

NICE-PAKE: On the Security of KEM-Based PAKE Constructions without Ideal Ciphers

Nouri Alnahawi^{1,4}, Jacob Alperin-Sheriff², Daniel Apon³, and Alexander Wiesmaier^{1,4}

¹ Hochschule Darmstadt. {nouri.alnahawi,alexander.wiesmaier}@h-de.de

² Independent Researcher. jacobmas@gmail.com

³ The MITRE Corporation. crypto@mitre.org

⁴ National Research Center for Applied Cybersecurity ATHENE.

Abstract. The interest in realizing generic PQC KEM-based PAKEs has increased significantly in the last few years. One such PAKE is the CAKE protocol, proposed by Beguinet et al. (ACNS '23). However, despite its simple design based on the well-studied PAKE protocol EKE by Bellare and Merritt (IEEE S&P '92), both CAKE and its variant OCAKE do not fully protect against quantum adversaries, as they rely on the Ideal Cipher (IC) model. Related and follow-up works, including Pan and Zeng (ASIACRYPT '23), Dos Santos et al. (EUROCRYPT '23), Alnahawi et al. (CANS '24), and Arragia et al. (IACR '24/308) although touching on that issue, still rely on an IC. Considering the lack of a quantum IC model and the difficulty of using the classical IC to achieve secure instantiations on public keys in general and PQC in particular, we set out to eliminate it from PAKE design. In this paper, we present the **No IC Encryption (NICE)-PAKE**, a (semi)-generic PAKE framework providing a quantum-safe alternative for the IC, utilizing simpler cryptographic components for the authentication step. To give a formal proof for our construction, we introduce the notions of A-Part-Secrecy (**A-SEC-CCA**), Splittable Collision Freeness (**A-CFR-CCA**) and Public Key Uniformity (**SPLIT-PKU**) for splittable LWE KEMs. We show the relation of the former to the Non-uniform LWE and the Weak Hint LWE assumptions, as well as its application to ring and module LWE. Notably, this side quest led to some surprising discoveries, as we concluded that the new notion is not directly interchangeable between the LWE variants, or at least not in a straightforward manner. Further, we show that our approach requires some tedious tweaking for the parameter choices in both FrodoKEM and CRYSTALS-Kyber to obtain a secure PAKE construction. We also address some fundamental issues with the common IC usage and identify differences between lattice KEMs regarding their suitability for generic PQC PAKEs, especially regarding the structure of their public keys. We believe that this work marks a further step towards achieving complete security against quantum adversaries in PQC PAKEs.

Keywords: Password Authenticated Key Exchange · PAKE · Key Encapsulation Mechanism · KEM · Post-Quantum Cryptography · PQC · Learning with Errors · LWE · Ideal Cipher Model · IC

1 Introduction

As the looming threat of large-scale quantum computers endangering public-key cryptography becomes more evident, the search for quantum-safe replacement primitives and schemes in cryptographic protocols and constructions has also begun to gain more importance. The efforts made towards this goal can be seen, e.g., in the work done throughout the NIST Post-Quantum Cryptography (PQC) standardization process [54], which mainly focuses on Key Encapsulation Mechanisms (KEM) and Digital Signatures. In 2023, a set of candidates, including CRYSTALS-Kyber [18,55] (now referred to as ML-KEM¹), were announced, and their standardization was recently finalized². These candidates are mainly meant to replace classical asymmetric schemes based on the discrete logarithm and prime integer factorizing assumptions. Their applications are not limited to, but primarily found in security protocols such as Transport Layer Security (TLS) and Internet Key Exchange version 2 (IKEv2), and thus, have been analyzed and tested extensively in the last decade for that purpose [5]. The security of Password Authenticated Key Exchange (PAKE) protocols is no exception here since they very often rely on the hardness assumptions of classical primitives such as the renowned Diffie-Hellman (DH) key agreement [30] and its variants. For that matter, the more significant amount of work done on quantum-resistant PAKEs is based on direct constructions from PQC primitives, in addition to a fewer number based on PQC KEMs (cf. Sec. 2).

Motivation: In 2023 and 2024, five generic PAKE constructions based on PQC KEMs were proposed [11,42,4,47,10], with a noticeable focus on lattice-based schemes from the Learning with Errors (LWE) and Module LWE assumptions (Sec. 2). Despite some differences, these works have in common that they rely on the *Ideal Cipher* (IC) model in their formal security analysis. However, one should consider the disadvantages of modeling the encryption of public keys in PAKE protocols as an IC for three main reasons (Sec. 2 and Sec. 4): First, it is non-trivial to instantiate an IC as a block cipher over a group domain [10] (e.g., over classical finite fields or post-quantum lattices and isogenies). Second, in the specific case of ML-KEM, the public keys are only computationally indistinguishable from uniform, but not statistically close to uniform in a space of fixed-length random bit strings [31]. Finally, there are still no known adaptations for the IC model to deal with adversaries with quantum capabilities [53].

Main Contribution: In this work, we present the **No IC Encryption (NICE)** PAKE protocol (depicted in Fig. 1) to address the issues arising from the use of the IC model in the formal security analysis and concrete instantiations of

¹ <https://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.203.pdf>

² <https://www.federalregister.gov/documents/2024/08/14/2024-17956/announcing-issuance-of-federal-information-processing-standards-fips-fips-203-module-lattice-based>

PQC PAKEs. To overcome this obstacle, we provide an alternative for the symmetric (block cipher) encryption of the public key. We realize the PAKE public key authentication step in the form of a bit-wise XOR operation. Inspired by [47,10], our construction incorporates LWE KEMs with splittable public keys of the form $(\mathbf{A}, \mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e})$, where the lattice base \mathbf{A} is sampled from a public fixed-length bit string seed. As the seed is usually either appended or prepended to the plain text representation of the public key, we mask (hide) it using the PAKE password. Thus, we allow perfect secrecy schemes to be applied on random bit strings of fixed length without resorting to block ciphers. To support our formal analysis, we introduce three new properties for LWE (RLWE and MLWE) KEMs (Sec. 5) and refer to them as *\mathbf{A} -Part Secrecy under Chosen Ciphertext Attacks* (A-SEC-CCA), *Splittable Collision Freeness* (A-CFR-CCA), and *Public Key Uniformity* (SPLIT-PKU) for splittable KEMs. We provide a formal security analysis for our protocol in the extended RoR (*Real-or-Random*) BPR (Bellare-Pointcheval-Rogaway) model [12] (Sec. 6) and discuss concrete instantiations from secure PQC KEMs (Sec. 7). As a result, we conclude that KEMs built from LWE and variants are not seamlessly swappable across cryptographic constructions, especially in the case of PAKEs.

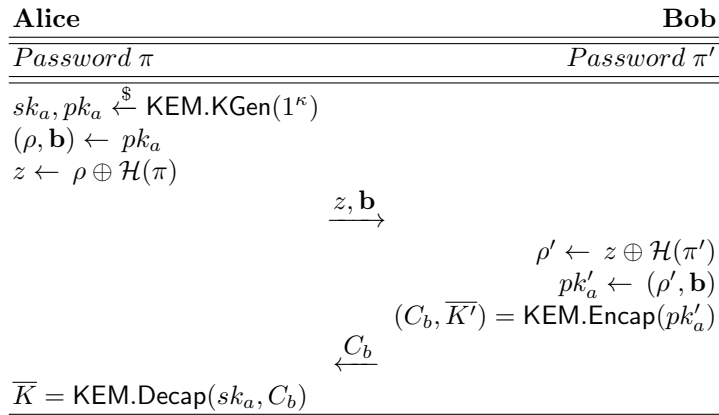


Fig. 1. The NICE-PAKE protocol utilizing a KEM with splittable keys using a random seed for the uniform sampling of the \mathbf{A} -part of the public key. Without knowledge of the correct password π , it is not possible to retrieve the seed required for sampling the matching lattice base used to invoke the encapsulation routine, since it internally requires both parts of the public key. Ciphertexts generated from a key with a non-matching \mathbf{A} -part cannot be decapsulated at all. Thus, only with a password known for the two parties is it possible to successfully execute the KEM. The protocol only provides implicit mutual authentication. A version with explicit authentication can nevertheless be obtained through adding a key confirmation round.

Bonus Contribution: Our proof for the newly introduced A-SEC-CCA property (Sec. 5) is based on less famous LWE assumptions. We show that this property requires only a straight-forward reduction to schemes built from plain (unstructured lattices) LWE assumptions, yet comes with a costly penalty leading to a significant loss in bit security for the concrete LWE scheme FrodoKEM. Nevertheless, we present a way to regain a sufficient amount of the lost bits for FrodoKEM using a non-standard value for the Gaussian variance parameter (Sec. 7). Further, we conclude that this method is not directly inter-changeable across other schemes from the Ring or Module LWE variants (e.g., NewHope and CRYSTALS-Kyber), and could lead to a complete break in security using the standard parameters (Sec. 7). Hence, we suggest that with a secret of higher entropy for the encapsulator’s randomness, the A-SEC-CCA property could also be applied to MLWE schemes with an acceptable loss in bit security resulting from reducing the dimension of the module lattice.

Limitations: Unlike augmented (asymmetric) PAKEs where a server only holds a one-way transformation of a user’s password, a rather known and common deficiency in balanced (symmetric) PAKEs including our construction is that they are, strictly speaking, vulnerable to pre-computation attacks and server compromise. Basically, a stolen (compromised) password could be used directly in impersonation attacks. Moreover, if the ephemeral randomness (of the underlying KEM) used by either party in the protocol is revealed, it would allow pre-computation attacks against the password. Both of these scenarios can neither be detected nor prevented by the threat model used in the formal analysis. In other words, the correctness of the provided proof in this paper relies strongly on the fact that an adversary is restricted by the capabilities allowed in the BPR framework, and especially that they are limited to random guessing on the correct password. That being said, an analysis done in the BPR model essentially only covers high-entropy password distributions. Moreover, a difficult issue is the inability to extract password guesses from an adversary’s interaction with the protocol, which ultimately leads to rather less-tight security bounds compared to proofs done in the Universal Composability (UC) model.

A Note on Idealized Objects (ROM vs. QROM): We note that we make use of only one idealized object, the classical ROM, in our analysis. Hence, we do not require the QROM, as the ROM is only used to model queries on a pre-defined and offline accessible password dictionary. This claim can be justified as follows:

1. It is safe to assume that an adversary with unbounded computational power can execute offline dictionary attacks. Hence, they can easily prepare a list of all hash outputs for all possible passwords in the password dictionary. This task is feasible even in classical settings, where no quantum computers are needed. In our construction, an adversary can indeed only make use of such queries before committing to a password guess, and cannot place adaptive password guesses via following `Send` queries in the BPR framework.
2. It is not plausible to model the behavior of internal functionalities from the KEM’s underlying PKE while maintaining a generic black-box construction

(e.g., the Fujisaki Okamoto Transform or the re-encryption routine). Thus, it is quite infeasible to extract the adversary’s password guesses from their interactions with instances and sessions in the protocol execution. Although our construction and its informal security arguments may look very similar to well known discrete-log based schemes, it is not possible to rely on interactive assumptions combined with the ROM to extract password guesses in this case. Therefore, there is no need to rely on a QROM for such queries, as they do not appear anywhere in the hybrids of the game-based proof.

A Note on CPA vs. CCA Security: Initially, we built our proof strategy (cf. App B) based on the CPA variants of the notions for key security (IND-CPA) and anonymity (ANO-CPA). Nevertheless, we opted to use the weaker assumptions IND-CCA and ANO-CCA within the actual formal analysis for two main reasons.

1. Considering the fact that the *Anonymity* property aims at protecting the identity of the public key used in an encapsulation to generate a ciphertext, it should be sufficient to rely on CPA security for the interaction between an honest user and an adversary within one session. However, anonymity not only restricts the adversary, but also the simulation to not being able to identify public keys used to generate ciphertexts. This leads to the simulation’s inability to extract password guesses from ciphertexts forwarded by an adversary to sessions that are connected with the instance, with whom the adversary is interacting. This means that the adversary may trick the simulation into decrypting malicious ciphertexts and then revealing the decapsulated keys, which means that CCA security is indeed required.
2. Similarly for *indistinguishability*, consider an adversary that sees the public key and the ciphertext in an honest or simulated protocol execution. The adversary’s view on one connected session can be leveraged to forward values to other sessions using the same password and tricking the simulation into decrypting their adaptively chosen ciphertexts and then revealing the decapsulated keys, which means that CCA security is also required.

2 Related Work

Katz and Vaikuntanathan [34] introduced one of the earliest works on PQC PAKEs in 2009, which may be considered the first quantum-safe PAKE based on the lattice LWE problem and post-quantum Adaptive Smooth Projective Hash (ASPH) systems. The construction of Zhu et al. [32] in 2014 may similarly be considered the first PQC PAKE based on isogenies. However, the Ring (R)LWE PQC-PAK by Ding et al. [26] published in 2017, arguably marks the emergence of several other PAKE proposals from PQC primitives, more precisely based on lattice LWE (e.g. [50,56,29,23]) and isogenies (e.g. [52,51,1]). To our knowledge, the only PQC PAKE with security analysis in the Quantum Random Oracle Model (QROM) so far was given in 2024 by Lyu et al. [39]. The authors proposed the first PAKE protocol from isogeny assumptions in the Universal

Composability (UC) framework and the Random Oracle Model (ROM), as well as two PAKE protocols from lattice LWE and the group-action decisional Diffie-Hellman (GA-DDH) in the QROM. Their presented construction is based on lossy public key encryption (LPKE) and is extended to security in the QROM by replacing the basic LPKE with an extractable LPKE (eLPKE). Moreover, the authors apply a Fujisaki-Okamoto (FO) transformation to elevate the security of the chosen LPKE from IND-CPA to IND-CCA. The aforementioned works utilize PQC asymmetric primitives directly, i.e., they do not utilize PQC KEMs in a non-modified black-box manner. In the following, we relate to generic and black-box KEM-based designs in more detail.

In 2023, Beguinet et al. [11] presented the first generic PQC PAKEs, the CAKE and OCAKE protocols, and provided a security analysis of their constructions in the UC framework utilizing the IC and ROM models. The CAKE suite is a generic transformation from KEM to PAKE based on the classical PAKEs EKE and OEKE. The high-level idea is to encrypt the public key and the ciphertext using the password to provide explicit mutual authentication in the CAKE variant. Alternatively, in OCAKE, the ciphertext is authenticated with a key confirmation tag only, which provides explicit authentication for the receiver. CAKE and OCAKE require that the underlying used KEM fulfills the notions of Key Indistinguishability (IND-CPA), public key fuzziness (Fuzzy KEM), and ciphertext anonymity (Anonymous KEM). Pan and Zeng [42] as well as Alnahawi et al. [4] presented further security analysis for CAKE and OCAKE respectively. Unlike the UC proof of [11], the two additional security proofs were presented in the BPR model. Pan and Zeng [42] suggested the notion of Anonymity under Plaintext Checking Oracles (ANO-PCA) for the chosen KEM and extended the security proof to handle multi-user challenges. The authors of [4] also adapted similar anonymity and multi-user notions and formulated the notion of Public Key Uniformity (KEM-PKU) as a replacement for the Fuzzy-KEM property. Additionally, Alnahawi et al. provided mutual explicit authentication by adding a key confirmation round, and showed how to formally handle password guesses in a detailed game-based proof.

Dos Santos et al. [47] presented a new way to construct a UC-secure PAKE protocol under a relaxation of the IC called Half-Ideal Cipher (HIC). Their EKE-KEM protocol utilizes a KEM and a modified 2-round Feistel construction, which they call m2F. The m2F avoids using an IC over a group through defining an IC over a fixed-length bit-string domain. Following that, Arriaga et al. [10] introduced the Compact HIC (CHIC) protocol, which improves the construction of EKE-KEM [47] in computation and communication costs. The CHIC protocol utilizes the m2F construction in a white-box manner by using the public seed of a splittable KEM public key as the ephemeral randomness input used in the m2F construction of [47]. The authors in [10] also address the required KEM properties and define the notions of Passive One-Way Security (OW-CPA), Pseudo-Uniformity of Public Keys (UNI-PK), and Anonymity under Plaintext-Checkable Attacks (ANO-iPCA) for a chosen KEM.

Considering the IC model itself, relying on a quantum equivalent is not yet an option. In principle, it is impossible to utilize the IC capabilities (cf. Sec. 4) due to the quantum no-cloning theorem and the infeasibility of rewinding or back-patching. Very few works in the literature [2,33,48] address the notion of the *Quantum Ideal Cipher* (QIC) model, yet do not show how to fully match the classical one or the lazy sampling technique. This is mainly because they focus on one-way functions and non-invertible permutations. The work by Unruh [53] builds upon the idea of compressed function oracles (CFO) and takes a step forward in modeling keyed invertible permutations (i.e., IC) in quantum settings by introducing compressed permutation oracles (CPO). Nevertheless, despite the proposed approach’s plausibility, it is not yet formally proven that a CPO is indistinguishable from a truly random permutation [53].

3 Preliminaries

In this section, we present the definitions of the main building blocks of our protocol and describe their respective security properties.

3.1 Password Authenticated Key Exchange

Unlike key negotiation and agreement that are found in security protocols such as TLS and IKEv2, Password Authenticated Key Exchange (PAKE) protocols, as the name suggests, aim at establishing a session key between communication parties over insecure (not-trusted) channels without static public keys or certificates. Thus, a PAKE provides an alternative to static authentication using a low entropy *long-lived-key* (i.e., a password) known to communication parties as a pre-shared secret. Among several methods to perform this task [30], the very first PAKE construction, Encrypted Key Exchange (EKE) [13], relies on the simple idea of (symmetrically) encrypting the protocol initiator’s public key using the password, or a password derived symmetric key. The encrypted public key can only be decrypted and, hence, correctly used by a receiver who has the same password. Thus, it serves as authentication for both sides of the key agreement. However, PAKEs are vulnerable to two types of attacks: passive offline-dictionary attacks aiming to guess the correct password from a protocol transcript, and active attacks, where adversaries must be limited to online guessing on the low entropy password with a probability at most equal to $\frac{2}{E(\pi)}$.

3.2 Lattice Based Cryptography and LWE Related Problems

A lattice is a discrete subgroup under addition of $(\mathbb{R}^n, +)$ and can be described as a set of points in n -dimensional space with a periodic structure. Formally, given n -linearly independent vectors $v_1, \dots, v_n \in \mathbb{R}^n$ the generated lattice is the set of vectors $L(v_1, \dots, v_n) := \{\sum_{i=1}^n \alpha_i v_i \mid \alpha_i \in \mathbb{Z}\}$ constructed by all possible integer linear combinations. The basis of a lattice L can also be denoted by a base matrix \mathbf{A} representing its basis vectors. Lattices are of great interest to cryptographers

since there are several classical computational problems in lattices from which cryptosystems can be built [43]. The most prominent are the Shortest Vector Problem (SVP), Closest Vector Problem (CVP), the Smallest Integer Solution (SIS), and the Learning with Errors (LWE) problem. Nevertheless, public key cryptosystems from lattice problems are mostly based on worst-case reductions, making their choice of parameters stricter than average-case reductions [45]. Ultimately, even though lattices are defined in the Euclidean vector space \mathbb{R}^n , from a computational viewpoint, these problems are defined on integral lattices, whose representation is a matrix of integers. Hence, an irreducible polynomial f can replace a matrix base, and thus, a lattice can be defined as a special subset where all vectors form an ideal in a certain ring $\mathbb{Z}[x]/\langle f \rangle$. This led to extensive work on optimizing their key sizes in particular, and as a result several variations of lattice-based schemes dominated the NIST PQC process. We refer the reader to App. C for more details on LWE and its Ring and Module variants.

Non-uniform Learning with Errors (NLWE) Boneh et al. [16] introduced a variant of the learning with errors (LWE) problem in which the columns of \mathbf{A} (i.e. the LWE sample points) are sampled from a *non-uniform* distribution η over \mathbb{Z}_q^n called the *Non-Uniform Learning with Errors* (NLWE) problem, and showed that for suitable parameters, it is as hard as the basic LWE problem. In what follows, let k denote the dimension of the NLWE problem and let n denote the dimension of the LWE problem. Also, write η^m to denote m independent samples from the distribution η .

Definition 1 (Non-uniform Learning with Errors). *For an integer $q = q(k) \geq 2$, a noise distribution $\chi = \chi(k)$ over \mathbb{Z}_q , and a distribution η over \mathbb{Z}_q^k , the $\text{NLWE}_{\mathbb{Z}_q, k, \chi, \eta}$ problem is to distinguish between two distributions:*

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

where $m = \text{poly}(k)$, $\mathbf{A} \leftarrow \eta^m : \mathbf{A} \in \mathbb{Z}_q^{k \times m}$, $\mathbf{s} \leftarrow \mathbb{Z}_q^k$, $\mathbf{e} \leftarrow \chi^m$, and $\mathbf{u} \leftarrow \mathbb{Z}_q^m$.

Generally speaking, Boneh et al. [16] showed that, for any Probabilistic Polynomial Time (PPT) adversary, NLWE is as hard as LWE for any distribution η that is *coset samplable* defined as follows.

Definition 2 (Coset Samplable Distributions [16]). *For integers $q = q(k)$ and $n = n(k)$ we say that a distribution $\eta = \eta(k)$ over \mathbb{Z}_q^k is n -coset samplable if there are two PPT algorithms (MatrixGen, SamplePre) such that:*

- MatrixGen($1^k, n, q$) outputs a matrix $\mathbf{M} \in \mathbb{Z}_q^{n \times k}$ and auxiliary data \mathbf{T} .
- SamplePre($\mathbf{z} \in \mathbb{Z}_q^n, \mathbf{T}$) outputs a $\mathbf{y} \in \mathbb{Z}_q^k$ satisfying $\mathbf{M}\mathbf{y} = \mathbf{z}$.

Moreover, if \mathbf{z} is distributed *uniformly* in \mathbb{Z}_q^n , then the output of SamplePre(\mathbf{z}, \mathbf{T}) is distributed statistically close to η . That is, we have the following theorem: $\text{NLWE}_{\mathbb{Z}_q, k, \chi, \eta}$ is as hard as $\text{LWE}_{\mathbb{Z}_q, n, \chi}$ for any n -coset samplable distribution η .

Theorem 1 ([16]). *Let $\eta = \eta(k)$ be an n -coset samplable distribution. Suppose there is a PPT algorithm \mathcal{A} that decides $\text{NLWE}_{\mathbb{Z}_q, k, \chi, \eta}$ with advantage $\varepsilon = \varepsilon(k)$. Then, there is a PPT algorithm \mathcal{B} that decides $\text{LWE}_{\mathbb{Z}_q, n, \chi}$ with the same advantage $\varepsilon = \varepsilon(k)$.*

Following [16, Remark 4.4], we especially point out that the above holds even when \mathbf{s} is distributed according to the error distribution rather than uniform.

Weak Hint Learning with Errors (whLWE) Cheon et al. [24] (resp. Lee et al. [36]) first defined Hint LWE (hLWE) and Weak-Hint LWE (whLWE) – with Liu et al. [38] providing a valuable discretization step in a recent preprint – as a potential enhancement of eLWE-type leakage-security properties, based on the following statistical insight about conditional Gaussian distributions:

Lemma 1 ([36], Lemma 4.8). *Let D_s denote a continuous Gaussian distribution with variance s . Let $D_{c,s}$ denote a continuous Gaussian distribution with center c and variance s . Then, for real numbers $\sigma_1, \sigma_2 > 0$, let e and f be random variables distributed as the Gaussian distributions D_{σ_1} and D_{σ_2} , respectively. Let $\sigma = \sqrt{\sigma_1^2 + \sigma_2^2}$, then the tuple $(e + f, e|(e + f))$ is distributed as the joint conditional Gaussian distribution $(D_\sigma, D_{L\sigma_1^2/\sigma^2, \sigma_1\sigma_2/\sigma})$ with L denoting $e + f$.*

We define the Weak-Hint LWE (whLWE) problem following [24,36]:

Definition 3 (Weak-Hint LWE). *Let n, q , and k be positive integers, $\sigma_1, \sigma_2 > 0$ be real numbers, \mathbf{z} be a vector in Domain (with Domain arbitrary – but looking forward, we will choose $\text{Domain} := \mathbb{Z}_q^k$) and S be a matrix in $\mathbb{Z}_q^{n \times k}$. The Weak-Hint LWE distribution, denoted $A_{n,q,\sigma_1,\sigma_2,k}^{\text{whLWE}}(\mathbf{z}, S)$ is the distribution of $(\mathbf{a}, S^t \mathbf{a} + \mathbf{e}, \langle \mathbf{z}, \mathbf{e} \rangle + f) \in \mathbb{Z}_q^n \times \mathbb{R}_q^k \times \mathbb{R}_q$, where $\mathbf{a} \leftarrow \mathbb{Z}_q^n$, $\mathbf{e} \leftarrow D_{\sigma_q}^k$, and $f \leftarrow D_{\sigma_2}$, and where D is the secret distribution. Then the Weak-Hint LWE problem $\text{whLWE}_{n,q,\sigma_1,\sigma_2}^k(D)$ is to distinguish, given arbitrarily many independent samples $\mathbf{z} \leftarrow \text{Domain}$ chosen by an adversary, between $A_{n,q,\sigma_1,\sigma_2,k}^{\text{whLWE}}(\mathbf{z}, S)$ for a fixed $S \leftarrow D$ and the distribution of $(\mathbf{a}, \mathbf{u}, \langle \mathbf{z}, \mathbf{e} \rangle + f)$ where $\mathbf{u} \leftarrow \mathbb{R}_q$.*

Using Lemma 1, Cheon et al. [24,36] are able to show the following security theorem, which we will leverage in what follows.

Theorem 2. *Let n, q, k be positive integers. Let $\sigma_1, \sigma'_1, \sigma'_2$ be positive real numbers that satisfy $\sigma_1 = \sigma'_1 \sigma'_2 / \sqrt{(\sigma'_1)^2 + (\sigma'_2)^2}$. Finally, let D be a distribution over $\mathbb{Z}_q^{n \times k}$. Then there exists a polynomial-time reduction from $\text{LWE}_{n,q,\sigma_1}^k(D)$ to $\text{whLWE}_{n,q,\sigma'_1 \sqrt{k} \sigma'_2}^k(D)$, which exactly preserves the adversary's advantage.*

3.3 Key Encapsulation Mechanism (KEM)

A KEM allows a key agreement initiator (encapsulator) to convey a shared secret key K via a ciphertext C to the receiver (decapsulator). A KEM can then be used in black-box (generic) manner through its provided interfaces as follows.

Definition 4 (Key encapsulation mechanism (KEM)). A KEM is a triple $(KGen, Encap, Decap)$:

- $KGen(1^\kappa) \rightarrow (pk, sk)$: On input security parameter κ **return** key pair pk, sk (probabilistic algorithm),
- $Encap(pk) \rightarrow (C, K)$: On input pk **return** ciphertext $C \in \mathcal{C}$ and key $K \in \mathcal{K}$ (probabilistic algorithm).
- $Decap(sk, C) \rightarrow K/\perp$: On input sk and ciphertext C **return** key $K \in \mathcal{K}$ or rejection $\perp \notin \mathcal{K}$ (deterministic algorithm).

Correctness of KEM: The correctness of a KEM means that the decapsulation algorithm $KEM.Decap$ recovers the same shared key K produced by the encapsulation algorithm $KEM.Encap$ for a public key generated by the key generation algorithm $KEM.KGen$. That is, however, except for a small probability over the space of key generation and encapsulation.

Definition 5 (KEM Correctness). We say a KEM is $(1-\delta)$ correct, if for every key pair $(pk, sk) \leftarrow \$ KGen(1^\kappa)$ and every encapsulation $(C, K) \leftarrow \$ Encap(pk)$ we have:

$$\Pr[K' \neq K | K' \leftarrow Decap(sk, C)] \leq \delta.$$

A KEM is perfectly correct if $\delta = 0$.

Security of KEM The basic security of a KEM is defined in terms of indistinguishability of encapsulated keys from random keys. The anonymity, robustness and collision freeness of a KEM also play an essential role in the context of protocol design, as they capture properties beyond key security and protect against other chosen ciphertext attacks. Further, distinguishing honestly generated public keys from uniformly random ones for a KEM was introduced in [11] and [42] as the *Fuziness* property. It was formalized in [4] as the notion of public key uniformity, which we also adopt in this work. The KEM security properties relevant for our construction are described in the experiments $\text{Exp}_{KEM}^{IND-CCA}$, Exp_{KEM}^{PKU} , $\text{Exp}_{KEM}^{ANO-CCA}$ and $\text{Exp}_{KEM}^{SCFR-CCA}$ respectively in Fig. 2.

Definition 6 (Key Uniformity of KEM:). For a KEM with public key space \mathcal{PK} , we define the advantage of an adversary \mathcal{A} in distinguishing honestly generated public keys from uniform random ones against the KEM as

$$\text{Adv}_{KEM}^{PKU}(\mathcal{A}) := |\Pr[PKU^0(\mathcal{A})] - \Pr[PKU^1(\mathcal{A})]|$$

where $\text{Exp}_{KEM}^{PKU}(\mathcal{A})$ is a security game relative to a challenge bit b (Fig. 2).

Definition 7 (IND-CCA security of KEM). For a KEM with key space \mathcal{K} , we define the advantage of an adversary \mathcal{A} in distinguishing honestly encapsulated keys from uniform random ones against the KEM as

$$\text{Adv}_{KEM}^{IND-CCA}(\mathcal{A}) := |\Pr[IND-CCA^0(\mathcal{A})] - \Pr[IND-CCA^1(\mathcal{A})]|$$

where $\text{Exp}_{KEM}^{IND-CCA}(\mathcal{A})$ is a security game relative to a challenge bit b (Fig. 2).

$\text{Exp}_{\text{KEM}}^{\text{PKU}^b}(\mathcal{A})$	$\text{Exp}_{\text{KEM}}^{\text{IND-CCA}^b}(\mathcal{A})$
1 : $(pk_0, sk_0) \leftarrow \$ \text{KGen}(1^\kappa)$	1 : $(pk, sk) \leftarrow \$ \text{KGen}(1^\kappa)$
2 : $pk_1 \leftarrow \$ \mathcal{PK}$	2 : $(C^*, K_0) \leftarrow \$ \text{Encap}(pk)$
3 : $b \leftarrow \$ \{0, 1\}$	3 : $K_1 \leftarrow \$ \mathcal{K}$
4 : $b' \leftarrow \mathcal{A}(pk_b)$	4 : $b \leftarrow \$ \{0, 1\}$
5 : return $b = b'$	5 : $b' \leftarrow \mathcal{A}^{\text{D}(\cdot, \cdot)}(pk, C^*, K_b)$
	6 : return $b = b'$
$\text{Exp}_{\text{KEM}}^{\text{ANO-CCA}^b}(\mathcal{A})$	$\text{Exp}_{\text{KEM}}^{\text{SCFR-CCA}}(\mathcal{A})$
1 : $(pk_0, sk_0) \leftarrow \$ \text{KGen}(1^\kappa)$	1 : $(pk_0, sk_0) \leftarrow \$ \text{KGen}(1^\kappa)$
2 : $(pk_1, sk_1) \leftarrow \$ \text{KGen}(1^\kappa)$	2 : $(pk_1, sk_1) \leftarrow \$ \text{KGen}(1^\kappa)$
3 : $b \leftarrow \$ \{0, 1\}$	3 : $C \leftarrow \$ \mathcal{A}^{\text{D}(\cdot, \cdot)}(pk_0, pk_1)$
4 : $(C^*, K^*) \leftarrow \$ \text{Encap}(pk_b)$	4 : $K_0 \leftarrow \text{Decap}(pk_0, sk_0, C)$
5 : $b' \leftarrow \mathcal{A}^{\text{D}(\cdot, \cdot)}(pk_0, pk_1, (C^*, K^*))$	5 : $K_1 \leftarrow \text{Decap}(pk_1, sk_1, C)$
6 : return $b = b'$	6 : return $K_0 = K_1 \neq \perp$

Fig. 2. KEM Security Experiments $\text{Exp}_{\text{KEM}}^{\text{PKU}}$, $\text{Exp}_{\text{KEM}}^{\text{IND-CCA}}$, $\text{Exp}_{\text{KEM}}^{\text{ANO-CCA}}$ and $\text{Exp}_{\text{KEM}}^{\text{SCFR-CCA}}$. The experiments for ANO-CCA and SCFR-CCA are adopted from [28], where the decryption oracle D takes a ciphertext and a bit that selects which secret key to use and disallows queries on the challenge ciphertext in the ANO-CCA experiment. The Decap routine in $\text{Exp}_{\text{KEM}}^{\text{SCFR-CCA}}$ explicitly takes an input parameter pk_i as in [25].

Definition 8 (ANO-CCA security of KEM). For a KEM with public key space \mathcal{PK} , we define the advantage of an adversary \mathcal{A} in distinguishing honest public keys used to probabilistically encapsulate keys into ciphertexts against the KEM as

$$\mathbf{Adv}_{\text{KEM}}^{\text{ANO-CCA}}(\mathcal{A}) := |\Pr[\text{ANO-CCA}^0(\mathcal{A})] - \Pr[\text{ANO-CCA}^1(\mathcal{A})]|$$

where $\text{Exp}_{\text{KEM}}^{\text{ANO-CCA}}(\mathcal{A})$ is a security game relative to a challenge bit b (Fig. 2).

Definition 9 (SCFR-CCA security of KEM). For a KEM with public key space \mathcal{PK} and key space \mathcal{K} , we define the advantage of an adversary \mathcal{A} in probabilistically generating a ciphertext C that decapsulates correctly under two unique public keys and their corresponding secret keys against the KEM as

$$\mathbf{Adv}_{\text{KEM}}^{\text{SCFR-CCA}}(\mathcal{A}) := \Pr[\text{SCFR-CCA}(\mathcal{A})]$$

where $\text{Exp}_{\text{KEM}}^{\text{SCFR-CCA}}(\mathcal{A})$ is a collision-finding security game (Fig. 2).

KEM with Implicit Rejection: Most NIST PQC KEMs use variants of the Fujisaki–Okamoto (FO) transformation with implicit rejection, where the decryption algorithm, unlike in explicitly rejecting KEMs, still outputs a random

key from the same key space, even if the decryption fails. Here we note that KEMs with implicit rejection satisfying the notion of SCFR-CCA cannot satisfy the robustness property SROB-CCA [28,25]. Although the notion of robustness is defined via an almost identical security experiment as collision freeness, the fact that the latter rules out rejecting a decapsulated key through the provided KEM interface cannot be interpreted in the same way as in the robustness property. Whereas SROB-CCA guarantees that a ciphertext cannot decapsulate correctly under two different key pairs, SCFR-CCA only ensures that a ciphertext cannot produce the same key under two different key pairs. Additionally, CPA versions of the aforementioned properties can be obtained through omitting the decryption oracles from the experiments.

4 The Problem with IC and KEM Public Key Issues

The ideal cipher model arguably dates back to Shannon [49] and was formalized in PAKE settings in the BPR model [12]. It has been widely used for security proofs of cryptographic protocols instantiated with block ciphers (cf. Sec. 2). Thus, an IC serves for modeling block ciphers (e.g., AES) as idealized objects similar to hash functions in the ROM with some exceptions [14]. Its main advantage is defining the behavior of a cipher, where each encryption maps to an independently random permutation that belongs to the same set of possible input values. In other words, it models a random block cipher as being chosen uniformly from the set of all possible block ciphers [14].

Definition 10 (Ideal Cipher). *IC is a function $\mathcal{C} : \{0,1\}^\kappa \times \{0,1\}^n \rightarrow \{0,1\}^n$, where each key $k \in \{0,1\}^\kappa$ defines a unique and independent random permutation $\mathcal{C}_k = \mathcal{C}(k, \cdot)$ on $\{0,1\}^n$.*

Similar to a RO, an IC provides oracle access for forward queries of the form \mathcal{C} . However, the main difference here is that an IC also provides oracle access for backward queries of the form \mathcal{C}^{-1} . Thus, the oracle answers queries to both encryption and decryption and models invertible random permutations through an interface accessible by both adversaries and challengers in a protocol instance. Queries and responses to and from the IC oracle are modeled via lazy sampling and are stored in a list, which can be used in steps of a protocol analysis to determine further action. This behavior is crucial in formulating sound proofs for three main reasons: 1. Keeping a record of both honestly and maliciously generated public values, ciphertexts, and secrets, 2. replacing values generated in a simulation with random values generated by a challenger, and 3. performing consistency checks across proof steps in both the UC and BPR frameworks.

IC Issues in PAKEs: Narrowing down the issue at hand, and ignoring the QIC for a moment, the usage of IC for public key authentication in a PAKE models random permutations as values in the set of all possible public keys, such that $IC : \mathcal{PK} \rightarrow \mathcal{PK}$. This behavior ensures that decryption always yields a valid (i.e., possibly honestly generated) public key, regardless of the correctness of

the password used for encrypting or decrypting this key. Further, the usage of IC in PAKE ensures well-formatted public keys, with the consequence that adversaries cannot use malicious keys to attack the security of the underlying asymmetric scheme. This results in problems when instantiating an IC with real-world ciphers (such as AES) to encrypt structured keys (such as MLWE-based) as follows: 1. Encrypting (structured) keys under one key and decrypting them under another key may very well result in values that do not exist in the key space (i.e., invalid keys). 2. Malicious (non-well-formed) keys may very well be used by adversaries to break the underlying scheme.

Practical Example: The following describes these problems in more detail. We devise a very simple PAKE construction as shown in Fig. 3, which resembles EKE or CAKE, but with no IC involved for that matter. In the authentication step, one would simply extend the password using a hash function to a fixed length bit string and then XOR the output with the public key. At first glance, such constructions seem very attractive and promise a smooth security proof relying solely on well-studied KEM properties (e.g., IND-CCA and ANO-CPA). However, not relying on IC properties (e.g., bijection and record keeping) quickly reveals that this is at least questionable.




Fig. 3. Simple NICE-PAKE

We consider a malicious *Bob* who receives a masked pk as z as shown in Fig. 3. *Bob* may guess a password π' and observe the result of the unmasking of z . With probability bound to the password dictionary size, it is likely that *Bob* guessed wrong. Assuming that not all z_i map to a valid pk_i under all π_i , he may observe an invalid pk' , which may be detectable, e.g., by checking the (wrong) structure. Then, *Bob* can exclude this one guess from an offline dictionary and take another guess in the same active session, and thus instantly break the PAKE security. In the specific case of Kyber, the the $\mathbf{KEM.KGen}$ algorithm generates a public key that mathematically resembles a vector of polynomials (i.e., the \mathbf{b} -part of the public key). Each polynomial has coefficients with values in \mathbb{Z}_q where

$q = 3329$ and is stored as 16-bit integer. This vector representation of the public key is then serialized into a byte array that maps two coefficients to three 8-bit integers. Unmasking the \mathbf{b} -part may however yield a byte array that does not follow this structure, and thus to a polynomial with coefficients outside the range defined by the scheme parameters, which in turn leads to invalid public keys and thus enables detecting wrong password guesses. Considering FrodoKEM, this is not the case, as its public keys are (close to) indistinguishable from random bit strings and do not yield a detectable mathematical structure. Now, we consider a malicious *Alice* and a KEM where all z_i map to valid pk_i under all π_i . Note that a malicious *Alice* is neither bound to honestly use KEM.KGen, nor to honestly calculate z . Thus, she may not necessarily be bound by the security properties of the KEM that may rely on honest key generation. Then, she could generate one or many malicious (non-well-formed) key pairs pk_i, sk_i and/or one or many malicious z_i and other data she may use to extract information from C_i . Note that these real-world problems vanish when modeling the key authentication using IC. An approach to prevent both attacks without an IC is restricting the PAKE to KEMs with splittable keys of the form $(\mathbf{A}, \mathbf{b}) : \mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e}$ (e.g., LWE) and enforcing an undetectable mismatch between \mathbf{A} and \mathbf{b} under wrong guesses of π for a malicious Bob, and yielding an unusable C_b for a malicious Alice. In the following, we investigate how this mismatch can be achieved.

5 Lattice KEMs with Splittable Public Keys

KEMs based on the lattice LWE problem and its variants (e.g. RLWE and MLWE) have public keys of the form $pk : (\mathbf{A}, \mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e})$, where the matrix \mathbf{A} defines the base of the lattice, and is sampled from a random seed (fixed-length bit string usually referred to as ρ) that is appended (or prepended) to the public key. Based on the work in [10], we adopt the notion of KEM with splittable and pseudo-uniform public keys (UNI-PK) and adapt it to our use case as SPLIT-PKU in the following:

Definition 11 (KEM with Splittable Public Keys). *We say a KEM scheme with public keys of the form $pk : (\mathbf{A}, \mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e})$ is a KEM with uniform splittable public keys if the public key is defined over the domains $\mathcal{PK}, \mathcal{G}$, and the split is an invertible map such that: 1) $G \in \mathcal{G}$ is a hash function modeled as a RO, 2) $\forall (pk, sk) \leftarrow KGen : pk \in \mathcal{PK}$, and 3) $(\mathbf{A}, \mathbf{b}) \leftarrow pk : \mathbf{A} \leftarrow G(\rho)$ are identical, using the same random bit string seed ρ , and \mathbf{A} is generated uniformly from $G(\rho)$ using an extended output function (XOF).*

Key Uniformity of KEM with Splittable Public Keys: Based on the general KEM public key uniformity (cf. Sec. 3.3 Def. 6), we derive a similar notion for KEM with splittable public keys, the security experiment of which is shown in Fig. 4. We claim with this extension of the original notion, and the UNI-PK notion of Arriaga et al. [10], that an adversary viewing only a \mathbf{b} -part of the public key is not able to distinguish the \mathbf{A} -part embedded in it as $\mathbf{A}\mathbf{s} + \mathbf{e}$ from one

that is chosen uniformly at random. As the seed is simply a random bit string, and the expansion functions and rejection sampling (in Kyber) are modelled as ROs, the keys in splittable KEMs can be viewed as uniform under the standard decisional LWE, RLWE, and MLWE assumptions in the ROM [10]. We also note, that FrodoKEM does not require any rejection sampling in its matrix expansion, which makes it even simpler [10]. The proof of this property for MLWE (e.g., Kyber-like) schemes follows from [10], and can be similarly shown for LWE (e.g., FrodoKEM) from standard LWE assumption and is therefore omitted.

Definition 12 (Key Uniformity of KEM with Splittable Public Keys):
For a KEM with splittable public key of the form $(\mathbf{A}, \mathbf{b}) : \mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e}$ and \mathbf{A} -part space \mathbb{A} , we define the advantage of an adversary \mathcal{A} in distinguishing honestly embedded \mathbf{A} -part in public keys from uniformly random ones as

$$\mathbf{Adv}_{\text{KEM}}^{\text{SPLIT-PKU}}(\mathcal{A}) := |\Pr[\text{SPLIT-PKU}^0(\mathcal{A})] - \Pr[\text{SPLIT-PKU}^1(\mathcal{A})]|.$$

where $\text{Exp}_{\text{KEM}}^{\text{SPLIT-PKU}}(\mathcal{A})$ is a security game relative to a challenge bit b (Fig. 4).

$\text{Exp}_{\text{KEM, Split}}^{\text{UNI-PK}^b}(\mathcal{A})$	$\text{Exp}_{\text{KEM}}^{\text{SPLIT-PKU}^b}(\mathcal{A})$
1 : $(pk_0, -) \leftarrow \$ \text{KGen}(1^\kappa)$	1 : $(pk_0, sk_0) \leftarrow \$ \text{KGen}(1^\kappa)$
2 : $(r_0, M_0) \leftarrow \text{Split}(pk_0)$	2 : $(\mathbf{A}_0, \mathbf{b}_0) \leftarrow pk_0$
3 : $(r_1, M_1) \leftarrow N_\kappa \times G_\kappa$	3 : $\mathbf{A}_1 \leftarrow \$ \mathbb{A}$
4 : $b \leftarrow \$ \{0, 1\}$	4 : $b \leftarrow \$ \{0, 1\}$
5 : $b' \leftarrow \mathcal{A}(r_b, M_b)$	5 : $b' \leftarrow \mathcal{A}(\mathbf{A}_b, \mathbf{b}_b)$
6 : return $b == b'$	6 : return $b == b'$

Fig. 4. Key Uniformity for KEM and Splittable KEM - UNI-PK is adopted from [10] and preserving the original notation where N maps a bit string of a fixed length r to an element M in $G = R_q^k$. \mathbb{A} in SPLIT-PKU denotes the set of uniform distributed lattice matrices (i.e., the lattice base space).

A-Part-Secrecy for KEM with Splittable Public Keys For a given KEM with splittable public keys, the \mathbf{A} -part is used by the encapsulator both indirectly for encrypting a message m and directly for encrypting a randomness \mathbf{r} . We recall that such KEMs (essentially LWE, RLWE, and MLWE based KEMs) internally use a PKE for encryption and decryption (Fig. 5).

KGen(1^κ)	Encap(pk)	Decap(sk', C)
$(pk, sk) \leftarrow \text{\$} \text{KGen}(1^\kappa)$	$m \leftarrow \text{\$} \mathcal{M}$	$m' \leftarrow \text{PKE.Dec}(C, sk)$
$s \leftarrow \text{\$} \{0, 1\}^{ s }$	$(r, k) \leftarrow G_2(G_1(pk), m)$	$(r', k') \leftarrow G_2(G_1(pk), m')$
$sk' \leftarrow (s, sk, pk, G_1(pk))$	$C \leftarrow \text{PKE.Enc}(pk, m; r)$	$C' \leftarrow \text{PKE.Enc}(pk, m'; r')$
return (pk, sk')	$K \leftarrow F(k, C)$	if $C' = C$ then
	return (C, K)	return $K' \leftarrow F(k', C)$
		else return $K' \leftarrow F(s, C)$

Fig. 5. Construction of IND-CCA KEM from IND-CPA PKE with message space \mathcal{M} and FO Transform using the hash functions G_1, G_2 and F - Adapted from [17] and [28].

When decapsulating a ciphertext C to obtain K , the decapsulator also learns the message m that the underlying PKE encrypted. This is unavoidable for any PKE-based KEM with splittable public keys because the ciphertext that contains m needs to be decrypted before deriving K from m . Further, the encryption yields a ciphertext of the form (\mathbf{u}, \mathbf{v}) , where \mathbf{v} represents the encrypted message m chosen by the encapsulator using the \mathbf{b} -part, and \mathbf{u} is the encryption of the encapsulators own uniformly chosen randomness. To decrypt (\mathbf{u}, \mathbf{v}) , the decapsulator must use the same \mathbf{A} -part as the encapsualtor, otherwise they cannot decrypt the message m . We rely on the fact that the \mathbf{u} -part in C resembles a uniformly chosen sample that belongs to the same hardness assumption of the chosen KEM, so that it is not possible to determine which \mathbf{A} was used to create \mathbf{u} . This assumption plays an essential role in the security proof of our protocol, as described in Sec. 6. Formally, we propose the assumption that using a different (non-matching) yet uniform \mathbf{A} -part for the direct encryption of \mathbf{r} results in ciphertexts that cannot be decrypted correctly and therefore cannot be used to reveal a non-related uniformly sampled \mathbf{A} . We formalize this assumption as the notion of KEM \mathbf{A} -Part Secrecy (A-SEC-CCA) as depicted in Fig. 6, and provide proof of its correctness, but rather tailored to our use case, namely in a PAKE.

Definition 13 (Splittable KEM \mathbf{A} -Part-Secrecy:). For a KEM with public key space \mathcal{PK} and \mathbf{A} -part space \mathbb{A} , we define the advantage of an adversary \mathcal{A} in distinguishing a random \mathbf{A} -part of a known public key used to probabilistically generate a ciphertext as

$$\mathbf{Adv}_{\text{KEM}}^{\text{A-SEC-CCA}}(\mathcal{A}) := |\Pr[\text{A-SEC-CCA}^0(\mathcal{A})] - \Pr[\text{A-SEC-CCA}^1(\mathcal{A})]|$$

where $\text{Exp}_{\text{KEM}}^{\text{A-SEC-CCA}}(\mathcal{A})$ is a security game relative to a challenge bit b (Fig. 6).

Provable Concrete Security of a FrodoKEM Instantiation In the following, we show that the \mathbf{A} -Part-Secrecy (A-SEC-CCA) property applies to plain LWE and give concrete numbers for bit security using FrodoKEM.

$\text{Exp}_{\text{KEM}}^{\text{A-SEC-CCA}^b}(\mathcal{A})$	$\text{Exp}_{\text{KEM}}^{\text{A-SEC-CCA}^{(b, \mathcal{D})}}(\mathcal{A})$
1 : $(pk, sk) \leftarrow \mathcal{A}$	1 : $(pk, sk) \leftarrow \mathcal{A}$
2 : $\mathbf{A}_0 \leftarrow \$ \mathbb{A}$	2 : $\mathbf{A}_0 \leftarrow \$ \mathcal{D}$
3 : $\mathbf{A}_1 \leftarrow \$ \mathbb{A}$	3 : $\mathbf{A}_1 \leftarrow \$ \mathbb{A}$
4 : $b \leftarrow \$ \{0, 1\}$	4 : $b \leftarrow \$ \{0, 1\}$
5 : $(C, K) \leftarrow \text{KEM.Encap}_{\mathbf{A}_b}(pk)$	5 : $(C, K) \leftarrow \text{KEM.Encap}_{\mathbf{A}_b}(pk)$
6 : $b' \leftarrow \mathcal{A}(C, pk, sk, \mathbf{A}_0, \mathbf{A}_1)$	6 : $b' \leftarrow \mathcal{A}(C, pk, sk, \mathbf{A}_0, \mathbf{A}_1)$
7 : return $b == b'$	7 : return $b == b'$

Fig. 6. A-Part-Secrecy Experiment (Left) and w.r.t. Password Dictionary (Right), where \mathbb{A} and \mathcal{D} denote the lattice matrix space and password dictionary space respectively, and $\text{KEM.Encap}_{\mathbf{A}}(pk)$ denotes invoking `Encap` explicitly with a matrix \mathbf{A}_b other than the one embedded in the adversary’s pk .

Theorem 3. *Assuming the hardness of both the Non-uniform LWE and Weak-Hint LWE problems and modeling SHAKE as a random oracle, a FrodoKEM-style Plain LWE KEM (with appropriate parameters as discussed in Sec. 7) is A-SEC-CCA-secure.*

Remark on property idea and purpose. Intuitively, this notion would yield a security experiment similar to that of ANO-CCA (Fig. 2). However, we want to explicitly capture the case where an adversary controls the key generation process, yet is restricted to forwarding a masked seed ρ bound by the password space for sampling an \mathbf{A} -part. By doing so, we show that even with a prepared list of all possible passwords and their hash outputs, an adversary cannot leverage a brute-force or an offline dictionary attack to guess the correct password of an honest user (i.e., the challenger in this experiment). The trick here is three-fold: First, the masked seed is an unstructured random bit string, and when unmasked, will always lead back to a valid, however not necessarily correct, password from a well spread and close to uniform distribution \mathcal{D} . Second, the \mathbf{A} -part sampled afterwards is used in the encapsulation regardless of the password guess (assuming the adversary guessed wrong), and thus creating a mismatch in the KEM encapsulation with overwhelming probability bound by collisions on password hash outputs and the hardness of LWE. Finally, an adversary cannot determine the \mathbf{A} -part in an honest ciphertext, and hence cannot connect it to a seed and its corresponding password.

Remark on concrete A-SEC-CCA-security. We will show that the concrete security loss (see also, Section 7) using FrodoKEM, as presented in NIST PQC Round 3, is only a few bits of security. Specifically, it depends on concrete parameters of the associated whLWE problem. Full NIST Category 1, 3, and 5 concrete security levels can be achieved by mild re-parameterizations of FrodoKEM (within its well-defined design space) inducing a minimal loss in practical efficiency.

Remark on terminology. We will refer to FrodoKEM in what follows, with the understanding that we consider either using FrodoKEM as defined but lose a few bits of security, or that we use an alternative “GimliKEM,”³ with slightly larger keys and ciphertexts, where the security proof equally applies but with no loss in concrete security.

Proof. To show A-SEC-CCA security, we use the common game-based approach, starting with the original security game \mathbf{G}_0 , and transitioning through a series of experiments $\mathbf{G}_1, \mathbf{G}_2, \dots$ while bounding the adversary’s advantage between each game. The final game will be independent of the challenge bit b , giving the theorem. Specifically, we will aim to swap the portion of the ciphertext C that depends on \mathbf{A}_b to uniform.

Game \mathbf{G}_0 : This is the original game. The adversary chooses pk, sk arbitrarily. While $pk = (\mathbf{A}, \mathbf{b})$, only \mathbf{b} is used in constructing C . Then, \mathbf{A}_0 is sampled from distribution \mathcal{D} , which is a well-spread depending on the entropy of the password π with support in the \mathbf{A} -space of FrodoKEM, and \mathbf{A}_1 is sampled uniformly from the \mathbf{A} -space of FrodoKEM. The challenger flips a bit b then constructs the challenge ciphertext C (ignoring the message component) of the form:

$$\begin{aligned} C &= (u, v) \\ u &= \mathbf{A}_b \mathbf{r} + \mathbf{e}' \\ v &= \mathbf{b} \mathbf{r} + \mathbf{e}'' \end{aligned}$$

The adversary’s view in \mathbf{G}_0 is $(C, pk, sk, \mathbf{A}_0, \mathbf{A}_1)$ and outputs a guess b' as to which matrix \mathbf{A}_b was used in constructing C .

Game \mathbf{G}_1 : In this hybrid, we swap \mathbf{A}_0 to be sampled uniformly from the \mathbf{A} -space of FrodoKEM. The rest of the experiment is the same.

Lemma 2 ($\mathbf{G}_0 \stackrel{\text{comp}}{\approx} \mathbf{G}_1$). *If SHAKE is modeled as a random oracle and Non-uniform LWE is $(1 - \text{negl}(1^\kappa))$ -hard, then $\Pr_{\mathbf{G}_0}[\mathcal{A} \text{ wins}] \leq \Pr_{\mathbf{G}_1}[\mathcal{A} \text{ wins}] + \text{negl}(1^\kappa)$.*

Proof (Proof of Lemma 2). The difference in the games is the distribution of \mathbf{A}_0 . In the first game, \mathcal{A} ’s view of \mathbf{A}_0 is no worse than having sampled \mathbf{A}_0 as the output of a random oracle, which was given as input a randomly sampled password π . We claim that, conditioned on the entropy of π being well-spread and uniform distributed in \mathcal{D} , this distribution is n -coset samplable in a straightforward way. Following [16, Remark 4.4]: even when \mathbf{r} is distributed according to the error distribution, NLWE is hard. Yet, the adversary additionally has leakage on \mathbf{r} in the form of $v = \mathbf{b} \mathbf{r} + \mathbf{e}''$. However, since \mathbf{b} is chosen by the adversary *non-adaptively* – that is, before seeing $(\mathbf{A}_0, \mathbf{A}_1)$ – this leakage is independent of the n -coset samplability of \mathbf{A}_0 . Therefore, n -coset samplability of \mathbf{A}_0 follows solely from modeling SHAKE as a random oracle, so distinguishing \mathbf{G}_0 and \mathbf{G}_1 is as hard as NLWE.

³ Gimli was another member of the Fellowship, who was slightly larger than Frodo but had a nice axe.

Game \mathbf{G}_2 : In this hybrid, we swap the u -part of the challenge ciphertext C to uniform. The rest of the experiment is the same.

Lemma 3 ($\mathbf{G}_1 \stackrel{\text{comp}}{\approx} \mathbf{G}_2$). *For $\varepsilon_{\text{whLWE}} < 1$, if Weak-Hint LWE is $(1 - \varepsilon_{\text{whLWE}})$ -hard, then $\Pr_{\mathbf{G}_1}[\mathcal{A} \text{ wins}] \leq \Pr_{\mathbf{G}_2}[\mathcal{A} \text{ wins}] + \varepsilon_{\text{whLWE}}$.*

Proof (Proof of Lemma 3). The difference between the games is the distribution of u . In \mathbf{G}_1 , we have $u = \mathbf{A}_b \mathbf{r} + \mathbf{e}'$, and in \mathbf{G}_2 , we have u sampled uniformly at random. The adversary's view of the challenge ciphertext $C = (u, v)$ includes the hint $v = \mathbf{b} \mathbf{r} + \mathbf{e}''$ on the ciphertext's secret \mathbf{r} that defines the LWE instance considered, where \mathbf{b} is *arbitrarily chosen* by the adversary. The adversary chooses \mathbf{b} and hands it to the challenger. The challenger flips a coin and returns either

$$\begin{array}{ccc} C = (u, v) & & C = (u, v) \\ u = \mathbf{A}_b \mathbf{r} + \mathbf{e}' & \text{or} & u = \text{uniform} \\ v = \mathbf{b} \mathbf{r} + \mathbf{e}'' & & v = \mathbf{b} \mathbf{r} + \mathbf{e}'' \end{array}$$

In the first case, the challenger simulates \mathbf{G}_1 ; in the second case, the challenger simulates \mathbf{G}_2 . Therefore, the adversary's advantage in distinguishing the hybrids is its advantage $\varepsilon_{\text{whLWE}}$ against the Weak-Hint LWE problem with appropriate parameters (concrete parameters discussed in Sec. 7).

Completing the proof of Theorem 3. Finally, in \mathbf{G}_2 , we note that the adversary's view is independent of the bit b in the A-SEC-CCA $^{(b, \mathcal{D})}$ security experiment. Therefore, $\Pr_{\mathbf{G}_2}[\mathcal{A} \text{ wins}] = 1/2$. Combining Lemmas 2 and 3, we have the probability that the adversary wins in the original A-SEC-CCA security game is $\Pr_{\mathbf{G}_0}[\mathcal{A} \text{ wins}] \leq 1/2 + \text{negl}(1^\kappa) + \varepsilon_{\text{whLWE}}$, which gives Theorem 3. \square

Collision Freeness for KEM with Splittable Public Keys Analogous to the formerly introduced A-SEC-CCA property, we suggest the notion of A-Part-Collision-Freeness A-CFR-CCA, or Splittable Collision Freeness.

Definition 14 (Splittable KEM Collision Freeness). *For a KEM with public key space \mathcal{PK} and \mathbf{A} -part space \mathbb{A} , we define the advantage of an adversary \mathcal{A} in probabilistically generating a ciphertext C that decapsulates correctly under two unique \mathbf{A} -parts under the same \mathbf{b} -part and the same secret key sk as*

$$\mathbf{Adv}_{\text{KEM}}^{\text{A-CFR-CCA}}(\mathcal{A}) := \Pr[\text{A-CFR-CCA}(\mathcal{A})]$$

where $\text{Exp}_{\text{KEM}}^{\text{A-CFR-CCA}}(\mathcal{A})$ is a collision-finding security game (Fig. 7).

Theorem 4. *Assuming the hardness of SCFR-CCA, Non-uniform LWE, and modeling SHAKE as a random oracle, an implicit rejecting KEM with splittable public keys (and with appropriate parameters as discussed in Sec. 7) is A-CFR-CCA-secure.*

$\text{Exp}_{\text{KEM}}^{\text{A-CFR-CCA}^b}(\mathcal{A})$	$\text{Exp}_{\text{KEM}}^{\text{A-CFR-CCA}^{(b, \mathcal{D})}}(\mathcal{A})$
1: $(pk, sk) \leftarrow \$ \text{KGen}(1^\kappa)$	1: $(pk, sk) \leftarrow \$ \text{KGen}(1^\kappa)$
2: $(\mathbf{A}, \mathbf{b}) \leftarrow pk$	2: $(\mathbf{A}, \mathbf{b}) \leftarrow pk$
3: $\mathbf{A}^* \leftarrow \$ \mathbb{A}$	3: $\mathbf{A}^* \leftarrow \$ \mathcal{D}$
4: $C \leftarrow \$ \mathcal{A}^{\text{D}(\cdot, \cdot)}(\mathbf{b}, \mathbf{A}, \mathbf{A}^*)$	4: $C \leftarrow \$ \mathcal{A}^{\text{D}(\cdot, \cdot)}(\mathbf{b}, \mathbf{A}, \mathbf{A}^*)$
5: $K \leftarrow \text{Decap}(\mathbf{b}, \mathbf{A}, sk, C)$	5: $K \leftarrow \text{Decap}(\mathbf{b}, \mathbf{A}, sk, C)$
6: $K^* \leftarrow \text{Decap}(\mathbf{b}, \mathbf{A}^*, sk, C)$	6: $K^* \leftarrow \text{Decap}(\mathbf{b}, \mathbf{A}^*, sk, C)$
7: return $K = K^* \neq \perp$	7: return $K = K^* \neq \perp$

Fig. 7. KEM \mathbf{A} -Part Collision Freeness (Left) and w.r.t. Password Dictionary (Right), where \mathbb{A} and \mathcal{D} denote the lattice matrix space and password dictionary space respectively, and $\text{KEM.Decap}(\mathbf{b}, \mathbf{A}, sk, C)$ denotes invoking Decap explicitly with a chosen matrix \mathbf{A} not necessarily matching the one embedded in \mathbf{b} .

Remark on the semantics of Collision Freeness (and its \mathbf{A} -part variant). As shown by Cremers et al. [25], SCRf-CCA for a KEM means that together, the output key and the ciphertext bind the public key; i.e., changing the public key changes the other two components by a sufficient distance, so that they both become statistically independent across different public keys. We argue that the \mathbf{A} -part of the public key has enough entropy to change the output key in both KEM.Encap and KEM.Decap routines, so that the resulting keys are not systematically relatable to each other. We note that an adversary can clearly forward arbitrary (non-related) ciphertexts in a real protocol execution, in which case no collisions can naturally occur for a KEM satisfying SCRf-CCA [25] (Fig. 2). However, an adversary’s goal in the A-CFR-CCA game (Fig. 7) is to enforce a collision on ciphertexts so that they decapsulate to matching (i.e., identical) keys although encapsulated with different \mathbf{A} -parts.

Proof. In order to show $\text{A-CFR-CCA}^{(b, \mathcal{D})}$ security, we use the common game based approach, starting with the original security game \mathbf{G}_0 and transitioning to a new experiment \mathbf{G}_1 while bounding the adversary’s advantage. This latter game will be shown hard based on the hardness of SCRf-CCA.

Game \mathbf{G}_0 : This is the original game. Here, the challenger first generates pk, sk honestly. \mathbf{A} is obtained from splitting \mathbf{b} , and is thus uniformly chosen from the \mathbf{A} -space of the chosen KEM. Then, \mathbf{A}^* is sampled from distribution \mathcal{D} , which is a well-spread distribution depending on the entropy of the password π with support in the \mathbf{A} -space of the KEM. While $pk = (\mathbf{A}, \mathbf{b})$, only the \mathbf{b} term is fixed. The adversary can thus freely choose whether to encapsulate using \mathbf{A} or \mathbf{A}^* , and they use one of those with \mathbf{b} to construct C . The challenger receives C from the adversary and invokes KEM.Decap twice using both \mathbf{A} -parts and obtains K and K^* respectively. A standard encapsulation routine of a PKE-based LWE KEM

(e.g., FrodoKEM as shown in Fig. 5) is then:

```

(k, r) = G2(G1(pk)||m)
C ← PKE.Enc(m, pk; r)
K ← F(C||k)
return (C, K)

```

where G_1 , G_2 and F are instantiated with *SHAKE*. The adversary's view in \mathbf{G}_0 is $(C, \mathbf{b}, \mathbf{A}, \mathbf{A}^*)$ and the challenger checks: K and K^* are equal and not rejected. For KEMs with implicit rejection the challenger only checks if $K = K^*$.

Game \mathbf{G}_1 : In this step, we swap the \mathbf{A} -part in pk to be sampled uniformly from the \mathbf{A} -space of the KEM. The rest of the experiment is the same.

Lemma 4 ($\mathbf{G}_0 \stackrel{\text{comp}}{\approx} \mathbf{G}_1$). *If *SHAKE* is modeled as a random oracle and Non-uniform *LWE* is $(1 - \text{negl}(1^\kappa))$ -hard, then*

$$\Pr_{\mathbf{G}_0}[\mathcal{A} \text{ wins}] \leq \Pr_{\mathbf{G}_1}[\mathcal{A} \text{ wins}] + \text{negl}(1^\kappa)$$

Proof (Proof of Lemma 4). The difference in the games is the distribution of \mathbf{A} . In the first game, \mathcal{A} 's view of \mathbf{b} does not allow him to identify the \mathbf{A} -part embedded in \mathbf{b} by the hardness of Non-uniform *LWE* (similar to the proof of Lemma 2). Thus, with the output of the function G_1 being modeled as a random oracle, inputting a pseudo-random pk as per Definition 12 will yield a statistically independent value, with probability bound by the birthday bound for a RO modeled hash function or XOF. For a PAKE, we claim that conditioned on the entropy of π being well-spread and uniform distributed in \mathcal{D} , the output of G_1 is also bound by the password space, which resembles the basic attacks on the password dictionary in the context of a PAKE. Hence we derive

$$\Pr_{\mathbf{G}_0}[\mathcal{A} \text{ wins}] \leq \Pr_{\mathbf{G}_1}[\mathcal{A} \text{ wins}] + |\mathcal{D}|^2 \cdot 2^{-|\rho|}$$

where ρ is the seed for sampling the \mathbf{A} -part and \mathcal{D} is the password dictionary.

Completing the proof of Theorem 4. Now, suppose there is an adversary \mathcal{A} that wins with non-negligible probability in the \mathbf{A} -CFR-CCA^(\mathbf{b}, \mathcal{D}) game. We show an adversary \mathcal{B} that succeeds against the SCFR-CCA security experiment with $(1 - \text{negl}(1^\kappa))$ -close to the same probability. The reduction is trivial: \mathcal{B} runs SCFR-CCA to line 3, then hands $(\mathbf{A}_0, \mathbf{A}_1, \mathbf{b}_0)$ to \mathcal{A} . When \mathcal{A} outputs C , \mathcal{B} outputs C in line 3 of the SCFR-CCA game. From the hardness of standard *LWE*, their probabilities of success are $\text{negl}(1^\kappa)$ close. From the hardness of SCFR-CCA via Grubbs et al. [28] and Cremers et al. [25], we get Theorem 4. \square

6 Security Analysis of NICE-PAKE

Definition 15 (The NICE-PAKE Protocol). *Let \mathcal{D} be a dictionary of possible passwords $\pi_i \in \mathcal{D}$, \mathcal{H} be a hash function or an XOF, and **KEM** be a KEM with key space \mathcal{K} . The protocol NICE-PAKE (denoted by Π) is depicted Fig. 1.*

6.1 Security Model

The security analysis follows the extended BPR [12] model for authenticated key exchange (AKE) [12]. It models an adversary in full control of the network, but without knowledge of the password, nor the secrets of honest users (e.g., the secret key in a KEM). The goal of an adversary is to break a protocol instance by guessing a test bit (i.e., the final session key) via a so-called test query. We define the original attack on the protocol as a security experiment (game) denoted by \mathbf{G}_0 . This game resembles (P)AKE security of BPR in the Real-or-Random model (i.e., multiple test queries are allowed in the whole game). As we believe most readers are familiar with this model, we opted to include the full model description in App. A. We also refer non-experts on PAKE security to App. B for a textual proof sketch and remarks on CPA vs. CCA security.

6.2 Correctness

The protocol Π is correct if an honest party outputs a shared secret that matches the shared secret of another honest party in a partnered session. This implies that the correctness of Π depends on correctness of the chosen KEM and that the password matches between the two honest parties.

Definition 16 (Correctness). *The protocol Π is δ -correct if for every pair of honest and partnered users $(\mathcal{U}_i, \mathcal{U}_j)$ with matching passwords $(\pi_{\mathcal{U}_i}, \pi_{\mathcal{U}_j}) \in \mathcal{D}$ in an honest protocol execution denoted by $\Pi(\mathcal{U}_i, \mathcal{U}_j)$ we have:*

$$\Pr[K' \neq K | (K, K') \leftarrow \Pi(\mathcal{U}_i, \mathcal{U}_j) : \pi_{\mathcal{U}_i} = \pi_{\mathcal{U}_j}] \leq \delta$$

where δ define the correctness of the chosen KEM. It follows that the correctness of Π depends on the correctness of KEM and is perfectly correct if $\delta = 0$.

6.3 Formal Security Analysis

Theorem 5. *Let KEM be IND-CCA, SPLIT-PKU, ANO-CCA, A-SEC-CCA and A-CFR-CCA with the respective advantages $\mathbf{Adv}_{KEM}^{PROP}$, and assuming that KEM is δ -correct, and for a well-spread and uniform distributed password dictionary \mathcal{D} and any chosen $\pi \in \mathcal{D}$, and assuming that \mathcal{H} is a hash function modeled as a random oracle. For the given protocol Π and an adversary \mathcal{A} , we define \mathcal{A} 's advantage in the BPR model for AKE security with respect to Π as $\mathbf{Adv}(\mathcal{A}, \Pi)$, and with respect to Q as the number of queries made by an adversary for each operation. We have:*

$$\begin{aligned} \mathbf{Adv}_{\Pi}^A(Q) &\leq 2 \cdot (q_s^2 \cdot 2^{-|\rho|}) + |\mathcal{D}|^2 \cdot 2^{-|\rho|} + 4q_s^2 \cdot \mathbf{Adv}_{KEM}^{IND-CCA}(\mathcal{B}) \\ &\quad + 2q_s^2 \cdot |\mathcal{D}|^{-1} \cdot (\mathbf{Adv}_{KEM}^{A-SEC-CC}(\mathcal{B}) + \mathbf{Adv}_{KEM}^{SPLIT-PKU}(\mathcal{B})) \\ &\quad + 2q_s^2 \cdot (\mathbf{Adv}_{KEM}^{ANO-CCA}(\mathcal{B}) + \mathbf{Adv}_{KEM}^{A-CFR-CCA}(\mathcal{B})) \end{aligned}$$

Where $|\mathcal{D}|$ is the password dictionary size, $|\rho|$ is the KEM **A**-part seed space size, and $Q := (q_e, q_s)$ denoting *Execute* and *Send* queries respectively.

Proof: We provide a security proof via the common game-based approach. We gradually modify the original security game \mathbf{G}_0 through a series of experiments $\mathbf{G}_1, \mathbf{G}_2, \dots$ showing that the success probability of an adversary cannot be significantly larger than $\frac{1}{2}$. The final game will be independent of the challenge bit b , giving the theorem. Specifically, we will aim to randomize the values transmitted on the network, which are visible for an adversary \mathcal{A} . We say the protocol Π is secure if the probability of \mathcal{A} winning is bounded to a negligible quantity. We denote the probability of \mathcal{A} winning in game \mathbf{G}_i as $\Pr[\mathbf{G}_i]$.

Game \mathbf{G}_0 : The original attack on the protocol.

Eliminating Collisions (\mathbf{G}_1 - \mathbf{G}_3)

Game \mathbf{G}_1 (Hash Collisions): Change Game \mathbf{G}_0 declaring the adversary to lose if there are distinct passwords $\pi \neq \pi' \in \mathcal{D}$ such that their hash values collide, $\mathcal{H}(\pi) = \mathcal{H}(\pi')$. The probability of this happening is bounded by the birthday bound as

$$\Pr[\mathbf{G}_0] \leq \Pr[\mathbf{G}_1] + |\mathcal{D}|^2 \cdot 2^{-|\rho|}.$$

This, in particular, means that we can assume that for each fixed value ρ , the function values $\rho \oplus \mathcal{H}(\pi)$, when varying over π , are all distinct.

Game \mathbf{G}_2 (Unique pk): Modify Game \mathbf{G}_1 by declaring the adversary to lose if there are two honest Alice-sessions (possibly with different passwords) which create the same initial message (z, \mathbf{b}) . Note that each honest Alice-session picks a fresh random value ρ in each execution. Hence, the probability that this random value (shifted by $H(\pi)$ for Alice's password π in this session) matches the value $\rho' \oplus H(\pi')$ in another honest Alice-session is at most $q_s^2 \cdot 2^{-|\rho|}$. We derive

$$\Pr[\mathbf{G}_1] \leq \Pr[\mathbf{G}_2] + q_s^2 \cdot 2^{-|\rho|}.$$

Game \mathbf{G}_3 (Unique C_b): Modify Game \mathbf{G}_2 by declaring the adversary to lose if there are two honest Bob-sessions (possibly with different passwords) which create the same response message C_b . Note that the ciphertext C_b encapsulates a random message m from $\{0, 1\}^{|\rho|}$. This implies that the ciphertext space is of size at least $2^{|\rho|}$ and the ciphertext is chosen uniformly among such encryptions of messages m . Hence, the probability that a ciphertext C_b matches any of the $i \leq q_s$ previously chosen ciphertexts is at most $i \cdot 2^{-|\rho|}$ and thus

$$\Pr[\mathbf{G}_2] \leq \Pr[\mathbf{G}_3] + q_s^2 \cdot 2^{-|\rho|}.$$

Simulating the KEM (\mathbf{G}_4 - \mathbf{G}_5) Modify Game \mathbf{G}_3 by having any honest Bob-session receiving a message (z, \mathbf{b}) created by an honest Alice-session use $C_{b,\text{sim}}$ instead of ciphertext C_b , where $(C_{b,\text{sim}}, \bar{K}'_{\text{sim}}) \leftarrow \text{KEM.Encap}(pk_{\text{sim}})$ for a freshly generated key pair $(sk_{\text{sim}}, pk_{\text{sim}}) \leftarrow \text{KEM.KGen}()$, for an independently

sampled \mathbf{A}_{sim} . When the honest Alice-session which has sent (z, \mathbf{b}) (and which is unique by Game \mathbf{G}_2) receives this ciphertext $C_{b,\text{sim}}$, it uses the original key \overline{K}' that Bob had encapsulated instead of decapsulating $C_{b,\text{sim}}$; the honest Bob-session also uses the key \overline{K}' for the further steps. The indistinguishability of the hops from \mathbf{G}_3 to \mathbf{G}_5 is shown in two substeps.

Game \mathbf{G}_4 (Simulate KEM Key): First, we replace C_b by $C_{b,\text{sim}}$ for honest Bob-sessions (with an honest Alice-session partner) *but use the key $\overline{K}'_{\text{sim}}$ generated by Bob in the further steps.* We claim that this modification is indistinguishable according to ANO-CCA. First, note that there can be at most q_s honest Bob-sessions with an honest Alice-session, and we can apply the above modification step-by-step. Second, ANO-CCA ensures that adversaries cannot gain information by forwarding ciphertexts to other connected sessions. Via a hybrid argument, we lose a factor q_s in the argument but can, from now on, focus on a single pair of sessions in which we replace Bob's values. It then follows via a straightforward reduction \mathcal{B} to ANO-CCA, simulating the entire Game \mathbf{G}_4 (also picking the passwords of parties), and injecting the challenge key pairs (sk_0, pk_0) , (sk_1, pk_1) from the anonymity game into the Alice-session, using pk_0 as pk_a resp. using pk_1 as pk_{sim} in the Bob-session. The latter loses a factor q_s in the security bound to guess the correct Alice-session. If this happens, we immediately know the honest Bob-session communicating with this unique Alice-session. The anonymity game also gives us a challenge ciphertext C^* and key K^* , created either under pk_0 or under pk_1 . We use C^* as $C_{b,\text{sim}}$, and both parties use the key K^* as the session key. The reduction against anonymity outputs 1 if and only if adversary \mathcal{A} wins the simulated game. Note that if C^*, K^* in the anonymity game are created under the public key pk_0 , then the simulation corresponds perfectly to game \mathbf{G}_4 . If (C^*, K^*) is created under pk_1 then the simulation corresponds precisely to our intermediate game. It follows

$$\Pr[\mathbf{G}_3] \leq \Pr[\mathbf{G}_4] + 2q_s^2 \cdot \mathbf{Adv}_{\text{KEM}}^{\text{ANO-CCA}}(\mathcal{B}),$$

for some adversary \mathcal{B} , where the factor 2 stems from switching from a left-or-right game ANO-CCA to a comparison between games \mathbf{G}_3 and \mathbf{G}_4 .

Game \mathbf{G}_5 (Use Original KEM Key): The next step now is to switch back to the original keys \overline{K}' instead of $\overline{K}'_{\text{sim}}$ in such honest Alice-Bob-interactions. This follows now from the IND-CCA security of the KEM, saying that one cannot distinguish which of the two keys is encapsulated in $C_{b,\text{sim}}$. It follows as in the anonymity case that

$$\Pr[\mathbf{G}_4] \leq \Pr[\mathbf{G}_5] + 2q_s^2 \cdot \mathbf{Adv}_{\text{KEM}}^{\text{IND-CCA}}(\mathcal{B})$$

for some adversary \mathcal{B} . Now, in honest Alice-Bob interactions, the key \overline{K} is entirely independent of the communication transcript. Further, the exchanged messages between Alice and Bob in such sessions are distributed independently of the password π of both parties: Alice sends a random string $\rho \oplus H(\pi)$ and Bob

responds with a ciphertext computed under a fresh public key. It remains to look into sessions where (a) Alice is malicious and interacts with an honest Bob-session, and (b) Bob is malicious but communicates with an honest Alice-session.

Randomizing the Simulation (\mathbf{G}_6 - \mathbf{G}_{10}) In the following, we consider two *bad* events that the simulation cannot prevent. The first event corresponds to the case where an adversary \mathcal{A} forwards a pair (z, \mathbf{b}) that contains a correct password guess on z . Intuitively, an adversary who guessed the password correctly would be able to distinguish between messages created by honest instances and ones randomized by the simulation. However, we make use of the A-SEC-CCA property to show that \mathcal{A} 's view on a pk used by an honest instance for encapsulating a key is indistinguishable from random. Therefore, we continue the simulation of the protocol and derive that \mathcal{A} 's advantage is bound to a factor of the number of queried sessions and random guessing over the password dictionary size $|\mathcal{D}|$. The second event corresponds to a `Corrupt` query placed by an adversary on connected honest instances.

Game \mathbf{G}_6 (Randomize C_b): Alter Game \mathbf{G}_5 by letting an honest Bob-session receiving (z, \mathbf{b}) which has not been created by an honest Alice-session use $C_{b,\text{sim}}$ instead of ciphertext C_b , where $(C_{b,\text{sim}}, \overline{K}'_{\text{sim}}) \leftarrow \text{KEM.Encap}_{\mathbf{A}_{\text{sim}}}(pk)$ for an independently sampled \mathbf{A}_{sim} . The honest Bob-session also uses the key \overline{K}' for further steps. We claim that this modification is indistinguishable according to A-SEC-CCA. It then follows via a straightforward reduction \mathcal{B} to A-SEC-CCA simulating the entire game \mathbf{G}_{6a} and injecting the challenge \mathbf{A} -parts $(\mathbf{A}_0, \mathbf{A}_1)$ for the received \mathbf{b} from the A-Part-Secrecy game into the Alice-session, using \mathbf{A}_0 as \mathbf{A}_a resp. using \mathbf{A}_1 as \mathbf{A}_{sim} in the Bob-session. Here, we lose the factor of ignoring trivial password guessing, as the Alice-session has an sk that actually decrypts under an encapsulation invoked by the \mathbf{A} -part obtained from the correct password. By doing so, an Alice-session cannot decapsulate the randomized ciphertext $C_{b,\text{sim}}$, which in turn does not reveal the randomized key $\overline{K}'_{\text{sim}}$, as `KEM.Decap` rejects implicitly by outputting a random key. On the other hand, the A-SEC-CCA game does not reveal the \mathbf{A} -part used for the encapsulation. The reduction against A-Part-Secrecy outputs 1 if and only if adversary \mathcal{A} wins the simulated game. Hence, if the pair (C^*, K^*) in the A-Part-Secrecy game is created under the public key with \mathbf{A}_0 , then the simulation corresponds perfectly to game \mathbf{G}_5 . If (C^*, K^*) is created under the public key with \mathbf{A}_1 , then the simulation corresponds to this game. It follows

$$\Pr[\mathbf{G}_5] \leq \Pr[\mathbf{G}_6] + 2q_s^2 \cdot \text{Adv}_{\text{KEM}}^{\text{A-SEC-CCA}}(\mathcal{B}) \cdot |\mathcal{D}|^{-1}$$

for some adversary \mathcal{B} , where we lose a factor 2 for switching from a left-or-right game A-SEC-CCA to the comparison between games \mathbf{G}_5 and \mathbf{G}_6 . Note that a `Reveal` query on an honest Bob-session does not carry any significance in this scenario, since \mathcal{A} cannot decapsulate $C_{b,\text{sim}}$, and did not guess the correct password. Thus they cannot compare Bob's key \overline{K}' to a meaningful value.

Game \mathbf{G}_7 (Corrupt Bob): Declare the adversary to lose, if an honest Bob-session is marked *unfresh* after receiving (z, \mathbf{b}) which has not been created by a connected honest Alice-session. This game simulates the case where an adversary \mathcal{A} obtains the password through *corrupting* an honest Bob instance in a protocol session. We define an *unbounded adversary*, whose interaction in this step does not affect the previous advantage. Thus, this game change is conceptual, and does not affect the security bound of the protocol. First, we recall that no test query can be placed by \mathcal{A} on *unfresh* sessions as per the security model. Second, although an honest Bob would have otherwise unknowingly generated a ciphertext C_b based on (z, \mathbf{b}) and their honest password, obtaining this password via corruption will not aid \mathcal{A} in decapsulating a ciphertext $C_{b,\text{sim}}$ with a non-matching $(sk, (\mathbf{A}_{\text{sim}}, \mathbf{b}))$ pair. That is since the ciphertext is simulated using an independently sampled \mathbf{A}_{sim} as per \mathbf{G}_6 , and we had hence already accounted for the adversarial advantage stemming from A-SEC-CCA. This is justified through observing that \mathcal{A} 's advantage in \mathbf{G}_6 is bounded by the reduction to A-SEC-CCA and a multiplicative factor of guessing the attacked session over the password dictionary size. Thus, restricting \mathcal{A} 's interaction to a specific and known session (marked unfresh), eliminates the factors in the previous advantage term. Consequently, the adversary interacts with a single instance of the A-SEC-CCA experiment. It follows

$$\Pr[\mathbf{G}_6] = \Pr[\mathbf{G}_7]$$

Game \mathbf{G}_8 (Randomize z): Change Game \mathbf{G}_6 by letting an honest Alice-session initiating the protocol use a randomly generated z_{sim} instead of an honestly masked seed ρ . Since we already accounted for collisions on z in game \mathbf{G}_2 , we claim that this modification is indistinguishable according to SPLIT-PKU that an adversary \mathcal{A} cannot distinguish between \mathbf{A} -parts resulting from splitting pk_a and ones embedded within the \mathbf{b} -part of the same public key. It then follows via a simple reduction \mathcal{B} to SPLIT-PKU simulating game \mathbf{G}_8 and injecting the challenge \mathbf{A} -part \mathbf{A}_1 for the received \mathbf{b} -part from the SPLIT-PKU security game into the Bob-session. Here as well, we lose a factor for ignoring trivial password guessing over the password dictionary size. It follows

$$\Pr[\mathbf{G}_7] \leq \Pr[\mathbf{G}_8] + 2q_s^2 \cdot \mathbf{Adv}_{\text{KEM}}^{\text{SPLIT-PKU}}(\mathcal{B}) \cdot |\mathcal{D}|^{-1}$$

for some adversary \mathcal{B} , where we lose a factor 2 for switching from a left-or-right game SPLIT-PKU to the comparison between games \mathbf{G}_6 and \mathbf{G}_8 .

Game \mathbf{G}_9 (Randomize pk_a): Change Game \mathbf{G}_8 by letting an honest Alice-session initiating the protocol use a randomly generated key pair $(sk_{\text{sim}}, pk_{\text{sim}})$ instead of an honestly generated one, where $(\mathbf{b}_{\text{sim}}, \rho_{\text{sim}}) \leftarrow pk_{\text{sim}}$ and an independently generated z_{sim} from \mathbf{G}_8 . Upon receiving a response from a Bob-session, the Alice-session uses the previously randomized key pair to invoke KEM.Decap. Since \mathcal{A} encapsulated their own key $K_{\mathcal{A}}$ into $C_{\mathcal{A}}$, it will not match the decapsulated key \bar{K} for an Alice-Session. This change is indistinguishable for \mathcal{A} according to the

collision freeness of KEM (except for the case of a `Reveal` query addressed in a following step). It then follows via a straightforward reduction \mathcal{B} to A-CFR-CCA simulating game \mathbf{G}_9 and injecting the challenge \mathbf{A} -parts $(\mathbf{A}_0, \mathbf{A}_1)$ for the received public keys from the A-CFR-CCA game into the session using \mathbf{A}_0 as \mathbf{A}_a resp. using \mathbf{A}_1 as \mathbf{A}_{sim} . Note that we don't lose a factor for trivial password guessing, as we already accounted for it in game \mathbf{G}_8 . However, we still lose a factor q_s for guessing the correct session. The reduction against A-CFR-CCA outputs 1 if and only if \mathcal{A} wins the simulated game. Hence, if the pair (C^*, K^*) in the A-CFR-CCA game is created under the public key with \mathbf{A}_0 , then the simulation corresponds to game \mathbf{G}_8 . If (C^*, K^*) is created under the public key with \mathbf{A}_1 , then the simulation corresponds to this game. We derive

$$\Pr[\mathbf{G}_8] \leq \Pr[\mathbf{G}_9] + 2q_s^2 \cdot \mathbf{Adv}_{KEM}^{\text{A-CFR-CCA}}(\mathcal{B})$$

Game \mathbf{G}_{10} (Corrupt Alice): Declare the adversary to lose, if an honest Alice-session is marked *unfresh* after sending (z, \mathbf{b}) to a Bob-Session. This game simulates the case where an adversary \mathcal{A} obtains the password through *corrupting* an honest Alice instances in a protocol session. As in \mathbf{G}_7 , we define an *unbounded adversary*, whose interaction in this step does not affect the previous advantage. Thus, this game change is conceptual, and does not affect the security bound of the protocol. Similarly, we first recall that no test query can be placed by \mathcal{A} on *unfresh* sessions as per the security model. Second, although an honest Alice would have otherwise sent a z that indeed reveals the correct seed for sampling the honest \mathbf{A} -part in her pk , and thus allowing \mathcal{A} to generate a valid ciphertext that normally decapsualtes the same key for both users, obtaining the correct password via a corruption will not aid \mathcal{A} in obtaining an identical key K^* with a non-matching (pk_a) . That is since Alice's pk and her masked seed z are simulated using independently sampled values as per \mathbf{G}_8 and \mathbf{G}_9 , we had already ruled out trivial password guesses on z , \mathcal{A} commits to a single (non-adaptive) response $C_{\mathcal{A}}$, and we have already accounted for the adversarial advantage stemming from both SPLIT-PKU and A-CFR-CCA respectively. This is justified through observing that \mathcal{A} 's advantage in \mathbf{G}_8 and \mathbf{G}_9 is bounded by the respective reductions and a multiplicative factor for guessing the attacked session (and the password dictionary size in \mathbf{G}_8). Thus restricting \mathcal{A} 's interaction to a specific and known session (marked unfresh) eliminates the factors in the previous advantage terms. Consequently, the adversary interacts with a single instance of SPLIT-PKU and A-CFR-CCA. It follows

$$\Pr[\mathbf{G}_9] = \Pr[\mathbf{G}_{10}]$$

Randomizing Session Keys (\mathbf{G}_{11}) We had shown that the remaining sessions should either be unfresh and cannot be tested, or if the sessions are fresh, then they must have already been replaced by some independent data on the network. It now remains to randomize the final sessions keys such that they yields random (unknown) values indistinguishable from real keys.

Game \mathbf{G}_{11} (Randomize K): Finally, we replace the final session key with a key chosen independently at random from the key space \mathcal{K} . That is, on all connected instances and for all **Execute** and **Send** queries, rendering the keys independent of all previous messages and the password. We claim that this change is indistinguishable from games \mathbf{G}_1 through \mathbf{G}_{10} for \mathcal{A} . In other words, a **Test** query may be placed by \mathcal{A} at any point of the simulation (except on *unfresh* instances), and the tested key will always be indistinguishable based on the key security of KEM. Hence, the adversary’s advantage is negligible and bounded by IND-CCA security of the KEM and the number of placed queries. It follows

$$\Pr[\mathbf{G}_{10}] \leq \Pr[\mathbf{G}_{11}] + 2q_s^2 \cdot \mathbf{Adv}_{KEM}^{\text{IND-CCA}}(\mathcal{B})$$

Where \mathcal{B} is some adversary playing the indistinguishability game of KEM. Note that **Reveal** queries on Alice do not need to be explicitly handled by the simulation. Since the used KEM is implicitly rejecting, decapsulating on honest values will always yield a valid key $K^* \in \mathcal{K}$ that is with negligible probability bound to the key space size indistinguishable from a real key (covered by IND-CCA). Having randomized all values involved in encapsulating and decapsulating a session key over the network, and since a **Test** query is rendered unavailable for an *unfresh* instance, \mathcal{A} ’s view on honest keys as well as their their own key is dependent only on their password guess being trivially correct (handled previously); and is indistinguishable from the randomization within the simulation.

Putting It All Together: In the final game, all keys returned through a **Test** query are random values independent from the protocol simulation. Thus, an adversary cannot distinguish real keys from random ones. Their advantage is therefore upper-bounded by $\frac{1}{2}$ and the numbers of protocol executions, and hence their overall advantage is bounded by $\frac{1}{2}$ plus any losses collected throughout the games, which gives Theorem 5. \square

7 Remarks on Instantiations using Frodo and Kyber

We note that Appendices D, E and F contain basics on some of the arguments and parameter choices we present in the following.

Remarks on the concrete value of ϵ_{whLWE} in Theorem 3. We first mention why alternative reductions in the literature do not work here. There are two incomparable reductions for Extended LWE applicable to our use-case: one from O’Neill et al. in [41] that yields a $(1/q)$ multiplicative advantage loss per application, and one from Brakerski et al. [21] that yields a negligible advantage loss, up to a mild change in parameters. In the first case [41], we cannot afford to iterate a reduction with $(1/q)$ multiplicative loss over a column of v in FrodoKEM (that is an 8×8 matrix) for a total multiplicative loss of at least $(1/8q^8)$, as this would give concrete loss of bits of security at least $3 \cdot 16 \cdot 8 = 384$. In the second case [21], we do not see (currently) how to prove the analogue of their Claim

4.6: that is, constructing an unimodular \mathbf{U} (with small largest singular value) so that removing \mathbf{U} 's left column yields that all the remaining columns are orthogonal to an arbitrary element of \mathbb{Z}_q^k . If such \mathbf{U} can be demonstrated, this would provide an alternative, concretely-effective proof of security. We also do not see how to easily use the “noise lossiness” techniques of Brakerski and Döttling [19], as we require decisional hardness, but their proof only guarantees search hardness in the case of a (non-prime) power-of-2 integer modulus q as in FrodoKEM. However, using the techniques of Cheon et al. [24,36] and Liu et al. [38], we can calculate a small loss in concrete bit-security for a FrodoKEM-instantiation as follows: Following Theorem 2, the noise and error distributions of FrodoKEM are identical, so we have $\sigma'_1 = \sigma'_2$. Then, $\sigma_1 = \sigma'_1/\sqrt{2} = \sigma'_2/\sqrt{2}$. Here, $k = 8$, so we have concrete security from Plain LWE, but with a $\sqrt{8} \cdot \sqrt{2} = 4$ multiplicative loss in variance. This yields that an instantiation of A-SEC-CCA security in our protocol from FrodoKEM-640 is as hard as if the variance were reduced from 2.8 to 0.7; or from FrodoKEM-1344 if the variance were reduced from 1.4 to 0.35. In the first case (Frodo-640), this is close to (but actually, slightly higher entropy) having a uniform $\{0, 1\}$ -valued marginal distribution of the secret key and error terms in the view of the adversary after leakage. In the second case, it's slightly lower entropy than uniform-binary secrets and errors, but the additional dimension of FrodoKEM-1344 may make up for this. In general, it would be better to re-parameterize a FrodoKEM-style Plain LWE KEM (call it “GimliKEM”) with either slightly higher noise variance or slightly higher dimension (or both). One can consider such trade-offs in practice by using Albrecht et al.'s regularly maintained LWE Estimator tool [3] to find the optimal choices.

A complete break (provable insecurity) of a Kyber-instantiation. We describe three different (very efficient) attacks by a Malicious Alice against a Kyber-instantiation of the protocol, which result in either a major leakage on or **complete recovery** of Bob's ciphertext randomness \mathbf{r} in any given Kyber session. In what follows, recall the concrete Kyber set-up. Let $R = \mathbb{Z}[X]/(X^{256} + 1)$ and $q = 3329$. Honest public keys are $(\mathbf{A}, \mathbf{b}) \in R_q^{k \times k} \times R_q^k$ for $k \in \{2, 3, 4\}$, where \mathbf{A} is unrolled via SHAKE from a random seed $\rho_{\mathbf{A}}$ and so can be treated as having the uniform distribution over the coordinates of the Chinese Remainder Theorem (CRT) embedding, respectively the coefficient embedding, of each polynomial $\mathbf{A}_{i,j \in [k]} \in R_q$, and where $\mathbf{b} := \mathbf{A}\mathbf{s} + \mathbf{e}$ for \mathbf{s}, \mathbf{e} drawn coefficient-wise from a binomial distribution η with parameter 2 (or parameter 3 for the secret \mathbf{s} for Kyber-512) and thus having support in $\{-2, -1, 0, +1, +2\}$. Since $n = 256, q = 3329$, we have $n|(q - 1)$ and so the ring R_q factors into $n/2$ distinct quadratic factors, meaning there is a CRT coordinate system defined that is isomorphic to $(\mathbb{Z}_q[X]/(X^2 + 1))^{n/2}$. Recall that $X^2 + 1$ is the 4th cyclotomic polynomial and that elements of the quadratic ring $\mathbb{Z}_q[X]/(X^2 + 1)$ are represented by two integers modulo q with slot-wise addition and simple, grade-school convolution-multiplication. One can map efficiently from the coefficient representation of polynomials in R_q to their CRT coordinate system and back using the Number Theoretic Transform (NTT) and inverse-NTT.

A first attack: When honest \mathbf{A} contains “correlated” zero divisors. Recall that the protocol’s security proof contemplates hints to leak on the secret \mathbf{r} of the challenge ciphertext’s MLWE secret in the form of v . To demonstrate that such hints “only reveal unstructured entropy,” we generically need that \mathbf{A} is invertible, so that we can consider the expression $\mathbf{r} + \mathbf{A}^{-1}\mathbf{e}'$ when attempting to swap u to uniform in order to show anonymity. While FrodoKEM’s \mathbf{A} -part is invertible except with negligible probability, this is not the case for Kyber. In particular, if there exists an CRT-coordinate index $z \in [n/2]$ and fixed column index $j \in [k]$ so that for all R_q polynomials $\mathbf{A}_{i,j}$ (i.e. for all row indices $i \in [k]$) we have that

$$\begin{aligned} \text{NTT}(\mathbf{A}_{i,j})[z] &= 0 \in \mathbb{Z}_q[X]/(X^2 + 1) \\ \text{that is, } \text{CRT}(\mathbf{A}_{i,j})[z] &= (0, 0) \in (\mathbb{Z}_q)^2, \end{aligned}$$

then \mathbf{A} is not invertible in $R_q^{k \times k}$. In this event, the leakage produced by the hint v is “unexpectedly” algebraically-structured and will convey information about \mathbf{r} and \mathbf{e}'' that is constrained to non-maximal ideals of the ring R_q . This event occurs with probability at least $\frac{k \cdot n/2}{q^{2 \cdot k}}$, which is noticeable since k is constant. For example, for Kyber-1024 this event occurs with probability at least $\frac{4 \cdot 256/2}{3329^{2 \cdot 4}} \approx 2^{-84.6} \gg 2^{-256}$.

“Patching” the first attack: Check \mathbf{A} -invertibility. A simple patch for this issue is for Honest Bob to attempt to invert any Kyber \mathbf{A} before using it to construct his ciphertext in the protocol. While relatively efficient (if costing an undesirable practical slowdown), this especially has the downside of forcing the protocol to be “aware” of the particular lattice KEM being used.

A second attack: When Alice chooses \mathbf{b} (resp. each \mathbf{b}_i) to be a unit in the ring R_q . A larger concern is if malicious Alice arbitrarily chooses \mathbf{b} to be a “degenerate” element of R_q , such as a unit – for example, $1 \in R_q$. In this case, \mathbf{b} vanishes completely in the expression $v = \mathbf{b}\mathbf{r} + \mathbf{e}''$, resulting in the hint $v = \sum_{i \in [k]} \mathbf{r}_i + \mathbf{e}''_i$, which dangerously exposes much of the entropy in \mathbf{r} and \mathbf{e}'' .

“Patching” the second attack: Check if \mathbf{b} is a unit (per slot). Again, there is a simple patch. Before Honest Bob constructs a ciphertext in the protocol, he can test if elements of \mathbf{b} are units in R_q with a straightforward calculation. However, concern should continue growing: How many bad cases could there be? When have you caught them all?

A third attack: When Alice chooses \mathbf{b} to be a “gadget vector” with moderately-sized radix. Finally, consider when Malicious Alice chooses \mathbf{b} to be a non-unit scalar. This is simplest to see in the Ring-LWE (i.e., rank 1) “case,” when \mathbf{b} , \mathbf{r} , and \mathbf{e}'' are all simply polynomials in R_q . Since the support of each coefficient of \mathbf{b} and \mathbf{e}'' are very small – e.g. in $\{-2, -1, 0, +1, +2\}$ – and the modulus $q = 3329$ is relatively large by comparison, consider if \mathbf{b} is chosen as (say) $64 \in R_q$. Then, the hint $v = 64\mathbf{r} + \mathbf{e}''$, and the coefficients of \mathbf{r} and \mathbf{e}'' can be read, directly, from

the bit representation of the coefficients of v . In the case of Kyber-1024, Alice can choose \mathbf{b} as a “gadget vector” such as $(4, 4^2, 4^3, 4^4) = (4, 16, 64, 256) \in R_{3329}^4$. Then, even with the vector-wise addition over rank 4, each of the coordinates of each of the \mathbf{r}_i can still be read off directly from the bit representation of v , independent of \mathbf{e}'' – this is a complete break.

How to (speculatively) tweak Kyber / Module-LWE to achieve A-SEC-CCA. Intuitively, the failure of A-SEC-CCA security for Kyber is due to the existence of many choices of $\mathbf{b} \in R_q^k$ such that computing the term $v = \mathbf{b}\mathbf{r} + \mathbf{e}''$, for \mathbf{r} and \mathbf{e}'' supported on a small range like $[-2, +2]$, is *very far* from inducing wrap-around in the arithmetic modulo q . When reduction modulo q does not occur, then releasing an R_q -element as a hint can completely determine the secret values $(\mathbf{r}, \mathbf{e}'')$ of $\mathbf{b}\mathbf{r} + \mathbf{e}''$. Alternatively, if (sufficient) wrap-around modulo q occurs when computing $\mathbf{b}\mathbf{r} + \mathbf{e}''$, such an R_q -element cannot – “on its own” – convey the full entropy that went into the sampling of $(\mathbf{r} = (\mathbf{r}_1, \dots, \mathbf{r}_k); \mathbf{e}'' = (\mathbf{e}''_1, \dots, \mathbf{e}''_k)) \in R_q^k \times R_q^k$. Of course, in the algebraically-structured case, one not only needs to ensure reduction modulo q , but also that there is no special algebraic structure (e.g., confinement of the arithmetic to non-maximal ideals) that can be adversarially abused. For the sake of the science, we *speculate* that changing the distribution of Kyber (respectively: Module-LWE with rank $k \geq 2$ or preferably even higher rank) to use much larger supports and much higher entropy distributions for \mathbf{r} and \mathbf{e}'' (as well as for \mathbf{e}' in the u -part of the ciphertext, etc.) might lead to a A-SEC-CCA-secure implementation. Concretely, to ensure that the arithmetic in each coordinate wraps around modulo q for every non-adaptive choice of (not easily rejection-samplable) \mathbf{b} that would otherwise “separate the coordinates” as in our gadget-based attack, it seems one needs to sample secret and error coordinates with weight at least up to $\approx \pm\sqrt{q}$, or around $\{-60, \dots, +60\}$ for Kyber’s modulus $q = 3329$. Note that this is a requirement of *larger secrets* than the asymptotically-suggested $\approx \sqrt{n}$ magnitude given by traditional security theorems in the literature [9]. The reason for this is that the typical Hermite Normal Form style of proof (that one can use the (M)LWE error distribution for the secret) inherently leverages that (M)LWE samples are built from uniform \mathbf{A} , which is not our case here. We emphasize that we have no security proof for this idea. In fact – we lack a theory for how to prove its security. While Plain LWE is known to be fairly robust against leakage [27,19] due to being able to use a leftover hash lemma argument that “plays with the dimension,” proving a practically useful form of entropic security for algebraically-structured lattice cryptography is notoriously difficult [15,37]. We re-highlight this gap in the theory as an interesting open problem area, especially when considering practical instantiations similar to Kyber, which we leave for future research.

8 Conclusion

In this work, we addressed the issues accompanying quantum-resistant PAKE designs and constructions from lattice-based PAC KEMs and ideal ciphers. We

presented our PAKE construction **NICE-PAKE**, which eliminates the usage of an IC in the public key authentication step through masking (XOR-ing) a purely uniform part of the splittable key in LWE KEMs. Our security proof relies on standard KEM properties and assumptions, as well as newly introduced ones (A-SEC-CCA, A-CFR-CCA, SPLIT-PKU). However, our construction indeed suffers from the lack of existence of standard LWE and MLWE KEMs, with which NICE-PAKE can be directly instantiated. To overcome this issue, we presented a discussion on concrete modifications for LWE Frodo-style and possible tweaks for MLWE Kyber-style KEMs.

Our findings through the process of devising and proving the security of this construction did indeed answer many questions regarding the possibility of replacing troublesome idealized objects (i.e., the IC) in PAKE designs. Most importantly, we are now confident that the IC can be completely eliminated. We learned that this would even be straightforward if PQC KEMs were slightly more tailored for PAKE design. Eliminating the IC highlighted the fact that neither uniform nor non-uniform public keys are per se favorable for building PAKEs. Both come with a respective (specific) security issue, that is idealized away by IC. From our perspective, FrodoKEM remarkably showed its strength being built from unstructured lattices and standard plain LWE assumptions, which in turn might motivate further work on optimizing the performance of the scheme through HW-SW designs.

An open question, and apparently a first order of business is to investigate the consequences of adjusting the variance value in Frodo’s error distribution. That is both for the core security and the performance of the scheme. The same goes also for MLWE Kyber, where we suggest that using larger supports and higher entropy in the encapsulators secrets could lead to secure splittable key applications. Nonetheless a similar question arises regarding the concrete overall security and performance of the scheme, should such adjustments be taken. Regarding the usage of KEMs for protocol design, and disregarding size and performance differences, we found that there are major differences regarding the effective remaining security of different LWE-based KEMs in protocol contexts; i.e., the loss of bits of security under the different KEM properties (such as A-SEC-CCA) in the security proof substantially varies between the schemes.

We deem deeper research on the security effects of swapping (even highly related) KEMs in higher level cryptographic constructions (such as PAKEs) an important future work. Further, observing the ongoing interest in purely generic KEM-based PAKE designs, it is also worth investigating, whether our first attempt **Simple NICE PAKE** (cf. Sec. 4) could actually be achieved with a formal analysis based only on inherent KEM properties, which we aim at deriving in future work as well. In order to do so, we suggest tackling properties directly aimed at the security of KEM public keys, such that some structured keys are statically far from uniform, yet still computationally close to uniform. However, and to the best of our knowledge, the only key property addressed so far in the literature concerns mainly the uniformity of keys.

Acknowledgments

We thank Marc Fischlin, Hart Montgomery, and Scott Simon for useful discussions. This research work has been partly funded by the German Federal Ministry of Education and Research and the Hessian State Ministry for Higher Education, Research and the Arts within their joint support of the National Research Center for Applied Cybersecurity ATHENE. Daniel Apon’s affiliation with The MITRE Corporation is provided for identification purposes only, and is not intended to convey or imply MITRE’s concurrence with, or support for, the positions, opinions, or viewpoints expressed by the author. This work was supported in part by the MITRE Independent Research and Development Program. Approved for Public Release. Distribution Unlimited. Case Number 24-3386.

References

1. Abdalla, Michel and Eisenhofer, Thorsten and Kiltz, Eike and Kunzweiler, Sabrina and Riepel, Doreen: Password-Authenticated Key Exchange from Group Actions. In: Dodis, Y., Shrimpton, T. (eds.) *Advances in Cryptology – CRYPTO 2022*. Lecture Notes in Computer Science, Springer Nature Switzerland, Cham (2022). https://doi.org/10.1007/978-3-031-15979-4_24
2. Alagic, G., Russell, A.: Quantum-secure symmetric-key cryptography based on hidden shifts. In: Annual international conference on the theory and applications of cryptographic techniques. pp. 65–93. Springer (2017)
3. Albrecht, M.R., Player, R., Scott, S.: On the concrete hardness of learning with errors. *Journal of Mathematical Cryptology*. Volume 9, Issue 3, Pages 169–203, ISSN (Online) 1862-2984, ISSN (Print) 1862-2976 DOI: 10.1515/jmc-2015-0016, October 2015
4. Alnahawi, N., Hövelmanns, K., Hülsing, A., Ritsch, S.: Towards post-quantum secure pake - a tight security proof for ocake in the bpr model. In: *Cryptology and Network Security*. pp. 191–212. Springer Nature Singapore (2024)
5. Alnahawi, N., Müller, J., Oupický, J., Wiesmaier, A.: A comprehensive survey on post-quantum TLS. *IACR Communications in Cryptology* 1(2) (2024). <https://doi.org/10.62056/ahee0iuc>
6. Alperin-Sheriff, J., Apon, D.: Dimension-preserving reductions from LWE to LWR. *Cryptology ePrint Archive*, Paper 2016/589 (2016), <https://eprint.iacr.org/2016/589>
7. Alperin-Sheriff, J., Peikert, C.: Circular and kdm security for identity-based encryption. In: *International Workshop on Public Key Cryptography*. pp. 334–352. Springer (2012)
8. Apon, D., Fan, X., Liu, F.H.: Deniable attribute based encryption for branching programs from lwe. In: *Theory of Cryptography: 14th International Conference, TCC 2016-B, Beijing, China, October 31–November 3, 2016, Proceedings, Part II* 14. pp. 299–329. Springer (2016)
9. Applebaum, B., Cash, D., Peikert, C., Sahai, A.: Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In: Halevi, S. (ed.) *Advances in Cryptology - CRYPTO 2009*. pp. 595–618. Springer Berlin Heidelberg, Berlin, Heidelberg (2009)

10. Arriaga, A., Barbosa, M., Jarecki, S., Skrobot, M.: C'est très chic: A compact password-authenticated key exchange from lattice-based kem. *Cryptology ePrint Archive*, Paper 2024/308 (2024), <https://eprint.iacr.org/2024/308>
11. Beguinet, H., Chevalier, C., Pointcheval, D., Ricosset, T., Rossi, M.: Get a cake: Generic transformations from key encapsulation mechanisms to password authenticated key exchanges. In: *Applied Cryptography and Network Security: 21st International Conference, ACNS 2023, Kyoto, Japan, June 19–22, 2023, Proceedings, Part II*. p. 516–538. Springer-Verlag, Berlin, Heidelberg (2023). https://doi.org/10.1007/978-3-031-33491-7_19
12. Bellare, Mihir and Pointcheval, David and Rogaway, Phillip: Authenticated Key Exchange Secure against Dictionary Attacks. In: *Advances in Cryptology – EUROCRYPT 2000*, vol. 1807. Springer Berlin Heidelberg, Berlin, Heidelberg (2000). https://doi.org/10.1007/3-540-45539-6_11, series Title: Lecture Notes in Computer Science
13. Bellare, S.M. and Merritt, M.: Encrypted key exchange: password-based protocols secure against dictionary attacks. In: *Proceedings 1992 IEEE Computer Society Symposium on Research in Security and Privacy (May 1992)*. <https://doi.org/10.1109/RISP.1992.213269>
14. Black, J.: The ideal-cipher model, revisited: An uninstantiable blockcipher-based hash function. In: *Fast Software Encryption: 13th International Workshop, FSE 2006, Graz, Austria, March 15–17, 2006, Revised Selected Papers 13*. pp. 328–340. Springer (2006)
15. Bolboceanu, M., Brakerski, Z., Perlman, R., Sharma, D.: Order-lwe and the hardness of ring-lwe with entropic secrets. In: Galbraith, S.D., Moriai, S. (eds.) *Advances in Cryptology – ASIACRYPT 2019*. pp. 91–120. Springer International Publishing, Cham (2019)
16. Boneh, D., Lewi, K., Montgomery, H., Raghunathan, A.: Key homomorphic prfs and their applications. In: Canetti, R., Garay, J.A. (eds.) *Advances in Cryptology – CRYPTO 2013*. pp. 410–428. Springer Berlin Heidelberg, Berlin, Heidelberg (2013)
17. Bos, J., Costello, C., Ducas, L., Mironov, I., Naehrig, M., Nikolaenko, V., Raghunathan, A., Stebila, D.: Frodo: Take off the ring! practical, quantum-secure key exchange from lwe. In: *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*. pp. 1006–1018 (2016)
18. Bos, Joppe and Ducas, Leo and Kiltz, Eike and Lepoint, T and Lyubashevsky, Vadim and Schanck, John M. and Schwabe, Peter and Seiler, Gregor and Stehle, Damien: CRYSTALS - Kyber: A CCA-Secure Module-Lattice-Based KEM. In: *2018 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, London (Apr 2018). <https://doi.org/10.1109/EuroSP.2018.00032>
19. Brakerski, Z., Döttling, N.: Hardness of LWE on General Entropic Distributions. In: Canteaut, A., Ishai, Y. (eds.) *Advances in Cryptology – EUROCRYPT 2020*. pp. 551–575. Springer International Publishing, Cham (2020)
20. Brakerski, Z., Gentry, C., Vaikuntanathan, V.: (leveled) fully homomorphic encryption without bootstrapping. *ACM Transactions on Computation Theory (TOCT)* **6**(3), 1–36 (2014)
21. Brakerski, Z., Langlois, A., Peikert, C., Regev, O., Stehlé, D.: Classical hardness of learning with errors. In: *Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing*. p. 575–584. STOC '13, Association for Computing Machinery, New York, NY, USA (2013), <https://doi.org/10.1145/2488608.2488680>

22. Canetti, R., Chen, Y.: Constraint-Hiding Constrained PRFs for NC^1 from LWE. In: Coron, J.S., Nielsen, J.B. (eds.) *Advances in Cryptology – EUROCRYPT 2017*. pp. 446–476. Springer International Publishing, Cham (2017)
23. Chaudhary, Dharminder and Kumar, Uddeshaya and Saleem, Kashif: A Construction of Three Party Post Quantum Secure Authenticated Key Exchange Using Ring Learning with Errors and ECC Cryptography. *IEEE Access* (2023)
24. Cheon, J.H., Kim, D., Kim, D., Lee, J., Shin, J., Song, Y.: Lattice-based secure biometric authentication for hamming distance. In: *Information Security and Privacy: 26th Australasian Conference, ACISP 2021, Virtual Event, December 1–3, 2021, Proceedings*. p. 653–672. Springer-Verlag, Berlin, Heidelberg (2021). https://doi.org/10.1007/978-3-030-90567-5_33
25. Cremers, C., Dax, A., Medinger, N.: Keeping up with the KEMs: Stronger security notions for KEMs and automated analysis of KEM-based protocols. *Cryptology ePrint Archive, Paper 2023/1933* (2023), <https://eprint.iacr.org/2023/1933>
26. Ding, Jintai and Alsayigh, Saed and Lancrenon, Jean and Rv, Saraswathy and Snook, Michael: Provably Secure Password Authenticated Key Exchange Based on RLWE for the Post-Quantum World. In: *Topics in Cryptology – CT-RSA 2017*, vol. 10159. Springer International Publishing, Cham (2017). https://doi.org/10.1007/978-3-319-52153-4_11, series Title: *Lecture Notes in Computer Science*
27. Goldwasser, S., Kalai, Y.T., Peikert, C., Vaikuntanathan, V.: Robustness of the learning with errors assumption. In: *International Conference on Supercomputing (2010)*, <https://api.semanticscholar.org/CorpusID:6166048>
28. Grubbs, P., Maram, V., Paterson, K.G.: Anonymous, robust post-quantum public key encryption. In: *Advances in Cryptology – EUROCRYPT 2022: 41st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Trondheim, Norway, May 30 – June 3, 2022, Proceedings, Part III*. p. 402–432. Springer-Verlag, Berlin, Heidelberg (2022). https://doi.org/10.1007/978-3-031-07082-2_15
29. Guo, S., Song, Y., Guo, S., Yang, Y., Song, S.: Three-party password authentication and key exchange protocol based on mlwe. *Symmetry* **15**, 1750 (09 2023). <https://doi.org/10.3390/sym15091750>
30. Hao, F., van Oorschot, P.C.: Sok: Password-authenticated key exchange – theory, practice, standardization and real-world lessons. In: *Proceedings of the 2022 ACM on Asia Conference on Computer and Communications Security*. p. 697–711. ASIA CCS ’22, Association for Computing Machinery, New York, NY, USA (2022). <https://doi.org/10.1145/3488932.3523256>
31. Hesse, J., Rosenberg, M.: PAKE combiners and efficient post-quantum instantiations. *Cryptology ePrint Archive, Paper 2024/1621* (2024), <https://eprint.iacr.org/2024/1621>
32. Hongfeng Zhu and Xin Hao and Yang Sun: Elliptic Curve Isogenies-Based Three-party Password Authenticated Key Agreement Scheme towards Quantum-Resistant. *J. Inf. Hiding Multim. Signal Process.* **5** (2014), <https://api.semanticscholar.org/CorpusID:16988408>
33. Hosoyamada, A., Yasuda, K.: Building quantum-one-way functions from block ciphers: Davies-meyer and merkle-damgård constructions. In: *Advances in Cryptology–ASIACRYPT 2018: 24th International Conference on the Theory and Application of Cryptology and Information Security, Brisbane, QLD, Australia, December 2–6, 2018, Proceedings, Part I* 24. pp. 275–304. Springer (2018)
34. Katz, Jonathan and Vaikuntanathan, Vinod: Smooth Projective Hashing and Password-Based Authenticated Key Exchange from Lattices. In: *Advances in Crypt-*

- tology – ASIACRYPT 2009, vol. 5912. Springer Berlin Heidelberg, Berlin, Heidelberg (2009). https://doi.org/10.1007/978-3-642-10366-7_37, series Title: Lecture Notes in Computer Science
35. Langlois, A., Stehlé, D.: Worst-case to average-case reductions for module lattices. *Designs, Codes and Cryptography* **75**(3), 565–599 (2015). <https://doi.org/10.1007/s10623-014-9938-4>
 36. Lee, J., Kim, D., Kim, D., Song, Y., Shin, J., Cheon, J.H.: Instant privacy-preserving biometric authentication for hamming distance. *Cryptology ePrint Archive*, Paper 2018/1214 (2018), <https://eprint.iacr.org/2018/1214>
 37. Liu, F.H., Wang, Z.: Rounding in the rings. In: Micciancio, D., Ristenpart, T. (eds.) *Advances in Cryptology – CRYPTO 2020*. pp. 296–326. Springer International Publishing, Cham (2020)
 38. Liu, Z., Sotiraki, K., Tromer, E., Wang, Y.: Snake-eye resistance from LWE for oblivious message retrieval and robust encryption. *Cryptology ePrint Archive*, Paper 2024/510 (2024), <https://eprint.iacr.org/2024/510>
 39. Lyu, Y., Liu, S., Han, S.: Universal composable password authenticated key exchange for the post-quantum world. In: *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. pp. 120–150. Springer (2024)
 40. Lyubashevsky, V., Peikert, C., Regev, O.: On ideal lattices and learning with errors over rings. In: *Advances in Cryptology–EUROCRYPT 2010: 29th Annual International Conference on the Theory and Applications of Cryptographic Techniques, French Riviera, May 30–June 3, 2010*. Proceedings 29. pp. 1–23. Springer (2010)
 41. O’Neill, A., Peikert, C., Waters, B.: Bi-deniable public-key encryption. In: *Advances in Cryptology–CRYPTO 2011: 31st Annual Cryptology Conference, Santa Barbara, CA, USA, August 14–18, 2011*. Proceedings 31. pp. 525–542. Springer (2011)
 42. Pan, J., Zeng, R.: A generic construction of tightly secure password-based authenticated key exchange. In: *Advances in Cryptology – ASIACRYPT 2023*. pp. 143–175. Springer Nature Singapore, Singapore (2023)
 43. Ravi, P., Howe, J., Chattopadhyay, A., Bhasin, S.: Lattice-based key-sharing schemes: A survey. *ACM Comput. Surv.* **54**(1) (2021). <https://doi.org/10.1145/3422178>
 44. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. In: *Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing*. p. 84–93. STOC ’05, Association for Computing Machinery, New York, NY, USA (2005). <https://doi.org/10.1145/1060590.1060603>
 45. Regev, O.: Lattice-based cryptography. In: *Annual International Cryptology Conference*. pp. 131–141. Springer (2006)
 46. Regev, O.: The learning with errors problem. *Invited survey in CCC* **7**(30), 11 (2010)
 47. Santos, B.F.D., Gu, Y., Jarecki, S.: Randomized half-ideal cipher on groups with applications to uc (a) pake. In: *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. pp. 128–156. Springer (2023)
 48. Sato, S., Shikata, J.: So-cca secure pke in the quantum random oracle model or the quantum ideal cipher model. In: *Cryptography and Coding: 17th IMA International Conference, IMACC 2019, Oxford, UK, December 16–18, 2019*, Proceedings 17. pp. 317–341. Springer (2019)
 49. Shannon, C.E.: Communication theory of secrecy systems. *The Bell system technical journal* **28**(4), 656–715 (1949)

50. Tang, Yongli and Li, Ying and Zhao, Zongqu and Zhang, Jing and Ren, Lina and Li, Yuanhong: Improved Verifier-Based Three-Party Password-Authenticated Key Exchange Protocol from Ideal Lattices. *Security and Communication Networks* **2021** (Nov 2021). <https://doi.org/10.1155/2021/6952869>, publisher: Hindawi
51. Taraskin, Oleg and Soukharev, Vladimir and Jao, David and LeGrow, Jason T.: Towards Isogeny-Based Password-Authenticated Key Establishment. *Journal of Mathematical Cryptology* **15**(1) (2020). <https://doi.org/10.1515/jmc-2020-0071>
52. Terada, S., Yoneyama, K.: Password-based authenticated key exchange from standard isogeny assumptions. In: *Provable Security*. pp. 41–56. Springer International Publishing, Cham (2019)
53. Unruh, D.: Towards compressed permutation oracles. In: *International Conference on the Theory and Application of Cryptology and Information Security*. pp. 369–400. Springer (2023)
54. U.S. National Institute of Standards and Technology (NIST): Submission Requirements and Evaluation Criteria for the Post-Quantum Cryptography Standardization Process. <https://csrc.nist.gov/CSRC/media/Projects/Post-Quantum-Cryptography/documents/call-for-proposals-final-dec-2016.pdf> (2016)
55. U.S. National Institute of Standards and Technology (NIST): FIPS 203: Module-lattice-based key-encapsulation mechanism standard. <https://csrc.nist.gov/pubs/fips/203/final> (2024)
56. Wang, Jinhua and Chen, Ting and Liu, Yanyan and Zhou, Yu and Dong, Xinfeng: Efficient Two-Party Authentication Key Agreement Protocol Using Reconciliation Mechanism from Lattice. In: *International Conference on Security and Privacy in New Computing Environments*. Springer (2022)

A The Bellare-Pointcheval-Rogaway (BPR) Model

An adversary \mathcal{A} is given a number of capabilities, or better said actions, called queries. \mathcal{A} interacts with the protocol through these queries with so-called oracles that implement the protocol flow on behalf of honest parties called users (our friends *Alice* and *Bob*). \mathcal{A} may assume any role they wish, i.e., either an active man-in-the-middle (malicious *Mallory*), or a passive connector forwarding messages between honest parties (eavesdropping *Eve*).

Sessions, Users, and Initialization: The game consists of an unlimited number of sessions between protocol instances. These sessions can be initiated by \mathcal{A} in parallel or sequentially whenever they want. Each of these sessions chooses, upon initialization, two protocol instances $\mathcal{U}_i, \mathcal{V}_j$. Instances are users in a set \mathcal{U} and can be viewed as initiators and receivers. An initiator \mathcal{U}_i is assigned their own unique password $\pi_{\mathcal{U}_i} \in \mathcal{D}$ that is also known to a receiver \mathcal{V}_j .

Termination, Accepting, and Partnering: An instance may terminate either in an accepting state (key agreement is successful), or with no output in the case of rejection (key agreement is unsuccessful). A terminated instance may refuse to participate further in a protocol session. Two instances are partnered if they both terminate in accepting state with the same session key. An instance can however be terminated or accepting without being partnered. That is if an instance is ready to use a session key, but does not necessarily have any accepting partner.

Adversarial Model: \mathcal{A} 's goal is to distinguish between real and random session keys. At any point, \mathcal{A} choose one session to test, which then outputs a bit b (random challenge). \mathcal{A} outputs another bit b' , and wins if they guessed correctly, i.e., if $b' = b$. We say that the adversary wins if on issuing a **Test** query for a user \mathcal{U}_i that has terminated in accepting state (i.e. has a session key K) and no **Reveal** or **Corrupt** query has been issued to this user or any user partnered with them, \mathcal{A} correctly guesses the bit selected in the **Test** query. Hence, the security here, same as in KEX security, lies within the ability of the protocol to guarantee the indistinguishability of honest session keys from random ones against any efficient adversary. \mathcal{A} is given access to the following queries, in addition to one oracle modeled as a RO for the hash function \mathcal{H} .

- **Execute**($\mathcal{U}_i, \mathcal{V}_j$): Execute an honest instance of the protocol that terminates in accepting state. \mathcal{A} is given a full transcript of the execution (models eavesdropping).
- **Send**(\mathcal{U}_i, m): Send a message to an honest user that causes them to proceed depending on their state (models impersonating attack)
- **Reveal**(\mathcal{U}_i): Get the final session key for an accepting user, or the rejection symbol \perp for a non-accepting user (models key leaking).
- **Corrupt**(\mathcal{U}_i): Get the password $\pi_{\mathcal{U}_i}$ of a user in the weak-corruption model. In the strong-corruption model \mathcal{A} may also view the internal state of the user (total break of an honest user).
- **Test**(\mathcal{U}_i): Output a real or a random session key based on a bit flip made by the challenger. If a user is unrefresh or not in accepting state, the test returns the rejection symbol \perp . Otherwise, if $b = 1$ it returns the real key K , else if $b = 0$ it returns a random key K' .

Without loss of generality, we assume that connected users in a session have the same password π from dictionary \mathcal{D} , and that the oracle chooses that same password upon initiation of a protocol instance Π_{ij} between two users $(\mathcal{U}_i, \mathcal{V}_j)$.

Freshness An instance or user \mathcal{U}_i is called **fresh**, if neither a **Reveal** query, nor a **Corrupt** query was placed by \mathcal{A} upon it, or upon any other partnered user. Otherwise an instance is then called **unfresh**. This notion disallows \mathcal{A} from winning the AKE security game trivially by testing sessions they previously revealed or corrupted their instances, and any instances partnered with them. The restriction made i.r.w the **Corrupt** query allows for modeling (perfect) forward secrecy (PFS) and adaptive corruption.

BPR AKE Security For a protocol Π and an adversary \mathcal{A} , we define the advantage of \mathcal{A} against the AKE security of Π as $\mathbf{Adv}_{\mathcal{A}}^{\Pi} = (\Pr[b' = b] - \frac{1}{2})$, where \mathcal{A} 's goal is to distinguish between real and random session keys determined by the bit b for an accepting and fresh user. The protocol Π is secure in the BPR model if for all efficient adversaries \mathcal{A} :

$$\mathbf{Adv}_{\mathcal{A}}^{\Pi} \leq \frac{q_s}{|\mathcal{D}|} + \varepsilon(1^\kappa)$$

where q_s is the number of sessions \mathcal{A} actively interacts with, \mathcal{D} is a password dictionary, ε is a negligible function, and κ is the security parameter of the underlying cryptographic primitive.

Definition 17 (Security of Authenticated Key Exchange (AKE)). *For a protocol Π and an adversary \mathcal{A} , we define \mathcal{A} 's advantage with respect to Π as $\text{Adv}_{\Pi}^{\mathcal{A}}$. We say that the adversary wins if on issuing a `Test` query for a user \mathcal{U}_j that has terminated in accepting state (i.e. is in possession of a session key K) and no `Reveal` or `Corrupt` query has been issued to this user, \mathcal{A} correctly guesses the bit selected in the test query. We say the protocol Π is secure if the probability of \mathcal{A} winning is bounded to a negligible quantity. We denote the probability of \mathcal{A} winning in game G_i as $\text{Pr}[G_i]$.*

B Proof Strategy

Passive Security \mathcal{A} 's goal is to compromise the final session key K . Without actively interfering with protocol executions, they only have one option. That is to observe a session transcript attempt to guess K through attacking the KEM. The oracle `Execute` implements an honest protocol session between two connected users. Through the provided interface, \mathcal{A} may view the transmitted pk in its splitted form (\mathbf{b}, z) , and the correspondingly transmitted ciphertext C_b . Considering the available `Test` interface, \mathcal{A} may also query either the real key or a random key resulting from one execution. Therefore, the security aimed at here is similar to the experiment of IND-CCA for a chosen KEM, however, without access to a decryption oracle. Hence, we claim that a KEM with only IND-CPA should also suffice. Intuitively, one would like the query `Test`($\mathcal{U}_i, \mathcal{V}_j$) to always output a random key through replacing the real session key K with a random one K^* chosen uniformly from the same key space \mathcal{K} . We will show that the \mathcal{A} 's advantage in distinguishing between the two keys is (1) determined through outputting a bit b (2) bound to the number of honestly executed sessions q_e they passively observe (3) bound to random password guessing on dictionary size plus a negligible advantage based on the IND-CPA security of the chosen KEM denoted by $\text{Adv}_{\text{KEM}}^{\text{IND-CPA}}(\mathcal{A})$.

Active Security We would also like to show that \mathcal{A} cannot distinguish between real or random keys in sessions they actively interact with. Here, they may, analogously to the passive case, also test a final session key for a partnered instance or for a terminated one in an accepting state. However, \mathcal{A} has the possibility of choosing z, b (i.e. (pk)) and/or C at will through forwarding them via `Send` queries to initiated instances. Hence, we would like to especially capture the cases where they may attempt to trick an honest user to use a malicious pk , and then try to guess the correct password by relating C to the pk it was encapsulated with. The other possibility an adversary has is through choosing C and forwarding it to an honest user. Lastly, we recall that \mathcal{A} may also place a `Reveal` and/or a `Corrupt` query on an instance. We primarily differentiate between the two cases here informally.

Adversary Impersonates Alice: Should \mathcal{A} choose to impersonate *Alice*, they may initialize the protocol and start a receiving instance \mathcal{V}_i via a $\text{Send}(\mathbf{b}, z)$ query with a pk of their choosing. In this case, they may guess (randomly choose) any π^* and use it with the seed ρ of their chosen pk as input for the authentication function f producing z , a quasi encryption of ρ under π^* . Intuitively, an honest Bob decrypts z , however, with the real password π , and gets ρ' to sample \mathbf{A}' , which is with probability depending on the dictionary size $|\mathcal{D}|$ not the same as \mathcal{A} 's choice. Bob then encapsulates pk using his reconstructed \mathbf{A}' and gets (K, C_b) and sends C_b back. Since active instances only accept one pk per session, \mathcal{A} is restricted to one password guess on that instance. Thus, we would first like to show that the advantage of \mathcal{A} choosing the correct password on random guessing is negligible (trivial). Second, we will show that neither a malicious key pair (sk, pk) nor a received ciphertext C or the decapsulated key K could raise \mathcal{A} 's advantage in determining an honest pk leading them to the correct π .

Anonymity: Limiting \mathcal{A} 's advantage on relating C to any pk (observing decapsulations on ciphertexts without prior manipulation of a key pair) is somewhat simple and can be based on the ANO-CPA property of the chosen KEM. In other words, should \mathcal{A} be able to distinguish between public keys used for encapsulating a shared key into a chosen ciphertext with non-negligible advantage, it would break the anonymity of the chosen KEM. Therefore, \mathcal{A} 's advantage is bounded by a probability not better than random guessing denoted by $\text{Adv}_{\text{KEM}}^{\text{ANO-CPA}}(\mathcal{A})$.

A-Part-Secrecy: We would also like to account for the case where \mathcal{A} chooses the public key pair at will, while also utilizing an offline dictionary attack on the public key space that is connected to the password dictionary. Here, we need to show that the chosen KEM does not allow randomly or maliciously chosen secret keys sk to correctly decapsulate on non-matching A-parts for known public keys pk . We rely on the assumption, that \mathcal{A} is not capable of finding a second key pair that will decrypt a valid ciphertext and consequently decapsulate into a valid shared key correctly for a non-matching A-part resulting from reconstruction under the correct password, regardless of their choice for the key pair (sk, pk) . Thus, regardless of the chosen sk , a decryption will always fail if a non-matching A-part was used by \mathcal{A} and Bob. Here, \mathcal{A} has the following options: 1) Iteratively change their chosen b part after receiving C in order to go through all possible A-part values resulting from reconstruction over the password dictionary, 2) iteratively changing sk used for decryption, and 3) combining both iterations to test all possible A-parts with all possible secret keys. Option 1 will yield successful with negligible probability bound by random password guessing, as the b-part contains the sk originally chosen by \mathcal{A} . Option 2 will yield unsuccessful, as the decryption will fail for all non-matching A-parts. Option 3 seems less promising, as the number of possible combinations will increase exponentially for all possible sk values iterated over password dictionary sizes. Nonetheless, such arguments break the purpose of a generic construction as they will force treating the KEM in white-box manner to formulate a reduction based on the encryption and decryption functions of the underlying PKE. To overcome this

obstacle, we rely on the notion of A-part-secrecy (A-SEC-CCA) for a KEM with splittable public keys. Using this property of KEM, we can argue about the case where \mathcal{A} freely chooses a key pair without opening up the KEM to apply a security reduction in the proof based on the KEM PKE, and thus maintain the generic construction. The reduction to this property will follow from the fact that using the seed in the authentication step binds an adversary to a subset of all possible key pairs limited by the dictionary size for reconstructing all possible $\mathbf{A}_i \in \mathbf{A}_\pi$. Meaning, that the adversary will commit to a password guess upon sending the value z , which will limit their choice for creating key pairs by the possible values of \mathbf{A} that can be used in a pk , which in turn leads to the impossibility of choosing secret keys that are advantageous for the adversary in decapsulation and re-encryption.

Adversary Impersonates Bob: Should \mathcal{A} impersonate Bob, they receive (\mathbf{b}, z) from Alice. Not knowing π , \mathcal{A} executes an offline dictionary attack to reconstruct all possible ρ_i values. \mathcal{A} then sends some C to Alice, who in turn decapsulates C with her honest sk . Here, a ciphertext can be any of the following: *Correct, incorrect, valid, or invalid*. Correctness means that C is a ciphertext that can be generated by the probabilistic encapsulation algorithm. Validity means that C will decrypt using the secret key corresponding to the public key used for encapsulation, without resulting in non-matching keys between Alice and Bob. Since the chosen KEM provides implicit rejection, Alice will get $K \neq K'$ and \mathcal{A} would have failed in creating the same session key. \mathcal{A} may try all possible \mathbf{A}_i values. However, \mathcal{A} will fail to create an identical pair (C, K) under any number of different keys, since the KEM encapsulation is probabilistic, and the KEM provides strong collision freeness. In other words, \mathcal{A} will fail at creating a C that decapsulates into the same K as an honest Bob. Since we adopt a modified version of collision freeness, this attack is bound to the splittable collision freeness of the chosen KEM denoted by $\text{Adv}_{\text{KEM}}^{\text{A-CFR-CCA}}(\mathcal{A})$. Should \mathcal{A} choose to place either a **Reveal** or **Corrupt** query on Alice after sending C , they can obviously learn the final key decapsulated by Alice or her state, including her key pair. However, learning the final key has no additional value for \mathcal{A} , as the keys are expected to not match on non-matching A-parts. Further, corrupting an instance renders \mathcal{A} unable to test the session connected to this instance, and therefore, they cannot win the security game by guessing the test bit b .

Adaptive Corruption and Perfect Forward Secrecy \mathcal{A} may leverage the **Corrupt** query to view the password of an honest user (weak corruption), or the complete internal state of an instance including the password (strong corruption). However, we recall that \mathcal{A} can only win if they guess the test bit in a fresh accepting instance. Corrupting an instance will render it unfresh, and thus disqualify it as a possibly winnable one. It remains to discuss, if corruption could aid \mathcal{A} in compromising fresh, yet non-partnered sessions. In the previously described scenarios, \mathcal{A} could place such a query and easily see the honest KEM key pair of Alice. But even if they know how C was created, they will not be able to

compromise other sessions since the KEM key pairs are always ephemeral (i.e., are newly generated for each fresh instance in a unique session) and the encapsulation routine is always probabilistic (i.e., each encapsulation yields statistically independent session keys and ciphertexts with overwhelming probability). They will also fail to forward ciphertexts to other connected sessions as per the KEM anonymity, and they will fail to create identical session keys for different sessions, due to the robustness of the KEM. Hence, any later on compromised protocol instance with an ephemeral KEM key pair will not affect previous ones, and thus Π ensures PFS.

C Plain, Ring and Module LWE

The decisional LWE problem [46] is basically to distinguish between (or find in the search variant) random linear equations from uniform equations after applying a small amount of noise. The Regev LWE-based public cryptosystem [44] is parametrized by a security parameter n , two integers (m, q) , and a probability distribution \mathcal{X} over \mathbb{Z}_q . It defines the private key sk as a small vector $\mathbf{s} \in \mathbb{Z}_q^n$. The public key pk consists of LWE samples $(\mathbf{a}_i, b_i)_{i=1}^m$ from the LWE distribution such that $\mathbf{b} \equiv \mathbf{a}\mathbf{s} + \mathbf{e} \pmod{q}$ with the secret \mathbf{s} , modulus q and the small noise (error parameter) $\mathbf{e} \in \mathbb{Z}_q$. Without an error \mathbf{e} finding a secret \mathbf{s} would be easy using Gaussian elimination, hence the error, which according to Regev resembles the problem of decoding random linear codes [46], yet was proven secure through a reduction to the SVP problem in [44].

In the ring variant (RLWE), the ring $R_q = \mathbb{Z}[x]/\langle x^n + 1 \rangle$ is defined as the ring of integer polynomials modulo $f(x)$ with n being a power of 2 and q a prime modulus so that $q = 1 \pmod{2n}$. Consequently, the elements of R_q are residues modulo both $f(x)$ and q . They can thus be viewed as integer polynomials of degree less than n with coefficients in a set of canonical representatives in \mathbb{Z}_q . An RLWE sample is chosen as (\mathbf{a}, \mathbf{b}) with $\mathbf{b} \equiv \mathbf{a}\mathbf{s} + \mathbf{e} \pmod{q}$ where $\mathbf{a} \in R_q$ is chosen uniformly, $\mathbf{s} \in R$ is a fixed secret chosen together with $\mathbf{e} \in R$ from an error distribution [46]. Hence, an RLWE sample $(\mathbf{a}, \mathbf{b}) \in R_q \times R_q$ can replace n standard LWE samples $(\mathbf{a}, \mathbf{b}) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$, which reduces the size of the public key. The Lyubashevsky-Peikert-Regev RLWE cryptosystem [40] from ideal lattices encrypts a message $z \in \{0, 1\}^n$ by using its bits as $\{0, 1\}$ coefficients of a polynomial and choosing random elements $\mathbf{r}, \mathbf{e}_1, \mathbf{e}_2 \in R$ and outputting the encryption of z as a pair $(\mathbf{u}, \mathbf{v}) \in R_q^2$ where $\mathbf{u} = \mathbf{a}\mathbf{r} + \mathbf{e}_1 \pmod{q}$ and $\mathbf{v} = \mathbf{b}\mathbf{r} + \mathbf{e}_2 + \lfloor q/2 \rfloor \cdot z \pmod{q}$. The decryption computes $\mathbf{v} - (\mathbf{u} \cdot \mathbf{s}) = (\mathbf{r} \cdot \mathbf{e} - \mathbf{s} \cdot \mathbf{e}_1 + \mathbf{e}_2) + \lfloor q/2 \rfloor \cdot m \pmod{q}$. The coefficients of the small terms $(\mathbf{r} \cdot \mathbf{e} - \mathbf{s} \cdot \mathbf{e}_1 + \mathbf{e}_2)$ are in the range $-q/4, q/4$. Thus, a bit in z is then decrypted as 0 if $\mathbf{b} - \langle \mathbf{a}, \mathbf{s} \rangle$ is closer to 0 than to $\lfloor q/2 \rfloor \pmod{q}$, otherwise it is 1.

The module (MLWE) variant was first defined by Brakerski et al. [20] and further studied by Langlois and Stehlé [35]. In essence, MLWE also replaces the integers in \mathbb{Z} by a ring of algebraic integers R of a number field K . The new component here is $M \subseteq K^d$, a module of R . Therefore, the parameter n is introduced as the degree of the number field, and the integer d denotes the

module rank. With M being a rank d module and K of degree n , the resulting module lattice has the dimension $N = nd$. Hence, the MLWE problem generalizes both LWE and RLWE, as the RLWE variant is obtained if the module rank is $d = 1$. As a consequence, an MLWE sample is defined as (\mathbf{a}, \mathbf{b}) with $\mathbf{b} \equiv \mathbf{a}\mathbf{s} + \mathbf{e} \pmod R$ where $\mathbf{a} \in R_q^d$ is chosen uniformly, $\mathbf{s} \in R_q^d$ is a fixed secret, and \mathbf{e} is sampled from an error distribution. The matrix representation for the lattice base is used when the number of samples m is fixed, which is denoted by $\mathbf{A} \in R_q^{m \times d}$.

D Extended Learning with Errors (eLWE)

The extended Learning with Errors assumption (eLWE) was first defined by O’Neill et al. in [41] in the context of proving security of deniable encryption from lattices (cf. Apon et al. [8]). eLWE has additionally been used by Alperin-Sheriff and Peikert [7] to show security theorems for circular and key-dependent message security of lattice identity-based encryption, as well as to show the classical hardness of LWE itself by Brakerski et al. [21]. The basic idea is to *extend* LWE security proofs, like

$$(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}) \stackrel{\text{comp}}{\approx} (\mathbf{A}, \mathbf{u}),$$

to a setting with additional auxiliary information z ; i.e.,

$$(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}, z) \stackrel{\text{comp}}{\approx} (\mathbf{A}, \mathbf{u}, z).$$

In other words, eLWE claims that standard LWE computational hardness assumptions still hold, up to some very small concrete loss in bit-security — even in the presence of certain, additional (potentially non-uniform) “hints” z on the hidden terms (\mathbf{s}, \mathbf{e}) of an LWE equation. The typical form of such “eLWE leakage” z on the hidden values (\mathbf{s}, \mathbf{e}) takes the form of a chosen *inner product* $(\mathbf{z}, \langle \mathbf{z}, \mathbf{e} \rangle)$, where some vector \mathbf{z} is *adversarially chosen* and $\langle \mathbf{z}, \mathbf{e} \rangle$ is the derived hint on honestly generated error \mathbf{e} . Intuitively, inner products reveal very little information on the LWE secret (\mathbf{s}, \mathbf{e}) .

Formally, define the torus $\mathbb{T} = \mathbb{R}/\mathbb{Z}$, then we have:

Definition 18 (Extended (Plain) Learning with Errors (eLWE)). *For $n, m, q, t \geq 1, \mathcal{Z} \subseteq \mathbb{Z}^m$, and a distribution χ over $\frac{1}{q}\mathbb{Z}^m$, the $\text{eLWE}_{n,m,q,\chi,\mathcal{Z}}$ problem is as follows. The algorithm gets to choose $\mathbf{z} \in \mathcal{Z}$ and then receives the tuple*

$$(\mathbf{A}, \mathbf{b}, \langle \mathbf{e}, \mathbf{z} \rangle) \in \mathbb{T}_q^{n \times m} \times \mathbb{T}_q^m \times (1/q)\mathbb{Z}.$$

Its goal is to distinguish between two cases: First, $\mathbf{A} \in \mathbb{T}_q^{n \times m}$ is chosen uniformly, $\mathbf{e} \in (1/q)\mathbb{Z}^m$ is chosen from χ , and $\mathbf{b} = \mathbf{A}^t \mathbf{s} + \mathbf{e} \pmod 1$ where $\mathbf{s} \in \{0, \dots, q-1\}^n$ are chosen uniformly. The second case is identical, except that \mathbf{b} is chosen uniformly in \mathbb{T}_q^m independently of everything else.

Prior works define a multi-hint version of eLWE, but we will only need the single-hint version as stated above. The hardness of eLWE follows from LWE, summarized as follows.

Theorem 6. For any $n \geq 2, q \geq 1, \varepsilon \in (0, 1/2)$, and $\alpha, r \geq (\ln(2m(1+1/\varepsilon))/\pi)^{1/2}/q$, there is a reduction from $\text{LWE}_{n+1,m,q,\alpha}$ to $\text{eLWE}_{n,m,q,(\alpha^2\xi^2+r^2)^{1/2},\mathcal{Z}}$ that reduces the advantage by at most $33\varepsilon/2$, where ξ is a small constant factor.

Next, we make a simple observation, that Plain eLWE – originally designated to account for leakage on the error term e – also handles the case of leakage on the secret vector s in certain scenarios.

Definition 19 (Extended Learning With Errors (eLWE, alt version: secret leakage)). For $n, q, t \geq 1, \mathcal{Z} \subseteq \mathbb{Z}^n$, and a distribution χ over $\frac{1}{q}\mathbb{Z}^n$, the $\text{eLWE}_{n,q,\chi,\mathcal{Z}}$ problem is as follows. The algorithm gets to choose $\mathbf{z} \in \mathcal{Z}$ and then receives the tuple

$$(\mathbf{A}, \mathbf{b}, \langle \mathbf{s}, \mathbf{z} \rangle) \in \mathbb{T}_q^{n \times n} \times \mathbb{T}_q^n \times (1/q)\mathbb{Z}.$$

Its goal is to distinguish between two cases: First, $\mathbf{A} \in \mathbb{T}_q^{n \times n}$ is chosen uniformly (conditioned on being invertible), $\mathbf{s}, \mathbf{e} \in (1/q)\mathbb{Z}^n$ is chosen from χ , and $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e} \bmod 1$. The second case is identical, except that \mathbf{b} is chosen uniformly in \mathbb{T}_q^n independently of everything else.

We claim that this version of eLWE (with leakage on the secret) asymptotically follows from standard eLWE (with leakage on the error) under the condition that \mathbf{A} is invertible (modulo q), which requires $n = m$ in the previous statement of eLWE. In this case, $\mathbf{A}\mathbf{s} + \mathbf{e} = \mathbf{s} + \mathbf{A}^{-1}\mathbf{e}$ exactly, and the “alt” theorem follows by swapping the roles of $(\mathbf{A}, \mathbf{A}^{-1})$ and (\mathbf{s}, \mathbf{e}) . The distinction between $\mathbf{s} \leftarrow U(\mathbb{Z}_q^n), \mathbf{e} \leftarrow \chi$ and $\mathbf{e}, \mathbf{s} \leftarrow \chi$ leads to at most a $(1/q)$ multiplicative loss in security (cf. [41]), which is clearly not tight. In practical instantiations from FrodoKEM, the proofs from the literature give a range of between $\ll 1$ and < 15 bits of security lost, and optimizing this concrete analysis is left to future work.

E Non-uniform Module Learning with Errors (NMLWE)

For completeness, we define a module-lattice version of NLWE, adapted from Canetti and Chen [22, Lemma 2.13]

Definition 20 (Non-uniform Module Learning with Errors). Let $\lambda \in \mathbb{N}$ be the security parameter. Let $m, m, q, p \in \mathbb{N}, \sigma$ s.t. $0 < \sigma < q$. Let $R = \mathbb{Z}[X]/(x^n + 1)$ for n a power of 2, γ_σ be a distribution over $R^{m \times m}$ parameterized by σ , and χ_σ be distribution over R parameterized by σ , with $\|\gamma_\sigma\|, \|\chi_\sigma\| \leq \sigma\sqrt{m}$.

The NMLWE problem asks to distinguish between samples $(\mathbf{D}, \mathbf{K}\mathbf{D} + \mathbf{E}) \in (R^{m \times m} \times R^{1 \times m})$ from $(\gamma \times U(R_q^{1 \times m}))$, where $\mathbf{D} \leftarrow \gamma_\sigma$ is possibly non-uniform, $\mathbf{K} \leftarrow U(R_q^{1 \times m}), \mathbf{E} \leftarrow \chi_\sigma^{1 \times m}$.

We cite Canetti and Chen’s [22, Lemma 2.13] security claim for NMLWE.

Theorem 7. For $R = \mathbb{Z}[X]/(X^n + 1)$ where n is a power of 2, set parameters $m \geq 2n \log(q), \sigma = \omega(\sqrt{n \log(q)})$, and set discrete Gaussian distributions $\gamma_\sigma = D_{R^m, \sigma}^{1 \times m}, \chi_\sigma = D_{R, \sigma}$.

Then the hardness of NMLWE follows from the hardness of Plain LWE.

On the weakness of NMLWE security proofs. We remark that – to date – NMLWE has only been shown secure in contexts where one can choose a *high rank*; that is, $m \geq 2n \log(q)$. When the module-rank is appreciably lower than linear in the security parameter (especially in the case of Kyber, which uses rank 2, 3, or 4), known security reduction are vacuous.

To wit, we speculate that improving such security proofs to allow for arbitrarily small, constant rank would imply general-purpose program obfuscation $i\mathcal{O}$ for P/poly from standard lattice assumptions. We view this as negative evidence for obtaining such proof in the near term.

F Extended Module Learning with Errors (eMLWE)

For completeness, following Alperin-Sheriff and Apon [6], we define a module-lattice version of eLWE as follows:

Definition 21 (Extended Module Learning with Errors (eMLWE)). For security parameter $1^\kappa \in \mathbb{N}$, let $n = n(\lambda)$ be an integer dimension, let $f(x) = x^d + 1$ where $d = d(1^\kappa)$ is a power of 2, let $q = q(1^\kappa) > 2$ be an integer, let $R = \mathbb{Z}[x]/(f(x))$ and $R_q = R/qR$, and let $\chi = \chi(1^\kappa)$ be a distribution over R . Then, the $\text{eMLWE}_{n,d,f,q,\chi}$ problem is to distinguish between the following two distributions:

In the first distribution, one samples (\mathbf{a}_i, b_i) uniformly from R_q^{n+1} . The adversary chooses $z_i \in R_q^n$ and then, for u_i uniformly from R_q^n , is presented with:

$$(\mathbf{a}_i, b_i)_i, (z_i)_i, \text{Tr}(\langle z_i, u_i \rangle)_i.$$

In the second distribution, one first draws $\mathbf{s} \leftarrow R_q^n$ uniformly, and samples $(\mathbf{a}_i, b_i) \in R_q^{n+1}$ by sampling $\mathbf{a}_i \leftarrow R_q^n$ uniformly, $e_i \leftarrow \chi$, and setting $b_i = \langle \mathbf{a}_i, \mathbf{s} \rangle + e_i$. The adversary chooses $z_i \in R_q^n$ and then is presented with:

$$(\mathbf{a}_i, b_i)_i, (z_i)_i, \text{Tr}(\langle z_i, \mathbf{s} \rangle)_i,$$

where the use of the trace function $\text{Tr}(\cdot)$ explicitly permits the adversary to request any \mathbb{Q} -linear function of the secret vector, as viewed in the coefficient embedding, for its hints.

We cite Alperin-Sheriff and Apon’s [6, Lemma 3.3] security claim for eMLWE.

Theorem 8. *There is a reduction from $\text{LWE}_{d,w,q,\alpha}$ to either of:*

1. *Following [7]: $\text{eMLWE}_{d,w,q,\alpha,k}$ that reduces the advantage by at most q^k , multiplicatively.*
2. *Following [21]: $\text{eMLWE}_{d+k,w,q,(\alpha^2+r^2)^{1/2},k}$, where $r \geq \omega(\sqrt{\log(w)})$, which reduces the advantage by at most a negligible amount, additively – given the change in dimension and error rate.*

On the weakness of eMLWE security proofs. A key gap in the security offered by eMLWE is that it considers leakage of (potentially, traces of) inner products $\langle z_i, u_i \rangle$ over the various coordinates i of the vectors of polynomials in the problem statement. Fundamentally, there is no compression of dimension as in (Plain) eLWE, so an entire ring element leaks (compared to a single integer modulo q for Plain eLWE), which – to date – prevents many applications in entropic security arguments.