

The Direction of Updatable Encryption *Does* Matter

Ryo Nishimaki¹

¹ NTT Social Informatics Laboratories, Tokyo, Japan
ryo.nishimaki@ntt.com

Abstract

We introduce a new definition for key updates, called backward-leak uni-directional key updates, in updatable encryption (UE). This notion is a variant of uni-directional key updates for UE. We show that existing secure UE schemes in the bi-directional key updates setting are not secure in the backward-leak uni-directional key updates setting. Thus, security in the backward-leak uni-directional key updates setting is *strictly* stronger than security in the bi-directional key updates setting. This result is in sharp contrast to the equivalence theorem by Jiang (Asiacrypt 2020), which says security in the bi-directional key updates setting is equivalent to security in the existing uni-directional key updates setting. We call the existing uni-directional key updates “forward-leak uni-directional” key updates to distinguish two types of uni-directional key updates in this paper.

We also present a UE schemes that is post-quantum secure in the backward-leak uni-directional key updates setting under the learning with errors assumption.

Keywords: updatable encryption, key update, lattice

Erratum and Remark

After the publication of this work at PKC 2022, Takahiro Matsuda pointed out an error in the UE scheme based on indistinguishability obfuscation, which was claimed to be secure in the no-directional key updates setting. Hence, we retracted the obfuscation-based scheme. Other results (the relation between definitions and the LWE-based scheme) are not affected. Note that, thanks to the subsequent work by Jiang Galteland and Pan [GP23], security in the backward-leak uni-directional key updates setting is equivalent to the no-directional one. Hence, our LWE-based scheme is also secure in the no-directional key updates setting.

Contents

1	Introduction	1
1.1	Background	1
1.2	Our Contribution	2
1.3	Related Work	2
1.4	Technical Overview	3
2	Preliminaries	5
2.1	Lattice Preliminaries	6
3	Updatable Encryption	7
4	Backward-Leak Uni-Directional Key Update and Relations	11
4.1	Definition	11
4.2	Observations on Definitions	12
4.3	Relationships	14
5	Construction with Backward-Leak Uni-Directional Key Update	16
5.1	Scheme Description and Design Idea	17
5.2	Correctness	19
5.3	Confidentiality	20
A	More Discussion on Key Updates	28
B	Other Security Definitions for UE	29

1 Introduction

1.1 Background

Updatable Encryption. Updatable encryption (UE) is a variant of secret key encryption (SKE) where we can periodically update a secret key and a ciphertext. More specifically, a secret key k_e is generated at each period, called epoch. Here, e denotes an index of an epoch. We can generate a conversion key Δ_{e+1} that converts a ciphertext under k_e (key at epoch e) to one under k_{e+1} (key at epoch $e + 1$). Such a conversion key is called update token and generated from two successive secret keys k_e, k_{e+1} . Roughly speaking, UE security guarantees that confidentiality holds even after some old (and even new) keys and tokens are corrupted as long as trivial winning conditions are not triggered. Adversaries trivially win if a target secret key is corrupted or a target ciphertext can be converted into a ciphertext under a corrupted secret key. In this study, we focus on ciphertext-independent updates UE, where we can generate an update token only from two secret keys [LT18, KLR19, BDGJ20, Jia20].¹

A serious threat to encryption is key leakage. In that case, no security is guaranteed by standard encryption. Key updating is a standard solution to guarantee security even after key leakage. However, the issue is how to update a ciphertext generated by an old key. A naive solution is decrypting all ciphertexts by the old key and re-encrypt them by a new key. However, it incurs significant efficiency loss. Moreover, if we save encrypted data in outsourced storage such as cloud servers, we need to download all ciphertexts from the server, decrypt and re-encrypt them, and upload them again to keep the new key secret. Update tokens of UE solve this problem since if we provide the server with an update token, it can directly convert old ciphertexts into new ones without the new key.

Confidentiality is the primary concern in UE. Confidentiality of UE has been improved to capture realistic attack models [EPRS17, LT18, KLR19, BDGJ20, CLT20] since after UE was introduced [BLMR13]. In particular, Lehman and Tackmann formalized trivially leaked information from corrupted keys and tokens as *the direction of key updates* [LT18]. Although previous works proposed UE schemes with improved confidentiality, most do not focus on preventing information leakage from corrupted keys and tokens. We will explain the detail of the information leakage below. In this work, we focus on the direction of key updates and try to minimize leaked information from update tokens to improve UE confidentiality.

Direction of key and ciphertext updates. Directions of key updates describe information leakage that UE schemes cannot avoid. If an adversary has Δ_{e+1} and k_e , it might be able to obtain k_{e+1} . Most existing UE schemes cannot prevent this attack. In particular, in all existing (ciphertext-independent) UE schemes, we cannot avoid leaking a secret key from both directions [LT18, KLR19, BDGJ20, Jia20]. That is, we can extract k_{e+1} (resp. k_e) from Δ_{e+1} and k_e (resp. k_{e+1}). This setting is defined as *bi-directional* key updates [EPRS17, LT18]. Lehman and Tackmann also defined *uni-directional* key updates, where we can extract k_{e+1} from k_e and Δ_{e+1} (forward direction inference). In other words, this setting means adversaries might not be able to infer k_e from k_{e+1} and Δ_{e+1} . Uni-directional key updates are more preferable than bi-directional ones since a token leaks less information. More information leakage triggers more trivial winning conditions in confidentiality games for UE.

At first glance, secure UE with uni-directional key updates is stronger than one with bi-directional key updates. However, Jiang proved that secure UE with bi-directional key updates *is equivalent to one* with uni-directional key updates [Jia20] (we call Jiang’s equivalence theorem in this paper). Jiang also presented the first post-quantum UE scheme with bi-directional key updates [Jia20].

A natural question is: Why do we consider only one-way uni-directional key updates? That is, we can consider a variant of uni-directional key updates where we can extract k_e from k_{e+1} and Δ_{e+1} (backward direction inference). To distinguish two versions of uni-directional key updates, we call the existing

¹The other variant is ciphertext-dependent updates UE, where we need not only two secret keys but also a part of ciphertext (called header) to generate a token [BLMR13, EPRS17, BEKS20]. Ciphertext-independent updates UE is more efficient.

definition *forward-leak uni-directional* key updates and our new one *backward-leak uni-directional* key updates. The backward-leak uni-directional key updates setting has never been studied in the UE literature, but it seems to be a valid setting. It is natural to think the latest key is the most important since the reason why we update keys is that the current and older keys might be leaked. In the forward-leak setting, we must protect older keys to protect newer keys *even if older ciphertexts are deleted*. This is undesirable. However, in the backward-leak setting, we need to protect only the latest key if older ciphertexts are properly deleted. Therefore, the backward-leak key updates are more suitable for UE than the forward-leak key updates.

A related issue is the direction of ciphertext updates. It describes whether we can convert ciphertext into one in an older epoch (downgrading ciphertext) by using an update token or not. If we can both update and downgrade ciphertexts by using a token, we say a UE scheme provides bi-directional ciphertext updates. If we can update but cannot downgrade ciphertexts by using a token, we say a UE scheme provides uni-directional ciphertext updates. UE with uni-directional ciphertext updates is more desirable since older epoch keys might be leaked, and downgrading ciphertexts leaks more information. However, all existing (ciphertext-independent) UE schemes provide bi-directional ciphertext updates.

Thus, the first main question of this study is as follows.

Q1. *Is UE with backward-leak uni-directional key updates strictly stronger than UE with bi-directional key updates?*

We affirmatively answer the first question in this work. Then, the next natural question is as follows.

Q2. *Can we achieve a (post-quantum) UE scheme with backward-leak uni-directional key updates and uni-directional ciphertext updates?*

We also affirmatively answer the second question.

1.2 Our Contribution

The first contribution of our work is a definitional work. We define a new definition of key updates, which we call backward-leak uni-directional key updates. In addition, we prove that UE with backward-leak uni-directional key updates is *strictly stronger* than bi-directional key updates (and forward-leak uni-directional key updates). More specifically, we show that there are UE schemes with bi-directional key updates that are not secure in the *backward-leak* uni-directional key updates setting. This is in sharp contrast to Jiang’s equivalence theorem [Jia20] explained above.

The second contribution is that we present a new constructions of UE. Our scheme is a UE scheme with backward-leak uni-directional key updates and secure under the learning with errors (LWE) assumption, which is known as a post-quantum assumption. This scheme satisfies confidentiality against CPA and ciphertext updates are randomized. This is the first UE schemes with stronger key updates. Note that our scheme provides uni-directional ciphertext updates (i.e., cannot downgrade ciphertext into older epoch ones). Our scheme is implementable since it is directly constructed from lattices. The proposed scheme satisfies r-IND-UE-CPA security, which was defined by Boyd, Davies, Gjøsteen, and Jiang [BDGJ20]. However, we consider the backward-leak uni-directional setting. See Section 3 for the definitions.

1.3 Related Work

We often use “forward-leak uni-/backward-leak uni-/bi-/no-directional UE” to refer to UE with forward-leak uni-/backward-leak uni-/bi-/no-directional key updates in this paper. No-directional key updates mean that even if adversaries have k_{e+1} (resp. k_e) and Δ_{e+1} , they cannot infer k_e (resp. k_{e+1}) [Jia20].

Ciphertext-independent updates UE. Lehman and Tackmann introduce post-compromise security for UE and refine previous security notions. Those are close to the definitions in this paper. They also present an efficient *bi-directional* UE scheme based on the DDH assumption [LT18]. Klooß, Lehmann, and Rupp present a CCA-secure *bi-directional* UE scheme based on the DDH assumption in the ROM and RCCA-secure *bidirectional* UE schemes based on the SXDH assumption [KLR19]. Boyd et al. integrate and refine previous security notions and present CCA-secure *bi-directional* UE schemes with deterministic ciphertext updates based on the DDH assumption in the ideal cipher model [BDGJ20]. Jiang studies relationships among various models for UE and presents a *bi-directional* UE scheme based on the LWE assumption [Jia20]. All these schemes provide bi-directional ciphertext updates (a token enables us to update and downgrade a ciphertext).

Ciphertext-dependent updates UE. Boneh, Lewi, Montgomery, and Raghunathan introduce the notion of UE in the ciphertext-dependent updates setting and present a *bi-directional* UE scheme based on key homomorphic PRFs [BLMR13]. Everspaugh, Paterson, Ristenpart, and Scott define stronger security notions for UE and present *bi-directional* UE schemes that satisfy those notions [EPRS17]. Chen, Li, and Tang introduce a stronger CCA security notion by considering malicious re-encryption attacks and present *bi-directional* UE schemes that satisfy the stronger CCA security [CLT20]. Boneh, Eskandarian, Kim, and Shih improve security notions by Everspaugh et al. [EPRS17] and present efficient *bi-directional* UE schemes [BEKS20].

UE in constructive cryptography. Levy-dit-Vehel and Roméas study security notions for UE in the constructive cryptography framework and explore the right security notion for UE [LR21]. Fabrega, Maurer, and Mularczyk also study security notions for UE in the constructive cryptography framework, generalize previous definitions, and discover new security-efficiency trade-offs. [FMM21].

Concurrent and independent work. Slamanig and Striecks [SS21] concurrently and independently proposed a UE scheme.² Their scheme is a pairing-based no-directional scheme. They define a stronger model for UE, where we can set an expiry epoch e_{\perp} to a ciphertext. If we update a ciphertext with expiry epoch e_{\perp} by using a token Δ_{e+1} such that $e + 1 > e_{\perp}$, the updated ciphertext can no longer be decrypted. Due to this stronger model, Jiang’s equivalence theorem [Jia20] does not necessarily hold. The scheme provides uni-directional ciphertext updates. The sharp differences between their work and ours are as follows. Let T be the maximum number of epochs. Their no-directional scheme is secure *in the expiry model under the SXDH assumption*, and the ciphertext and key size are $O(\log^2 T)$ and $O(\log^2 T)$, respectively. Our uni-directional scheme is *post-quantum secure with backward-leak* key updates, and the ciphertext and key size do *not depend on T* .

1.4 Technical Overview

In this section, we present a high-level overview of our technique.

Direction of key updates. As we introduce in Section 1.1, we can consider two types of uni-directional tokens, forward-leak and backward-leak uni-directional tokens. If we can infer in both directions, we call bi-directional token. In the definitions of confidentiality for UE, trivial winning conditions of adversaries depend on those token variations.

We show the following adversary against existing bi-directional UE schemes: (1) s/he triggers the trivial winning condition of the forward-leak uni-directional key updates setting. (2) s/he does not trigger

²Their paper [SS21] appeared on Cryptology ePrint archive right after the initial version of this paper (<https://eprint.iacr.org/2021/221/20210311:210911>) appeared on Cryptology ePrint archive. The comparison here is based on the latest versions of their and our papers.

the trivial winning condition of the backward-leak uni-directional key updates. (3) s/he trivially breaks confidentiality of the schemes in the backward-leak uni-directional key updates. Therefore, existing bi-directional UE schemes are not secure in the backward-leak uni-directional key updates setting. The best way to understand the separation result is looking at an example described in Section 4.3.

In this section, we explain the source of the difference between the two settings. First, we recall that UE needs the power of public key encryption (PKE) such as the DDH assumption. We can find this fact in all existing ciphertext-independent UE schemes [LT18, KLR19, BDGJ20, Jia20]. Alapati, Montgomery, and Patranabis [AMP19] prove that ciphertext-independent UE implies PKE. By this fact, we can assume that an epoch key k_e consists of a secret part sk_e and a public key part pk_e . As an example, in RISE scheme [LT18], $sk_e = x_e \in \mathbb{Z}_p$, $pk_e = g^{x_e} \in \mathbb{G}$, and $\Delta_{e+1} = x_{e+1}/x_e$ where g is a generator of a prime-order group \mathbb{G} . It is easy to see the token is a bi-directional token.

The direction of key updates depends on how to generate a token. A simple but crucial observation is that we must use sk_e to generate Δ_{e+1} . Otherwise, Δ_{e+1} does not have the power of decrypting and converting a ciphertext at epoch e . On the other hand, we do not necessarily need sk_{e+1} to generate Δ_{e+1} since we can generate a ciphertext at epoch $e + 1$ by using pk_{e+1} .

The relation between the direction types and how to generate a token is as follows. A forward-leak uni-directional token means Δ_{e+1} explicitly contains information about sk_{e+1} . By combining the observation above, Δ_{e+1} should contain information about sk_e and sk_{e+1} in the forward-leak uni-directional key updates setting. In addition, we can update an older epoch ciphertext into a newer epoch ciphertext and attack the new one if the newer epoch key is revealed. In other words, we can attack older epoch ciphertext even if older epoch keys are not revealed (backward-leak inference is not possible in this setting). The key inference direction could be the same as the ciphertext update direction. By this observation, it is natural that Jiang’s equivalence theorem holds.

On the other hand, a backward-leak uni-directional token means Δ_{e+1} explicitly contains information about sk_e . It is possible to generate Δ_{e+1} from sk_e and pk_{e+1} based on the observations so far. Thus, a backward-leak uni-directional token could hide information about sk_{e+1} and prevent the forward inference. In addition, this property prevents downgrading a ciphertext into an older epoch ciphertext. Thus, even if an older epoch key is revealed, we cannot necessarily attack the newer epoch ciphertexts since downgrading ciphertext and forward-leak inference are impossible. The key inference direction is opposite to the ciphertext update direction. This property is in sharp contrast to the forward-leak setting. Therefore, triggers of trivial winning conditions are different in these two settings. An intuition behind our separation result is based on those observations. See Section 4.3 for the detail. Those observations are the starting points of our UE scheme in the backward-leak uni-directional key updates setting. See the next paragraph for an overview.

Our backward-leak uni-directional key updates scheme. Roughly speaking, a token Δ_{e+1} is a homomorphic encryption of sk_e under a public key pk_{e+1} in our backward-leak uni-directional UE scheme. To update a ciphertext $ct_e \leftarrow \text{Enc}(pk_e, \mu)$ at epoch e , we homomorphically decrypt ct_e by using $\Delta_{e+1} = \text{Enc}(pk_{e+1}, sk_e)$ and obtain $\text{Enc}(pk_{e+1}, \mu)$. It is easy to see that if we have Δ_{e+1} and sk_{e+1} , we can obtain sk_e by decryption. However, it is difficult to infer sk_{e+1} from Δ_{e+1} and sk_e since sk_{e+1} is not used to generate Δ_{e+1} . By the security of PKE, it is difficult to obtain sk_{e+1} from pk_{e+1} . To achieve confidentiality for UE, we need to re-randomize tokens and updated ciphertext. This is also possible by using the homomorphic property. Although we use the homomorphic property of lattice-based encryption in our construction, we do not need fully homomorphic encryption (FHE). We use the key-switching technique [BV14, BV11] and the noise smudging technique [AJL⁺12] to directly achieve secure UE from the LWE assumption. This idea is inspired by uni-directional proxy re-encryption schemes based on lattices [Gen09, ABPW13, CCL⁺14, NX15].

To prove confidentiality, we need to erase information about sk_{e^*} where e^* is the target epoch (otherwise, we cannot use confidentiality under pk_{e^*}). However, secret keys are linked to update tokens.

Thus, we need to gradually erase secret keys in update tokens from new ones to old ones. That is, we change $\text{Enc}(\text{pk}_{e+1}, \text{sk}_e)$ into $\text{Enc}(\text{pk}_{e+1}, 0^{|\text{sk}_e|})$. Once this change is done, we can change $\text{Enc}(\text{pk}_e, \text{sk}_{e-1})$ into $\text{Enc}(\text{pk}_e, 0^{|\text{sk}_{e-1}|})$, and so forth. Note that there exists an epoch e_r where Δ_{e_r+1} is not corrupted such that $e^* \leq e_r$ as long as adversaries do not trigger the trivial winning conditions. We can start the erasing process from e_r since sk_{e_r} is not used anywhere. This proof outline is reminiscent of the proof technique for multi-hop universal proxy re-encryption [DN21].

Organization. In Section 2, we provide preliminaries and basic definitions. In Section 3, we review the syntax and security definitions of UE. Section 4 defines a new definition of uni-directional key updates (backward-leak uni-directional key updates) and shows that it is strictly stronger than those of bi-directional and forward-leak uni-directional key updates. In Section 5, we present our UE scheme with backward-leak uni-directional key updates based on the LWE problem and prove its security.

2 Preliminaries

We define some notations and introduce cryptographic notions in this section.

Notations. In this paper, $x \leftarrow X$ denotes selecting an element from a finite set X uniformly at random, and $y \leftarrow A(x)$ denotes assigning to y the output of a probabilistic or deterministic algorithm A on an input x . When we explicitly show that A uses randomness r , we write $y \leftarrow A(x; r)$. For a finite set S , $U(S)$ denotes the uniform distribution over S . For strings x and y , $x||y$ denotes the concatenation of x and y . Let $[\ell]$ and $[\ell, r]$ denote the set of integers $\{1, \dots, \ell\}$ and $\{\ell, \dots, r\}$, respectively, λ denote a security parameter, and $y := z$ denote that y is set, defined, or substituted by z .

We say that a distribution over \mathbb{R} is B -bounded if a sample from the distribution is in $[-B, B]$ with overwhelming probability. For $x \in \mathbb{R}$, let $\lfloor x \rfloor := \lfloor x - 1/2 \rfloor$. For $\mathbf{x} = (x_1, \dots, x_\ell) \in \mathbb{R}^\ell$, let $\lfloor \mathbf{x} \rfloor := (\lfloor x_1 \rfloor, \dots, \lfloor x_\ell \rfloor) \in \mathbb{Z}^\ell$. For any integer $q \geq 2$, we write \mathbb{Z}_q for the ring $\{\lfloor -q/2 \rfloor, \dots, -1, 0, 1, \dots, \lfloor q/2 \rfloor\}$ with addition and multiplication modulo q .

For two matrices $\mathbf{X} \in \mathbb{R}^{m \times n_1}$ and $\mathbf{Y} \in \mathbb{R}^{m \times n_2}$, $[\mathbf{X} \mid \mathbf{Y}] \in \mathbb{R}^{m \times (n_1+n_2)}$ denotes the concatenation of the columns of \mathbf{X} and \mathbf{Y} . For two matrices $\mathbf{X} \in \mathbb{R}^{m_1 \times n}$ and $\mathbf{Y} \in \mathbb{R}^{m_2 \times n}$, $[\mathbf{X}; \mathbf{Y}] \in \mathbb{R}^{(m_1+m_2) \times n}$ denotes the concatenation of the rows of \mathbf{X} and \mathbf{Y} . For a vector $\mathbf{x} \in \mathbb{R}^m$, $\|\mathbf{x}\|_p$ denotes the ℓ_p norm of \mathbf{x} . We omit subscript if $p = 2$ for simplicity. PPT stands for probabilistic polynomial time.

Basic concepts.

- A function $f : \mathbb{N} \rightarrow \mathbb{R}$ is a negligible function if for any constant c , there exists $\lambda_0 \in \mathbb{N}$ such that for any $\lambda > \lambda_0$, $f(\lambda) < \lambda^{-c}$. We write $f(\lambda) \leq \text{negl}(\lambda)$ to denote $f(\lambda)$ being a negligible function.
- If $\mathcal{X}^{(b)} = \{X_\lambda^{(b)}\}_{\lambda \in \mathbb{N}}$ for $b \in \{0, 1\}$ are two ensembles of random variables indexed by $\lambda \in \mathbb{N}$, we say that $\mathcal{X}^{(0)}$ and $\mathcal{X}^{(1)}$ are computationally indistinguishable if for any PPT distinguisher \mathcal{D} , there exists a negligible function $\text{negl}(\lambda)$, such that

$$\Delta := |\Pr[\mathcal{D}(X_\lambda^{(0)}) = 1] - \Pr[\mathcal{D}(X_\lambda^{(1)}) = 1]| \leq \text{negl}(\lambda).$$

We write $\mathcal{X}^{(0)} \stackrel{s}{\approx} \mathcal{X}^{(1)}$ to denote that the advantage Δ is negligible.

- The statistical distance between $\mathcal{X}^{(0)}$ and $\mathcal{X}^{(1)}$ over a countable set S is defined as $\Delta_s(\mathcal{X}^{(0)}, \mathcal{X}^{(1)}) := \frac{1}{2} \sum_{\alpha \in S} |\Pr[X_\lambda^{(0)} = \alpha] - \Pr[X_\lambda^{(1)} = \alpha]|$. We say that $\mathcal{X}^{(0)}$ and $\mathcal{X}^{(1)}$ are statistically/perfectly indistinguishable (denoted by $\mathcal{X}^{(0)} \stackrel{s}{\approx} \mathcal{X}^{(1)}/\mathcal{X}^{(0)} \stackrel{p}{\approx} \mathcal{X}^{(1)}$) if $\Delta_s(\mathcal{X}^{(0)}, \mathcal{X}^{(1)}) \leq \text{negl}(\lambda)$ and $\Delta_s(\mathcal{X}^{(0)}, \mathcal{X}^{(1)}) = 0$, respectively.

2.1 Lattice Preliminaries

Distributions. We review the basic definitions of the distributions in lattice-based cryptography. Let $N(0, \sigma^2)$ denote the Gaussian distribution whose mean is 0 and variance is σ^2 . The Gaussian distribution is defined by density function $(1/\sigma\sqrt{2\pi}) \cdot \exp(-x^2/2\sigma^2)$ over \mathbb{R} . We also define the followings.

- *Discretized Gaussian* Ψ_α : For $\alpha \in (0, 1)$ and a positive integer q , we sample x from $N(0, \alpha^2/2\pi)$ and output $\lfloor qx \rfloor \bmod q$.
- *Discrete Gaussian*: For a positive real s , the n -dimensional Gaussian function is defined as $\rho_s(\mathbf{x}) = \exp(-\pi\|\mathbf{x}\|^2/s^2)$.
- *Discrete Gaussian distribution* $D_{A,s}$: For a positive real s and a countable set A , $D_{A,s}(\mathbf{x}) := \frac{\rho_s(\mathbf{x})}{\sum_{\mathbf{y} \in A} \rho_s(\mathbf{y})}$.

Gentry, Peikert, and Vaikuntanathan [GPV08] gave an efficient sampler, SampleD, for $D_{\mathbb{Z},s}$.

We use the following statistical properties in this paper.

Lemma 2.1 ([Ban93, Lemma 1.5], [LP11, Lemma 2.1]). *Let $c \geq 1$, $C = c \cdot \exp((1 - c^2)/2)$. For any real $s > 0$ and any integer $n \geq 1$, we have that*

$$\Pr_{\mathbf{e} \leftarrow D_{\mathbb{Z}^n, s}} [\|\mathbf{e}\|_2 \geq cs\sqrt{n/(2\pi)}] \leq C^n.$$

In particular, letting $c = \sqrt{2\pi}$ and $C < 1/4$, we have that $\Pr_{\mathbf{e} \leftarrow D_{\mathbb{Z}^n, s}} [\|\mathbf{e}\|_2 \geq s\sqrt{n}] < 2^{-2n}$.

Lemma 2.2 ([AJL⁺12, Lemma 1 (smudging lemma)]. *Let $B_1 = B_1(\lambda)$, and $B_2 = B_2(\lambda)$ be positive integers and let $e_1 \in [-B_1, B_1]$ be a fixed integer. Let $e_2 \leftarrow [-B_2, B_2]$ be chosen uniformly at random. If $B_1/B_2 = \text{negl}(\lambda)$, the distribution of e_2 is statistically indistinguishable from that of $e_2 + e_1$.*

Lemma 2.3 (Adapted version of the leftover hash lemma). *Let q be an odd prime. Let D be a distribution over \mathbb{Z}_q^m of min-entropy at least $(n + \ell) \lg q + g(n)$. Then, we have that $\Delta((A, \mathbf{e}A), (A, \mathbf{u}^\top)) \leq 2^{-g(n)/2}$, where $A \leftarrow \mathbb{Z}_q^{m \times (n+\ell)}$, $\mathbf{e} \leftarrow D$, and $\mathbf{u} \leftarrow \mathbb{Z}_q^{n+\ell}$*

This lemma holds if we use $E \leftarrow D^k$ and $\mathbf{U} \leftarrow \mathbb{Z}_q^{k \times (n+\ell)}$ instead of \mathbf{e} and \mathbf{u} where $k = \text{poly}(n)$.

Learning with Errors. Regev introduced the LWE problem [Reg09]. Let $A(\mathbf{s}, \chi)$ be a distribution over $\mathbb{Z}_q^n \times \mathbb{Z}_q$ defined as follows. For a vector $\mathbf{s} \in \mathbb{Z}_q^n$ and a distribution χ over \mathbb{Z}_q , we sample $\mathbf{a} \leftarrow \mathbb{Z}_q^n$ and $x \leftarrow \chi$, and output $(\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle + x)$.

Definition 2.4 (The LWE problem and assumption). *For an integer $q = q(n)$, and distributions χ over \mathbb{Z}_q and ψ over \mathbb{Z}_q^n , the learning with errors problem, $\text{LWE}(n, q, \chi)$ for the distribution ψ , is distinguishing oracle $A(\mathbf{s}, \chi)$ from oracle $A(\mathbf{s}, U(\mathbb{Z}_q))$, where $\mathbf{s} \leftarrow \psi$. We say the $\text{LWE}(n, q, \chi)$ assumption holds for ψ if for any PPT adversary \mathcal{A} , its advantage*

$$\text{Adv}_{\mathcal{A}, (n, q, \chi, \psi)}^{\text{lwe}}(n) = \left| \Pr[\mathcal{A}^{A(\mathbf{s}, \chi)}(1^n) = 1 \mid \mathbf{s} \leftarrow \psi] - \Pr[\mathcal{A}^{A(\mathbf{s}, U(\mathbb{Z}_q))}(1^n) = 1 \mid \mathbf{s} \leftarrow \psi] \right|$$

is negligible in n where $\mathbf{s} \leftarrow \psi$.

We note that $A(\mathbf{s}, U(\mathbb{Z}_q)) = U(\mathbb{Z}_q^n \times \mathbb{Z}_q)$ for any $\mathbf{s} \in \mathbb{Z}_q^n$. We also note that even if we use matrices $A \leftarrow \mathbb{Z}_q^{m \times n}$ and $S \in \mathbb{Z}_q^{n \times \ell}$ instead of vectors \mathbf{a} and \mathbf{s} , respectively, the assumption holds (if the original assumption holds) by a simple hybrid argument. In that case, we omit parameter m, ℓ from the notation for simplicity. Applebaum, Cash, Peikert, and Sahai [ACPS09] showed that we can use χ^n instead of $U(\mathbb{Z}_q)$ for the distribution of \mathbf{s} .

Theorem 2.5 (Adapted version of [ACPS09, Lemma 2]). *Let $q = p^e$ be a prime power. If the $\text{LWE}(n, q, \chi)$ assumption holds for $U(\mathbb{Z}_q^n)$, then the $\text{LWE}(n, q, \chi)$ assumption holds for χ^n .*

Solving the LWE problem with $\chi = \bar{\Psi}_\alpha$ or $D_{\mathbb{Z},s}$ on average is as hard as the worst case of the approximation version of the shortest independent vector problem, SIVP_γ , and the decision version of the shortest vector problem, GapSVP_γ , under a classical/quantum reduction, where γ is an approximation factor [Reg09, Pei09, BLP⁺13].

3 Updatable Encryption

In this section, we briefly review the syntax and definitions of UE.

Syntax.

Definition 3.1. *An updatable encryption scheme UE for message space \mathcal{M} consists of a tuple of PPT algorithms $(\text{UE.Setup}, \text{UE.KeyGen}, \text{UE.Enc}, \text{UE.Dec}, \text{UE.TokGen}, \text{UE.Upd})$.*

$\text{UE.Setup}(1^\lambda) \rightarrow \text{pp}$: *The setup algorithm takes as input the security parameter and outputs a public parameter pp. (This algorithm is an option for UE.)*

$\text{UE.KeyGen}(\text{pp}) \rightarrow k_e$: *The key generation algorithm takes as input the public parameter and outputs an epoch key k_e .*

$\text{UE.Enc}(k, \mu) \rightarrow \text{ct}$: *The encryption algorithm takes as input an epoch key and a plaintext μ and outputs a ciphertext ct.*

$\text{UE.Dec}(k, \text{ct}) \rightarrow \mu'$: *The decryption algorithm takes as input an epoch key and a ciphertext and outputs a plaintext μ' or \perp .*

$\text{UE.TokGen}(k_e, k_{e+1}) \rightarrow \Delta_{e+1}$: *The token generation algorithm takes as input two keys of successive epochs e and $e + 1$ and outputs a token Δ_{e+1} .*

$\text{UE.Upd}(\Delta_{e+1}, \text{ct}_e) \rightarrow \text{ct}_{e+1}$: *The update algorithm takes as input a token Δ_{e+1} and a ciphertext ct_e and outputs a ciphertext ct_{e+1} .*

Let T be the maximum number of the epoch.

Security experiments. We review security definitions for UE in this section.

Definition 3.2 (Correctness). *For any $\mu \in \mathcal{M}$, for $0 \leq e_1 \leq e_2 \leq T$, it holds that*

$$\Pr[\text{UE.Dec}(k_{e_2}, \text{ct}_{e_2}) \neq \mu] \leq \text{negl}(\lambda),$$

where $\text{pp} \leftarrow \text{UE.Setup}(1^\lambda)$, $k_{e_1}, \dots, k_{e_2} \leftarrow \text{UE.KeyGen}(\text{pp})$, $\text{ct}_{e_1} \leftarrow \text{UE.Enc}(k_{e_1}, \mu)$, and $\Delta_{i+1} \leftarrow \text{UE.TokGen}(k_i, k_{i+1})$, $\text{ct}_{i+1} \leftarrow \text{UE.Upd}(\Delta_{i+1}, \text{ct}_i)$ for $i \in [e_1, e_2 - 1]$.

Definition 3.3 (Confidentiality for Updatable Encryption [BDGJ20, Jia20]). *For $x \in \{d, r\}$, $\text{atk} \in \{\text{cpa}, \text{cca}\}$, the game $\text{Exp}_{\Sigma, \mathcal{A}}^{x\text{-ind-ue-atk}}(\lambda, b)$ is formalized as follows.*

- Invoke Setup and set $\text{phase} := 0$.
- Let $\mathcal{O} := \mathcal{O}.\{\text{Enc}, \text{Next}, \text{Upd}, \text{Corr}, \text{Chall}, \text{Upd}\tilde{\mathcal{C}}\}$ if $\text{atk} = \text{cpa}$. If $\text{atk} = \text{cca}$, $\mathcal{O}.\text{Dec}$ is also added in \mathcal{O} .

- Run $\text{coin}' \leftarrow \mathcal{A}^{\mathcal{O}}(1^\lambda)$.
- If $((\mathcal{K}^* \cap \mathcal{C}^* \neq \emptyset) \vee (x = d \wedge (e^* \in \mathcal{T}^* \vee \mathcal{O}.\text{Upd}(\overline{\text{ct}})$ is invoked))) then $\text{twf} := 1$
- If $\text{twf} = 1$ then $\text{coin}' \leftarrow \{0, 1\}$
- return coin'

We say a UE scheme is x -IND-UE-atk secure if it holds

$$\text{Adv}_{\Sigma, \mathcal{A}}^{x\text{-ind-ue-atk}}(\lambda) := |\Pr[\text{Exp}_{\Sigma, \mathcal{A}}^{x\text{-ind-ue-atk}}(\lambda, 0) = 1] - \Pr[\text{Exp}_{\Sigma, \mathcal{A}}^{x\text{-ind-ue-atk}}(\lambda, 1) = 1]| \leq \text{negl}(\lambda).$$

The definitions of oracles are described in Figure 1.

The prefix d and r in the definition above indicate that we consider UE schemes with *deterministic* and *randomized* update algorithms, respectively.

Leakage sets. We introduce leakage sets. Adversaries can obtain secret keys, update tokens, challenge-equal ciphertexts from oracles. We record epochs in the following sets to maintain which epoch key/token/challenge-equal-ciphertext was given to adversaries.

- \mathcal{K} : Set of epochs where \mathcal{A} corrupted the epoch key via $\mathcal{O}.\text{Corr}$.
- \mathcal{T} : Set of epochs where \mathcal{A} corrupted the update token via $\mathcal{O}.\text{Corr}$.
- \mathcal{C} : Set of epochs where \mathcal{A} obtained a challenge-equal ciphertext via $\mathcal{O}.\text{Chall}$ or $\mathcal{O}.\text{Upd}\tilde{\mathcal{C}}$.

We also record ciphertexts given via oracles to maintain which (updated) ciphertexts adversaries obtained.

- \mathcal{L} : Set of non-challenge ciphertexts $(\text{cnt}, \text{ct}, e; \mu)$ returned via $\mathcal{O}.\text{Enc}$ or $\mathcal{O}.\text{Upd}$, where cnt is a query index incremented by each invocation of $\mathcal{O}.\text{Enc}$, ct is the given ciphertext, e is the epoch where the query happens, and μ is the queried plaintext or the plaintext in the queried ciphertext.
- $\tilde{\mathcal{L}}$: Set of challenge-equal ciphertexts (ct_e^*, e) returned via $\mathcal{O}.\text{Chall}$ or $\mathcal{O}.\text{Upd}\tilde{\mathcal{C}}$, where ct_e^* is the given challenge-equal ciphertext and e is the epoch where the query happens.

In the deterministic update setting, where algorithm Upd is deterministic, an updated ciphertext is uniquely determined by a token and a ciphertext. Thus, we consider extended ciphertext sets \mathcal{L}^* and $\tilde{\mathcal{L}}^*$ inferred from \mathcal{L} and $\tilde{\mathcal{L}}$, respectively, by using \mathcal{T} . Regarding \mathcal{L}^* , we only need information about the ciphertext and epoch. That is, \mathcal{L}^* consists of sets of a ciphertext and an epoch index.

In the randomized update setting, where algorithm Upd is probabilistic, an update ciphertext is not uniquely determined. Thus, we consider sets of plaintexts of which adversaries have ciphertexts.

- \mathcal{Q}^* : Set of plaintexts (μ, e) such that the adversary obtained or could generate a ciphertext of μ at epoch e .
- $\tilde{\mathcal{Q}}^*$: Set of challenge plaintexts $\{(\bar{\mu}, e), (\bar{\mu}_1, e)\}$, where $(\bar{\mu}, \overline{\text{ct}})$ is the query to $\mathcal{O}.\text{Chall}$ and $\bar{\mu}_1$ is the plaintext in $\overline{\text{ct}}$. The adversary obtained or could generate a challenge-equal ciphertext of $\bar{\mu}$ or $\bar{\mu}_1$ at epoch e .

Inferred leakage sets. Lehman and Tackmann [LT18] presented the bookkeeping technique to analyze the epoch leakage sets. We maintain leaked information by the technique in security games.

Setup(1^λ):

- $k_0 \leftarrow \text{UE.KeyGen}(1^\lambda)$
- $\Delta_0 := \perp; e, \text{cnt}, \text{twf} := 0$
- $\mathcal{L}, \tilde{\mathcal{L}}, \mathcal{C}, \mathcal{K}, \mathcal{T} := \emptyset$

$\mathcal{O}.\text{Enc}(\mu)$:

- $\text{cnt} := \text{cnt} + 1$
- $\text{ct} \leftarrow \text{UE.Enc}(k_e, \mu)$
- $\mathcal{L} := \mathcal{L} \cup \{(\text{cnt}, \text{ct}, e; \mu)\}$
- **return** ct

$\mathcal{O}.\text{Dec}(\text{ct})$:

- $\mu' / \perp \leftarrow \text{UE.Dec}(k_e, \text{ct})$
- **if** $\left((x = d \wedge (\text{ct}, e) \in \tilde{\mathcal{L}}^*) \vee (x = r \wedge (\mu', e) \in \tilde{\mathcal{Q}}^*) \right)$
then $\text{twf} := 1$
- **return** μ' **or** \perp

$\mathcal{O}.\text{Next}()$:

- $e := e + 1$
- $k_e \leftarrow \text{UE.KeyGen}(1^\lambda)$
- $\Delta_e \leftarrow \text{UE.TokGen}(k_{e-1}, k_e)$
- **if** $\text{phase} = 1$
then $\text{ct}_e^* \leftarrow \text{UE.Upd}(\Delta_e, \text{ct}_{e-1}^*)$

$\mathcal{O}.\text{Upd}(\text{ct}_{e-1})$:

- **if** $(j, \text{ct}_{e-1}, e - 1; \mu) \notin \mathcal{L}$
then return \perp
- $\text{ct}_e \leftarrow \text{UE.Upd}(\Delta_e, \text{ct}_{e-1})$
- $\mathcal{L} := \mathcal{L} \cup \{(\text{cnt}, \text{ct}_e, e; \mu)\}$
- **return** ct_e

$\mathcal{O}.\text{Corr}(\text{mode}, \hat{e})$:

- **if** $\hat{e} > e$ **then return** \perp
- **if** $\text{mode} = \text{key}$
then $\mathcal{K} := \mathcal{K} \cup \{\hat{e}\}$
return $k_{\hat{e}}$
- **if** $\text{mode} = \text{token}$
then $\mathcal{T} := \mathcal{T} \cup \{\hat{e}\}$
return $\Delta_{\hat{e}}$

$\mathcal{O}.\text{Chall}(\bar{\mu}, \bar{\text{ct}})$:

- **if** $\text{phase} = 1$ **then return** \perp
- $\text{phase} := 1; e^* := e$
- **if** $(\cdot, \bar{\text{ct}}, e^* - 1; \bar{\mu}_1) \notin \mathcal{L}$
then return \perp
- **if** $b = 0$
then $\text{ct}_{e^*}^* \leftarrow \text{UE.Enc}(k_{e^*}, \bar{\mu})$
else $\text{ct}_{e^*}^* \leftarrow \text{UE.Upd}(\Delta_{e^*}, \bar{\text{ct}})$
- $\mathcal{C} := \mathcal{C} \cup \{e^*\}$
- $\tilde{\mathcal{L}} := \tilde{\mathcal{L}} \cup \{(\text{ct}_{e^*}^*, e^*)\}$
- **return** $\text{ct}_{e^*}^*$

$\mathcal{O}.\text{Upd}\tilde{\mathcal{C}}()$:

- **if** $\text{phase} \neq 1$ **then return** \perp
- $\mathcal{C} := \mathcal{C} \cup \{e\}$
- $\tilde{\mathcal{L}} := \tilde{\mathcal{L}} \cup \{(\text{ct}_e^*, e)\}$
- **return** ct_e^*

$\mathcal{O}.\text{Try}(\text{ct}^*)$:

- $\mu' / \perp \leftarrow \text{UE.Dec}(k_e, \text{ct}^*)$
- **if** $(e \in \mathcal{K}^* \vee (\text{atk} = \text{ctxt} \wedge (\text{ct}^*, e) \in \mathcal{L}^*) \vee (\text{atk} = \text{ptxt} \wedge (\mu', e) \in \mathcal{Q}^*))$
then $\text{twf} := 1$
- **if** $\mu' \neq \perp$ **then** $\text{win} := 1$

Figure 1: The behavior of oracles in security experiments for updatable encryption. Leakages sets $\mathcal{L}, \tilde{\mathcal{L}}, \mathcal{L}^*, \tilde{\mathcal{L}}^*, \mathcal{C}, \mathcal{K}, \mathcal{K}^*, \mathcal{T}, \mathcal{T}^*, \mathcal{Q}, \mathcal{Q}^*, \tilde{\mathcal{Q}}^*$ are defined in Section 3.

Key leakage. Adversaries can infer some information from leakage sets \mathcal{K} and \mathcal{T} . Here, “infer” means that adversaries can trivially extract some secret information from given keys and tokens. For example, in the ElGamal-based UE scheme by Lehman and Tackmann (called RISE) [LT18], a secret key at epoch e is $k_e \in \mathbb{Z}_p$ where p is a prime and a token is $\Delta_{e+1} = k_{e+1}/k_e \in \mathbb{Z}_p$. Thus, we can easily extract k_e from Δ_{e+1} and k_{e+1} (and vice versa).

Inferred information depends on the direction of key updates. In previous works on UE, there are three types of directions of key updates, called bi/uni/no-directional key updates. Formally, for $\text{kk} \in \{\text{no}, \text{uni}, \text{bi}\}$, we consider the following kk -directional key update setting.

Definition 3.4 (Direction of Key Update). We define inferred leakage key sets. The sets depend on the setting of key updates.

- No-directional key updates: $\mathcal{K}_{\text{no}}^* := \mathcal{K}$.
- Uni-directional key updates:

$$\mathcal{K}_{\text{uni}}^* := \{e \in [0, \ell] \mid \text{CorrK}(e) = \text{true}\}$$

$$\text{where } \text{CorrK}(e) = \text{true} \Leftrightarrow (e \in \mathcal{K}) \vee (\text{CorrK}(e-1) \wedge e \in \mathcal{T})$$

- Bi-directional key updates:

$$\mathcal{K}_{\text{bi}}^* := \{e \in [0, \ell] \mid \text{CorrK}(e) = \text{true}\}$$

$$\text{where } \text{CorrK}(e) = \text{true} \Leftrightarrow (e \in \mathcal{K}) \vee (\text{CorrK}(e-1) \wedge e \in \mathcal{T}) \vee (\text{CorrK}(e+1) \wedge e+1 \in \mathcal{T})$$

Token leakage. If two successive keys are leaked, a token generated from those keys is also inferred.

Definition 3.5 (Inferred Token Sets). For $\text{kk} \in \{\text{no}, \text{uni}, \text{bi}\}$,

$$\mathcal{T}_{\text{kk}}^* := \{e \in [0, \ell] \mid (e \in \mathcal{T}) \vee (e \in \mathcal{K}_{\text{kk}}^* \wedge e-1 \in \mathcal{K}_{\text{kk}}^*)\}$$

Challenge-equal ciphertext leakage. We can update ciphertexts by using tokens. That is, we can obtain updated ciphertexts generated from a challenge ciphertext via leaked tokens. To check whether a challenge ciphertext can be converted into a ciphertext under a corrupted key, we maintain challenge-equal ciphertext epochs defined below.

Definition 3.6 (Direction of Ciphertext Update). We define two types of challenge-equal ciphertext epoch sets. For $\text{kk} \in \{\text{no}, \text{uni}, \text{bi}\}$,

- Uni-directional ciphertext updates:

$$\mathcal{C}_{\text{kk}, \text{uni}}^* := \{e \in [0, \ell] \mid \text{ChallEq}(e) = \text{true}\}$$

$$\text{where } \text{ChallEq}(e) = \text{true} \Leftrightarrow (e \in \mathcal{C}) \vee (\text{ChallEq}(e-1) \wedge e \in \mathcal{T}_{\text{kk}}^*)$$

- Bi-directional ciphertext updates:

$$\mathcal{C}_{\text{kk}, \text{bi}}^* := \{e \in [0, \ell] \mid \text{ChallEq}(e) = \text{true}\}$$

$$\text{where } \text{ChallEq}(e) = \text{true} \Leftrightarrow (e \in \mathcal{C}) \vee (\text{ChallEq}(e-1) \wedge e \in \mathcal{T}_{\text{kk}}^*) \vee (\text{ChallEq}(e+1) \wedge e+1 \in \mathcal{T}_{\text{kk}}^*)$$

By considering directions of key/ciphertext updates, we can consider variants of security notions for UE [Jia20].

Definition 3.7 ((kk, cc)-variant of confidentiality [Jia20]). Let UE be a UE scheme. Then the (kk, cc)-notion advantage, for $\text{kk} \in \{\text{no}, \text{uni}, \text{bi}\}$, $\text{cc} \in \{\text{uni}, \text{bi}\}$ and $\text{notion} \in \{\text{r-ind-ue-cpa}, \text{d-ind-ue-cpa}, \text{r-ind-ue-cca}, \text{d-ind-ue-cca}\}$, of an adversary \mathcal{A} against UE is defined as

$$\text{Adv}_{\text{UE}, \mathcal{A}}^{(\text{kk}, \text{cc})\text{-notion}}(1^\lambda) := |\Pr[\text{Exp}_{\text{UE}, \mathcal{A}}^{(\text{kk}, \text{cc})\text{-notion}}(\lambda, 0) = 1] - \Pr[\text{Exp}_{\text{UE}, \mathcal{A}}^{(\text{kk}, \text{cc})\text{-notion}}(\lambda, 1) = 1]|,$$

where $\text{Exp}_{\text{UE}, \mathcal{A}}^{(\text{kk}, \text{cc})\text{-notion}}(\lambda, b)$ is the same as the experiment $\text{Exp}_{\text{UE}, \mathcal{A}}^{\text{notion}}(\lambda, b)$ in Definition 3.3 except for all leakage sets are both in the kk-directional key updates and cc-directional ciphertext updates.

Trivial winning condition. Adversaries trivially win the security game if we can convert a challenge ciphertext into a ciphertext under a corrupted key. Thus, we need to define trivial winning conditions.

For all confidentiality games in Definition 3.3, the trivial winning condition $\mathcal{K}^* \cap \mathcal{C}^* \neq \emptyset$ is checked since if the condition holds, adversaries can win the game by decrypting a challenge-equal ciphertext by using a corrupted key.

For all confidentiality games for deterministic update UE, the trivial winning condition $\tilde{e} \in \mathcal{T}^* \vee \text{“}\mathcal{O}.\text{Upd}(\bar{ct}) \text{ is queried”}$ is checked since if the condition holds, adversaries can win the game by checking the challenge ciphertext is equal to an updated ciphertext generated from the token and a queried ciphertext to $\mathcal{O}.\text{Chall}$.

We need to consider other trivial winning conditions in the CCA setting (both for randomized and deterministic updates) and integrity setting. However, we do not consider these settings in this work. We do not explain those conditions. See the paper by Jiang [Jia20] for the detail.

Firewall and insulated region.

Definition 3.8 (Firewall [LT18, KLR19, BDGJ20, Jia20]). *An insulated region with firewalls fwl and fwr is a consecutive sequence of epochs $[\text{fwl}, \text{fwr}]$ for which:*

- *No key in the sequence of epochs $[\text{fwl}, \text{fwr}]$ is corrupted. That is, it holds $[\text{fwl}, \text{fwr}] \cap \mathcal{K} = \emptyset$.*
- *The tokens Δ_{fwl} and $\Delta_{\text{fwr}+1}$ are not corrupted if they exist. That is, it holds $\text{fwl}, \text{fwr} + 1 \notin \mathcal{T}$.*
- *All tokens $(\Delta_{\text{fwl}+1}, \dots, \Delta_{\text{fwr}})$ are corrupted. That is, $[\text{fwl} + 1, \text{fwr}] \subseteq \mathcal{T}$.*

Definition 3.9 (Insulated Region [LT18, KLR19, BDGJ20, Jia20]). *The union of all insulated regions is defined as $\mathcal{IR} := \bigcup_{[\text{fwl}, \text{fwr}] \in \mathcal{FW}}$ $[\text{fwl}, \text{fwr}]$, where \mathcal{FW} is the set of insulated region with firewalls.*

On security definitions. In the r-IND-UE-CPA security experiment, adversaries try to distinguish a ciphertext of $\bar{\mu}$ from an updated ciphertext of \bar{ct} . This is different from the standard indistinguishability-based definition, r-IND-ENC-CPA, where adversaries try to distinguish a ciphertext of μ_0 from a ciphertext of μ_1 . Boyd et al. [BDGJ20] prove that r-IND-UE-CPA implies both r-IND-ENC-CPA and r-IND-UPD-CPA, where adversaries try to distinguish an updated ciphertext of \bar{ct}_0 from an updated ciphertext of \bar{ct}_1 . That is, r-IND-UE-CPA guarantees the standard CPA security and unlinkability of updated ciphertexts. In this paper, we present UE schemes with randomized update algorithms. Thus, we focus on the r-IND-UE-CPA security in this paper.

Note that a similar implication holds for the deterministic update algorithm case, but a d-IND-UE-CPA secure scheme must satisfy a few additional properties for the implication. See the paper by Boyd et al. [BDGJ20] for the detail.

4 Backward-Leak Uni-Directional Key Update and Relations

4.1 Definition

We introduce a new notion for the direction of key updates in this section. The notion is categorized in uni-directional key updates, but the direction is the opposite of the uni-directional key updates in Definition 3.4.

Definition 4.1 (Uni-Directional Key Update (revisited)). *We define two types of uni-directional key updates. One is the same as that in Definition 3.4. To distinguish two types of uni-directional key updates, we rename the original one in Definition 3.4 to forward-leak uni-directional key updates. The definitions of two notions are as follows.*

- forward-leak uni-directional key updates: $\mathcal{K}_{\text{f-uni}}^* := \mathcal{K}_{\text{uni}}^*$.
- backward-leak uni-directional key updates:

$$\mathcal{K}_{\text{b-uni}}^* := \{e \in [0, \ell] \mid \text{CorrK}(e) = \text{true}\}$$

where $\text{CorrK}(e) = \text{true} \Leftrightarrow (e \in \mathcal{K}) \vee (\text{CorrK}(e+1) \wedge e+1 \in \mathcal{T})$

By using the definition above, we can consider Definitions 3.5 and 3.6 for $\text{kk} \in \{\text{no}, \text{f-uni}, \text{b-uni}, \text{bi}\}$. We illustrate leaked information in the setting of forward/backward-leak uni-directional key updates settings in Figure 2.

set	e - 1	e	e + 1	set	e - 1	e	e + 1
$\mathcal{K}_{\text{f-uni}}^*$	×	✓	inferred	$\mathcal{K}_{\text{b-uni}}^*$	inferred	✓	×
$\mathcal{T}_{\text{f-uni}}^*$		✓	✓	$\mathcal{T}_{\text{b-uni}}^*$		✓	✓

Figure 2: Inferred keys in the forward-leak/backward-leak uni-directional key updates settings. Symbol ✓ means the key/token was given via $\mathcal{O}.\text{Corr}$. Symbol × means we cannot trivially obtain the information. The text “inferred” means we can trivially extract the information from given values.

4.2 Observations on Definitions

On the meaningfulness of backward-leak uni-directional key updates. First of all, all ciphertext-independent UE schemes rely on public key encryption power in some sense [LT18, BDGJ20, Jia20].³ This fact is endorsed by the result by Alarnati, Montgomery, and Patranabis [AMP19], which shows any ciphertext-independent UE scheme that is forward and post-compromise secure implies PKE. Thus, we can assume that an epoch key consists of a secret key part sk_e and a public key part pk_e .

To achieve the ciphertext update mechanism of UE, a token Δ_{e+1} must include information about sk_e since an update algorithm essentially decrypts a ciphertext at epoch e and generates a ciphertext for epoch $e+1$. The question is: “Do we really need sk_{e+1} for updating a ciphertext from e to $e+1$?”. The answer is no. The point is that we need only the public key part of an epoch key to generate a ciphertext in most existing ciphertext-independent UE schemes. Thus, we might be able to construct an update token by using only sk_e and pk_{e+1} . More specifically, we might be able to transform a ciphertext for epoch e by using encryption of sk_e under pk_{e+1} and homomorphic properties. This is what we do in Section 5. This insight comes from a few constructions of uni-directional proxy re-encryption [Gen09, ABPW13, CCL⁺14, NX15].

Based on the observations above, we can say the backward-leak uni-directional key updates setting is natural. If a token Δ_{e+1} is generated by using $(\text{sk}_e, \text{pk}_{e+1})$, it is likely we can infer sk_e from Δ_{e+1} and sk_{e+1} (our backward-leak uni-directional scheme is an example). However, it might be difficult to extract information about sk_{e+1} from sk_e and Δ_{e+1} since only pk_{e+1} is embedded in Δ_{e+1} . In fact, it is difficult in our backward-leak uni-directional scheme.

In the forward-leak uni-directional key updates setting, we assume that it is easy to infer sk_{e+1} from Δ_{e+1} and sk_e . In some sense, this says sk_{e+1} is directly embedded in Δ_{e+1} . We might be able to execute bi-directional key/ciphertext updates if a token enables us to update a ciphertext (in the forward direction). Here, “directly embedded” means that a secret key is not encrypted. In fact, in all existing

³Everspaugh et al. [EPRS17] presented a ciphertext-independent UE scheme from authenticated encryption (AE). However, they assume an AE scheme is secure against related key attacks. So far, it seems that we need the power of public key encryption (such as DDH) to achieve related key secure AE [HLL16]. In addition, Everspaugh et al. retracted the ciphertext-independent construction in their full version paper (<https://eprint.iacr.org/2017/527/20180903:192110>).

UE schemes bi-directional (and forward-leak uni-directional) key updates, sk_{e+1} is directly embedded in Δ_{e+1} [LT18, KLR19, BDGJ20, Jia20]. In addition, generating a token Δ_{e+1} from sk_{e+1} and pk_e is unnatural since it is unlikely such Δ_{e+1} can update a ciphertext under pk_e .

Note that the argument above does not consider obfuscation [BGI⁺12]. If we can somehow obfuscate secret keys in a token, it could be difficult to infer secret keys in the token even if we use those secret keys to generate the token.

As we argue in Section 1.1, backward-leak uni-directional key updates are more suitable than forward-leak ones in practice. In fact, we prove that confidentiality in the backward-leak uni-directional key updates setting is *strictly* stronger than that in the forward-leak uni-directional key updates setting.

On meaningful combination with bi/uni-directional ciphertext updates. For ciphertext updates, it is natural to consider only the uni-directional ciphertext updates in Definition 3.6 since updating ciphertext should go forward direction due to the nature of UE. Of course, we can define another uni-directional ciphertext updates (called “backward uni-directional” or “downgrade-only” ciphertext updates), but it is not meaningful.

Jiang considered a setting where key updates are uni-directional (this is forward-leak uni-directional by our definition) and ciphertext updates are bi-directional. This is meaningful only in the forward-leak uni-directional key updates since forward-leak uni-directional and bi-directional key updates are equivalent by Jiang’s result. However, it is unnatural to consider bi-directional ciphertext updates with *backward-leak* uni-directional key updates. This is because we show that backward-leak uni-directional key updates are strictly stronger than bi-directional key updates. In addition, it is difficult to use Δ_{e+1} to convert a ciphertext under k_{e+1} into one under k_e in the backward-leak uni-directional key updates setting. This observation affects a theorem proved by Jiang [Jia20, Theorem 3.2 in the ePrint ver.] (Theorem 4.9 in this paper), which we explain later.

Relaxed firewall. As we observed above, it is natural to consider uni-directional ciphertext updates in the backward uni-directional key updates setting. In this setting, adversaries cannot convert a ciphertext at the challenge epoch into a ciphertext at an older epoch by using tokens. Thus, even if a token Δ_{fwl} at a left firewall fwl is given to adversaries when a challenge epoch is in between fwl and fwr , adversaries cannot obtain a challenge-equal ciphertext at an epoch whose secret key is corrupted. We define this modified firewall notion as relaxed firewall below.

Definition 4.2 (Relaxed Firewall). *A relaxed insulated region with relaxed firewalls fwl and fwr is a consecutive sequence of epochs $[fwl, fwr]$ for which:*

- *No key in the sequence of epochs $[fwl, fwr]$ is corrupted. That is, it holds $[fwl, fwr] \cap \mathcal{K} = \emptyset$.*
- *The token Δ_{fwr+1} is not corrupted if they exist. That is, it holds $fwr + 1 \notin \mathcal{T}$.*
- *All tokens $(\Delta_{fwl}, \dots, \Delta_{fwr})$ can be corrupted. That is, $[fwl, fwr] \subseteq \mathcal{T}$.*

The difference from Definition 3.8 is that Δ_{fwl} can be corrupted.

Definition 4.3 (Relaxed Insulated Region). *The union of all relaxed insulated regions is defined as $r\mathcal{IR} := \bigcup_{[fwl, fwr] \in r\mathcal{FW}} [fwl, fwr]$, where $r\mathcal{FW}$ is the set of relaxed insulated region with relaxed firewalls.*

As we will see in the proof of Theorem 4.8, there exists an epoch such that it is set as the challenge ciphertext epoch (does not trigger the trivial winning condition), but not in a firewall area under Definition 3.8 (the original definition of firewall). In the example in Figure 3, which will appear later, epoch $\{5\}$ is such an area. Therefore, we introduce the modified notion.⁴

⁴In the previous version of this paper, we introduced extrusion area of firewall. However, we found that the relaxed firewall notion is a more appropriate notion. Thus, we introduce the relaxed firewall.

More observation. We present a minor observation in Appendix A.

Summary of observations. We summarize possible combinations for token generation and directions of key and ciphertext updates in Table 1. Note that we do not consider using obfuscation in this table. In each field, possible types are written. In the key update column, “forward-leak? or bi?” means that it can be forward-leak, but in this case, it might not be able to update a ciphertext in the forward direction. If it can update, it essentially includes sk_e and should be bi-directional. In the ciphertext update column, “backward-leak? or bi?” means that it can be backward, but it does not fit the nature of UE, and if it can be forward, it essentially has the power of bi-directional updates. That is, the second-row case could collapse to the first-row case in Table 1 if the second case works as UE (ciphertext updates are in the forward direction). Lastly, “?” means that we do not know whether this type can update a ciphertext or not (or it is unlikely that the type can update a ciphertext).

All previous ciphertext-independent updates UE schemes fall into the first row category. Our scheme in Section 5 falls into the third row category. There might be a hope that we can achieve a no-directional UE scheme by using obfuscation-like techniques. It is an interesting open question.

Table 1: Possible combinations for token generation from pk or sk and its relationship to possible directions of key updates and ciphertext updates.

use pk or sk	key update type	ct update type
$\text{TokGen}(sk_e, sk_{e+1})$	bi	bi
$\text{TokGen}(pk_e, sk_{e+1})$	forward-leak? or bi?	backward? or bi?
$\text{TokGen}(sk_e, pk_{e+1})$	backward-leak	forward
$\text{TokGen}(pk_e, pk_{e+1})$	no	?

4.3 Relationships

We show that bi-directional key updates does not imply backward-leak uni-directional key updates in this section. More precisely, we prove the following

Theorem 4.4. *There exist secure r -IND-UE-CPA UE schemes in the bi-directional key updates setting that are not r -IND-UE-CPA in the backward-leak uni-directional key updates setting.*

On the equivalence between bi-directional and uni-directional key updates. First, we review a simple fact. It is easy to see that the following theorem holds by the definition of confidentiality (Definition 3.3).

Theorem 4.5. *If a UE scheme is r -IND-UE-CPA in the backward-leak uni-directional, forward-leak uni-directional, or no-directional key updates setting, it is also r -IND-UE-CPA secure in the bi-directional key updates setting.*

Next, we review Jiang’s equivalence theorem.

Theorem 4.6 ([Jia20, Theorem 2]). *Let UE be an UE scheme and notion $\in \{d\text{-ind-ue-cpa}, r\text{-ind-ue-cpa}, d\text{-ind-ue-cca}, r\text{-ind-ue-cca}, \text{int-ctxt}, \text{int-ptxt}\}$. For any $kk, kk' \in \{f\text{-uni}, \text{bi}\}$, $cc, cc' \in \{\text{uni}, \text{bi}\}$, and any (kk, cc) -notion adversary \mathcal{A} against UE, there exists a (kk', cc') -notions adversary \mathcal{B} against UE such that*

$$\text{Adv}_{\text{UE}, \mathcal{A}}^{(kk, cc)\text{-notion}}(1^\lambda) = \text{Adv}_{\text{UE}, \mathcal{B}}^{(kk', cc')\text{-notion}}(1^\lambda).$$

The key lemma for proving Jiang’s theorem (Theorem 4.6) for the confidentiality case is the following.

Lemma 4.7 ([Jia20, Lemma 6]). For any $\mathcal{K}, \mathcal{T}, \mathcal{C}$, we have $\mathcal{K}_{f\text{-uni}}^* \cap \mathcal{C}_{f\text{-uni,uni}}^* \neq \emptyset \Leftrightarrow \mathcal{K}_{bi}^* \cap \mathcal{C}_{bi,bi}^* \neq \emptyset$.

See Definitions 3.6 and 4.1 for the sets in the lemma. Note that this lemma holds for *forward-leak* uni-directional key updates. We show a counterexample to this lemma (for confidentiality) in the case of the *backward-leak* uni-directional key updates setting.

Counterexample in backward-leak uni-directional key updates setting. Looking at an example is the best thing to understand relationships. We consider an example of epoch key leakage sets in Figure 3.

	0	{1}	2	3	4	5	{6	7}	8
\mathcal{K}	✓	×	×	×	✓	×	×	×	✓
\mathcal{T}	×	×	×	✓	✓	✓	×	✓	×
\mathcal{K}_{bi}^*	✓	×	✓	✓	✓	✓	×	×	✓
\mathcal{T}_{bi}^*	×	×	×	✓	✓	✓	×	✓	×
$\mathcal{K}_{f\text{-uni}}^*$	✓	×	×	✓	✓	✓	×	×	✓
$\mathcal{T}_{f\text{-uni}}^*$	×	×	×	✓	✓	✓	×	✓	×
$\mathcal{K}_{b\text{-uni}}^*$	✓	×	✓	✓	✓	×	×	×	✓
$\mathcal{T}_{b\text{-uni}}^*$	×	×	×	✓	✓	✓	×	✓	×

Figure 3: Example of leakage sets in the setting of bi/forward/backward-leak uni-directional key updates where $\mathcal{K} := \{0, 4, 8\}$, $\mathcal{T} := \{3, 4, 5, 7\}$, $\mathcal{IR} = \{1, 6, 7\}$. Here, \times and \checkmark indicates an epoch key or token is not corrupted and corrupted, respectively. The boldface check mark \checkmark indicates an epoch key or token is inferred from other corrupted keys/tokens.

In the example in Figure 3, the firewall area is $\mathcal{IR} = \{1, 6, 7\}$. The difference between the bi-directional setting and forward-leak uni-directional setting is the epochs 2 and 3. The difference between the bi-directional setting and backward-leak uni-directional setting is the epoch 5. (Both differences are underlined in Figure 3.) We investigate each difference in the forward/backward-leak uni-directional settings.

The case of bi/forward-leak uni-directional key updates: First, we consider the bi/forward-leak uni-directional key updates settings. If we set $\mathcal{C} = \{3\}$, it holds $\mathcal{C}_{bi,bi}^* = \{2, 3, 4, 5\}$ and $\mathcal{C}_{f\text{-uni,uni}}^* = \{3, 4, 5\}$. Thus, $\mathcal{K}_{bi}^* \cap \mathcal{C}_{bi,bi}^* = \{2, 3, 4, 5\} \neq \emptyset$ and $\mathcal{K}_{f\text{-uni}}^* \cap \mathcal{C}_{f\text{-uni,uni}}^* = \{4, 5\} \neq \emptyset$. If we set $\mathcal{C} = \{5\}$, it holds that $\mathcal{K}_{bi}^* \cap \mathcal{C}_{bi,bi}^* = \{2, 3, 4, 5\} \neq \emptyset$ and $\mathcal{K}_{f\text{-uni}}^* \cap \mathcal{C}_{f\text{-uni,uni}}^* = \{5\} \neq \emptyset$. This is consistent with Lemma 4.7 (Jiang’s Lemma 6 [Jia20]). Note that if we set $\mathcal{C} = \{2\}$, we obtain a similar result to $\mathcal{C} = \{3\}$.

The case of bi/backward-leak uni-directional key updates: Next, we consider the bi/backward-leak uni-directional key updates settings. If we set $\mathcal{C} = \{3\}$, it holds $\mathcal{C}_{bi,bi}^* = \{2, 3, 4, 5\}$ and $\mathcal{C}_{b\text{-uni,uni}}^* = \{3, 4, 5\}$ since Δ_5 is given even though k_5 is not given in the backward-leak uni-directional setting. Thus, it holds $\mathcal{K}_{bi}^* \cap \mathcal{C}_{bi,bi}^* = \{2, 3, 4, 5\} \neq \emptyset$ and $\mathcal{K}_{b\text{-uni}}^* \cap \mathcal{C}_{b\text{-uni,uni}}^* = \{3, 4\} \neq \emptyset$. However, if we set $\mathcal{C} = \{5\}$, the difference between forward/backward directional key updates is clear. Now, $\mathcal{K}_{bi}^* \cap \mathcal{C}_{bi,bi}^* = \{2, 3, 4, 5\} \neq \emptyset$, but $\mathcal{K}_{b\text{-uni}}^* \cap \mathcal{C}_{b\text{-uni,uni}}^* = \emptyset$ since we cannot infer k_5 (the key at epoch 5) due to the definition of backward-leak uni-directional key updates (we cannot go to forward direction even if we are given k_4 and Δ_5). This means that even if we set $\mathcal{C} = \{5\}$, the trivial winning condition is not triggered in the backward-leak uni-directional setting. However, the trivial winning condition in the bi-directional setting is triggered. Therefore, this is a counterexample to Lemma 4.7 (Jiang’s Lemma 6 [Jia20]) when we use the definition of *backward-leak* uni-directional key updates.

By using the example above, we immediately obtain the following theorem.

Theorem 4.8. *The ciphertext-independent UE schemes Lehman and Tackmann [LT18], Boyd et al. [BDGJ20], and Jiang [Jia20] do not satisfy confidentiality in the backward-leak uni-directional setting.*

Proof. We use the leakage sets example \mathcal{K} and \mathcal{T} in Figure 3 and set $\mathcal{C} = \{5\}$. This does not trigger the trivial winning condition in the backward-leak uni-directional setting. However, an adversary can infer k_5 by using k_4 and Δ_5 in the *bi-directional key updates* schemes described in the theorem statement. Thus, the adversary trivially wins the confidentiality game in the backward-leak uni-directional setting since a challenge ciphertext is encrypted under k_5 . ■

By Theorem 4.8 and the results by Lehman and Tackmann [LT18], Boyd et al. [BDGJ20], and Jiang [Jia20], we immediately obtain Theorem 4.4 since they show that their schemes satisfy confidentiality in the bi-directional key updates setting. Therefore, surprisingly (or unsurprisingly), UE with backward-leak uni-directional (and no-directional) key updates is *strictly stronger* than UE with bi-directional key updates by Theorems 4.4 and 4.5.

On equivalence between no/uni/bi-directional key updates in bi-directional ciphertext update setting.

We give an observation on the equivalence theorem about no-directional key updates. Jiang also proves the following theorem.

Theorem 4.9 ([Jia20, Theorem 3.2 in the ePrint ver.]). *Let UE be an UE scheme and notion $\in \{\text{d-ind-ue-cpa}, \text{r-ind-ue-cpa}, \text{d-ind-ue-cca}, \text{r-ind-ue-cca}\}$. For any (no, bi)-notion adversary \mathcal{A} against UE, there exists a (f-uni, bi)-notions adversary \mathcal{B} against UE such that*

$$\text{Adv}_{\text{UE}, \mathcal{A}}^{(\text{no, bi})\text{-notion}}(1^\lambda) = \text{Adv}_{\text{UE}, \mathcal{B}}^{(\text{f-uni, bi})\text{-notion}}(1^\lambda).$$

This theorem seems to contradict our conclusion above, which says UE with no-directional key updates is strictly stronger than UE with forward-leak uni-directional key updates. Recall that no-directional key updates is stronger than backward-leak uni-directional key updates. We also note that bi-directional key updates and forward-leak uni-directional key updates are equivalent.

The source of the puzzle above comes from the fact that the theorem holds for *bi-directional ciphertext updates*. The key lemma for proving Jiang’s theorem above (Theorem 4.9) is the following.

Lemma 4.10 ([Jia20, Lemma 3.15 in the ePrint ver.]). *For any $\mathcal{K}, \mathcal{T}, \mathcal{C}$, we have $\mathcal{K}_{\text{f-uni}}^* \cap \mathcal{C}_{\text{f-uni, bi}}^* \neq \emptyset \Rightarrow \mathcal{K}_{\text{no}}^* \cap \mathcal{C}_{\text{no, bi}}^* \neq \emptyset$.*

The proof of the lemma above heavily relies on the bi-directional ciphertext update setting. As we argued in Section 4.2, it is unnatural to consider bi-directional ciphertext updates with backward-leak uni-directional (and no-directional) key updates. Thus, if we exclude such an unnatural or artificial setting, the equivalence theorem above (Theorem 4.9), which is counterintuitive, does not hold in the case of the backward-leak uni-directional key updates setting.

5 Construction with Backward-Leak Uni-Directional Key Update

In this section, we present a backward-leak uni-directional key update scheme from the LWE assumption.

5.1 Scheme Description and Design Idea

We present a UE scheme with backward-leak uni-directional key updates based on the Regev PKE scheme [Reg09], and denoted by RtR. A proxy re-encryption scheme by Nishimaki and Xagawa [NX15] inspired this construction idea.

The ciphertext update technique is based on the key-switching technique [BV14, BV11, BGV14]. In particular, we use that for multi-bit plaintexts [BGH13]. In the following, we denote a plaintext by $\mu \in \{0, 1\}^\ell$ and error distributions by χ and χ_{ns} .

A variant of Regev PKE scheme. We review a variant of Regev PKE scheme [Reg09] in the multi-user settings.

- Setup(1^λ): Choose $A \leftarrow \mathbb{Z}_q^{m \times n}$ and output $\text{pp} := (A, 1^\lambda, 1^n, 1^m, 1^\ell, q, \chi, \chi_{\text{ns}})$.
- Reg.Gen(pp): Choose $S \leftarrow \mathbb{Z}_q^{n \times \ell}$ and $X \leftarrow \chi^{m \times \ell}$, compute $B := AS + X \in \mathbb{Z}_q^{m \times \ell}$, and outputs $\text{pk} = B$ and $\text{sk} = S$.
- Reg.Enc(pk, μ): Choose $r \leftarrow \{-1, +1\}^m$ and $e' \leftarrow \chi_{\text{ns}}^\ell$ and output $(u, c) := (rA, rB + e' + \lfloor q/2 \rfloor \mu)$.
- Reg.Dec(sk, (u, c)) Compute $d := c - uS$ and output $\mu := \lfloor (2/q)d \rfloor \bmod 2$.

Key-switching technique. We review the key-switching technique in the multi-bit version for our update algorithm. Let $\eta := \lceil \lg q \rceil$. We give the definitions of the binary-decomposition algorithm $\text{BD}(\cdot)$ and the powers-of-2 algorithm $\text{P2}(\cdot)$.

- $\text{BD}(x \in \mathbb{Z}_q^n)$: It decomposes $x = \sum_{k=1}^{\eta} 2^{k-1} u_k$, where $u_k \in \{0, 1\}^n$, and outputs $(u_1, u_2, \dots, u_\eta) \in \{0, 1\}^{n\eta}$.
- $\text{P2}(s \in \mathbb{Z}_q^{n \times 1})$: It outputs $[1, 2, \dots, 2^{\eta-1}]^\top \otimes s = [s; 2s; \dots; 2^{\eta-1}s] \in \mathbb{Z}_q^{n\eta \times 1}$, where \otimes denotes the standard tensor product. We extend the domain of P2 by setting $\text{P2}([s_1 \dots s_\ell] \in \mathbb{Z}_q^{n \times \ell}) = [\text{P2}(s_1) \dots \text{P2}(s_\ell)] \in \mathbb{Z}_q^{n\eta \times \ell}$.

By the definition, it holds that $\text{BD}(x) \cdot \text{P2}(S) = x \cdot S \in \mathbb{Z}_q^\ell$ for any $x \in \mathbb{Z}_q^n$ and $S \in \mathbb{Z}_q^{n \times \ell}$.

Let $S_e, S_{e+1} \in \mathbb{Z}_q^{n \times \ell}$ be two secret keys at epoch $e, e+1$, respectively. The key-switching technique enables us to homomorphically decrypt a ciphertext at epoch e and obtain a ciphertext at epoch $e+1$ by using encryption of S_e under the key at epoch $e+1$. More formally, the key-switching matrix M_{e+1} is $[A' \mid A'S_{e+1} + Y] + [O \mid -\text{P2}(S_e)]$, where $A' \leftarrow \mathbb{Z}_q^{m\eta \times n}$, $Y \leftarrow \chi^{m\eta \times \ell}$. To update a ciphertext (u, c) under S_e to one under S_{e+1} , we compute $(u', c') = (u, c) + \text{BD}(u)M_{e+1}$. By simple calculation, we have that

$$\begin{aligned} (u', c') &= (u, c) + \text{BD}(u) ([A' \mid A'S_{e+1} + Y] + [O \mid -\text{P2}(S_e)]) \\ &= (\text{BD}(u)A', c - uS_e + \text{BD}(u)A'S_{e+1} + \text{BD}(u) \cdot Y). \end{aligned}$$

To decrypt ciphertext by secret key S_{e+1} , we compute

$$\begin{aligned} c' - u'S_{e+1} &= c - uS_e + \text{BD}(u)A'S_{e+1} + \text{BD}(u) \cdot Y - \text{BD}(u)A'S_{e+1} \\ &= c - uS_e + \text{BD}(u) \cdot Y. \end{aligned}$$

Thus, the decryption is correct if the magnitude of additional noises $\text{BD}(u) \cdot Y$ is small.

backward-leak uni-directional update. In fact, we do not need the secret key S_{e+1} at epoch $e + 1$ for update. We set $B_{e+1} = AS_{e+1} + Y_{e+1}$, which we call the public key part of the key at epoch $e + 1$. We choose $R_{e+1} \leftarrow \{-1, +1\}^{n\eta \times m}$ and compute an update token

$$\begin{aligned} M_{e+1} &= R_{e+1}[A \mid B_{e+1}] + [O \mid -P2(S_e)] \\ &= [A' \mid A'S_{e+1} + Y'] + [O \mid -P2(S_e)], \end{aligned}$$

where $A' = R_{e+1}A$ and $Y' = R_{e+1}Y_j$. By using M_{e+1} , we can update ciphertext (u, c) at epoch e . Thus, even if given the key S_e at epoch e and the token M_{e+1} , we cannot infer S_{e+1} since only the public key part B_{e+1} (this is pseudorandom by the LWE assumption) of the key at epoch $e + 1$ is embedded in M_{e+1} . Note that S_e and S_{e+1} are independently chosen. However, if given the key S_{e+1} at epoch $e + 1$ and the token M_{e+1} , we can easily infer S_e since S_e is encrypted under S_{e+1} . Thus, this update mechanism is a backward-leak uni-directional key update and uni-directional ciphertext update.

How to achieve randomized update. The update algorithm above is deterministic. To re-randomize an updated ciphertext, we set the update token as M_{e+1} and B_{e+1} , which is the public key part at epoch $e + 1$. First, we convert ciphertext (u, c) at epoch e into (u', c') using M_{e+1} as above and masking (u', c') with a new ciphertext $(\tilde{u}, \tilde{v}) := \tilde{r}[A \mid B_{e+1}]$ of the plaintext $\mathbf{0}$. This is not enough for confidentiality since it includes information about B_{e+1} and is not random. To overcome this issue, we randomize $[A \mid B_{e+1}]$ into $N_{e+1} = R'_{e+1} \cdot [A \mid B_{e+1}]$, where $R'_{e+1} \leftarrow \{-1, +1\}^{m \times m}$ and add it to Δ_{e+1} . Since the matrix N_{e+1} consists of m ciphertexts of the message $\mathbf{0}$, this is pseudorandom. The update token consists of key-switching matrix M_{e+1} and randomized matrix N_{e+1} .

Backward-leak uni-directional key update scheme. A UE scheme, RtR, is defined as follows:

Setup(1^λ):

1. Choose $A \leftarrow \mathbb{Z}_q^{m \times n}$.
2. Output pp := $(A, 1^\lambda, 1^n, 1^m, 1^\ell, q, \chi, \chi_{ns})$.

Gen(pp):

1. Generate $(B_e, S_e) \leftarrow \text{Reg.Gen}(1^\lambda)$.
2. Output $k_e := (\text{sk}_e, \text{pk}_e) := (S_e, B_e)$.

Enc($k_e, \mu \in \{0, 1\}^\ell$):

1. Parse $k_e = (S_e, B_e)$.
2. Generate $(u, c) \leftarrow \text{Reg.Enc}(B_e, \mu)$.
3. Output ct := $(u, c) \in \mathbb{Z}_q^n \times \mathbb{Z}_q^\ell$.

Dec(k_e, ct):

1. Parse $k_e = (S_e, B_e)$ ct = (u, c) .
2. Compute and output $\mu \leftarrow \text{Reg.Dec}(S_e, \text{ct})$.

TokGen(k_e, k_{e+1}):

1. Parse $k_e = (S_e, B_e)$ and $k_{e+1} = (S_{e+1}, B_{e+1})$.
2. Compute $M_{e+1} := R_{e+1} \cdot [A \mid B_{e+1}] + [O \mid -P2(S_e)]$, where $R_{e+1} \leftarrow \{-1, +1\}^{n\eta \times m}$.
3. Compute $N_{e+1} := R'_{e+1} \cdot [A \mid B_{e+1}]$, where $R'_{e+1} \leftarrow \{-1, +1\}^{m \times m}$.

4. Output $\Delta_{e+1} := (M_{e+1}, N_{e+1})$.

Upd($\Delta_{e+1}, \text{ct}_e$):

1. Parse $\Delta_{e+1} = (M_{e+1}, N_{e+1})$ and $\text{ct}_e = (\mathbf{u}_e, \mathbf{c}_e)$.
2. Compute $(\mathbf{u}', \mathbf{c}') := \text{BD}(\mathbf{u}_e)M_{e+1}$;
3. Compute $(\tilde{\mathbf{u}}, \tilde{\mathbf{v}}) := \tilde{\mathbf{r}} \cdot N_{e+1}$, where $\tilde{\mathbf{r}} \leftarrow \{-1, +1\}^m$;
4. Output $\text{ct}_{e+1} := (\bar{\mathbf{u}}, \bar{\mathbf{c}}) := (\mathbf{u}' + \tilde{\mathbf{u}}, \mathbf{c}' + \tilde{\mathbf{v}})$.

For notational convenience, we call $\text{pk}_e = \mathbf{B}_e$ and $\text{sk}_e = \mathbf{S}_e$ public key and secret key of epoch e , respectively. Note that we can run Enc without $\text{sk}_e = \mathbf{S}_e$ (we need only $\text{pk}_e = \mathbf{B}_e$). We also note that we can run TokGen(k_e, k_{e+1}) without sk_{e+1} (we need only pk_{e+1} and sk_e).

The scheme is correct and r-IND-UE-CPA secure. We prove the following theorems in Sections 5.2 and 5.3. Let T be the maximum number of the epoch.

Theorem 5.1. *Let χ and χ_{ns} be B -bounded and B' -bounded distributions, respectively, such that $B/B' = \text{negl}(\lambda)$ and $m = 2n \lg q + \omega(\sqrt{\lg \lambda})$. Suppose that $(1 + n\eta + m)mB + B' \leq q/4T$. Then RtR is correct.*

Theorem 5.2. *Suppose that $m \geq (n + \ell) \lg q + \omega(\lg \lambda)$. Under the $\text{LWE}(n, q, \chi)$ assumption, RtR is r -IND-UE-CPA secure in the backward-leak uni-directional setting. That is, $\text{Adv}_{\text{RtR}, \mathcal{A}}^{(\text{b-uni, uni})\text{-r-ind-ue-cpa}}(1^\lambda) \leq \text{negl}(\lambda)$.*

5.2 Correctness

We give rough estimations on B -bounded and B' -bounded distributions χ and χ_{ns} , respectively, for simplicity. However, if we set $\chi = \bar{\Psi}_\alpha$ or $D_{\mathbb{Z}, s}$, we can obtain tighter bounds.

Proof of Theorem 5.1. The theorem follows from Propositions 5.3 and 5.4 below. ■

Proposition 5.3. *The scheme is correct for the encryption algorithm if $mB + B' < q/4$.*

Proposition 5.4. *The scheme is correct for the update algorithm if $(1 + n\eta + m)mB + B' < q/4T$.*

Those correctness easily follows from the proof by Regev [Reg09]. For completeness, we include the proofs.

Proof of Proposition 5.3. Let $(\mathbf{S}, \mathbf{B}) = (\mathbf{S}, \mathbf{AS} + \mathbf{X})$, where $\mathbf{S} \leftarrow \mathbb{Z}_q^{n \times \ell}$ and $\mathbf{X} \leftarrow \chi^{m \times \ell}$ be a key k_e of epoch e . To decrypt a ciphertext $(\mathbf{u}, \mathbf{c}) = (\mathbf{rA}, \mathbf{rB} + \lfloor q/2 \rfloor \boldsymbol{\mu})$ of the message $\boldsymbol{\mu} \in \{0, 1\}^\ell$ with \mathbf{S} under k_e , we compute

$$\begin{aligned} \mathbf{d} &= \mathbf{c} - \mathbf{uS} = \lfloor q/2 \rfloor \boldsymbol{\mu} + \mathbf{e}' + \mathbf{rB} - \mathbf{rAS} \\ &= \lfloor q/2 \rfloor \boldsymbol{\mu} + \mathbf{e}' + \mathbf{rX}. \end{aligned}$$

If $\|\mathbf{e}'\|_\infty + \|\mathbf{rX}\|_\infty < q/4$, we can obtain $\boldsymbol{\mu} = \lfloor 2\mathbf{d}/q \rfloor \bmod 2$. Since $\mathbf{r} \in \{-1, +1\}^m$ and \mathbf{X} consists of samples from a B -bounded distribution, $\|\mathbf{rX}\|_\infty$ is at most mB . In addition, \mathbf{e}' are from a B' -bounded distribution, $\|\mathbf{e}'\|_\infty$ is at most B' . From the parameter setting of Proposition 5.3, we have $\|\mathbf{e}' + \mathbf{rX}\|_\infty \leq \|\mathbf{e}'\|_\infty + \|\mathbf{rX}\|_\infty \leq mB + B' < q/4T$. This completes the proof. ■

Proof of Proposition 5.4. Let $\mathbf{B}_e = \mathbf{A}\mathbf{S}_e + \mathbf{X}_e$ and $\mathbf{B}_{e+1} = \mathbf{A}\mathbf{S}_{e+1} + \mathbf{X}_{e+1}$ be public keys of epoch e and $e + 1$, respectively, where $\mathbf{S}_e, \mathbf{S}_{e+1} \leftarrow \mathbb{Z}_q^{n \times \ell}$ and $\mathbf{X}_e, \mathbf{X}_{e+1} \leftarrow \chi^{m \times \ell}$. The update token from e to $e + 1$ is generated as

$$\mathbf{M}_{e+1} = \mathbf{R}_{e+1}[\mathbf{A} \mid \mathbf{B}_{e+1}] + [\mathbf{O} \mid -\text{P2}(\mathbf{S}_e)] \text{ and } \mathbf{N}_{e+1} = \mathbf{R}'_{e+1} \cdot [\mathbf{A} \mid \mathbf{B}_{e+1}],$$

where $\mathbf{R}_{e+1} \leftarrow \{-1, +1\}^{n\eta \times m}$ and $\mathbf{R}'_{e+1} \leftarrow \{-1, +1\}^{m \times m}$. We consider a ciphertext $(\mathbf{u}, \mathbf{c}) = \mathbf{r}[\mathbf{A} \mid \mathbf{B}_e] + \mathbf{e}'_e + (\mathbf{0}, \lfloor q/2 \rfloor \boldsymbol{\mu})$ of the message $\boldsymbol{\mu} \in \{0, 1\}^\ell$, and an updated ciphertext $(\tilde{\mathbf{u}}, \tilde{\mathbf{c}}) = (\mathbf{u}' + \tilde{\mathbf{u}}, \mathbf{c} + \mathbf{c}' + \tilde{\mathbf{c}}) = (\mathbf{0}, \mathbf{c}) + \text{BD}(\mathbf{u}) \cdot \mathbf{M}_{e+1} + \tilde{\mathbf{r}} \cdot \mathbf{N}_{e+1}$, where $\tilde{\mathbf{r}} \leftarrow \{-1, +1\}^m$.

The decryption algorithm $\text{Dec}(\mathbf{S}_{e+1}, (\tilde{\mathbf{u}}, \tilde{\mathbf{c}}))$ computes

$$\begin{aligned} \mathbf{d} &= \tilde{\mathbf{c}} - \tilde{\mathbf{u}}\mathbf{S}_{e+1} \\ &= \mathbf{c} + \text{BD}(\mathbf{u})\mathbf{M}_{e+1}[-\mathbf{S}_{e+1}; \mathbf{I}_\ell] + \tilde{\mathbf{r}}\mathbf{N}_{e+1}[-\mathbf{S}_{e+1}; \mathbf{I}_\ell] \\ &= \mathbf{c} + \text{BD}(\mathbf{u}) (\mathbf{R}_{e+1}(\mathbf{B}_{e+1} - \mathbf{A}\mathbf{S}_{e+1}) - \text{P2}(\mathbf{S}_e)) + \tilde{\mathbf{r}}\mathbf{R}'_{e+1}(\mathbf{B}_{e+1} - \mathbf{A}\mathbf{S}_{e+1}) \\ &= \mathbf{c} - \mathbf{u}\mathbf{S}_e + (\text{BD}(\mathbf{u})\mathbf{R}_{e+1} + \tilde{\mathbf{r}}\mathbf{R}'_{e+1})\mathbf{X}_{e+1} \\ &= \lfloor q/2 \rfloor \boldsymbol{\mu} + \underbrace{\mathbf{e}'_e + \mathbf{r}\mathbf{X}_e + (\text{BD}(\mathbf{u})\mathbf{R}_{e+1} + \tilde{\mathbf{r}}\mathbf{R}'_{e+1})\mathbf{X}_{e+1}}_{=: \delta}, \end{aligned}$$

where we used the equation $\mathbf{B}_{e+1} = \mathbf{A}\mathbf{S}_{e+1} + \mathbf{X}_{e+1}$. As in the proof of Proposition 5.3, if $\|\delta\|_\infty < q/4$, we can correctly decrypt the update ciphertext into $\boldsymbol{\mu}$ by \mathbf{S}_{e+1} . We have that following bounds by routine calculation:

$$\|\mathbf{e}'_e\|_\infty \leq B', \|\mathbf{r}\mathbf{X}_e\|_\infty \leq mB, \|\text{BD}(\mathbf{u})\mathbf{R}_{e+1}\mathbf{X}_{e+1}\|_\infty \leq n\eta \cdot mB, \text{ and } \|\tilde{\mathbf{r}}\mathbf{R}'_{e+1}\mathbf{X}_{e+1}\|_\infty \leq m^2B.$$

From the triangle inequality and the parameter setting, we have that

$$\begin{aligned} \|\delta\|_\infty &\leq \|\mathbf{e}'_e\|_+ + \|\mathbf{r}\mathbf{X}_e\|_\infty + \|\text{BD}(\mathbf{u})\mathbf{R}_{e+1}\mathbf{X}_{e+1}\|_\infty + \|\tilde{\mathbf{r}}\mathbf{R}'_{e+1}\mathbf{X}_{e+1}\|_\infty \\ &\leq B' + mB + n\eta mB + m^2B = (1 + n\eta + m)mB + B' \\ &< q/4T \leq q/4. \end{aligned}$$

Thus, when we update a ciphertext once, a noise is added and its magnitude is at most $q/4T$. This means that we can update a ciphertext T times. We complete the proof. ■

5.3 Confidentiality

We show RtR is r-IND-UE-CPA in the backward-leak uni-directional setting. Although it is trivial that RtR satisfies uni-directional ciphertext updates from its security, we confirm it below.

Lemma 5.5. *If (Setup, Reg.Gen, Reg.Enc, Reg.Dec) is IND-CPA secure PKE, adversaries cannot convert a ciphertext under a public key pk_{e+1} into one under a public key pk_e even if they are given Δ_{e+1} .*

Proof. We construct an algorithm \mathcal{B} that breaks IND-CPA security under pk_{e+1} by using an adversary \mathcal{D} that converts a ciphertext under pk_{e+1} into one under pk_e by using $(\text{pk}_e, \text{sk}_e)$, pk_{e+1} , and Δ_{e+1} .

First, \mathcal{B} is given pk_{e+1} . \mathcal{B} generates $(\text{pk}_e, \text{sk}_e)$ and $\Delta_{e+1} \leftarrow \text{TokGen}(\text{sk}_e, \text{pk}_{e+1})$, selects any (m_0, m_1) , sends (m_0, m_1) to its challenger, and receives a target ciphertext $\text{ct}^* \leftarrow \text{Reg.Enc}(\text{pk}_{e+1}, m_b)$ where $b \leftarrow \{0, 1\}$. Next, \mathcal{B} sends $((\text{pk}_e, \text{sk}_e), \Delta_{e+1}, \text{ct}^*)$ to \mathcal{D} . \mathcal{D} outputs a ciphertext ct' under pk_e . Then, \mathcal{B} computes $m' \leftarrow \text{Reg.Dec}(\text{sk}_e, \text{ct}')$ by using sk_e and if $m' = m_{b'}$, it outputs b' .

It is easy to see that if \mathcal{D} can convert ct^* into a ciphertext under pk_e , \mathcal{B} outputs $b' = b$. This completes the proof. ■

Second, we look at the detail of the update procedure again. By simple calculation, we obtain

$$\begin{aligned}
(\bar{\mathbf{u}}, \bar{\mathbf{c}}) &= (\mathbf{u}' + \tilde{\mathbf{u}}, \mathbf{c}_e + \mathbf{c}' + \tilde{\mathbf{c}}) \\
&= (\mathbf{0}, \mathbf{c}_e) + \text{BD}(\mathbf{u}_e) \cdot \mathbf{M}_{e+1} + \tilde{\mathbf{r}} \cdot \mathbf{N}_{e+1} \\
&= ((\text{BD}(\mathbf{u}_e)\mathbf{R}_{e+1} + \tilde{\mathbf{r}}\mathbf{R}'_{e+1})\mathbf{A}, \\
&\quad (\text{BD}(\mathbf{u}_e)\mathbf{R}_{e+1} + \tilde{\mathbf{r}}\mathbf{R}'_{e+1})\mathbf{A}\mathbf{S}_{e+1} + (\text{BD}(\mathbf{u}_e)\mathbf{R}_{e+1} + \tilde{\mathbf{r}}\mathbf{R}'_{e+1})\mathbf{X}_{e+1} + \mathbf{e}'_e + \mathbf{r}\mathbf{X}_e + \lfloor q/2 \rfloor \boldsymbol{\mu}) \\
&= (\mathbf{r}^\dagger \mathbf{A}, \mathbf{r}^\dagger \mathbf{B}_{e+1} + \mathbf{e}'_e + \mathbf{r}\mathbf{X}_e + \lfloor q/2 \rfloor \boldsymbol{\mu}) \text{ where } \mathbf{r}^\dagger := \text{BD}(\mathbf{u}_e)\mathbf{R}_{e+1} + \tilde{\mathbf{r}}\mathbf{R}'_{e+1} \\
&\stackrel{\S}{\approx} (\mathbf{r}^\dagger \mathbf{A}, \mathbf{r}^\dagger \mathbf{B}_{e+1} + \mathbf{e}'_e + \lfloor q/2 \rfloor \boldsymbol{\mu}). \tag{1}
\end{aligned}$$

The last equation (statistical indistinguishability) holds by Lemma 2.2. This equation shows that we can simulate an update ciphertext by using the original ciphertext, its plaintext and randomness, the new epoch public key, and *randomness* for generating the token Δ_{e+1} (not the token itself).

To show the security, we define auxiliary algorithms for simulation.

Hyb.Upd($\text{ct}_e, \mathbf{B}_{e+1}, \boldsymbol{\mu}; \mathbf{e}'_e, (\mathbf{R}_{e+1}, \mathbf{R}'_{e+1})$):

- Parse $\text{ct}_e = (\mathbf{u}_e, \mathbf{c}_e)$.
- Choose $\tilde{\mathbf{r}} \leftarrow \{-1, +1\}^m$.
- Set $\mathbf{r}^\dagger := \text{BD}(\mathbf{u}_e)\mathbf{R}_{e+1} + \tilde{\mathbf{r}}\mathbf{R}'_{e+1}$.
- Set $\text{ct}_{e+1} := (\bar{\mathbf{u}}, \bar{\mathbf{c}}) := (\mathbf{r}^\dagger \mathbf{A}, \mathbf{r}^\dagger \mathbf{B}_{e+1} + \mathbf{e}'_e + \lfloor q/2 \rfloor \boldsymbol{\mu})$.
- Output $(\text{ct}_{e+1}; \mathbf{e}'_e)$.

Sim.Gen(pp):

- Choose $\mathbf{B}_e^+ \leftarrow \mathbb{Z}_q^{m \times \ell}$.
- Output $\text{pk}_e := \mathbf{B}_e^+$.

Sim.TokGen(pp):

- Choose $\mathbf{M}_{e+1}^+ \leftarrow \mathbb{Z}_q^{n\eta \times (n+\ell)}$ and $\mathbf{N}_{e+1}^+ \leftarrow \mathbb{Z}_q^{m \times (n+\ell)}$.
- Output $\Delta_{e+1}^+ := (\mathbf{M}_{e+1}^+, \mathbf{N}_{e+1}^+)$.

Sim.Upd(pp):

- Choose $(\bar{\mathbf{u}}, \bar{\mathbf{c}}) \leftarrow \mathbb{Z}_q^n \times \mathbb{Z}_q^\ell$.
- Output $\text{ct}_{e+1} := (\bar{\mathbf{u}}, \bar{\mathbf{c}})$.

Sim.Enc(pp):

- Choose $(\bar{\mathbf{u}}, \bar{\mathbf{c}}) \leftarrow \mathbb{Z}_q^n \times \mathbb{Z}_q^\ell$.
- Output $\text{ct}_e := (\bar{\mathbf{u}}, \bar{\mathbf{c}})$.

Lemma 5.6. $\text{Upd}(\Delta_{e+1}, \text{ct}_e) \stackrel{\S}{\approx} \text{Hyb.Upd}(\text{ct}_e, \mathbf{B}_{e+1}, \boldsymbol{\mu}; \mathbf{e}'_e, (\mathbf{R}_{e+1}, \mathbf{R}'_{e+1}))$

By Equation (1), Lemma 5.6 immediately holds. That is, we can simulate $\mathcal{O}.\text{Upd}(\text{ct}_e)$ by using $\text{Hyb.Upd}(\text{ct}_e, \mathbf{B}_{e+1}, \boldsymbol{\mu}; \mathbf{e}'_e, (\mathbf{R}_{e+1}, \mathbf{R}'_{e+1}))$.

We follow the firewall technique [LT18, KLR19, BDGJ20, Jia20] to prove security, but we use the relaxed firewall notion in Definition 4.2.

Proof of Theorem 5.2. Let T be the upper bound of the number of epoch. We consider a sequence of hybrid games. First, we define the following hybrid game:

$\text{Hyb}_i(b)$: This is the same as $\text{Exp}_{\text{RtR}, \mathcal{A}}^{(\text{b-uni, uni})\text{-r-ind-ue-cpa}}(\lambda, b)$ except the following difference: When the adversary sends a query $(\bar{\mu}, \bar{\text{ct}})$ to $\mathcal{O}.\text{Chall}$ or an empty query to $\mathcal{O}.\text{Upd}\tilde{\text{C}}$ at epoch j ,

- for $j < i$, return an honestly generated challenge-equal ciphertext. That is, if $b = 0$, $\text{UE.Enc}(k_{\bar{e}}, \bar{\mu})$ else $\text{UE.Upd}(\Delta_{\bar{e}}, \bar{\text{ct}})$.
- for $j \geq i$, return a random ciphertext.

It is easy to see that $\text{Hyb}_{T+1}(b)$ is the same as the original r-INE-UE-CPA game in the backward-leak uni-directional setting $\text{Exp}_{\text{RtR}, \mathcal{A}}^{(\text{b-uni, uni})\text{-r-ind-ue-cpa}}(\lambda, b)$. Let $U(\lambda)$ be a random variable distributed uniformly in $[0, T]$, by the standard hybrid argument, we have

$$\begin{aligned}
& \text{Adv}_{\text{RtR}, \mathcal{A}}^{\text{r-ind-ue-cpa}}(\lambda) \\
&= |\Pr[\text{Hyb}_{T+1}(1) = 1] - \Pr[\text{Hyb}_{T+1}(0) = 1]| \\
&= \left| \sum_{i=0}^T \{(\Pr[\text{Hyb}_{i+1}(1) = 1] - \Pr[\text{Hyb}_i(1) = 1]) - (\Pr[\text{Hyb}_{i+1}(0) = 1] - \Pr[\text{Hyb}_i(0) = 1])\} \right| \quad (2) \\
&= \left| \sum_{i=0}^T (\Pr[\text{Hyb}_{U(\lambda)+1}(1) = 1 \mid U(\lambda) = i] - \Pr[\text{Hyb}_i(1) = 1 \mid U(\lambda) = i]) \right. \\
&\quad \left. - \sum_{i=0}^T (\Pr[\text{Hyb}_{i+1}(0) = 1 \mid U(\lambda) = i] - \Pr[\text{Hyb}_i(0) = 1 \mid U(\lambda) = i]) \right| \\
&= (T+1) \left| \sum_{i=0}^T (\Pr[\text{Hyb}_{U(\lambda)+1}(1) = 1 \wedge U(\lambda) = i] - \Pr[\text{Hyb}_{U(\lambda)}(1) = 1 \wedge U(\lambda) = i]) \right. \\
&\quad \left. - \sum_{i=0}^T (\Pr[\text{Hyb}_{U(\lambda)+1}(0) = 1 \wedge U(\lambda) = i] - \Pr[\text{Hyb}_{U(\lambda)}(0) = 1 \wedge U(\lambda) = i]) \right| \quad (3) \\
&\leq (T+1) |\Pr[\text{Hyb}_{U(\lambda)+1}(1) = 1] - \Pr[\text{Hyb}_{U(\lambda)}(1) = 1]| \\
&\quad + (T+1) |\Pr[\text{Hyb}_{U(\lambda)+1}(0) = 1] - \Pr[\text{Hyb}_{U(\lambda)}(0) = 1]|,
\end{aligned}$$

where we use $\Pr[U(\lambda) = i] = 1/(T+1)$ for Equation (3). Note that $\text{Hyb}_0(0) = \text{Hyb}_0(1)$ trivially holds since all challenge-equal ciphertexts are random ciphertexts. We use this fact in Equation (2). Thus, our goal is to prove $|\Pr[\text{Hyb}_{U(\lambda)+1}(b) = 1] - \Pr[\text{Hyb}_{U(\lambda)}(b) = 1]| \leq \text{negl}(\lambda)$ for $b \in \{0, 1\}$.

Hereafter, we write $\text{Hyb}_i(b)$ instead of $\text{Hyb}_{U(\lambda)}(b)$ for simplicity. Next, we define the following hybrid game:

$\text{Hyb}'_i(b)$: This is the same as $\text{Hyb}_i(b)$ except that the game chooses $\text{fwl}, \text{fwr} \leftarrow [0, T]$. If the adversary corrupts k_j such that $j \in [\text{fwl}, \text{fwr}]$ or $\Delta_{\text{fwr}+1}$, the game aborts.

The guess is correct with probability $1/(T+1)^2$. We have

$$|\Pr[\text{Hyb}_i(b) = 1] - \Pr[\text{Hyb}_{i-1}(b) = 1]| \leq (T+1)^2 |\Pr[\text{Hyb}'_i(b) = 1] - \Pr[\text{Hyb}'_{i-1}(b) = 1]|.$$

If $|\Pr[\text{Hyb}'_{U(\lambda)+1}(b) = 1] - \Pr[\text{Hyb}'_{U(\lambda)}(b) = 1]| \leq \text{negl}(\lambda)$, we complete the proof of Theorem 5.2. ■

Lemma 5.7. *If the LWE assumption holds, it holds that $|\Pr[\text{Hyb}'_{i+1}(b) = 1] - \Pr[\text{Hyb}'_i(b) = 1]| \leq \text{negl}(\lambda)$.*

Proof. Note that the difference between these two games appears when the challenge query is sent at epoch i , so we can assume $\tilde{e} = i$. We start from $\text{Hyb}'_{i+1}(b)$ and gradually change it to $\text{Hyb}'_i(b)$. We define another sequence of games.

$\text{Hyb}_i^r(b)$: This is the same as $\text{Hyb}'_i(b)$ except that we use the hybrid update algorithm Hyb.Upd to simulate $\mathcal{O}.\text{Upd}$. More precisely, $\mathcal{O}.\text{Upd}(\text{ct}_{e-1})$ act as follows:

- If $(\cdot, \text{ct}_{e-1}, e-1; \mathbf{e}'_{e-1}; \boldsymbol{\mu}) \notin \mathcal{L}$, then return \perp
- Otherwise, $(\text{ct}_e, \mathbf{e}'_e) \leftarrow \text{Hyb.Upd}(\text{ct}_{e-1}, \mathbf{B}_e, \boldsymbol{\mu}; \mathbf{e}'_{e-1}, (\mathbf{R}_e, \mathbf{R}'_e))$.
- $\mathcal{L} := \mathcal{L} \cup \{(\cdot, \text{ct}_e, e; \mathbf{e}'_e, \boldsymbol{\mu})\}$.

Note that \mathbf{R}_e and \mathbf{R}'_e are randomness used in TokGen , so anyone can choose them. Simulators internally choose and record them.

Proposition 5.8. $|\Pr[\text{Hyb}'_i(b) = 1] - \Pr[\text{Hyb}_i^r(b) = 1]| \leq \text{negl}(\lambda)$.

It is easy to see Proposition 5.8 holds by Lemma 5.6. The next goal is proving $|\Pr[\text{Hyb}'_{i+1}(b) = 1] - \Pr[\text{Hyb}_i^r(b) = 1]| \leq \text{negl}(\lambda)$. We define the following games.

$\text{G}_j(i, b)$: This is the same as $\text{Hyb}_i^r(b)$ except the following difference.

- For $i \leq k < j$, pk_k and Δ_k are honestly generated as in the real.
- For $\text{fwr} \geq k \geq j$, pk_k and Δ_k are uniformly random.

That is, we gradually erase information about UE secret keys from newer epochs to older epochs. We note that $j \in [i, \text{fwr} + 1]$ and i is fixed. By the definition, we have

$$\text{G}_{\text{fwr}+1}(i+1, b) = \text{Hyb}'_{i+1}(b) \text{ and } \text{G}_{\text{fwr}+1}(i, b) = \text{Hyb}_i^r(b). \quad (4)$$

We prove that

$$|\Pr[\text{G}_{j+1}(i+1, b) = 1] - \Pr[\text{G}_j(i+1, b) = 1]| \leq \text{negl}(\lambda) \text{ for } j \in [i, \text{fwr}] \quad (5)$$

$$|\Pr[\text{G}_i(i+1, b) = 1] - \Pr[\text{G}_i(i, b) = 1]| \leq \text{negl}(\lambda) \quad (6)$$

$$|\Pr[\text{G}_{j+1}(i, b) = 1] - \Pr[\text{G}_j(i, b) = 1]| \leq \text{negl}(\lambda) \text{ for } j \in [i, \text{fwr}]. \quad (7)$$

From these equations, we immediately obtain

$$|\Pr[\text{G}_{\text{fwr}+1}(i+1, b) = 1] - \Pr[\text{G}_{\text{fwr}+1}(i, b)]| \leq \text{negl}(\lambda).$$

By combining this with Proposition 5.8 and Equation (4), we obtain what we want to prove (Lemma 5.7). Thus, all we must do is proving Equations (5) to (7).

First, we prove Equation (5). We define a few hybrid games as follows.

- **Game-0(b)**: This is the same as $\text{G}_{j+1}(i+1, b)$. At this point, public keys and tokens of epochs in $[i, j]$ are real values while those at epochs in $[j+1, \text{fwr}]$ are already random values.
- **Game-1(b)**: This is the same as **Game-0(b)** except that we modify the public key part of epoch j . We use $\mathbf{B}_j^+ \leftarrow \mathbb{Z}_q^{m \times \ell}$ instead of \mathbf{B}_j such that $(\mathbf{S}_j, \mathbf{B}_j) \leftarrow \text{Reg.Gen}(1^\lambda)$. Note that we do not use the secret key \mathbf{S}_j of epoch j anywhere in this game since Δ_{j+1} is already a random value.
- **Game-2(b)**: This is the same as **Game-1(b)** except that we modify the token generation algorithm for token Δ_j . We use $\Delta_j := (\mathbf{M}_j^+, \mathbf{N}_j^+) \leftarrow \mathbb{Z}_q^{n\eta \times (n+\ell)} \times \mathbb{Z}_q^{m \times (n+\ell)}$ instead of $(\mathbf{M}_j, \mathbf{N}_j) \leftarrow \text{TokGen}(k_{j-1}, k_j)$.

Obviously, $\text{Game-2}(b)$ is the same as $G_j(i+1, b)$. It is easy to see if we prove the following, we complete the proof of Equation (5).

Proposition 5.9. *If the LWE assumption holds, it holds that $|\Pr[\text{Game-1}(b) = 1] - \Pr[G_{j+1}(i+1, b) = 1]| \leq \text{negl}(\lambda)$.*

Proposition 5.10. *It holds that $|\Pr[\text{Game-2}(b) = 1] - \Pr[\text{Game-1}(b) = 1]| \leq \text{negl}(\lambda)$.*

We will prove these propositions above later.

Next, we prove Equation (6). The only difference between $G_i(i+1, b)$ and $G_i(i, b)$ is the challenge-equal ciphertext at epoch i . That is, $G_i(i, b)$ is the same as $G_i(i+1, b)$ except that we modify the challenge-equal ciphertext for b at epoch i . We use $(\bar{\mathbf{u}}, \bar{\mathbf{c}}) \leftarrow \mathbb{Z}_q^n \times \mathbb{Z}_q^\ell$ instead of $(\bar{\mathbf{u}}, \bar{\mathbf{c}}) \leftarrow \text{Upd}(\Delta_i^+, \bar{\mathbf{c}})$ (the case $b = 1$) or $(\bar{\mathbf{u}}, \bar{\mathbf{c}}) \leftarrow \text{Enc}(k_i, \bar{\mu}_0)$ (the case $b = 0$). We prove the following proposition later.

Proposition 5.11. *It holds that $|\Pr[G_i(i+1, b) = 1] - \Pr[G_i(i, b) = 1]| \leq \text{negl}(\lambda)$.*

Lastly, we prove Equation (7). Once the challenge-equal ciphertext at epoch i becomes random, we need to go back to games where public keys and tokens are real. In $G_j(i, b)$ for $j \in [i, \text{fwr}]$, public keys and tokens (from epochs j to fwr) are also random. We need to change them from random to real since we need to arrive at Hyb_i^r , where public keys and tokens are real (but ciphertext at epoch i is random). Thus, we need to prove Equation (7). These backward transitions are possible by using the proof of Equation (5) in a reverse manner. We summarize how public keys, update tokens, and challenge-equal ciphertexts at epoch i are generated in Figure 4.

Value	$G_{i+1}(i+1, b)$	Game-1	Game-2 = $G_i(i+1, b)$	$G_i(i, b)$
pk_i	$\text{Reg.Gen}(1^\lambda)$	$\text{Sim.Gen}(\text{pp})$	$\text{Sim.Gen}(\text{pp})$	$\text{Sim.Gen}(\text{pp})$
Δ_i	$\text{TokGen}(\text{sk}_{i-1}, \text{pk}_i)$	$\text{TokGen}(\text{sk}_{i-1}, \text{pk}_i)$	$\text{Sim.TokGen}(\text{pp})$	$\text{Sim.TokGen}(\text{pp})$
$\text{ct}_{i,1}^*$	$\text{Upd}(\Delta_i, \text{ct}_{i-1})$	$\text{Upd}(\Delta_i, \text{ct}_{i-1})$	$\text{Upd}(\Delta_i^+, \text{ct}_{i-1})$	$\text{Sim.Upd}(\text{pp})$
$\text{ct}_{i,0}^*$	$\text{Enc}(\text{pk}_i, \bar{\mu}_0)$	$\text{Enc}(\text{pk}_i, \bar{\mu}_0)$	$\text{Enc}(\text{pk}_i, \bar{\mu}_0)$	$\text{Sim.Enc}(\text{pp})$

Figure 4: The differences of public keys, update tokens, challenge-equal ciphertexts at epoch i in hybrid games. We focus the case where $i = \tilde{e}$.

Thus, we complete the proof of Lemma 5.7 if we prove Propositions 5.9 to 5.11. We write those proofs below. ■

Proofs of core propositions. We give the proofs of Propositions 5.9 to 5.11.

Proof of Proposition 5.9. We construct a reduction \mathcal{B} that solves the LWE problem by using the distinguisher \mathcal{A} for the two games.

Recall that the key k_j of epoch j consists of $(\text{sk}_j, \text{pk}_j)$. \mathcal{B} is given an LWE instance (A, B) and set $B_j := B$. That is, B is used as the public key pk_j of epoch j . Note that \mathcal{B} can simulate all values in epoch $k \in [0, T] \setminus [\text{fwl}, \text{fwr}]$ since all values in epoch k (outside the firewall) are independent of the secret key of epoch j . (Note that such values may be related to the public key of epoch j via tokens.) That is, \mathcal{B} can choose the secret key S_k . We also note that \mathcal{B} can simulate $\mathcal{O}.\text{Upd}$ by using $\text{Hyb}.\text{Upd}$. In $[\text{fwl}, \text{fwr}]$, values are related to the secret key S behind B . However, in $G_{j+1}(i+1, b)$ (and $\text{Game-1}(b)$), all values in $[j+1, \text{fwr}]$ are uniformly random values. Note that the original update token Δ_{j+1} needs sk_j and pk_{j+1} . However, Δ_{j+1} was already changed to Δ_{j+1}^+ , which is uniformly random value, and we do not need sk_j .

Thus, the issue is how to simulate values in epoch j' such that $j' \in [\text{fwl}, j]$ (including the case where $\text{fwl} = j$). As we see in the definition of TokGen , we do not need sk_j to generate Δ_j and \mathcal{B} can simulate Δ_j .

Therefore, \mathcal{B} can also simulate $\text{ct}_{j,b}^*$ for both $b = 0, 1$. For $j'' \in [\text{fwl}, j - 1]$, public keys and tokens are not related to sk_j . Thus, \mathcal{B} chooses $S_{j''}$ and can simulate all values $(\text{pk}_{j''}, \Delta_{j''}, \text{ct}_{j'',b}^*)$ by using the normal algorithms.

If $\mathbf{B} = \mathbf{A}\mathbf{S} + \mathbf{X}$ where $\mathbf{S} \leftarrow \mathbb{Z}_q^{n \times \ell}$ and $\mathbf{X} \leftarrow \chi^{m \times \ell}$, the distribution is the same as $G_{i+1}(i + 1, b)$. If \mathbf{B} is uniformly random, the distribution is the same as $\text{Game-1}(b)$. Therefore, \mathcal{B} distinguish the instance if \mathcal{A} distinguishes the two games. This completes the proof. ■

Proof of Proposition 5.10. The difference between these two games is as follows:

Game-1(b): $\Delta_j = (\mathbf{M}_j, \mathbf{N}_j)$:

$$\mathbf{M}_j := \mathbf{R}_j \cdot [\mathbf{A} \mid \mathbf{B}_j] + [\mathbf{O} \mid -\text{P2}(S_{j-1})], \mathbf{N}_j := \mathbf{R}'_j \cdot [\mathbf{A} \mid \mathbf{B}_j],$$

$$\text{where } \mathbf{R}_j \leftarrow \{-1, +1\}^{n\eta \times m}, \mathbf{R}'_j \leftarrow \{-1, +1\}^{m \times m}.$$

Game-2(b): $\Delta_j^+ = (\mathbf{M}_j^+, \mathbf{N}_j^+)$:

$$(\mathbf{M}_j^+, \mathbf{N}_j^+) \leftarrow \mathbb{Z}_q^{n\eta \times (n+\ell)} \times \mathbb{Z}_q^{m \times (n+\ell)}.$$

In Game-1(b) and Game-2(b), the public key $\mathbf{B}_j \leftarrow \mathbb{Z}_q^{m \times \ell}$ is uniformly random. Thus, we can apply the leftover hash lemma (Lemma 2.3) and these differences are statistically indistinguishable. This completes the proof. ■

Proof of Proposition 5.11. The difference between these two games is as follows: For $b = 1$,

$G_i(i + 1, 1)$: $\text{ct}_{i,1}^* = (\tilde{\mathbf{u}}, \tilde{\mathbf{c}})$:

$$(\mathbf{u}' + \tilde{\mathbf{u}}, \mathbf{c}_{i-1} + \mathbf{c}' + \tilde{\mathbf{v}}) = (\mathbf{0}, \mathbf{c}_{i-1}) + \text{BD}(\mathbf{u}_i)\mathbf{M}_i^+ + \tilde{\mathbf{r}}\mathbf{N}_i^+,$$

$$\text{where } \tilde{\mathbf{r}} \leftarrow \{-1, +1\}^m.$$

$G_i(i, 1)$: $\text{ct}_{i,1}^* = (\tilde{\mathbf{u}}, \tilde{\mathbf{c}})$:

$$(\tilde{\mathbf{u}}, \tilde{\mathbf{c}}) \leftarrow \mathbb{Z}_q^n \times \mathbb{Z}_q^\ell.$$

In $G_i(i + 1, b)$ and $G_i(i, b)$, \mathbf{N}_i^+ is uniformly random. Thus, we can apply the leftover hash lemma (Lemma 2.3) and these differences are statistically indistinguishable. For $b = 0$,

$G_i(i + 1, 0)$: $\text{ct}_{i,0}^* = (\mathbf{u}, \mathbf{c})$:

$$(\mathbf{r}\mathbf{A}_i, \mathbf{r}\mathbf{B}_i^+ + \mathbf{e}' + \lfloor q/2 \rfloor \boldsymbol{\mu}_0),$$

$$\text{where } \mathbf{A} \leftarrow \mathbb{Z}_q^{m \times \ell}, \mathbf{r} \leftarrow \{-1, +1\}^m, \mathbf{e}' \leftarrow \chi_{\text{ns}}^\ell, \text{ and } \mathbf{B}_i^+ \leftarrow \mathbb{Z}_q^{m \times \ell}.$$

$G_i(i, 0)$: $\text{ct}_{i,0}^* = (\mathbf{u}, \mathbf{c})$:

$$(\mathbf{u}, \mathbf{c}) \leftarrow \mathbb{Z}_q^n \times \mathbb{Z}_q^\ell.$$

In $G_i(i + 1, b)$ and $G_i(i, b)$, the public key $\mathbf{B}_i^+ \leftarrow \mathbb{Z}_q^{m \times \ell}$ is uniformly random. Thus, we can apply the leftover hash lemma (Lemma 2.3) and these differences are statistically indistinguishable. This completes the proof. ■

Acknowledgments

The author would like to thank Takahiro Matsuda for finding an error in the UE scheme based on indistinguishability obfuscation proposed in the previous version. The author would like to thank Fuyuki Kitagawa for giving a pointer to the right reference about an RKA-secure AE scheme. The author also thanks anonymous reviewers of PKC 2022 for useful comments.

References

- [ABPW13] Yoshinori Aono, Xavier Boyen, Le Trieu Phong, and Lihua Wang. Key-private proxy re-encryption under LWE. In Goutam Paul and Serge Vaudenay, editors, *INDOCRYPT 2013*, volume 8250 of *LNCS*, pages 1–18. Springer, Heidelberg, December 2013. (Cited on page 4, 12.)
- [ACPS09] Benny Applebaum, David Cash, Chris Peikert, and Amit Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In Shai Halevi, editor, *CRYPTO 2009*, volume 5677 of *LNCS*, pages 595–618. Springer, Heidelberg, August 2009. (Cited on page 6, 7.)
- [AJL⁺12] Gilad Asharov, Abhishek Jain, Adriana López-Alt, Eran Tromer, Vinod Vaikuntanathan, and Daniel Wichs. Multiparty computation with low communication, computation and interaction via threshold FHE. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 483–501. Springer, Heidelberg, April 2012. (Cited on page 4, 6.)
- [AMP19] Navid Alamati, Hart Montgomery, and Sikhar Patranabis. Symmetric primitives with structured secrets. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part I*, volume 11692 of *LNCS*, pages 650–679. Springer, Heidelberg, August 2019. (Cited on page 4, 12.)
- [Ban93] Wojciech Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 4(296):625–635, 1993. (Cited on page 6.)
- [BDGJ20] Colin Boyd, Gareth T. Davies, Kristian Gjøsteen, and Yao Jiang. Fast and secure updatable encryption. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part I*, volume 12170 of *LNCS*, pages 464–493. Springer, Heidelberg, August 2020. (Cited on page 1, 2, 3, 4, 7, 11, 12, 13, 16, 21, 29.)
- [BEKS20] Dan Boneh, Saba Eskandarian, Sam Kim, and Maurice Shih. Improving speed and security in updatable encryption schemes. In Shiho Moriai and Huaxiong Wang, editors, *ASIACRYPT 2020, Part III*, volume 12493 of *LNCS*, pages 559–589. Springer, Heidelberg, December 2020. (Cited on page 1, 3.)
- [BGH13] Zvika Brakerski, Craig Gentry, and Shai Halevi. Packed ciphertexts in LWE-based homomorphic encryption. In Kaoru Kurosawa and Goichiro Hanaoka, editors, *PKC 2013*, volume 7778 of *LNCS*, pages 1–13. Springer, Heidelberg, February / March 2013. (Cited on page 17.)
- [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. *Journal of the ACM*, 59(2):6:1–6:48, 2012. (Cited on page 13.)
- [BGV14] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (leveled) fully homomorphic encryption without bootstrapping. *ACM Trans. Comput. Theory*, 6(3):13:1–13:36, 2014. (Cited on page 17.)
- [BLMR13] Dan Boneh, Kevin Lewi, Hart William Montgomery, and Ananth Raghunathan. Key homomorphic PRFs and their applications. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 410–428. Springer, Heidelberg, August 2013. (Cited on page 1, 3.)

- [BLP⁺13] Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of learning with errors. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th ACM STOC*, pages 575–584. ACM Press, June 2013. (Cited on page 7.)
- [BV11] Zvika Brakerski and Vinod Vaikuntanathan. Fully homomorphic encryption from ring-LWE and security for key dependent messages. In Phillip Rogaway, editor, *CRYPTO 2011*, volume 6841 of *LNCS*, pages 505–524. Springer, Heidelberg, August 2011. (Cited on page 4, 17.)
- [BV14] Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. *SIAM Journal on Computing*, 43(2):831–871, 2014. (Cited on page 4, 17.)
- [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 Hugo Krawczyk, editor, *PKC 2014*, volume 8383 of *LNCS*, pages 95–112. Springer, Heidelberg, March 2014. (Cited on page 4, 12.)
- [CLT20] Long Chen, Yanan Li, and Qiang Tang. CCA updatable encryption against malicious re-encryption attacks. In Shiho Moriai and Huaxiong Wang, editors, *ASIACRYPT 2020, Part III*, volume 12493 of *LNCS*, pages 590–620. Springer, Heidelberg, December 2020. (Cited on page 1, 3.)
- [DN21] Nico Döttling and Ryo Nishimaki. Universal proxy re-encryption. In Juan Garay, editor, *PKC 2021, Part I*, volume 12710 of *LNCS*, pages 512–542. Springer, Heidelberg, May 2021. (Cited on page 5.)
- [EPRS17] Adam Everspaugh, Kenneth G. Paterson, Thomas Ristenpart, and Samuel Scott. Key rotation for authenticated encryption. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part III*, volume 10403 of *LNCS*, pages 98–129. Springer, Heidelberg, August 2017. (Cited on page 1, 3, 12.)
- [FMM21] Andrés Fabrega, Ueli Maurer, and Marta Mularczyk. A fresh approach to updatable symmetric encryption. Cryptology ePrint Archive, Report 2021/559, 2021. <https://eprint.iacr.org/2021/559>. (Cited on page 3.)
- [Gen09] Craig Gentry. *A fully homomorphic encryption scheme*. PhD thesis, Stanford University, 2009. crypto.stanford.edu/craig. (Cited on page 4, 12.)
- [GP23] Yao Jiang Galteland and Jiaxin Pan. Backward-leak uni-directional updatable encryption from (homomorphic) public key encryption. In Alexandra Boldyreva and Vladimir Kolesnikov, editors, *PKC 2023, Part II*, volume 13941 of *LNCS*, pages 399–428. Springer, Heidelberg, May 2023. (Cited on page 2.)
- [GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 197–206. ACM Press, May 2008. (Cited on page 6.)
- [HLL16] Shuai Han, Shengli Liu, and Lin Lyu. Efficient KDM-CCA secure public-key encryption for polynomial functions. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *ASIACRYPT 2016, Part II*, volume 10032 of *LNCS*, pages 307–338. Springer, Heidelberg, December 2016. (Cited on page 12.)

- [Jia20] Yao Jiang. The direction of updatable encryption does not matter much. In Shiho Moriai and Huaxiong Wang, editors, *ASIACRYPT 2020, Part III*, volume 12493 of *LNCS*, pages 529–558. Springer, Heidelberg, December 2020. (Cited on page 1, 2, 3, 4, 7, 10, 11, 12, 13, 14, 15, 16, 21, 29.)
- [KLR19] Michael Kloof, Anja Lehmann, and Andy Rupp. (R)CCA secure updatable encryption with integrity protection. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part I*, volume 11476 of *LNCS*, pages 68–99. Springer, Heidelberg, May 2019. (Cited on page 1, 3, 4, 11, 13, 21.)
- [LP11] Richard Lindner and Chris Peikert. Better key sizes (and attacks) for LWE-based encryption. In Aggelos Kiayias, editor, *CT-RSA 2011*, volume 6558 of *LNCS*, pages 319–339. Springer, Heidelberg, February 2011. (Cited on page 6.)
- [LR21] Françoise Levy-dit-Vehel and Maxime Roméas. A composable look at updatable encryption. Cryptology ePrint Archive, Report 2021/538, 2021. <https://eprint.iacr.org/2021/538>. (Cited on page 3.)
- [LT18] Anja Lehmann and Björn Tackmann. Updatable encryption with post-compromise security. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part III*, volume 10822 of *LNCS*, pages 685–716. Springer, Heidelberg, April / May 2018. (Cited on page 1, 3, 4, 8, 9, 11, 12, 13, 16, 21.)
- [NX15] Ryo Nishimaki and Keita Xagawa. Key-private proxy re-encryption from lattices, revisited. *IEICE Transactions*, 98-A(1):100–116, 2015. (Cited on page 4, 12, 17.)
- [Pei09] Chris Peikert. Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In Michael Mitzenmacher, editor, *41st ACM STOC*, pages 333–342. ACM Press, May / June 2009. (Cited on page 7.)
- [Reg09] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *Journal of the ACM*, 56(6):34:1–34:40, 2009. (Cited on page 6, 7, 17, 19.)
- [SS21] Daniel Slamanig and Christoph Striecks. Puncture ’em all: Updatable encryption with no-directional key updates and expiring ciphertexts. Cryptology ePrint Archive, Report 2021/268, 2021. <https://eprint.iacr.org/2021/268>. (Cited on page 3.)

A More Discussion on Key Updates

Why does the direction matter? The question is: “What is the source of the difference between forward and backward-leak uni-directional key updates?” The point is that ciphertext updates can always go forward but cannot necessarily go backward. Look at the example in Figure 3 again. If we set $\mathcal{C} = \{3\}$ in the forward-leak uni-directional setting, we can go forward by using corrupted tokens Δ_4, Δ_5 . Thus, even if epoch 3 is not corrupted, we can convert a ciphertext under k_3 into one under k_5 , which is corrupted. However, in the backward-leak uni-directional setting, if we set $\mathcal{C} = \{5\}$, we cannot go backward even if we use Δ_4, Δ_5 . Note that we focus on the uni-directional ciphertext updates setting here. It is natural to consider only uni-directional ciphertext updates in the backward-leak uni-directional key updates since we do not need sk_{e+1} to generate Δ_{e+1} as we observed so far. This asymmetry incurs the difference.

Based on this observation, if we consider “backward uni-directional ciphertext updates” (or “downgrade only updates”), a similar equivalence result to Jiang’s one holds even in the backward-leak uni-directional key updates setting. However, note that the backward uni-directional ciphertext updates setting is quite unnatural and artificial, as we discussed in Section 4.2.

B Other Security Definitions for UE

Definition B.1 (Integrity for Updatable Encryption [BDGJ20, Jia20]). For $\text{atk} \in \{\text{ctxt}, \text{ptxt}\}$, the game $\text{Exp}_{\Sigma, \mathcal{A}}^{\text{int-atk}}(\lambda, b)$ is formalized as follows.

- Invoke Setup and set $\text{win} := 0$.
- Let $\mathcal{O} := \mathcal{O}.\{\text{Enc}, \text{Next}, \text{Upd}, \text{Corr}, \text{Try}\}$.
- Run $\mathcal{A}^{\mathcal{O}}(1^\lambda)$.
- If $\text{twf} = 1$ then $\text{win} := 1$
- return win

We say a UE scheme is INT-atk secure if it holds

$$\text{Adv}_{\Sigma, \mathcal{A}}^{\text{int-atk}}(\lambda) := |\Pr[\text{Exp}_{\Sigma, \mathcal{A}}^{\text{int-atk}}(\lambda, 0) = 1] - \Pr[\text{Exp}_{\Sigma, \mathcal{A}}^{\text{int-atk}}(\lambda, 1) = 1]| \leq \text{negl}(\lambda).$$

The definitions of oracles are described in Figure 1.

Definition B.2 ((kk, cc)-variant of integrity [Jia20]). Let UE be a UE scheme. Then the (kk, cc)-notion advantage, for $\text{kk} \in \{\text{no}, \text{uni}, \text{bi}\}$, $\text{cc} \in \{\text{uni}, \text{bi}\}$ and $\text{notion} \in \{\text{int-ctxt}, \text{int-ptxt}\}$, of an adversary \mathcal{A} against UE is defined as

$$\text{Adv}_{\text{UE}, \mathcal{A}}^{(\text{kk}, \text{cc})\text{-notion}}(1^\lambda) := |\Pr[\text{Exp}_{\text{UE}, \mathcal{A}}^{(\text{kk}, \text{cc})\text{-notion}}(\lambda, 0) = 1] - \Pr[\text{Exp}_{\text{UE}, \mathcal{A}}^{(\text{kk}, \text{cc})\text{-notion}}(\lambda, 1) = 1]|,$$

where $\text{Exp}_{\text{UE}, \mathcal{A}}^{(\text{kk}, \text{cc})\text{-notion}}(\lambda, b)$ is the same as the experiment $\text{Exp}_{\text{UE}, \mathcal{A}}^{\text{notion}}(\lambda, b)$ in Definition B.1 except for all leakage sets are both in the kk-directional key updates and cc-directional ciphertext updates.