


A Lattice-Based IND-CCA Threshold KEM from the BCHK+ Transform

Oleksandra Lapiha¹ and Thomas Prest² 

¹ Royal Holloway, University of London
sasha.lapiha.2021@live.rhul.ac.uk

² PQShield
thomas.prest@pqshield.com

Abstract. We present a simple IND-CCA lattice-based threshold KEM. At a high level, our design is based on the BCHK transform (Canetti et al., EUROCRYPT 2004), which we adapt to the lattice setting by combining it with the FO transform (Fujisaki and Okamoto, PKC 1999) in order to achieve decryption consistency.

As for the BCHK transform, our construction requires a threshold identity-based encryption (TIBE) scheme with suitable properties. We build such an IBE by combining the ABB IBE (Agrawal, Boneh, Boyen, EUROCRYPT 2010) with recent advances in lattice threshold cryptography, such as the threshold-friendly signature Plover (Esgin et al., EUROCRYPT 2024) and a variant of the Threshold Raccoon scheme (Katsumata et al., CRYPTO 2024).

The security proof of our scheme relies on a new assumption which we call the Coset-Hint-MLWE assumption, and which is a natural generalisation of the Hint-MLWE assumption (Kim et al., CRYPTO 2023). We prove the hardness of Coset-Hint-MLWE under standard assumptions. We believe this new assumption may be of independent interest.

Unlike prior works on IND-CCA lattice-based threshold KEMs, our construction only relies on simple algorithmic tools and does not use heavy machinery such as multi-party computation or threshold fully homomorphic encryption.

1 Introduction

Lattices are one of the most prominent mathematical families for building post-quantum cryptosystems, thanks to their versatility and efficiency. In 2024, NIST has released their first lattice-based cryptographic standards, the ML-DSA signature scheme and the ML-KEM [NIS24] key encapsulation mechanism (KEM). Based on Kyber [SAB⁺22], ML-KEM realises the standard security notion of semantic security against chosen-ciphertext attacks (IND-CCA), which is considered necessary for real-life deployment. The standard approach for achieving IND-CCA, which is also used by ML-KEM, is the Fujisaki-Okamoto (FO) transform. This transform relies heavily on the evaluation by the decrypter of hash functions (modelled as random oracles) on private input (the decrypted message).

Threshold decryption. In some applications, it is desirable to split the power of decryption across several parties. In a threshold KEM, the decapsulation key would be split across N parties, of which at least $T \leq N$ would be needed to be able to decrypt valid ciphertexts. Most threshold cryptosystems share the secret key using linear secret sharing (such as Shamir’s), and perform sensitive operations using (linearly) homomorphic properties.

In e-voting protocols, threshold cryptosystems have been suggested as a solution to protect the confidentiality of the tally until the close of voting, see for example the seminal work of Cramer et al. [CGS97], or Aranha et al. [ABGS23] for a lattice-based protocol. More recently, they have been advertised as a way to enhance the privacy of mempools (a mempool is a queue of pending and unconfirmed transactions for a cryptocurrency network node), see for example Choudhuri et al. [CGPP24]. We expect that NIST’s upcoming call for threshold cryptosystems [LTANB25] will foster more applications.

Limitations of the Fujisaki-Okamoto transform. Unfortunately, the FO transform is difficult to perform in a distributed manner: this would require expensive hash function evaluation using multi-party computation (MPC).

While a few IND-CCA lattice-based threshold decryption schemes have been proposed in the literature, they often use heavy machinery such as non-interactive zero-knowledge arguments (NIZKs), threshold fully homomorphic encryption (ThFHE) or MPC; as a result, their efficiency is unclear. It is therefore highly desirable to look into alternative designs for lattice-based IND-CCA threshold decryption.

1.1 Our Contributions

We propose a simple lattice-based threshold KEM that achieves IND-CCA security. At a high level, our construction is based on the Boneh-Canetti-Halevi-Katz transform [CHK04, BCHK07], or BCHK transform, which has been shown by Boneh, Boyen and Halevi [BBH06] to allow constructing IND-CCA TPKEs under suitable conditions. Unfortunately, in the case of lattices, the BCHK transform cannot be used as-is, and in particular it does not guarantee decryption/decapsulation consistency.

Our first contribution is to provide a variant of the BCHK transform that is tailored for lattices. This construction, which we call BCHK+, yields a threshold KEM with decapsulation consistency in addition to IND-CCA security.

Equipped with our new BCHK+ transform, we build an IND-CCA TKEM. One of the key ingredients is a threshold identity-based encryption (TIBE) with suitable properties. As no such TIBE exists in the lattice world, we construct our own lattice-based TIBE. This construction ended up being rather technical, and is our second contribution.

A third contribution, which might be of independent interest, is the introduction of the Coset-Hint-MLWE assumption, a natural generalisation of the

Hint-MLWE assumption [KLSS23]. This assumption underlies the security of our TIBE, and therefore of our TKEM. Finally, we show that the hardness of Hint-MLWE and the existence of trapdoor sampling implies the hardness of Coset-Hint-MLWE.

1.2 Technical Overview

The BCHK+ transform (Section 4.1). Our starting point is the BCHK transform, a semi-generic transform by Boneh, Canetti, Halevi and Katz [CHK04, BCHK07] that (i) takes as input an IBE with suitable properties and a one-time signature scheme, and (ii) outputs a IND-CCA PKE. In a follow-up, Boneh, Boyen and Halevi [BBH06] modify the BCHK transform to (i) take as input a *threshold* IBE with suitable properties and a one-time signature scheme, and (ii) output an IND-CCA *threshold* PKE. In the instantiations of [BBH06] the validity test is performed on public values, and therefore does not need to be thresholdised.

Unfortunately, the BCHK transform cannot be applied as-is to our setting. Indeed, lattice ciphertexts are inherently noisy and, while this is not an issue for IND-CCA security, this precludes *decryption consistency*, a necessary condition in [BBH06] that stipulates that the same ciphertext cannot be decrypted by two sets of decryptions to two different messages.

We resolve this by combining the BCHK transform with the FO transform. Note that this does not contradict our prior claim that the FO transform is not compatible with thresholdisation. In schemes that achieve IND-CCA security via the FO transform alone, the input to re-encryption is sensitive and therefore needs to be thresholdised, as it is otherwise susceptible to re-encryption attacks [UXT⁺22]. In our case, IND-CCA is ensured by the BCHK transform, and FO is only required for decapsulation consistency. In particular, the input to FO re-encryption is no longer sensitive and does not need to be thresholdised. We call this new transform the BCHK+ transform, and it can be summarized as:

$$\text{BCHK+} = \text{BCHK (IND-CCA)} + \text{FO (decapsulation consistency)}.$$

Constructing a lattice TIBE (Section 4.2). The next step is to instantiate the BCHK+ transform with a TIBE. To our knowledge, no efficient lattice-based TIBE is described in the literature (in the work of Albrecht et al., [ALLW25] the public parameters scale quadratically with the decryption threshold). Starting from the seminar work of Gentry, Peikert and Vaikuntanathan [GPV08], Several lattice IBEs have been proposed: see for example Gentry, Peikert and Vaikuntanathan [GPV08], Agrawal, Boneh and Boyen [ABB10], Cash et al. [CHKP10], Ducas, Lyubashevsky and Prest [DLP14], and Katsumata and Yamada [KY16]; however these schemes all rely on Gaussian (trapdoor) sampling, which is notoriously difficult to thresholdise. These IBEs can be seen as extensions of lattice hash-then-sign schemes, where (i) the message to be signed becomes the identity, and (ii) the signature becomes the identity's decryption key. This correspondence is made explicit in [GPV08, ABB10, DLP14, BFRS18].

On the other hand, efficient lattice threshold signatures were recently proposed, based on either Fiat-Shamir “Raccoon” signatures [DKM⁺24, KRT24], or in the hash-then-sign paradigm [ENP24]. A common pattern of these schemes is to replace Gaussian sampling by the much threshold-friendlier *noise flooding*, and to quantify security using the recent Hint-MLWE assumption by Kim et al. [KLSS23]. For example, the Plover signature [EEN⁺24] and its threshold variant Pelican [ENP24] are noise-flooded variants of the trapdoor sampling-based Eagle [YJW23].

Given the correspondence between hash-then-sign and IBE schemes, a natural idea is to apply noise flooding to an existing IBE in order to make it *threshold-friendly*, then apply thresholdisation techniques from e.g. [DKM⁺24, KRT24] to obtain a proper TIBE. Unfortunately, this is not necessarily secure. Given an IBE, let us note \mathbf{A} the master public key, id an identity, and sk_{id} a valid decryption key for id . IBEs such as the ones in [GPV08, DLP14] satisfy $\mathbf{A} \cdot \text{sk}_{\text{id}} = H(\text{id})$. Thus an adversary could query two decryption keys $\text{sk}_{\text{id}}^1, \text{sk}_{\text{id}}^2$ for the same identity and compute a short $\mathbf{s} = \text{sk}_{\text{id}}^1 - \text{sk}_{\text{id}}^2$ that satisfies $\mathbf{A} \cdot \mathbf{s} = \mathbf{0}$. This breaks the security proof and, with a bit more effort, can be converted into a full key-recovery attack.

Instead, we start from the Agrawal-Boneh-Boyen IBE [ABB10]. In this IBE, a valid identity key satisfies $[\mathbf{A} \mid \mathbf{B}(\text{id})] \cdot \text{sk}_{\text{id}}$, where $\mathbf{B}(\text{id})$ is a matrix that depends on id . While the same attack as before applies, it only gives a trapdoor for $[\mathbf{A} \mid \mathbf{B}(\text{id})]$, which does not lead to a trapdoor for \mathbf{A} , and important properties such as selective ID security are preserved. Thus we select this IBE and convert it into a *threshold-friendly* IBE via noise flooding. We then convert this threshold-friendly IBE into a proper TIBE using the thresholdisation techniques described in [DKM⁺24, KRT24], in particular the $\text{TRaccoon}_{3-\text{rnd}}^{\text{sel}}$ scheme from [KRT24].

The security proof (Section 5) and Coset-Hint-MLWE (Section 3). The security proof (for the selective ID security) of our TIBE relies on a new assumption which we call the Coset-Hint-MLWE assumption (or Coset-Hint-MLWE), and which we believe might be of independent interest.

The MLWE assumption states that the pair $(\mathbf{A}, [\mathbf{A} \mid \mathbf{I}] \mathbf{s})$ is pseudorandom when $\mathbf{A} \in R_q^{m \times n}$ is uniformly random and $\mathbf{s} \in R^{n+m}$ is sampled from an adequate (short) distribution. The Hint-MLWE assumption [KLSS23] posits that this is the case even given ℓ hints of the form $(c_i, \mathbf{z}_i = c_i \cdot \mathbf{s} + \mathbf{r}_i)$, where all the $(c_i, \mathbf{r}_i)_i \in R \times R^{n+m}$ are sampled from an adequate (short) distribution. When $\mathbf{s}, (\mathbf{r}_i)_i$ are sampled from Gaussian distributions, there exist efficient dimension-preserving reductions from MLWE to Hint-MLWE [KLSS23, EEN⁺24, ENP24], and these will also be applicable to the parameters used in the paper. Hint-MLWE appears as a natural by-product of lattice signatures based on noise flooding, and has been used to build efficient lattice-based threshold signatures.

In the case of our TIBE, the threshold decryption procedure produces hints $(c_i, \mathbf{z}_i = c_i \cdot \mathbf{s} + \mathbf{r}_i)$ as in the usual Hint-MLWE, except that the (Gaussian) noise \mathbf{r}_i is no longer sampled from R^{n+m} , but from a public coset $\mathbf{a}_i + \Lambda \subseteq R^{n+m}$, for Λ a lattice and \mathbf{a}_i an offset. For this reason we call this new assumption

Coset-Hint-MLWE. One can see that it generalises Hint-MLWE, since taking $\Lambda = R^{n+m}$ recovers Hint-MLWE. While this is a natural generalisation, it could be unclear whether this is secure, as the noise is sampled from a sparser space.

We show that this is the case: for appropriate parameters, we can embed a trapdoor in \mathbf{A} to simulate the sampling of the noise in the coset. The reduction loses a number of dimensions that matches the dimension gap between Λ and R^{n+m} , which is intuitively what one would expect. For our concrete instantiation, we use Module-NTRU trapdoors [CPS⁺20]. The analysis of Coset-Hint-MLWE is detailed in Section 3.

1.3 Related Works

Lattice threshold signatures. Several threshold signatures schemes (TSS) based on lattices were recently proposed. In most of these, the base signature scheme is based on noise flooding and leverages Hint-MLWE assumption [KLSS23] for the security proof, see e.g. Raccoon [dPKPR24] and Plover [EEN⁺24]. The linear structure of these schemes facilitate their thresholdisation.

Regarding the linearisation tools, TSS based on fully homomorphic encryption have been proposed [ASY22, GKS24]. However, most schemes are based on standard tools, which gives 4-round [KRT24, ENP24], 3-round [DKM⁺24, KRT24] or 2-round [EKT24, CATZ24, BKL⁺25] TSS.

Lattice threshold cryptosystems. Compared to TSS, few lattice threshold public key encryption schemes (TPKEs, TIBEs) and threshold KEMs (TKEMs) have been proposed. Two recent papers [BS23, MS24] propose IND-CPA TPKEs based on noise-flooded variants of the Lindner-Peikert [LP11] and Regev [Reg09].

We have found two results in the area of lattice-based TIBEs. The work of [BKP13] builds a threshold version of GPV IBE with a proof in UC and adaptive corruptions. It is robust, and has potential for more compact sizes compared to our scheme. However, it requires generic MPC techniques for Gaussian sampling which typically has higher running time and round complexity.

The other TIBE construction is proved in [ALLW25]. It is expected to be relatively fast, since it only requires matrix vector multiplications, it has identifiable aborts and only 1 round. It is proved secure under weaker highly-selective security where the Adversary declares all oracle queries before receiving the public keys. The size of their master public key and ciphertexts are $O(T^2)$ and $O(T)$ respectively, whereas ours are essentially constant in T . Although [ALLW25] does not provide parameter estimations, we expect our concrete sizes to also be smaller.

IND-CCA constructions. IND-CCA security in a TPKE or TKEM remains challenging, as the usual reencryption approach by Fujisaki and Okamoto [FO99, FO13] does not thresholdise well.

The universal thresholdiser by Boneh et al. [BGG⁺18] uses ThFHE to thresholdise any functionality, and use it to build an IND-CCA TPKE. The scheme

is robust, has game-based selective security and only has one round. However to our knowledge, this work remains theoretical, and no parameter set has been proposed. It mentions the slow runtime as an open problem.

Cong et al. [CCMS21] proposed an IND-CCA TKEM called Gladius. On one hand, it also has the potential for smaller public key and ciphertext sizes. On the other hand, their design requires to evaluate a hash function using MPC. This entails an extremely high number of rounds (136491 rounds for 3 users), and also makes the security argument heuristic, since one cannot represent a random oracle using a circuit (which is needed for the MPC evaluation).

Finally, Devevey et al [DLN⁺21] propose a lattice-based INC-CCA TPKE achieving advanced notions such as adaptive security and robustness. However, it relies on heavy machinery such as lossy encryption schemes and NIZKs based on correlation-intractable hashes. Again, no parameter set has been proposed.

Compared to the above, our TIBE and TKEM are the only ones that simultaneously (i) do not require FHE/MPC/NIZK, (ii) have sizes essentially independent of T and N . As a result, they have small communication complexity and we expect that they would be fast when implemented. On the other hand, our schemes do not currently achieve robustness, have a trusted setup, and are less compact than [CCMS21]. We provide a summary of the related work in Table 1.

Table 1: Related work comparison for TIBE and TKEM. Sizes of $|ek|$ and $|ct|$ are given in KiB. N/A = Not Available, UC = Universal Composability, IA = Identifiable Aborts.

Reference	Security	Decrypt	$ ek $	$ ct $	Rounds	Robust	Tools
TIBE							
Bendlin et al. [BKP13]	Adaptive, UC	N/A	N/A	N/A	N/A	Yes	MPC
Albrecht et al. [ALLW25]	Highly-selective	N/A	$O(T^2)$	$O(T)$	1	IA	Generic
This work	Selective	N/A	50	450	3	No	Generic
IND-CCA TKEM							
Boneh et al. [BGG ⁺ 18]	Selective, UC	N/A	N/A	N/A	1	Yes	FHE
Kraitsberg et al. [KLO ⁺ 19]	Heuristic	4.342 sec. ³	4	4	≥ 114	No	MPC
Devevey et al. [DLN ⁺ 21]	Adaptive	N/A	N/A	N/A	1	Yes	NIZK
Cong et al. [CCMS21]	Selective, UC	4.99 sec. ⁴	N/A	N/A	136498	IA	MPC
This work	Selective	N/A	50	450	3	No	Generic

2 Preliminaries

2.1 Notations

We denote the security parameter as κ . We denote **assert** the function that aborts the procedure and returns \perp if the input condition is false. We note \ln the natural logarithm and \log the logarithm in basis 2.

³ In the 3-out-of-3 setting.

⁴ Ignoring latency.

Distributions. For an integer $N > 0$, we note $[N] = \{0, \dots, N-1\}$. To denote the assign operation, we use $y := f(x)$ when f is a deterministic and $y \leftarrow f(x)$ when randomized. When S is a finite set, we note $\mathcal{U}(S)$ the uniform distribution over S , and shorthand $x \leftarrow S$ for $x \leftarrow \mathcal{U}(S)$. We note $\text{SD}(X, Y)$ the statistical distance between two distributions X and Y . Given two distributions $\mathcal{P}_0, \mathcal{P}_1$ and a PPT adversary algorithm \mathcal{A} , we define the distinguishing advantage as $\text{Adv}_{\mathcal{A}}(\mathcal{P}_0, \mathcal{P}_1) := |\Pr(\mathcal{A}(\mathcal{P}_0) = 1) - \Pr(\mathcal{A}(\mathcal{P}_1) = 1)|$.

Norms. We note $\|\cdot\| = \|\cdot\|_2$ the Euclidean norm and $\|\cdot\|_\infty$ the infinity norm of a vector. Given a matrix $\mathbf{M} \in \mathbb{R}^{n \times m}$, we note $\|\mathbf{M}\| := \max_{\mathbf{x} \neq 0} \frac{\|\mathbf{M} \cdot \mathbf{x}\|}{\|\mathbf{x}\|}$ its spectral norm. $\|\mathbf{M}\|$ is also equal to the square root of the largest eigenvalue of $\mathbf{M}^t \cdot \mathbf{M}$. When \mathbf{M} is symmetric, $\|\mathbf{M}\|$ is equal to the largest eigenvalue of \mathbf{M} . The spectral norm is also called the largest singular value of \mathbf{M} , which is noted $s_1(\mathbf{M})$.

The Frobenius norm of a matrix $\mathbf{M} = (m_{i,j})_{i,j}$ is defined as $\|\mathbf{M}\|_F := \sqrt{\text{Trace}(\mathbf{M}^t \cdot \mathbf{M})} = \sqrt{\sum_{i,j} |m_{i,j}|^2}$. We recall that $\|\mathbf{M}\| \leq \|\mathbf{M}\|_F$. Lastly, we denote the Gram-Schmidt orthogonalisation of $\mathbf{M} \in \mathbb{R}^{n \times m}$ as matrix $\tilde{\mathbf{M}} \in \mathbb{R}^{n \times m}$ with columns $\tilde{\mathbf{m}}_i, i \in [m]$. Then the GS norm of a matrix $\mathbf{M} \in \mathbb{R}^{n \times m}$ is defined as $\|\mathbf{M}\|_{GS} = \max_i \|\tilde{\mathbf{m}}_i\|$.

Number fields. We consider cyclotomic number fields $K = \mathbb{Q}[\zeta_{2d}]$ for d a power-of-two. Alternatively, it can be written as $K = \mathbb{Q}[X]/(X^d + 1)$. Its ring of integers is of the form $R = \mathbb{Z}[X]/(X^d + 1)$. We choose q to be a prime modulus such that $R_q = R/qR$ is isomorphic to a direct product of two fields of order $q^{d/2}$. By Lemma 1 it suffices to pick $q = 5 \bmod 8$.

For $f = \sum_{i \in [d]} f_i X^i \in K$ we define the coefficient embedding as $\tau(f) = (f_0, \dots, f_{d-1}) \in \mathbb{Q}^d$. We note $\text{rot}(f)$ the matrix in $\mathbb{Q}^{d \times d}$ which i -th column (starting from 0) is the vector $\tau(X^i \cdot f)$. We extend the rot operator to matrices by entry-wise application. We also extend the spectral norm operator to any matrix $\mathbf{M} \in K^{n \times m}$: $\|\mathbf{M}\| = \|\text{rot}(\mathbf{M})\|$. For a vector $\mathbf{v} \in K^n$ we extend the norms as $\|\mathbf{v}\| = \|\tau(\mathbf{v})\|$ and $\|\mathbf{v}\|_\infty = \|\tau(\mathbf{v})\|_\infty$. We denote $S_\eta = \{c \in R \mid \|c\|_\infty \leq \eta\}$.

Lemma 1 ([LS18, Corollary 1.2]). *Let $2d > k > 0$ be powers of 2. If q is a prime such that $q = 2k + 1 \bmod 4k$ then there exist distinct $r_i \in \mathbb{Z}_q^*$ such that: (i) $X^{d/k} - r_i$ are irreducible in $\mathbb{Z}_q[X]$, and (ii) we can factor $X^d + 1$ as:*

$$X^d + 1 = \prod_{i=1}^k (X^{d/k} - r_i) \bmod q$$

Lemma 2 ([AL21, Proposition 2]). *Let R be a power-of-2 cyclotomic ring of degree d . Then*

$$\max_{a_0, a_1 \in R} \frac{\|a_0 \cdot a_1\|_\infty}{\|a_0\|_\infty \cdot \|a_1\|_\infty} \leq d$$

Decomposition function. For an odd prime q and β that is a power of 2 we define the function $\text{Decomp}_\beta : \mathbb{Z}_q \rightarrow \mathbb{Z} \times \mathbb{Z}$ that maps $x \in \{-\frac{q-1}{2}, \dots, \frac{q-1}{2}\}$ to the unique pair (c_0, c_1) such that $x = c_0 \cdot \beta + c_1$, $|c_1| \leq \beta/2$ and $|c_0| \leq \lceil \frac{q-1}{2\beta} \rceil$. We extend the function to a ring R and its quotient $R_q = R/qR$ as $\text{Decomp}_\beta : R_q \rightarrow R \times R$ by decomposing the elements coefficient-wise. Then $\|c_1\|_\infty \leq \beta/2$ and $\|c_0\|_\infty \leq \lceil \frac{q-1}{2\beta} \rceil$.

Shamir Secret Sharing. Let \mathbb{F} be a finite field. Given $a \in \mathbb{F}$, $1 \leq T$ and $\mathcal{S} \subseteq \mathbb{F}^*$, a (T, \mathcal{S}) -sharing of a is obtained by doing the following:

1. Generate $P \in \mathbb{F}_{<T}[x]$ uniformly at random, conditioned to $P(0) = a$.
2. For $s \in \mathcal{S}$, compute $\llbracket a \rrbracket_s^P = P(s)$.
3. The output is the indexed tuple $\llbracket a \rrbracket_{\mathcal{S}}^P = (\llbracket a \rrbracket_s^P)_{s \in \mathcal{S}}$.

Given (a, P, \mathcal{S}) , $\llbracket a \rrbracket_{\mathcal{S}}^P$ is uniquely defined. We say that an indexed tuple $\llbracket a \rrbracket$ is a valid T -sharing of a if there exists a polynomial $P \in \mathbb{F}_{<T}[x]$ and an evaluation set $\mathcal{S} \subseteq \mathbb{F}^*$ such that $\llbracket a \rrbracket = \llbracket a \rrbracket_{\mathcal{S}}^P$. If $|\mathcal{S}| \geq T$, then for any indexed tuple $b = (b_s)_{s \in \mathcal{S}}$ there exists at most a single pair (a, P) such that $b = \llbracket a \rrbracket_{\mathcal{S}}^P$.

When P and/or \mathcal{S} are clear from context, we may omit them and note $\llbracket a \rrbracket^P$, $\llbracket a \rrbracket_{\mathcal{S}}$ or $\llbracket a \rrbracket$. The set of valid (T, \mathcal{S}) -sharings is a \mathbb{F} -linear space of dimension T . Indeed, given $\lambda, \mu \in \mathbb{F}$:

$$\lambda \llbracket a \rrbracket_{\mathcal{S}}^P + \mu \llbracket b \rrbracket_{\mathcal{S}}^Q = \llbracket \lambda a \rrbracket_{\mathcal{S}}^{\lambda P} + \llbracket \mu b \rrbracket_{\mathcal{S}}^{\mu Q} = \llbracket \lambda a + \mu b \rrbracket_{\mathcal{S}}^{\lambda P + \mu Q} \quad (1)$$

Consider a (T, \mathcal{S}) -sharing \mathbf{a} . If $|\mathcal{S}| < T$, then \mathbf{a} leaks no information about the underlying shared value. If $|\mathcal{S}| = T$, then there exists a unique $a \in \mathbb{F}$ such that $\mathbf{a} = \llbracket a \rrbracket_{\mathcal{S}}$, and we can recover it using Lagrange interpolation:

$$a = \sum_{s \in \mathcal{S}} \lambda_{s, \mathcal{S}} \llbracket a \rrbracket_s, \quad \text{where} \quad \lambda_{s, \mathcal{S}} = \prod_{s' \in \mathcal{S} \setminus s} \frac{s'}{s' - s}. \quad (2)$$

For a prime q we extend the secret sharing to elements $a = \sum_{i \in [d]} a_i \cdot X^i \in R_q \sim \mathbb{Z}_q[X]/(X^d + 1)$ by sharing them coefficient-wise over the field \mathbb{Z}_q .

2.2 Cryptographic Primitives

In this section we define the syntax and the security properties of all cryptographic primitive used in the paper.

Definition 1 (Digital Signature Scheme). A digital signature scheme is a triple of PPT algorithms $\text{SIG} = (\text{SIG.Keygen}, \text{SIG.Sign}, \text{SIG.Verify})$ where

1. $\text{SIG.Keygen}(1^\kappa) \rightarrow (\text{sk}, \text{vk})$ takes as input the security parameter and outputs the signing key sk and the verification key vk .
2. $\text{SIG.Sign}(\text{sk}, \text{msg}) \rightarrow \text{sig}$ takes as input a message $\text{msg} \in \{0, 1\}^*$ and the signing key sk and outputs a signature sig .

3. $\text{SIG.Verify}(\text{vk}, \text{msg}, \text{sig}) \rightarrow b$ takes as input the verification key vk , a message $\text{msg} \in \{0, 1\}^*$ and a signature sig , and outputs a bit $b = 1$ if the signature and message pass the verification and $b = 0$ if they do not.

Definition 2 (IBE with Threshold Key Generation). An Identity-based Encryption Scheme with Threshold Key Generation is a tuple of PPT algorithms $\text{TIBE} = (\text{TIBE.Setup}, \text{TIBE.Encrypt}, (\text{TIBE.ShareExtract}_r)_r, \text{TIBE.Combine})$ s.t.:

1. $\text{TIBE.Setup}(1^\kappa) \rightarrow (\text{ek}, \{\text{dk}_i\}_{i \in N})$ takes as input the security parameter, and outputs the encryption key ek and N decryption key shares $\{\text{dk}_i\}_{i \in N}$.
2. $\text{TIBE.Encrypt}(\text{ek}, \text{id}, \text{msg}) \rightarrow \text{ct}$ takes as input an encryption key ek , identity id and message msg . It outputs a ciphertext ct .
3. For $r \in \{1, \dots, \text{rounds}\}$, $\text{TIBE.ShareExtract}_r(\text{id}, \text{act}, \text{dk}_i, \{\text{contrib}_{k,j}\}_{j < r}^{k \in \text{act}}) \rightarrow \text{contrib}_{i,r}$ takes as input an identity id , the set of active participants act , a decryption key share dk_i and contributions from previous rounds $\{\text{contrib}_{k,j}\}_{j < r}^{k \in \text{act}}$. It outputs a contribution of the current round $\text{contrib}_{i,r}$.
4. $\text{TIBE.Combine}(\text{ct}, \text{id}, \text{act}, \{\text{contrib}_{i,r}\}_{i \in \text{act}}) \rightarrow \text{msg}/\perp$ takes as input a set of active participants act , their key share contributions $\{\text{contrib}_{i,r}\}_{i \in \text{act}, r \leq \text{rounds}}$, an identity id and a ciphertext ct . It outputs a message msg or an error \perp .

Definition 3 (Threshold Key Encapsulation Mechanism). A threshold key encapsulation mechanism (TKEM) is a tuple of PPT algorithms $\text{TKEM} = (\text{TKEM.Keygen}, \text{TKEM.Encaps}, (\text{TKEM.ShareDecaps}_r)_r, \text{TKEM.Combine})$ s.t.:

1. $\text{TKEM.Keygen}(1^\kappa) \rightarrow (\text{ek}, \{\text{dk}_i\}_{i \in N})$ takes as input the security parameter. It outputs the encapsulation key ek and N decapsulation key shares $\{\text{dk}_i\}_{i \in N}$.
2. $\text{TKEM.Encaps}(\text{ek}) \rightarrow \text{ct}, K$ takes as input an encapsulation key ek . It outputs a ciphertext ct and a session key from the key space $K \in \mathcal{K}$.
3. For $r \in \{1, \dots, \text{rounds}\}$, $\text{TKEM.ShareDecaps}_r(\text{ct}, \text{act}, \text{dk}_i, \{\text{contrib}_{i,j}\}_{j < r}) \rightarrow \text{contrib}_{i,r}$ takes as input a ciphertext ct , a set of active parties act , a decapsulation key share dk_i and all contributions by parties in act up to the current round r . It outputs a contribution $\text{contrib}_{i,r}$ for the current round.
4. $\text{TKEM.Combine}(\text{ct}, \text{act}, \{\text{contrib}_{i,r}\}_{i \in \text{act}}) \rightarrow K/\perp$ takes as input a ciphertext ct , a set of active parties act and their contributions $\{\text{contrib}_{i,r}\}_{i \in \text{act}, r \leq \text{rounds}}$. It outputs a key K or an error message \perp .

Remark 1. We assume that TIBE.Setup (or respectively TKEM.Keygen) is run by a trusted authority which then distributes the keys among the parties via a secure channel. Then TIBE.ShareExtract (TKEM.ShareDecaps) is run interactively and as a result each party computes a decryption (decapsulation) share. Lastly the decryption shares are either sent to a single party or broadcast to all participants to decrypt the message by running TIBE.Combine (TKEM.Combine).

Definition 4 ($((Q, \varepsilon)\text{-sEU-CMA Security})$. A signature scheme SIG satisfies strong EU-CMA if any PPT adversary \mathcal{A} that makes at most Q signing queries wins $\text{Exp}_{\mathcal{A}, \text{SIG}}^{\text{sEU-CMA}}(1^\kappa)$ (see Figure 1a) with probability at most ε . When $Q = 1$ we say that SIG is an $\varepsilon\text{-sEU-CMA}$ one-time signature scheme.

$\text{Exp}_{\mathcal{A}, \text{SIG}}^{\text{sEU-CMA}}(1^\kappa)$
$(\text{sk}, \text{vk}) \leftarrow \text{SIG.Keygen}(1^\kappa)$
$(\text{msg}^*, \text{sig}^*) \leftarrow \mathcal{A}^{\text{SIG.Sign}(\text{sk}, \cdot)}(\text{vk})$
Let $\{(\text{msg}_i, \text{sig}_i)\}_{i=0}^{Q-1}$ be the list of \mathcal{A} 's oracle queries
return $(\text{SIG.Verify}(\text{vk}, \text{msg}^*, \text{sig}^*) = 1) \wedge (\text{msg}^*, \text{sig}^*) \notin \{(\text{msg}_i, \text{sig}_i)\}_{i=0}^{Q-1}$

(a) The (Q, ε) -sEU-CMA security game

$\text{Exp}_{\mathcal{A}, \text{TIBE}}^{\text{sel-ID-IND}}(1^\kappa)$
$(\text{id}^*, \text{cor}) \leftarrow \mathcal{A}$
assert $ \text{cor} = T - 1$
$(\text{ek}, \{\text{dk}_i\}_{i \in N}) \leftarrow \text{TIBE.Setup}(1^\kappa)$
$(\text{msg}_0, \text{msg}_1) \leftarrow \mathcal{A}^{\text{TIBE.ShareExtract}(\cdot, \text{id}), \text{id} \neq \text{id}^*}(\{\text{dk}_i\}_{\text{cor}})$
$b \leftarrow \{0, 1\}$
$\text{ct}^* = \text{TIBE.Encrypt}(\text{ek}, \text{id}^*, \text{msg}_b)$
$b' \leftarrow \mathcal{A}^{\text{TIBE.ShareExtract}(\cdot, \text{id}), \text{id} \neq \text{id}^*}(\{\text{dk}_i\}_{\text{cor}})$
return $(b = b')$

(b) TIBE selective-ID security

$\text{Exp}_{\mathcal{A}, \text{TKEM}}^{\text{CCA2}}(1^\kappa)$
$\text{cor} \leftarrow \mathcal{A}$
assert $ \text{cor} = T - 1$
$(\text{ek}, \{\text{dk}_i\}_{i \in N}) \leftarrow \text{TKEM.Keygen}(1^\kappa)$
$1 \leftarrow \mathcal{A}^{\text{TKEM.ShareDecaps}(\cdot)}(\{\text{dk}_i\}_{\text{cor}})$
$(\text{ct}^*, K_0) = \text{TKEM.Encaps}(\text{ek})$
$K_1 \leftarrow \mathcal{K}$
$b \leftarrow \{0, 1\}$
$b' \leftarrow \mathcal{A}^{\text{TKEM.ShareDecaps}(\text{ct}, \cdot), \text{ct} \neq \text{ct}^*}((\text{ct}^*, K_b))$
return $(b = b')$

(c) TKEM selective security

$\text{Exp}_{\mathcal{A}, \text{TIBE}}^{\text{sel-ID-OW}}(1^\kappa)$
$(\text{id}^*, \text{cor}) \leftarrow \mathcal{A}$
assert $ \text{cor} = T - 1$
$(\text{ek}, \{\text{dk}_i\}_{i \in N}) \leftarrow \text{TIBE.Setup}(1^\kappa)$
$\mathcal{A}^{\text{TIBE.ShareExtract}(\cdot, \text{id}), \text{id} \neq \text{id}^*}(\{\text{dk}_i\}_{\text{cor}})$
$\text{msg}^* \leftarrow \{0, 1\}^d$
$\text{ct}^* = \text{TIBE.Encrypt}(\text{ek}, \text{id}^*, \text{msg}^*)$
$\text{msg} \leftarrow \mathcal{A}^{\text{TIBE.ShareExtract}(\cdot, \text{id}), \text{id} \neq \text{id}^*}(\{\text{dk}_i\}_{\text{cor}}, \text{ct}^*)$
return $(\text{msg} = \text{msg}^*)$

(d) TIBE Selective-ID One-Way Security

$\text{Exp}_{\mathcal{A}, \text{TKEM}}^{\text{CD}}(1^\kappa)$
$\text{cor} \leftarrow \mathcal{A}$
assert $ \text{cor} = T - 1$
$(\text{ek}, \{\text{dk}_i\}_{i \in N}) \leftarrow \text{TKEM.Keygen}(1^\kappa)$
$(\text{ct}, \text{act}, \text{act}', S, S') \leftarrow \mathcal{A}^{\text{TKEM.ShareDecaps}(\cdot)}(\{\text{dk}_i\}_{i \in \text{cor}})$
$K_0 = \text{TKEM.Combine}(\text{ct}, \text{act}, S)$
$K_1 = \text{TKEM.Combine}(\text{ct}, \text{act}', S')$
return $(\forall i : K_i \neq \perp) \wedge (K_0 \neq K_1)$

(e) TKEM Decapsulation Consistency

Fig. 1: Security games

Definition 5 (Selective-ID/CCA2 Security). We say that a TIBE (resp. a TKEM) scheme achieves selective-ID indistinguishability, selective-ID one-wayness (resp. selective-CCA2 security) if no efficient adversary against the game in Figure 1b, Figure 1d (resp. Figure 1c) has a non-negligible advantage. We denote the probability of winning this game for an adversary \mathcal{A} as $\text{Adv}_{\mathcal{A}, \text{TIBE}}^{\text{sel-ID-IND}}$, $\text{Adv}_{\mathcal{A}, \text{TIBE}}^{\text{sel-ID-OW}}$ (resp. $\text{Adv}_{\mathcal{A}, \text{TKEM}}^{\text{CCA2}}$).

Definition 6 (Correctness). A TKEM scheme is δ -correct if for every message msg , every key set $(\text{ek}, \{\text{dk}_i\}_{i \in N})$ and every set act of decapsulators over the randomness of the honest encapsulation and decapsulation procedures

$$\Pr_{\text{ct}, \text{contrib}} (\text{TKEM.Combine}(\text{ct}, \text{act}, \text{contrib}) = \text{msg}) \geq 1 - \delta$$

where we denote $\text{ct} := \text{TKEM.Encaps}(\text{ek}, \text{msg})$ and $\text{contrib} := \{\text{contrib}_{i,j}\}_{i \in \text{act}}^{j \leq r}$ from the partial decapsulation procedure.

Definition 7 (Decapsulation Consistency [BBH06]). We say that a TKEM scheme satisfies Decapsulation Consistency if no PPT adversary can win the CD experiment in Figure 1e with a non-negligible probability. We write $\text{Adv}_{\mathcal{A}, \text{TKEM}}^{\text{CD}}$ for the probability of winning the game for an adversary \mathcal{A} .

Remark 2. We slightly modify the definition of Decryption Consistency of Boneh, Boyen and Halevi [BBH06]. As opposed to [BBH06], now the Adversary only wins if they can decrypt the ciphertext to two different **valid** messages. Note that this change reflects that our construction is not robust, nor it has identifiable aborts. The adversary is always able to make the decryption fail by providing an invalid decryption share.

Definition 8 (γ -spread IBE adapted from [HHK17]). Let TIBE be a Threshold IBE scheme defined in Definition 2 and let $\mathcal{K}, \mathcal{M}, \mathcal{I}, \mathcal{C}$ denote the sets of valid encryption keys, messages, identities, and ciphertexts respectively. For given $\text{ek} \in \mathcal{E}$, $\text{msg} \in \mathcal{M}$, $\text{id} \in \mathcal{I}$ we define the min-entropy of the ciphertext as

$$\gamma(\text{msg}, \text{id}, \text{ek}) := -\log \max_{\text{ct} \in \mathcal{C}} \Pr_{\text{rand}}(\text{ct} = \text{TIBE.Encrypt}(\text{msg}, \text{id}, \text{ek}; \text{rand})).$$

For $\gamma > 0$ the TIBE is γ -spread if for every tuple $(\text{msg}, \text{id}, \text{ek})$: $\gamma(\text{msg}, \text{id}, \text{ek}) \geq \gamma$.

2.3 Lattices

A lattice Λ is a discrete additive subgroup of \mathbb{R}^m . It is usually defined as all integer linear combinations of a set of vectors $\mathbf{b}_1, \dots, \mathbf{b}_n \in \mathbb{R}^m$ that are linearly independent in \mathbb{R}^m . The matrix $\mathbf{B} = [\mathbf{b}_1, \dots, \mathbf{b}_n]$ is called the basis of the lattice and we write $\Lambda = \Lambda(\mathbf{B})$. For a full-rank lattice $\det(\Lambda) := \det(\mathbf{B})$. We call m the lattice dimension and n its rank. For a lattice Λ we define its dual lattice as $\Lambda^* = \{\mathbf{y} \in \text{span}_{\mathbb{R}}(\Lambda) : \langle \mathbf{y}, \mathbf{x} \rangle \in \mathbb{Z} \ \forall \mathbf{x} \in \Lambda\}$.

For $\mathbf{A} \in R_q^{n \times m}$ we also define the q -ary kernel lattice $\Lambda_q^\perp(\mathbf{A}) = \{\tau(\mathbf{y}) \mid \mathbf{y} \in R^m : \mathbf{A}\mathbf{y} = \mathbf{0} \bmod q\}$, for $\mathbf{t} \in R_q^n$ and $\mathbf{x}_t \in R^m : \mathbf{A} \cdot \mathbf{x}_t = \mathbf{t} \bmod q$ we denote a lattice coset that maps to \mathbf{t} as $\Lambda_q^{\mathbf{t}}(\mathbf{A}) := \Lambda_q^\perp(\mathbf{A}) + \tau(\mathbf{x})$. Lastly a q -ary span lattice is defined as $\Lambda_q(\mathbf{A}^T) = \{\tau(\mathbf{y}) \mid \mathbf{y} \in R^m : \exists \mathbf{s} \in R^n \text{ s.t. } \mathbf{A}^T \mathbf{s} = \mathbf{y} \bmod q\}$.

2.4 Gaussians

Let $\varsigma > 0$ and $\mathbf{x}, \mathbf{c} \in \mathbb{R}^n$. We define the Gaussian mass function as

$$\rho_{\varsigma, \mathbf{c}}(\mathbf{x}) = \exp(-\pi \cdot \|\mathbf{x} - \mathbf{c}\|^2 / \varsigma^2).$$

The Discrete Gaussian distribution over a full-rank lattice $\Lambda \subset \mathbb{R}^n$ with centre $\mathbf{c} \in \mathbb{R}^n$ and parameter $\varsigma > 0$ is defined as

$$\forall \mathbf{x} \in \Lambda : \mathcal{D}_{\Lambda, \varsigma, \mathbf{c}}(\mathbf{x}) = \frac{\rho_{\varsigma, \mathbf{c}}(\mathbf{x})}{\sum_{\mathbf{z} \in \Lambda} \rho_{\varsigma, \mathbf{c}}(\mathbf{z})}.$$

Lemma 3 ([Ban93, Lemma 1.5]). *Let $m > 0$ be an integer, $c > 1/\sqrt{2\pi}$, $\varsigma > 0$ and Λ be an m -dimensional lattice. Let $\mathbf{x} \leftarrow \mathcal{D}_{\Lambda, \varsigma}$. Then*

$$\Pr(\|\mathbf{x}\| > c \cdot \varsigma \cdot \sqrt{m}) \leq C^m$$

where $C := c \cdot \sqrt{2\pi e} \cdot e^{-\pi c^2} < 1$. In particular for $c = 1$: $\Pr(\|\mathbf{x}\| > \varsigma \cdot \sqrt{m}) \leq 2^{-2m}$.

Lemma 4 (Adapted [Pei08, Corollary 5.3]). *Let Λ be an m -dimensional lattice. Then for any $\varsigma > 0$ it holds that*

$$\Pr(\|\mathbf{x}\|_{\infty} > \varsigma \cdot k \mid \mathbf{x} \leftarrow \mathcal{D}_{\Lambda, \varsigma}) \leq 2m \cdot e^{-\pi k^2}$$

Lemma 5 (Implicit in [MR07, Lemma 4.4]). *Let $n \geq 1$ and Λ be a full-rank n -dimensional lattice. If $\varsigma > \eta_{\varepsilon}(\Lambda)$, then $\rho_{\varsigma}(\Lambda) \geq \frac{\varsigma^n}{\det(\Lambda)} \cdot (1 - \varepsilon)$.*

For $\varepsilon > 0$ and a lattice Λ , the smoothing parameter $\eta_{\varepsilon}(\Lambda)$, first defined by Micciancio and Regev [MR07], is the smallest $\varsigma > 0$ such that $\rho_{1/\varsigma}(\Lambda^*) \leq 1 + \varepsilon$.

Lemma 6 ([GPV08, Lemma 3.1]). *For $\varepsilon > 0$ and a full-rank basis $\mathbf{B} \in \mathbb{R}^{n \times n}$:*

$$\eta_{\varepsilon}(\Lambda(\mathbf{B})) \leq \|\mathbf{B}\|_{GS} \cdot \sqrt{\ln(2n(1 + 1/\varepsilon)/\pi)}$$

Lemma 7 (Adapted [GPV08, Lemma 5.2 (eprint)]). *Let $0 < n < m$ and q be integers and let $\varsigma, \varepsilon > 0$. Sample $\bar{\mathbf{A}} \leftarrow R_q^{n \times m-n}$ and set $\mathbf{A} = [\mathbf{I} \mid \bar{\mathbf{A}}]$. Assume that $\varsigma > \eta_{\varepsilon}(\Lambda_q^{\perp}(\mathbf{A}))$. Then for $\mathbf{e} \leftarrow \mathcal{D}_{R^m, \varsigma}$ the distribution of $\mathbf{u} = \mathbf{A} \cdot \mathbf{e} \bmod q$ is within statistical distance 2ε of uniform. Moreover, for fixed $\mathbf{u} \in R_q^n$ and $\mathbf{A} \in R_q^{n \times m}$. We have*

$$\{\mathbf{e} \leftarrow \mathcal{D}_{R^m, \varsigma} \mid \mathbf{A} \cdot \mathbf{e} = \mathbf{u} \bmod q\} = \mathcal{D}_{\Lambda_q^{\perp}(\mathbf{A}), \varsigma}.$$

Lemma 8 (Explicit form of [LPR13, Theorem 7.4 (eprint)]). *Let $\mathbf{F} = \begin{bmatrix} 1 & a & b \end{bmatrix}$ where $(a, b) \leftarrow R_q^2$, $q \geq 2$ and $\delta > 0$ s.t. $2^{2\delta} \geq 2^{-d+1} + 6q^2$. Then with probability at least $1 - 2^{-d+\delta}$:*

$$\eta_{2^{-d+\delta}}(\Lambda_q^{\perp}(\mathbf{F})) \leq 2d \cdot q^{1/3+1/3d}.$$

Lemma 9 (Implicit in [CHKP10]). Let $\varepsilon \in (0, 1/2)$, and $n, m_0, m_1 > 0$ be integers such that $n < m_0 + m_1$. Let $\mathbf{F} = [\mathbf{F}_0 \mid \mathbf{F}_1] \in R_q^{n \times (m_0 + m_1)}$, $\mathbf{t} \in R_q^n$. Let $\Sigma_0 \in \mathbb{R}^{dm_0 \times dm_0}$, $\Sigma_1 \in \mathbb{R}^{dm_1 \times dm_1}$ be positive definite such that $s_{\min}(\Sigma_1) \geq \eta_\varepsilon(\mathbf{F}_1)$. Define $\Sigma := \begin{bmatrix} \Sigma_0 & \mathbf{0} \\ \mathbf{0} & \Sigma_1 \end{bmatrix}$. Then the following distributions are 2ε -statistically close:

$$\begin{aligned} \mathcal{D}_0 &= \left\{ (\mathbf{z}, \mathbf{x}) \leftarrow \mathcal{D}_{\Lambda_q^{\mathbf{t}}(\mathbf{F}), \Sigma} \right\} \\ \mathcal{D}_1 &= \left\{ (\mathbf{z}, \mathbf{x}) \mid \mathbf{z} \leftarrow \mathcal{D}_{R^3, \Sigma_0}, \mathbf{x} \leftarrow \mathcal{D}_{\Lambda_q^{\mathbf{t} - \mathbf{F}_0 \mathbf{z}}(\mathbf{F}_1), \Sigma_1} \right\} \end{aligned}$$

Lemma 10 (Adapted [MP12, Theorem. 5.5 (eprint)]). Let $0 < n < m$ be integers, $\varepsilon \in (0, 1/2)$ and q be a prime number. Let $\mathbf{A}_0 \in R_q^{n \times m}$, $\mathbf{H} \in (R_q^{n \times n})^*$ and $\mathbf{R} \in R_q^{m \times m}$. Let $\mathbf{G} \in R_q^{n \times m}$ be the gadget matrix, $\mathbf{F} = [\mathbf{A}_0 \mid -\mathbf{A}_0 \cdot \mathbf{R} + \mathbf{H} \cdot \mathbf{G}]$ and $\mathbf{t} \in R_q^n$. And

$$\varsigma^2 > 27 \cdot (\|\mathbf{R}\|^2 + 1) \cdot \eta_\varepsilon(R^m).$$

Then there exists an efficient algorithm we denote $\text{SamplePreMP}(\mathbf{A}_0, \mathbf{R}, \mathbf{H}, \mathbf{G}, \mathbf{t}, \varsigma)$ that returns \mathbf{x} such that

$$\text{SD}(\mathbf{x}, \mathcal{D}_{\Lambda_q^{\mathbf{t}}(\mathbf{F}), \varsigma}) \leq 4\varepsilon.$$

Lemma 11 ([GPV08, Theorem 4.1 (eprint)]). Let $0 < n < m$ be integers, $\varepsilon \in (0, 1/2)$ and q be a prime number. Let $\mathbf{A} \in R_q^{n \times m}$ and $\mathbf{T}_A \in \mathbb{Z}^{dm \times dm}$ s.t. $\text{rot}(\mathbf{A}) \cdot \mathbf{T}_A = \mathbf{0} \bmod q$. Let $\mathbf{t} \in R_q^n$. And

$$\varsigma > \|\mathbf{T}_A\|_{GS} \cdot \eta_\varepsilon(R^m).$$

Then there exists an efficient algorithm we denote $\text{SamplePre}(\mathbf{A}, \mathbf{T}_A, \mathbf{t}, \varsigma)$ that returns \mathbf{x} such that

$$\text{SD}(\mathbf{x}, \mathcal{D}_{\Lambda_q^{\mathbf{t}}(\mathbf{F}), \varsigma}) \leq O(\varepsilon).$$

Lemma 12 (Adapted [KY16, Lemma 1]). Let $m > 0$ be an integer, $\varepsilon \in (0, 1/2)$ and q be a prime number. Let $\varsigma > \sqrt{2} \cdot \eta_\varepsilon(R^m)$ and $\varsigma' > 2\|\mathbf{R}\| \cdot \varsigma$. Let $\mathbf{b} \in R^m$ and $\mathbf{R} \in R^{m \times m}$ be arbitrary and let $\mathbf{x} \leftarrow \mathcal{D}_{R^m, \varsigma}$. Then there exists an efficient algorithm $\text{ReRand}(\mathbf{R}, \mathbf{b} + \mathbf{x}, \varsigma', \varsigma)$ that outputs $\mathbf{b}' = \mathbf{R}\mathbf{b} + \mathbf{x}'$ where $\text{SD}(\mathbf{x}', \mathcal{D}_{R^m, \varsigma'}) \leq 6\varepsilon$.

2.5 Hardness Assumptions

We recall the Module-NTRU assumption by Chuengsatiansup et al [CPS⁺20]. While [CPS⁺20] did not formally define it, Definition 9 formalizes a number of assumptions made in [CPS⁺20].

Definition 9 (Module-NTRU, adapted from [CPS⁺20]). Let $\text{GS_SLACK} \geq 1$. Let \mathcal{D} be the distribution over $R^m \times (R^{m \times m})^\times$ defined as follows: $\forall 1 \leq i \leq m$, the i -th row of $[\mathbf{g} \mid \mathbf{F}] \sim \mathcal{D}$ is sampled from a discrete Gaussian D_{R^{m+1}, ς_i} with $\varsigma_i = \frac{\text{GS_SLACK} \sqrt{2\pi}}{\sqrt{d(m+2-i)}} q^{1/(m+1)}$, conditioned on \mathbf{F} being invertible. A $\text{ModNTRU}_{R,q,m}$ instance is a vector $\mathbf{h} \in R_q^m$ s.t. $\mathbf{h} = \mathbf{F}^{-1} \cdot \mathbf{g} \bmod q$ for $(\mathbf{g}, \mathbf{F}) \leftarrow \mathcal{D}$. The $\text{ModNTRU}_{R,q,m}$ assumption stipulates:

1. A $\text{ModNTRU}_{R,q,m}$ instance is pseudorandom over R_q^m .
2. With probability $\Omega(1)$, a pair $(\mathbf{g}, \mathbf{F}) \leftarrow \mathcal{D}_0$ can be completed into a trapdoor \mathbf{T}_h for $[1 \ \mathbf{h}^t]$ that satisfies:

$$\|\mathbf{T}_h\|_{GS} \leq \text{GS_SLACK} \cdot q^{1/(m+1)}, \quad (3)$$

$$[1 \ \mathbf{h}^t] \cdot \mathbf{T}_h = \mathbf{0} \bmod q. \quad (4)$$

In particular, this implies the existence of a PPT algorithm $\text{TrapGenNTRU}(\lambda)$ that generates a pseudorandom $\text{ModNTRU}_{R,q,m}$ instance along with a trapdoor satisfying Item 2. Note that [CPS⁺20] (implicitly) assumes that the $\text{ModNTRU}_{R,q,m}$ assumption holds for $(m, \text{GS_SLACK}) \in \{(1, 1.17), (2, 1.17), (3, 1.24)\}$.

Remark 3. Combined with Lemma 11, the $\text{ModNTRU}_{R,q,m}$ assumption implies that \mathbf{T}_h is an Ajtai trapdoor for the matrix $\mathbf{H} = [1 \ \mathbf{h}^t]$ and the $\text{SamplePre}(\cdot)$ algorithm allows one to efficiently sample $\mathbf{x} \sim_{O(\varepsilon)} \mathcal{D}_{\Lambda_q^t(\mathbf{H}), \Sigma}$ for any $t \in R_q$ and Σ such that $s_{\min}(\Sigma) \geq \text{GS_SLACK} \cdot q^{1/(m+1)} \cdot \eta_\varepsilon(R^3)$.

We rely on the Hint-MLWE assumption, introduced by Kim et al. [KLSS23]. More specifically, we use a variant by Egin et al. [EEN⁺24], recalled in Definition 10. While the distribution of the \mathbf{c}_i 's in Definition 10 may seem artificial, it is a natural byproduct of the signing procedure used in [EEN⁺24], which we re-use in this work.

Definition 10 (Hint-MLWE). Let $m > 0$, $\ell \geq 0$ and q be integers, and $\varsigma, \varsigma_p > 0$. Let $\mathbf{A} \leftarrow R_q^{m \times n}$. The Hint-MLWE $_{R,n,m,q,\varsigma}^{\ell, \varsigma_p, \beta}$ problem requires the Adversary to distinguish distributions $\mathcal{D}_{\text{real}}, \mathcal{D}_{\text{ideal}}$ where $\mathbf{s} \leftarrow \mathcal{D}_{R^n, \varsigma}, \mathbf{e} \leftarrow \mathcal{D}_{R^m, \varsigma}$ and:

$$\begin{aligned} \mathcal{D}_{\text{real}} &= \left\{ (\mathbf{A}, \mathbf{b} = \mathbf{A} \cdot \mathbf{s} + \mathbf{e}, \{\mathbf{c}_i, \mathbf{z}_i\}_{i=0}^{\ell-1}) \mid \{\mathbf{c}_i, \mathbf{z}_i\}_i \leftarrow \mathcal{P}(\mathbf{s}, \mathbf{e}) \right\}, \\ \mathcal{D}_{\text{ideal}} &= \left\{ (\mathbf{A}, \mathbf{b}, \{\mathbf{c}_i, \mathbf{z}_i\}_{i=0}^{\ell-1}) \mid \mathbf{b} \leftarrow R_q^m, \{\mathbf{c}_i, \mathbf{z}_i\}_i \leftarrow \mathcal{P}(\mathbf{s}, \mathbf{e}) \right\}, \end{aligned}$$

where the distribution $\mathcal{P}(\mathbf{s}, \mathbf{e})$ is described in Figure 2a.

When the adversary receives no hints ($\ell = 0$) in Definition 10, we recover the $\text{MLWE}_{R,n,m,q,\varsigma}$ problem. We note $\text{Hint-MLWE}_{R,1,m,q,\varsigma}^{\ell, \varsigma_p, \beta} = \text{Hint-RLWE}_{R,m,q,\varsigma}^{\ell, \varsigma_p, \beta}$. When $(\ell, n) = (0, 1)$, we recover the Ring-LWE problem $\text{RLWE}_{R,m,q,\varsigma}$. In [KLSS23, EEN⁺24] the standard hardness of Hint-RLWE for $\ell \geq 1$ is proven. We recall the reduction in Lemma 13. Note that there is a factor 2 in the formula for $\frac{1}{\varsigma_0^2}$, which seems to be an artifact of the proof, as is clearly the case when taking $\ell = 0$.

Lemma 13 (Theorem 1 and Lemma 2 from [EEN⁺24]). Assume $\ell = \omega(\kappa d \log d)^2$, and let $B_{\text{hint}} = \frac{\ell d M^2}{12}$, where $M = 2 \left\lceil \frac{q-1}{2\beta} \right\rceil + 1$. Let $\varsigma_0, \varsigma, \varsigma_p > 0$ such that $\frac{1}{\varsigma_0^2} = 2 \left(\frac{1}{\varsigma^2} + \frac{B_{\text{hint}}(1+o(1))}{\varsigma_p^2} \right)$. If $\varsigma_0 \geq \sqrt{2} \eta_\epsilon(\mathbb{Z}^d)$ for $0 < \epsilon \leq 1/2$, where $\eta_\epsilon(\mathbb{Z}^d)$ is the smoothing parameter of \mathbb{Z}^d , then there is an efficient reduction from $\text{RLWE}_{R,m,q,\varsigma_0}$ to $\text{Hint-RLWE}_{R,m,q,\varsigma}^{\ell, \varsigma_p, \beta}$ that reduces the advantage by at most 4ϵ .

$\mathcal{P}(\mathbf{s}, \mathbf{e})$ <hr/> $\forall i : u_i \leftarrow R_q$ $(c_{0,i}, c_{1,i}) := \text{Decomp}_\beta(u_i)$ $(\mathbf{p}_{0,i}, \mathbf{p}_{1,i}) \leftarrow \mathcal{D}_{R^{n+m}, \varsigma_p}$ $\mathbf{z}_i = \begin{bmatrix} \mathbf{s} \\ \mathbf{e} \end{bmatrix} \cdot c_{0,i} + \mathbf{p}_i$ $\text{return } \{(\mathbf{c}_i = (c_{0,i}, c_{1,i}), \mathbf{z}_i)\}_i$	$\mathcal{P}_{\text{cond}}(\mathbf{S}, \mathbf{E})$ <hr/> $\forall i : \mathbf{u}_i \leftarrow R_q^m$ $(\mathbf{c}_{0,i}, \mathbf{c}_{1,i}) := \text{Decomp}_\beta(\mathbf{u}_i)$ $(\mathbf{p}_i, \mathbf{x}_i) \leftarrow \mathcal{D}_{A_q^{\mathbf{r}_i - \mathbf{A}_0(\mathbf{B}) \cdot \mathbf{T} \cdot \mathbf{c}_i}([\mathbf{A}_0(\mathbf{B}) \mathbf{A}_2]), \varsigma_p}$ $\mathbf{z}_i = \mathbf{T} \cdot \mathbf{c}_i + \mathbf{p}_i$ $\text{return } \{\mathbf{c}_i = (\mathbf{c}_{0,i}, \mathbf{c}_{1,i}), (\mathbf{z}_i, \mathbf{x}_i)\}_i$
--	--

(a) Hint distribution for Hint-MLWE (b) Hint distribution for Coset-Hint-MLWE

Fig. 2: Hint distributions

3 The Coset-Hint-MLWE Assumption

In this section, we formally define the Coset-Hint-MLWE assumption. Intuitively, the hints component $\mathbf{p}_{j,i}$ in this new problem is sampled from a lattice coset dependent on \mathbf{c}_i , as opposed to the whole ring. One might think that this leaks more information about the MLWE secret, we prove that this is not the case by constructing reductions from the Coset-Hint-MLWE to standard assumptions.

Definition 11 (Coset-Hint-MLWE Assumption). Let $0 < n \leq m \leq m'$ and $\ell \geq 0$ be integers, q an integer and $\varsigma, \varsigma_p > 0$. Let $\mathbf{A} \leftarrow R_q^{m \times n}$ and $\mathbf{A}_2 \leftarrow R_q^{m \times m'}$. Let $\{\mathbf{r}_i\}_{i=0}^{\ell-1}$ be arbitrary vectors in R_q^m . The Coset-Hint-MLWE $_{R,n,m,m',q,\varsigma}^{\ell, \{\mathbf{r}_i\}, \varsigma_p, \beta}$ problem requires the Adversary to distinguish $\mathcal{D}_{\text{real}}$ and $\mathcal{D}_{\text{ideal}}$, where $\mathbf{S} \leftarrow \mathcal{D}_{R^{n \times m}, \varsigma}$, $\mathbf{E} \leftarrow \mathcal{D}_{R^{m \times m}, \varsigma}$ and:

$$\mathcal{D}_{\text{real}} = \left\{ (\mathbf{A}, \mathbf{B} = \mathbf{A} \cdot \mathbf{S} + \mathbf{E}, \mathbf{A}_2, \{\mathbf{c}_i, (\mathbf{z}_i, \mathbf{x}_i)\}_{i=0}^{\ell-1}) \mid \{\mathbf{c}_i, (\mathbf{z}_i, \mathbf{x}_i)\}_i \leftarrow \mathcal{P}_{\text{cond}}(\mathbf{S}, \mathbf{E}) \right\},$$

$$\mathcal{D}_{\text{ideal}} = \left\{ (\mathbf{A}, \mathbf{B}, \mathbf{A}_2, \{\mathbf{c}_i, (\mathbf{z}_i, \mathbf{x}_i)\}_{i=0}^{\ell-1}) \mid \mathbf{B} \leftarrow R_q^{m \times m}, \{\mathbf{c}_i, (\mathbf{z}_i, \mathbf{x}_i)\}_i \leftarrow \mathcal{P}_{\text{cond}}(\mathbf{S}, \mathbf{E}) \right\},$$

where $\mathbf{A}_0(\mathbf{B}) = [\mathbf{I}_m \mid \mathbf{A} \mid \beta \cdot \mathbf{I}_m - \mathbf{B}]$, $\mathbf{T}^T = \begin{bmatrix} \mathbf{E}^T & \mathbf{S}^T & \mathbf{I}_m \\ \mathbf{I}_m & \mathbf{0} & \mathbf{0} \end{bmatrix}$, and the distribution $\mathcal{P}_{\text{cond}}(\mathbf{S}, \mathbf{E})$ is described in Figure 2b.

Remark 4. In fact, to construct the matrix $\mathbf{A}_0(\mathbf{B})$ we require a multi-secret version of MLWE with $\mathbf{S} \in R^{n \times m}$ and $\mathbf{E} \in R^{m \times m}$. In our TKEM construction we use $n = m = 1$.

Remark 5. Note that the hint distributions of the two problems differ only in the sampling of \mathbf{p}_i . Now if we sample vectors $\{\mathbf{r}_i\}_i$ in Coset-Hint-MLWE uniformly at random we get distributions

$$\mathcal{D}_p := \left\{ \{(\mathbf{p}_i, \mathbf{x}_i), \mathbf{r}_i\}_{i \in [\ell]} \mid \mathbf{r}_i \leftarrow R_q^n, (\mathbf{p}_i, \mathbf{x}_i) \leftarrow \mathcal{D}_{A_q^{\mathbf{r}_i}([\mathbf{A}_0(\mathbf{B})|\mathbf{A}_2]), \varsigma_p} \right\}$$

$$\mathcal{D}'_p := \left\{ \{(\mathbf{p}_i, \mathbf{x}_i), \mathbf{r}_i\}_{i \in [\ell]} \mid (\mathbf{p}_i, \mathbf{x}_i) \leftarrow \mathcal{D}_{R^{2m+n+m'}, \varsigma_p}, \mathbf{r}_i = [\mathbf{A}_0(\mathbf{B})|\mathbf{A}_2] \cdot \begin{pmatrix} \mathbf{p}_i \\ \mathbf{x}_i \end{pmatrix} \right\}$$

for \mathbf{p}_i in $\mathcal{P}_{\text{cond}}$ and \mathcal{P} accordingly. If $\varsigma_p \geq \eta_\varepsilon(\Lambda_q^\perp([\mathbf{A}_0(\mathbf{B}) \mid \mathbf{A}_2]))$, by Leftover Hash Lemma in Lemma 7 the distributions \mathcal{D}_p and \mathcal{D}'_p are statistically close. Hence, we recover the usual Hint-MLWE assumption from Coset-Hint-MLWE up to a negligible statistical distance when $\{\mathbf{r}_i\}_i$ are uniform⁵.

When $\{\mathbf{r}_i\}_i$ are fixed as in Definition 11 we need a more complex argument to reduce Coset-Hint-MLWE to standard assumptions.

Theorem 1 (Hint-MLWE + Trapdoor Sampling \leq Coset-Hint-MLWE).

Let $0 < n \leq m \leq m'$, $\ell \geq 0$ and q be integers. Let $\varsigma, B > 0$ and $\varepsilon \in (0, 1/2)$. Assume that there exists a PPT algorithm $(\mathbf{A}_2, \mathbf{T}_{\mathbf{A}_2}) \leftarrow \text{TrapGen}(1^\lambda)$ such that $\mathbf{A}_2 \in R_q^{m \times m'}$, $\mathbf{T}_{\mathbf{A}_2} \in R^{m' \times m'}$ is full-rank, $\mathbf{A}_2 \cdot \mathbf{T}_{\mathbf{A}_2} = \mathbf{0} \bmod q$ and $\|\mathbf{T}_{\mathbf{A}_2}\|_{GS} \leq B$. Let $\varsigma_p > B \cdot \eta_\varepsilon(R^{m'})$. For any PPT adversary \mathcal{A} :

$$\text{Adv}_{\mathcal{A}}(\mathcal{D}_{\text{real}}, \mathcal{D}_{\text{ideal}}) \leq 2 \cdot \text{Adv}_{\mathcal{A}_2}(\mathbf{A}_2, \mathcal{U}(R_q^{m \times m'})) + m \cdot \text{Adv}_{\mathcal{A}_3}^{\text{Hint-MLWE}_{R, n, m, q, \varsigma}^{\ell, \varsigma_p, \beta}} + \ell O(\varepsilon)$$

where algorithms \mathcal{A}_2 and \mathcal{A}_3 run in approximately the same time as \mathcal{A} .

We instantiate the statement above with the ModNTRU assumption.

Corollary 1 (Hint-MLWE + ModNTRU \leq Coset-Hint-MLWE). Let $n = m = 1$, $m' = 3$, $\ell \geq 0$, q be integers, and $\varsigma > 0$, $\varepsilon \in (0, 1/2)$. Let $\varsigma_p > 1.17 \cdot q^{1/3} \cdot \eta_\varepsilon(R^3)$. Then for a PPT adversary \mathcal{A} :

$$\text{Adv}_{\mathcal{A}}(\mathcal{D}_{\text{real}}, \mathcal{D}_{\text{ideal}}) \leq 2 \text{Adv}_{\mathcal{A}_2}^{\text{ModNTRU}_{R, q, 2}}(\kappa) + \text{Adv}_{\mathcal{A}_3}^{\text{Hint-RLWE}_{R, 1, q, \varsigma_a}^{\ell, \varsigma_p, \beta}}(\kappa) + \ell \cdot O(\varepsilon)$$

where \mathcal{A}_2 and \mathcal{A}_3 run in approximately the same time as \mathcal{A} .

Proof (of Theorem 1). We consider a sequence of six hybrids. Differences between consecutive hybrids are highlighted in green. \mathbf{T} is defined as in Definition 11.

$\text{Hyb}_0 = \mathcal{D}_{\text{real}}$: The original distribution can be written as follows.

$$\begin{aligned}
\mathcal{D}_{\text{real}} : \quad & \mathbf{A} \leftarrow R_q^{n \times m}, \mathbf{A}_2 \leftarrow R_q^{m \times m'} \\
& \mathbf{S} \leftarrow \mathcal{D}_{R^{n \times m}, \varsigma}, \mathbf{E} \leftarrow \mathcal{D}_{R^{m \times m}, \varsigma}, \mathbf{B} = \mathbf{A} \cdot \mathbf{S} + \mathbf{E} \\
& \mathbf{A}_0 = \mathbf{A}_0(\mathbf{B}) := [\mathbf{I}_m \mid \mathbf{A} \mid \beta \cdot \mathbf{I}_m - \mathbf{B}] \\
& \mathbf{F} := [\mathbf{A}_0 \mid \mathbf{A}_2] \\
\forall i : \quad & \mathbf{u}_i \leftarrow R_q^m, (\mathbf{c}_{0,i}, \mathbf{c}_{1,i}) := \text{Decomp}_\beta(\mathbf{u}_i) \\
& (\mathbf{p}_i, \mathbf{x}_i) \leftarrow \mathcal{D}_{A_q^{\mathbf{r}_i - \mathbf{A}_0 \cdot \mathbf{T} \cdot \mathbf{c}_i}(\mathbf{F}), \varsigma_p} \\
& \mathbf{z}_i = \mathbf{T} \cdot \mathbf{c}_i + \mathbf{p}_i
\end{aligned}$$

⁵ The values $\mathbf{A}_2, \{\mathbf{x}_i\}$ do not contain any additional information on the secret anymore and can be dropped from $\mathcal{D}_{\text{real}}$ and $\mathcal{D}_{\text{ideal}}$.

Hyb₁ : In the first hybrid we introduce the trapdoor for \mathbf{A}_2 . By definition:

$$\text{Adv}(\mathcal{D}_{\text{real}}, \text{Hyb}_1) \leq \text{Adv}(\mathbf{A}_2, \mathcal{U}(R_q^{m \times m'})).$$

$\begin{aligned} \text{Hyb}_1 : \quad & \mathbf{A} \leftarrow R_q^{n \times m}, (\mathbf{A}_2, \mathbf{T}_{\mathbf{A}_2}) \leftarrow \text{TrapGen}(1^\lambda) \\ & \mathbf{S} \leftarrow \mathcal{D}_{R^{n \times m}, \varsigma}, \mathbf{E} \leftarrow \mathcal{D}_{R^{m \times m}, \varsigma}, \mathbf{B} = \mathbf{A} \cdot \mathbf{S} + \mathbf{E} \\ & \mathbf{A}_0 = \mathbf{A}_0(\mathbf{B}) := [\mathbf{I}_m \mid \mathbf{A} \mid \beta \cdot \mathbf{I}_m - \mathbf{B}] \\ & \mathbf{F} := [\mathbf{A}_0 \mid \mathbf{A}_2] \\ \forall i : \quad & \mathbf{u}_i \leftarrow R_q^m, (\mathbf{c}_{0,i}, \mathbf{c}_{1,i}) := \text{Decomp}_\beta(\mathbf{u}_i) \\ & (\mathbf{p}_i, \mathbf{x}_i) \leftarrow \mathcal{D}_{\Lambda_q^{\mathbf{r}_i - \mathbf{A}_0 \cdot \mathbf{T} \cdot \mathbf{c}_i}(\mathbf{F}), \varsigma_p} \\ & \mathbf{z}_i = \mathbf{T} \cdot \mathbf{c}_i + \mathbf{p}_i \end{aligned}$

Hyb₂ : We now sample \mathbf{p}_i as a spherical Gaussian. Since $\varsigma_p > \|\mathbf{T}_{\mathbf{A}_2}\|_{GS} \cdot \eta_\varepsilon(R^{m'})$, it follows from Lemma 6 and Lemma 9 that:

$$\text{SD}(\text{Hyb}_1, \text{Hyb}_2) \leq 2\ell\varepsilon.$$

$\begin{aligned} \text{Hyb}_2 : \quad & \mathbf{A} \leftarrow R_q^{n \times m}, (\mathbf{A}_2, \mathbf{T}_{\mathbf{A}_2}) \leftarrow \text{TrapGen}(1^\lambda) \\ & \mathbf{S} \leftarrow \mathcal{D}_{R^{n \times m}, \varsigma}, \mathbf{E} \leftarrow \mathcal{D}_{R^{m \times m}, \varsigma}, \mathbf{B} = \mathbf{A} \cdot \mathbf{S} + \mathbf{E} \\ & \mathbf{A}_0 = \mathbf{A}_0(\mathbf{B}) := [\mathbf{I}_m \mid \mathbf{A} \mid \beta \cdot \mathbf{I}_m - \mathbf{B}] \\ & \mathbf{F} := [\mathbf{A}_0 \mid \mathbf{A}_2] \\ \forall i : \quad & \mathbf{u}_i \leftarrow R_q^m, (\mathbf{c}_{0,i}, \mathbf{c}_{1,i}) := \text{Decomp}_\beta(\mathbf{u}_i) \\ & \mathbf{p}_i \leftarrow \mathcal{D}_{R^{2m+n}, \varsigma_p}, \mathbf{x}_i \leftarrow \mathcal{D}_{\Lambda_q^{\mathbf{r}_i - \mathbf{A}_0 \cdot (\mathbf{T} \cdot \mathbf{c}_i - \mathbf{p}_i)}(\mathbf{A}_2), \varsigma_p} \\ & \mathbf{z}_i = \mathbf{T} \cdot \mathbf{c}_i + \mathbf{p}_i \end{aligned}$

Hyb₃ : We now sample the \mathbf{x}_i using the trapdoor for \mathbf{A}_2 . Note that now all random variables in Hyb₃ can be sampled in polynomial time. We have:

$$\text{SD}(\text{Hyb}_2, \text{Hyb}_3) \leq 2\ell O(\varepsilon).$$

$\begin{aligned} \text{Hyb}_3 : \quad & \mathbf{A} \leftarrow R_q^{n \times m}, (\mathbf{A}_2, \mathbf{T}_{\mathbf{A}_2}) \leftarrow \text{TrapGen}(1^\lambda) \\ & \mathbf{S} \leftarrow \mathcal{D}_{R^{n \times m}, \varsigma}, \mathbf{E} \leftarrow \mathcal{D}_{R^{m \times m}, \varsigma}, \mathbf{B} = \mathbf{A} \cdot \mathbf{S} + \mathbf{E} \\ & \mathbf{A}_0 = \mathbf{A}_0(\mathbf{B}) := [\mathbf{I}_m \mid \mathbf{A} \mid \beta \cdot \mathbf{I}_m - \mathbf{B}] \\ & \mathbf{F} := [\mathbf{A}_0 \mid \mathbf{A}_2] \\ \forall i : \quad & \mathbf{u}_i \leftarrow R_q^m, (\mathbf{c}_{0,i}, \mathbf{c}_{1,i}) := \text{Decomp}_\beta(\mathbf{u}_i) \\ & \mathbf{p}_i \leftarrow \mathcal{D}_{R^{2m+n}, \varsigma_p}, \mathbf{x}_i \leftarrow \text{SamplePre}(\mathbf{A}_2, \mathbf{T}_{\mathbf{A}_2}, \mathbf{r}_i - \mathbf{A}_0 \cdot (\mathbf{T}\mathbf{c}_i + \mathbf{p}_i), \varsigma_p) \\ & \mathbf{z}_i = \mathbf{T} \cdot \mathbf{c}_i + \mathbf{p}_i \end{aligned}$

Hyb₄ : Between Hyb₃ and Hyb₄ we invoke $\text{Hint-MLWE}_{R,n,m,q,\varsigma}^{\ell, \varsigma_p, \beta}$ m times.

$$\begin{array}{l}
\text{Hyb}_4 : \quad \mathbf{A} \leftarrow R_q^{n \times m}, (\mathbf{A}_2, \mathbf{T}_{\mathbf{A}_2}) \leftarrow \text{TrapGen}(1^\lambda) \\
\quad \mathbf{B} \leftarrow R_q^{m \times m}, \mathbf{S} \leftarrow \mathcal{D}_{R^{n \times m}, \varsigma}, \mathbf{E} \leftarrow \mathcal{D}_{R^{m \times m}, \varsigma} \\
\quad \mathbf{A}_0 = \mathbf{A}_0(\mathbf{B}) := [\mathbf{I}_m \mid \mathbf{A} \mid \beta \cdot \mathbf{I}_m - \mathbf{B}] \\
\quad \mathbf{F} := [\mathbf{A}_0 \mid \mathbf{A}_2] \\
\forall i : \quad \mathbf{u}_i \leftarrow R_q^m, (\mathbf{c}_{0,i}, \mathbf{c}_{1,i}) := \text{Decomp}_\beta(\mathbf{u}_i) \\
\quad \mathbf{p}_i \leftarrow \mathcal{D}_{R^{2m+n}, \varsigma_p}, \mathbf{x}_i \leftarrow \text{SamplePre}(\mathbf{A}_2, \mathbf{T}_{\mathbf{A}_2}, \mathbf{r}_i - \mathbf{A}_0 \cdot (\mathbf{T}\mathbf{c}_i + \mathbf{p}_i), \varsigma_p) \\
\quad \mathbf{z}_i = \mathbf{T} \cdot \mathbf{c}_i + \mathbf{p}_i
\end{array}$$

For the j -th column of matrix \mathbf{B} , $j \in [m]$ we define the following transform. Before the transformation the j -th column is equal to $\mathbf{b}^{(j)} = \mathbf{A} \cdot \mathbf{s}^{(j)} + \mathbf{e}^{(j)}$ where $\mathbf{s}^{(j)}$ and $\mathbf{e}^{(j)}$ are the j -th columns of \mathbf{S} and \mathbf{E} accordingly. After the transformation $\mathbf{b}^{(j)} \leftarrow R_q^m$ and the remaining values stay unchanged.

This jump can be reduced to solving a $\text{Hint-MLWE}_{R,n,m,q,\varsigma}^{\ell,\varsigma_p,\beta}$ instance. Assume that $\{\tilde{\mathbf{A}}, \tilde{\mathbf{b}}^{(j)}, \{(\tilde{\mathbf{c}}_i^j = (\tilde{c}_{0,i}^{(j)}, \tilde{c}_{1,i}^{(j)}), \tilde{\mathbf{z}}_i^{(j)} \in R^2)\}_i\}$ is the Hint-MLWE instance received from a challenger. We set $\mathbf{A} := \tilde{\mathbf{A}}, \mathbf{b}^{(j)} := \tilde{\mathbf{b}}^{(j)}$, the first coordinates of $\mathbf{c}_{0,i}$ and $\mathbf{c}_{1,i}$ are now equal to $\tilde{c}_{0,i}^{(j)}$ and $\tilde{c}_{1,i}^{(j)}$ accordingly the remaining coordinates are sampled as before. For the matrix \mathbf{T} we set the j -th column to $\mathbf{0}$ and sample all other columns as before. Then $\mathbf{z}_i := \mathbf{T} \cdot \mathbf{c}_i + \tilde{\mathbf{z}}_i^{(j)} + (\mathbf{0}^{m+n} \parallel \mathbf{p}_2)$ with $\mathbf{p}_2 \leftarrow \mathcal{D}_{R^m, \varsigma_p}$. Following the definition of the Hint-MLWE problem if $\tilde{\mathbf{b}}^{(j)}$ is an LWE sample we have the distribution before the transformation and if $\tilde{\mathbf{b}}^{(j)}$ is uniform it is the distribution after. After m such transformation for $j \in [m]$ we get Hyb_4 distribution.

$$\text{Adv}(\text{Hyb}_3, \text{Hyb}_4) \leq m \cdot \text{Hint-MLWE}_{R,n,m,q,\varsigma}^{\ell,\varsigma_p,\beta}.$$

Hyb_5 : Lastly, we remove the trapdoor from \mathbf{A}_2 , the simulation of the Gaussian distribution and the split sampling to get back to $\text{Hyb}_5 = \mathcal{D}_{\text{ideal}}$.

$$\begin{array}{l}
\text{Hyb}_5 : \quad \mathbf{A} \leftarrow R_q^{n \times m}, \mathbf{A}_2 \leftarrow R_q^{m \times m'} \\
\quad \mathbf{B} \leftarrow R_q^{m \times m}, \mathbf{S} \leftarrow \mathcal{D}_{R^{n \times m}, \varsigma}, \mathbf{E} \leftarrow \mathcal{D}_{R^{m \times m}, \varsigma} \\
\quad \mathbf{A}_0 = \mathbf{A}_0(\mathbf{B}) := [\mathbf{I}_m \mid \mathbf{A} \mid \beta \cdot \mathbf{I}_m - \mathbf{B}] \\
\quad \mathbf{F} := [\mathbf{A}_0 \mid \mathbf{A}_2] \\
\forall i : \quad \mathbf{u}_i \leftarrow R_q^m, (\mathbf{c}_{0,i}, \mathbf{c}_{1,i}) := \text{Decomp}_\beta(\mathbf{u}_i), \alpha_i := \mathbf{A}_0 \mathbf{T} \mathbf{c}_i \\
\quad (\mathbf{p}_i, \mathbf{x}_i) \leftarrow \mathcal{D}_{\Lambda_q^{r_i - \alpha_i}([\mathbf{A}_0(\mathbf{B}) \mid \mathbf{A}_2]), \varsigma_p} \\
\quad \mathbf{z}_i = \mathbf{T} \cdot \mathbf{c}_i + \mathbf{p}_i
\end{array}$$

We have: $\text{Adv}(\text{Hyb}_4, \text{Hyb}_5) \leq \text{Adv}(\mathbf{A}_2, \mathcal{U}(R_q^{m \times m'})) + 2\ell O(\varepsilon) + 2\ell \varepsilon.$ \square

4 Our TKEM Construction

In this section we describe the modified BCHK transform and its required building blocks. Our TIBE builds on the Agrawal-Boneh-Boyen IBE [ABB10] and on the Plover signature [EEN⁺24]. For the one-time signature we take WOTS⁺

defined in Section A. The security proof of the TIBE uses techniques from del Pino et al. [DKM⁺24] and improvements by Katsumata, Reichle and Takemure [KRT24].

4.1 The BCHK+ Transform

We describe the BCHK+ transform in Figure 3 and highlight the tweaks we make to the transform of [BBH06]. The tweaks are based on the Fujisaki-Okamoto transform [FO13] and its variants by Hofheinz, Hövelmanns and Kiltz [HHK17]. They are applied to guarantee the Decapsulation Consistency of the scheme, the CCA security also holds without them.

TKEM.Keygen (1^κ) <hr/> $(\{dk_i\}_{i \in [N]}, ek) \leftarrow \text{TIBE.Setup}(1^\kappa)$ return $(\{dk_i\}_{i \in [N]}, ek)$	TKEM.ShareDecaps $_j(dk_i, (ct, vk, sig)), j \leq r$ <hr/> assert $\text{SIG.Verify}((vk, ct, sig)) = 1$ $contrib_{i,j} = \text{TIBE.ShareExtract}_j(dk_i, id = vk)$ return $contrib_{i,j}$
TKEM.Encaps (ek) <hr/> $msg \leftarrow \{0, 1\}^d$ $rand = G_{fo}(msg)$ $(sk, vk) \leftarrow \text{SIG.Keygen}(1^\kappa)$ $ct = \text{TIBE.Encrypt}(ek, id = vk, msg; rand)$ $sig \leftarrow \text{SIG.Sign}(sk, ct)$ $K = H_{fo}(msg ct)$ return $(ct, vk, sig), K$	TKEM.Combine ($\{contrib_{i,r}\}_{i \in \text{act}}, (ct, vk, sig)$) <hr/> assert $\text{SIG.Verify}((vk, ct, sig)) = 1$ $msg = \text{TIBE.Combine}(ct, vk, \text{act}, \{contrib_{i,r}\}_i)$ $rand = G_{fo}(msg)$ assert $ct = \text{TIBE.Encrypt}(ek, vk, msg; rand)$ $K = H_{fo}(msg ct)$ return K

Fig. 3: Our BCHK+ transform. The major differences with the original BCHK transform [CHK04, BCHK07, BBH06] are highlighted in green.

Theorem 2 (Adapted [BBH06, Theorem 2]). *Let TIBE be a threshold identity-based encryption scheme satisfying Definition 5. Let \mathcal{M} and \mathcal{R} denote the message space and randomness space of the TIBE, respectively. Let SIG be a one-time signature satisfying $\varepsilon_{\text{SIG-sEU-CMA}}$ and $G_{fo} : \{0, 1\}^* \rightarrow \mathcal{R}$ and $H_{fo} : \{0, 1\}^* \rightarrow \mathcal{K}$ for some set \mathcal{K} modelled as random oracles. Then TKEM as described in Figure 3 is a CCA-secure TKEM with:*

$$\text{Adv}_{\mathcal{A}, \text{TKEM}}^{\text{CCA2}}(\kappa) \leq \varepsilon_{\text{SIG}} + (Q_{G_{fo}} + Q_{H_{fo}}) \cdot \text{Adv}_{\mathcal{A}, \text{TIBE}}^{\text{sel-ID-IND}} + \frac{(Q_{G_{fo}} + Q_{H_{fo}})}{|\mathcal{M}|}.$$

Proof. First, we modify the CCA experiment in Figure 1c to add an abort condition. We denote the new adversary's advantage as ε_1 . Here the encryption randomness $\text{rand}^* = G_{f_0}(\text{msg}^*)$ in the challenge ciphertext is sampled as $\text{rand}^* \leftarrow \mathcal{R}$ instead. These distributions are exactly the same as long as msg^* has not been queried to the G_{f_0} random oracle. Given ct^* the probability of querying msg^* to the G_{f_0} oracle is equivalent to breaking sel-ID-OW security of the underlying TIBE. Hence, the advantage is bounded as

$$\text{Adv}_{\mathcal{A}, \text{TKEM}}^{\text{CCA2}}(\kappa) \leq \varepsilon_1 + Q_{G_{f_0}} \cdot \text{Adv}_{\mathcal{A}, \text{TIBE}}^{\text{sel-ID-OW}}.$$

We claim that $\varepsilon_1 \leq \varepsilon_{\text{SIG}} + Q_{H_{f_0}} \cdot \text{Adv}_{\mathcal{A}, \text{TIBE}}^{\text{sel-ID-OW}}$. Using \mathcal{A} we build an adversary \mathcal{B} against TIBE. For the random oracles they are simulated by \mathcal{B} that samples the answers at random and maintains a list of queried inputs. In the experiment \mathcal{B} first samples $(\text{sk}^*, \text{vk}^*) \leftarrow \text{SIG.Keygen}(1^\kappa)$ and sets vk^* as the challenge identity. Then \mathcal{B} relays the messages between \mathcal{A} and the Challenger on the corrupt parties and their keys with no changes. For every decryption query $(\text{ct}, \text{vk}, \text{sig})$, act of \mathcal{A} if $\text{SIG.Verify}(\text{vk}, \text{ct}, \text{sig}) = 0$ then \mathcal{B} replies with \perp to \mathcal{A} . Moreover, if a query has $\text{vk} = \text{vk}^*$, \mathcal{B} aborts the experiment. Otherwise, \mathcal{B} forwards $(\text{ct}, \text{id} = \text{vk}, \text{act})$ to the Challenger and returns the reply to \mathcal{A} .

When \mathcal{A} is ready to receive the challenge \mathcal{B} gets $\text{ct}^* = \text{TIBE.Encrypt}(\text{ek}, \text{msg}^*)$ corresponding to some $\text{msg}^* \leftarrow \mathcal{M}$ from the Challenger. Then \mathcal{B} computes $\text{ct}_{\text{TKEM}}^* = (\text{ct}^*, \text{vk}^*, \text{SIG.Sign}(\text{sk}^*, \text{ct}^*))$ and $K^* \leftarrow \mathcal{K}$ and sends $(\text{ct}_{\text{TKEM}}^*, K^*)$ to \mathcal{A} . The decryption queries after the challenge are simulated in the same way as before. Lastly, \mathcal{A} returns a bit b then \mathcal{B} checks the list of queries in H_{f_0} and forwards msg^* to the Challenger if $(\text{msg}^*, \text{ct}^*)$ was queried.

We analyse the success probability of \mathcal{B} . The decryption queries are simulated perfectly unless \mathcal{B} issues an abort. The first type of abort happens when \mathcal{A} makes a decryption query $(\text{ct}, \text{vk}, \text{sig})$ for $\text{vk} = \text{vk}^*$ either before or after the challenge. It implies that sig is a valid signature for vk . In both cases this constitutes a one-time signature forgery. Hence, the probability of this type of abort is bounded by ε_{SIG} .

If the abort condition does not occur the challenge value $\text{ct}_{\text{TKEM}}^*$ is simulated perfectly. For the value K^* the view of the Adversary is independent of whether $K^* = H_{f_0}(\text{msg}^* || \text{ct}^*)$ or $K^* \leftarrow \mathcal{K}$ as long as H_{f_0} is not queried on $\text{msg}^*, \text{ct}^*$. Hence, winning the is equivalent to querying the pair $\text{msg}^*, \text{ct}^*$ and that happens with probability $Q_{H_{f_0}} \cdot \text{Adv}_{\mathcal{A}, \text{TIBE}}^{\text{sel-ID-OW}}$. Then

$$\varepsilon_1 \leq Q_{H_{f_0}} \cdot \text{Adv}_{\mathcal{A}, \text{TIBE}}^{\text{sel-ID-OW}} + \Pr(\text{abort}) \leq Q_{H_{f_0}} \cdot \text{Adv}_{\mathcal{A}, \text{TIBE}}^{\text{sel-ID-OW}} + \varepsilon_{\text{SIG}}$$

Lastly, we use a folklore fact that IND-CPA security implies OW-CPA by sampling the pair of challenge messages at random (see e.g. [Pas24, Tutorial 8]). This technique applies to TIBE directly hence $\text{Adv}_{\mathcal{A}, \text{TIBE}}^{\text{sel-ID-OW}} \leq \text{Adv}_{\mathcal{A}, \text{TIBE}}^{\text{sel-ID-IND}} + 1/|\mathcal{M}|$. Adding all advantage changes together we get the statement. \square

We include the additional integrity checks after the reconstruction of the message in order to satisfy Decapsulation Consistency. They follow ideas of the FO transform variants described in [HHK17]. Intuitively, even though FO is difficult

to thresholdise, after we have the message in the clear, the complications of the threshold setting no longer apply.

Lemma 14 (Decapsulation Consistency). *Let $G_{\text{fo}} : \{0,1\}^* \rightarrow \{0,1\}^\kappa$ be modelled as a random oracle. Let TKEM be the scheme obtained by applying Figure 3. If the underlying TIBE is γ -spread, then for any PPT adversary \mathcal{A} :*

$$\text{Adv}_{\mathcal{A}, \text{TKEM}}^{\text{CD}} \leq \frac{Q_{G_{\text{fo}}} \cdot (Q_{G_{\text{fo}}} - 1)}{2} \cdot 2^{-\gamma}.$$

where $Q_{G_{\text{fo}}}$ is the number of queries \mathcal{A} makes to the G_{fo} random oracle.

Proof. If the Adversary wins the game in Figure 1e with a tuple $((\text{ct}, \text{vk}, \text{sig}), \text{act}, \text{act}', S, S')$. Then the decapsulation returns two valid keys $K_0 \neq K_1$. If the messages reconstructed as $\text{msg}_1 := \text{TIBE.Combine}(\text{ct}, \text{vk}, \text{act}', S')$ and $\text{msg}_0 := \text{TIBE.Combine}(\text{ct}, \text{vk}, \text{act}, S)$ are the same, by Figure 3 the corresponding keys are the same as well. Since $K_i = H_{\text{fo}}(\text{msg}_i || \text{ct})$, this implies that there exists a pair of valid messages $\text{msg}_0 \neq \text{msg}_1$ for which the following holds:

$$\begin{aligned} \text{msg}_1 &:= \text{TIBE.Combine}(\text{ct}, \text{vk}, \text{act}', S') \\ \text{msg}_0 &:= \text{TIBE.Combine}(\text{ct}, \text{vk}, \text{act}, S) \\ \text{ct} &:= \text{TIBE.Encrypt}(\text{ek}, \text{vk}, \text{msg}_0; G_{\text{fo}}(\text{msg}_0)) \\ \text{ct} &:= \text{TIBE.Encrypt}(\text{ek}, \text{vk}, \text{msg}_1; G_{\text{fo}}(\text{msg}_1)). \end{aligned}$$

In particular it implies that $\mathcal{A}^{G_{\text{fo}}}$ can compute $(\text{ct}, \text{vk}, \text{msg}_0, \text{msg}_1)$ such that

$$\begin{aligned} \text{ct} &= \text{TIBE.Encrypt}(\text{ek}, \text{vk}, \text{msg}_0; G_{\text{fo}}(\text{msg}_0)) \\ &= \text{TIBE.Encrypt}(\text{ek}, \text{vk}, \text{msg}_1; G_{\text{fo}}(\text{msg}_1)). \end{aligned} \tag{5}$$

We simplify the winning condition of the Adversary to only require a tuple $(\text{ct}, \text{vk}, \text{msg}_0, \text{msg}_1)$ that satisfies Equation (5), this can only increase their advantage. For a given tuple $(\text{ct}, \text{vk}, \text{msg}_0, \text{msg}_1)$, we consider two possible cases:

1. If either of the messages (e.g. msg_0) has not been queried to the G_{fo} oracle the winning probability is bounded as

$$\Pr(\text{TIBE.Encrypt}(\text{ek}, \text{vk}, \text{msg}_0; G_{\text{fo}}(\text{msg}_0)) = \text{ct} \mid G_{\text{fo}}(\text{msg}_0) \leftarrow \{0,1\}^\kappa) \leq 2^{-\gamma}.$$

2. Otherwise, if both messages were queried before the probability that at least two of them have colliding ciphertexts satisfying Equation (5) is bounded by

$$\Pr(\exists i \neq j \in [Q_{G_{\text{fo}}}] \text{ s.t. } \text{ct}_i = \text{ct}_j) \leq \sum_{i \in [Q_{G_{\text{fo}}}] i \cdot 2^{-\gamma} = \frac{Q_{G_{\text{fo}}} \cdot (Q_{G_{\text{fo}}} - 1)}{2} \cdot 2^{-\gamma}.$$

The Adversary's advantage is bounded by the maximum of both probabilities. \square

4.2 Our TIBE construction

Let $H_{\text{cmt}}(\cdot) : R_q \rightarrow \{0,1\}^{2\kappa}$ and $H_{\text{mask}}(\cdot) : \{0,1\}^* \rightarrow R_q^2$ be hash functions modelled as random oracles.

Public parameters. These include $(\mathbf{A}_0, \mathbf{A}_1, \mathbf{A}_2, \mathbf{G}, r, N, T)$, where N is the overall number of parties in the protocol and T is the threshold on the number of parties required for decryption. We explain how each matrix is used in more detail.

1. \mathbf{A}_0 is a matrix with trapdoor *à la* Eagle [YJW23] or Plover [EEN+24], multiplied by an invertible element of R_q . We need d_0 since \mathbf{A}_0 is a part of the MLWE sample that masks messages and has to look uniformly random. The parameter $\beta \in \mathbb{Z}$ can be set arbitrarily and impacts the quality of the trapdoor.

$$\mathbf{A}_0 = [1 \ a_0 \ b_0] \cdot d_0, \quad \text{where} \quad d_0, a_0 \leftarrow \mathcal{R}_q^\times \times \mathcal{R}_q, \quad s_a, e_a \leftarrow \mathcal{D}_{R, s_a} \quad (6)$$

and $b_0 = (a_0 \cdot s_a + e_a) - \beta$ for some $\beta \in \mathbb{Z}$

2. $\mathbf{A}_1 \leftarrow \mathcal{R}_q^3$ is a uniformly random matrix.
3. \mathbf{A}_2 is uniformly random with the first element equal to 1. This is because we embed a ModNTRU trapdoor in \mathbf{A}_2 in the proof. We further multiply it by a unit tag d_2 since \mathbf{A}_2 is also used as part of the MLWE sample that masks messages.

$$\mathbf{A}_2 = [1 \ a_2 \ b_2] \cdot d_2, \quad \text{where} \quad d_2, a_2, b_2 \leftarrow \mathcal{R}_q^\times \times \mathcal{R}_q^2 \quad (7)$$

4. \mathbf{G} is a gadget matrix as introduced by Micciancio and Peikert [MP12].

$$\mathbf{G} = [1 \ g \ g^2], \quad \text{where} \quad g \approx q^{1/3} \quad (8)$$

5. $r \leftarrow R_q$ is used as a target for preimage sampling.

Identity Embedding. The set of identities of our TIBE scheme coincides with the set of signature verification keys of the corresponding TKEM we denote S_{vk} ⁶. Let $\mathbf{E} : S_{\text{vk}} \rightarrow R_q$ denote an embedding map from the set of identities S_{vk} of the TIBE into ring elements. For the security proof to work this map has to satisfy $\forall \text{vk}_0 \neq \text{vk}_1 : \mathbf{E}(\text{vk}_0) - \mathbf{E}(\text{vk}_1) \in R_q^\times$. Hence, we choose the ring R and modulus q such that $R_q \sim \mathbb{F}_{d/2} \times \mathbb{F}_{d/2}$, let $f : R_q \rightarrow \mathbb{F}_{d/2} \times \mathbb{F}_{d/2}$ denote the said isomorphism.

We require that $|S_{\text{vk}}| < q^{d/2}$, then we pick an arbitrary efficient embedding E_F from S_{vk} to $\mathbb{F}_{d/2} \setminus \{0\}$ e.g using binary representations of the polynomial coefficients. For $\text{vk} \in S_{\text{vk}}$ we define $\mathbf{E}(\text{vk}) := f^{-1}(E_F(\text{vk}), E_F(\text{vk}))$. Then $\forall \text{vk}_0 \neq \text{vk}_1 : E_F(\text{vk}_0) - E_F(\text{vk}_1) \neq 0$ implying that $\mathbf{E}(\text{vk}_0) - \mathbf{E}(\text{vk}_1)$ is a unit in R_q .

Encryption procedure. In the encryption procedure, the encrypter constructs a public matrix $[\mathbf{F}_{\text{id}} \ r]$ for the identity id . Then they generate a Lindner-Peikert [LP11] style ciphertext (\mathbf{u}, v) for the public key $[\mathbf{F}_{\text{id}} \ r]$ and the message msg . This is described in Algorithm 4.

⁶ When instantiated with WOTS⁺ the verification key is a binary string of fixed length.

Decryption procedure. In the decryption procedure, the parties jointly compute a vector $\bar{\mathbf{z}} = [\mathbf{z} \ \mathbf{x}_0 \ \mathbf{x}_1]$ such that $\mathbf{F}_{\text{id}} \cdot \bar{\mathbf{z}}^T = r$. They then use $\bar{\mathbf{z}}$ to perform a Lindner-Peikert style decryption procedure, by computing $v - \bar{\mathbf{z}} \cdot \mathbf{u}$ and decoding it to recover the message. This is described in Figure 5.

The TIBE is formally described in Figures 4 and 5. For clarity of reading, we assume that the state of party i contains (i) the secret key share of i , (ii) all contributions of all parties up to the current round, (iii) the values privately generated by i during previous rounds.

5 Threshold-CCA Security Proof

We use the matrix \mathbf{R} sampled as below to introduce a Micciancio-Peikert trapdoor [MP12] into $[\mathbf{A}_0 \mid \mathbf{A}_1]$ as in Lemma 10. Hence, the quality of preimages sampled using \mathbf{R} depends on its norm we upper bound below.

Lemma 15 (Norm bound on \mathbf{R}). *Let $R = \mathbb{Z}[X]/(X^d + 1)$. Let*

$$\mathbf{R} = \begin{bmatrix} e_1 & e_2 & e_3 \\ s_1 & s_2 & s_3 \\ 0 & 0 & 0 \end{bmatrix} \text{ where } e_i, s_i \leftarrow \mathcal{D}_{R, \varsigma_{\text{RLWE}}}$$

Then:

$$\Pr(\|\mathbf{R}\| \geq \sqrt{6} \cdot d \cdot \varsigma_{\text{RLWE}}) \leq 6 \cdot 2^{-12d}.$$

Proof. The spectral norm of \mathbf{R} does not change if we remove the last all-zero row of the matrix. Denote $\mathbf{R}' = \begin{bmatrix} e_1 & e_2 & e_3 \\ s_1 & s_2 & s_3 \end{bmatrix}$. Then:

$$\|\mathbf{R}'\| \leq \|\text{rot}(\mathbf{R}')\|_F = \sqrt{d \left(\sum_{i=1}^3 \|e_i\|^2 + \sum_{i=1}^3 \|s_i\|^2 \right)} \leq \sqrt{d \cdot \varsigma_{\text{RLWE}}^2 \cdot 6d} = \sqrt{6} \cdot d \cdot \varsigma_{\text{RLWE}}$$

The last inequality holds by Lemma 3 with probability at least $1 - 6 \cdot 2^{-12d}$. \square

We prove the Selective-ID Security of the TIBE scheme in Section 4. Then the CCA Security of the TKEM follows by Theorem 2.

Theorem 3. *Let TIBE be the scheme defined in Figure 4 and Figure 5. Suppose that the space of identities is bounded as $|S_{\text{id}}| = |S_{\text{vk}}| < q^{d/2}$. Let $R = \mathbb{Z}[X]/(X^d + 1)$ be a ring of degree d , q be a prime number s.t. $R_q \sim \mathbb{F}_{d/2} \times \mathbb{F}_{d/2}$. Let $\varsigma, \varsigma_{\text{RLWE}}, \varsigma_a, \beta > 0$, $\varepsilon \in (0, 1/2)$, $\delta > 0$ s.t. $2^{2\delta} \geq 2^{-d+1} + 6q^2$. Suppose that:*

$$\varsigma > \eta_\varepsilon(R^3) \tag{9}$$

$$\varsigma' > 2\sqrt{6} \cdot d \cdot \varsigma_{\text{RLWE}} \cdot \varsigma \tag{10}$$

$$\varsigma_p > 2d \cdot q^{1/3+1/3d} \tag{11}$$

$$\varsigma_p > 1.17q^{1/3} \cdot \eta_\varepsilon(R^3) \tag{12}$$

Algorithm 1 $\text{TIBE.Setup}(1^\kappa) \rightarrow \text{ek}, \{\text{dk}_i\}_{i \in [N]}$

```

1:  $d_0, a_0 \leftarrow \mathcal{R}_q^\times \times \mathcal{R}_q$ 
2:  $(s_a, e_a) \leftarrow \mathcal{D}_{R^2, \varsigma_a}$ 
3:  $\mathbf{A}_0 := \begin{bmatrix} 1 & a_0 & b_0 \end{bmatrix} \cdot d_0$ , where  $b_0 = (a_0 \cdot s_a + e_a) - \beta$ 
4:  $\mathbf{A}_1 \leftarrow \mathcal{R}_q^3$ 
5:  $\mathbf{A}_2 := \begin{bmatrix} 1 & a_2 & b_2 \end{bmatrix} \cdot d_2$ , where  $d_2, a_2, b_2 \leftarrow \mathcal{R}_q^\times \times \mathcal{R}_q^2$ 
6:  $\mathbf{G} := \begin{bmatrix} 1 & g & g^2 \end{bmatrix}$  where  $g \approx q^{1/3}$ 
7:  $r \leftarrow R_q$ 
8:  $\text{ek} := (\mathbf{A}_0, \mathbf{A}_1, \mathbf{A}_2, \mathbf{G}, r)$ 
9:  $\forall i \in [N] : \text{dk}_i := (\llbracket s_a \rrbracket_i, \llbracket e_a \rrbracket_i)$   $\triangleright (T, [N])$ -sharing of  $s_a, e_a$ 
10: return  $\text{ek}, \{\text{dk}_i\}_{i \in [N]}$ 

```

Algorithm 2 $\text{Encode}(\text{msg} = (b_0, \dots, b_{d-1})) \rightarrow M \in R_q$

```

1: return  $M = \sum_{i \in [d]} b_i \cdot \lfloor q/2 \rfloor \cdot X^i$ 

```

Algorithm 3 $\text{Decode}(M = \sum_{i \in [d]} \alpha_i \cdot X^i) \rightarrow \text{msg} \in \{0, 1\}^d / \perp$

```

1: return  $\text{msg} = \sum_{i \in [d]} \lfloor \alpha_i \rfloor_q X^i$ 

```

Algorithm 4 $\text{TIBE.Encrypt}(\text{msg} \in \{0, 1\}^d, \text{id} = \text{vk}, \text{ek}) \rightarrow \text{ct}$

```

1:  $\mathbf{F}_{\text{vk}} := [\mathbf{A}_0 \mid \mathbf{A}_1 - \mathbf{E}(\text{vk}) \cdot \mathbf{G} \mid \mathbf{A}_2]$   $\triangleright \mathbf{F}_{\text{vk}} \in \mathcal{R}_q^{1 \times 9}$ 
2:  $s \leftarrow D_{R, \varsigma}$ 
3:  $\mathbf{e} \leftarrow D_{R^3, \varsigma} \times D_{R^3, \varsigma'} \times D_{R^3, \varsigma}$   $\triangleright \mathbf{e} \in \mathcal{R}_q^{1 \times 9}$ 
4:  $\mathbf{e}' \leftarrow D_{R, \varsigma}$ 
5:  $\mathbf{u} := \mathbf{F}_{\text{vk}}^T \cdot s + \mathbf{e}$   $\triangleright \mathbf{u} \in \mathcal{R}_q^{1 \times 9}$ 
6:  $v := r \cdot s + \mathbf{e}' + \text{Encode}(\text{msg})$ 
7: return  $\text{ct} := (\mathbf{u}, v)$ 

```

Fig. 4: TIBE algorithms for key generation, encryption and message encoding.

Algorithm 5 $\text{TIBE.ShareExtract}_0(\text{ct}, \text{vk}, \text{act}, \text{state}_i) \rightarrow \text{contrib}_{i,0}/\perp$

```

1:  $(\mathbf{x}_{i,0}, \mathbf{x}_{i,1}, \mathbf{p}_i) \leftarrow (D_{R^3, \varsigma_p})^3$ 
2:  $y_{i,0} = a_0 \cdot p_{i,0} + p_{i,1} \bmod q$ 
3:  $y_{i,1} = (\mathbf{A}_1 - \mathbf{E}(\text{vk}) \cdot \mathbf{G}) \cdot \mathbf{x}_{i,0} \bmod q$ 
4:  $y_{i,2} = \mathbf{A}_2 \cdot \mathbf{x}_{i,1} \bmod q$ 
5:  $w_i = y_{i,0} + y_{i,1} + y_{i,2}$ 
6:  $\text{cmt}_i = H_{\text{cmt}}(w_i)$ 
7: return  $\text{contrib}_{i,0} := \text{cmt}_i$ 

```

Algorithm 6 $\text{TIBE.ShareExtract}_1(\text{ct}, \text{vk}, \text{act}, \text{state}_i) \rightarrow \text{contrib}_{i,1}/\perp$

```

1: return  $\text{contrib}_{i,1} := w_i$ 

```

Algorithm 7 $\text{TIBE.ShareExtract}_2(\text{ct}, \text{vk}, \text{act}, \text{state}_i) \rightarrow \text{contrib}_{i,2}/\perp$

```

1: assert  $\forall j \in \text{act} \setminus \{i\} : H_{\text{cmt}}(w_j) = \text{cmt}_j$ 
2:  $\text{cnt} = \text{act} \parallel \text{ct} \parallel \{\text{cmt}_j, w_j\}_{j \in \text{act}}$ 
3:  $\mathbf{m}_i := \sum_{j \in \text{act}} H_{\text{mask}}(\text{seed}_{i \rightarrow j}, \text{cnt}) - \sum_{j \in \text{act}} H_{\text{mask}}(\text{seed}_{j \rightarrow i}, \text{cnt})$ 
4:  $w := \sum_{j \in \text{act}} w_j$ 
5:  $(c_0, c_1) := \text{Decomp}_\beta(d_0^{-1} \cdot (r - w))$ 
6:  $z_{i,0} = p_{i,0} + c_0 \cdot \lambda_{i,\text{act}} \cdot \llbracket e_a \rrbracket_i + m_{i,0}$ 
7:  $z_{i,1} = p_{i,1} + c_0 \cdot \lambda_{i,\text{act}} \cdot \llbracket s_a \rrbracket_i + m_{i,1}$ 
8:  $z_{i,2} = p_{i,2}$ 
9: return  $\text{contrib}_{i,2} := (\mathbf{z}_i, \mathbf{x}_{i,0}, \mathbf{x}_{i,1})$ 

```

Algorithm 8 $\text{TIBE.Combine}(\text{ct}, \text{vk}, \text{act}, \{(\mathbf{z}_i, \mathbf{x}_i)\}_{i \in \text{act}}) \rightarrow \text{msg}/\perp$

```

1:  $z_0 = \sum_{i \in \text{act}} z_{i,0} + c_1$ 
2:  $z_1 = \sum_{i \in \text{act}} z_{i,1}$ 
3:  $z_2 = \sum_{i \in \text{act}} z_{i,2} + c_0$ 
4:  $\mathbf{x}_0 = \sum_{i \in \text{act}} \mathbf{x}_{i,0}$ 
5:  $\mathbf{x}_1 = \sum_{i \in \text{act}} \mathbf{x}_{i,1}$ 

6:  $\mathbf{F}_{\text{vk}} := [\mathbf{A}_0 \mid \mathbf{A}_1 - \mathbf{E}(\text{vk}) \cdot \mathbf{G} \mid \mathbf{A}_2]$ 
7: assert  $\mathbf{F}_{\text{vk}} \cdot [\mathbf{z} \mid \mathbf{x}_0 \mid \mathbf{x}_1]^T = r \wedge \|\mathbf{F}_{\text{vk}} \cdot [\mathbf{z} \mid \mathbf{x}_0 \mid \mathbf{x}_1]\|_\infty < B$ 
8: return  $\text{msg} := \text{Decode}(v - [\mathbf{z} \mid \mathbf{x}_0 \mid \mathbf{x}_1]^T \cdot \mathbf{u})$ 

```

Fig. 5: TIBE algorithms for threshold decryption.

Then for any PPT adversary \mathcal{A} :

$$\begin{aligned} \text{Adv}_{\mathcal{A}, \text{TIBE}}^{\text{sel-ID-IND}}(\kappa) &\leq 3 \cdot \text{Adv}_{\mathcal{A}_1}^{\text{RLWE}_{R,1,q,\varsigma_{\text{RLWE}}}}(\kappa) + 2\text{Adv}_{\mathcal{A}_2}^{\text{ModNTRU}_{R,q,2}}(\kappa) \\ &\quad + \text{Adv}_{\mathcal{A}_3}^{\text{Hint-RLWE}_{R,1,q,\varsigma_a}^{Q_{\text{Dec}}, \varsigma_p, \beta}}(\kappa) + \text{Adv}_{\mathcal{A}_4}^{\text{RLWE}_{R,7,q,\varsigma}}(\kappa) \\ &\quad + 6\varepsilon + Q_{\text{Dec}} O(\varepsilon) + 9 \cdot 2^{-d+\delta} \cdot Q_{\text{Dec}} + \frac{(Q_{\text{Dec}} + Q_{H_{\text{cmt}}})^2}{2^{2\kappa}} \\ &\quad + \frac{Q_{\text{Dec}}(Q_{\text{Dec}} + Q_{H_{\text{cmt}}})}{2^{-d+\delta+2}} + 2 \cdot q^{-d/2} + \frac{Q_{H_{\text{mask}}}}{2^\kappa} + 6 \cdot 2^{-12d}. \end{aligned}$$

where adversaries \mathcal{A}_4 to \mathcal{A}_1 run in approximately the same time as \mathcal{A} .

We give an overview of the proof hybrid by hybrid. The full proof can be found in Section B.

Hyb₀ Identical to the real experiment.

Hyb₁ In this transition, the challenger changes the distribution of \mathbf{A}_1 from uniformly random to $\mathbf{A}_1 = \mathbf{A}_0 \cdot \mathbf{R} + \mathbf{E}(\text{vk}^*) \cdot \mathbf{G}$ based on the challenge identity vk^* . This introduces a gadget trapdoor for the matrix $[\mathbf{A}_0 \mid \mathbf{A}_1 - \mathbf{E}(\text{vk}) \cdot \mathbf{G}]$ for every $\text{vk} \neq \text{vk}^*$ by the properties of the embedding function. This transition relies on $\text{RLWE}_{R,1,q,\varsigma_{\text{RLWE}}}$.

Hyb₂ We pick an arbitrary honest party h which we call the “last acting” honest party. In this hybrid, the challenger delays sampling of $(\mathbf{p}_h, \mathbf{x}_{h,0}, \mathbf{x}_{h,1})$ for h until the end of **TIBE.ShareExtract₁**. This is done by programming the random oracle $H_{\text{cmt}}(w_h)$. The programming is possible since w_h has enough entropy, which is guaranteed by Equation (11) and Lemma 8. This transition is statistical.

Hyb_{3,4} The following hybrids change the distribution of the masks \mathbf{m}_i to perfectly uniform for all honest parties except h . This transition relies on Lemma 17, that we borrow from [KRT24, Theorem 4.1]. The lemma requires all commitments to have a unique opening which we enforce in **Hyb₃**. We prove that the statistical distance between these hybrids is negligible.

Hyb₅ In this hybrid we change the simulation of honest party contributions in **TIBE.ShareExtract₂**. The shares of regular honest parties are replaced with uniform values and the mask are computed accordingly. This relies on the distribution of the masks from **Hyb₄**. For the last acting honest party we rephrase its contributions as a function of the main secrets s_a, e_a removing the Shamir sharing. This change is syntactical and does not change the Adversary’s advantage.

Hyb₆₋₈ In **Hyb₆** – **Hyb₈**, we remove the trapdoor $b = a_0 s_a + e_a$ from \mathbf{A}_0 using the Coset-Hint-MLWE assumption.

- In **Hyb₆**, we prepare the values dependent on (s_a, e_a) to follow the Coset-Hint-MLWE distributions exactly. To do that we change the order of sampling (c_0, c_1) and $(\mathbf{p}_h, \mathbf{x}_{h,0}, \mathbf{x}_{h,1})$ and sample the latter vector in two parts.

- In Hyb_7 , we apply Coset-Hint-MLWE assumption and sample $b \leftarrow R_q$. It is easy to see that the distributions in Hyb_6 and Hyb_7 can be simulated efficiently given a Coset-Hint-MLWE instance. Using Equation (12), we also leverage Corollary 1 to break down Coset-Hint-MLWE in two more standard assumptions, ModNTRU and Hint-MLWE.
- In Hyb_8 we change the sampling of $(\mathbf{p}_h, \mathbf{x}_{h,0}, \mathbf{x}_{h,1})$ again to efficiently simulate the preimages only given the MP12 trapdoor \mathbf{R} for every $\mathbf{vk} \neq \mathbf{vk}^*$. In this hybrid and the next one, we require $|S_{\mathbf{vk}}| < q^{d/2}$ as enforced in the conditions.

Note that now the matrix \mathbf{A}_0 is perfectly random.

- Hyb_9 In this hybrid, thanks to Equations (9) and (10), we can use Lemma 12 to change the noise distribution of the challenge ciphertext in order to embed an LWE sample with respect to $\mathbf{A}_0, \mathbf{A}_2$ and r . The norm of the new noise follows from Lemma 15. This change is statistical.
- Hyb_{10} In the last hybrid we apply the RLWE assumption to the challenge ciphertext. Since now the challenge plaintext is masked with a uniformly random value the Adversary's advantage in this experiment is 0.

We conclude the argument with an upper bound on the \mathcal{A} 's advantage in the initial experiment by tracing the change in the advantage after every transition. The main security theorem is a direct corollary of Theorems 2, 3 and 6.

Theorem 4. *The TKEM defined in Figures 3 to 5 is CCA secure with*

$$\begin{aligned}
\text{Adv}_{\mathcal{A}}^{\text{CCA2}}(\kappa) &\leq 2^{-128} + 3 \cdot \text{Adv}_{\mathcal{A}_1}^{\text{RLWE}_{R,1,q,\varsigma_{\text{RLWE}}}}(\kappa) + 2\text{Adv}_{\mathcal{A}_2}^{\text{ModNTRU}_{R,q,2}}(\kappa) \\
&\quad + \text{Adv}_{\mathcal{A}_3}^{\text{Hint-RLWE}_{R,1,q,\varsigma_a}^{Q_{\text{Dec}}, \varsigma_p, \beta}}(\kappa) + \text{Adv}_{\mathcal{A}_4}^{\text{RLWE}_{R,7,q,\varsigma}}(\kappa) + 6\varepsilon \\
&\quad + Q_{\text{Dec}} O(\varepsilon) + 9 \cdot 2^{-d+\delta} \cdot Q_{\text{Dec}} + \frac{(Q_{\text{Dec}} + Q_{H_{\text{cmt}}})^2}{2^{2\kappa}} \\
&\quad + \frac{Q_{\text{Dec}}(Q_{\text{Dec}} + Q_{H_{\text{cmt}}})}{2^{-d+\delta+2}} + 2 \cdot q^{-d/2} + \frac{Q_{H_{\text{mask}}}}{2^\kappa} + 6 \cdot 2^{-12d}
\end{aligned}$$

where adversaries \mathcal{A}_4 to \mathcal{A}_1 run in approximately the same time as \mathcal{A} .

6 Decapsulation Consistency

In a TKEM, we expect that if a valid message was recovered from the decapsulation algorithm, it is equal to the message that was encapsulated in the given ciphertext. There are 4 possible configurations for proving this property:

1. Honest ciphertext, and honest decapsulation shares;
2. Honest ciphertext, and malicious decapsulation shares;
3. Malicious ciphertext, and honest decapsulation shares;
4. Malicious ciphertext, and malicious decapsulation shares.

Note that since our scheme is not robust, in any of the malicious settings (Items 2 to 4) the Adversary can force the decapsulation procedure to output \perp . When both the ciphertext and the shares are honest (Item 1), we recover the Correctness property from Definition 6. We prove that the TIBE in Figures 4 and 5 fulfils this property with overwhelming probability.

Theorem 5 (Correctness). *Let $B = d\varsigma_a \sqrt{d} \left\lceil \frac{q-1}{2\beta} \right\rceil + \beta/2 + \varsigma_p \sqrt{Td}$, and $\varsigma' = 2\sqrt{6} \cdot d \cdot \varsigma_{\text{RLWE}} \cdot \varsigma$. If:*

$$\varsigma\sqrt{d} + 3d\sqrt{d} \left(\varsigma B + (\varsigma + \varsigma') \varsigma_p \sqrt{Td} \right) \leq q/4, \quad (13)$$

then the TIBE in Figures 4 and 5 is correct except with probability $O(de^{-\pi d})$.

Proof. For any message msg and honestly generated ciphertext ct :

$$\text{ct} = (\mathbf{u} := \mathbf{F}_{\text{vk}} \cdot s + \mathbf{e}, v := r \cdot s + e' + \text{Encode}(\text{msg})), \text{vk}, \text{sig}),$$

at the end of the honest decryption procedure (Figure 5), **TIBE.Combine** computes a vector $\mathbf{y} = (z_0, z_1, z_2, \mathbf{x}) \in R^9$ where:

$$\begin{aligned} z_0 &= c_0 \cdot s_a + c_1 + \sum_{j \in \text{act}} p_{j,0}, & z_1 &= c_0 \cdot e_a + \sum_{j \in \text{act}} p_{j,1}, & z_2 &= c_0 + \sum_{j \in \text{act}} p_{j,2} \\ \mathbf{x} &= \sum_{j \in \text{act}} \mathbf{x}_j. \end{aligned}$$

Recall from Figures 4 and 5 that $(s_a, e_a) \sim \mathcal{D}_{R^2, \varsigma_a}$ and $(\mathbf{x}_{i,0}, \mathbf{x}_{i,1}, \mathbf{p}_i) \sim (D_{R^3, \varsigma_p})^3$. It follows from (i) the triangle inequality and (ii) Lemmas 2 and 4 that:

$$\|z_0\|_\infty \leq \|c_0 \cdot s_a\|_\infty + \|c_1\|_\infty + \left\| \sum_{j \in \text{act}} p_{j,0} \right\|_\infty \quad (14)$$

$$\leq d\varsigma_a \sqrt{d} \left\lceil \frac{q-1}{2\beta} \right\rceil + \beta/2 + \varsigma_p \sqrt{Td} = B \quad (15)$$

except with probability $\leq 4de^{-\pi d}$. The same tail bound holds for $\|z_1\|_\infty$ and $\|z_2\|_\infty$. Similarly, we obtain $\|\mathbf{x}\|_\infty \leq \varsigma_p \sqrt{Td}$ with the same probability. Then for the decryption, by applying the triangle inequality in Equation (17), followed by Lemmas 2 and 4 in Equation (17), we get:

$$\begin{aligned} \|v - \mathbf{u}^T \cdot \mathbf{y} - \text{Encode}(\text{msg})\|_\infty &= \|e' - \mathbf{e}^T \cdot \mathbf{y}\|_\infty \\ &\leq \|e'\|_\infty + \|\mathbf{e}^T \cdot \mathbf{y}\|_\infty \end{aligned} \quad (16)$$

$$\begin{aligned} &\leq \|e'\|_\infty + \|\mathbf{e}_0^T \cdot \mathbf{z}\|_\infty + \|\mathbf{e}_1^T \cdot \mathbf{x}_0\|_\infty + \|\mathbf{e}_2^T \cdot \mathbf{x}_1\|_\infty \\ &\leq \varsigma\sqrt{d} + 3d\sqrt{d} \left(\varsigma B + (\varsigma + \varsigma') \varsigma_p \sqrt{Td} \right) \\ &\leq q/4 \end{aligned} \quad (17)$$

except with probability $O(de^{-\pi d})$. Equation (13) allows us to conclude. \square

Theorem 5 also covers the setting of Item 2. Indeed, the checks of **TKEM.Combine** imply that the shares reconstruct to \mathbf{y} s.t. $\mathbf{F}_{\mathbf{vk}} \cdot \mathbf{y} = r \bmod q$ and $\|\mathbf{y}\|_\infty \leq B$. When these two conditions are satisfied, Theorem 5 guarantees correct decryption with overwhelming probability.

The settings in Items 3 and 4 translate into the Decapsulation Consistency formally defined in Figure 1e. To prove it we lower bound the min-entropy of the honest ciphertext in Lemma 16. Then Decapsulation Consistency of the scheme follows from Lemma 14.

Intuitively, the BCHK part of our transform ensures that the adversary cannot tamper with received ciphertexts, since they do not own the corresponding signature keys. However, they are still able to craft their own malformed ciphertexts (e.g. sending random (\mathbf{u}, v) instead of the LWE samples) and this will be detected by the honest parties. Hence, we additionally check the ciphertexts that correspond to valid messages using the FO part of our BCHK+ transform. Since this check is only preformed after obtaining a valid message, it does not need to be thresholdised.

Lemma 16 (Honest Ciphertext Min-Entropy). *Let TIBE be the scheme described in Figures 4 and 5. If $\varsigma, \varsigma' \geq \eta_\varepsilon(R)$, it is γ -spread with $\gamma = \log(\varsigma^{7d} \cdot \varsigma'^{3d} \cdot (1 - \varepsilon)^{10})$.*

Proof. For a given $\text{msg} \in \{0, 1\}^d$, $\mathbf{vk} \in S_{\mathbf{vk}}$ and the public parameters \mathbf{ek} , the ciphertext is of the form $(\mathbf{u} = \mathbf{F}_{\mathbf{vk}}^T \cdot s + \mathbf{e}, v = r \cdot s + e' + \text{Encode}(\text{msg}))$, where $s \leftarrow D_{R, \varsigma}$, $\mathbf{e} \leftarrow D_{R^3, \varsigma} \times D_{R^3, \varsigma'} \times D_{R^3, \varsigma}$, $e' \leftarrow D_{R, \varsigma}$ are sampled at random. Let $p := \Pr_{\mathbf{e}, s, e'}((\mathbf{u}, v) = \text{TIBE.Encrypt}(\text{msg}, \mathbf{vk}, \mathbf{ek}))$. For any $(\mathbf{u}, v) \in R_q^9 \times R_q$ over the randomness of \mathbf{e}, s, e' we have

$$\begin{aligned} p &= \sum_{s \in R} \Pr_{\mathbf{e}, e'}(\mathbf{e} = \mathbf{u} - \mathbf{F}_{\mathbf{vk}}^T \cdot s, e' = v - r \cdot s - \text{Encode}(\text{msg})) \cdot \Pr(s \leftarrow D_{R, \varsigma}) \\ &\leq \sum_{s \in R} \Pr_{\mathbf{e}, e'}(\mathbf{e} = \mathbf{0}, e' = 0) \cdot \Pr(s \leftarrow D_{R, \varsigma}) = \Pr_{\mathbf{e}, e'}(\mathbf{e} = \mathbf{0}, e' = 0) \\ &= \frac{1}{\rho_\varsigma(R)^7 \cdot \rho_{\varsigma'}(R)^3} \leq \frac{1}{\varsigma^{7d} \cdot \varsigma'^{3d} \cdot (1 - \varepsilon)^{10}} \end{aligned}$$

where the last inequality comes from $\det(R) = 1$ and Lemma 5. \square

7 Instantiation and Parameter Sets

Assuming the conditions of Theorem 3 are met, we recall that $\text{Adv}_{\mathcal{A}, \text{TIBE}}^{\text{sel-ID-IND}}(\kappa)$ is upper-bounded as in Theorem 4. We now study each term independently.

RLWE. The hardnesses of $\text{RLWE}_{R, 1, q, \varsigma_{\text{RLWE}}}$ and $\text{RLWE}_{R, 7, q, \varsigma}$ can be estimated in a straightforward way using standard tools such as the lattice estimator. If (κ, d, q) are fixed, then this fixes ς and ς_{RLWE} as well.

Hint-MLWE. Next, we estimate the hardness of $\text{Hint-MLWE}_{R,1,1,q,\varsigma_a}^{Q_{\text{Dec}},\varsigma_p,\beta}$. Following Lemma 13, $\text{Hint-MLWE}_{R,1,1,q,\varsigma_a}^{Q_{\text{Dec}},\varsigma_p,\beta} \geq \text{RLWE}_{R,1,q,\varsigma_0}$, for $\frac{1}{\varsigma_0} \sim 2\left(\frac{1}{\varsigma_a} + \frac{B_{\text{hint}}}{\varsigma_p^2}\right)$ and $B_{\text{hint}} \sim \frac{Q_{\text{Dec}} d q^2}{12 \beta^2}$. This encourages us to select β rather large, since the reduction requires that $Q_{\text{Dec}} = O\left(\frac{\beta^2 \varsigma_p^2}{d q^2}\right)$ if we want $\varsigma_0 = \Omega(1)$. We ignore the factor 2 in the reduction, since it seems to be an artifact of the proof. We then study $\text{RLWE}_{R,1,q,\varsigma_0}$ as above.

ModNTRU. The best known attack against $\text{ModNTRU}_{R,q,m=2}$ is described in [CPS⁺20]. Due to the extremely large modulus q , one may wonder whether over-stretched NTRU attacks [ABD16, KF17, DvW21] may apply. A straightforward application would not work, since existing works focused on the NTRU case ($m = 1$). Assessing the applicability of this class of attacks to our case would require additional work, which we believe to be interesting and relevant future work.

Correctness. Finally, we study correctness via Theorem 5. Since the Hint-RLWE reduction forces β to be very large, it holds that $B \sim \beta/2$. For fixed d , the term ς' is essentially fixed by Equation (10). Since ς is also fixed (see RLWE paragraph), we may simply increase ς_q and T as long as Equation (13) is satisfied.

Putting everything together gives us the parameter sets in Table 2.

Table 2: Parameter sets.

κ	T	Q_{Dec}	q	d	β	ς_a	ς	ς_p	ς_{RLWE}	$ \text{ek} $	$ \text{ct} $
128	32	2^{40}	2^{100}	4096	2^{77}	8	4	2^{47}	4	50 KiB	450 KiB

Acknowledgements

We would like the reviewers of ASIACRYPT 2025 for their thoughtful feedback, as well as Martin R. Albrecht, Katharina Boudgoust and Rafaël del Pino for useful discussions throughout the project.

Oleksandra Lapiha is supported by the EPSRC and the UK Government as a student of the Centre for Doctoral Training in Cyber Security for the Everyday at Royal Holloway, University of London (EP/S021817/1). Part of this work was conducted while she was interning at PQShield.

References

- ABB10. Shweta Agrawal, Dan Boneh, and Xavier Boyen. Efficient lattice (H)IBE in the standard model. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 553–572. Springer, Berlin, Heidelberg, May / June 2010. 1.2, 4

- ABD16. Martin R. Albrecht, Shi Bai, and Léo Ducas. A subfield lattice attack on overstretched NTRU assumptions - cryptanalysis of some FHE and graded encoding schemes. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part I*, volume 9814 of *LNCS*, pages 153–178. Springer, Berlin, Heidelberg, August 2016. [7](#)
- ABGS23. Diego F. Aranha, Carsten Baum, Kristian Gjøsteen, and Tjerdand Silde. Verifiable mix-nets and distributed decryption for voting from lattice-based assumptions. In Weizhi Meng, Christian Damsgaard Jensen, Cas Cremers, and Engin Kirda, editors, *ACM CCS 2023*, pages 1467–1481. ACM Press, November 2023. [1](#)
- AL21. Martin R. Albrecht and Russell W. F. Lai. Subtractive sets over cyclotomic rings - limits of Schnorr-like arguments over lattices. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part II*, volume 12826 of *LNCS*, pages 519–548, Virtual Event, August 2021. Springer, Cham. [2](#)
- ALLW25. Martin R. Albrecht, Russell W. F. Lai, Oleksandra Lapiha, and Ivy K. Y. Woo. Partial lattice trapdoors: How to split lattice trapdoors, literally. Cryptology ePrint Archive, Paper 2025/367, 2025. [1.2](#), [1.3](#), [1](#)
- ASY22. Shweta Agrawal, Damien Stehlé, and Anshu Yadav. Round-optimal lattice-based threshold signatures, revisited. In Mikolaj Bojanczyk, Emanuela Merelli, and David P. Woodruff, editors, *ICALP 2022*, volume 229 of *LIPICs*, pages 8:1–8:20. Schloss Dagstuhl, July 2022. [1.3](#)
- Ban93. W. Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 1(296):625–635, 1993. [3](#)
- BBH06. Dan Boneh, Xavier Boyen, and Shai Halevi. Chosen ciphertext secure public key threshold encryption without random oracles. In David Pointcheval, editor, *CT-RSA 2006*, volume 3860 of *LNCS*, pages 226–243. Springer, Berlin, Heidelberg, February 2006. [1.1](#), [1.2](#), [7](#), [2](#), [4.1](#), [3](#), [2](#)
- BCHK07. Dan Boneh, Ran Canetti, Shai Halevi, and Jonathan Katz. Chosen-ciphertext security from identity-based encryption. *SIAM J. Comput.*, 36(5):1301–1328, 2007. [1.1](#), [1.2](#), [3](#)
- BFRS18. Pauline Bert, Pierre-Alain Fouque, Adeline Roux-Langlois, and Mohamed Sabt. Practical implementation of ring-SIS/LWE based signature and IBE. In Tanja Lange and Rainer Steinwandt, editors, *Post-Quantum Cryptography - 9th International Conference, PQCrypto 2018*, pages 271–291. Springer, Cham, 2018. [1.2](#)
- BGG⁺18. Dan Boneh, Rosario Gennaro, Steven Goldfeder, Aayush Jain, Sam Kim, Peter M. R. Rasmussen, and Amit Sahai. Threshold cryptosystems from threshold fully homomorphic encryption. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part I*, volume 10991 of *LNCS*, pages 565–596. Springer, Cham, August 2018. [1.3](#), [1](#)
- BKL⁺25. Cecilia Boschini, Darya Kaviani, Russell Lai, Giulio Malavolta, Akira Takahashi, and Mehdi Tibouchi. Ringtail: Practical Two-Round Threshold Signatures from Learning with Errors. In *2025 IEEE Symposium on Security and Privacy (SP)*, pages 149–164, Los Alamitos, CA, USA, May 2025. IEEE Computer Society. [1.3](#)
- BKP13. Rikke Bendlin, Sara Krehbiel, and Chris Peikert. How to share a lattice trapdoor: Threshold protocols for signatures and (H)IBE. In Michael J. Jacobson, Jr., Michael E. Locasto, Payman Mohassel, and Reihaneh Safavi-Naini, editors, *ACNS 13 International Conference on Applied Cryptography and Network Security*, volume 7954 of *LNCS*, pages 218–236. Springer, Berlin, Heidelberg, June 2013. [1.3](#), [1](#)

- BS23. Katharina Boudgoust and Peter Scholl. Simple threshold (fully homomorphic) encryption from LWE with polynomial modulus. In Jian Guo and Ron Steinfeld, editors, *ASIACRYPT 2023, Part I*, volume 14438 of *LNCS*, pages 371–404. Springer, Singapore, December 2023. [1.3](#)
- CATZ24. Rutchathon Chairattana-Apirom, Stefano Tessaro, and Chenzhi Zhu. Partially non-interactive two-round lattice-based threshold signatures. In Kai-Min Chung and Yu Sasaki, editors, *ASIACRYPT 2024, Part IV*, volume 15487 of *LNCS*, pages 268–302. Springer, Singapore, December 2024. [1.3](#)
- CCMS21. Kelong Cong, Daniele Cozzo, Varun Maram, and Nigel P. Smart. Gladius: LWR based efficient hybrid public key encryption with distributed decryption. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part IV*, volume 13093 of *LNCS*, pages 125–155. Springer, Cham, December 2021. [1.3](#), [1](#)
- CGPP24. Arka Rai Choudhuri, Sanjam Garg, Julien Piet, and Guru-Vamsi Policharla. Mempool privacy via batched threshold encryption: Attacks and defenses. In Davide Balzarotti and Wenyuan Xu, editors, *USENIX Security 2024*. USENIX Association, August 2024. [1](#)
- CGS97. Ronald Cramer, Rosario Gennaro, and Berry Schoenmakers. A secure and optimally efficient multi-authority election scheme. In Walter Fumy, editor, *EUROCRYPT’97*, volume 1233 of *LNCS*, pages 103–118. Springer, Berlin, Heidelberg, May 1997. [1](#)
- CHK04. Ran Canetti, Shai Halevi, and Jonathan Katz. Chosen-ciphertext security from identity-based encryption. In Christian Cachin and Jan Camenisch, editors, *EUROCRYPT 2004*, volume 3027 of *LNCS*, pages 207–222. Springer, Berlin, Heidelberg, May 2004. [1.1](#), [1.2](#), [3](#)
- CHKP10. David Cash, Dennis Hofheinz, Eike Kiltz, and Chris Peikert. Bonsai trees, or how to delegate a lattice basis. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 523–552. Springer, Berlin, Heidelberg, May / June 2010. [1.2](#), [9](#)
- CPS⁺20. Chitchanok Chuengsatiansup, Thomas Prest, Damien Stehlé, Alexandre Wallet, and Keita Xagawa. ModFalcon: Compact signatures based on module-NTRU lattices. In Hung-Min Sun, Shiuh-Pyng Shieh, Guofei Gu, and Giuseppe Ateniese, editors, *ASIACCS 20*, pages 853–866. ACM Press, October 2020. [1.2](#), [2.5](#), [9](#), [9](#), [7](#)
- DKM⁺24. Rafaël Del Pino, Shuichi Katsumata, Mary Maller, Fabrice Mouhartem, Thomas Prest, and Markku-Juhani O. Saarinen. Threshold raccoon: Practical threshold signatures from standard lattice assumptions. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part II*, volume 14652 of *LNCS*, pages 219–248. Springer, Cham, May 2024. [1.2](#), [1.3](#), [4](#)
- DLN⁺21. Julien Devevey, Benoît Libert, Khoa Nguyen, Thomas Peters, and Moti Yung. Non-interactive CCA2-secure threshold cryptosystems: Achieving adaptive security in the standard model without pairings. In Juan Garay, editor, *PKC 2021, Part I*, volume 12710 of *LNCS*, pages 659–690. Springer, Cham, May 2021. [1.3](#), [1](#)
- DLP14. Léo Ducas, Vadim Lyubashevsky, and Thomas Prest. Efficient identity-based encryption over NTRU lattices. In Palash Sarkar and Tetsu Iwata, editors, *ASIACRYPT 2014, Part II*, volume 8874 of *LNCS*, pages 22–41. Springer, Berlin, Heidelberg, December 2014. [1.2](#)
- dPKPR24. Rafaël del Pino, Shuichi Katsumata, Thomas Prest, and Mélissa Rossi. Raccoon: A masking-friendly signature proven in the probing model. In

- Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part I*, volume 14920 of *LNCS*, pages 409–444. Springer, Cham, August 2024. [1.3](#)
- DvW21. Léo Ducas and Wessel P. J. van Woerden. NTRU fatigue: How stretched is overstretched? In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part IV*, volume 13093 of *LNCS*, pages 3–32. Springer, Cham, December 2021. [7](#)
- EEN⁺24. Muhammed F. Esgin, Thomas Espitau, Guilhem Niot, Thomas Prest, Amin Sakzad, and Ron Steinfeld. Plover: Masking-friendly hash-and-sign lattice signatures. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part VII*, volume 14657 of *LNCS*, pages 316–345. Springer, Cham, May 2024. [1.2](#), [1.2](#), [1.3](#), [2.5](#), [2.5](#), [13](#), [4](#), [1](#)
- EKT24. Thomas Espitau, Shuichi Katsumata, and Kaoru Takemure. Two-round threshold signature from algebraic one-more learning with errors. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part VII*, volume 14926 of *LNCS*, pages 387–424. Springer, Cham, August 2024. [1.3](#)
- ENP24. Thomas Espitau, Guilhem Niot, and Thomas Prest. Flood and submerse: Distributed key generation and robust threshold signature from lattices. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part VII*, volume 14926 of *LNCS*, pages 425–458. Springer, Cham, August 2024. [1.2](#), [1.2](#), [1.3](#)
- FO99. Eiichiro Fujisaki and Tatsuaki Okamoto. How to enhance the security of public-key encryption at minimum cost. In Hideki Imai and Yuliang Zheng, editors, *PKC’99*, volume 1560 of *LNCS*, pages 53–68. Springer, Berlin, Heidelberg, March 1999. [1.3](#)
- FO13. Eiichiro Fujisaki and Tatsuaki Okamoto. Secure integration of asymmetric and symmetric encryption schemes. *Journal of Cryptology*, 26(1):80–101, January 2013. [1.3](#), [4.1](#)
- GKS24. Kamil Doruk Gür, Jonathan Katz, and Tjerand Silde. Two-round threshold lattice-based signatures from threshold homomorphic encryption. In Markku-Juhani Saarinen and Daniel Smith-Tone, editors, *Post-Quantum Cryptography - 15th International Workshop, PQCrypto 2024, Part II*, pages 266–300. Springer, Cham, June 2024. [1.3](#)
- 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. [1.2](#), [6](#), [7](#), [11](#)
- HHK17. Dennis Hofheinz, Kathrin Hövelmanns, and Eike Kiltz. A modular analysis of the Fujisaki-Okamoto transformation. In Yael Kalai and Leonid Reyzin, editors, *TCC 2017, Part I*, volume 10677 of *LNCS*, pages 341–371. Springer, Cham, November 2017. [8](#), [4.1](#), [4.1](#)
- HRS16. Andreas Hülsing, Joost Rijneveld, and Fang Song. Mitigating multi-target attacks in hash-based signatures. In Chen-Mou Cheng, Kai-Min Chung, Giuseppe Persiano, and Bo-Yin Yang, editors, *PKC 2016, Part I*, volume 9614 of *LNCS*, pages 387–416. Springer, Berlin, Heidelberg, March 2016. [12](#), [7](#)
- Hül13. Andreas Hülsing. W-OTS+ - shorter signatures for hash-based signature schemes. In Amr Youssef, Abderrahmane Nitaj, and Aboul Ella Hassanien, editors, *AFRICACRYPT 13*, volume 7918 of *LNCS*, pages 173–188. Springer, Berlin, Heidelberg, June 2013. [12](#), [6](#), [6](#)

- KF17. Paul Kirchner and Pierre-Alain Fouque. Revisiting lattice attacks on over-stretched NTRU parameters. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017, Part I*, volume 10210 of *LNCS*, pages 3–26. Springer, Cham, April / May 2017. [7](#)
- KLO⁺19. Michael Kraitsberg, Yehuda Lindell, Valery Osheter, Nigel P. Smart, and Younes Talibi Alaoui. Adding distributed decryption and key generation to a ring-LWE based CCA encryption scheme. In Julian Jang-Jaccard and Fuchun Guo, editors, *ACISP 19*, volume 11547 of *LNCS*, pages 192–210. Springer, Cham, July 2019. [1](#)
- KLSS23. Duhyeon Kim, Dongwon Lee, Jinyeong Seo, and Yongsoo Song. Toward practical lattice-based proof of knowledge from hint-MLWE. In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part V*, volume 14085 of *LNCS*, pages 549–580. Springer, Cham, August 2023. [1.1](#), [1.2](#), [1.2](#), [1.3](#), [2.5](#), [2.5](#)
- KRT24. Shuichi Katsumata, Michael Reichle, and Kaoru Takemure. Adaptively secure 5 round threshold signatures from MLWE/MSIS and DL with rewinding. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part VII*, volume 14926 of *LNCS*, pages 459–491. Springer, Cham, August 2024. [1.2](#), [1.3](#), [4](#), [5](#), [B](#)
- KY16. Shuichi Katsumata and Shota Yamada. Partitioning via non-linear polynomial functions: More compact IBEs from ideal lattices and bilinear maps. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *ASIACRYPT 2016, Part II*, volume 10032 of *LNCS*, pages 682–712. Springer, Berlin, Heidelberg, December 2016. [1.2](#), [12](#)
- 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, Berlin, Heidelberg, February 2011. [1.3](#), [4.2](#)
- LPR13. Vadim Lyubashevsky, Chris Peikert, and Oded Regev. A toolkit for ring-LWE cryptography. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 35–54. Springer, Berlin, Heidelberg, May 2013. [8](#)
- LS18. Vadim Lyubashevsky and Gregor Seiler. Short, invertible elements in partially splitting cyclotomic rings and applications to lattice-based zero-knowledge proofs. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part I*, volume 10820 of *LNCS*, pages 204–224. Springer, Cham, April / May 2018. [1](#)
- LTANB25. René Peralta Luís T. A. N. Brandão. NIST First Call for Multi-Party Threshold Schemes, 2025. [1](#)
- MP12. Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 700–718. Springer, Berlin, Heidelberg, April 2012. [10](#), [4](#), [5](#)
- MR07. Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on gaussian measures. *SIAM J. Comput.*, 37(1):267–302, 2007. [5](#), [2.4](#)
- MS24. Daniele Micciancio and Adam Suhl. Simulation-secure threshold PKE from LWE with polynomial modulus. *IACR Commun. Cryptol.*, 1(4):2, 2024. [1.3](#)
- NIS24. NIST. FIPS 203: Module-Lattice-Based Key-Encapsulation Mechanism Standard, 2024. [1](#)

- Pas24. Alain Passelégue. Lecture notes in cryptography and security, March 2024. https://perso.ens-lyon.fr/alain.passelegue/teaching/docs/M1_2024/TD8_2024_corrected.pdf, last accessed on 16.05.2025. 4.1
- Pei08. Chris Peikert. Limits on the hardness of lattice problems in lp norms. *computational complexity*, 17(2):300–351, 2008. 4
- Reg09. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6), September 2009. 1.3
- SAB⁺22. Peter Schwabe, Roberto Avanzi, Joppe Bos, Léo Ducas, Eike Kiltz, Tancrede Lepoint, Vadim Lyubashevsky, John M. Schanck, Gregor Seiler, Damien Stehlé, and Jintai Ding. CRYSTALS-KYBER. Technical report, National Institute of Standards and Technology, 2022. available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022>. 1
- UXT⁺22. Rei Ueno, Keita Xagawa, Yutaro Tanaka, Akira Ito, Junko Takahashi, and Naofumi Homma. Curse of re-encryption: A generic power/EM analysis on post-quantum KEMs. *IACR TCHES*, 2022(1):296–322, 2022. 1.2
- YJW23. Yang Yu, Huiwen Jia, and Xiaoyun Wang. Compact lattice gadget and its applications to hash-and-sign signatures. In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part V*, volume 14085 of *LNCS*, pages 390–420. Springer, Cham, August 2023. 1.2, 1

A Additional Preliminaries

Definition 12 (WOTS⁺ Signature Scheme [Hül13], [HRS16]). Let n be an integer indication the message length in bits. Let $w > 1$ be the integer Winternitz parameter. Let $l_1 = \left\lceil \frac{n}{\log w} \right\rceil$, $l_2 = \left\lfloor \frac{\log(l_1 \cdot (w-1))}{\log w} \right\rfloor + 1$ and $l = l_1 + l_2$, $a > 0$. Let $\mathcal{F}_n = \{f_k : \{0, 1\}^n \rightarrow \{0, 1\}^n \mid k \in \{0, 1\}^n\}$ and $\mathcal{H}_n = \{\text{PRF}_{\text{seed}} : \{0, 1\}^a \rightarrow \{0, 1\}^n \mid \text{seed} \in \{0, 1\}^n\}$ be two families of functions. Let $H_{\text{msg}} : \{0, 1\}^* \rightarrow \{0, 1\}^n$ and $H_{\text{key}} : \{0, 1\}^* \rightarrow \{0, 1\}^n$ be collision-resistant hash functions. Define the chaining function for $n > 0, x \in \{0, 1\}^n, k = (k_0, \dots, k_{j-1}) \in \{0, 1\}^{n \cdot j}, r = (r_0, \dots, r_{j-1}) \in \{0, 1\}^{n \cdot j}$ and $i \in [j]$

$$\begin{aligned} \text{for } i = 0 : c_k^i(x, r) &= x, \\ \text{for } i > 0 : c_k^i(x, r) &= f_k(c_{k_i}^{i-1}(x, r) \oplus r_i). \end{aligned}$$

We define the $\text{WOTS.Keygen}(1^\kappa)$, $\text{WOTS.Sign}(\text{sk}, \text{msg})$, and $\text{WOTS.Verify}(\text{vk}, \text{msg}, \text{sig})$ algorithms of the WOTS⁺ signature scheme in Figure 6.

Theorem 6 (Implicit in [Hül13, Theorem 1]). For parameters $n = 256, w = 16, \kappa = 128$ instantiate

$$\begin{aligned} f_k(x) &= \text{SHA2-256}(\text{toByte}(0, 32) || k || x) \\ \text{PRF}_{\text{seed}}(y) &= \text{SHA2-256}(\text{toByte}(1, 32) || \text{seed} || y) \\ H_{\text{msg}}(m) &= \text{SHA2-256}(\text{toByte}(2, 32) || m) \\ H_{\text{key}}(k) &= \text{SHA2-256}(\text{toByte}(3, 32) || k) \end{aligned}$$

WOTS.Keygen(1^κ)	WOTS.Sign($\text{sk}, \text{vk}, \text{msg} \in \{0, 1\}^*$)
$\text{sk} = (\text{sk}_1, \dots, \text{sk}_l) \leftarrow \{0, 1\}^{n \times l}$	$m = H_{\text{msg}}(\text{msg})$
$\text{seed} \leftarrow \{0, 1\}^{2^\kappa}$	$\text{Decomp}_w(m) = (m_1, \dots, m_{l_1})$
$\text{vk}_1 = H_{\text{key}}(\{c_k^{w-1}(\text{sk}_i, r)\}_{i=1}^l)$	$\text{chk} = \sum_{i=1}^{l_1} (w - 1 - m_i)$
return $(\text{sk}, \text{vk} = (\text{seed}, \text{vk}_1))$	$\text{Decomp}_w(\text{chk}) = (\text{chk}_1, \dots, \text{chk}_{l_1})$
	denote $B = (b_1, \dots, b_l) = m \parallel \text{chk}$
	$(r, k) \leftarrow \{\text{PRF}_{\text{seed}}(i)\}_{i=1}^{2 \cdot (w-1)}$
	$\text{sig} = (\text{sig}_1, \dots, \text{sig}_l) = \{c_k^{b_i}(\text{sk}_i, r)\}_{i=1}^l$
	return sig
WOTS.Verify($\text{vk}, \text{msg}, \text{sig}$)	
compute $B = (b_1, \dots, b_l)$ as in WOTS.Sign	
$(r, k) \leftarrow \{\text{PRF}_{\text{seed}}(i)\}_{i=1}^{2 \cdot (w-1)}$	
return $\left(\text{vk}_1 = H_{\text{key}} \left(\left\{ c_{k[b_i+1, w-1]}^{w-1-b_i}(\text{sig}_i, r[b_i+1, w-1]) \right\}_{i=1}^l \right) \right)$	

Fig. 6: WOTS⁺ Signature Scheme

Then WOTS⁺ Signature Scheme in Figure 6 is an ε -sEU-CMA one-time signature scheme with $\varepsilon = 2^{-110}$. The length of the resulting signature equals 2144 bytes.

Remark 6. Unlike the scheme in [Hül13] we hash the message before signing and instantiate H_{msg} with SHA2-256 assumed to be collision resistant. This reduces the security from 2^{-256} to 2^{-128} due to birthday paradox.

Remark 7. The security loss in $\varepsilon = 2^{-110}$ seems to be an artefact of the proof technique resolved in [HRS16, Theorem 2] within the proof of EU-CMA security of the XMSS signature. Hence, we assume that WOTS⁺ is ε -sEU-CMA secure with $\varepsilon = 2^{-128}$.

B Full proof of Theorem 3

Proof. We consider a sequence of hybrids $(\text{Hyb}_i)_i$. For each i , we note ε_i the advantage of an adversary against Hyb_i .

Hyb_0 (sel-ID-IND game): the original sel-ID-IND experiment (Figure 1b) against TIBE (Figures 4 and 5).

Hyb_1 (Rerandomisation of \mathbf{A}_1): At the start of the sel-ID-IND experiment, the Challenger receives the challenge identity vk^* from the Adversary. We then

change the **TIBE.Setup**(1^κ) algorithm as follows. Sample

$$\mathbf{R} \leftarrow \begin{bmatrix} e_1 & e_2 & e_3 \\ s_1 & s_2 & s_3 \\ 0 & 0 & 0 \end{bmatrix} \text{ where } (e_i, s_i) \leftarrow \mathcal{D}_{R^2, \varsigma_{\text{RLWE}}}.$$

If $\|\mathbf{R}\| > \sqrt{6} \cdot d \cdot \varsigma_{\text{RLWE}}$, abort the experiment with a uniformly random answer bit $\leftarrow \{1, 0\}$. By Lemma 15, this event occurs with probability at most $6 \cdot 2^{-12d}$. Otherwise, set $\mathbf{A}_1 = \mathbf{A}_0 \cdot \mathbf{R} + \mathbf{E}(\text{vk}^*) \cdot \mathbf{G}$. Now

$$\mathbf{A}_0 \cdot \mathbf{R} = d_0 \cdot [a_0 \cdot s_1 + e_1, a_0 \cdot s_2 + e_2, a_0 \cdot s_3 + e_3]$$

contains 3 RLWE samples instead of uniformly random elements. The preimage queries are still answered using the trapdoor (s_a, e_a) of \mathbf{A}_0 . Therefore, the only difference between Hyb_0 and Hyb_1 is the sampling of \mathbf{A}_1 and the abort condition. Then the adversary's advantage can only increase by

$$\varepsilon - \varepsilon_1 \leq 6 \cdot 2^{-12d} + 3 \cdot \text{Adv}_{\mathcal{A}_1}^{\text{RLWE}_{R,1,q,\varsigma_{\text{RLWE}}}}(\kappa).$$

Hyb₂ (Programming the commitment): For each decryption query of the adversary, the set of active participants contains at least one honest party: $\text{act} \cap \text{honest} \neq \emptyset$. We designate $h \in \text{act} \cap \text{honest}$ as the “last acting” honest party. In this hybrid, in **TIBE.ShareExtract₀** we set $\text{cmt}_h \leftarrow \{0, 1\}^{2\kappa}$. We also delay the sampling of $(\mathbf{p}_h, \mathbf{x}_{h,0}, \mathbf{x}_{h,1}) \leftarrow \mathcal{D}_{R^9, \varsigma_p}$ to the end of **TIBE.ShareExtract₁** when other parties have already revealed their contributions $\{w_i\}_{i \in \text{cor}}$, and we program the random oracle to $H_{\text{cmt}}(\mathbf{F}_{\text{vk}} \cdot (\mathbf{p}_h \parallel \mathbf{x}_{h,0} \parallel \mathbf{x}_{h,1})) := \text{cmt}_h$.

This change is detected only if w_h has been queried to $H_{\text{cmt}}(\cdot)$ before, we prove that the Adversary can only guess the value of w_h with negligible probability. By Lemma 8, with probability at least $1 - 2^{-d+\delta}$ it holds that

$$\varsigma_p > 2d \cdot q^{1/3+1/3d} \geq \eta_{2^{-d+\delta}}(\Lambda_q^\perp([1 \ b_2 \ d_2])) \geq \eta_{2^{-d+\delta}}(\Lambda_q^\perp(\mathbf{F}_{\text{vk}})).$$

Recall in Hyb_1 the commitment is generated honestly as $w_h \leftarrow \mathbf{F}_{\text{vk}} \cdot \mathcal{D}_{R^9, \varsigma_p}$. By Lemma 7, its distribution is statistically close (for a term at most $2 \cdot 2^{-d+\delta}$) to the uniform distribution over R_q . Then, the probability of sampling any fixed value of w_h is bounded by $q^{-d} + 2 \cdot 2^{-d+\delta} + 2^{-d+\delta} \leq 2^{-d+\delta+2}$ where the additional $2^{-d+\delta}$ comes from the smoothing condition on the matrix \mathbf{F}_{vk} described above. Hence, the probability that one of the Q_{Dec} honestly generated w_h 's collides with previous queries to $H_{\text{cmt}}(\cdot)$ is at most:

$$\varepsilon_1 - \varepsilon_2 \leq 2^{-d+\delta+2} \cdot Q_{\text{Dec}} \cdot (Q_{\text{Dec}} + Q_{H_{\text{cmt}}}).$$

Hyb₃ (Unique commitment opening): In this hybrid, the Challenger aborts if there is a collision in the output of the random oracle H_{cmt} . The number of queries to H_{cmt} is bounded by $Q_{\text{Dec}} + Q_{H_{\text{cmt}}}$, so the probability of the collision is bounded by $\frac{(Q_{\text{Dec}} + Q_{H_{\text{cmt}}})^2}{2^{2\kappa}}$. Then:

$$\varepsilon_2 - \varepsilon_3 \leq \frac{(Q_{\text{Dec}} + Q_{H_{\text{cmt}}})^2}{2^{2\kappa}}.$$

Hyb₄ (Uniform masking): We recall in Lemma 17 a fact proved within [KRT24, Theorem 4.1] in the transition from Game 3 to Game 8 in the authors' numbering. The proof idea of [KRT24] is to carefully trace the value of cnt in the honest executions of the protocol and conclude that the adversary can distinguish the distributions only if it queries H_{mask} using a correct seed of one of the honest parties.

Lemma 17. *With the notations of Figure 5, let $H_{\text{mask}} : \{0,1\}^* \rightarrow R_q^2$ be modelled as a random oracle and $\text{seed}_{i,j} \leftarrow \{0,1\}^\kappa$ be uniformly distributed. For every decryption query of the adversary, let cnt be uniquely defined by $\text{cnt}_w := \text{act} \parallel \text{ct} \parallel \{\text{cmt}_j\}_{j \in \text{act}}$. The adversary's advantage in distinguishing \mathcal{D}_m and \mathcal{D}'_m is at most $Q_{H_{\text{mask}}}/2^\kappa$, where:*

$$\mathcal{D}_m = \left\{ \mathbf{m}_i = \sum_{j \in \text{act}} H_{\text{mask}}(\text{seed}_{i \rightarrow j}, \text{cnt}) - \sum_{j \in \text{act}} H_{\text{mask}}(\text{seed}_{j \rightarrow i}, \text{cnt}), i \in \text{act} \right\}$$

$$\mathcal{D}'_m = \left\{ \begin{array}{l} \mathbf{m}_i = \sum_{j \in \text{act}} H_{\text{mask}}(\text{seed}_{i \rightarrow j}, \text{cnt}) - \sum_{j \in \text{act}} H_{\text{mask}}(\text{seed}_{j \rightarrow i}, \text{cnt}), i \in \text{cor}, \\ \mathbf{m}_i \leftarrow R_q^2, \quad i \in \text{honest} \setminus \{h\} \\ \mathbf{m}_h = - \sum_{i \in \text{act} \setminus \{h\}} \mathbf{m}_i \end{array} \right\}$$

The abort condition in Hyb₃ ensures that cmt_i corresponds to only one value of w_i . Hence, cnt_w uniquely defines cnt and the conditions of Lemma 17 are satisfied. As in \mathcal{D}'_m we set $\mathbf{m}_j \leftarrow R_q^2$ for $j \in \text{honest} \setminus \{h\}$. For the last honest signer h we set for consistency $\mathbf{m}_h = - \sum_{j \in \text{act} \setminus \{h\}} \mathbf{m}_j$. This jump is indistinguishable based on Lemma 17. The preimage value is not affected by the change, so we can still simulate all signing rounds using the \mathbf{A}_0 trapdoor. Therefore,

$$\varepsilon_3 - \varepsilon_4 \leq \frac{Q_{H_{\text{mask}}}}{2^\kappa}.$$

Hyb₅ (Secret sharing): In this hybrid we change the simulation of partial decryption contributions. First, when the Adversary chooses the set cor of $T-1$ corrupt parties, the reduction samples $\text{dk}_i = (\llbracket s_a \rrbracket_i, \llbracket e_a \rrbracket_i) \leftarrow R_q^2$ for $i \in \text{cor}$ and sends the keys dk_i and honestly sampled seeds $\{\text{seed}_{i \rightarrow j}\}_{j \in [N]}$ to $i \in \text{cor}$. This does not change the output distributions by the properties of Shamir secret sharing.

For all $j \in (\text{honest} \cap \text{act}) \setminus \{h\}$ we change the order of sampling \mathbf{m}_j and \mathbf{z}_j . We sample $\mathbf{z}_j \leftarrow R_q^2$ and then define:

$$\begin{bmatrix} m_{j,0} \\ m_{j,1} \end{bmatrix} = \begin{bmatrix} z_{j,0} \\ z_{j,1} \end{bmatrix} - \begin{bmatrix} p_{j,0} \\ p_{j,1} \end{bmatrix} - c_0 \begin{bmatrix} \lambda_{j,\text{act}} \cdot \llbracket s_a \rrbracket_j \\ \lambda_{j,\text{act}} \cdot \llbracket e_a \rrbracket_j \end{bmatrix}$$

This change does not affect the distributions, since \mathbf{m}_j and \mathbf{z}_j are in bijective correspondence. For the last honest signer h we write:

$$\begin{aligned} z_{h,0} &= c_0 \cdot \lambda_{h,\text{act}} \cdot \llbracket s_a \rrbracket_h + p_{h,0} + m_{h,0} \\ &= c_0 \left(s_a - \sum_{j \in \text{act} \setminus \{h\}} \lambda_{j,\text{act}} \cdot \llbracket s_a \rrbracket_j \right) + p_{h,0} - \sum_{j \in \text{act} \setminus \{h\}} m_{j,0} \end{aligned} \quad (18)$$

$$\begin{aligned} &= c_0 \left(s_a - \sum_{j \in \text{act} \setminus \{h\}} \lambda_{j,\text{act}} \cdot \llbracket s_a \rrbracket_j \right) + p_{h,0} - \sum_{j \in \text{cor} \cap \text{act}} m_{j,0} \quad (19) \\ &\quad - \sum_{j \in (\text{honest} \cap \text{act}) \setminus \{h\}} (z_{j,0} - c_0 \cdot \lambda_{j,\text{act}} \cdot \llbracket s_a \rrbracket_j - p_{j,0}) \\ &= (c_0 \cdot s_a + p_{h,0}) - \sum_{j \in \text{cor} \cap \text{act}} \lambda_{j,\text{act}} \cdot \llbracket s_a \rrbracket_j - \sum_{j \in \text{cor} \cap \text{act}} m_{j,0} \\ &\quad - \sum_{j \in (\text{honest} \cap \text{act}) \setminus \{h\}} (z_{j,0} - p_{j,0}). \end{aligned}$$

Equation (18) uses the Shamir secret sharing definition and the transformation of \mathbf{m}_h from Hyb_4 . Equation (19) follows from changing the order of sampling \mathbf{m}_j and \mathbf{z}_j for $j \in (\text{honest} \cap \text{act}) \setminus \{h\}$ stated above. In the end, to simulate $z_{h,0}$ we only need $c_0 s_a + p_{h,0}$ and other known values. Similarly, for $z_{h,1}$ we only need $c_0 e_a + p_{h,1}$ and known values. The reconstructed preimage looks like

$$\begin{aligned} z_0 &= \underbrace{p_{h,0} + c_0 s_a + c_1}_{\tilde{z}_0} + \sum_{j \in \text{act} \setminus \{h\}} p_{j,0} \\ z_1 &= \underbrace{p_{h,1} + c_0 e_a}_{\tilde{z}_1} + \sum_{j \in \text{act} \setminus \{h\}} p_{j,1} \\ z_2 &= \underbrace{p_{h,2} + c_0}_{\tilde{z}_2} + \sum_{j \in \text{act} \setminus \{h\}} p_{j,2} \\ \mathbf{x} &= \mathbf{x}_h + \sum_{j \in \text{act} \setminus \{h\}} \mathbf{x}_j \end{aligned}$$

and by construction $\mathbf{A}_0 \cdot \tilde{\mathbf{z}} = c_0 \beta + c_1 \bmod q$. Hence,

$$\begin{aligned} w_h &= d_0^{-1} \cdot r - \sum_{i \in \text{act} \setminus \{h\}} w_i - \mathbf{A}_0 \cdot \tilde{\mathbf{z}} \\ &= d_0^{-1} \cdot r - \sum_{i \in \text{act} \setminus \{h\}} w_i - c_0 \beta - c_1. \end{aligned}$$

So w_h can be simulated using public values as well. These changes in the distributions are only syntactical. However, they allow us to simulate the protocol execution only using the LWE sample $a_0 s_a + e_a$ for the matrix \mathbf{A}_0 and Q_{Dec} hints of the form $(c_0, c_1), (c_0 s_a + p_{h,0}, c_0 e_a + p_{h,1})$. For convenience, we define an offset target vector $r_1 := d_0^{-1} \cdot r - \sum_{i \in \text{act} \setminus \{h\}} w_i$. Since the described changes do not affect any distributions, the advantage of the adversary does not change $\varepsilon_5 = \varepsilon_4$.

Hyb₆ (*Coset-Hint-LWE preparation*): At this point, public parameters are sampled as:

$$\begin{cases} a_0, a_2, b_2 \leftarrow R_q, (s_a, e_a) \leftarrow \mathcal{D}_{R^2, \varsigma_a}, b = a_0 \cdot s_a + e_a \\ \mathbf{A}_0 := \begin{bmatrix} 1 & a_0 & \beta - b \end{bmatrix}, \mathbf{A}_1 := \mathbf{A}_0 \cdot \mathbf{R} + \mathbf{E}(\text{vk}^*) \cdot \mathbf{G} \\ \mathbf{A}_2 := \begin{bmatrix} 1 & a_2 & b_2 \end{bmatrix}, \mathbf{T}^T = \begin{bmatrix} e_a & s_a & 1 \\ 1 & 0 & 0 \end{bmatrix} \end{cases}$$

Denote $\mathbf{A}'_1 = \mathbf{A}_1 - \mathbf{E}(\text{vk}) \cdot \mathbf{G}$. Then $\mathbf{F}_{\text{vk}} = [\mathbf{A}_0 \mid \mathbf{A}'_1 \mid \mathbf{A}_2]$, and the response to a decryption query with respect to $\text{id} = \text{vk}$ looks as follows.

$$\begin{aligned} & \mathbf{p}_h \leftarrow \mathcal{D}_{R^3, \varsigma_p}, \mathbf{x}_{h,1} \leftarrow \mathcal{D}_{R^3, \varsigma_p}, \mathbf{x}_{h,2} \leftarrow \mathcal{D}_{R^3, \varsigma_p} \\ & w_h = \mathbf{A}_0 \cdot \mathbf{p}_h + \mathbf{A}'_1 \cdot \mathbf{x}_{h,1} + \mathbf{A}_2 \cdot \mathbf{x}_{h,2} \\ & (c_0, c_1) := \text{Decomp}_\beta(r_1 - w_h) \\ & \tilde{\mathbf{z}} = \mathbf{T} \cdot \mathbf{c} + \mathbf{p}_h \\ & \text{return } (\tilde{\mathbf{z}}, \mathbf{x}_h) \end{aligned}$$

In Hyb₆, for each decryption query we do a statistical jump to the distribution where we sample (c_0, c_1) as independent variables, and adjust the distribution of $(\mathbf{p}_h, \mathbf{x}_h)$ to satisfy the image constraint.

$$\begin{aligned} & u \leftarrow R_q, (c_0, c_1) := \text{Decomp}_\beta(u) \\ & \mathbf{x}_{h,1} \leftarrow \mathcal{D}_{R^3, \varsigma_p}, (\mathbf{p}_h, \mathbf{x}_{h,2}) \leftarrow \mathcal{D}_{\Lambda_q^{r_1 - u - \mathbf{A}'_1 \cdot \mathbf{x}_{h,1}}([\mathbf{A}_0 \mid \mathbf{A}_2]), \varsigma_p} \\ & \tilde{\mathbf{z}} = \mathbf{T} \cdot \mathbf{c} + \mathbf{p}_h \\ & \text{return } (\tilde{\mathbf{z}}, \mathbf{x}_h) \end{aligned}$$

By construction $w_h = r_1 - u$, hence the value \mathbf{c} satisfies the original formula $\beta c_0 + c_1 = r_1 - w = u$. By Lemma 8, with probability $\geq 1 - 2^{-d+\delta}$ it holds that:

$$\varsigma_p > 2d \cdot q^{1/3+1/3d} \geq \eta_{2^{-d+\delta}}(\Lambda_q^\perp([1 \ b_2 \ d_2])) \geq \eta_{2^{-d+\delta}}(\Lambda_q^\perp(\mathbf{F}_{\text{vk}})).$$

Then in Hyb₅ by Lemma 7 $u' = \mathbf{F}_{\text{vk}} \cdot [\mathbf{p}_h \ \mathbf{x}_h]^T = w_h \bmod q$ is distributed within $3 \cdot 2^{-d+\delta}$ of uniform. Therefore, $u = d_0^{-1} \cdot r - u' \bmod q$ is statistically close to $u \leftarrow R_q$ in Hyb₆.

As mentioned before, by Lemma 8 $\varsigma_p > \eta_{2^{-d+\delta}}(\Lambda_q^\perp(\mathbf{F}_{\text{vk}}))$ with probability at least $1 - 2^{-d+\delta}$. Then in Hyb₆ by Lemma 7 $u' = \mathbf{F}_{\text{vk}} \cdot [\mathbf{p}_h \ \mathbf{x}_h]^T = w_h \bmod q$ is distributed within $2 \cdot 2^{-d+\delta}$ of uniform. Therefore, $u = d_0^{-1} \cdot r - u' \bmod q$ is statistically close to uniform in Hyb₇.

Similarly, by Lemma 7 the distribution of $\mathbf{p}_h, \mathbf{x}_h$ conditioned on u is exactly equal to $\mathcal{D}_{\Lambda_q^{r_1 - u}(\mathbf{F}_{\text{vk}}), \varsigma_p}$. We sample it by parts using Lemma 9 and the smoothing condition. We get $\text{SD}(D_1, D_2) \leq 2 \cdot 2^{-d+\delta}$, where:

$$\begin{cases} D_1 = \left\{ \mathbf{x}_{h,1} \leftarrow \mathcal{D}_{R^3, \varsigma_p}, (\mathbf{p}_h, \mathbf{x}_{h,2}) \leftarrow \mathcal{D}_{\Lambda_q^{r_1 - u - \mathbf{A}'_1 \cdot \mathbf{x}_{h,1}}([\mathbf{A}_0 \mid \mathbf{A}_2]), \varsigma_p} \right\} \\ D_2 = \left\{ (\mathbf{p}_h, \mathbf{x}_{h,1}, \mathbf{x}_{h,2}) \leftarrow \mathcal{D}_{\Lambda_q^{r_1 - u}(\mathbf{F}_{\text{vk}}), \varsigma_p} \right\} \end{cases}$$

Hence, overall:

$$\varepsilon_5 - \varepsilon_6 \leq 5 \cdot 2^{-d+\delta} \cdot Q_{\text{Dec}}.$$

Hyb₇ (*Hint-LWE jump*): Now the distribution of $(\tilde{\mathbf{z}}_h, \mathbf{x}_{h,2})$ with target vector $r_2 := r_1 - \mathbf{A}'_1 \cdot \mathbf{x}_{h,1}$ is equal to $\mathcal{D}_{\text{real}}$ of Coset-Hint-MLWE, therefore in **Hyb₇** we transition to the distribution that contains $\mathcal{D}_{\text{ideal}}$ as:

$$\begin{aligned} & a_0, a_2, b_2 \leftarrow R_q, \mathbf{b} \leftarrow R_q, (s_a, e_a) \leftarrow \mathcal{D}_{R^2, \varsigma_a} \\ & \mathbf{A}_0 := \begin{bmatrix} 1 & a_0 & \beta - b \end{bmatrix}, \mathbf{A}_1 := \mathbf{A}_0 \cdot \mathbf{R} + \mathbf{E}(\mathbf{vk}^*) \cdot \mathbf{G} \\ & \mathbf{A}_2 := \begin{bmatrix} 1 & a_2 & b_2 \end{bmatrix}, \mathbf{T}^T = \begin{bmatrix} e_a & s_a & 1 \\ 1 & 0 & 0 \end{bmatrix} \\ & u \leftarrow R_q, (c_0, c_1) := \text{Decomp}_\beta(u), \alpha := \mathbf{A}_0 \mathbf{T} \mathbf{c} \\ & \mathbf{x}_{h,1} \leftarrow \mathcal{D}_{R^3, \varsigma_p}, (\mathbf{p}_h, \mathbf{x}_{h,2}) \leftarrow \mathcal{D}_{\Lambda_q^{r_2 - \alpha}([\mathbf{A}_0 | \mathbf{A}_2]), \varsigma_p} \\ & \tilde{\mathbf{z}} = \mathbf{T} \cdot \mathbf{c} + \mathbf{p}_h \\ & \text{return } (\tilde{\mathbf{z}}, \mathbf{x}_h) \end{aligned}$$

The advantage of the adversary changes according to Theorem 1:

$$\varepsilon_6 - \varepsilon_7 \leq 2 \text{Adv}_{\mathcal{A}_2}^{\text{ModNTRU}_{R_q, 2, \mathcal{D}}}(\kappa) + \text{Adv}_{\mathcal{A}_3}^{\text{Hint-RLWE}_{R, 1, q, \varsigma_a}^{Q_{\text{Dec}}, \varsigma_p, \beta}}(\kappa) + Q_{\text{Dec}} \cdot O(\varepsilon)$$

Hyb₈ (*Hint-LWE post-processing*): To get back to a distribution we can efficiently simulate we change the order of sampling again to:

$$\begin{aligned} & u \leftarrow R_q, (c_0, c_1) := \text{Decomp}_\beta(u), \alpha := \mathbf{A}_0 \mathbf{T} \mathbf{c} \\ & \mathbf{x}_{h,2} \leftarrow \mathcal{D}_{R^3, \varsigma_p}, (\mathbf{p}_h, \mathbf{x}_{h,1}) \leftarrow \mathcal{D}_{\Lambda_q^{r_1 - \alpha - \mathbf{A}_2 \cdot \mathbf{x}_{h,2}}([\mathbf{A}_0 | \mathbf{A}'_1]), \varsigma_p} \\ & \tilde{\mathbf{z}} = \mathbf{T} \cdot \mathbf{c} + \mathbf{p}_h \\ & \text{return } (\tilde{\mathbf{z}}, \mathbf{x}_h) \end{aligned}$$

This distribution can be simulated efficiently since for all $\mathbf{vk} \neq \mathbf{vk}^*$ we have $\mathbf{E}(\mathbf{vk}) - \mathbf{E}(\mathbf{vk}^*) \in R_q^\times$ and by Lemma 10 the matrix \mathbf{R} is sufficient to sample preimages of $[\mathbf{A}_0 | \mathbf{A}'_1] = [\mathbf{A}_0 | -\mathbf{A}_0 \cdot \mathbf{R} - (\mathbf{E}(\mathbf{vk}) - \mathbf{E}(\mathbf{vk}^*)) \cdot \mathbf{G}]$. The advantage changes by another

$$\varepsilon_7 - \varepsilon_8 \leq 3 \cdot 2^{-d+\delta} \cdot Q_{\text{Dec}}.$$

In this hybrid we have $\mathbf{F}_{\mathbf{vk}^*} = [\mathbf{A}_0 | -\mathbf{A}_0 \cdot \mathbf{R} | \mathbf{A}_2]$ and the challenge ciphertext is of the form

$$\mathbf{u}^* = \mathbf{F}_{\mathbf{vk}^*}^T \cdot s^* + \mathbf{e}^* = \begin{bmatrix} \mathbf{A}_0^T \cdot s^* + \mathbf{e}_0^* \\ -\mathbf{R}^T \mathbf{A}_0^T \cdot s^* + \mathbf{e}_1^* \\ \mathbf{A}_2^T \cdot s^* + \mathbf{e}_2^* \end{bmatrix} \quad (20)$$

$$v^* = r \cdot s^* + e_3^* + \mathbf{M}_b \cdot \left\lfloor \frac{q}{2} \right\rfloor \quad (21)$$

The distributions of errors are $(\mathbf{e}_0^*, \mathbf{e}_2^*, e_3^*) \leftarrow \mathcal{D}_{R^7, \varsigma}$, $\mathbf{e}_1^* \leftarrow \mathcal{D}_{R^3, \varsigma'}$.

Hyb₉ : (Noise convolution) The distribution of the challenge ciphertext in Hyb₈ is statistically close to

$$\mathbf{u}^* = \mathbf{F}_{ID^*}^T \cdot s^* + \mathbf{e}^* = \begin{bmatrix} \mathbf{A}_0^T \cdot s^* + \mathbf{e}_1^* \\ \text{ReRand}(-\mathbf{R}^T, \mathbf{A}_0^T \cdot s^* + \mathbf{e}_1^*, \varsigma', \varsigma) \\ \mathbf{A}_2^T \cdot s^* + \mathbf{e}_2^* \end{bmatrix}.$$

Note that the challenge ciphertext includes RLWE samples $\mathbf{A}_0^T \cdot s^* + \mathbf{e}_0^*$, $\mathbf{A}_2^T \cdot s^* + \mathbf{e}_2^*$ and $r \cdot s^* + e_3^*$. Then the advantage of the adversary can only increase by the negligible distance between distributions determined in Lemma 12 as $\varepsilon_8 - \varepsilon_9 \leq 6\varepsilon$, which is applicable since ς , as constrained in Equation (9), and ς' as constructed in Hyb₁ and constrained in Equation (10), satisfy the conditions of Lemma 12.

Hyb₁₀ : (LWE for the ciphertext) Finally the challenge ciphertext is sampled as $\text{ct}^* = (\mathbf{u}^*, v^*)$ where $\mathbf{u}^* = (\mathbf{b}_0, \text{ReRand}(-\mathbf{R}^T, \mathbf{b}_0, \varsigma', \varsigma), \mathbf{b}_2)$ and $v^* \leftarrow b_3 + \text{Encode}(\text{msg}_b)$ where $(\mathbf{b}_0, \mathbf{b}_2, b_3) \leftarrow R_q^7$. Since the message is now masked by a uniformly random value, the advantage of the adversary is $\varepsilon_{10} = 0$.

If the adversary can distinguish Hybrids 9 and 10 they can break $\text{RLWE}_{R,7,q,\varsigma}$ conditioned on the first matrix element being a unit. The advantage of the adversary in LWE with a unit in the challenge matrix can decrease at most by

$$\Pr(a \notin R_q^* \mid a \leftarrow R_q) = \frac{2q^{d/2} - 1}{q^d} \leq \frac{2}{q^{d/2}}$$

The LWE challenge is embedded into $\mathbf{A}_0^T \cdot s^* + \mathbf{e}_0^*$, $\mathbf{A}_2^T \cdot s^* + \mathbf{e}_2^*$ and $r \cdot s^* + e_3^*$. Therefore, $\varepsilon_9 - \varepsilon_{10} \leq \text{Adv}_{\mathcal{A}_4}^{\text{RLWE}_{R,7,q,\varsigma}}(\kappa) + 2 \cdot q^{-d/2}$. \square