciphertext ct_j .³*
- $\text{m}/\perp \leftarrow \text{Dec}(\text{sk}, \text{ct})$: *On input a secret key sk and a ciphertext ct , the decryption algorithm outputs either a plaintext message or an error symbol.*

Definition 4.2 (Correctness). *A PRE scheme $\text{PRE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{ReKeyGen}, \text{ReEnc}, \text{Dec})$ is said to be correct if for any $\text{pp} \stackrel{\$}{\leftarrow} \text{Setup}(1^\lambda)$, for any $\ell \geq 0$, for any $(\ell+1)$ key-pairs $(\text{pk}_0, \text{sk}_0), \dots,$*

²In a bidirectional PRE scheme, the re-key generation algorithm additionally takes as input the destination secret key sk_j , i.e., it takes as input $(\text{sk}_i, \text{pk}_i)$ and $(\text{sk}_j, \text{pk}_j)$, and outputs the re-encryption key $\text{rk}_{i,j}$.

³The re-encryption algorithm could be either deterministic or randomized; in this work, we assume throughout that the re-encryption algorithm is randomized.

<p><u>Setup(1^λ):</u></p> <pre> pp $\stackrel{\\$}{\leftarrow}$ Setup(1^λ) state := \perp $\mathcal{L}, \tilde{\mathcal{L}}, \mathcal{K}_{\text{Corrupt}}, \mathcal{K}_{\text{Honest}} := \emptyset$ return pp </pre> <p><u>$\mathcal{O}.$KeyGen(n, \mathcal{K}):</u></p> <pre> $(sk_1, pk_1), \dots, (sk_n, pk_n) \stackrel{\\$}{\leftarrow}$ KeyGen(1^λ) $\mathcal{K}_{\text{Corrupt}} := \mathcal{K}$ $\mathcal{K}_{\text{Honest}} := [n] \setminus \mathcal{K}$ state := $\{sk_i, pk_i\}_{i \in [n]}$ return $(\{sk_i\}_{i \in \mathcal{K}_{\text{Corrupt}}}, \{pk_i\}_{i \in [n]})$ </pre> <p><u>$\mathcal{O}.$enc(i, m)</u></p> <pre> ct $\stackrel{\\$}{\leftarrow}$ Enc(pk_i, m) $\mathcal{L} := \mathcal{L} \cup (i, ct)$ return ct </pre> <p><u>$\mathcal{O}.$ReKeyGen(i, j)</u></p> <pre> if $i \in \mathcal{K}_{\text{Honest}}$ and $j \in \mathcal{K}_{\text{Corrupt}}$ then return \perp $rk_{i,j} \stackrel{\\$}{\leftarrow}$ ReKeyGen(sk_i, pk_i, pk_j) return $rk_{i,j}$ </pre> <p><u>$\mathcal{O}.$ReEnc(i, j, ct)</u></p> <pre> if $i \in \mathcal{K}_{\text{Honest}}$ and $j \in \mathcal{K}_{\text{Corrupt}}$ then return \perp $rk_{i,j} \stackrel{\\$}{\leftarrow}$ ReKeyGen(sk_i, pk_i, pk_j) $ct' \stackrel{\\$}{\leftarrow}$ ReEnc($rk_{i,j}, ct$) return ct' </pre>	<p><u>$\mathcal{O}.$HonReEnc(i, j, ct)</u></p> <pre> if $(i, ct) \notin \mathcal{L}$, or $(i, ct) \notin \tilde{\mathcal{L}}$ then return \perp if $(i, ct) \in \tilde{\mathcal{L}}$ and $j \in \mathcal{K}_{\text{Corrupt}}$ then return \perp $rk_{i,j} \stackrel{\\$}{\leftarrow}$ ReKeyGen(sk_i, pk_i, pk_j) $ct' \stackrel{\\$}{\leftarrow}$ ReEnc($rk_{i,j}, ct$) if $(i, ct) \in \tilde{\mathcal{L}}$ then $\tilde{\mathcal{L}} := \tilde{\mathcal{L}} \cup \{(j, ct')\}$ else $\mathcal{L} := \mathcal{L} \cup \{(j, ct')\}$ return ct' </pre> <p><u>$\mathcal{O}.$chall-PRE(i, \bar{m}_0, \bar{m}_1)^a</u></p> <pre> if $i \in \mathcal{K}_{\text{Corrupt}}$ or $\bar{m}_0 \neq \bar{m}_1$ then return \perp $\tilde{ct} \stackrel{\\$}{\leftarrow}$ Enc(pk_i, \bar{m}_b) $\tilde{\mathcal{L}} := \tilde{\mathcal{L}} \cup \{(i, \tilde{ct})\}$ return \tilde{ct} </pre> <p><u>$\mathcal{O}.$chall-IND-PCS($i, j, \bar{ct}_0, \bar{ct}_1$)</u></p> <pre> if $(i, \bar{ct}_0) \notin \mathcal{L}$ or $(i, \bar{ct}_1) \notin \mathcal{L}$ then return \perp if $\bar{ct}_0 \neq \bar{ct}_1$ then return \perp if $j \in \mathcal{K}_{\text{Corrupt}}$ then return \perp $rk_{i,j} \stackrel{\\$}{\leftarrow}$ ReKeyGen(sk_i, pk_i, pk_j) $\tilde{ct} \stackrel{\\$}{\leftarrow}$ ReEnc($rk_{i,j}, \bar{ct}_b$) $\tilde{\mathcal{L}} := \tilde{\mathcal{L}} \cup \{(j, \tilde{ct})\}$ return \tilde{ct} </pre> <hr style="width: 20%; margin-left: auto; margin-right: auto;"/> <p>^aThis is the challenge oracle that is used in both the IND-CPA and the IND-HRA experiments for PRE.</p>
--	--

Figure 5: Oracles in security games for unidirectional PRE.

$(pk_\ell, sk_\ell) \stackrel{\$}{\leftarrow}$ KeyGen(pp), and for any plaintext message m , letting $ct_0 \stackrel{\$}{\leftarrow}$ Enc(pk_0, m), and letting for each $j \in [\ell]$

$$rk_j \stackrel{\$}{\leftarrow}$$
 ReKeyGen(sk_{j-1}, pk_{j-1}, pk_j), $ct_j \stackrel{\$}{\leftarrow}$ ReEnc(rk_j, ct_{j-1}),

we have $\text{Dec}(sk_\ell, ct_\ell) = m$ (with all but negligible probability).

Confidentiality. We recall the various security notions for unidirectional PRE, namely IND-CPA security, IND-HRA security (introduced by Cohen [Coh19]), and IND-PCS security (introduced by Davidson et al. [DDLM19]). Our definitions are game-based, where the game is played between a

challenger and a PPT adversary \mathcal{A} . We assume that the adversary \mathcal{A} in a PRE security game has access to (a subset of) the oracles defined in Figure 5. As with our UE definitions, we begin by introducing the following notations for our security definitions:

- \mathcal{L} : Set of non-challenge ciphertexts (i, ct_i) (i being the key-index under which ct_i is generated) produced by calls to the $\mathcal{O}.\text{enc}$ or $\mathcal{O}.\text{HonReEnc}$ oracle. $\mathcal{O}.\text{HonReEnc}$ only outputs a re-encryption of a non-challenge ciphertext provided that this ciphertext is currently available in \mathcal{L} .
- $\tilde{\mathcal{L}}$: Set of re-encrypted versions of the challenge ciphertext $\tilde{\text{ct}}$. $\tilde{\mathcal{L}}$ is initiated with the challenge ciphertext $(i, \tilde{\text{ct}})$ (i being the key-index under which $\tilde{\text{ct}}$ is generated). Any call to the $\mathcal{O}.\text{HonReEnc}$ oracle on a ciphertext in $\tilde{\mathcal{L}}$ results in a re-encrypted challenge ciphertext that gets added to $\tilde{\mathcal{L}}$.
- $\mathcal{K}_{\text{Corrupt}}$: Set of indices $i \in [n]$ for which the adversary has corrupted the secret key sk_i (using $\mathcal{O}.\text{KeyGen}$).
- $\mathcal{K}_{\text{Honest}}$: Set of indices $i \in [n]$ for which the adversary has not corrupted the secret key sk_i .

We now formally define the various security notions of unidirectional PRE. Note that the CPA and HRA proceed almost identically, with the major difference being the restrictions posed by their respective re-encryption oracles, $\mathcal{O}.\text{ReEnc}$ and $\mathcal{O}.\text{HonReEnc}$.

Definition 4.3 (IND-CPA/IND-HRA Security). *A unidirectional PRE scheme*

$$\text{PRE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{ReKeyGen}, \text{ReEnc}, \text{Dec})$$

is said to be IND-CPA-secure (IND-HRA-secure resp.) if for any security parameter λ and any (non-uniform) PPT adversary \mathcal{A} , it holds that

$$\left| \Pr \left[\text{Exp}_{\mathcal{A}, \text{PRE}}^{\text{IND-CPA/IND-HRA}}(1^\lambda) = 1 \right] - \frac{1}{2} \right| \leq \text{negl}(\lambda),$$

for some negligible function $\text{negl}(\cdot)$, where the experiment $\text{Exp}_{\mathcal{A}, \text{PRE}}^{\text{IND-CPA}}$ ($\text{Exp}_{\mathcal{A}, \text{PRE}}^{\text{IND-HRA}}$ resp.) is defined as below.

Experiment $\text{Exp}_{\mathcal{A}, \text{PRE}}^{\text{IND-CPA/IND-HRA}}(1^\lambda)$:

Run $\text{pp} \xleftarrow{\$} \text{Setup}(1^\lambda)$
 $(\text{state}, n, \mathcal{K}) \leftarrow \mathcal{A}(\text{pp})$
 Run $(\{\text{sk}_i\}_{i \in \mathcal{K}_{\text{Corrupt}}}, \{\text{pk}_i\}_{i \in [n]}) \xleftarrow{\$} \mathcal{O}.\text{KeyGen}(n, \mathcal{K})$
 $(\text{state}, i, \bar{\text{m}}_0, \bar{\text{m}}_1) \leftarrow \mathcal{A}^{\mathcal{O}.\text{enc}, \mathcal{O}.\text{ReKeyGen}, \mathcal{O}.\text{ReEnc}/\mathcal{O}.\text{HonReEnc}}(\text{state}, \{\text{sk}_i\}_{i \in \mathcal{K}_{\text{Corrupt}}}, \{\text{pk}_i\}_{i \in [n]})$
 $\text{b} \xleftarrow{\$} \{0, 1\}$
 $\tilde{\text{ct}} \leftarrow \mathcal{O}.\text{chall-PRE}(i, \bar{\text{m}}_0, \bar{\text{m}}_1)$
 Proceed only if $\tilde{\text{ct}} \neq \perp$
 $\text{b}' \leftarrow \mathcal{A}^{\mathcal{O}.\text{enc}, \mathcal{O}.\text{ReKeyGen}, \mathcal{O}.\text{ReEnc}/\mathcal{O}.\text{HonReEnc}}(\text{state}, \tilde{\text{ct}})$
return 1 if $\text{b} = \text{b}'$

Definition 4.4 (IND-PCS Security). *A unidirectional PRE scheme*

$$\text{PRE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{ReKeyGen}, \text{ReEnc}, \text{Dec})$$

is said to be IND-PCS-secure if for any security parameter λ and any (non-uniform) PPT adversary \mathcal{A} , it holds that

$$\left| \Pr \left[\text{Exp}_{\mathcal{A}, \text{PRE}}^{\text{IND-PCS}}(1^\lambda) = 1 \right] - \frac{1}{2} \right| \leq \text{negl}(\lambda),$$

for some negligible function $\text{negl}(\cdot)$, where the experiment $\text{Exp}_{\mathcal{A}, \text{PRE}}^{\text{IND-PCS}}$ is defined as below.

Experiment $\text{Exp}_{\mathcal{A}, \text{PRE}}^{\text{IND-PCS}}(1^\lambda)$:

Run $\text{pp} \xleftarrow{\$} \text{Setup}(1^\lambda)$
 $(\text{state}, n, \mathcal{K}) \leftarrow \mathcal{A}(\text{pp})$
 Run $(\{\text{sk}_i\}_{i \in \mathcal{K}_{\text{Corrupt}}}, \{\text{pk}_i\}_{i \in [n]}) \xleftarrow{\$} \mathcal{O}.\text{KeyGen}(n, \mathcal{K})$
 $(\text{state}, i, j, \overline{\text{ct}}_0, \overline{\text{ct}}_1) \leftarrow \mathcal{A}^{\mathcal{O}.\text{enc}, \mathcal{O}.\text{ReKeyGen}, \mathcal{O}.\text{HonReEnc}}(\text{state}, \{\text{sk}_i\}_{i \in \mathcal{K}_{\text{Corrupt}}}, \{\text{pk}_i\}_{i \in [n]})$
 $b \xleftarrow{\$} \{0, 1\}$
 $\tilde{\text{ct}} \leftarrow \mathcal{O}.\text{chall-IND-PCS}(i, j, \overline{\text{ct}}_0, \overline{\text{ct}}_1)$
 Proceed only if $\tilde{\text{ct}} \neq \perp$
 $b' \leftarrow \mathcal{A}^{\mathcal{O}.\text{enc}, \mathcal{O}.\text{ReKeyGen}, \mathcal{O}.\text{HonReEnc}}(\text{state}, \tilde{\text{ct}})$
return 1 if $b = b'$

4.2 HRA-Secure Unidirectional PRE

We show that any KPHE scheme with $2n$ -bit secret keys and plaintext messages that satisfies: (a) distributional semantic and *circular* security with respect to the distribution \mathcal{U}_n , and (b) blinding, implies the existence of a multi-hop IND-HRA secure unidirectional PRE scheme.

Construction. Given a KHPE scheme of the form

$$\text{KPHE} = (\text{KPHE.Setup}, \text{KPHE.SKGen}, \text{KPHE.PKGen}, \text{KPHE.Enc}, \text{KPHE.Dec}, \text{KPHE.Eval}),$$

with $2n$ -bit secret keys, we construct a unidirectional PRE scheme

$$\text{PRE} = (\text{PRE.Setup}, \text{PRE.KeyGen}, \text{PRE.Enc}, \text{PRE.ReKeyGen}, \text{PRE.ReEnc}, \text{PRE.Dec}),$$

with message space $\mathcal{M} = \{0, 1\}^{2n}$ as follows:

- $\text{PRE.Setup}(1^\lambda)$: Sample $\text{pp} \xleftarrow{\$} \text{KPHE.Setup}(1^\lambda)$ and output pp .
- $\text{PRE.KeyGen}(\text{pp})$: Sample and output (pk, sk) where

$$\text{sk} \xleftarrow{\$} \text{KPHE.SKGen}(\text{pp}, \mathcal{U}_n), \quad \text{pk} \xleftarrow{\$} \text{KPHE.PKGen}(\text{pp}, \text{sk}).$$

- $\text{PRE.Enc}(\text{pk}, m)$: Compute $\text{ctx}_0 \xleftarrow{\$} \text{KPHE.Enc}(\text{pk}, m)$ and output

$$\text{ct} = (0, (\text{pk}, \text{ctx}_0)).$$

- $\text{PRE.ReKeyGen}(\text{sk}_i, \text{pk}_i, \text{pk}_j)$: Output $\text{rk}_{i,j} = (\text{pk}_j, \text{ctx}_\Delta)$, where

$$\text{ctx}_\Delta \stackrel{\$}{\leftarrow} \text{KPHE.Enc}(\text{pk}_j, \text{sk}_i).$$

- $\text{PRE.ReEnc}(\text{rk}_{i,j}, \text{ct})$: Parse the update token and the ciphertext as

$$\text{rk}_{i,j} = (\text{pk}_j, \text{ctx}_\Delta), \quad \text{ct} = (t, (\overline{\text{pk}}_0, \overline{\text{ctx}}_0), \dots, (\overline{\text{pk}}_{t-1}, \overline{\text{ctx}}_{t-1}), (\widehat{\text{pk}}_t, \widehat{\text{ctx}}_t)),$$

for some $t \geq 0$. Sample a uniformly random permutation $\pi : [2n] \rightarrow [2n]$. Also, let $\pi_{\text{id}} : [2n] \rightarrow [2n]$ denote the *identity* permutation. Compute

$$(\overline{\text{pk}}_t, \overline{\text{ctx}}_t) \stackrel{\$}{\leftarrow} \text{KPHE.Eval}(\widehat{\text{pk}}_t, \widehat{\text{ctx}}_t, \pi, \pi_{\text{id}}),$$

$$(\widehat{\text{pk}}_{t+1}, \widehat{\text{ctx}}_{t+1}) \stackrel{\$}{\leftarrow} \text{KPHE.Eval}(\text{pk}_j, \text{ctx}_\Delta, \pi_{\text{id}}, \pi),$$

and output the updated ciphertext as:

$$\text{ct}' = (t + 1, (\overline{\text{pk}}_0, \overline{\text{ctx}}_0), \dots, (\overline{\text{pk}}_t, \overline{\text{ctx}}_t), (\widehat{\text{pk}}_{t+1}, \widehat{\text{ctx}}_{t+1})).$$

- $\text{PRE.Dec}(\text{sk}, \text{ct})$: Parse the ciphertext as

$$\text{ct} = (t, (\overline{\text{pk}}_0, \overline{\text{ctx}}_0), \dots, (\overline{\text{pk}}_{t-1}, \overline{\text{ctx}}_{t-1}), (\widehat{\text{pk}}_t, \widehat{\text{ctx}}_t)),$$

for some $t \geq 0$. Compute $\overline{\text{sk}}_{t-1} = \text{KPHE.Dec}(\text{sk}, \widehat{\text{ctx}}_t)$. Next, compute the following for each ℓ from $(t - 1)$ to 1 in decreasing order:

$$\overline{\text{sk}}_{\ell-1} = \text{KPHE.Dec}(\overline{\text{sk}}_\ell, \overline{\text{ctx}}_\ell).$$

Finally, output the message $\mathbf{m} = \text{KPHE.Dec}(\overline{\text{sk}}_0, \overline{\text{ctx}}_0)$.

We prove correctness in the following lemma:

Lemma 4.5 (Correctness). *If the underlying KPHE scheme is correct, then for any $\text{pp} \stackrel{\$}{\leftarrow} \text{PRE.Setup}(1^\lambda)$, for any sequence of $(n + 1)$ key-pairs (where $n \geq 0$):*

$$(\text{pk}_0, \text{sk}_0), (\text{pk}_1, \text{sk}_1), \dots, (\text{pk}_n, \text{sk}_n) \stackrel{\$}{\leftarrow} \text{PRE.KeyGen}(\text{pp}),$$

and for any plaintext message \mathbf{m} , letting $\text{ct}_0 \stackrel{\$}{\leftarrow} \text{PRE.Enc}(\text{pk}_1, \mathbf{m})$, and letting for each $j \in [n]$

$$\text{rk}_j \stackrel{\$}{\leftarrow} \text{PRE.ReKeyGen}(\text{sk}_{j-1}, \text{pk}_{j-1}, \text{pk}_j), \quad \text{ct}_j \stackrel{\$}{\leftarrow} \text{PRE.ReEnc}(\text{rk}_j, \text{ct}_{j-1}),$$

we must have $\text{PRE.Dec}(\text{sk}_n, \text{ct}_n) = \mathbf{m}$ (with all but negligible probability).

Proof. From the description of the re-encryption algorithm, it follows that

$$\text{ct}_n = (n, (\overline{\text{pk}}_0, \overline{\text{ctx}}_0), (\overline{\text{pk}}_1, \overline{\text{ctx}}_1), \dots, (\overline{\text{pk}}_{n-1}, \overline{\text{ctx}}_{n-1}), (\widehat{\text{pk}}_n, \widehat{\text{ctx}}_n)),$$

where for each $j \in [0, n - 1]$, letting $\pi_j : [2n] \rightarrow [2n]$ be a uniformly random permutation sampled during the execution of $\text{PRE.ReEnc}(\text{rk}_{j+1}, \text{ct}_j)$, and letting $\overline{\text{sk}}_j = \pi_j(\text{sk}_j)$, we have the following (with all but negligible probability) whenever the KPHE scheme is correct:

$$\begin{aligned} \text{KPHE.Dec}(\overline{\text{sk}}_0, \overline{\text{ct}}_{x_0}) &= m, \quad \{\text{KPHE.Dec}(\overline{\text{sk}}_j, \overline{\text{ct}}_{x_j}) = \overline{\text{sk}}_{j-1}\}_{j \in [n-1]}, \\ \text{KPHE.Dec}(\text{sk}_n, \widehat{\text{ct}}_{x_n}) &= \overline{\text{sk}}_{n-1}. \end{aligned}$$

It immediately follows that we have

$$\begin{aligned} &\text{PRE.Dec}(\text{sk}_t, \text{ct}_t) \\ &= \text{KPHE.Dec}(\dots (\text{KPHE.Dec}(\boxed{\text{KPHE.Dec}(\text{sk}_n, \widehat{\text{ct}}_{x_n})}, \overline{\text{ct}}_{x_{n-1}}), \dots), \overline{\text{ct}}_{x_0}) \\ &= \text{KPHE.Dec}(\dots (\text{KPHE.Dec}(\boxed{\overline{\text{sk}}_{n-1}}, \overline{\text{ct}}_{x_{n-1}}), \dots), \overline{\text{ct}}_{x_0}) \\ &\quad \vdots \\ &= \text{KPHE.Dec}(\boxed{\text{KPHE.Dec}(\overline{\text{sk}}_1, \overline{\text{ct}}_{x_1})}, \overline{\text{ct}}_{x_0}) \\ &= \text{KPHE.Dec}(\boxed{\overline{\text{sk}}_0}, \overline{\text{ct}}_{x_0}) \\ &= m, \end{aligned}$$

with all but negligible probability, as desired. \square

Theorem 4.6 (IND-HRA Security). *Assuming that KPHE is a KPHE scheme with 2n-bit secret keys that satisfies blinding, distributional circular security and distributional semantic security with respect to the distribution \mathcal{U}_n , PRE is a multi-hop IND-HRA secure unidirectional PRE scheme.*

IND-CPA Security (Warm-Up). As a warm-up, we begin by outlining the proof of IND-CPA security for our construction, which is significantly simpler than the proof of IND-HRA security. We subsequently present the proof of IND-HRA security (in Appendix C), which is significantly more nuanced since it requires careful simulation of adversarial re-encryption queries on honestly generated ciphertexts.

We assume throughout that the adversary corrupts keys *statically* at the beginning of the IND-CPA game, while the ReKeyGen queries may be issued adaptively. Finally, the adversary may issue the challenge encryption query adaptively on any honest key of its choice at any point during the IND-CPA game.

The proof proceeds through a sequence of hybrids. We use $[N]$ to denote the set of all keys (honest + corrupt), $\mathcal{K}_{\text{Honest}} \subset [N]$ to denote the set of honest keys, and $\mathcal{K}_{\text{Corrupt}} \subset [N]$ to denote the set of corrupt keys. As mentioned earlier, these sets are fixed (adversarially) at the beginning of the game, before the adversary is allowed to issue re-key generation queries (this remains unchanged in each hybrid).

Hyb₀: This hybrid is identical to the real IND-CPA security game between the challenger and the adversary.

Hyb₁: This hybrid is identical to **Hyb₀** except for the manner in which the challenger answers the ReKeyGen and ReEnc queries issued by the adversary. In particular, at setup, the challenger creates an $(n \times n)$ table T_{rk} (initially empty) and populates it as follows:

- If $i \in \mathcal{K}_{\text{Honest}}$ and $j \in \mathcal{K}_{\text{Corrupt}}$, set $T_{\text{rk}}[i, j] = \perp$.
- Else, set $T_{\text{rk}}[i, j] = (\text{pk}_j, \text{KPHE.Enc}(\text{pk}_j, \text{sk}_i))$.

Given a query of the form $\text{ReKeyGen}(i, j)$, the challenger responds with $T_{\text{rk}}[i, j]$. Additionally, given a query of the form $\text{ReEnc}(i, j, \text{ct})$, the challenger proceeds as follows:

- If $i \in \mathcal{K}_{\text{Honest}}$ and $j \in \mathcal{K}_{\text{Corrupt}}$, respond with \perp .
- Else, respond with the updated ciphertext $\text{ct}' \stackrel{\$}{\leftarrow} \text{ReEnc}(T_{\text{rk}}[i, j], \text{ct})$.

It is easy to see that Hyb_1 is identical to Hyb_0 .

Hyb₂: This hybrid is identical to the hybrid Hyb_1 except that the challenger pre-populates the table T_{rk} at setup as follows: if $i \in \mathcal{K}_{\text{Honest}}$ and $j \in \mathcal{K}_{\text{Honest}}$, set $T_{\text{rk}}[i, j] = (\text{pk}_j, \text{KPHE.Enc}(\text{pk}_j, 0^{2n}))$.

We argue that Hyb_2 is indistinguishable from Hyb_1 in a straightforward manner under the assumption that KPHE satisfies distributional circular security.

Remark 4.7. *Note that in Hyb_2 , the set of honest secret keys $\{\text{sk}_i\}_{i \in \mathcal{K}_{\text{Honest}}}$ is no longer used by the challenger when answering ReKeyGen and ReEnc queries. In other words, Hyb_2 allows the challenger to “forget” the honest secret keys.*

Hyb₃: This hybrid is identical to the hybrid Hyb_2 except for the manner in which the challenger answers the challenge encryption query issued by the adversary. In particular, given a query of the form $\mathcal{O}.\text{chall}(i, \bar{m}_0, \bar{m}_1)$, the challenger proceeds as follows:

- If $i \in \mathcal{K}_{\text{Corrupt}}$, respond with $\text{ct}^* = \perp$ (this is exactly as in Hyb_1).
- If $i \in \mathcal{K}_{\text{Honest}}$, respond with $\text{ct}^* = (0, \text{KPHE.Enc}(\text{pk}_i, 0^{2n}))$.

We argue that Hyb_3 is indistinguishable from Hyb_2 under the assumption that KPHE satisfies distributional semantic security.

We defer the detailed proof of IND-HRA security to Appendix C. Finally, we defer the construction and proof of IND-PCS secure unidirectional PRE from KPHE to Appendix D.

References

- [ACDT20] Joël Alwen, Sandro Coretti, Yevgeniy Dodis, and Yiannis Tselekounis. Security analysis and improvements for the IETF MLS standard for group messaging. In *CRYPTO*, 2020.
- [ACPS09] Benny Applebaum, David Cash, Chris Peikert, and Amit Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In *CRYPTO 2009*, volume 5677, pages 595–618, 2009.
- [AFGH06] Giuseppe Ateniese, Kevin Fu, Matthew Green, and Susan Hohenberger. Improved proxy re-encryption schemes with applications to secure distributed storage. *ACM Trans. Inf. Syst. Secur.*, 9(1):1–30, 2006.

- [AMP19] Navid Alamati, Hart Montgomery, and Sikhar Patranabis. Symmetric primitives with structured secrets. In *CRYPTO*, 2019.
- [BBB⁺12] Elaine Barker, William Barker, William Burr, William Polk, Miles Smid, Patrick D. Gallagher, and Under Secretary For. NIST Special Publication 800-57 Recommendation for Key Management – Part 1: General, 2012.
- [BBS98] Matt Blaze, Gerrit Bleumer, and Martin Strauss. Divertible protocols and atomic proxy cryptography. In *EUROCRYPT*, 1998.
- [BDGJ20] Colin Boyd, Gareth T. Davies, Kristian Gjøsteen, and Yao Jiang. Fast and secure updatable encryption. In *CRYPTO*, 2020.
- [BEKS20] Dan Boneh, Saba Eskandarian, Sam Kim, and Maurice Shih. Improving speed and security in updatable encryption schemes. In *ASIACRYPT 2020*, pages 559–589, 2020.
- [BF03] Dan Boneh and Matthew K. Franklin. Identity-based encryption from the weil pairing. *SIAM J. Comput.*, 32(3):586–615, 2003.
- [BGI⁺12] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. *J. ACM*, 59(2):6:1–6:48, 2012.
- [BHHO08] Dan Boneh, Shai Halevi, Michael Hamburg, and Rafail Ostrovsky. Circular-secure encryption from decision diffie-hellman. In *CRYPTO 2008*, volume 5157, pages 108–125, 2008.
- [BLMR13] Dan Boneh, Kevin Lewi, Hart William Montgomery, and Ananth Raghunathan. Key homomorphic prfs and their applications. In *CRYPTO*, 2013.
- [BRS02] John Black, Phillip Rogaway, and Thomas Shrimpton. Encryption-scheme security in the presence of key-dependent messages. In *SAC 2002*, volume 2595, pages 62–75. Springer, 2002.
- [CCL⁺14] Nishanth Chandran, Melissa Chase, Feng-Hao Liu, Ryo Nishimaki, and Keita Xagawa. Re-encryption, functional re-encryption, and multi-hop re-encryption: A framework for achieving obfuscation-based security and instantiations from lattices. In *PKC 2014*, pages 95–112, 2014.
- [CL01] Jan Camenisch and Anna Lysyanskaya. An efficient system for non-transferable anonymous credentials with optional anonymity revocation. In *EUROCRYPT 2001*, volume 2045, pages 93–118. Springer, 2001.
- [Coh19] Aloni Cohen. What about bob? the inadequacy of CPA security for proxy reencryption. In *PKC*, 2019.
- [DDL19] Alex Davidson, Amit Deo, Ela Lee, and Keith Martin. Strong post-compromise secure proxy re-encryption. In *ACISP 2019*, pages 58–77, 2019.
- [DKW21] Yevgeniy Dodis, Harish Karthikeyan, and Daniel Wichs. Updatable public key encryption in the standard model. In *TCC*, 2021.

- [EPRS17] Adam Everspaugh, Kenneth G. Paterson, Thomas Ristenpart, and Samuel Scott. Key rotation for authenticated encryption. In *CRYPTO*, 2017.
- [FKKP19] Georg Fuchsbauer, Chethan Kamath, Karen Klein, and Krzysztof Pietrzak. Adaptively secure proxy re-encryption. In *PKC*, 2019.
- [FL17] Xiong Fan and Feng-Hao Liu. Proxy re-encryption and re-signatures from lattices. *IACR Cryptology ePrint Archive*, 2017:456, 2017.
- [Gam85] Taher El Gamal. A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE Trans. Inf. Theory*, 31(4):469–472, 1985.
- [Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In *ACM STOC 2009*, pages 169–178, 2009.
- [GHV10] Craig Gentry, Shai Halevi, and Vinod Vaikuntanathan. *i*-hop homomorphic encryption and rerandomizable yao circuits. In *CRYPTO 2010*, volume 6223, pages 155–172, 2010.
- [GKPV10] Shafi Goldwasser, Yael Tauman Kalai, Chris Peikert, and Vinod Vaikuntanathan. Robustness of the learning with errors assumption. In *ICS 2010*, pages 230–240. Tsinghua University Press, 2010.
- [GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *ACM STOC 2008*, pages 197–206. ACM, 2008.
- [ID03] Anca-Andreea Ivan and Yevgeniy Dodis. Proxy cryptography revisited. In *NDSS*, 2003.
- [Jia20] Yao Jiang. The direction of updatable encryption does not matter much. In *ASIACRYPT 2020*, pages 529–558, 2020.
- [JMM19] Daniel Jost, Ueli Maurer, and Marta Mularczyk. Efficient ratcheting: Almost-optimal guarantees for secure messaging. In *EUROCRYPT*, 2019.
- [Kir14] Elena Kirshanova. Proxy re-encryption from lattices. In *PKC 2014*, pages 77–94, 2014.
- [KLR19] Michael Kloof, Anja Lehmann, and Andy Rupp. (R)CCA secure updatable encryption with integrity protection. In *EUROCRYPT*, 2019.
- [LT18] Anja Lehmann and Björn Tackmann. Updatable encryption with post-compromise security. In *EUROCRYPT*, 2018.
- [LV11] Benoît Libert and Damien Vergnaud. Unidirectional chosen-ciphertext secure proxy re-encryption. *IEEE Trans. Inf. Theory*, 57(3):1786–1802, 2011.
- [NAL15] David Nuñez, Isaac Agudo, and Javier López. Ntrurencrypt: An efficient proxy re-encryption scheme based on NTRU. In *ACM ASIA CCS*, 2015.
- [Nis21] Ryo Nishimaki. The direction of updatable encryption does matter. *Cryptology ePrint Archive*, Report 2021/221, 2021. <https://eprint.iacr.org/2021/221>.
- [NS12] Moni Naor and Gil Segev. Public-key cryptosystems resilient to key leakage. *SIAM J. Comput.*, 41(4):772–814, 2012.

- [NX15] Ryo Nishimaki and Keita Xagawa. Key-private proxy re-encryption from lattices, revisited. *IEICE Transactions*, 98-A(1):100–116, 2015.
- [Pay18] Payment Card Industry (PCI). Data Security Standard – Version 3.2.1, May 2018.
- [PRSV17] Yuriy Polyakov, Kurt Rohloff, Gyana Sahu, and Vinod Vaikuntanathan. Fast proxy re-encryption for publish/subscribe systems. *ACM Trans. Priv. Secur.*, 20(4):14:1–14:31, 2017.
- [PWA⁺16] Le Trieu Phong, Lihua Wang, Yoshinori Aono, Manh Ha Nguyen, and Xavier Boyen. Proxy re-encryption schemes with key privacy from LWE. *IACR Cryptol. ePrint Arch.*, page 327, 2016.
- [Reg09] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6):34:1–34:40, 2009.
- [SD19] Vipin Singh Sehrawat and Yvo Desmedt. Bi-homomorphic lattice-based prfs and unidirectional updatable encryption. In *CANS*, volume 11829, pages 3–23. Springer, 2019.
- [SS21] Daniel Slamanig and Christoph Striecks. Puncture ’em all: Stronger updatable encryption with no-directional key updates. *Cryptology ePrint Archive*, Report 2021/268, 2021. <https://eprint.iacr.org/2021/268>.

A Hybrid Arguments in Proof of IND-ENC Security for UE Scheme (Theorem 3.5)

In this section, we present the proofs for the hybrid arguments in Theorem 3.5, which establishes the IND-ENC security of our UE construction detailed in Section 3.2. In particular, we formally prove Lemmas 3.6 and 3.7.

A.1 Proof of Lemma 3.6

We first present the detailed proof of Lemma 3.6.

Proof. Let Q be the upper bound on the total number of $\text{UE.upd}(\Delta_e, \text{ct}_{e-1})$ and $\text{UE.upd}(\Delta_e, \tilde{\text{ct}}_{e-1})$ outputs. We construct a series of Q intermediate hybrids $\{\text{Hyb}_{0,q}\}_{q \in [Q]}$ where in each hybrid we only change the q -th output of UE.upd (one by one from the first to the last) as described in Hyb_1 . Note that $\text{Hyb}_0 = \text{Hyb}_{0,0}$ and $\text{Hyb}_1 = \text{Hyb}_{0,Q}$. In the following we argue that changing the q -th UE.upd output is computationally indistinguishable to any PPT adversary, namely $\text{Hyb}_{0,q-1} \stackrel{c}{\approx} \text{Hyb}_{0,q}$ for all $q \in [Q]$.

For the q -th output of UE.upd , without loss of generality we assume it is $\text{UE.upd}(\Delta_e, \text{ct}_{e-1})$ handled in $\mathcal{O.upd}$.

Intermediate hybrid $\text{Hyb}'_{0,q-1}$. We first construct an intermediate hybrid $\text{Hyb}'_{0,q-1}$ between $\text{Hyb}_{0,q-1}$ and $\text{Hyb}_{0,q}$ that computes $\text{UE.upd}(\Delta_e, \text{ct}_{e-1})$ as follows:

- Let $\mathbf{k}_{e-1} = (\text{pp}, \text{sk}_{e-1})$.
- Parse the ciphertext ct_{e-1} as

$$(t, (\overline{\text{pk}}_{e-1-t}, \overline{\text{ctx}}_{e-1-t}), \dots, (\overline{\text{pk}}_{e-2}, \overline{\text{ctx}}_{e-2}), (\text{pk}_{e-1}, \text{ctx}_{e-1})),$$

where $\text{ctx}_{e-1} = \text{KPHE.Enc}(\text{pk}_{e-1}, x)$. Note that if $t = 0$, then $x = \mathbf{m}$ for some message, otherwise $x = \overline{\text{sk}}_{e-2}$ that is the KPHE secret key corresponding to $\overline{\text{pk}}_{e-2}$.

- Sample a uniform random permutation $\pi : [2n] \rightarrow [2n]$, let $\overline{\text{sk}}_{e-1} = \pi(\text{sk}_{e-1})$, and sample $\overline{\text{pk}}_{e-1} \stackrel{\$}{\leftarrow} \text{KPHE.PKGen}(\text{pp}, \overline{\text{sk}}_{e-1})$. Compute $\overline{\text{ctx}}_{e-1} \stackrel{\$}{\leftarrow} \text{KPHE.Enc}(\overline{\text{pk}}_{e-1}, x)$.
- $(\text{pk}_e, \text{ctx}_e)$ are computed same as in $\text{Hyb}_{0,q-1}$ using π and π_{id} . That is, let $\Delta_e = (\text{pk}_\Delta, \text{ctx}_\Delta)$, compute $(\text{pk}_e, \text{ctx}_e) \stackrel{\$}{\leftarrow} \text{KPHE.Eval}(\text{pk}_\Delta, \text{ctx}_\Delta, \pi_{\text{id}}, \pi)$.
- Let ct_e be

$$(t+1, (\overline{\text{pk}}_{e-1-t}, \overline{\text{ctx}}_{e-1-t}), \dots, (\overline{\text{pk}}_{e-1}, \overline{\text{ctx}}_{e-1}), (\text{pk}_e, \text{ctx}_e)).$$

$\text{Hyb}_{0,q-1} \stackrel{c}{\approx} \text{Hyb}'_{0,q-1}$. We first argue $\text{Hyb}_{0,q-1} \stackrel{c}{\approx} \text{Hyb}'_{0,q-1}$ by the blinding property of KPHE. Assume for the purpose of contradiction that there exists a PPT adversary \mathcal{A} that can distinguish the two hybrids. Then we construct a PPT adversary \mathcal{B} that breaks the blinding property of KPHE. The adversary \mathcal{B} first receives $(\text{pp}, \text{sk}, \text{pk})$ from the challenger in the KPHE blinding experiment. Then \mathcal{B} plays the UE game with \mathcal{A} as a challenger in $\text{Hyb}_{0,q-1}$.

Let E be the upper bound on the number of epochs during the UE game. \mathcal{B} randomly guesses e^* from $[E]$ as the epoch where the target $\text{UE.upd}(\Delta_e, \text{ct}_{e-1})$ computation occurs. \mathcal{B} uses pp to generate UE keys and update tokens as in $\text{Hyb}_{0,q-1}$ except that in epoch $(e^* - 1)$, it uses (pp, sk) as \mathbf{k}_{e^*-1} .

Let C be the upper bound on the number of ciphertexts generated from UE.enc and UE.upd during epoch $(e^* - 1)$. \mathcal{B} randomly guesses c from $[C]$ as the target ciphertext ct_{e^*-1} . In the computation of ct_{e^*-1} , first set $\text{pk}_{e^*-1} = \text{pk}$ (from the KPHE blinding experiment). Assume ctx_{e^*-1} should be an encryption of x under pk_{e^*-1} for some x . \mathcal{B} sends $\mathbf{m} = y$ to the challenger in the KPHE blinding experiment and gets back a ciphertext ct . Let $\text{ctx}_{e^*-1} := \text{ct}$.

If the q -th UE.upd computation does not occur in epoch e^* or if it is not computed on the above ct_{e^*-1} , namely \mathcal{B} 's guess is wrong, then \mathcal{B} aborts the UE game and outputs a random bit $b \in \{0, 1\}$. If \mathcal{B} 's guess is correct, then \mathcal{B} computes the target $\text{UE.upd}(\Delta_{e^*}, \text{ct}_{e^*-1})$ as follows:

- First note that $\mathbf{k}_{e^*-1} = (\text{pp}, \text{sk})$ received in the KPHE blinding experiment.
- Parse the ciphertext ct_{e^*-1} as

$$(t, (\overline{\text{pk}}_{e^*-1-t}, \overline{\text{ctx}}_{e^*-1-t}), \dots, (\overline{\text{pk}}_{e^*-2}, \overline{\text{ctx}}_{e^*-2}), (\text{pk}_{e^*-1}, \text{ctx}_{e^*-1})),$$

where $\text{ctx}_{e^*-1} = \text{KPHE.Enc}(\text{pk}_{e^*-1}, x)$, which is computed by the challenger in the KPHE blinding experiment.

- Sample a uniform random permutation $\pi : [2n] \rightarrow [2n]$. Send (π, π_{id}) to the challenger in the KPHE blinding experiment and get back $(\overline{\text{pk}}_{e^*-1}, \overline{\text{ctx}}_{e^*-1})$.
- $(\text{pk}_{e^*}, \text{ctx}_{e^*})$ are computed same as in $\text{Hyb}_{0,q-1}$ using π and π_{id} . That is, let $\Delta_{e^*} = (\text{pk}_{\Delta}, \text{ctx}_{\Delta})$, compute $(\text{pk}_{e^*}, \text{ctx}_{e^*}) \stackrel{\$}{\leftarrow} \text{KPHE.Eval}(\text{pk}_{\Delta}, \text{ctx}_{\Delta}, \pi_{\text{id}}, \pi)$.
- Let ct_{e^*} be

$$(t + 1, (\overline{\text{pk}}_{e^*-1-t}, \overline{\text{ctx}}_{e^*-1-t}), \dots, (\overline{\text{pk}}_{e^*-1}, \overline{\text{ctx}}_{e^*-1}), (\text{pk}_{e^*}, \text{ctx}_{e^*})).$$

Finally, \mathcal{B} outputs whatever \mathcal{A} outputs. Note that if the challenger in the KPHE blinding experiment responds to \mathcal{B} with the output of KPHE.Eval , then the game is identical to $\text{Hyb}_{0,q-1}$ to \mathcal{A} ; otherwise the game is identical to $\text{Hyb}'_{0,q-1}$ to \mathcal{A} . If \mathcal{A} can distinguish between the two hybrids with non-negligible probability, then \mathcal{B} can break the blinding property of KPHE with non-negligible probability, which leads to contradiction.

$\text{Hyb}'_{0,q-1} \stackrel{c}{\approx} \text{Hyb}_{0,q}$. We next argue $\text{Hyb}'_{0,q-1} \stackrel{c}{\approx} \text{Hyb}_{0,q}$ also by the blinding property of KPHE. Assume for the purpose of contradiction that there exists a PPT adversary \mathcal{A} that can distinguish between the two hybrids. Then we construct a PPT adversary \mathcal{B} that breaks the blinding property of KPHE. The adversary \mathcal{B} first receives $(\text{pp}, \text{sk}, \text{pk})$ from the challenger in the KPHE blinding experiment. Then \mathcal{B} plays the UE game with \mathcal{A} as a challenger in $\text{Hyb}'_{0,q-1}$.

Let E be the upper bound on the number of epochs during the UE game. \mathcal{B} randomly guesses e^* from $[E]$ as the epoch where the target $\text{UE.upd}(\Delta_e, \text{ct}_{e-1})$ computation occurs. \mathcal{B} uses pp to generate UE keys and update tokens as in $\text{Hyb}'_{0,q-1}$ except that in epoch e^* , it uses (pp, sk) as \mathbf{k}_{e^*} . Additionally, \mathcal{B} sends $\mathbf{m} = \text{sk}_{e^*-1}$ in the KPHE blinding experiment and receives ctx . Then \mathcal{B} sets $\Delta_{e^*} = (\text{pk}, \text{ctx})$.

If the q -th UE.upd computation does not occur in epoch e^* , namely \mathcal{B} 's guess is wrong, then \mathcal{B} aborts the UE game and outputs a random bit $b \in \{0, 1\}$. If \mathcal{B} 's guess is correct, then \mathcal{B} computes the target UE.upd($\Delta_{e^*}, \text{ct}_{e^*-1}$) as follows:

- Let $\mathbf{k}_{e^*-1} = (\mathbf{pp}, \mathbf{sk}_{e^*-1})$. Note that $\Delta_{e^*} = (\mathbf{pk}, \text{ct})$, where $\text{ct} = \text{KPHE.Enc}(\mathbf{pk}, \mathbf{sk}_{e^*-1})$ computed by the challenger in the KPHE blinding experiment.
- Parse the ciphertext ct_{e^*-1} as

$$(t, (\overline{\mathbf{pk}}_{e^*-1-t}, \overline{\text{ctx}}_{e^*-1-t}), \dots, (\overline{\mathbf{pk}}_{e^*-2}, \overline{\text{ctx}}_{e^*-2}), (\mathbf{pk}_{e^*-1}, \text{ctx}_{e^*-1})),$$

where $\text{ctx}_{e^*-1} = \text{KPHE.Enc}(\mathbf{pk}_{e^*-1}, x)$. Note that if $t = 0$, then $x = \mathbf{m}$ for some message, otherwise $x = \overline{\mathbf{sk}}_{e^*-2}$ that is the KPHE secret key corresponding to $\overline{\mathbf{pk}}_{e^*-2}$.

- Sample a uniform random permutation $\pi : [2n] \rightarrow [2n]$, let $\overline{\mathbf{sk}}_{e^*-1} = \pi(\mathbf{sk}_{e^*-1})$, and sample $\overline{\mathbf{pk}}_{e^*-1} \stackrel{\$}{\leftarrow} \text{KPHE.PKGen}(\mathbf{pp}, \overline{\mathbf{sk}}_{e^*-1})$. Compute $\overline{\text{ctx}}_{e^*-1} \stackrel{\$}{\leftarrow} \text{KPHE.Enc}(\overline{\mathbf{pk}}_{e^*-1}, x)$.
- Send (π_{id}, π) to the challenger in the KPHE blinding experiment and get back $(\mathbf{pk}_{e^*}, \text{ctx}_{e^*})$.
- Let ct_{e^*} be

$$(t+1, (\overline{\mathbf{pk}}_{e^*-1-t}, \overline{\text{ctx}}_{e^*-1-t}), \dots, (\overline{\mathbf{pk}}_{e^*-1}, \overline{\text{ctx}}_{e^*-1}), (\mathbf{pk}_{e^*}, \text{ctx}_{e^*})).$$

Finally, \mathcal{B} outputs whatever \mathcal{A} outputs. Note that if the challenger in the KPHE blinding experiment responds to \mathcal{B} with the output of KPHE.Eval , then the game is identical to $\text{Hyb}'_{0,q-1}$ to \mathcal{A} ; otherwise the game is identical to $\text{Hyb}_{0,q}$ to \mathcal{A} . If \mathcal{A} can distinguish between the two hybrids with non-negligible probability, then \mathcal{B} can break the blinding property of KPHE with non-negligible probability, which leads to contradiction. This concludes our proof. \square

A.2 Proof of Lemma 3.7

We next present the detailed proof of Lemma 3.7.

Proof. Between Hyb_4 and Hyb_5 , we construct a series of intermediate hybrids $\text{Hyb}_{4,\bar{e}}, \text{Hyb}_{4,\bar{e}-1}, \dots, \text{Hyb}_{4,\tilde{e}+1}$ where in each hybrid we change a single Δ_e from $(\widehat{\mathbf{pk}}_e, \text{KPHE.Enc}(\widehat{\mathbf{pk}}_e, \mathbf{sk}_{e-1}))$ to $(\widehat{\mathbf{pk}}_e, \text{KPHE.Enc}(\widehat{\mathbf{pk}}_e, 0^{2n}))$, one by one from \bar{e} down to $\tilde{e}+1$. Note that $\text{Hyb}_4 = \text{Hyb}_{4,\bar{e}+1}$ and $\text{Hyb}_5 = \text{Hyb}_{4,\tilde{e}+1}$. In the following we argue that changing a single Δ_e is computationally indistinguishable to any PPT adversary, namely $\text{Hyb}_{4,e+1} \stackrel{c}{\approx} \text{Hyb}_{4,e}$ for all $\bar{e} \geq e \geq \tilde{e}+1$.

Assume for the purpose of contradiction that there exists a PPT adversary \mathcal{A} that can distinguish between $\text{Hyb}_{4,e+1}$ and $\text{Hyb}_{4,e}$ for some $\bar{e} \geq e \geq \tilde{e}+1$. Then we construct a PPT adversary \mathcal{B} that breaks the leakage-resilient semantic security of KPHE. The adversary \mathcal{B} first receives $(\mathbf{pp}, \mathbf{pk})$ from the challenger in the KPHE semantic security game. Then \mathcal{B} plays the UE game with \mathcal{A} as a challenger in $\text{Hyb}_{4,e+1}$. \mathcal{B} uses \mathbf{pp} to generate UE keys and update tokens as in $\text{Hyb}_{4,e+1}$ except that for epoch e , the UE key \mathbf{k}_e is unknown. In addition, \mathcal{B} sends $(\mathbf{m}_0 = \mathbf{sk}_{e-1}, \mathbf{m}_1 = 0^{2n})$ to the KPHE challenger and gets back a ciphertext ctx . Then \mathcal{B} sets $\Delta_e = (\mathbf{pk}, \text{ctx})$. Note that \mathcal{B} doesn't need to know \mathbf{k}_e because it is never used in $\text{Hyb}_{4,e+1}$ or $\text{Hyb}_{4,e}$. In particular, \mathcal{B} can use \mathbf{pk} to compute all the $\text{KPHE.Enc}(\mathbf{k}_e, \cdot)$. Finally, \mathcal{B} outputs whatever \mathcal{A} outputs.

Note that if the challenger in the KPHE leakage resilience experiment responds to \mathcal{B} with an encryption of m_0 , then the UE game is identical to $\text{Hyb}_{4,e+1}$ to \mathcal{A} ; otherwise the UE game is identical to $\text{Hyb}_{4,e}$ to \mathcal{A} . If \mathcal{A} can distinguish between the two hybrids with non-negligible probability, then \mathcal{B} can break the leakage-resilient semantic security of KPHE with non-negligible probability, which leads to contradiction. This concludes our proof. \square

B Proof of IND-UPD Security (Theorem 3.9)

In this section, we prove Theorem 3.9, which formally establishes the IND-UPD security of our UE construction in Section 3.3.1.

Proof. The proof proceeds via a hybrid argument.

Hyb_0 The challenger plays the real game with the adversary.

Hyb_1 Same as Hyb_0 but for $\text{UE.upd}(\Delta_e, \text{ct}_{e-1})$ in $\mathcal{O}.\text{upd}$ and $\text{UE.upd}(\Delta_e, \tilde{\text{ct}}_{e-1})$ in $\mathcal{O}.\text{next}$, do the following:

- Let $\mathbf{k}_{e-1} = (\text{pp}, \text{sk}_{e-1})$ and $\mathbf{k}_e = (\text{pp}, \text{sk}_e)$.
- Parse the ciphertext ct_{e-1} or $\tilde{\text{ct}}_{e-1}$ as

$$(t, (\overline{\text{pk}}_{e-1-t}, \overline{\text{ctx}}_{e-1-t}), \dots, (\overline{\text{pk}}_{e-2}, \overline{\text{ctx}}_{e-2}), (\text{pk}_{e-1}, \text{ctx}_{e-1})),$$

where $\text{KPHE.Dec}(\overline{\text{sk}}_{e-1-t}, \overline{\text{ctx}}_{e-1-t}) = m$, $\text{KPHE.Dec}(\overline{\text{sk}}_{e-t}, \overline{\text{ctx}}_{e-t}) = \overline{\text{sk}}_{e-1-t}, \dots, \text{KPHE.Dec}(\overline{\text{sk}}_{e-2}, \overline{\text{ctx}}_{e-2}) = \overline{\text{sk}}_{e-3}$, $\text{KPHE.Dec}(\text{sk}_{e-1}, \text{ctx}_{e-1}) = \overline{\text{sk}}_{e-2}$.

- Sample $(t+1)$ uniform random permutations $\pi_{e-1-t}, \dots, \pi_{e-1} : [2n] \rightarrow [2n]$. Also, let $\pi_{\text{id}} : [2n] \rightarrow [2n]$ denote the identity permutation.
- For each $i \in \{e-1-t, \dots, e-2\}$, let $\tilde{\text{sk}}_i = \pi_i(\overline{\text{sk}}_i)$ and sample $\tilde{\text{pk}}_i \xleftarrow{\$} \text{KPHE.PKGen}(\text{pp}, \tilde{\text{sk}}_i)$. Additionally, let $\overline{\text{sk}}_{e-1} = \pi_{e-1}(\text{sk}_{e-1})$ and sample $\overline{\text{pk}}_{e-1} \xleftarrow{\$} \text{KPHE.PKGen}(\text{pp}, \overline{\text{sk}}_{e-1})$. Sample $\text{pk}_e \xleftarrow{\$} \text{KPHE.PKGen}(\text{pp}, \text{sk}_e)$.
- For each $i \in \{e-t, \dots, e-2\}$, compute $\tilde{\text{ctx}}_i \xleftarrow{\$} \text{KPHE.Enc}(\tilde{\text{pk}}_i, \tilde{\text{sk}}_{i-1})$. Additionally, compute

$$\tilde{\text{ctx}}_{e-1-t} \xleftarrow{\$} \text{KPHE.Enc}(\tilde{\text{pk}}_{e-1-t}, m),$$

$$\overline{\text{ctx}}_{e-1} \xleftarrow{\$} \text{KPHE.Enc}(\overline{\text{pk}}_{e-1}, \tilde{\text{sk}}_{e-2}),$$

$$\text{ctx}_e \xleftarrow{\$} \text{KPHE.Enc}(\text{pk}_e, \overline{\text{sk}}_{e-1}),$$

- Let ct_e or $\tilde{\text{ct}}_e$ be

$$(t+1, (\tilde{\text{pk}}_{e-1-t}, \tilde{\text{ctx}}_{e-1-t}), \dots, (\tilde{\text{pk}}_{e-2}, \tilde{\text{ctx}}_{e-2}), (\overline{\text{pk}}_{e-1}, \overline{\text{ctx}}_{e-1}), (\text{pk}_e, \text{ctx}_e)).$$

This hybrid is computationally indistinguishable from Hyb_0 to any PPT adversary by the blinding property of KPHE. We omit the detailed proof here, but it follows similarly to the proof of Lemma 3.6.

Hyb₂ Same as Hyb₁ but for $\text{UE.upd}(\Delta_e, \text{ct}_{e-1})$ in $\mathcal{O}.\text{upd}$ and $\text{UE.upd}(\Delta_e, \tilde{\text{ct}}_{e-1})$ in $\mathcal{O}.\text{next}$, sample each $\tilde{\text{sk}}_{e-1-t}, \dots, \tilde{\text{sk}}_{e-2}, \overline{\text{sk}}_{e-1}$ from the distribution \mathcal{U}_n . This hybrid is statistically identical to Hyb₁.

Hyb₃ Let \tilde{e} be the challenge epoch, and let \bar{e} be the last epoch where the adversary corrupts continuous update tokens from \tilde{e} , namely the adversary corrupts $\Delta_{\tilde{e}+1}, \Delta_{\tilde{e}+2}, \dots, \Delta_{\bar{e}}$ but not $\Delta_{\bar{e}+1}$. This hybrid is the same as Hyb₃ except that the challenger guesses \tilde{e}^* and \bar{e}^* at the beginning of the game and aborts the game if guessing incorrectly. Let E be the upper bound on the number of epochs during the game. If the challenger does not abort, then this hybrid is identical to Hyb₂, which happens with probability at least $\frac{1}{E^2}$. In the remaining hybrids, we assume for simplicity that the challenger guesses \tilde{e} and \bar{e} correctly.

Hyb₄ Same as Hyb₃ except that for each $\mathbf{k}_e = (\text{pp}, \text{sk}_e)$, generate a single public key $\widehat{\text{pk}}_e \stackrel{\$}{\leftarrow} \text{KPHE.PKGen}(\text{pp}, \text{sk}_e)$. Then whenever $\text{KPHE.Enc}(\text{pk}_e, x)$ is computed for a freshly generated pk_e and some x , compute it as $\text{KPHE.Eval}(\widehat{\text{pk}}_e, \text{KPHE.Enc}(\widehat{\text{pk}}_e, x), \pi_{\text{id}}, \pi_{\text{id}})$. That is, instead of generating a fresh pk_e from sk_e every time, use the same $\widehat{\text{pk}}_e$ to encrypt x and use then KPHE.Eval to re-randomize it.

This hybrid is computationally indistinguishable from Hyb₃ by the blinding property of KPHE. We omit the detailed reduction here, but it is similar to the reduction in the proof of Lemma 3.6.

Hyb₅ Same as Hyb₄ except that for all $\tilde{e}+1 \leq e \leq \bar{e}$, $\text{UE.next}(\mathbf{k}_{e-1})$ is computed as follows. Generate $\text{sk}_e \stackrel{\$}{\leftarrow} \text{KPHE.SKGen}(\text{pp}, \mathcal{U}_n)$ and let $\widehat{\text{pk}}_e \stackrel{\$}{\leftarrow} \text{KPHE.PKGen}(\text{pp}, \text{sk}_e)$ be the single public key for \mathbf{k}_e (that will be used for every KPHE.Enc). Output

$$\mathbf{k}_e = (\text{pp}, \text{sk}_e), \quad \Delta_e = (\widehat{\text{pk}}_e, \text{KPHE.Enc}(\widehat{\text{pk}}_e, 0^{2n})).$$

This hybrid is computationally indistinguishable from Hyb₄ based on the distributional semantic security of KPHE. We omit the detailed proof here, but it follows similarly to the proof of Lemma 3.7.

Hyb₆ Same as Hyb₅ except that for each $\mathbf{k}_e = (\text{pp}, \text{sk}_e)$, generate a single public key $\widehat{\text{pk}}_e \stackrel{\$}{\leftarrow} \text{KPHE.PKGen}(\text{pp}, \text{sk}_e)$ and use $\text{KPHE.Eval}(\widehat{\text{pk}}_e, \text{KPHE.Enc}(\widehat{\text{pk}}_e, \cdot), \pi_{\text{id}}, \pi_{\text{id}})$ for all the computation of $\text{KPHE.Enc}(\mathbf{k}_e, \cdot)$ (including the computation of Δ_e). The only exception is in the computation of the challenge ciphertext $\tilde{\text{ct}}_{\tilde{e}} \stackrel{\$}{\leftarrow} \text{UE.upd}(\Delta_{\tilde{e}}, \overline{\text{ct}}_b)$, which is in the form of

$$\tilde{\text{ct}}_{\tilde{e}} = (t+1, (\tilde{\text{pk}}_{\tilde{e}-1-t}, \tilde{\text{ctx}}_{\tilde{e}-1-t}), \dots, (\tilde{\text{pk}}_{\tilde{e}-2}, \tilde{\text{ctx}}_{\tilde{e}-2}), (\overline{\text{pk}}_{\tilde{e}-1}, \overline{\text{ctx}}_{\tilde{e}-1})(\text{pk}_{\tilde{e}}, \text{ctx}_{\tilde{e}})),$$

where the last ciphertext is computed from $\widehat{\text{pk}}_{\tilde{e}}$ directly, namely $\text{pk}_{\tilde{e}} = \widehat{\text{pk}}_{\tilde{e}}$ and $\text{ctx}_{\tilde{e}} \stackrel{\$}{\leftarrow} \text{KPHE.Enc}(\widehat{\text{pk}}_{\tilde{e}}, \overline{\text{sk}}_{\tilde{e}-1})$.

This hybrid is computationally indistinguishable from Hyb₅ by the blinding property of KPHE. We omit the detailed reduction here, but it is similar to the reduction in the proof of Lemma 3.6.

Hyb₇ Same as Hyb₆ but the challenge ciphertext $\tilde{\text{ct}}_{\tilde{e}}$ is computed as follows.

- Let $\mathbf{k}_{\tilde{e}} = (\mathbf{pp}, \mathbf{sk}_{\tilde{e}})$ and let $\widehat{\mathbf{pk}}_{\tilde{e}}$ be its corresponding public key.
- Parse the ciphertext $\overline{\mathbf{ct}}_0$ as

$$(t, (\overline{\mathbf{pk}}_{\tilde{e}-1-t}, \overline{\mathbf{ctx}}_{\tilde{e}-1-t}), \dots, (\overline{\mathbf{pk}}_{\tilde{e}-2}, \overline{\mathbf{ctx}}_{\tilde{e}-2}), (\mathbf{pk}_{\tilde{e}-1}, \mathbf{ctx}_{\tilde{e}-1})),$$

Note that $\overline{\mathbf{ct}}_1$ is of the same form because $|\overline{\mathbf{ct}}_0| = |\overline{\mathbf{ct}}_1|$.

- For each $i \in \{\tilde{e}-1-t, \dots, \tilde{e}-1\}$, sample $\tilde{\mathbf{sk}}_i$ from \mathcal{U}_n and sample $\tilde{\mathbf{pk}}_i \xleftarrow{\$} \text{KPHE.PKGen}(\mathbf{pp}, \tilde{\mathbf{sk}}_i)$.
- For each $i \in \{\tilde{e}-1-t, \dots, \tilde{e}-1\}$, compute $\tilde{\mathbf{ctx}}_i \xleftarrow{\$} \text{KPHE.Enc}(\tilde{\mathbf{pk}}_i, 0^{2n})$. Additionally, compute $\mathbf{ctx}_{\tilde{e}} \xleftarrow{\$} \text{KPHE.Enc}(\widehat{\mathbf{pk}}_{\tilde{e}}, 0^{2n})$.
- Let the challenge ciphertext $\tilde{\mathbf{ct}}_{\tilde{e}}$ be

$$(t+1, (\tilde{\mathbf{pk}}_{\tilde{e}-1-t}, \tilde{\mathbf{ctx}}_{\tilde{e}-1-t}), \dots, (\tilde{\mathbf{pk}}_{\tilde{e}-1}, \tilde{\mathbf{ctx}}_{\tilde{e}-1}), (\widehat{\mathbf{pk}}_{\tilde{e}}, \mathbf{ctx}_{\tilde{e}})).$$

We prove in Lemma B.1 that this hybrid is computationally indistinguishable from Hyb_6 based on the distributional semantic security of KPHE.

Notice that in the final hybrid Hyb_7 , $\text{UE.upd}(\Delta_{\tilde{e}}, \overline{\mathbf{ct}}_b)$ is computed in the exact same way for $b = 0$ and $b = 1$. This concludes our proof.

Lemma B.1. $\text{Hyb}_6 \stackrel{c}{\approx} \text{Hyb}_7$ in the proof of Theorem 3.9.

Proof. Between Hyb_6 and Hyb_7 , we construct a series of intermediate hybrids $\text{Hyb}_{6,\tilde{e}}, \text{Hyb}_{6,\tilde{e}-1}, \dots, \text{Hyb}_{6,\tilde{e}-1-t}$ where in each hybrid we change a single \mathbf{ctx} (in $\tilde{\mathbf{ct}}_{\tilde{e}}$) to a KPHE encryption of 0^{2n} , one by one from \tilde{e} down to $\tilde{e}-1-t$. That is, changing $\mathbf{ctx}_{\tilde{e}}, \overline{\mathbf{ctx}}_{\tilde{e}-1}, \tilde{\mathbf{ctx}}_{\tilde{e}-2}, \dots, \tilde{\mathbf{ctx}}_{\tilde{e}-1-t}$ to encryptions of 0^{2n} one by one in each hybrid. Note that $\text{Hyb}_6 = \text{Hyb}_{6,\tilde{e}+1}$ and $\text{Hyb}_7 = \text{Hyb}_{6,\tilde{e}-1-t}$. In the following we argue that changing a single \mathbf{ctx} is computationally indistinguishable to any PPT adversary, namely $\text{Hyb}_{6,\tilde{e}+1} \stackrel{c}{\approx} \text{Hyb}_{6,\tilde{e}}$ for all $\tilde{e} \geq \tilde{e} \geq \tilde{e}-1-t$.

$\text{Hyb}_{6,\tilde{e}+1} \stackrel{c}{\approx} \text{Hyb}_{6,\tilde{e}}$. Assume for the purpose of contradiction that there exists a PPT adversary \mathcal{A} that can distinguish between $\text{Hyb}_{6,\tilde{e}+1} (= \text{Hyb}_6)$ and $\text{Hyb}_{6,\tilde{e}}$. Then we construct a PPT adversary \mathcal{B} that breaks the distributional semantic security of KPHE. The adversary \mathcal{B} first receives $(\mathbf{pp}, \mathbf{pk})$ from the challenger in the KPHE semantic security game. Then \mathcal{B} plays the UE game with \mathcal{A} as a challenger in Hyb_6 . \mathcal{B} uses \mathbf{pp} to generate UE keys and update tokens as in Hyb_6 except that for epoch \tilde{e} , the UE key $\mathbf{k}_{\tilde{e}}$ is unknown, but \mathcal{B} uses \mathbf{pk} as $\widehat{\mathbf{pk}}_{\tilde{e}}$ for all the computation of $\text{KPHE.Enc}(\mathbf{k}_{\tilde{e}}, \cdot)$. In addition, the challenge ciphertext $\tilde{\mathbf{ct}}_{\tilde{e}}$ is computed as follows.

- Parse the ciphertext $\overline{\mathbf{ct}}_0$ as

$$(t, (\overline{\mathbf{pk}}_{\tilde{e}-1-t}, \overline{\mathbf{ctx}}_{\tilde{e}-1-t}), \dots, (\overline{\mathbf{pk}}_{\tilde{e}-2}, \overline{\mathbf{ctx}}_{\tilde{e}-2}), (\mathbf{pk}_{\tilde{e}-1}, \mathbf{ctx}_{\tilde{e}-1})),$$

- For each $i \in \{\tilde{e}-1-t, \dots, \tilde{e}-1\}$, sample $\tilde{\mathbf{sk}}_i$ from \mathcal{U}_n and sample $\tilde{\mathbf{pk}}_i \xleftarrow{\$} \text{KPHE.PKGen}(\mathbf{pp}, \tilde{\mathbf{sk}}_i)$.
- For each $i \in \{\tilde{e}-1-t, \dots, \tilde{e}-1\}$, compute $\tilde{\mathbf{ctx}}_i$ same as in Hyb_6 .
- Send $(\mathbf{m}_0 = \tilde{\mathbf{sk}}_{\tilde{e}-1}, \mathbf{m}_1 = 0^{2n})$ to the KPHE challenger and get back \mathbf{ctx} , and set $\mathbf{ctx}_{\tilde{e}} := \mathbf{ctx}$.

- Let the challenge ciphertext $\widetilde{\text{ct}}_{\tilde{e}}$ be

$$(t + 1, (\widetilde{\text{pk}}_{\tilde{e}-1-t}, \widetilde{\text{ctx}}_{\tilde{e}-1-t}), \dots, (\widetilde{\text{pk}}_{\tilde{e}-1}, \widetilde{\text{ctx}}_{\tilde{e}-1}), (\widehat{\text{pk}}_{\tilde{e}}, \text{ctx}_{\tilde{e}})).$$

Note that \mathcal{B} doesn't need to know $\mathbf{k}_{\tilde{e}}$ because it is never used in Hyb_6 or $\text{Hyb}_{6,\tilde{e}}$. In particular, \mathcal{B} can use pk to compute all the $\text{KPHE.Enc}(\mathbf{k}_{\tilde{e}}, \cdot)$. Finally, \mathcal{B} outputs whatever \mathcal{A} outputs.

Note that if the challenger in the KPHE leakage resilience experiment responds to \mathcal{B} with an encryption of \mathbf{m}_0 , then the UE game is identical to Hyb_6 to \mathcal{A} ; otherwise the UE game is identical to $\text{Hyb}_{6,\tilde{e}}$ to \mathcal{A} . If \mathcal{A} can distinguish between the two hybrids with non-negligible probability, then \mathcal{B} can break the distributional semantic security of KPHE with non-negligible probability, which leads to contradiction.

$\text{Hyb}_{6,e+1} \stackrel{c}{\approx} \text{Hyb}_{6,e}$ **for all** $\tilde{e} - 1 \geq e \geq \tilde{e} - 1 - t$ Assume for the purpose of contradiction that there exists a PPT adversary \mathcal{A} that can distinguish between $\text{Hyb}_{6,e+1}$ and $\text{Hyb}_{6,e}$ for some $\tilde{e} - 1 \geq e \geq \tilde{e} - 1 - t$. Then we construct a PPT adversary \mathcal{B} that breaks the distributional semantic security of KPHE. The adversary \mathcal{B} first receives (pp, pk) from the challenger in the KPHE semantic security game. Then \mathcal{B} plays the UE game with \mathcal{A} as a challenger in $\text{Hyb}_{6,e+1}$. \mathcal{B} uses pp to generate UE keys and update tokens as in $\text{Hyb}_{6,e+1}$ except that the challenge ciphertext $\widetilde{\text{ct}}_{\tilde{e}}$ is computed as follows.

- Let $\mathbf{k}_{\tilde{e}} = (\text{pp}, \text{sk}_{\tilde{e}})$ and let $\widehat{\text{pk}}_{\tilde{e}}$ be its corresponding public key.

- Parse the ciphertext $\overline{\text{ct}}_0$ as

$$(t, (\overline{\text{pk}}_{\tilde{e}-1-t}, \overline{\text{ctx}}_{\tilde{e}-1-t}), \dots, (\overline{\text{pk}}_{\tilde{e}-2}, \overline{\text{ctx}}_{\tilde{e}-2}), (\text{pk}_{\tilde{e}-1}, \text{ctx}_{\tilde{e}-1})),$$

- For each $i \in \{\tilde{e}-1-t, \dots, \tilde{e}-1\} \setminus \{e\}$, sample $\widetilde{\text{sk}}_i$ from \mathcal{U}_n and sample $\widetilde{\text{pk}}_i \stackrel{\$}{\leftarrow} \text{KPHE.PKGen}(\text{pp}, \widetilde{\text{sk}}_i)$. Set $\widetilde{\text{pk}}_i := \text{pk}$ (received from the KPHE challenger).

- For each $i \in \{\tilde{e} - 1 - t, \dots, e - 1\}$, compute $\widetilde{\text{ctx}}_i \stackrel{\$}{\leftarrow} \text{KPHE.Enc}(\widetilde{\text{pk}}_i, \widetilde{\text{sk}}_{i-1})$. For each $i \in \{e + 1, \dots, \tilde{e} - 1\}$, compute $\widetilde{\text{ctx}}_i \stackrel{\$}{\leftarrow} \text{KPHE.Enc}(\widetilde{\text{pk}}_i, 0^{2n})$. Additionally, compute $\text{ctx}_{\tilde{e}} \stackrel{\$}{\leftarrow} \text{KPHE.Enc}(\widehat{\text{pk}}_{\tilde{e}}, 0^{2n})$.

- Send $(\mathbf{m}_0 = \widetilde{\text{sk}}_{e-1}, \mathbf{m}_1 = 0^{2n})$ to the KPHE challenger and get back a ciphertext ctx , and let $\widetilde{\text{ctx}}_e := \text{ctx}$.

- Let the challenge ciphertext $\widetilde{\text{ct}}_{\tilde{e}}$ be

$$(t + 1, (\widetilde{\text{pk}}_{\tilde{e}-1-t}, \widetilde{\text{ctx}}_{\tilde{e}-1-t}), \dots, (\widetilde{\text{pk}}_{\tilde{e}-1}, \widetilde{\text{ctx}}_{\tilde{e}-1}), (\widehat{\text{pk}}_{\tilde{e}}, \text{ctx}_{\tilde{e}})).$$

Finally, \mathcal{B} outputs whatever \mathcal{A} outputs. Note that if the challenger in the KPHE leakage resilience experiment responds to \mathcal{B} with an encryption of \mathbf{m}_0 , then the UE game is identical to $\text{Hyb}_{6,e+1}$ to \mathcal{A} ; otherwise the UE game is identical to $\text{Hyb}_{6,e}$ to \mathcal{A} . If \mathcal{A} can distinguish between the two hybrids with non-negligible probability, then \mathcal{B} can break the distributional semantic security of KPHE with non-negligible probability, which leads to contradiction. □

This concludes the proof of Theorem 3.9. □

C Proof of IND-HRA Security (Theorem 4.6)

In this section, we prove Theorem 4.6, which establishes the IND-HRA security for our PRE construction in Section 4.

Proof. The proof follows a sequence of hybrids as outlined below. We use $[N]$ to denote the set of all keys (honest + corrupt), $\mathcal{K}_{\text{Honest}} \subset [N]$ to denote the set of honest keys, and $\mathcal{K}_{\text{Corrupt}} \subset [N]$ to denote the set of corrupt keys. These sets are fixed (adversarially) at the beginning of the game, before the adversary is allowed to issue re-key generation queries (this remains unchanged in each hybrid).

Hyb₀: This hybrid is identical to the real IND-HRA security game between the challenger and the adversary.

Hyb₁: This hybrid is identical to the hybrid **Hyb₀** except that the challenger locally maintains two additional tables T_0 and T_1 (initially empty), and does the following:

- For each honest encryption query of the form $\text{Enc}(i, \bar{m})$, the challenger adds to the local table T_0 an entry of the form (i, \bar{ct}, \bar{m}) , where \bar{ct} is generated as $\bar{ct} = \text{KPHE.Enc}(\text{pk}_i, \bar{m})$.
- For each honest re-encryption query of the form $\text{ReEnc}(i, j, \bar{ct}_i)$ to which the challenger does not respond with \perp , it fetches the entry (i, \bar{ct}_i) (from either \mathcal{L} or $\tilde{\mathcal{L}}$) and proceeds as follows.

Suppose \bar{ct}_i is of the form:

$$\bar{ct}_i = (0, (\text{pk}_i, \widehat{\text{ctx}}_0)).$$

Recover the plaintext message $\bar{m} = \text{KPHE.Dec}(\text{sk}_i, \widehat{\text{ctx}}_0)$. Next, sample a uniform permutation $\pi : [2n] \rightarrow [2n]$ and set

$$\bar{\text{sk}}_0 = \pi(\text{sk}_i), \quad \bar{\text{pk}}_0 = \text{KPHE.PKGen}(\text{pp}, \bar{\text{sk}}_0).$$

Also set

$$\begin{aligned} \bar{\text{ctx}}_0 &\stackrel{\$}{\leftarrow} \text{KPHE.Enc}(\bar{\text{pk}}_0, \bar{m}), & \bar{\text{ctx}}_1 &\stackrel{\$}{\leftarrow} \text{KPHE.Enc}(\text{pk}_j, \bar{\text{sk}}_0), \\ (\widehat{\text{pk}}_1, \widehat{\text{ctx}}_1) &\stackrel{\$}{\leftarrow} \text{KPHE.Eval}(\text{pk}_j, \bar{\text{ctx}}_1, \pi_{\text{id}}, \pi_{\text{id}}), \end{aligned}$$

where $\pi_{\text{id}} : [2n] \rightarrow [2n]$ is the identity permutation. Output the re-encrypted ciphertext as:

$$\bar{ct}_j = (1, (\bar{\text{pk}}_0, \bar{\text{ctx}}_0), (\widehat{\text{pk}}_1, \widehat{\text{ctx}}_1)).$$

Also, add to the local table T_1 an entry of the form $(j, \bar{ct}_j, \bar{\text{sk}}_0)$.

Alternatively, suppose \bar{ct}_i is of the form (for some $t > 0$):

$$\bar{ct}_i = (t, (\bar{\text{pk}}_0, \bar{\text{ctx}}_0), \dots, (\bar{\text{pk}}_{t-1}, \bar{\text{ctx}}_{t-1}), (\widehat{\text{pk}}_t, \widehat{\text{ctx}}_t)).$$

Recover the intermediate secret key $\bar{\text{sk}}_{t-1} = \text{KPHE.Dec}(\text{sk}_i, \bar{\text{ctx}}_t)$. Next, sample a uniform permutation $\pi_t : [2n] \rightarrow [2n]$ and set

$$\bar{\text{sk}}_t = \pi_t(\text{sk}_i), \quad \bar{\text{pk}}_t \stackrel{\$}{\leftarrow} \text{KPHE.PKGen}(\text{pp}, \bar{\text{sk}}_t).$$

Also set

$$\begin{aligned} \overline{\text{ctx}}_t &\stackrel{\$}{\leftarrow} \text{KPHE.Enc}(\overline{\text{pk}}_t, \overline{\text{sk}}_{t-1}), & \overline{\text{ctx}}_{t+1} &\stackrel{\$}{\leftarrow} \text{KPHE.Enc}(\text{pk}_j, \overline{\text{sk}}_t), \\ & & (\widehat{\text{pk}}_{t+1}, \widehat{\text{ctx}}_{t+1}) &\stackrel{\$}{\leftarrow} \text{KPHE.Eval}(\text{pk}_j, \overline{\text{ctx}}_{t+1}, \pi_{\text{id}}, \pi_{\text{id}}), \end{aligned}$$

where $\pi_{\text{id}} : [2n] \rightarrow [2n]$ is the identity permutation. Output the re-encrypted ciphertext as:

$$\overline{\text{ct}}_j = ((t+1), (\overline{\text{pk}}_0, \overline{\text{ctx}}_0), \dots, (\overline{\text{pk}}_{t-1}, \overline{\text{ctx}}_{t-1}), (\overline{\text{pk}}_t, \overline{\text{ctx}}_t), (\widehat{\text{pk}}_{t+1}, \widehat{\text{ctx}}_{t+1})).$$

Also, add to the local table T_1 an entry of the form $(j, \overline{\text{ct}}_j, \overline{\text{sk}}_t)$.

The view of the IND-HRA adversary in this hybrid is computationally indistinguishable from that in the hybrid Hyb_0 under the assumption the KPHE scheme satisfies public key and ciphertext blinding. We prove this formally in Lemma C.4.

Hyb₂: This hybrid is identical to the hybrid Hyb_1 except that the challenger does the following (it still locally maintains the tables T_0 and T_1 as in Hyb_1): for each (valid) honest re-encryption query of the form $\text{ReEnc}(i, j, \overline{\text{ct}}_i)$, suppose $\overline{\text{ct}}_i$ is of the form:

$$\overline{\text{ct}}_i = (0, (\text{pk}_i, \widehat{\text{ctx}}_0)).$$

Look up the local table T_0 for an entry of the form $(i, \overline{\text{ct}}_i, \overline{\text{m}}^*)$. Such an entry is guaranteed to exist. Set $\overline{\text{m}} = \overline{\text{m}}^*$. The rest of the simulation proceeds as in hybrid Hyb_1 .

Alternatively, suppose $\overline{\text{ct}}_i$ is of the form (for some $t > 0$):

$$\overline{\text{ct}}_i = (t, (\overline{\text{pk}}_0, \overline{\text{ctx}}_0), \dots, (\overline{\text{pk}}_{t-1}, \overline{\text{ctx}}_{t-1}), (\widehat{\text{pk}}_t, \widehat{\text{ctx}}_t)).$$

Look up the local table T_1 for an entry of the form $(i, \overline{\text{ct}}_i, \text{sk}^*)$. Such an entry is guaranteed to exist. Set the intermediate secret key $\overline{\text{sk}}_{t-1} = \text{sk}^*$. The rest of the simulation proceeds as in hybrid Hyb_1 .

It is easy to see that, assuming that the KPHE scheme is correct with overwhelmingly large probability, the view of the IND-HRA adversary in this hybrid is statistically indistinguishable from that in the hybrid Hyb_1 .

Hyb₃: This hybrid is identical to the hybrid Hyb_2 except that the challenger does the following (it still locally maintains the tables T_0 and T_1 as in Hyb_2): for each (valid) honest re-encryption query of the form $\text{ReEnc}(i, j, \overline{\text{ct}}_i)$, suppose $\overline{\text{ct}}_i$ is of the form:

$$\overline{\text{ct}}_i = (0, (\text{pk}_i, \widehat{\text{ctx}}_0)).$$

Look up the local table T_0 for an entry of the form $(i, \overline{\text{ct}}_i, \overline{\text{m}}^*)$. Such an entry is guaranteed to exist. Set $\overline{\text{m}} = \overline{\text{m}}^*$. Next, set

$$\overline{\text{sk}}_0 \stackrel{\$}{\leftarrow} \mathcal{U}_n, \quad \overline{\text{pk}}_0 = \text{KPHE.PKGen}(\text{pp}, \overline{\text{sk}}_0).$$

The rest of the simulation proceeds as in hybrid Hyb_2 .

Alternatively, suppose $\overline{\text{ct}}_i$ is of the form (for some $t > 0$):

$$\overline{\text{ct}}_i = (t, (\overline{\text{pk}}_0, \overline{\text{ctx}}_0), \dots, (\overline{\text{pk}}_{t-1}, \overline{\text{ctx}}_{t-1}), (\widehat{\text{pk}}_t, \widehat{\text{ctx}}_t)).$$

Look up the local table T_1 for an entry of the form $(i, \overline{ct}_i, sk^*)$. Such an entry is guaranteed to exist. Set the intermediate secret key $\overline{sk}_{t-1} = sk^*$. Next, set

$$\overline{sk}_t \stackrel{\$}{\leftarrow} \mathcal{U}_n, \quad \overline{pk}_t = \text{KPHE.PKGen}(\text{pp}, \overline{sk}_t).$$

The rest of the simulation proceeds as in hybrid Hyb_2 .

It is easy to see that the view of the IND-HRA adversary in this hybrid is identical to that in the hybrid Hyb_2 , since the distributions of \overline{sk}_0 and \overline{pk}_0 (resp., \overline{sk}_t and \overline{pk}_t) remain unchanged from hybrid Hyb_2 .

Remark C.1. *Note that in Hyb_3 , the set of secret keys $\{sk_i\}_{i \in [N]}$ (including both corrupt and honest secret keys) is no longer used by the challenger when answering ReEnc queries. In particular, all ReEnc queries are answered using only the knowledge of the public keys $\{pk_i\}_{i \in [N]}$, and the local tables T_0 and T_1 maintained by the challenger.*

Hyb₄: This hybrid is identical to the hybrid Hyb_3 except for the manner in which the challenger answers the ReKeyGen queries issued by the adversary. In particular, given a query of the form $\text{ReKeyGen}(i, j)$ such that $i \in \mathcal{K}_{\text{Honest}}$ and $j \in \mathcal{K}_{\text{Honest}}$, the challenger responds with $rk_{i,j} = (pk_j, \text{KPHE.Enc}(pk_j, 0^{2n}))$.

We argue that Hyb_4 is indistinguishable from Hyb_3 in a straightforward manner under the assumption that KPHE satisfies distributional circular security. This argument relies *crucially* on the fact that, as in Hyb_3 , Hyb_4 allows the challenger to answer all ReEnc queries using only the knowledge of the public keys $\{pk_i\}_{i \in [N]}$, and the local tables T_0 and T_1 maintained by the challenger. We prove this formally in Lemma C.5.

Remark C.2. *Note that in Hyb_4 , the set of honest secret keys $\{sk_i\}_{i \in \mathcal{K}_{\text{Honest}}}$ is no longer used by the challenger when answering ReKeyGen queries. Also, as in Hyb_3 , the set of honest secret keys $\{sk_i\}_{i \in \mathcal{K}_{\text{Honest}}}$ are also not used by the challenger when answering ReEnc queries. In other words, Hyb_4 allows the challenger to entirely “forget” the set of honest secret keys.*

Hyb₅: This hybrid is identical to the hybrid Hyb_4 except for the manner in which the challenger answers the challenge encryption query issued by the adversary. In particular, given a query of the form $\mathcal{O}.\text{chall}(i, \overline{m}_0, \overline{m}_1)$, the challenger proceeds as follows:

- If $i \in \mathcal{K}_{\text{Corrupt}}$, respond with $\overline{ct}^* = \perp$ (this is exactly as in Hyb_4).
- If $i \in \mathcal{K}_{\text{Honest}}$, respond with $\overline{ct}^* = (0, \text{KPHE.Enc}(pk_i, 0^{2n}))$.

We argue that Hyb_5 is indistinguishable from Hyb_4 under the assumption that KPHE satisfies distributional semantic security. This argument relies *crucially* on the fact that, as in Hyb_4 , Hyb_5 allows the challenger to answer all ReKeyGen queries and ReEnc queries using only the knowledge of the corrupt secret keys $\{sk_j\}_{j \in \mathcal{K}_{\text{Corrupt}}}$, the set of all public keys $\{pk_i\}_{i \in [N]}$, and the local tables T_0 and T_1 maintained by the challenger. We prove this formally in Lemma C.6.

Remark C.3. *Note that in Hyb_5 , the challenge ciphertext \overline{ct}^* is independent of $(\overline{m}_0, \overline{m}_1)$, and hence the adversary’s advantage in winning the IND-HRA game is zero.*

Lemma C.4. $\text{Hyb}_0 \stackrel{c}{\approx} \text{Hyb}_1$ in the proof of Theorem 4.6.

Lemma C.5. $\text{Hyb}_3 \stackrel{c}{\approx} \text{Hyb}_4$ in the proof of Theorem 4.6.

Lemma C.6. $\text{Hyb}_4 \stackrel{c}{\approx} \text{Hyb}_5$ in the proof of Theorem 4.6.

In what follows, we prove these lemmas formally. These proofs complete the overall proof of Theorem 4.6. \square

C.1 Proof of Lemma C.4

We first present the detailed proof of Lemma C.4. The proof is essentially identical to the proof of Lemma 3.6 in the proof of IND-ENC security of our UE construction, with only minor syntactic changes to account for the differences between the UE and the PRE schemes. Let Q be the upper bound on the total number of $\text{ReEnc}(i, j, \overline{\text{ct}}_i)$ queries issued by the adversary \mathcal{A} in hybrid Hyb_0 . We construct a series of Q intermediate hybrids $\{\text{Hyb}_{0,q}\}_{q \in [Q]}$ where in each hybrid we only change the q -th output of ReEnc (one by one from the first to the last) as described in Hyb_1 . Note that $\text{Hyb}_0 = \text{Hyb}_{0,0}$ and $\text{Hyb}_1 = \text{Hyb}_{0,Q}$. In the following we argue that changing the q -th ReEnc output is computationally indistinguishable to any PPT adversary, namely $\text{Hyb}_{0,q-1} \stackrel{c}{\approx} \text{Hyb}_{0,q}$ for all $q \in [Q]$.

Intermediate Hybrid. To argue that $\text{Hyb}_{0,q-1} \stackrel{c}{\approx} \text{Hyb}_{0,q}$ for all $q \in [Q]$, we first define an intermediate hybrid $\text{Hyb}'_{0,q-1}$ as follows: this hybrid is identical to $\text{Hyb}_{0,q-1}$, except that the challenger answers the q -th ReEnc query as follows: suppose that the q -th honest re-encryption query of the form $\text{ReEnc}(i, j, \overline{\text{ct}}_i)$ to which the challenger does not respond with \perp . The challenger fetches the entry $(i, \overline{\text{ct}}_i)$ (from either \mathcal{L} or $\tilde{\mathcal{L}}$) and proceeds as follows.

- Suppose $\overline{\text{ct}}_i$ is of the form:

$$\overline{\text{ct}}_i = (0, (\overline{\text{pk}}_i, \widehat{\text{ctx}}_0)).$$

Recover the plaintext message $\overline{\text{m}} = \text{KPHE.Dec}(\text{sk}_i, \widehat{\text{ctx}}_0)$, sample a uniform permutation $\pi : [2n] \rightarrow [2n]$ and do the following:

- Set $(\overline{\text{pk}}_0, \overline{\text{ctx}}_0)$ as in hybrid $\text{Hyb}_{0,q}$.
- Set $(\widehat{\text{pk}}_1, \widehat{\text{ctx}}_1)$ as in hybrid $\text{Hyb}_{0,q-1}$.
- Output the re-encrypted ciphertext as:

$$\overline{\text{ct}}_j = (1, (\overline{\text{pk}}_0, \overline{\text{ctx}}_0), (\widehat{\text{pk}}_1, \widehat{\text{ctx}}_1)).$$

- Alternatively, suppose $\overline{\text{ct}}_i$ is of the form (for some $t > 0$):

$$\overline{\text{ct}}_i = (t, (\overline{\text{pk}}_0, \overline{\text{ctx}}_0), \dots, (\overline{\text{pk}}_{t-1}, \overline{\text{ctx}}_{t-1}), (\widehat{\text{pk}}_t, \widehat{\text{ctx}}_t)).$$

Recover the intermediate secret key $\overline{\text{sk}}_{t-1} = \text{KPHE.Dec}(\text{sk}_i, \overline{\text{ctx}}_t)$, sample a uniform permutation $\pi_t : [2n] \rightarrow [2n]$ and do the following:

- Set $(\overline{\text{pk}}_t, \overline{\text{ctx}}_t)$ as in hybrid $\text{Hyb}_{0,q}$.
- Set $(\widehat{\text{pk}}_{t+1}, \widehat{\text{ctx}}_{t+1})$ as in hybrid $\text{Hyb}_{0,q-1}$.
- Output the re-encrypted ciphertext as:

$$\overline{\text{ct}}_j = ((t+1), (\overline{\text{pk}}_0, \overline{\text{ctx}}_0), \dots, (\overline{\text{pk}}_{t-1}, \overline{\text{ctx}}_{t-1}), (\overline{\text{pk}}_t, \overline{\text{ctx}}_t), (\widehat{\text{pk}}_{t+1}, \widehat{\text{ctx}}_{t+1})).$$

$\text{Hyb}_{0,q-1} \stackrel{c}{\approx} \text{Hyb}'_{0,q-1}$. We argue that $\text{Hyb}_{0,q-1} \stackrel{c}{\approx} \text{Hyb}'_{0,q-1}$ for all $q \in [Q]$.

Proof. Suppose that there exists some PPT adversary \mathcal{A} that can efficiently distinguish $\text{Hyb}_{0,q-1}$ and $\text{Hyb}'_{0,q-1}$ with non-negligible advantage ϵ . We construct a PPT algorithm \mathcal{B} that breaks the blinding security of the KPHE scheme with non-negligible advantage ϵ' . The algorithm \mathcal{B} receives from the challenger in the blinding security game the public parameters pp , a secret key sk^* , and a public key pk^* , and proceeds as follows.

- \mathcal{B} uses pp to set up the PRE keys (both honest and corrupt) to be provided to the adversary \mathcal{A} exactly as in Hyb_4 except that it sets $\text{sk}_i := \text{sk}^*$ and $\text{pk}_i := \text{pk}^*$ for some uniform $i \xleftarrow{\$} \mathcal{K}_{\text{Honest}}$. Since pk^* is uniformly random, the view of the adversary \mathcal{A} with respect to the distribution of the keys (both honest and corrupt) remains identical.
- Suppose that the q -th honest re-encryption query of the form $\text{ReEnc}(i', j, \overline{\text{ct}}_{i'})$ to which the challenger does not respond with \perp . If $i \neq i'$, \mathcal{B} outputs \perp . Otherwise, it fetches the entry $(i, \overline{\text{ct}}_i)$ (from either \mathcal{L} or $\tilde{\mathcal{L}}$) and proceeds as follows.
- Suppose $\overline{\text{ct}}_i$ is of the form:

$$\overline{\text{ct}}_i = (0, (\text{pk}_i, \widehat{\text{ctx}}_0)).$$

Recover the plaintext message $\overline{\text{m}} = \text{KPHE.Dec}(\text{sk}_i, \widehat{\text{ctx}}_0)$ and provide to the challenger in the blinding security game the message $\overline{\text{m}}$. Next, sample a uniform permutation $\pi : [2n] \rightarrow [2n]$ and provide to the challenger the pair of permutations

$$(T, T') = (\pi, \pi_{\text{id}}).$$

Receive from the challenger in the blinding security game a tuple of the form $(\overline{\text{pk}}^*, \overline{\text{ctx}}^*)$, and set

$$\overline{\text{pk}}_0 = \overline{\text{pk}}^*, \quad \overline{\text{ctx}}_0 = \overline{\text{ctx}}^*.$$

Set $(\widehat{\text{pk}}_1, \widehat{\text{ctx}}_1)$ as in hybrid $\text{Hyb}_{0,q-1}$, and output the re-encrypted ciphertext as:

$$\overline{\text{ct}}_j = (1, (\overline{\text{pk}}_0, \overline{\text{ctx}}_0), (\widehat{\text{pk}}_1, \widehat{\text{ctx}}_1)).$$

- Alternatively, suppose $\overline{\text{ct}}_i$ is of the form (for some $t > 0$):

$$\overline{\text{ct}}_i = (t, (\overline{\text{pk}}_0, \overline{\text{ctx}}_0), \dots, (\overline{\text{pk}}_{t-1}, \overline{\text{ctx}}_{t-1}), (\widehat{\text{pk}}_t, \widehat{\text{ctx}}_t)).$$

Recover the intermediate secret key $\overline{\text{sk}}_{t-1} = \text{KPHE.Dec}(\text{sk}_i, \overline{\text{ctx}}_t)$ and provide to the challenger in the blinding security game the message $\overline{\text{sk}}_{t-1}$. Next, sample a uniform permutation $\pi_t : [2n] \rightarrow [2n]$ and provide to the challenger the pair of permutations

$$(T, T') = (\pi_t, \pi_{\text{id}}).$$

Receive from the challenger in the blinding security game a tuple of the form $(\overline{\text{pk}}^*, \overline{\text{ctx}}^*)$, and set

$$\overline{\text{pk}}_t = \overline{\text{pk}}^*, \quad \overline{\text{ctx}}_t = \overline{\text{ctx}}^*.$$

Set $(\widehat{\text{pk}}_{t+1}, \widehat{\text{ctx}}_{t+1})$ as in hybrid $\text{Hyb}_{0,q-1}$, and output the re-encrypted ciphertext as:

$$\overline{\text{ct}}_j = ((t+1), (\overline{\text{pk}}_0, \overline{\text{ctx}}_0), \dots, (\overline{\text{pk}}_{t-1}, \overline{\text{ctx}}_{t-1}), (\overline{\text{pk}}_t, \overline{\text{ctx}}_t), (\widehat{\text{pk}}_{t+1}, \widehat{\text{ctx}}_{t+1})).$$

- Eventually the adversary \mathcal{A} outputs a bit b' . \mathcal{B} outputs the same bit b' .

It immediately follows \mathcal{B} has the same advantage ϵ as \mathcal{A} in breaking the blinding security of the KPHE scheme whenever $i = i'$ (i.e., when \mathcal{B} does not abort the simulation). Hence, the overall advantage of \mathcal{B} in breaking the blinding security of the KPHE scheme is

$$\epsilon' \geq \epsilon \cdot \Pr[i = i'] \geq \epsilon/n,$$

which is non-negligible whenever ϵ is non-negligible, for any $n = \text{poly}(\lambda)$. This leads to a contradiction, and concludes the proof of $\text{Hyb}_{0,q-1} \stackrel{c}{\approx} \text{Hyb}'_{0,q-1}$ for all $q \in [Q]$. \square

$\text{Hyb}'_{0,q-1} \stackrel{c}{\approx} \text{Hyb}_{0,q}$. We now argue that $\text{Hyb}'_{0,q-1} \stackrel{c}{\approx} \text{Hyb}_{0,q}$ for all $q \in [Q]$.

Proof. Suppose that there exists some PPT adversary \mathcal{A} that can efficiently distinguish $\text{Hyb}'_{0,q-1}$ and $\text{Hyb}_{0,q}$ with non-negligible advantage ϵ . We construct a PPT algorithm \mathcal{B} that breaks the blinding security of the KPHE scheme with non-negligible advantage ϵ' . The algorithm \mathcal{B} receives from the challenger in the blinding security game the public parameters pp , a secret key sk^* , and a public key pk^* , and proceeds as follows.

- \mathcal{B} uses pp to set up the PRE keys (both honest and corrupt) to be provided to the adversary \mathcal{A} exactly as in Hyb_4 except that it sets $\text{sk}_j := \text{sk}^*$ and $\text{pk}_j := \text{pk}^*$ for some uniform $j \xleftarrow{\$} \mathcal{K}_{\text{Honest}}$. Since pk^* is uniformly random, the view of the adversary \mathcal{A} with respect to the distribution of the keys (both honest and corrupt) remains identical.
- Suppose that the q -th honest re-encryption query of the form $\text{ReEnc}(i, j', \overline{\text{ct}}_i)$ to which the challenger does not respond with \perp . If $j \neq j'$, \mathcal{B} outputs \perp . Otherwise, it fetches the entry $(i, \overline{\text{ct}}_i)$ (from either \mathcal{L} or $\tilde{\mathcal{L}}$) and proceeds as follows.
- Suppose $\overline{\text{ct}}_i$ is of the form:

$$\overline{\text{ct}}_i = (0, (\text{pk}_i, \widehat{\text{ctx}}_0)).$$

Provide to the challenger in the blinding security game the message sk_i . Next, sample a uniform permutation $\pi : [2n] \rightarrow [2n]$ and provide to the challenger the pair of permutations

$$(T, T') = (\pi_{\text{id}}, \pi).$$

Receive from the challenger in the blinding security game a tuple of the form $(\overline{\text{pk}}^*, \overline{\text{ctx}}^*)$. Set $(\overline{\text{pk}}_0, \overline{\text{ctx}}_0)$ as in hybrid $\text{Hyb}_{0,q}$, and

$$\widehat{\text{pk}}_1 = \overline{\text{pk}}^*, \quad \widehat{\text{ctx}}_1 = \overline{\text{ctx}}^*.$$

Output the re-encrypted ciphertext as:

$$\overline{\text{ct}}_j = (1, (\overline{\text{pk}}_0, \overline{\text{ctx}}_0), (\widehat{\text{pk}}_1, \widehat{\text{ctx}}_1)).$$

- Alternatively, suppose $\overline{\text{ct}}_i$ is of the form (for some $t > 0$):

$$\overline{\text{ct}}_i = (t, (\overline{\text{pk}}_0, \overline{\text{ctx}}_0), \dots, (\overline{\text{pk}}_{t-1}, \overline{\text{ctx}}_{t-1}), (\widehat{\text{pk}}_t, \widehat{\text{ctx}}_t)).$$

Provide to the challenger in the blinding security game the message sk_i . Next, sample a uniform permutation $\pi_t : [2n] \rightarrow [2n]$ and provide to the challenger the pair of permutations

$$(T, T') = (\pi_{\text{id}}, \pi_t).$$

Receive from the challenger in the blinding security game a tuple of the form $(\overline{\text{pk}}^*, \overline{\text{ctx}}^*)$. Set $(\overline{\text{pk}}_t, \overline{\text{ctx}}_t)$ as in hybrid $\text{Hyb}_{0,q}$, and

$$\widehat{\text{pk}}_{t+1} = \overline{\text{pk}}^*, \quad \widehat{\text{ctx}}_{t+1} = \overline{\text{ctx}}^*.$$

Output the re-encrypted ciphertext as:

$$\overline{\text{ct}}_j = ((t+1), (\overline{\text{pk}}_0, \overline{\text{ctx}}_0), \dots, (\overline{\text{pk}}_{t-1}, \overline{\text{ctx}}_{t-1}), (\overline{\text{pk}}_t, \overline{\text{ctx}}_t), (\widehat{\text{pk}}_{t+1}, \widehat{\text{ctx}}_{t+1})).$$

- Eventually the adversary \mathcal{A} outputs a bit b' . \mathcal{B} outputs the same bit b' .

It immediately follows \mathcal{B} has the same advantage ϵ as \mathcal{A} in breaking the blinding security of the KPHE scheme whenever $j = j'$ (i.e., when \mathcal{B} does not abort the simulation). Hence, the overall advantage of \mathcal{B} in breaking the blinding security of the KPHE scheme is

$$\epsilon' \geq \epsilon \cdot \Pr[j = j'] \geq \epsilon/n,$$

which is non-negligible whenever ϵ is non-negligible, for any $n = \text{poly}(\lambda)$. This leads to a contradiction, and concludes the proof of $\text{Hyb}'_{0,q-1} \stackrel{c}{\approx} \text{Hyb}_{0,q}$ for all $q \in [Q]$. □

Putting these together, we have $\text{Hyb}_{0,q-1} \stackrel{c}{\approx} \text{Hyb}_{0,q}$ for all $q \in [Q]$. Finally, a simple hybrid argument over each $q \in [Q]$ completes the proof of Lemma C.4.

C.2 Proof of Lemma C.5

We now present the detailed proof of Lemma C.5.

Proof. Suppose that there exists some PPT adversary \mathcal{A} that can efficiently distinguish Hyb_3 and Hyb_4 in the proof of Theorem 4.6 with non-negligible advantage ϵ . We construct a PPT algorithm \mathcal{B} that breaks the distributional circular security of the KPHE scheme with non-negligible advantage ϵ' . The algorithm \mathcal{B} receives from the challenger in the distributional circular security game the public parameters pp proceeds as follows.

- Suppose that the adversary \mathcal{A} outputs the honest and corrupt sets $\mathcal{K}_{\text{Honest}}$ and $\mathcal{K}_{\text{Corrupt}}$ of its choice. \mathcal{B} provides to the challenger in the distributional circular security game with $n' = |\mathcal{K}_{\text{Honest}}|$, and receives the set of KPHE public keys, say, $\{\text{pk}_j\}_{j \in \mathcal{K}_{\text{Honest}}}$.
- In its simulation, \mathcal{B} uses the KPHE public keys $\{\text{pk}_j\}_{j \in \mathcal{K}_{\text{Honest}}}$ to simulate the honest PRE keys to be provided to the adversary \mathcal{A} . It additionally samples the corrupt keys to be provided to the adversary \mathcal{A} on its own. Note that the view of the adversary \mathcal{A} with respect to the distribution of the keys (both honest and corrupt) remains identical to that in Hyb_3 .

- \mathcal{B} simulates the challenger in answering all re-encryption queries (on honestly generated ciphertexts) issued by \mathcal{A} exactly as in Hyb_3 . Note that \mathcal{B} can simulate the challenger in Hyb_3 perfectly in answering such re-encryption queries, even without the knowledge of the secret keys corresponding to the honest public keys, since the challenger in Hyb_3 does not require any knowledge of the honest secret keys to answer re-encryption queries on honestly generated ciphertexts.
- \mathcal{B} receives an ensemble of ciphertexts $\{\text{ct}_{i,j}^*\}_{i,j \in \mathcal{K}_{\text{Honest}}}$ from the challenger in the distributional circular security game. Upon receipt of a query from \mathcal{A} of the form $\text{ReKeyGen}(i, j)$ such that $i \in \mathcal{K}_{\text{Honest}}$ and $j \in \mathcal{K}_{\text{Honest}}$, \mathcal{B} responds with $\text{rk}_{i,j} = (\text{pk}_j, \text{ct}_{j,i}^*)$.
- \mathcal{B} simulates the challenger in answering the challenge encryption query issued by \mathcal{A} exactly as in Hyb_3 . Note that \mathcal{B} can simulate the challenger in Hyb_3 perfectly in answering the challenge encryption query, even without the knowledge of the secret keys corresponding to the honest public keys, since the challenger in Hyb_3 does not require any knowledge of the honest secret keys to answer the challenge encryption query.
- Eventually the adversary \mathcal{A} outputs a bit b' . \mathcal{B} outputs the same bit b' .

It immediately follows \mathcal{B} has the same advantage ϵ as \mathcal{A} in breaking the distributional circular security of the KPHE scheme. This leads to a contradiction, and concludes the proof of Lemma C.5. \square

C.3 Proof of Lemma C.6

Finally, we present the detailed proof of Lemma C.6.

Proof. Suppose that there exists some PPT adversary \mathcal{A} that can efficiently distinguish Hyb_4 and Hyb_5 in the proof of Theorem 4.6 with non-negligible advantage ϵ . We construct a PPT algorithm \mathcal{B} that breaks the distributional semantic security of the KPHE scheme with non-negligible advantage ϵ' . The algorithm \mathcal{B} receives from the challenger in the distributional semantic security game the public parameters pp and a public key pk^* , and proceeds as follows.

- \mathcal{B} uses pp to set up the PRE keys (both honest and corrupt) to be provided to the adversary \mathcal{A} exactly as in Hyb_4 except that it sets $\text{pk}_{i^*} := \text{pk}^*$ for some uniform $i^* \xleftarrow{\$} \mathcal{K}_{\text{Honest}}$. Since pk^* is uniformly random, the view of the adversary \mathcal{A} with respect to the distribution of the keys (both honest and corrupt) remains identical.
- \mathcal{B} then proceeds to simulate the challenger exactly in Hyb_4 . Note that \mathcal{B} can simulate the challenger in Hyb_4 perfectly, without the knowledge of the secret key sk^* corresponding to pk^* , since the challenger in Hyb_4 does not require any knowledge of the honest secret keys.
- Suppose that in the challenge phase of hybrid Hyb_4 , the adversary issues a challenge query of the form $\mathcal{O}.\text{chall}(i, \overline{\text{m}}_0, \overline{\text{m}}_1)$. If $i \neq i^*$, \mathcal{B} outputs \perp and aborts.
- Otherwise, \mathcal{B} samples $b \in \{0, 1\}$ and outputs the challenge messages $(\overline{\text{m}}_b, 0^{2n})$ to the challenger in the distributional semantic security game. Upon receipt of the challenge ciphertext $\widehat{\text{ct}}^*$, \mathcal{B} outputs $\overline{\text{ct}}^* = (0, (\text{pk}^*, \widehat{\text{ct}}^*))$ to the adversary \mathcal{A} .

- Eventually the adversary \mathcal{A} outputs a bit b' . \mathcal{B} outputs the same bit b' .

It immediately follows when $i = i^*$ (i.e., when \mathcal{B} does not abort the simulation), \mathcal{B} has the same advantage ϵ as \mathcal{A} in breaking the distributional semantic security of the KPHE scheme. Hence, the overall advantage of \mathcal{B} in breaking the distributional semantic security of the KPHE scheme is

$$\epsilon' \geq \epsilon \cdot \Pr[i = i^*] \geq \epsilon/n,$$

which is non-negligible whenever ϵ is non-negligible, for any $n = \text{poly}(\lambda)$. This leads to a contradiction, and concludes the proof of Lemma C.6. \square

D Strongly Post-Compromise Secure Unidirectional PRE

We augment the construction in Section 4 to achieve a strongly post-compromise secure unidirectional PRE scheme (as per the notion of strong post-compromise security introduced in a recent work by Davidson et al. [DDL19]), without assuming any additional properties on the underlying KPHE scheme. In particular, we only make the following modification to the re-encryption algorithm PRE.ReEnc from the IND-HRA construction described above:

$\text{PRE.ReEnc}(\text{rk}_{i,j}, \text{ct})$: Parse the update token and the ciphertext as

$$\text{rk}_{i,j} = (\text{pk}_j, \text{ctx}_\Delta), \quad \text{ct} = (t, (\overline{\text{pk}}_0, \overline{\text{ctx}}_0), \dots, (\overline{\text{pk}}_{t-1}, \overline{\text{ctx}}_{t-1}), (\widehat{\text{pk}}_t, \widehat{\text{ctx}}_t)),$$

for some $t \geq 0$. Sample $(t + 1)$ uniform permutations $\pi_0, \dots, \pi_t : [2n] \rightarrow [2n]$. Also, let $\pi_{\text{id}} : [2n] \rightarrow [2n]$ denote the *identity* permutation. For each $\ell \in [1, t]$, compute

$$(\widetilde{\text{pk}}_\ell, \widetilde{\text{ctx}}_\ell) \stackrel{\$}{\leftarrow} \text{KPHE.Eval}(\overline{\text{pk}}_\ell, \overline{\text{ctx}}_\ell, \pi_\ell, \pi_{\ell-1}).$$

Additionally compute

$$\begin{aligned} (\widetilde{\text{pk}}_0, \widetilde{\text{ctx}}_0) &\stackrel{\$}{\leftarrow} \text{KPHE.Eval}(\overline{\text{pk}}_0, \overline{\text{ctx}}_0, \pi_0, \pi_{\text{id}}), \\ (\widehat{\text{pk}}_{t+1}, \widehat{\text{ctx}}_{t+1}) &\stackrel{\$}{\leftarrow} \text{KPHE.Eval}(\text{pk}_j, \text{ctx}_\Delta, \pi_{\text{id}}, \pi_t), \end{aligned}$$

and output the updated ciphertext as:

$$\text{ct}' = (t + 1, (\widetilde{\text{pk}}_0, \widetilde{\text{ctx}}_0), \dots, (\widetilde{\text{pk}}_{t-1}, \widetilde{\text{ctx}}_{t-1}), (\widetilde{\text{pk}}_t, \widetilde{\text{ctx}}_t), (\widehat{\text{pk}}_{t+1}, \widehat{\text{ctx}}_{t+1})).$$

Correctness. The correctness of the IND-PCS secure PRE scheme follows from essentially the same arguments as its IND-HRA counterpart, and is hence not detailed.

Theorem D.1 (IND-PCS Security). *Assuming that KPHE is a KPHE scheme with $2n$ -bit secret keys that satisfies blinding, distributional circular security and distributional semantic security with respect to the distribution \mathcal{U}_n , PRE is a multi-hop strongly post-compromise secure unidirectional PRE scheme.*

Proof. The proof follows a sequence of hybrids as outlined below. As before, we use $[N]$ to denote the set of all keys (honest + corrupt), $\mathcal{K}_{\text{Honest}} \subset [N]$ to denote the set of honest keys, and $\mathcal{K}_{\text{Corrupt}} \subset [N]$ to denote the set of corrupt keys. As mentioned earlier, these sets are fixed (adversarially) at the beginning of the game, before the adversary is allowed to issue re-key generation queries (this remains unchanged in each hybrid). We additionally let \mathcal{T} denote the set of all re-encryption keys received by the adversary in response to the (valid) ReKeyGen queries issued by it during the game (where these queries can be issued adaptively).

Hyb₀: This hybrid is identical to the real IND-PCS security game between the challenger and the adversary.

Hyb₁: This hybrid is identical to the hybrid Hyb₀ except that the challenger locally maintains an additional table T (initially empty), and does the following:

- For each honest encryption query of the form $\text{Enc}(i, \bar{m})$, the challenger adds to the local table T an entry of the form (i, \bar{ct}, \bar{m}) , where \bar{ct} is generated as $\bar{ct} = \text{KPHE.Enc}(\text{pk}_i, \bar{m})$.
- For each honest re-encryption query of the form $\text{ReEnc}(i, j, \bar{ct}_i)$ to which the challenger does not respond with \perp , it fetches the entry (i, \bar{ct}_i) (from either \mathcal{L} or $\tilde{\mathcal{L}}$) and proceeds as follows.

Suppose \bar{ct}_i is of the form:

$$\bar{ct}_i = (t, (\bar{\text{pk}}_0, \bar{\text{ctx}}_0), \dots, (\bar{\text{pk}}_{t-1}, \bar{\text{ctx}}_{t-1}), (\widehat{\text{pk}}_t, \widehat{\text{ctx}}_t)).$$

for some $t \geq 0$. The challenger recovers the following (in order):

$$\begin{aligned} \bar{\text{sk}}_{t-1} &= \text{KPHE.Dec}(\text{sk}_i, \widehat{\text{ctx}}_t), \\ \bar{\text{sk}}_{t-2} &= \text{KPHE.Dec}(\bar{\text{sk}}_{t-1}, \bar{\text{ctx}}_{t-1}), \\ &\vdots \\ \bar{\text{sk}}_0 &= \text{KPHE.Dec}(\bar{\text{sk}}_1, \bar{\text{ctx}}_1), \\ \bar{m} &= \text{KPHE.Dec}(\bar{\text{sk}}_0, \bar{\text{ctx}}_0). \end{aligned}$$

The challenger samples uniform permutations $\pi_0, \dots, \pi_t : [2n] \rightarrow [2n]$ and sets the following:

$$\tilde{\text{sk}}_\ell = \pi_\ell(\bar{\text{sk}}_\ell), \quad \tilde{\text{pk}}_\ell \stackrel{\$}{\leftarrow} \text{KPHE.PKGen}(\text{pp}, \tilde{\text{sk}}_\ell) \text{ for each } \ell \in [0, t-1],$$

$$\bar{\text{sk}}_t = \pi_t(\text{sk}_i), \quad \bar{\text{pk}}_t \stackrel{\$}{\leftarrow} \text{KPHE.PKGen}(\text{pp}, \bar{\text{sk}}_t).$$

Also set

$$\begin{aligned} \bar{\text{ctx}}_t &\stackrel{\$}{\leftarrow} \text{KPHE.Enc}(\bar{\text{pk}}_t, \bar{\text{sk}}_{t-1}), \quad \bar{\text{ctx}}_{t+1} \stackrel{\$}{\leftarrow} \text{KPHE.Enc}(\text{pk}_j, \bar{\text{sk}}_t), \\ (\widehat{\text{pk}}_{t+1}, \widehat{\text{ctx}}_{t+1}) &\stackrel{\$}{\leftarrow} \text{KPHE.Eval}(\text{pk}_j, \bar{\text{ctx}}_{t+1}, \pi_{\text{id}}, \pi_{\text{id}}), \end{aligned}$$

where $\pi_{\text{id}} : [2n] \rightarrow [2n]$ is the identity permutation. Finally, the challenger sets the following:

$$\tilde{\text{ctx}}_0 \stackrel{\$}{\leftarrow} \text{KPHE.Enc}(\tilde{\text{pk}}_0, \bar{m}),$$

$$\widetilde{\text{ctx}}_\ell \stackrel{\$}{\leftarrow} \text{KPHE.Enc}(\widetilde{\text{pk}}_\ell, \widetilde{\text{sk}}_{\ell-1}) \text{ for each } \ell \in [1, t-1].$$

Output the re-encrypted ciphertext as:

$$\overline{\text{ct}}_j = ((t+1), (\widetilde{\text{pk}}_0, \widetilde{\text{ctx}}_0), \dots, (\widetilde{\text{pk}}_{t-1}, \widetilde{\text{ctx}}_{t-1}), (\overline{\text{pk}}_t, \overline{\text{ctx}}_t), (\widehat{\text{pk}}_{t+1}, \widehat{\text{ctx}}_{t+1})).$$

Also, add to the local table T an entry of the form $(j, \overline{\text{ct}}_j, \widetilde{\text{sk}}_0, \dots, \widetilde{\text{sk}}_{t-1}, \overline{\text{sk}}_t)$.

The view of the IND-PCS adversary in this hybrid is computationally indistinguishable from that in the hybrid Hyb_0 under the assumption the KPHE scheme satisfies public key and ciphertext blinding. The proof of this indistinguishability claim follows from a hybrid argument very similar to that used in the proof of Lemma C.4, and is hence not detailed separately.

Hyb₂: This hybrid is identical to the hybrid Hyb_1 except that the challenger does the following (it still locally maintains the table T as in Hyb_1): for each honest re-encryption query of the form $\text{ReEnc}(i, j, \overline{\text{ct}}_i)$ to which the challenger does not respond with \perp , it fetches the entry $(i, \overline{\text{ct}}_i)$ (from either \mathcal{L} or $\widetilde{\mathcal{L}}$) and proceeds as follows.

Suppose $\overline{\text{ct}}_i$ is of the form:

$$\overline{\text{ct}}_i = (t, (\overline{\text{pk}}_0, \overline{\text{ctx}}_0), \dots, (\overline{\text{pk}}_{t-1}, \overline{\text{ctx}}_{t-1}), (\widehat{\text{pk}}_t, \widehat{\text{ctx}}_t)).$$

for some $t \geq 0$. The challenger looks up the local table T for an entry of the form

$$(i, \overline{\text{ct}}_i, \overline{\text{m}}, \overline{\text{sk}}_0, \dots, \overline{\text{sk}}_{t-1}),$$

and uses these secret keys for the simulation instead of decrypting and recovering them as in hybrid Hyb_1 . Such an entry is guaranteed to exist. The rest of the simulation proceeds exactly as in hybrid Hyb_1 .

It is easy to see that view of the IND-PCS adversary in this hybrid is identical to that in the hybrid Hyb_1 .

Hyb₃: This hybrid is identical to the hybrid Hyb_2 except that the challenger does the following (it still locally maintains the table T as in Hyb_2): for each honest re-encryption query of the form $\text{ReEnc}(i, j, \overline{\text{ct}}_i)$ to which the challenger does not respond with \perp , it fetches the entry $(i, \overline{\text{ct}}_i)$ (from either \mathcal{L} or $\widetilde{\mathcal{L}}$) and proceeds as follows.

Suppose $\overline{\text{ct}}_i$ is of the form:

$$\overline{\text{ct}}_i = (t, (\overline{\text{pk}}_0, \overline{\text{ctx}}_0), \dots, (\overline{\text{pk}}_{t-1}, \overline{\text{ctx}}_{t-1}), (\widehat{\text{pk}}_t, \widehat{\text{ctx}}_t)),$$

for some $t \geq 0$. The challenger looks up the local table T for an entry of the form

$$(i, \overline{\text{ct}}_i, \overline{\text{m}}, \overline{\text{sk}}_0, \dots, \overline{\text{sk}}_{t-1}),$$

and uses these secret keys for the simulation instead of decrypting and recovering them as in hybrid Hyb_1 . Such an entry is guaranteed to exist. Next, the challenger samples the following:

$$\widetilde{\text{sk}}_\ell \stackrel{\$}{\leftarrow} \mathcal{U}_n, \quad \widetilde{\text{pk}}_\ell \stackrel{\$}{\leftarrow} \text{KPHE.PKGen}(\text{pp}, \widetilde{\text{sk}}_\ell) \text{ for each } \ell \in [0, t-1].$$

$$\overline{\text{sk}}_t \stackrel{\$}{\leftarrow} \mathcal{U}_n, \quad \overline{\text{pk}}_t \stackrel{\$}{\leftarrow} \text{KPHE.PKGen}(\text{pp}, \overline{\text{sk}}_t).$$

The rest of the simulation proceeds exactly as in hybrid Hyb_2 .

It is again easy to see that the view of the IND-PCS adversary in this hybrid is identical to that in the hybrid Hyb_2 , since the distributions of the intermediate secret keys in the re-encrypted ciphertexts remain identical.

Remark D.2. *Note that in Hyb_3 , the set of secret keys $\{\text{sk}_i\}_{i \in [N]}$ (including both corrupt and honest secret keys) is no longer used by the challenger when answering ReEnc queries. In particular, all ReEnc queries are answered using only the knowledge of the public keys $\{\text{pk}_i\}_{i \in [N]}$, and the local table T maintained by the challenger.*

Hyb_4 : This hybrid is identical to the hybrid Hyb_3 except for the manner in which the challenger answers the ReKeyGen queries issued by the adversary. In particular, given a query of the form $\text{ReKeyGen}(i, j)$ such that $i \in \mathcal{K}_{\text{Honest}}$ and $j \in \mathcal{K}_{\text{Honest}}$, the challenger responds with $\text{rk}_{i,j} = (\text{pk}_j, \text{KPHE.Enc}(\text{pk}_j, 0^{2n}))$.

We argue that Hyb_4 is indistinguishable from Hyb_3 in a straightforward manner under the assumption that KPHE satisfies distributional circular security. This argument is identical to the argument in the proof of Lemma C.5, and is hence not detailed.

Remark D.3. *Note that in Hyb_4 , the set of honest secret keys $\{\text{sk}_i\}_{i \in \mathcal{K}_{\text{Honest}}}$ is no longer used by the challenger when answering ReKeyGen queries. Also, as in Hyb_3 , the set of honest secret keys $\{\text{sk}_i\}_{i \in \mathcal{K}_{\text{Honest}}}$ are also not used by the challenger when answering ReEnc queries. In other words, Hyb_4 allows the challenger to entirely “forget” the set of honest secret keys.*

Hyb_5 : This hybrid is identical to the hybrid Hyb_4 except for the manner in which the challenger answers the challenge re-encryption query issued by the adversary. In particular, given a challenge query of the form $\mathcal{O}.\text{chall-IND-PCS}(i, j, \overline{\text{ct}}_0^*, \overline{\text{ct}}_1^*)$, the challenger proceeds as follows:

- If $(i, \overline{\text{ct}}_0) \notin \mathcal{L}$ or $(i, \overline{\text{ct}}_1) \notin \mathcal{L}$, respond with $\overline{\text{ct}}^* = \perp$ (this is exactly as in Hyb_4).
- If $|\overline{\text{ct}}_0^*| \neq |\overline{\text{ct}}_1^*|$ or $j \in \mathcal{K}_{\text{Corrupt}}$, respond with $\overline{\text{ct}}^* = \perp$ (this is exactly as in Hyb_4).
- Otherwise, proceed as follows: let b be the choice bit of the challenger and suppose $\overline{\text{ct}}_b^*$ is of the form:

$$\overline{\text{ct}}_b^* = (t, (\overline{\text{pk}}_0, \overline{\text{ctx}}_0), \dots, (\overline{\text{pk}}_{t-1}, \overline{\text{ctx}}_{t-1}), (\widehat{\text{pk}}_t, \widehat{\text{ctx}}_t)),$$

for some $t \geq 0$. Look up the local table T for an entry of the form

$$(i, \overline{\text{ct}}_b^*, \overline{\text{m}}, \overline{\text{sk}}_0, \dots, \overline{\text{sk}}_{t-1}).$$

Such an entry is guaranteed to exist. Next, sample uniform permutations $\pi_0, \dots, \pi_t : [2n] \rightarrow [2n]$ and set the following:

$$\widetilde{\text{sk}}_\ell = \pi_\ell(\overline{\text{sk}}_\ell), \quad \widetilde{\text{pk}}_\ell \stackrel{\$}{\leftarrow} \text{KPHE.PKGen}(\text{pp}, \widetilde{\text{sk}}_\ell) \text{ for each } \ell \in [0, t-1],$$

$$\overline{\text{sk}}_t = \pi_t(\text{sk}_i), \quad \overline{\text{pk}}_t \stackrel{\$}{\leftarrow} \text{KPHE.PKGen}(\text{pp}, \overline{\text{sk}}_t).$$

Also set

$$\overline{\text{ctx}}_t \stackrel{\$}{\leftarrow} \text{KPHE.Enc}(\overline{\text{pk}}_t, \overline{\text{sk}}_{t-1}), \quad \overline{\text{ctx}}_{t+1} \stackrel{\$}{\leftarrow} \text{KPHE.Enc}(\text{pk}_j, \overline{\text{sk}}_t),$$

$$(\widehat{\text{pk}}_{t+1}, \widehat{\text{ctx}}_{t+1}) \stackrel{\$}{\leftarrow} \text{KPHE.Eval}(\text{pk}_j, \overline{\text{ctx}}_{t+1}, \pi_{\text{id}}, \pi_{\text{id}}),$$

where $\pi_{\text{id}} : [2n] \rightarrow [2n]$ is the identity permutation. Finally, the challenger sets the following:

$$\begin{aligned} \widetilde{\text{ctx}}_0 &\stackrel{\$}{\leftarrow} \text{KPHE.Enc}(\widetilde{\text{pk}}_0, \overline{\text{m}}), \\ \widetilde{\text{ctx}}_\ell &\stackrel{\$}{\leftarrow} \text{KPHE.Enc}(\widetilde{\text{pk}}_\ell, \widetilde{\text{sk}}_{\ell-1}) \text{ for each } \ell \in [1, t-1]. \end{aligned}$$

Output the re-encrypted ciphertext as:

$$\overline{\text{ct}}_j^* = ((t+1), (\widetilde{\text{pk}}_0, \widetilde{\text{ctx}}_0), \dots, (\widetilde{\text{pk}}_{t-1}, \widetilde{\text{ctx}}_{t-1}), (\overline{\text{pk}}_t, \overline{\text{ctx}}_t), (\widehat{\text{pk}}_{t+1}, \widehat{\text{ctx}}_{t+1})).$$

Also, add to the local table T_1 an entry of the form $(j, \overline{\text{ct}}_j^*, \widetilde{\text{sk}}_0, \dots, \widetilde{\text{sk}}_{t-1}, \overline{\text{sk}}_t)$.

We argue that the view of the adversary in Hyb_5 is computationally indistinguishable from that in Hyb_4 under the assumption the KPHE scheme satisfies public key and ciphertext blinding.

Hyb₆: This hybrid is identical to the hybrid Hyb_5 except for the manner in which the challenger answers the challenge re-encryption query issued by the adversary. In particular, when responding to the challenge query, the challenger sets:

$$\begin{aligned} \widetilde{\text{sk}}_\ell &\stackrel{\$}{\leftarrow} \mathcal{U}_n, \quad \widetilde{\text{pk}}_\ell \stackrel{\$}{\leftarrow} \text{KPHE.PKGen}(\text{pp}, \widetilde{\text{sk}}_\ell) \text{ for each } \ell \in [0, t-1]. \\ \overline{\text{sk}}_t &\stackrel{\$}{\leftarrow} \mathcal{U}_n, \quad \overline{\text{pk}}_t \stackrel{\$}{\leftarrow} \text{KPHE.PKGen}(\text{pp}, \overline{\text{sk}}_t). \end{aligned}$$

The rest of the simulation proceeds exactly as in hybrid Hyb_5 .

It is again easy to see that the view of the adversary in Hyb_6 is identical to that in Hyb_5 , since the distributions of the intermediate secret keys in the challenge ciphertext do not change.

Hyb₇: This hybrid is identical to the hybrid Hyb_6 except for the manner in which the challenger answers the challenge re-encryption query issued by the adversary. In particular, when responding to the challenge query, the challenger sets:

$$\overline{\text{ctx}}_t \stackrel{\$}{\leftarrow} \text{KPHE.Enc}(\overline{\text{pk}}_t, 0^{2n}), \quad \overline{\text{ctx}}_{t+1} \stackrel{\$}{\leftarrow} \text{KPHE.Enc}(\text{pk}_j, 0^{2n}),$$

$$(\widehat{\text{pk}}_{t+1}, \widehat{\text{ctx}}_{t+1}) \stackrel{\$}{\leftarrow} \text{KPHE.Eval}(\text{pk}_j, \overline{\text{ctx}}_{t+1}, \pi_{\text{id}}, \pi_{\text{id}}),$$

where $\pi_{\text{id}} : [2n] \rightarrow [2n]$ is the identity permutation. Finally, the challenger sets the following:

$$\begin{aligned} \widetilde{\text{ctx}}_0 &\stackrel{\$}{\leftarrow} \text{KPHE.Enc}(\widetilde{\text{pk}}_0, 0^{2n}), \\ \widetilde{\text{ctx}}_\ell &\stackrel{\$}{\leftarrow} \text{KPHE.Enc}(\widetilde{\text{pk}}_\ell, 0^{2n}) \text{ for each } \ell \in [1, t-1]. \end{aligned}$$

The rest of the simulation proceeds exactly as in hybrid Hyb_6 , i.e., the challenger outputs the re-encrypted ciphertext as:

$$\overline{\text{ct}}_j^* = ((t+1), (\widetilde{\text{pk}}_0, \widetilde{\text{ctx}}_0), \dots, (\widetilde{\text{pk}}_{t-1}, \widetilde{\text{ctx}}_{t-1}), (\overline{\text{pk}}_t, \overline{\text{ctx}}_t), (\widehat{\text{pk}}_{t+1}, \widehat{\text{ctx}}_{t+1})).$$

We argue that Hyb_7 is indistinguishable from Hyb_6 under the assumption that KPHE satisfies distributional semantic security. The argument follows via a hybrid argument over each of the intermediate secret keys in the challenge ciphertext in reverse order, i.e., starting with $\overline{\text{sk}}_t$, and then in order from $\overline{\text{sk}}_{t-1}$ through $\overline{\text{sk}}_0$, where each hybrid argument is very similar to the argument used in the proof of Lemma C.6. The overall argument is also very similar to the proof of Lemma B.1 in the proof of IND-UPD security for our UE construction in Section 3.3.1. Hence, the argument is not detailed separately.

Remark D.4. *Note that in Hyb_7 , the challenge ciphertext $\overline{\text{ct}}^*$ is independent of $(\overline{\text{ct}}_0^*, \overline{\text{ct}}_1^*)$, and hence the adversary’s advantage in winning the IND-PCS game is zero.*

□

E KPHE from Concrete Assumptions

In this section, we show how to instantiate a KPHE scheme as defined in Section 2 from concrete assumptions such as DDH and LWE.

E.1 KPHE from the DDH Assumption

In [BHHO08], Boneh *et al.* showed that assuming that the decisional Diffie-Hellman (DDH) assumption holds (over any prime-order group), there exists a KPHE scheme with $2n$ -bit secret keys that satisfies circular security with respect to the distribution \mathcal{U}_n , as well as public key and ciphertext blinding as described in Section 2. In addition, Naor and Segev [NS12] showed that the scheme in [BHHO08] satisfies distributional semantic security, while Gentry *et al.* [GHV10] pointed out that this scheme additionally satisfies distributional circular security.

For the sake of completeness, we recall the DDH assumption, as well as the main results from [BHHO08, NS12, GHV10].

The DDH assumption. Let \mathbb{G} be a cyclic group of prime order p , and let g be any uniformly sampled generator for \mathbb{G} . The decisional Diffie-Hellman (DDH) assumption is that for all PPT algorithms \mathcal{A} , we have

$$\left| \Pr \left[\mathcal{A} \left(g, g^\alpha, g^\beta, g^{\alpha\beta} \right) = 1 \right] - \Pr \left[\mathcal{A} \left(g, g^\alpha, g^\beta, g^\gamma \right) = 1 \right] \right| \leq \text{negl}(\lambda),$$

where $\alpha, \beta, \gamma \xleftarrow{R} \mathbb{Z}_p^*$.

Theorem E.1 (Imported from [BHHO08, NS12, GHV10]). *Assuming that the DDH assumption holds over some cyclic group \mathbb{G} of prime order p , there exists a KPHE scheme with $2n$ -bit secret keys that satisfies distributional semantic security with respect to the distribution \mathcal{U}_n , distributional circular security with respect to the distribution \mathcal{U}_n , and blinding.*

In particular, Boneh *et al.* explicitly showed in Section 3.2 of [BHHO08] that an “expanded version” of their circular-secure PKE from DDH is a KPHE scheme that satisfies the required properties summarized above.

E.2 KPHE from the LWE Assumption

Based on existing results on public-key encryption from the learning with errors (LWE) assumption [Reg09, ACPS09, GKPV10], it follows that assuming LWE, there exists a KPHE scheme with $2n$ -bit secret keys that satisfies distributional (semantic and circular) security with respect to the distribution \mathcal{U}_n , as well as public key and ciphertext blinding as defined in Section 2. The construction is essentially a simple variant of the seminal PKE scheme due to Regev [Reg09]; nonetheless, we recall it here for completeness.

E.2.1 The LWE Assumption

We first review the learning with errors (LWE) assumption. Let $n, m, q \in \mathbb{N}$ be positive integers and let χ be a noise distribution over \mathbb{Z}_q . In the $\text{LWE}(n, m, q, \chi)$ problem, the adversary's goal is to distinguish between the two distributions:

$$(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}) \quad \text{and} \quad (\mathbf{A}, \mathbf{u}),$$

where $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{m \times n}$, $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n$, $\mathbf{e} \xleftarrow{\$} \chi^m$, and $\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^m$. The corresponding $\text{LWE}(n, m, q, \chi)$ assumption (informally) states that the $\text{LWE}(n, m, q, \chi)$ problem is computationally intractable. More formally, the $\text{LWE}(n, m, q, \chi)$ assumption states that for $m, n, q = \text{poly}(\lambda)$ and for any PPT adversary \mathcal{A} , we have

$$|\Pr[\mathcal{A}(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e})] - \Pr[\mathcal{A}(\mathbf{A}, \mathbf{u})]| < \text{negl}(\lambda),$$

where $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{m \times n}$, $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n$, $\mathbf{e} \xleftarrow{\$} \chi^m$, and $\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^m$.

The Binary-Secret LWE Assumption. We also review the learning with errors (LWE) assumption for *binary secrets*. Let $n, m, q \in \mathbb{N}$ be positive integers and let χ be a noise distribution over \mathbb{Z}_q . In the binary-secret $\text{LWE}(n, m, q, \chi)$ problem, the adversary's goal is to distinguish between the two distributions:

$$(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}) \quad \text{and} \quad (\mathbf{A}, \mathbf{u}),$$

where $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{m \times n}$, $\mathbf{s} \xleftarrow{\$} \{0, 1\}^n$, $\mathbf{e} \xleftarrow{\$} \chi^m$, and $\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^m$. The corresponding binary-secret $\text{LWE}(n, m, q, \chi)$ assumption (informally) states that the binary-secret $\text{LWE}(n, m, q, \chi)$ problem is computationally intractable. More formally, the binary-secret $\text{LWE}(n, m, q, \chi)$ assumption states that for $m, n, q = \text{poly}(\lambda)$ and for any PPT adversary \mathcal{A} , we have

$$|\Pr[\mathcal{A}(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e})] - \Pr[\mathcal{A}(\mathbf{A}, \mathbf{u})]| < \text{negl}(\lambda),$$

where $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{m \times n}$, $\mathbf{s} \xleftarrow{\$} \{0, 1\}^n$, $\mathbf{e} \xleftarrow{\$} \chi^m$, and $\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^m$.

The Distributional Binary-Secret LWE Assumption. We also consider a variant of the binary-secret LWE assumption the secret vector \mathbf{s} is not necessarily a uniformly random bit-string, but is sampled from some min-entropy distribution \mathcal{D} over $\{0, 1\}^n$. Formally, the $\text{LWE}(n, m, q, \chi, \mathcal{D})$ assumption states that for $m, n, q = \text{poly}(\lambda)$ and for any PPT adversary \mathcal{A} , we have

$$|\Pr[\mathcal{A}(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e})] - \Pr[\mathcal{A}(\mathbf{A}, \mathbf{u})]| < \text{negl}(\lambda),$$

where \mathcal{D} is some distribution over $\{0, 1\}^n$ and we have $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{m \times n}$, $\mathbf{s} \xleftarrow{\$} \mathcal{D}^n$, $\mathbf{e} \xleftarrow{\$} \chi^m$, and $\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^m$.

E.2.2 The KPHE Construction

In this section, we show a KPHE construction from the distributional binary-secret LWE assumption with super-polynomial modulus-to-noise ratio.

Setup: On input the security parameter λ , fix and output the parameters $m, n, q = \text{poly}(\lambda)$, where $m > 6n \log q + \omega(\sqrt{\log \lambda})$, as well as error distributions χ, χ' over $[-B, B]$ and $[-B', B']$, respectively, for some $B, B' \in \mathbb{N}$ such that

$$B/B' = \text{negl}(\lambda), \quad B'/q = \text{negl}(\lambda),$$

where T is the number of Output the public parameters

$$\text{pp} = (m, n, q, \chi, \chi').$$

SKGen: On input the public parameter $\text{pp} = (m, n, q, \chi)$, sample a vector $\mathbf{s} \in \{0, 1\}^{2n} \xleftarrow{\$} \mathcal{U}_n$ output the secret key $\text{sk} = \mathbf{s}$.

PKGen: On input the public parameter $\text{pp} = (m, n, q, \chi)$ and a secret key vector $\mathbf{s} \in \{0, 1\}^{2n}$, sample a matrix $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{m \times 2n}$ and an error vector $\mathbf{e} \xleftarrow{\$} \chi^m$, and output a public key pk generated as follows:

$$\text{pk} = (\mathbf{A}, \mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e}).$$

Enc: On input a public key $\text{pk} = (\mathbf{A}, \mathbf{b})$ and a message vector $\vec{\mathbf{m}} \in \{0, 1\}^{2n}$, sample a uniformly random binary matrix $\mathbf{R} \xleftarrow{\$} \{0, 1\}^{2n \times m}$ and an error vector $\mathbf{e}' \xleftarrow{\$} (\chi')^{2n}$, and output the ciphertext $\text{ct} = (\mathbf{U}, \mathbf{v})$ where

$$\mathbf{U} = \mathbf{R}\mathbf{A}, \quad \mathbf{v} = \mathbf{R}\mathbf{b} + \mathbf{e}' + \lfloor q/2 \rfloor \vec{\mathbf{m}}.$$

Dec: On input a secret key $\text{sk} = \mathbf{s}$ and a ciphertext $\text{ct} = (\mathbf{U}, \mathbf{v})$, output the decrypted message vector $\vec{\mathbf{m}}'$ as

$$\vec{\mathbf{m}}' = \text{decode}_q(\mathbf{v} - \mathbf{U}\mathbf{s}),$$

where for any vector $\mathbf{y} = (y_1, \dots, y_n) \in \mathbb{Z}_q^{2n}$, we have $\text{decode}_q(\mathbf{y}) = (z_1, \dots, z_n) \in \{0, 1\}^{2n}$ such that for each $i \in [n]$, we have

$$z_i = \begin{cases} 0 & \text{if } y_i < q/4, \\ 1 & \text{otherwise.} \end{cases}$$

Eval: On input a public key $\text{pk} = (\mathbf{A}, \mathbf{b})$, a ciphertext $\text{ct} = (\mathbf{U}, \mathbf{v})$, and a pair of permutations $\pi, \pi' : [2n] \rightarrow [2n]$, proceed via the following steps:

- **Step-1: Message Rotation.** This step takes as input the ciphertext $\text{ct} = (\mathbf{U}, \mathbf{v})$ that encrypts some message vector $\vec{\mathbf{m}} \in \{0, 1\}^{2n}$ under a secret key $\mathbf{s} \in \{0, 1\}^{2n}$, and produces a ciphertext $\text{ct}' = (\mathbf{U}', \mathbf{v}')$ that encrypts the message vector $\vec{\mathbf{m}}' = \pi'(\vec{\mathbf{m}})$ under the same secret key \mathbf{s} . Concretely, parse the matrix \mathbf{U} and the vector \mathbf{v} as:

$$\mathbf{U} = \begin{bmatrix} \mathbf{u}_1^t \\ \vdots \\ \mathbf{u}_{2n}^t \end{bmatrix}, \quad \mathbf{v} = \begin{bmatrix} v_1 \\ \vdots \\ v_{2n} \end{bmatrix}.$$

and output $\text{ct}' = (\mathbf{U}', \mathbf{v}')$ where \mathbf{U}' and \mathbf{v}' are *row-rotated* versions of \mathbf{U} and \mathbf{v} , respectively, under the permutation π' , i.e., we have

$$\mathbf{U}' = \begin{bmatrix} \mathbf{u}_{\pi'(1)}^t \\ \vdots \\ \mathbf{u}_{\pi'(2n)}^t \end{bmatrix}, \quad \mathbf{v}' = \begin{bmatrix} v_{\pi'(1)} \\ \vdots \\ v_{\pi'(2n)} \end{bmatrix}.$$

- **Step-2: Key Rotation (Public Key).** This step takes the public key $\text{pk} = (\mathbf{A}, \mathbf{b})$ corresponding to a secret key $\mathbf{s} \in \{0, 1\}^{2n}$, and produces a public key $\text{pk}' = (\mathbf{A}', \mathbf{b}')$ corresponding to a rotated secret key $\mathbf{s}' = \pi(\mathbf{s})$. Concretely, parse the matrix \mathbf{A} as:

$$\mathbf{A} = [\mathbf{a}_1 \mid \dots \mid \mathbf{a}_{2n}],$$

and output $\text{pk}' = (\mathbf{A}', \mathbf{b}')$ where \mathbf{A}' is a *column-rotated* version of \mathbf{A} under the permutation π , i.e., we have

$$\mathbf{A}' = [\mathbf{a}_{\pi(1)} \mid \dots \mid \mathbf{a}_{\pi(2n)}].$$

- **Step-3: Key Rotation (Ciphertext).** This step takes the ciphertext $\text{ct}' = (\mathbf{U}', \mathbf{v}')$ that encrypts some message vector $\vec{\mathbf{m}}' \in \{0, 1\}^{2n}$ under a secret key $\mathbf{s} \in \{0, 1\}^{2n}$, and produces a ciphertext $\text{ct}'' = (\mathbf{U}'', \mathbf{v}'')$ that encrypts the same message vector $\vec{\mathbf{m}}'$ under a rotated secret key $\mathbf{s}' = \pi(\mathbf{s})$. Concretely, parse the matrix \mathbf{U}' as:

$$\mathbf{U}' = [\mathbf{u}'_1 \mid \dots \mid \mathbf{u}'_{2n}],$$

and output $\text{ct}'' = (\mathbf{U}'', \mathbf{v}'')$ where \mathbf{U}'' is a *column-rotated* version of \mathbf{U}' under the permutation π , i.e., we have

$$\mathbf{U}'' = [\mathbf{u}'_{\pi(1)} \mid \dots \mid \mathbf{u}'_{\pi(2n)}].$$

Step-4: Public Key Blinding: This step takes the public key $\text{pk}' = (\mathbf{A}', \mathbf{b}')$ corresponding to the secret key \mathbf{s}' , and produces a public key $\widetilde{\text{pk}} = (\widetilde{\mathbf{A}}, \widetilde{\mathbf{b}})$ that is computationally indistinguishable from a uniformly random public key corresponding to \mathbf{s}' . Concretely, sample a uniformly random binary matrix $\mathbf{T} \xleftarrow{\$} \{0, 1\}^{m \times m}$, an error matrix $\widetilde{\mathbf{E}}_0 \xleftarrow{\$} \chi^{m \times 2n}$, and an error vector $\widetilde{\mathbf{e}}_1 \xleftarrow{\$} (\chi')^m$, and output $\widetilde{\text{pk}} = (\widetilde{\mathbf{A}}, \widetilde{\mathbf{b}})$, where

$$\widetilde{\mathbf{A}} = \mathbf{T}\mathbf{A}' + \widetilde{\mathbf{E}}_0, \quad \widetilde{\mathbf{b}} = \mathbf{T}\mathbf{b}' + \widetilde{\mathbf{e}}_1.$$

Step-5: Ciphertext Blinding: This step takes as input:

- the public key $\widetilde{\text{pk}} = (\widetilde{\mathbf{A}}, \widetilde{\mathbf{b}})$ corresponding to the secret key \mathbf{s}' , and
- the ciphertext $\text{ct}'' = (\mathbf{U}'', \mathbf{v}'')$ that encrypts the message vector $\vec{\mathbf{m}}'$ under the secret key \mathbf{s}' ,

and produces a ciphertext $\widetilde{\text{ct}} = (\widetilde{\mathbf{U}}, \widetilde{\mathbf{v}})$ that is computationally indistinguishable from a uniformly random ciphertext encrypting the same message $\vec{\mathbf{m}}'$ under the same secret key \mathbf{s}' . Concretely, sample a uniformly random binary matrix $\widetilde{\mathbf{R}} \xleftarrow{\$} \{0, 1\}^{2n \times m}$ and an error vector $\widetilde{\mathbf{e}} \xleftarrow{\$} (\chi')^{2n}$, and output $\widetilde{\text{ct}} = (\widetilde{\mathbf{U}}, \widetilde{\mathbf{v}})$, where

$$\widetilde{\mathbf{U}} = \mathbf{U}'' + \widetilde{\mathbf{R}}\mathbf{A}', \quad \widetilde{\mathbf{v}} = \mathbf{v}'' + \widetilde{\mathbf{R}}\mathbf{b}' + \widetilde{\mathbf{e}}.$$

Finally output the transformed public key-ciphertext pair $(\widetilde{\text{pk}}, \widetilde{\text{ct}})$, where

$$\widetilde{\text{pk}} = (\widetilde{\mathbf{A}}, \widetilde{\mathbf{b}}), \quad \widetilde{\text{ct}} = (\widetilde{\mathbf{U}}, \widetilde{\mathbf{v}}),$$

generated by following the steps as described above.

E.2.3 Correctness

We prove correctness of the steps in the evaluation algorithm individually. Together, this establishes correctness of the evaluation algorithm. Throughout, we assume that

$$B/B' = \text{negl}(\lambda), \quad B'/q = \text{negl}(\lambda).$$

Step-1: Message Rotation. Let $\text{ct} = (\mathbf{U}, \mathbf{v})$ be a ciphertext that encrypts some message vector $\vec{\mathbf{m}} = (m_1, \dots, m_n)$ under a secret key $\mathbf{s} \in \{0, 1\}^{2n}$, and suppose that step-1 produces a ciphertext $\text{ct}' = (\mathbf{U}', \mathbf{v}')$. We prove that ct' encrypts the message vector $\vec{\mathbf{m}}' = \pi'(\vec{\mathbf{m}})$ under the same secret key \mathbf{s} . To see this, observe that we have:

$$\mathbf{U} = \mathbf{R}\mathbf{A}, \quad \mathbf{v} = \mathbf{R}\mathbf{b} + \mathbf{e}' + \lfloor q/2 \rfloor \vec{\mathbf{m}} = \mathbf{U}\mathbf{s} + \hat{\mathbf{e}} + \lfloor q/2 \rfloor \vec{\mathbf{m}}.$$

Hence, we have

$$\mathbf{U} = \begin{bmatrix} \mathbf{u}_1^t \\ \vdots \\ \mathbf{u}_{2n}^t \end{bmatrix}, \quad \mathbf{v} = \begin{bmatrix} v_1 = \mathbf{u}_1^t \mathbf{s} + e_1 + m_1 \\ \vdots \\ v_{2n} = \mathbf{u}_{2n}^t \mathbf{s} + e_{2n} + m_{2n} \end{bmatrix}.$$

Also, recall that $\text{ct}' = (\mathbf{U}', \mathbf{v}')$ where \mathbf{U}' and \mathbf{v}' are *row-rotated* versions of \mathbf{U} and \mathbf{v} , respectively, under the permutation π' . Hence, we have

$$\mathbf{U}' = \begin{bmatrix} \mathbf{u}_{\pi'(1)}^t \\ \vdots \\ \mathbf{u}_{\pi'(2n)}^t \end{bmatrix}, \quad \mathbf{v}' = \begin{bmatrix} v_{\pi'(1)} \\ \vdots \\ v_{\pi'(2n)} \end{bmatrix} = \begin{bmatrix} \mathbf{u}_{\pi'(1)}^t \mathbf{s} + e_{\pi'(1)} + m_{\pi'(1)} \\ \vdots \\ \mathbf{u}_{\pi'(2n)}^t \mathbf{s} + e_{\pi'(2n)} + m_{\pi'(2n)} \end{bmatrix}.$$

This in turn implies that we have

$$\mathbf{v}' = \mathbf{U}'\mathbf{s} + \pi'(\hat{\mathbf{e}}) + \lfloor q/2 \rfloor \pi'(\vec{\mathbf{m}}) = \mathbf{U}'\mathbf{s} + \hat{\mathbf{e}}' + \lfloor q/2 \rfloor \vec{\mathbf{m}}',$$

as desired. □

Step-2: Key Rotation (Public Key). Suppose that the public key $\text{pk} = (\mathbf{A}, \mathbf{b})$ corresponds to a secret key $\mathbf{s} \in \{0, 1\}^{2n}$, and step-2 produces a public key $\text{pk}' = (\mathbf{A}', \mathbf{b})$. We prove that pk' is a valid public key corresponding to the rotated secret key $\mathbf{s}' = \pi(\mathbf{s})$. Concretely, recall that we have

$$\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e}.$$

In other words, letting $\mathbf{s} = (s_1, \dots, s_{2n})$, we have:

$$\mathbf{A} = [\mathbf{a}_1 \mid \dots \mid \mathbf{a}_{2n}], \quad \mathbf{b} = [\mathbf{a}_1 \mid \dots \mid \mathbf{a}_{2n}] \begin{bmatrix} s_1 \\ \vdots \\ s_{2n} \end{bmatrix} + \mathbf{e} = \sum_{i=1}^{2n} s_i \mathbf{a}_i + \mathbf{e}.$$

Also, suppose we have $\text{pk}' = (\mathbf{A}', \mathbf{b})$ where \mathbf{A}' is a *column-rotated* version of \mathbf{A} under the permutation π , i.e., we have

$$\mathbf{A}' = [\mathbf{a}_{\pi(1)} \mid \dots \mid \mathbf{a}_{\pi(2n)}].$$

Now, observe the following:

$$\mathbf{b} = \sum_{i=1}^{2n} s_i \mathbf{a}_i + \mathbf{e} = \sum_{i=1}^{2n} s_{\pi(i)} \mathbf{a}_{\pi(i)} + \mathbf{e} = [\mathbf{a}_{\pi(1)} \mid \dots \mid \mathbf{a}_{\pi(2n)}] \begin{bmatrix} s_{\pi(1)} \\ \vdots \\ s_{\pi(2n)} \end{bmatrix} + \mathbf{e}.$$

In other words, we have

$$\mathbf{b} = \mathbf{A}' \mathbf{s}' + \mathbf{e},$$

as desired. \square

Step-3: Key Rotation (Ciphertext). Suppose that the ciphertext $\mathbf{ct}' = (\mathbf{U}', \mathbf{v}')$ encrypts some message vector $\vec{\mathbf{m}}' \in \{0, 1\}^{2n}$ under a secret key $\mathbf{s} \in \{0, 1\}^{2n}$, and suppose that step-3 produces a ciphertext $\mathbf{ct}'' = (\mathbf{U}'', \mathbf{v}'')$. We prove that \mathbf{ct}'' encrypts the same message vector $\vec{\mathbf{m}}'$ under a rotated secret key $\mathbf{s}' = \pi(\mathbf{s})$. Concretely, recall that we have

$$\mathbf{v}' = \mathbf{U}' \mathbf{s} + \widehat{\mathbf{e}}' + \lfloor q/2 \rfloor \vec{\mathbf{m}}'.$$

In other words, letting $\mathbf{s} = (s_1, \dots, s_{2n})$, we have:

$$\mathbf{U}' = [\mathbf{u}'_1 \mid \dots \mid \mathbf{u}'_{2n}], \quad \mathbf{v}' = [\mathbf{u}'_1 \mid \dots \mid \mathbf{u}'_{2n}] \begin{bmatrix} s_1 \\ \vdots \\ s_{2n} \end{bmatrix} + \widehat{\mathbf{e}}' + \lfloor q/2 \rfloor \vec{\mathbf{m}}' = \sum_{i=1}^{2n} s_i \mathbf{u}'_i + \widehat{\mathbf{e}}' + \lfloor q/2 \rfloor \vec{\mathbf{m}}'.$$

Also, suppose that $\mathbf{ct}'' = (\mathbf{U}'', \mathbf{v}'')$ where \mathbf{U}'' is a *column-rotated* version of \mathbf{U}' under the permutation π , i.e., we have

$$\mathbf{U}'' = [\mathbf{u}'_{\pi(1)} \mid \dots \mid \mathbf{u}'_{\pi(2n)}].$$

Now, observe the following:

$$\mathbf{v}' = \sum_{i=1}^{2n} s_i \mathbf{u}'_i + \widehat{\mathbf{e}}' + \lfloor q/2 \rfloor \vec{\mathbf{m}}' = \sum_{i=1}^{2n} s_{\pi(i)} \mathbf{u}'_{\pi(i)} + \widehat{\mathbf{e}}' + \lfloor q/2 \rfloor \vec{\mathbf{m}}',$$

i.e., we have

$$\mathbf{v}' = [\mathbf{u}'_{\pi(1)} \mid \dots \mid \mathbf{u}'_{\pi(2n)}] \begin{bmatrix} s_{\pi(1)} \\ \vdots \\ s_{\pi(2n)} \end{bmatrix} + \widehat{\mathbf{e}}' + \lfloor q/2 \rfloor \vec{\mathbf{m}}'.$$

In other words, we have

$$\mathbf{v}' = \mathbf{U}'' \mathbf{s}' + \widehat{\mathbf{e}}' + \lfloor q/2 \rfloor \vec{\mathbf{m}}',$$

as desired. \square

Step-4: Public Key Blinding: Let $\text{pk}' = (\mathbf{A}', \mathbf{b}')$ be a public key corresponding to the secret key \mathbf{s}' , and suppose that step-4 produces a public key $\widetilde{\text{pk}} = (\widetilde{\mathbf{A}}, \widetilde{\mathbf{b}})$. We prove that $\widetilde{\text{pk}}$ is a valid public key corresponding to \mathbf{s}' . Recall that we have:

$$\widetilde{\mathbf{A}} = \mathbf{T}\mathbf{A}' + \widetilde{\mathbf{E}}_0, \quad \widetilde{\mathbf{b}} = \mathbf{T}\mathbf{b}' + \widetilde{\mathbf{e}}_1.$$

Also, recall that we have

$$\mathbf{b}' = \mathbf{A}'\mathbf{s}' + \mathbf{e}'.$$

Hence, we have

$$\widetilde{\mathbf{b}} = \mathbf{T}\mathbf{b}' + \widetilde{\mathbf{e}}_1 = \mathbf{T}(\mathbf{A}'\mathbf{s}' + \mathbf{e}') + \widetilde{\mathbf{e}}_1 = \mathbf{T}\mathbf{A}'\mathbf{s}' + \mathbf{T}\mathbf{e}' + \widetilde{\mathbf{e}}_1 = \widetilde{\mathbf{A}}\mathbf{s}' - \mathbf{T}\widetilde{\mathbf{E}}_0\mathbf{s}' + \mathbf{T}\mathbf{e}' + \widetilde{\mathbf{e}}_1.$$

In other words, we have

$$\widetilde{\mathbf{b}} = \widetilde{\mathbf{A}}\mathbf{s}' + \widetilde{\mathbf{e}}^*,$$

where the vector $\widetilde{\mathbf{e}}^*$ has low norm, as desired. \square

Step-5: Ciphertext Blinding: Suppose that step-5 takes as input

- the public key $\widetilde{\text{pk}} = (\widetilde{\mathbf{A}}, \widetilde{\mathbf{b}})$ corresponding to the secret key \mathbf{s}' , and
- the ciphertext $\text{ct}'' = (\mathbf{U}'', \mathbf{v}'')$ that encrypts the message vector $\vec{\mathbf{m}}'$ under the secret key \mathbf{s}' ,

and produces a ciphertext $\widetilde{\text{ct}} = (\widetilde{\mathbf{U}}, \widetilde{\mathbf{v}})$. We prove that $\widetilde{\text{ct}}$ encrypts the same message $\vec{\mathbf{m}}'$ under the same secret key \mathbf{s}' . Recall that we have $\widetilde{\text{ct}} = (\widetilde{\mathbf{U}}, \widetilde{\mathbf{v}})$, where

$$\widetilde{\mathbf{U}} = \mathbf{U}'' + \widetilde{\mathbf{R}}\widetilde{\mathbf{A}}, \quad \widetilde{\mathbf{v}} = \mathbf{v}'' + \widetilde{\mathbf{R}}\widetilde{\mathbf{b}} + \widetilde{\mathbf{e}}.$$

Also, recall that we have

$$\widetilde{\mathbf{b}} = \widetilde{\mathbf{A}}\mathbf{s}' + \widetilde{\mathbf{e}}^*, \quad \mathbf{v}'' = \mathbf{U}''\mathbf{s}' + \widehat{\mathbf{e}}' + \lfloor q/2 \rfloor \vec{\mathbf{m}}'.$$

It is easy to see that for appropriate parameters – need to be specified, we have

$$\widetilde{\mathbf{v}} = \mathbf{v}'' + \widetilde{\mathbf{R}}\widetilde{\mathbf{b}} + \widetilde{\mathbf{e}} = \widetilde{\mathbf{U}}\mathbf{s}' + \widehat{\mathbf{e}}' + \lfloor q/2 \rfloor \vec{\mathbf{m}}',$$

where the vector $\widehat{\mathbf{e}}'$ has low norm, as desired. \square

E.2.4 Security

We recall certain theorems from prior work that allow us to establish the security of the aforementioned KPHE scheme from the LWE assumption.

Theorem E.2 (Imported from [Reg09, GKPV10]). *Assuming that the LWE assumption holds, the distributional binary secret LWE assumption holds for $\mathcal{D} = \mathcal{U}_n$.*

Theorem E.3 (Imported from [Reg09, GKPV10]). *Assuming that the distributional binary secret LWE assumption holds with respect to the distribution $\mathcal{D} = \mathcal{U}_n$, the aforementioned KPHE scheme satisfies distributional semantic security with respect to the same distribution $\mathcal{D} = \mathcal{U}_n$.*

Theorem E.4 (Imported from [ACPS09,GKPV10]). *Assuming that the distributional binary secret LWE assumption holds with respect to the distribution $\mathcal{D} = \mathcal{U}_n$, the aforementioned KPHE scheme satisfies distributional circular security with respect to the same distribution $\mathcal{D} = \mathcal{U}_n$.*

Theorem E.5. *Assuming that the distributional binary secret LWE assumption holds with respect to the distribution $\mathcal{D} = \mathcal{U}_n$, the aforementioned KPHE scheme satisfies blinding.*

Proof. The aforementioned KPHE scheme satisfies statistical ciphertext blinding (this follows immediately by observation) and computational public blinding (this follows from the distributional binary secret LWE assumption via a simple hybrid argument over the rows of the matrix \mathbf{T} and the error matrix \mathbf{E}_0 used in the public key blinding step). \square

Putting everything together, we get the following corollary:

Corollary E.6. *Assuming that the LWE assumption holds, there exists a KPHE scheme with $2n$ -bit secret keys that satisfies distributional semantic security with respect to the distribution \mathcal{U}_n , distributional circular security with respect to the distribution \mathcal{U}_n , and blinding.*