

Fully Secure Searchable Encryption from PRFs, Pairings, and Lattices

Hiroto Shinoki*, Hisayoshi Sato, and Masayuki Yoshino

Hitachi, Ltd., Japan

October 11, 2024

Abstract

Searchable encryption is a cryptographic primitive that allows us to perform searches on encrypted data. Searchable encryption schemes require that ciphertexts do not leak information about keywords. However, most of the existing schemes do not achieve the security notion that trapdoors do not leak information. Shen et al. (TCC 2009) proposed a security notion called full security, which includes both ciphertext privacy and trapdoor privacy, but there are few fully secure constructions. Full security is defined for the secret key settings since it is known that public key schemes cannot achieve the trapdoor privacy in principle.

In this paper, we construct a query-bounded fully secure scheme from pseudorandom functions. In addition, we propose two types of efficient (unbounded) fully secure schemes, each of which is based on bilinear groups and lattices respectively. We then analyze the existing constructions. First, we simplify the Cheng et al. scheme (Information Sciences 2023) and prove its security. This scheme had not been proved to be secure. Second, we show that the Li-Boyen pairing-based scheme (IACR CiC 2024) does not achieve the trapdoor privacy, not as claimed.

1 Introduction

Searchable encryption is a cryptosystem that allows keyword search on encrypted data. There are two types of searchable encryption: Searchable Symmetric Encryption (SSE) [33] and Public key Encryption with Keyword Search (PEKS) [7]. The secret key is used for encryption in SSE, while the public key is used in PEKS. To search for a keyword w , one generates a trapdoor T_w using the secret key. Then, by running the test algorithm, anyone can check whether the given ciphertext is associated with w .

Searchable encryption requires that ciphertexts leak no information on the keywords. However, most of the existing schemes do not satisfy the security notion that trapdoors leak no information. With respect to the public key setting, it is known that PEKS cannot achieve trapdoor privacy in principle. In fact, Byun et al. [10] showed that PEKS is vulnerable to the Keyword Guessing Attack (KGA) and the associated keyword is leaked from the trapdoor when the keyword space is small. In KGA, the attacker tries to find which keyword corresponds to the given trapdoor. The attacker can choose a keyword, encrypt it, and check whether the chosen keyword corresponds to the trapdoor by running the test algorithm. If the number of candidate keywords is small, the attacker can easily find the correct keyword. SSE can achieve the trapdoor privacy, but many SSE schemes use deterministic trapdoor generation algorithms. In these cases, anyone can know whether the given two trapdoors are associated with the same keyword.

In 2009, Shen et al. [32] proposed the security notion called predicate privacy for Predicate Encryption (PE) in the secret key setting. Predicate-private PE for equality is equivalent to

*hiroto.shinoki.sw@hitachi.com

trapdoor-private SSE. They proposed an inner product predicate encryption with selective predicate privacy using the composite order bilinear groups. They also formulated the notion of full security, which implies both ciphertext privacy and trapdoor privacy. These security notions can be generalized to Functional Encryption (FE), and there has been a lot of research on function-private FE schemes [2, 3, 4, 6, 9, 13, 20, 22, 34]. However, these schemes have different functionalities from SSE or have complex structure to support rich functionalities. Until recently, there has been no research on the construction of simple schemes that support the equality predicate.

In 2024, Li and Boyen [21] proposed two efficient trapdoor-private SSE schemes. These schemes are based on pairings and lattices respectively. They also proposed a generic construction of Public key Authenticated Encryption with Keyword Search (PAEKS) from trapdoor-private SSE and Non-Interactive Key Exchange (NIKE). PAEKS is a variant primitive of PEKS, which was introduced by Huang and Li [19] in 2017 to prevent KGA. As mentioned above, PEKS is vulnerable to KGA in principle. In PAEKS, the data sender has its own secret key to authenticate the ciphertext. Since the test algorithm works correctly only if the ciphertext is authenticated, PAEKS has the potential to prevent KGA. Many PAEKS schemes have been proposed [11, 12, 16, 23, 28, 29, 30], but most PAEKS schemes achieve only limited trapdoor privacy. Besides the Li-Boyen constructions, there are two PAEKS with (unlimited) trapdoor privacy [11, 12]. These schemes imply trapdoor-private SSE.

1.1 Our Contributions

In this paper, we propose three types of fully secure SSE.

The first construction uses a pseudorandom function. This construction achieves the bounded version of full security. In the bounded full security setting, the upper bound of the number of queries must be determined before the key generation phase. Although this setting is a weaker variant, the security can be proved without any additional assumption. The other schemes we propose achieve (unbounded) full security by introducing additional assumptions.

The second construction uses a pseudorandom function and bilinear groups. This construction is based on the Uniform Matrix Decisional Diffie-Hellman assumption, which we write as U_k -MDDH. This assumption is parameterized by an integer $k \geq 1$. Increasing k makes the assumption weaker, but makes the scheme less efficient. Note that the U_k -MDDH assumption is weaker than the k -Lin assumption.

The third construction uses a pseudorandom function and lattices. This construction is based on the (Ring) Learning with Errors assumption. In addition, it is quite simple compared to the existing schemes. In fact, it does not involve complicated lattice algorithms such as the preimage sampling.

The comparison of SSE with trapdoor privacy is summarized in Table 1. CM22 and CQFM23 are PAEKS with trapdoor privacy. It has been claimed that the security of CQFM23 follows from the Computational Oracle Diffie-Hellman (CODH) assumption in [12], but Li and Boyen [21] pointed out that there is an error in the security analysis. It was also shown that CODH is insufficient and that at least the Decisional Bilinear Diffie-Hellman (DBDH) assumption is required. Thus, no security proof has been given for CQFM23. Note that it seems difficult to prove the security directly from the DBDH assumption. This is because the bilinear map only appears in the search algorithm and does not appear in the security analysis. In this paper, we simplify CQFM23 and show that its security is based on the Decisional Linear (DLin) assumption. This scheme can be seen as a variant of our U_2 -MDDH-based construction and can be extended to k -Lin-based construction. We also show that CQFM23 can be generalized to the asymmetric bilinear group setting. In this case, the full security is based on the bilateral k -Lin assumption. We also analyze the existing constructions by Li and Boyen [21]. In particular, their security proof for the pairing-based scheme is incorrect, and the trapdoor privacy can be broken.

Ours1 is the first SSE from symmetric primitives that achieves (at least) bounded trapdoor privacy. Ours2 is more efficient than the existing pairing-based constructions when $k = 1$. For lattice-based constructions, Ours3 is much more efficient compared to CM22 and LB24-L. Ours3 is simple and does not use the lattice trapdoor generation algorithm [5, 26] or the preimage

sampling algorithm [18]. Compared to LB24-N, Ours3 is less efficient in terms of space efficiency, but it has the following advantages.

- LB24-N uses complicated algorithms such as the NTRU lattice generation and the preimage sampling in the encryption algorithm. The encryption algorithm of Ours3 consists of simple calculations.
- The underlying assumption is weaker.

Finally, we note that ciphertext privacy and trapdoor privacy have been considered in the previous research, but full security has not. Full security implies “ciphertext privacy and trapdoor privacy,” but not vice versa. In this paper, we prove that our constructions achieve full security instead of proving ciphertext privacy and trapdoor privacy.

Table 1: Comparison of SSE with Trapdoor Privacy

	Ciphertext	Trapdoor	Assumption	Remarks
CM22 [11]	$\Omega(\kappa_{\text{ct}} n \log^2 q)$	$\Omega(\kappa_{\text{td}} n \log^2 q)$	LWE	
CQFM23 [12]	$4 G $	$4 G $	PRF, DLin	Slightly modified
LB24-P [21]	$2 G_1 + G_T $	$2 G_2 $	(PRF, DBDH)	Broken
LB24-L [21]	$\Omega(\kappa_{\text{ct}} n \log^2 q)$	$\Omega(\kappa_{\text{td}} n \log^2 q)$	PRF, LWE	
LB24-N [21]	$2N \log q$	$2N \log q$	PRF, RLWE, NTRU	
Ours1	$Q_{\text{ct}} + Q_{\text{td}} \cdot \mathbb{F} $	$Q_{\text{td}} + Q_{\text{ct}} \cdot \mathbb{F} $	PRF	Q -bounded
Ours2	$2k G_1 $	$2k G_2 $	PRF, U_k -MDDH	
Ours3 (LWE)	$\kappa_{\text{ct}} \cdot 3n \log q$	$\kappa_{\text{td}} \cdot 3n \log q$	PRF, LWE	
Ours3 (RLWE)	$3N \log q$	$3N \log q$	PRF, RLWE	
Gen. CQFM23	$2k G_1 $	$2k G_2 $	PRF, bil- k -Lin	

For a finite set S , $|S|$ denotes the bit size of a random element in S . CQFM23 uses a symmetric bilinear map from $G \times G$. LB24-P, Ours2, and Generalized CQFM23 use a bilinear map from $G_1 \times G_2$ to G_T . Q_{ct} and Q_{td} denote upper bounds of the number of ciphertext queries and trapdoor queries respectively. For simplicity, it is assumed that Q_{ct} and Q_{td} are large enough in Ours1.

1.2 Paper Organization

In Section 2, we summarize the basic definitions used in this paper. In Section 3, 4, and 5, we propose the PRF-based bounded construction (Ours1), the pairing-based construction (Ours2), and the lattice-based construction (Ours3) respectively. In Section 6, we analyze the Cheng et al. scheme (CQFM23) and show that the security can be reduced to the (bil-)DLin assumption. In Section 7, we analyze the Li-Boyer schemes (LB24-P, LB24-N).

2 Preliminaries

The basic notations used in this paper are summarized here. “Probabilistic Polynomial-Time” is abbreviated to “PPT”. For a finite set S , $x \xleftarrow{\$} S$ means sampling x uniformly at random from S , and $\text{Unif}(S)$ denotes the uniform distribution over S . We say that a function $f : \mathbb{N} \rightarrow \mathbb{R}$ is negligible and write $f(n) = \text{negl}(n)$ if for any positive integer k there exists an integer n_k such that $|f(n)| < n^{-k}$ for any $n > n_k$. We say that a probability $p(n)$ is overwhelming if $1 - p(n)$ is negligible. Let \mathbb{Z}_q denote the quotient ring $\mathbb{Z}/q\mathbb{Z}$. Let $\log x$ denote $\log_2 x$. We use the following notations for matrix concatenation:

$$[A|B] = \begin{pmatrix} A & B \end{pmatrix}, \quad [A; C] = \begin{pmatrix} A \\ C \end{pmatrix}.$$

2.1 Pseudorandom Function (PRF)

Let $\{\mathcal{K}_\lambda\}_{\lambda \in \mathbb{N}}$, $\{\mathcal{X}_\lambda\}_{\lambda \in \mathbb{N}}$, and $\{\mathcal{Y}_\lambda\}_{\lambda \in \mathbb{N}}$ be families of finite sets. A family of functions $F = \{F_\lambda : \mathcal{K}_\lambda \times \mathcal{X}_\lambda \rightarrow \mathcal{Y}_\lambda\}_{\lambda \in \mathbb{N}}$ is said to be pseudorandom if for any PPT adversary \mathcal{A} ,

$$\text{Adv}_{\mathcal{A}, F}^{\text{PRF}}(\lambda) := \frac{1}{2} |\Pr[\mathbf{K} \xleftarrow{\$} \mathcal{K}_\lambda : 1 \leftarrow \mathcal{A}^{F_\lambda(\mathbf{K}, \cdot)}(1^\lambda)] - \Pr[f \xleftarrow{\$} \text{Func}[\mathcal{X}_\lambda, \mathcal{Y}_\lambda] : 1 \leftarrow \mathcal{A}^{f(\cdot)}(1^\lambda)]|$$

is negligible in λ . Here, $\text{Func}[A, B]$ denotes the set of all functions from A to B . In this paper, we often omit the description of λ and write $F_\lambda(\mathbf{K}, x)$ as $F(\mathbf{K}, x)$ or $F_{\mathbf{K}}(x)$.

2.2 Bilinear Groups

The bilinear group generator \mathcal{G} is a PPT algorithm that outputs $\mathbb{G} := (p, G_1, G_2, G_T, g_1, g_2, e) \leftarrow \mathcal{G}(1^\lambda)$. Here, p is a prime of $\Theta(\lambda)$ bits, G_1, G_2, G_T are cyclic groups of order p , g_1 is a generator of G_1 , g_2 is a generator of G_2 , and $e : G_1 \times G_2 \rightarrow G_T$ is a non-degenerate bilinear map. For matrices $A = (a_{i,j}) \in \mathbb{Z}_p^{m \times n}$, we write $[A]_1 := (g_1^{a_{i,j}}) \in G_1^{m \times n}$, $[A]_2 = (g_2^{a_{i,j}}) \in G_2^{m \times n}$, and $[A]_T = (e(g_1, g_2)^{a_{i,j}}) \in G_T^{m \times n}$. Also, for matrices A and B , we write $e([A]_1, [B]_2) := [AB]_T$. Note that $[AB]_T$ is efficiently computable from $([A]_1, [B]_2)$.

Definition 1 (Matrix Diffie-Hellman Assumption). Let $\ell > k$ and $\mathcal{D}_{\ell,k}$ be a distribution on $\mathbb{Z}_p^{\ell \times k}$. We say that the $\mathcal{D}_{\ell,k}$ -Matrix Decisional Diffie-Hellman ($\mathcal{D}_{\ell,k}$ -MDDH) assumption on G_i ($i = 1, 2$) holds if $(\mathbb{G}, [A]_i, [Ar]_i)$ is computationally indistinguishable from $(\mathbb{G}, [A]_i, [u]_i)$ where $A \leftarrow \mathcal{D}_{\ell,k}$, $r \xleftarrow{\$} \mathbb{Z}_p^k$, and $u \xleftarrow{\$} \mathbb{Z}_p^\ell$. We just say that the $\mathcal{D}_{\ell,k}$ -MDDH assumption holds if it holds for both G_1 and G_2 .

We say that the bilateral $\mathcal{D}_{\ell,k}$ -Matrix Decisional Diffie-Hellman (bil- $\mathcal{D}_{\ell,k}$ -MDDH) assumption holds if $(\mathbb{G}, [A]_1, [Ar]_1, [A]_2, [Ar]_2)$ is computationally indistinguishable from $(\mathbb{G}, [A]_1, [u]_1, [A]_2, [u]_2)$.

In this paper, we mainly consider the case that $\ell = k + 1$ and $\mathcal{D}_{\ell,k}$ is uniform. We call this case the U_k -MDDH assumption. It is known that U_k -MDDH is weaker than any $\mathcal{D}_{k+1,k}$ -MDDH. In particular, the k -Lin assumption implies the U_k -MDDH assumption. We use U_k -MDDH in the following form.

Lemma 1. For $i = 1, 2$ and $r \in \mathbb{Z}_p^k$, let $\mathcal{O}_{G_i}^r$ be an oracle that samples $a \xleftarrow{\$} \mathbb{Z}_p^k$ and outputs $([a]_i, [a^\top r]_i)$. Let \mathcal{O}_{G_i} be an oracle that samples $a \xleftarrow{\$} \mathbb{Z}_p^k$, $u \xleftarrow{\$} \mathbb{Z}_p$ and outputs $([a]_i, [u]_i)$. If the U_k -MDDH assumption on G_i holds, for any PPT algorithm \mathcal{A} ,

$$\text{Adv}_{i, \mathcal{A}}^{U_k\text{-MDDH}}(\lambda) := \frac{1}{2} |\Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}_{G_i}^r}(1^\lambda)] - \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}_{G_i}}(1^\lambda)]|$$

is negligible for a randomly chosen secret $r \xleftarrow{\$} \mathbb{Z}_p^k$.

Proof. Using a U_k -MDDH instance $([A]_i, [s]_i)$, the oracles in this lemma can be simulated in the following way: one samples $t \xleftarrow{\$} \mathbb{Z}_p^{k+1}$ and outputs $([t^\top A]_i, [t^\top s]_i)$. Since the bit length of p is $\Theta(\lambda)$, $A \in \mathbb{Z}_p^{(k+1) \times k}$ is rank k with an overwhelming probability. Thus, it simulates $\mathcal{O}_{G_i}^r$ when $s = Ar$. Similarly, when s is random, $[A]_i [s]_i \in \mathbb{Z}_p^{(k+1) \times (k+1)}$ is full-rank with an overwhelming probability. Thus, it simulates \mathcal{O}_{G_i} in this case. \square

For $m \geq 2$, m -instance version of U_k -MDDH is defined by replacing $r \xleftarrow{\$} \mathbb{Z}_p^k$ with $r \xleftarrow{\$} \mathbb{Z}_p^{k \times m}$ and $u \xleftarrow{\$} \mathbb{Z}_p$ with $u \xleftarrow{\$} \mathbb{Z}_p^m$ in the setting of Lemma 1. In this case, we write the advantage as $\text{Adv}_{i, \mathcal{A}}^{U_k\text{-MDDH}_m}(\lambda)$. Then, the following tight reduction holds.

Lemma 2 ([17]). For any PPT adversary \mathcal{A} , there exists a PPT algorithm \mathcal{B} such that

$$\text{Adv}_{i, \mathcal{A}}^{U_k\text{-MDDH}_m}(\lambda) \leq \text{Adv}_{i, \mathcal{B}}^{U_k\text{-MDDH}}(\lambda) + \frac{1}{2(p-1)}.$$

2.3 Learning with Errors (LWE)

We introduce the assumptions on lattices used in this paper.

Definition 2 (LWE Assumption). For each positive integer λ , let $n(\lambda)$ be a positive integer, $q(\lambda) \geq 3$ be an integer, and $\chi(\lambda)$ be an error distribution on \mathbb{Z}_q .

(Normal) LWE [31] For $s \in \mathbb{Z}_q^n$, let \mathcal{O}_s be the oracle that samples $a \xleftarrow{\$} \mathbb{Z}_q^n$, $x \leftarrow \chi$ and outputs $(a, a^\top s + x)$. Let \mathcal{O}_\S be the oracle that outputs $(a, b) \xleftarrow{\$} \mathbb{Z}_q^n \times \mathbb{Z}_q$. We say that the (n, q, χ) -LWE assumption holds if for any PPT algorithm \mathcal{A} ,

$$\text{Adv}_{\mathcal{A}}^{\text{LWE}_{(n,q,\chi)}}(\lambda) := \frac{1}{2} |\Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}_s}(1^\lambda)] - \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}_\S}(1^\lambda)]|$$

is negligible for a randomly chosen secret $s \xleftarrow{\$} \mathbb{Z}_q^n$. If \mathcal{A} makes at most m queries, we write the advantage as $\text{Adv}_{\mathcal{A}}^{\text{LWE}_{n,m,q,\chi}}(\lambda)$. We use similar notations for LWE-variants defined below.

Non-uniform LWE [8] Let η be a distribution on \mathbb{Z}_q^n . For $s \in \mathbb{Z}_q^n$, let \mathcal{O}_s^η be the oracle that samples $a \leftarrow \eta$, $x \leftarrow \chi$ and outputs $(a, a^\top s + x)$. Let \mathcal{O}_\S^η be the oracle that outputs $(a, b) \leftarrow \eta \times \text{Unif}(\mathbb{Z}_q)$. We say that the (n, q, χ, η) -NLWE assumption holds if for any PPT algorithm \mathcal{A} ,

$$\text{Adv}_{\mathcal{A}}^{\text{NLWE}_{(n,q,\chi,\eta)}}(\lambda) := \frac{1}{2} |\Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}_s^\eta}(1^\lambda)] - \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}_\S^\eta}(1^\lambda)]|$$

is negligible for a randomly chosen secret $s \xleftarrow{\$} \mathbb{Z}_q^n$.

Boneh et al. [8] proved that the (k, q, χ, η) -NLWE problem is as hard as the (n, q, χ) -LWE problem if η satisfies the property called coset sampleability.

Definition 3 (Coset Sampleability [8]). A distribution η on \mathbb{Z}_q^k is called n -coset sampleable if there exist two PPT algorithms `MatrixGen` and `SamplePre` such that:

- `MatrixGen`(1^λ) outputs $M \in \mathbb{Z}_q^{n \times k}$ and auxiliary data T .
- `SamplePre`($z \in \mathbb{Z}_q^n, M, T$) outputs $y \in \mathbb{Z}_q^k$ such that $My = z$. In addition, if z is sampled uniformly at random, the distribution of y is η .

Lemma 3 ([8]). If η is n -coset sampleable, the (k, q, χ, η) -NLWE problem is at least as hard as the (n, q, χ) -LWE problem. Namely, for any PPT adversary \mathcal{A} , there exists a PPT algorithm \mathcal{B} such that

$$\text{Adv}_{\mathcal{A}}^{\text{NLWE}_{k,m,q,\chi,\eta}}(\lambda) \leq \text{Adv}_{\mathcal{B}}^{\text{LWE}_{n,m,q,\chi}}(\lambda).$$

Finally, we review the notion of Ring-LWE (RLWE) assumption.

Definition 4 (RLWE Assumption). Let $N(\lambda)$ be a power of 2, $q(\lambda)$ be a positive integer, and $\mathcal{R}_q := \mathbb{Z}_q[X]/(X^N + 1)$. Let $\chi(\lambda)$ be an error distribution on \mathcal{R}_q . For $s \in \mathcal{R}_q$, let \mathcal{O}_s be the oracle that samples $a \xleftarrow{\$} \mathcal{R}_q$, $x \leftarrow \chi$ and outputs $(a, as + x)$. Let \mathcal{O}_\S be the oracle that outputs $(a, b) \xleftarrow{\$} \mathcal{R}_q^2$. We say that the (N, q, χ) -RLWE assumption holds if for any PPT algorithm \mathcal{A} ,

$$\text{Adv}_{\mathcal{A}}^{\text{RLWE}_{(N,q,\chi)}}(\lambda) := \frac{1}{2} |\Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}_s}(1^\lambda)] - \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}_\S}(1^\lambda)]|$$

is negligible for a randomly chosen secret $s \xleftarrow{\$} \mathcal{R}_q$.

2.4 Searchable Symmetric Encryption

The following definition of the Searchable Symmetric Encryption (SSE) is a special case of the symmetric predicate-only encryption proposed by Shen et al [32].

- $\text{KeyGen}(1^\lambda)$: Given a security parameter λ , it outputs a public parameter pp and a secret key K .
- $\text{Enc}(\text{pp}, \text{K}, w)$: Given a public parameter pp , a secret key K , and a keyword w , it outputs a ciphertext C .
- $\text{Trapdoor}(\text{pp}, \text{K}, w)$: Given a public parameter pp , a secret key K , and a keyword w , it outputs a trapdoor T .
- $\text{Test}(\text{pp}, T, C)$: Given a public parameter pp , a trapdoor T , and a ciphertext C , it outputs a bit $b \in \{0, 1\}$.

For simplicity, we often omit the description of pp from the input of algorithms.

We say that the SSE scheme is correct if $\text{Test}(T, C)$ almost always outputs 1 when T and C have been generated from the same keyword. We say that the SSE scheme is consistent if $\text{Test}(T, C)$ almost always outputs 0 when T and C have been generated from different keywords. Our formulations below are based on the definitions for PEKS [1].

Definition 5 (Correctness). We say that an SSE scheme is correct if the minimum value of

$$\Pr[(\text{pp}, \text{K}) \leftarrow \text{KeyGen}(1^\lambda); C \leftarrow \text{Enc}(\text{K}, w); T \leftarrow \text{Trapdoor}(\text{K}, w) : \text{Test}(T, C) = 1]$$

with respect to w is overwhelming in λ .

Definition 6 (Consistency). We say that an SSE scheme is (computationally) consistent if for any PPT algorithm \mathcal{A} ,

$$\Pr[(\text{pp}, \text{K}) \leftarrow \text{KeyGen}(1^\lambda); (w, w') \leftarrow \mathcal{A}(1^\lambda, \text{pp}); C \leftarrow \text{Enc}(\text{K}, w); \\ T \leftarrow \text{Trapdoor}(\text{K}, w') : \text{Test}(T, C) = 0 \vee w = w']$$

is overwhelming in λ .

We introduce the notion of full security for SSE. As weaker properties, the definitions of ciphertext privacy and trapdoor privacy are given in Appendix B.

Definition 7 (Full security [32]). We say that an SSE scheme Σ is fully secure if for any PPT algorithm \mathcal{A} , the advantage

$$\text{Adv}_{\mathcal{A}, \Sigma}^{\text{Full}}(\lambda) := \left| \Pr[(\text{pp}, \text{K}) \leftarrow \text{KeyGen}(1^\lambda); b \xleftarrow{\$} \{0, 1\}; \\ b' \leftarrow \mathcal{A}^{\mathcal{O}_{\text{ct}}^b(\text{K}, \cdot, \cdot), \mathcal{O}_{\text{td}}^b(\text{K}, \cdot, \cdot)}(1^\lambda, \text{pp}) : b = b'] - \frac{1}{2} \right|$$

is negligible in λ . Here, $\mathcal{O}_{\text{ct}}^b(\text{K}, \cdot, \cdot)$ outputs $C \leftarrow \text{Enc}(\text{K}, w_b)$ given (w_0, w_1) as input. $\mathcal{O}_{\text{td}}^b(\text{K}, \cdot, \cdot)$ outputs $T \leftarrow \text{Trapdoor}(\text{K}, w_b)$ given (w_0, w_1) as input. The restriction is that for any input (x_0, x_1) to $\mathcal{O}_{\text{ct}}^b$ and any input (y_0, y_1) to $\mathcal{O}_{\text{td}}^b$, $(x_0, x_1) = (y_0, y_1)$ or “ $x_0 \neq y_0 \wedge x_1 \neq y_1$ ” holds.

Let $Q_{\text{ct}}(\lambda)$ (resp. $Q_{\text{td}}(\lambda)$) be an upper bound polynomial of the number of ciphertext (resp. trapdoor) queries. In the “bounded version” of full security, Q_{td} and Q_{ct} are determined before KeyGen phase. We write $\text{KeyGen}(1^\lambda, Q_{\text{td}}, Q_{\text{ct}})$ instead of $\text{KeyGen}(1^\lambda)$ when considering bounded full security. We write the advantage of the bounded full security game as $\text{Adv}_{\mathcal{A}, \Sigma}^{\text{BFull}}(\lambda)$.

3 Bounded fully secure SSE from PRF

In this section, we propose an SSE construction from PRF, which we write as SSE1. In addition, we show that SSE1 satisfies correctness, consistency, and bounded full security.

3.1 Construction

Let \mathbb{F} be a finite field of order $q(\lambda) = 2^{\omega(\log \lambda)}$. Let \mathcal{W} denote the keyword space. The construction of SSE1 is as follows:

- **KeyGen**($1^\lambda, Q_{\text{td}}, Q_{\text{ct}}$):
 1. Sets $d_1 \geq Q_{\text{td}}$ and $d_2 \geq Q_{\text{ct}}$.
 2. Chooses a pseudorandom function $F : \mathcal{K} \times \mathcal{W} \rightarrow \mathbb{F}^{d_1 \times d_2}$.
 3. Chooses a distribution η_i on \mathbb{F}^{d_i} for $i = 1, 2$. For each η_i , it is required that d_i independent samples from η_i become linearly independent with an overwhelming probability.
 4. Samples $\mathbf{K} \xleftarrow{\$} \mathcal{K}$.
 5. Outputs the public parameter F, η_1, η_2 and the secret key \mathbf{K} .
- **Enc**(\mathbf{K}, w): Samples $x \leftarrow \eta_2$ and outputs $(c_1, c_2) := (x, F_{\mathbf{K}}(w)x)$.
- **Trapdoor**(\mathbf{K}, w): Samples $y \leftarrow \eta_1$ and outputs $(t_1, t_2) := (y, F_{\mathbf{K}}(w)^\top y)$.
- **Test**($(t_1, t_2), (c_1, c_2)$): Outputs 1 if $c_1^\top t_2 = c_2^\top t_1$, otherwise outputs 0.

$\eta_i = \text{Unif}(\mathbb{F}^{d_i})$ satisfies the above condition since $q = 2^{\omega(\log \lambda)}$. Instead, different distributions may be used for more efficient constructions. For example, when q is prime and $\eta_i = \text{Unif}(\{0, 1\}^{d_i})$, the probability that d_i samples are linearly independent is $1 - O(1/p) - 2^{-\Omega(d_i)}$ [25, 27]. By setting $d_i = \omega(\log \lambda)$, this example satisfies the condition.

3.2 Properties

We show that SSE1 satisfies correctness, consistency, and bounded full security.

Correctness easily follows from the definition. Consistency follows from the pseudorandomness of F .

Theorem 1. SSE1 is consistent provided that F is pseudorandom.

Proof. Suppose that the consistency adversary outputs (w, w') where $w \neq w'$. For $(c_1, c_2) \leftarrow \text{Enc}(\mathbf{K}, w)$ and $(t_1, t_2) \leftarrow \text{Trapdoor}(\mathbf{K}, w')$,

$$c_1^\top t_2 - c_2^\top t_1 = x^\top (F_{\mathbf{K}}(w') - F_{\mathbf{K}}(w))^\top y$$

holds. By the pseudorandomness of F , it is sufficient to show that

$$\Pr \left[x \leftarrow \eta_2, y \leftarrow \eta_1; M \xleftarrow{\$} \mathbb{F}^{d_1 \times d_2} : x^\top M^\top y = 0 \right]$$

is negligible. Since $x, y \neq 0$ with an overwhelming probability, the probability that $x^\top M^\top y = 0$ is negligible. \square

Bounded full security also follows from the pseudorandomness of F . We prove the following theorem in the next subsection.

Theorem 2. SSE1 is bounded fully secure provided that F is pseudorandom. Namely, for any PPT adversary \mathcal{A} that makes a bounded number of queries, there exists a PPT algorithm \mathcal{B} such that

$$\text{Adv}_{\mathcal{A}, \text{SSE1}}^{\text{BFull}}(\lambda) \leq 2\text{Adv}_{\mathcal{B}, F}^{\text{PRF}}(\lambda) + \text{negl}(\lambda).$$

3.3 Proof of Theorem 2

Let $\mathcal{O}_{\mathbf{K}}^b$ denote the pair of the oracles $(\mathcal{O}_{\text{ct}}^b(\mathbf{K}, \cdot, \cdot), \mathcal{O}_{\text{td}}^b(\mathbf{K}, \cdot, \cdot))$. Consider the following pair of stateful oracles $\mathcal{O}^M = (\mathcal{O}_{\text{ct}}^M(\cdot, \cdot), \mathcal{O}_{\text{td}}^M(\cdot, \cdot))$.

1. List = \emptyset at the beginning.

2. $\mathcal{O}_{\text{ct}}^M(w_0, w_1)$ samples $x \leftarrow \eta_2$. If $((w_0, w_1), M) \in \text{List}$ for some $M \in \mathbb{F}^{d_1 \times d_2}$, it returns (x, Mx) . Otherwise, it samples $M \xleftarrow{\$} \mathbb{F}^{d_1 \times d_2}$ and returns (x, Mx) . Then, $((w_0, w_1), M)$ is appended to List .
3. $\mathcal{O}_{\text{td}}^M(w_0, w_1)$ samples $y \leftarrow \eta_1$. If $((w_0, w_1), M) \in \text{List}$ for some $M \in \mathbb{F}^{d_1 \times d_2}$, it returns $(y, M^\top y)$. Otherwise, it samples $M \xleftarrow{\$} \mathbb{F}^{d_1 \times d_2}$ and returns $(y, M^\top y)$. Then, $((w_0, w_1), M)$ is appended to List .

We prove that any PPT adversary cannot distinguish \mathcal{O}_{K}^0 (or equivalently \mathcal{O}_{K}^1) with \mathcal{O}^M with a non-negligible advantage provided that F is pseudorandom.

Now, we define \mathcal{O}^0 by replacing the pseudorandom function in \mathcal{O}_{K}^0 with the randomly chosen function. Let $\text{FList} = \emptyset$ at the beginning. When the oracles compute $F_{\text{K}}(w)$ for the first time, it chooses $M \xleftarrow{\$} \mathbb{F}^{d_1 \times d_2}$ and appends (w, M) to FList . After that, the oracles use M instead of $F_{\text{K}}(w)$. Then, for any PPT adversary \mathcal{A} , there exists a PPT algorithm \mathcal{B} such that

$$|\Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}_{\text{K}}^0}(1^\lambda)] - \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}^0}(1^\lambda)]| \leq 2\text{Adv}_{\mathcal{B}, F}^{\text{PRF}}(\lambda).$$

Here, the description of pp is omitted from the input.

Next, we define the intermediate oracles between \mathcal{O}^0 and \mathcal{O}^M . $\mathcal{O}^{0,i}$ ($i = 0, 1, \dots, Q$) sets $\text{List} = \emptyset$ and $\text{FList} = \emptyset$ at the beginning. Then, $\mathcal{O}_{\text{ct}}^{0,i}(\cdot, \cdot)$ behaves as follows:

1. It samples $x \leftarrow \eta_2$.
2. If $((w_0, w_1), M) \in \text{List}$ for some $M \in \mathbb{F}^{d_1 \times d_2}$, it returns (x, Mx) .
3. If $|\text{List}| < i$, it samples $M \xleftarrow{\$} \mathbb{F}^{d_1 \times d_2}$ and returns (x, Mx) . Then, $((w_0, w_1), M)$ is appended to List .
4. If $(w_0, M) \in \text{FList}$ for some $M \in \mathbb{F}^{d_1 \times d_2}$, it returns (x, Mx) .
5. Otherwise, it samples $M \xleftarrow{\$} \mathbb{F}^{d_1 \times d_2}$ and returns (x, Mx) . Then, (w_0, M) is appended to FList .

$\mathcal{O}_{\text{td}}^{0,i}(\cdot, \cdot)$ are defined in the same way. They sample $y \leftarrow \eta_1$ and return $(y, M^\top y)$. We have $\mathcal{O}^0 = \mathcal{O}^{0,0}$ and $\mathcal{O}^M = \mathcal{O}^{0,Q}$.

We prove that $\mathcal{O}^{0,i}$ and $\mathcal{O}^{0,i+1}$ are computationally indistinguishable. Regarding $\mathcal{O}^{0,i}$, let Q_{i+1} denote the query in which the first element is appended to FList . In the case of $\mathcal{O}^{0,i+1}$, this corresponds to the query in which the $(i+1)$ -th element is appended to List . The difference between these oracles is the behaviour when receiving Q_{i+1} . Let $(w_{0,i+1}, w_{1,i+1})$ be the pair of keywords corresponding to Q_{i+1} .

Let E denote the event that $(w_{0,i+1}, w_{1,i+1})$ has not been queried both to the ciphertext oracle and the trapdoor oracle throughout the game. In the case of $\neg E$, queries of the form $(w_{0,i+1}, \cdot)$ except $(w_{0,i+1}, w_{1,i+1})$ are not allowed. Thus, seen from the outside, $\mathcal{O}^{0,i}$ and $\mathcal{O}^{0,i+1}$ behave in the same way. It follows that for any PPT adversary \mathcal{A} ,

$$\Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}^{0,i}}(1^\lambda) \wedge \neg E] = \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}^{0,i+1}}(1^\lambda) \wedge \neg E].$$

E can be divided into two cases: E_1 and E_2 . E_1 denotes the case that $(w_{0,i+1}, w_{1,i+1})$ has been queried only to the ciphertext oracle. E_2 denotes the case that $(w_{0,i+1}, w_{1,i+1})$ has been queried only to the trapdoor oracle.

By symmetry, we consider the case of E_1 . Then, trapdoor queries of the form $(w_{0,i+1}, \cdot)$ are not allowed. Let $\mathcal{O}^{0,i+1/3}$ denote the oracle constructed by modifying $\mathcal{O}^{0,i}$ in the following way:

For Q_{i+1} or ciphertext queries of the form $\text{input} = (w_{0,i+1}, \cdot)$ after Q_{i+1} ,

1. If $(\text{input}, M) \in \text{List}$ for some $M \in \mathbb{F}^{d_1 \times d_2}$, it returns (x, Mx) where $x \leftarrow \eta_2$.
2. Otherwise, it returns $(x, u) \leftarrow \eta_2 \times \text{Unif}(\mathbb{F}^{d_1})$.

In addition, let $\mathcal{O}^{0,2/3}$ be the oracle defined by applying the above modification only to Q_{i+1} .

The responses to the ciphertext queries on $(w_{0,i+1}, \cdot)$ that are not in List are summarized in Table 2. Q_{i+1}^+ denotes the next such ciphertext query of Q_{i+1} . $\text{List}(\cdot)$ denotes M such that

$(\cdot, M) \in \text{List}$, and $\text{FList}(\cdot)$ denotes M such that $(\cdot, M) \in \text{FList}$. “Random” means that the oracle returns a sample from $\eta_2 \times \text{Unif}(\mathbb{F}^{d_1})$.

Table 2: Responses to ciphertext queries on $(w_{0,i+1}, \cdot)$ that are not in List

	Q_{i+1}	Q_{i+1}^+	After Q_{i+1}^+
\mathcal{O}^i	uses $\text{FList}(w_{0,i+1}) \xleftarrow{\$} \mathbb{F}^{d_1 \times d_2}$	uses $\text{FList}(w_{0,i+1})$	uses $\text{FList}(w_{0,i+1})$
$\mathcal{O}^{i+1/3}$	random	random	random
$\mathcal{O}^{i+2/3}$	random	uses $\text{FList}(w_{0,i+1}) \xleftarrow{\$} \mathbb{F}^{d_1 \times d_2}$	uses $\text{FList}(w_{0,i+1})$
\mathcal{O}^{i+1}	uses $\text{List}(\text{input}) \xleftarrow{\$} \mathbb{F}^{d_1 \times d_2}$	uses $\text{FList}(w_{0,i+1}) \xleftarrow{\$} \mathbb{F}^{d_1 \times d_2}$	uses $\text{FList}(w_{0,i+1})$

In the case of \mathcal{O}^i , let x_1, x_2, \dots, x_k ($k \leq Q_{\text{ct}}$) be the set of x sampled in such queries. Since x_1, x_2, \dots, x_k are linearly independent with an overwhelming probability,

$$([x_1|x_2|\dots|x_k], \text{FList}(w_{0,i+1}) \cdot [x_1|x_2|\dots|x_k])$$

is statistically close to the distribution $\eta_2^k \times \text{Unif}(\mathbb{F}^{d_1 \times k})$. Since $\text{FList}(w_{0,i+1})$ is independent of the other view of \mathcal{A} , we have

$$|\Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}^i}(1^\lambda) \wedge E_1] - \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}^{0,i+1/3}}(1^\lambda) \wedge E_1]| = \text{negl}(\lambda).$$

In the same way, we have

$$|\Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}^{0,i+(t-1)/3}}(1^\lambda) \wedge E_1] - \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}^{0,i+t/3}}(1^\lambda) \wedge E_1]| = \text{negl}(\lambda)$$

for each $t = 2, 3$. Since this property holds also for E_2 , we have

$$|\Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}^i}(1^\lambda)] - \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}^{0,i+1}}(1^\lambda)]| = \text{negl}(\lambda).$$

Therefore,

$$|\Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}^k}(1^\lambda)] - \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}^M}(1^\lambda)]| \leq 2\text{Adv}_{\mathcal{B},F}^{\text{PRF}}(\lambda) + \text{negl}(\lambda)$$

and this completes the proof.

4 Pairing-based SSE with Full Security

In this section, we propose an SSE construction based on pairings, which we write as SSE2. In addition, we show that this construction satisfies correctness, consistency, and full security. Full security follows from the U_k -MDDH assumption.

4.1 Construction

Let \mathcal{G} be a bilinear group generation algorithm and k be a positive integer. For simplicity of description, we assume that the group order p can be regarded as a deterministic function of the security parameter λ . Let \mathcal{W} denote the keyword space and $F : \mathcal{K} \times \mathcal{W} \rightarrow \mathbb{Z}_p^{k \times k}$ be a pseudorandom function. The construction of SSE2 is as follows:

- **KeyGen**(1^λ):
 1. Runs $\mathbb{G} := (p, G_1, G_2, G_T, g_1, g_2, e) \leftarrow \mathcal{G}(1^\lambda)$.
 2. Samples $K \xleftarrow{\$} \mathcal{K}$.
 3. Outputs the public parameter \mathbb{G} and the secret key K .
- **Enc**(K, w): Samples $x \xleftarrow{\$} \mathbb{Z}_p^k$ and outputs $(c_1, c_2) := ([x]_1, [F_K(w)x]_1)$.

- $\text{Trapdoor}(\mathbf{K}, w)$: Samples $y \xleftarrow{\$} \mathbb{Z}_p^k$ and outputs $(t_1, t_2) := ([y]_2, [F_{\mathbf{K}}(w)^\top y]_2)$.
- $\text{Test}((t_1, t_2), (c_1, c_2))$: Outputs 1 if $e(c_1, t_2) = e(c_2, t_1)$, otherwise outputs 0.

k determines the underlying assumption of SSE2. The case of $k = 1$ is the most efficient. The underlying assumption U_1 -MDDH is the strongest, but still it is weaker than the SXDH assumption.

4.2 Properties

It is easy to see that SSE2 satisfies correctness. In fact, if a ciphertext (c_1, c_2) and a trapdoor (t_1, t_2) are associated with the same keyword w ,

$$e(c_1, t_2) = [x^\top F_{\mathbf{K}}(w)^\top y]_T = e(c_2, t_1).$$

Consistency follows from the pseudorandomness of F .

Theorem 3. SSE2 is consistent provided that F is pseudorandom.

Proof. Suppose that the consistency adversary outputs (w, w') where $w \neq w'$. For $(c_1, c_2) \leftarrow \text{Enc}(\mathbf{K}, w)$ and $(t_1, t_2) \leftarrow \text{Trapdoor}(\mathbf{K}, w')$,

$$e(c_1, t_2)/e(c_2, t_1) = [x^\top (F_{\mathbf{K}}(w') - F_{\mathbf{K}}(w))^\top y]_T$$

holds. By the pseudorandomness of F , it is sufficient to show that

$$\Pr \left[x, y \xleftarrow{\$} \mathbb{Z}_p^k; M \xleftarrow{\$} \mathbb{Z}_p^{k \times k} : x^\top M y = 0 \right]$$

is negligible. Since $x, y \neq 0$ with an overwhelming probability, the probability that $x^\top M y = 0$ is negligible. \square

Full security follows from the pseudorandomness of F and the U_k -MDDH assumption. We prove the following theorem in the next subsection.

Theorem 4. SSE2 is fully secure provided that F is pseudorandom and the U_k -MDDH assumption holds. Namely, for any PPT adversary \mathcal{A} that makes at most Q queries, there exist PPT algorithms $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3$ such that

$$\text{Adv}_{\mathcal{A}, \text{SSE2}}^{\text{Full}}(\lambda) \leq 2\text{Adv}_{\mathcal{B}_1, F}^{\text{PRF}}(\lambda) + 6Q \left(\text{Adv}_{1, \mathcal{B}_2}^{U_k\text{-MDDH}}(\lambda) + \text{Adv}_{2, \mathcal{B}_3}^{U_k\text{-MDDH}}(\lambda) \right) + \text{negl}(\lambda).$$

4.3 Proof of Theorem 4

The proof flow is the same as that of Theorem 2. Let Q be an upper bound of the number of queries. We define the oracles $\mathcal{O}_{\mathbf{K}}^0, \mathcal{O}^{0,t/3}(t = 0, 1, \dots, 3Q)$ and the events E, E_1, E_2 in the same way.

First, there exists a PPT algorithm \mathcal{B}_1 such that

$$\left| \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}_{\mathbf{K}}^0}(1^\lambda)] - \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}}(1^\lambda)] \right| \leq 2\text{Adv}_{\mathcal{B}_1, F}^{\text{PRF}}(\lambda).$$

Then, we evaluate the computational difference between $\mathcal{O}^{0,i}$ and $\mathcal{O}^{0,i+1}$ for $i = 0, 1, \dots, Q - 1$. We have

$$\Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}^{0,i}}(1^\lambda) \wedge \neg E] = \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}^{0,i+1}}(1^\lambda) \wedge \neg E],$$

so we consider the case of E_1 by symmetry. By treating $\text{FList}(w_{0,i+1})$ as the k -instance MDDH secret, it follows that there exists a PPT algorithm $\mathcal{B}_{2,1}$ such that

$$\left| \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}^{0,i}}(1^\lambda) \wedge E_1] - \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}^{0,i+1/3}}(1^\lambda) \wedge E_1] \right| \leq 2\text{Adv}_{1, \mathcal{B}_{2,1}}^{U_k\text{-MDDH}_k}(\lambda).$$

Similarly, there exist PPT algorithms $\mathcal{B}_{2,t}$ ($t = 2, 3$) such that

$$\left| \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}^{0,i+(t-1)/3}}(1^\lambda) \wedge E_1] - \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}^{0,i+t/3}}(1^\lambda) \wedge E_1] \right| \leq 2\text{Adv}_{1,\mathcal{B}_{2,t}}^{U_k\text{-MDDH}_k}(\lambda).$$

Thus, for some PPT algorithm \mathcal{B}_2 ,

$$\left| \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}^{0,i}}(1^\lambda) \wedge E_1] - \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}^{0,i+1}}(1^\lambda) \wedge E_1] \right| \leq 6\text{Adv}_{1,\mathcal{B}_2}^{U_k\text{-MDDH}_k}(\lambda)$$

holds. By symmetry, there exists a PPT algorithm \mathcal{B}_3 such that

$$\left| \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}^{0,i}}(1^\lambda) \wedge E_2] - \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}^{0,i+1}}(1^\lambda) \wedge E_2] \right| \leq 6\text{Adv}_{2,\mathcal{B}_3}^{U_k\text{-MDDH}_k}(\lambda).$$

By Lemma 2, the inequality in Theorem 4 holds.

5 Lattice-based SSE with Full Security

In this section, we propose an SSE construction with full security based on lattices. We describe an LWE-based construction, which we write as SSE3-L. Then, we show that the RLWE-based scheme SSE3-R can be constructed in the same way.

5.1 Revisiting Coset Sampleability

For constructing SSE, we use an n -coset sampleable distribution η on \mathbb{Z}_q^k that outputs short vectors. Boneh et al. [8] proposed two examples of such distributions.

In the first example, η is the uniform distribution on $\{0, 1\}^k$ where $k = n \lceil \log q \rceil$. q is restricted to a power of 2 if q is polynomially bounded. In the second example, η is a discrete Gaussian distribution with the deviation σ where $k = \Omega(n \log q)$ and $\sigma = \Omega(\sqrt{n \log q})$.

In both of these cases, $k \geq n \log q$ holds. We propose a generalized construction of the first example. For an integer $d(\lambda) \geq 2$, we define n -coset sampleable distributions η_d with the dimension $k = dn$ as follows. Let $q_0 = \lceil q^{1/d} \rceil$.

- **MatrixGen**(1^λ): It outputs

$$M = [q_0^{d-1} I_n | q_0^{d-2} I_n | \cdots | I_n] \in \mathbb{Z}_q^{n \times k}$$

and auxiliary data $T = \emptyset$.

- **SamplePre**($z \in \mathbb{Z}_q^n, M, T$): It treats z as $z \in (-q/2, q/2]^n$ and sets $y_{d-1} = \lfloor z/q_0^{d-1} \rfloor$, $z_{d-1} = z - q_0^{d-1} y_{d-1}$. Then, it inductively sets $y_i = \lfloor z_{i+1}/q_0^i \rfloor$ and $z_i = z_{i+1} - q_0^i y_i$ for $i = d-2, \dots, 0$. Finally, it outputs $[y_{d-1}; \dots; y_0] \in [-q_0/2, q_0/2]^k$.

Next, we try to apply the non-uniform setting to the RLWE assumption. However, the non-uniform RLWE assumption does not hold if η outputs sufficiently short vectors. This is because if $(a_1, a_1 s + x_1)$ and $(a_2, a_2 s + x_2)$ are non-uniform RLWE instances, $a_1(a_2 s + x_2) - a_2(a_1 s + x_1)$ is relatively short.

To construct non-uniform assumptions that are as weak as RLWE, we introduce the Non-uniform Module-LWE (NMLWE) assumption.

Definition 8. Let $k(\lambda)$ be a positive integer and $\eta(\lambda)$ be a distribution on \mathcal{R}_q^k . Let $\chi(\lambda)$ be an error distribution on \mathcal{R}_q . For $s \in \mathcal{R}_q^k$, let \mathcal{O}_s^η be the oracle that samples $a \leftarrow \eta$, $x \leftarrow \chi$ and outputs $(a, a^\top s + x)$. Let \mathcal{O}_s^η be the oracle that outputs $(a, b) \leftarrow \eta \times \text{Unif}(\mathcal{R}_q)$. We say that the (N, q, k, χ, η) -NMLWE assumption holds if for any PPT algorithm \mathcal{A} ,

$$\text{Adv}_{\mathcal{A}}^{\text{NMLWE}_{(N,q,k,\chi,\eta)}}(\lambda) := \frac{1}{2} \left| \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}_s^\eta}(1^\lambda)] - \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}_s^\eta}(1^\lambda)] \right|$$

is negligible for a randomly chosen secret $s \xleftarrow{\$} \mathcal{R}_q$.

As in the case of NLWE, the NMLWE problem is as hard as the standard MLWE problem when η achieves some properties. For simplicity of description, we only deal with the problems that are as hard as RLWE. We define \mathcal{R}_q -coset sampleability as follows.

Definition 9. A distribution η on \mathcal{R}_q^k is called \mathcal{R}_q -coset sampleable if there exist two PPT algorithms **MatrixGen** and **SamplePre** such that:

- **MatrixGen**(1^λ) outputs $v \in \mathcal{R}_q^k$ and auxiliary data T .
- **SamplePre**($z \in \mathcal{R}_q, v, T$) outputs $y \in \mathcal{R}_q^k$ such that $v^\top y = z$. In addition, if z is sampled uniformly at random, the distribution of y is η .

In a similar way to the case of n -coset sampleability, for an integer $k(\lambda) \geq 2$, we can construct \mathcal{R}_q -sampleable distributions $\eta_k^{\mathcal{R}}$. Let $q_0 = \lceil q^{1/k} \rceil$.

- **MatrixGen**(1^λ): It outputs

$$v = (q_0^{k-1}, q_0^{k-2}, \dots, 1)^\top \in \mathcal{R}_q^k$$

and auxiliary data $T = \emptyset$.

- **SamplePre**($z \in \mathcal{R}_q, v, T$): It treats coefficients of z as elements in $(-q/2, q/2]$, and sets $y_{k-1} = \lfloor z/q_0^{k-1} \rfloor$, $z_{k-1} = z - q_0^{k-1}y_{k-1}$. Then, it inductively defines $y_i = \lfloor z_{i+1}/q_0^i \rfloor$ and $z_i = z_{i+1} - q_0^i y_i$ for $i = k-2, \dots, 0$. Finally, it outputs $(y_{k-1}, \dots, y_0)^\top \in \mathcal{R}_q^k$ where each coefficient is in $[-q_0/2, q_0/2]$.

For \mathcal{R}_q -coset sampleable distributions, the similar reduction to Lemma 3 holds. The following lemma can be proved in the same way as Lemma 3.

Lemma 4. If η is \mathcal{R}_q -coset sampleable, the (N, q, k, χ, η) -NMLWE problem is at least as hard as the (N, q, χ) -RLWE problem. Namely, for any PPT adversary \mathcal{A} , there exists a PPT algorithm \mathcal{B} such that

$$\text{Adv}_{\mathcal{A}}^{\text{NMLWE}_{N,m,q,k,\chi,\eta}}(\lambda) \leq \text{Adv}_{\mathcal{B}}^{\text{RLWE}_{N,m,q,\chi}}(\lambda).$$

Proof. Let \mathcal{A} be any PPT solver of NMLWE. We construct an RLWE solver \mathcal{B} using \mathcal{A} . First, \mathcal{B} runs $(v, T) \leftarrow \text{MatrixGen}(1^\lambda)$ and samples $r \xleftarrow{\$} \mathcal{R}_q^k$. Suppose that \mathcal{B} has received an RLWE instance $(a, b) \in \mathcal{R}_q^2$. \mathcal{B} runs $a' \leftarrow \text{SamplePre}(a, v, T)$ and sets $b' = b + a'^\top r$. Finally, \mathcal{B} sends (a', b') to \mathcal{A} as an NMLWE instance.

If $b = as + x$ where s is the secret and $x \leftarrow \chi$,

$$b' = (a'^\top v)s + a'^\top r + x = a'^\top (sv + r) + x.$$

Note that $sv + r$ is fixed for all instances. If (a, b) is uniformly random, (a', b') is distributed as $\eta \times \text{Unif}(\mathcal{R}_q)$. Thus, (a', b') is a proper NMLWE instance. \square

5.2 LWE-based Construction

Let $m(\lambda), q(\lambda), \kappa_{\text{td}}(\lambda), \kappa_{\text{ct}}(\lambda)$ be positive integers such that $\kappa_{\text{td}}\kappa_{\text{ct}} = \omega(\log \lambda)$. Let $\chi(\lambda)$ be an error distribution on \mathbb{Z}_q and $\eta(\lambda)$ be a distribution that outputs short vectors. Let \mathcal{W} denote the keyword space and $F: \mathcal{K} \times \mathcal{W} \rightarrow \mathbb{Z}_q^{m \times m}$ be a pseudorandom function. The construction of SSE3-L is as follows:

- **KeyGen**(1^λ):
 1. Samples $K \xleftarrow{\$} \mathcal{K}$.
 2. Outputs the secret key K .
- **Enc**(K, w):
 1. Samples $c_j \leftarrow \eta$ and $x_j \leftarrow \chi^m$ for $j = 1, \dots, \kappa_{\text{ct}}$.
 2. Computes $d_j = F_K(w)c_j + x_j$ for each j .

3. Outputs $\{c_j, d_j\}_{j=1}^{\kappa_{\text{ct}}}$.
- **Trapdoor**(K, w):
 1. Samples $t_i \leftarrow \eta$ and $y_i \leftarrow \chi^m$ for $i = 1, \dots, \kappa_{\text{td}}$.
 2. Computes $u_i = F_K(w)^\top t_i + y_i$ for each i .
 3. Outputs $\{t_i, u_i\}_{i=1}^{\kappa_{\text{td}}}$.
 - **Test**($\{t_i, u_i\}_{i=1}^{\kappa_{\text{td}}}, \{c_j, d_j\}_{j=1}^{\kappa_{\text{ct}}}$):
 1. Computes $\alpha_{i,j} = c_j^\top u_i - d_j^\top t_i \in \mathbb{Z}_q$ for each $(i, j) \in [1, \kappa_{\text{td}}] \times [1, \kappa_{\text{ct}}]$.
 2. Treats $\alpha_{i,j}$ as an integer in $(-q/2, q/2]$. Outputs 1 if $-[q/4] < \alpha_{i,j} < [q/4]$ for every (i, j) , otherwise outputs 0.

5.3 Properties

First, we consider correctness of SSE3-L. Suppose that $\{t_i, u_i\}_{i=1}^{\kappa_{\text{td}}}$ and $\{c_j, d_j\}_{j=1}^{\kappa_{\text{ct}}}$ are generated from the same keyword. Then, for any (i, j) , $|\alpha_{i,j}| = |c_j^\top y_i - d_j^\top t_i|$ holds. This value is relatively small, so this scheme can achieve correctness in appropriate parameter settings.

For example, consider the case that $\eta = \eta_2$ and $m = 2n$. In order that the worst-case to average-case reduction works, let χ be the discrete Gaussian with the deviation $\sigma = \Theta(\sqrt{n})$. Correctness holds if

$$q = \sqrt{2n}q_0\sigma \cdot \omega(\sqrt{\log n}).$$

Thus, when $q = \omega(n^2 \log n)$, our scheme is correct. As another example, when $\eta = \eta_{\lceil \log q \rceil}$, smaller $q = \omega(n \log n)$ can be used. However, since m is large in this example, η_2 seems to be more attractive in terms of efficiency.

Consistency follows from the pseudorandomness of F and the NLWE assumption.

Theorem 5. SSE3-L is consistent provided that F is pseudorandom and the (m, q, χ, η) -NLWE assumption holds.

Proof. By the pseudorandomness of F and the (m, q, χ, η) -NLWE assumption, it is sufficient to show that

$$\begin{aligned} & \Pr[c_j \leftarrow \eta, d_j \xleftarrow{\$} \mathbb{Z}_q^m \text{ for } j = 1, \dots, \kappa_{\text{ct}}; \\ & \quad t_i \leftarrow \eta, u_i \xleftarrow{\$} \mathbb{Z}_q^m \text{ for } i = 1, \dots, \kappa_{\text{td}}; \\ & \quad 1 \leftarrow \text{Test}(\{t_i, u_i\}_{i=1}^{\kappa_{\text{td}}}, \{c_j, d_j\}_{j=1}^{\kappa_{\text{ct}}})] \end{aligned}$$

is negligible. Let $z_{i,j} \leftarrow \chi$ for each $(i, j) \in [1, \kappa_{\text{td}}] \times [1, \kappa_{\text{ct}}]$. Then, by the (m, q, χ, η) -NLWE assumption and the hybrid argument, $\{c_j^\top u_i + z_{i,j}\}_{i,j}$ is computationally indistinguishable from random. Moreover, since d_j, u_i are independent of $c_j, u_i, z_{i,j}$, $\{c_j^\top u_i - d_j^\top t_i + z_{i,j}\}_{i,j}$ is also computationally indistinguishable from random. Suppose that the output of χ is in $[-E, E]$ with an overwhelming probability. If **Test** outputs 1, $|c_j^\top u_i - d_j^\top t_i + z_{i,j}| < q/4 + E$ holds for any (i, j) . Thus, this probability is bounded above by $(1/2 + 2E/q)^{\kappa_{\text{td}}\kappa_{\text{ct}}} + \text{negl}(\lambda)$. When κ_{td} and κ_{ct} are large enough that $\kappa_{\text{td}}\kappa_{\text{ct}} = \omega(\log \lambda)$, SSE3-L is consistent. \square

Full security also follows from the pseudorandomness of F and the NLWE assumption. We prove the following theorem in the next subsection.

Theorem 6. SSE3-L is fully secure provided that F is pseudorandom and the (m, q, χ, η) -NLWE assumption holds. Namely, for any PPT adversary \mathcal{A} that makes at most Q queries, there exist PPT algorithms $\mathcal{B}_1, \mathcal{B}_2$ such that

$$\text{Adv}_{\mathcal{A}, \text{SSE3-L}}^{\text{Full}}(\lambda) \leq 2\text{Adv}_{\mathcal{B}_1, F}^{\text{PRF}}(\lambda) + 12mQ \cdot \text{Adv}_{\mathcal{B}_2}^{\text{NLWE}_{m, \kappa Q, q, \chi, \eta}}(\lambda)$$

where $\kappa := \max\{\kappa_{\text{td}}, \kappa_{\text{ct}}\}$.

5.4 Proof of Theorem 6

The proof flow is the same as that of Theorem 2. Let Q be an upper bound of the number of queries. We define the oracles $\mathcal{O}_K^0, \mathcal{O}^{0,t/3}(t = 0, 1, \dots, 3Q)$ and the events E, E_1, E_2 in the same way.

First, there exists a PPT algorithm \mathcal{B}_1 such that

$$|\Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}_K^0}(1^\lambda)] - \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}^0}(1^\lambda)]| \leq 2\text{Adv}_{\mathcal{B}_1, F}^{\text{PRF}}(\lambda).$$

Then, we evaluate the computational difference between $\mathcal{O}^{0,i}$ and $\mathcal{O}^{0,i+1}$ for $i = 0, 1, \dots, Q-1$. We have

$$\Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}^{0,i}}(1^\lambda) \wedge \neg E] = \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}^{0,i+1}}(1^\lambda) \wedge \neg E],$$

so we consider the case of E_1 by symmetry. By treating $\text{FList}(w_{0,i+1})$ as the m -instance NLWE secret and using the hybrid argument, it follows that there exists a PPT algorithm $\mathcal{B}_{2,1}$ such that

$$|\Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}^{0,i}}(1^\lambda) \wedge E_1] - \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}^{0,i+1/3}}(1^\lambda) \wedge E_1]| \leq 2m\text{Adv}_{\mathcal{B}_{2,1}}^{\text{NLWE}_{m, \kappa_{\text{ct}} Q, q, \chi, \eta}}(\lambda).$$

Similarly, there exist PPT algorithms $\mathcal{B}_{2,t}$ ($t = 2, 3$) such that

$$|\Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}^{0,i+(t-1)/3}}(1^\lambda) \wedge E_1] - \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}^{0,i+t/3}}(1^\lambda) \wedge E_1]| \leq 2m\text{Adv}_{\mathcal{B}_{2,t}}^{\text{NLWE}_{m, \kappa_{\text{ct}} Q, q, \chi, \eta}}(\lambda).$$

Thus, for some PPT algorithm \mathcal{B}_2 ,

$$|\Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}^{0,i}}(1^\lambda) \wedge E_1] - \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}^{0,i+1}}(1^\lambda) \wedge E_1]| \leq 6m\text{Adv}_{\mathcal{B}_2}^{\text{NLWE}_{m, \kappa_{\text{ct}} Q, q, \chi, \eta}}(\lambda)$$

holds. By symmetry, there exists a PPT algorithm \mathcal{B}'_2 such that

$$|\Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}^{0,i}}(1^\lambda) \wedge E_2] - \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}^{0,i+1}}(1^\lambda) \wedge E_2]| \leq 6m\text{Adv}_{\mathcal{B}'_2}^{\text{NLWE}_{m, \kappa_{\text{id}} Q, q, \chi, \eta}}(\lambda).$$

Therefore, the inequality in Theorem 6 holds.

5.5 RLWE-based Construction

We can construct the RLWE-based variant SSE3-R of SSE3-L. Let $k(\lambda)$ be a positive integer (e.g. $k = 2$). Let $\chi(\lambda)$ be an error distribution on \mathcal{R}_q and $\eta(\lambda)$ be a distribution on \mathcal{R}_q^k that outputs short vectors. Let \mathcal{W} denote the keyword space and $F : \mathcal{K} \times \mathcal{W} \rightarrow \mathcal{R}_q^{k \times k}$ be a pseudorandom function. The construction of SSE3-R is as follows:

- **KeyGen**(1^λ):
 1. Samples $K \xleftarrow{\$} \mathcal{K}$.
 2. Outputs the secret key K .
- **Enc**(K, w):
 1. Samples $c \leftarrow \eta$ and $x \leftarrow \chi^k$.
 2. Computes $d = F_K(w)c + x$.
 3. Outputs (c, d) .
- **Trapdoor**(K, w):
 1. Samples $t \leftarrow \eta$ and $y \leftarrow \chi^k$.
 2. Computes $u = F_K(w)^\top t + y$.
 3. Outputs (t, u) .
- **Test**($((t, u), (c, d))$):
 1. Computes $\alpha = c^\top u - d^\top t \in \mathcal{R}_q$.

2. Treats the coefficients of α as integers in $(-q/2, q/2]$. Outputs 1 if they are in $(- \lfloor q/4 \rfloor, \lfloor q/4 \rfloor)$ for $\Theta(\lambda)$ coefficients in a fixed position, otherwise outputs 0. Note that $N > \lambda$ holds for practical RLWE parameters.

Correctness, consistency, and full security can be proved in the same way as SSE3-L.

If (c, d) and (t, u) are generated from the same keyword, $c^\top u - d^\top t = c^\top y - x^\top t$ has relatively small coefficients. For example, when $\eta = \eta_2^{\mathcal{R}}$ and χ is the discrete Gaussian distribution with the deviation σ , correctness holds if

$$q = \sqrt{2N}q_0\sigma \cdot \omega(\sqrt{\log N}).$$

Consistency follows from the pseudorandomness of F and the (N, q, k, χ, η) -NMLWE assumption. The proof is almost the same as that of Theorem 5. By the NMLWE assumption, $c^\top u - d^\top t + z$ is computationally indistinguishable from random where $(c, d), (t, u) \leftarrow \eta \times \text{Unif}(\mathcal{R}_q)$ and $z \leftarrow \chi$.

Full security follows also from the pseudorandomness of F and the (N, q, k, χ, η) -NMLWE assumption. For any PPT adversary \mathcal{A} of SSE3-R that makes at most Q queries, there exist PPT algorithms $\mathcal{B}_1, \mathcal{B}_2$ such that

$$\text{Adv}_{\mathcal{A}, \text{SSE3-R}}^{\text{Full}}(\lambda) \leq 2\text{Adv}_{\mathcal{B}_1, F}^{\text{PRF}}(\lambda) + 12kQ \cdot \text{Adv}_{\mathcal{B}_2}^{\text{NMLWE}_{N, Q, q, k, \chi, \eta}}(\lambda).$$

6 Simplifying and Verifying Cheng et al. Scheme

In this section, we analyze the Cheng et al. PAEKS scheme [12]. This scheme is essentially the combination of NIKE and SSE, so we focus on the SSE part. It has been unknown whether the SSE part is secure since Li and Boyen [21] pointed out that the original proof is incomplete. We simplify and generalize this scheme. Then, we show that full security can be proved from the bilateral MDDH assumption.

6.1 Simplification of Cheng et al. PAEKS Scheme

First, we describe the algorithms of the Cheng et al. PAEKS [12] scheme. The syntax of PAEKS is given in Appendix A.

- **Setup**(1^λ):
 1. Runs the symmetric bilinear group generation algorithm $\mathbb{G} = (p, G, G_T, g, e) \leftarrow \mathcal{G}(1^\lambda)$.
 2. Chooses hash functions $H : \{0, 1\}^* \rightarrow G$ and $\hat{H} : G \rightarrow \{0, 1\}^\ell$.
 3. Samples $d \xleftarrow{\$} \mathbb{Z}_p$ and sets $h = g^d$.
 4. Outputs the public parameter $\text{pp} = (\mathbb{G}, h, H, \hat{H})$.
- **KeyGen_S**(pp): Samples $y \xleftarrow{\$} \mathbb{Z}_p$ and outputs $(\text{pk}_S, \text{sk}_S) = (g^y, y)$.
- **KeyGen_R**(pp): Samples $x \xleftarrow{\$} \mathbb{Z}_p$ and outputs $(\text{pk}_R, \text{sk}_R) = (g^x, x)$.
- **Enc**($\text{pk}_R, \text{pk}_S, \text{sk}_S, w$): Samples $r_1, r_2 \xleftarrow{\$} \mathbb{Z}_p$ and outputs the ciphertext (C_1, C_2, C_3, C_4) where

$$C_1 = H(\hat{H}(g^{xy}), w, \text{pk}_R, \text{pk}_S)^{r_1} \cdot h^{yr_2}, \quad C_2 = g^{xr_1}, \quad C_3 = h^{r_2}, \quad C_4 = g^{r_1}.$$

- **Trapdoor**($\text{pk}_R, \text{pk}_S, \text{sk}_R, w$): Samples $s_1, s_2 \xleftarrow{\$} \mathbb{Z}_p$ and outputs the trapdoor (T_1, T_2, T_3, T_4) where
$$T_1 = H(\hat{H}(g^{xy}), w, \text{pk}_R, \text{pk}_S)^{s_1} \cdot h^{xs_2}, \quad T_2 = g^{ys_1}, \quad T_3 = h^{s_2}, \quad T_4 = g^{s_1}.$$
- **Test**($(T_1, T_2, T_3, T_4), (C_1, C_2, C_3, C_4)$): Outputs 1 if $e(C_1, T_4) \cdot e(C_2, T_3) = e(T_1, C_4) \cdot e(T_2, C_3)$, otherwise outputs 0.

In this scheme, H and \hat{H} are used to reduce the security to the Computational Oracle Diffie-Hellman assumption. However, since this reduction is not correct as stated in [21], we replace this part with a pseudorandom function in order to remove the random oracle. Moreover, it is not necessary to use h since the distribution of (h^{yr_2}, h^{r_2}) (resp. (h^{xs_2}, h^{s_2})) is identical to that of (g^{yr_2}, g^{r_2}) (resp. (g^{xs_2}, g^{s_2})).

Now, we propose a simplified version of this scheme, which we write as CSSE. We describe an SSE scheme since SSE can be easily converted to PAEKS by using NIKE to compute the shared secret key [21].

Let $F : \mathcal{K} \times \mathcal{W} \rightarrow \mathbb{Z}_p$ be a pseudorandom function. For generality, we use the bilinear map $e : G_1 \times G_2 \rightarrow G_T$ where $G_1 = G_2$ does not necessarily hold.

- **KeyGen**(1^λ):

1. Runs the bilinear group generation algorithm $\mathbb{G} = (p, G_1, G_2, G_T, g_1, g_2, e) \leftarrow \mathcal{G}(1^\lambda)$.
2. Samples $x, y \xleftarrow{\$} \mathbb{Z}_p$.
3. Samples $K \xleftarrow{\$} \mathcal{K}$.
4. Outputs the public parameter \mathbb{G} and the secret key $K' = (K, x, y)$.

- **Enc**(K', w): Samples $r_1, r_2 \xleftarrow{\$} \mathbb{Z}_p$ and outputs a ciphertext (C_1, C_2, C_3, C_4) where

$$C_1 = g_1^{F_K(w)r_1 + yr_2}, C_2 = g_1^{xr_1}, C_3 = g_1^{r_2}, C_4 = g_1^{r_1}.$$

- **Trapdoor**(K', w): Samples $s_1, s_2 \xleftarrow{\$} \mathbb{Z}_p$ and outputs a trapdoor (T_1, T_2, T_3, T_4) where

$$T_1 = g_2^{F_K(w)s_1 + xs_2}, T_2 = g_2^{ys_1}, T_3 = g_2^{s_2}, T_4 = g_2^{s_1}.$$

- **Test**($(T_1, T_2, T_3, T_4), (C_1, C_2, C_3, C_4)$): Outputs 1 if $e(C_1, T_4) \cdot e(C_2, T_3) = e(C_4, T_1) \cdot e(C_3, T_2)$, otherwise outputs 0.

Note that instead of setting (x, y) as a part of the secret key, we can set $(g_1^x, g_1^y, g_2^x, g_2^y)$ as a part of the public parameter.

Let $G_{K'} : \mathcal{W} \rightarrow \mathbb{Z}_p^{2 \times 2}$ denote the function defined by

$$G_{K'}(w) := \begin{pmatrix} F_K(w) & y \\ x & 0 \end{pmatrix}.$$

Using $G_{K'}$, ciphertexts and the trapdoors can be written in the similar form to SSE2 based on the U_2 -MDDH assumption.

- **Ciphertext**: $([r]_1, [G_{K'}(w)r]_1)$ where $r \xleftarrow{\$} \mathbb{Z}_p^2$.
- **Trapdoor**: $([s]_2, [G_{K'}(w)^\top s]_2)$ where $s \xleftarrow{\$} \mathbb{Z}_p^2$.

Compared to U_2 -MDDH-based SSE2, the ciphertext size and the trapdoor size is the same. The advantage is the improved computational cost due to the smaller range of F . The disadvantage is that the security is reduced to a stronger assumption. We describe the details in the next subsection.

6.2 Security Analysis

Correctness and consistency can be easily proved in the same way as SSE2. We analyze the security of CSSE. In the (bil-) U_2 -MDDH assumption, the matrix distribution is $\text{Unif}(\mathbb{Z}_p^{3 \times 2})$. In the (bil-)DLin assumption, the matrix distribution is

$$A = \begin{pmatrix} a_1 & 0 \\ 0 & a_2 \\ 1 & 1 \end{pmatrix} \text{ where } a_1, a_2 \xleftarrow{\$} \mathbb{Z}_p.$$

Full security of CSSE is reduced to the bil-MDDH assumption such that the matrix distribution is

$$\begin{pmatrix} b_1 & b_2 \\ b_3 & b_4 \\ 0 & 1 \end{pmatrix} \text{ where } b_1, b_2, b_3, b_4 \stackrel{\$}{\leftarrow} \mathbb{Z}_p,$$

which we write as V_2 in this paper. We show that the DLin assumption implies the V_2 -MDDH assumption using the MDDH reduction method [17]. Let L, R be the matrices defined as

$$L = \begin{pmatrix} 1 & 0 & l_1 \\ 0 & 1 & l_2 \\ 0 & 0 & 1 \end{pmatrix} \text{ where } l_1, l_2 \stackrel{\$}{\leftarrow} \mathbb{Z}_p, \quad R = \begin{pmatrix} 1 & 0 \\ -1 & 1 \end{pmatrix}.$$

For a bil-DLin instance $([A]_1, [z]_1, [A]_2, [z]_2)$, the distribution of LAR is V_2 since

$$LAR = \begin{pmatrix} a_1 & l_1 \\ -a_2 & a_2 + l_2 \\ 0 & 1 \end{pmatrix}.$$

By the invertibility of R , $([LAR]_1, [LAz]_1, [LAR]_2, [LAz]_2)$ is a bil- V_2 -MDDH instance. Thus, the (bil-)DLin assumption implies the (bil-) V_2 -MDDH assumption.

We give the security proof in the next subsection.

Theorem 7. The modified Cheng et al. SSE scheme CSSE is fully secure provided that F is pseudorandom and the bil- V_2 -MDDH assumption holds.

This scheme can be generalized to the constructions based on the bilateral k -Lin assumption by setting $F : \mathcal{K} \times \mathcal{W} \rightarrow \mathbb{Z}_p^{(k-1) \times (k-1)}$, $x, y \stackrel{\$}{\leftarrow} \mathbb{Z}_p^k$, and

$$G_{\mathcal{K}'}(w) := \begin{pmatrix} F_{\mathcal{K}}(w) & y \\ x^\top & 0 \end{pmatrix}.$$

6.3 Proof of Theorem 7

Similarly to U_k -MDDH, the following lemma holds. The proof is the same as that of Lemma 1.

Lemma 5. For $r \in \mathbb{Z}_p^2$, let \mathcal{O}^r be an oracle that samples $a \stackrel{\$}{\leftarrow} \mathbb{Z}_p^2$ and outputs $([a]_1, [a^\top r]_1, [a]_2, [a^\top r]_2)$.

Let \mathcal{O} be an oracle that samples $a \stackrel{\$}{\leftarrow} \mathbb{Z}_p^2$, $u \stackrel{\$}{\leftarrow} \mathbb{Z}_p$ and outputs $([a]_1, [u]_1, [a]_2, [u]_2)$. If the bil- V_2 -MDDH assumption holds, for any PPT algorithm \mathcal{A} ,

$$\text{Adv}_{\mathcal{A}}^{\text{bil-}V_2\text{-MDDH}}(\lambda) := \frac{1}{2} |\Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}^r}(1^\lambda, [y]_1, [y]_2)] - \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}}(1^\lambda, [y]_1, [y]_2)]|$$

is negligible for a randomly chosen secret $r = (x, y)^\top \stackrel{\$}{\leftarrow} \mathbb{Z}_p^2$.

Proof. Using the bil- V_2 -MDDH instance $([A]_1, [s]_1, [A]_2, [s]_2)$, the view of \mathcal{A} in this lemma can be simulated. $[y]_1$ and $[y]_2$ are set as the third components of $[s]_1$ and $[s]_2$ respectively. The oracle simulator samples $t \stackrel{\$}{\leftarrow} \mathbb{Z}_p^3$ and outputs $([t^\top A]_1, [t^\top s]_1, [t^\top A]_2, [t^\top s]_2)$. Since the bit length of p is $\Theta(\lambda)$ and two rows of A are uniformly random, $A \in \mathbb{Z}_p^{3 \times 2}$ is rank 2 with an overwhelming probability. Thus, it simulates \mathcal{O}^r when $s = Ar$. When s is random, $[A|s] \in \mathbb{Z}_p^{3 \times 3}$ is full-rank with an overwhelming probability since two rows of $[A|s]$ are uniformly random. Thus, it simulates \mathcal{O} in this case. \square

Let Q be an upper bound of the number of queries. We define the oracles $\mathcal{O}_{\mathcal{K}'}^0, \mathcal{O}_{x,y}^{0,t/3}$ ($t = 0, 1, \dots, 3Q$) and the events E, E_1, E_2 in the same way as the proof of Theorem 2. Note that the secret key \mathcal{K}' contains x, y in CSSE, so the oracles depend on x, y .

First, there exists a PPT algorithm \mathcal{B}_1 such that

$$|\Pr[1 \leftarrow \mathcal{A}_{\mathcal{K}'}^{\mathcal{O}}(1^\lambda)] - \Pr[1 \leftarrow \mathcal{A}_{x,y}^{\mathcal{O}}(1^\lambda)]| \leq 2\text{Adv}_{\mathcal{B}_1, F}^{\text{PRF}}(\lambda).$$

In addition, we have

$$\Pr[1 \leftarrow \mathcal{A}_{x,y}^{\mathcal{O}_{x,y}^{0,i}}(1^\lambda) \wedge \neg E] = \Pr[1 \leftarrow \mathcal{A}_{x,y}^{\mathcal{O}_{x,y}^{0,i+1}}(1^\lambda) \wedge \neg E].$$

In the case of E_1 , consider treating $(\text{FList}(w_{0,i+1}), y)$ as the secret and constructing a bil- V_2 -MDDH solver using the oracle distinguisher of $\mathcal{O}_{x,y}^{0,i}$ and $\mathcal{O}_{x,y}^{0,i+1/3}$. The solver samples $x \xleftarrow{\$} \mathbb{Z}_p$. Since the solver knows $(x, [y]_1, [y]_2)$, it can correctly respond to queries from the distinguisher. Note that in this case, “random” in Table 2 means that the oracle returns $([(r_1, r_2)^\top]_1, [(r_3, r_1 x)^\top]_1)$ where $r_1, r_2, r_3 \xleftarrow{\$} \mathbb{Z}_p$. Thus, it follows that there exists a PPT algorithm $\mathcal{B}_{2,1}$ such that

$$|\Pr[1 \leftarrow \mathcal{A}_{x,y}^{\mathcal{O}_{x,y}^{0,i}}(1^\lambda) \wedge E_1] - \Pr[1 \leftarrow \mathcal{A}_{x,y}^{\mathcal{O}_{x,y}^{0,i+1/3}}(1^\lambda) \wedge E_1]| \leq 2\text{Adv}_{\mathcal{B}_{2,1}}^{\text{bil-}V_2\text{-MDDH}}(\lambda).$$

Similarly, there exist PPT algorithms $\mathcal{B}_{2,t}$ ($t = 2, 3$) such that

$$|\Pr[1 \leftarrow \mathcal{A}_{x,y}^{\mathcal{O}_{x,y}^{0,i+(t-1)/3}}(1^\lambda) \wedge E_1] - \Pr[1 \leftarrow \mathcal{A}_{x,y}^{\mathcal{O}_{x,y}^{0,i+t/3}}(1^\lambda) \wedge E_1]| \leq 2\text{Adv}_{\mathcal{B}_{2,t}}^{\text{bil-}V_2\text{-MDDH}}(\lambda).$$

Thus, for some PPT algorithm \mathcal{B}_2 ,

$$|\Pr[1 \leftarrow \mathcal{A}_{x,y}^{\mathcal{O}_{x,y}^{0,i}}(1^\lambda) \wedge E] - \Pr[1 \leftarrow \mathcal{A}_{x,y}^{\mathcal{O}_{x,y}^{0,i+1}}(1^\lambda) \wedge E]| \leq 12\text{Adv}_{\mathcal{B}_2}^{\text{bil-}V_2\text{-MDDH}}(\lambda).$$

holds. Therefore,

$$\text{Adv}_{\mathcal{A}, \text{CSSE}}^{\text{Full}}(\lambda) \leq 2\text{Adv}_{\mathcal{B}_1, F}^{\text{PRF}}(\lambda) + 12Q\text{Adv}_{\mathcal{B}_2}^{\text{bil-}V_2\text{-MDDH}}(\lambda)$$

holds and this completes the proof.

7 Comments on Li-Boyen Schemes

In this section, we show that the Li-Boyen pairing-based SSE [21] is not trapdoor-private. Then, we give some comments on the Li-Boyen NTRU-based SSE.

7.1 Attack on Li-Boyen Pairing-based Scheme

First, we describe the algorithms of the Li-Boyen pairing-based scheme [21]. Let $F : \mathcal{K} \times \mathcal{W} \rightarrow \mathbb{Z}_p$ be a pseudorandom function.

- **KeyGen**(1^λ):
 1. Runs $\mathbb{G} := (p, G_1, G_2, G_T, g_1, g_2, e) \leftarrow \mathcal{G}(1^\lambda)$.
 2. Samples $z \xleftarrow{\$} \mathbb{Z}_p$ and sets $u_1 = g_1^z, u_2 = g_2^z$.
 3. Samples $h_2 \xleftarrow{\$} G_2$.
 4. Samples $\mathcal{K} \xleftarrow{\$} \mathcal{K}$.
 5. Outputs the public parameter $(\mathbb{G}, u_1, u_2, h_2)$ and the secret key \mathcal{K} .
- **Enc**(\mathcal{K}, w): Samples $s \xleftarrow{\$} \mathbb{Z}_p$ and outputs $(c_1, c_2, c_3) := (e(g_1, h_2)^{F_{\mathcal{K}}(w) \cdot s}, u_1^s, g_1^s)$.
- **Trapdoor**(\mathcal{K}, w): Samples $r \xleftarrow{\$} \mathbb{Z}_p$ and outputs $(t_1, t_2) := (h_2^{F_{\mathcal{K}}(w)} u_2^r, g_2^r)$.
- **Test**($(t_1, t_2), (c_1, c_2, c_3)$): Outputs 1 if $e(c_3, t_1) = c_1 \cdot e(c_2, t_2)$, otherwise outputs 0.

They claimed that its ciphertext privacy can be reduced to the DBDH assumption, and its trapdoor privacy can be reduced to the DDH assumption on G_1 (See Appendix B for the definitions of ciphertext/trapdoor privacy). However, their proof for trapdoor privacy uses an incorrect DDH definition, and trapdoor privacy can be broken in the following way.

Theorem 8. The Li-Boyen pairing-based SSE does not satisfy trapdoor privacy.

Proof. The adversary chooses arbitrary challenge keywords w_0^*, w_1^* ($w_0^* \neq w_1^*$) and receives a challenge trapdoor $(t_1^*, t_2^*) = (h_2^{F_k(w_b^*)} u_2^{r^*}, g_2^{r^*})$ where $b \xleftarrow{\$} \{0, 1\}$. In addition, the adversary receives a trapdoor $(t_1, t_2) = (h_2^{F_k(w_0^*)} u_2^r, g_2^r)$ of w_0^* and a ciphertext $(c_1, c_2, c_3) = (e(g_1, h_2)^{F_k(w) \cdot s}, u_1^s, g_1^s)$ of an arbitrary keyword w ($w \notin \{w_0^*, w_1^*\}$). Then, we have $e(c_3, t_1^*)/e(c_2, t_2^*) = e(g_1, h_2)^{F_k(w_b^*) \cdot s}$ and $e(c_3, t_1)/e(c_2, t_2) = e(g_1, h_2)^{F_k(w_0^*) \cdot s}$. By checking whether they are equal, the adversary can guess whether $b = 0$ with an overwhelming probability. \square

In this attack, (u_1, g_1) can be used instead of (c_2, c_3) . However, we used (c_2, c_3) to indicate that keeping (u_1, g_1) as secret in this scheme does not affect the proof.

7.2 On Li-Boyen NTRU-based Scheme

We give the concrete description of the Li-Boyen NTRU-based scheme in Appendix C, which is not given in [21]. It is claimed in [21] that ciphertexts and trapdoors are in \mathcal{R}_q^3 , but Appendix C shows that they are actually in \mathcal{R}_q^2 .

In the Li-Boyen schemes based on ideal lattices, the modulus q is set to prime. In particular, the Li-Boyen NTRU-based scheme is based on the structure of Ducas et al. IBE [14], so q is a prime such that $q \equiv 1 \pmod{2N}$. However, their security proof uses a kind of regularity lemma for $q = 3^k$ [15]. Instead, we can use the regularity lemma by Lyubashevsky et al [24]. It ensures that for an overwhelming fraction of $h \in \mathcal{R}_q$, $a_1 + a_2 h$ is statistically close to random where $a_1, a_2 \in \mathcal{R}_q$ are sampled from the discrete Gaussian distribution with large enough deviation.

Consistency is not considered in [21], but it can be proved in a similar way to the proof for our lattice-based construction. In the consistency game, a ciphertext (c, d_2) and a trapdoor (t, u_2) can be regarded as independent samples from $\text{Unif}(\mathcal{R}_q) \times \mathcal{D}$ where \mathcal{D} denotes the discrete Gaussian distribution on \mathcal{R}_q . By the regularity lemma, $td_2 + d_1$ where $d_1 \leftarrow \mathcal{D}$ is statistically close to uniform. Since each coefficients of d_1 is relatively close to 0, the false-positive rate of this scheme is small.

References

- [1] Michel Abdalla, Mihir Bellare, Dario Catalano, Eike Kiltz, Tadayoshi Kohno, Tanja Lange, John Malone-Lee, Gregory Neven, Pascal Paillier, and Haixia Shi, “Searchable encryption revisited: Consistency properties, relation to anonymous IBE, and extensions,” CRYPTO 2005, pages 205-222, 2005.
- [2] Michel Abdalla, Dario Catalano, Dario Fiore, Romain Gay, and Bogdan Ursu, “Multi-input functional encryption for inner products: Function-hiding realizations and constructions without pairings,” CRYPTO 2018, pages 597-627, 2018.
- [3] Shashank Agrawal, Shweta Agrawal, Saikrishna Badrinarayanan, Abishek Kumarasubramanian, Manoj Prabhakaran, and Amit Sahai, “On the practical security of inner product functional encryption,” PKC 2015, pages 777-798, 2015.
- [4] Shweta Agrawal and Shota Yamada, “CP-ABE for circuits (and more) in the symmetric key setting,” TCC 2020, pages 117-148, 2020.
- [5] Joël Alwen and Chris Peikert, “Generating shorter bases for hard random lattices,” Theory of Computing Systems 48, pages 535-553, 2011.

- [6] Allison Bishop, Abhishek Jain, and Lucas Kowalczyk, “Function-hiding inner product encryption,” ASIACRYPT 2015, pages 470-491, 2015.
- [7] Dan Boneh, Giovanni Di Crescenzo, Rafail Ostrovsky, and Giuseppe Persiano, “Public key encryption with keyword search,” EUROCRYPT 2004, pages 506-522, 2004.
- [8] Dan Boneh, Kevin Lewi, Hart Montgomery, and Ananth Raghunathan, “Key homomorphic PRFs and their applications,” CRYPTO 2013, pages 410-428, 2013.
- [9] Zvika Brakerski and Gil Segev, “Function-private functional encryption in the private-key setting,” Journal of Cryptology 31, pages 202-225, 2018.
- [10] Jin Wook Byun, Hyun Suk Rhee, Hyun-A Park, and Dong Hoon Lee, “Off-line keyword guessing attacks on recent keyword search schemes over encrypted data,” SDM 2006, pages 75-83, 2006.
- [11] Leixiao Cheng and Fei Meng, “Public key authenticated encryption with keyword search from LWE,” ESORICS 2022, pages 303-324, 2022.
- [12] Leixiao Cheng, Jing Qin, Feng Feng, Fei Meng, “Security-enhanced public-key authenticated searchable encryption,” Information Sciences 647:119454, 2023.
- [13] Pratish Datta, Ratna Dutta, and Sourav Mukhopadhyay, “Functional encryption for inner product with full function privacy,” PKC 2016, pages 164-195, 2016.
- [14] Léo Ducas, Vadim Lyubashevsky, and Thomas Prest, “Efficient identity-based encryption over NTRU lattices,” ASIACRYPT 2014, pages 22-41, 2014.
- [15] Léo Ducas and Daniele Micciancio, “Improved short lattice signatures in the standard model,” CRYPTO 2014, pages 335-352, 2014.
- [16] Keita Emura, “Generic construction of public-key authenticated encryption with keyword search revisited: stronger security and efficient construction,” APKC 2022, pages 39-49, 2022.
- [17] Alex Escala, Gottfried Herold, Eike Kiltz, Carla Ràfols, and Jorge Villar, “An algebraic framework for Diffie–Hellman assumptions,” Journal of Cryptology 30, pages 242-288, 2017.
- [18] Craig Gentry, Chris Peikert, Vinod Vaikuntanathan, “Trapdoors for hard lattices and new cryptographic constructions,” ACM STOC 2008, pages 197-206, 2008.
- [19] Qiong Huang and Hongbo Li, “An efficient public-key searchable encryption scheme secure against inside keyword guessing attacks,” Information Sciences 403, pages 1-14, 2017.
- [20] Sam Kim, Kevin Lewi, Avradip Mandal, Hart Montgomery, Arnab Roy, and David J. Wu, “Function-hiding inner product encryption is practical,” SCN 2018, pages 544-562, 2018.
- [21] Qinyi Li and Xavier Boyen, “Public-key authenticated encryption with keyword search made easy,” IACR Communications in Cryptology, vol. 1, no. 2, 2024.
- [22] Huijia Lin, “Indistinguishability obfuscation from SXDH on 5-linear maps and locality-5 PRGs,” CRYPTO 2017, pages 599-629, 2017.
- [23] Zi-Yuan Liu, Yi-Fan Tseng, Raylin Tso, Masahiro Mambo, Yu-Chi Chen, “Public-key authenticated encryption with keyword search: Cryptanalysis, enhanced security, and quantum-resistant instantiation,” ASIACCS 2022, pages 423-436, 2022.
- [24] Vadim Lyubashevsky, Chris Peikert, and Oded Regev, “A toolkit for ring-LWE cryptography,” EUROCRYPT 2013, pages 35-54, 2013.
- [25] Kenneth Maples, “Singularity of random matrices over finite fields,” arXiv preprint arXiv:1012.2372, 2010.

- [26] Daniele Micciancio and Chris Peikert, “Trapdoors for lattices: Simpler, tighter, faster, smaller,” EUROCRYPT 2012, pages 700-718, 2012.
- [27] Hoi H. Nguyen and Elliot Paquette, “Surjectivity of near-square random matrices,” *Combinatorics, Probability and Computing* 29.2, pages 267-292, 2020.
- [28] Mahnaz Noroozi and Ziba Eslami, “Public key authenticated encryption with keyword search: revisited,” *IET Information Security* 13(4), pages 336-342, 2019.
- [29] Baodong Qin, Yu Chen, Qiong Huang, Ximeng Liu, Dong Zheng, “Public-key authenticated encryption with keyword search revisited: Security model and constructions,” *Information Sciences* 516, pages 515-528, 2020.
- [30] Baodong Qin, Hui Cui, Xiaokun Zheng, and Dong Zheng, “Improved security model for public-key authenticated encryption with keyword search,” *ProvSec 2021*, pages 19-38, 2021.
- [31] Oded Regev, “On lattices, learning with errors, random ear codes, and cryptography,” *Journal of the ACM* 56(6), pages 1-40, 2009.
- [32] Emily Shen, Elaine Shi, and Brent Waters, “Predicate privacy in encryption systems,” *TCC 2009*, pages 457-473, 2009.
- [33] Dawn Xiaodong Song, David Wagner, and Adrian Perrig, “Practical techniques for searches on encrypted data,” *IEEE S&P 2000*, pages 44-55, 2000.
- [34] Junichi Tomida, Masayuki Abe, and Tatsuaki Okamoto, “Efficient functional encryption for inner-product values with full-hiding security,” *ISC 2016*, pages 408-425, 2016.

A Public key Authenticated Encryption with Keyword Search (PAEKS)

We describe the syntax of Public key Authenticated Encryption with Keyword Search (PAEKS). PAEKS consists of the following six PPT algorithms.

- $\text{Setup}(1^\lambda)$: Given a security parameter λ , it outputs a public parameter pp .
- $\text{KeyGen}_S(\text{pp})$: Given a public parameter pp , it outputs a sender’s public key pk_S and a sender’s secret key sk_S .
- $\text{KeyGen}_R(\text{pp})$: Given a public parameter pp , it outputs a receiver’s public key pk_R and a receiver’s secret key sk_R .
- $\text{Enc}(\text{pp}, \text{pk}_R, \text{pk}_S, \text{sk}_S, w)$: Given a public parameter pp , a receiver’s public key pk_R , a sender’s public key pk_S , a sender’s secret key sk_S , and a keyword w , it outputs a ciphertext C .
- $\text{Trapdoor}(\text{pp}, \text{pk}_R, \text{pk}_S, \text{sk}_R, w)$: Given a public parameter pp , a receiver’s public key pk_R , a sender’s public key pk_S , a receiver’s secret key sk_R , and a keyword w , it outputs a trapdoor T .
- $\text{Test}(\text{pp}, T, C)$: Given a public parameter pp , a trapdoor T , and a ciphertext C , it outputs a bit $b \in \{0, 1\}$.

In the generic construction of PAEKS by Li and Boyen [21], Enc and Trapdoor compute the same secret value K using NIKE. Then, ciphertexts and trapdoors are generated by running $\text{SSE.Enc}(K, w)$ and $\text{SSE.Trapdoor}(K, w)$ respectively.

B Ciphertext/Trapdoor Privacy of SSE

We describe the definitions of ciphertext privacy and trapdoor privacy. Note that full security implies both of them.

Let $\mathcal{O}_{\text{ct}}(K, \cdot)$ be the oracle that outputs $C \leftarrow \text{Enc}(K, w)$ given a keyword w as input. Let $\mathcal{O}_{\text{td}}(K, \cdot)$ be the oracle that outputs $T \leftarrow \text{Trapdoor}(K, w)$ given a keyword w as input.

Definition 10 (Ciphertext Privacy). We say that an SSE scheme is ciphertext-private if for any PPT algorithm \mathcal{A} , the advantage

$$\left| \Pr[(\text{pp}, \mathbf{K}) \leftarrow \text{KeyGen}(1^\lambda); (w_0^*, w_1^*, \text{st}) \leftarrow \mathcal{A}^{\mathcal{O}_{\text{ct}}(\mathbf{K}, \cdot), \mathcal{O}_{\text{td}}(\mathbf{K}, \cdot)}(\text{find}, \text{pp}); b \stackrel{\$}{\leftarrow} \{0, 1\}; C^* \leftarrow \text{Enc}(\mathbf{K}, w_b^*); b' \leftarrow \mathcal{A}^{\mathcal{O}_{\text{ct}}(\mathbf{K}, \cdot), \mathcal{O}_{\text{td}}(\mathbf{K}, \cdot)}(\text{guess}, \text{st}, C^*) : b = b'] - \frac{1}{2} \right|$$

is negligible in λ . The challenge keyword w_0^*, w_1^* cannot be input to $\mathcal{O}_{\text{td}}(\mathbf{K}, \cdot)$.

Definition 11 (Trapdoor Privacy). We say that an SSE scheme is trapdoor-private if for any PPT algorithm \mathcal{A} , the advantage

$$\left| \Pr[(\text{pp}, \mathbf{K}) \leftarrow \text{KeyGen}(1^\lambda); (w_0^*, w_1^*, \text{st}) \leftarrow \mathcal{A}^{\mathcal{O}_{\text{ct}}(\mathbf{K}, \cdot), \mathcal{O}_{\text{td}}(\mathbf{K}, \cdot)}(\text{find}, \text{pp}); b \stackrel{\$}{\leftarrow} \{0, 1\}; T^* \leftarrow \text{Trapdoor}(\mathbf{K}, w_b^*); b' \leftarrow \mathcal{A}^{\mathcal{O}_{\text{ct}}(\mathbf{K}, \cdot), \mathcal{O}_{\text{td}}(\mathbf{K}, \cdot)}(\text{guess}, \text{st}, T^*) : b = b'] - \frac{1}{2} \right|$$

is negligible in λ . The challenge keyword w_0^*, w_1^* cannot be input to $\mathcal{O}_{\text{ct}}(\mathbf{K}, \cdot)$.

C Description of Li-Boyen NTRU-based Scheme

We describe the concrete construction of the Li-Boyen NTRU-based scheme. Let DLP.KeyGen denote the key generation algorithm of Ducas et al. IBE (DLP-IBE) [14] and SampleD denote the preimage sampler used in DLP-IBE. Let $\chi(\lambda)$ be an error distribution on \mathcal{R}_q , $\sigma(\lambda), \beta(\lambda)$ be positive real parameters, and $F : \mathcal{K} \times (\mathcal{W} \times \{0, 1\}) \rightarrow \{0, 1\}^\ell$ be a pseudorandom function.

- $\text{KeyGen}(1^\lambda)$:
 1. Samples $u \stackrel{\$}{\leftarrow} \mathcal{R}_q$.
 2. Samples $\mathbf{K} \stackrel{\$}{\leftarrow} \mathcal{K}$.
 3. Outputs the public parameter u and the secret key \mathbf{K} .
- $\text{Enc}(\mathbf{K}, w)$:
 1. Computes $\text{rnd}_0 = F_{\mathbf{K}}(w, 0), \text{rnd}_1 = F_{\mathbf{K}}(w, 1)$.
 2. Runs $(h, B) \leftarrow \text{DLP.KeyGen}(1^\lambda; \text{rnd}_0)$.
 3. Runs $(h', B') \leftarrow \text{DLP.KeyGen}(1^\lambda; \text{rnd}_1)$.
 4. Samples $s, x, y \leftarrow \chi$.
 5. Runs $(d_1, d_2) \leftarrow \text{SampleD}(B', us + x, \sigma)$.
(Then $d_1 + h'd_2 = us + x$ holds.)
 6. Computes $c \leftarrow hs + y$.
 7. Outputs the ciphertext $(c, d_2) \in \mathcal{R}_q^2$.
- $\text{Trapdoor}(\mathbf{K}, w)$:
 1. Computes $\text{rnd}_0 = F_{\mathbf{K}}(w, 0), \text{rnd}_1 = F_{\mathbf{K}}(w, 1)$.
 2. Runs $(h, B) \leftarrow \text{DLP.KeyGen}(1^\lambda; \text{rnd}_0)$.
 3. Runs $(h', B') \leftarrow \text{DLP.KeyGen}(1^\lambda; \text{rnd}_1)$.
 4. Samples $s', x', y' \leftarrow \chi$.
 5. Runs $(u_1, u_2) \leftarrow \text{SampleD}(B, us' + x', \sigma)$.
(Then $u_1 + hu_2 = us' + x'$ holds.)
 6. Computes $t \leftarrow h's' + y'$.
 7. Outputs the trapdoor $(t, u_2) \in \mathcal{R}_q^2$.
- $\text{Test}((t, u_2), (c, d_2))$: Outputs 1 if $\|td_2 - u_2c\| \leq \beta$, otherwise outputs 0.