Towards Leakage-Resistant Post-Quantum CCA-Secure Public Key Encryption

Clément Hoffmann[1], Benoît Libert[2,3], Charles Momin[1],
Thomas Peters[1], François-Xavier Standaert[1]

[1] Crypto Group, ICTEAM Institute, UCLouvain, Louvain-la-Neuve, Belgium.
[2] CNRS, Laboratoire LIP, France.
[3] ENS de Lyon, Laboratoire LIP (U. Lyon, CNRS, ENSL, Inria, UCBL), France.

Abstract. As for any cryptographic algorithm, the deployment of post-quantum CCA-secure public key encryption schemes may come with the need to be protected against side-channel attacks. For existing post-quantum schemes that have not been developed with leakage in mind, recent results showed that the cost of these protections can make their implementations more expensive by orders of magnitude. In this paper, we describe a new design, coined POLKA, that is specifically tailored for this purpose. It leverages various ingredients in order to enable efficient side-channel protected implementations such as: (i) the rigidity property (which intuitively means that de-randomized encryption and decryption are injective functions) to avoid the very leaky re-encryption step of the Fujisaki-Okamoto transform, (ii) the randomization of the decryption thanks to the incorporation of a dummy ciphertext, removing the adversary's control of its intermediate computations and making these computations ephemeral, (iii) key-homomorphic computations that can be masked against side-channel attacks with overheads that scale linearly in the number of shares, (iv) hard physical learning problem to argue about the security of some critical unmasked operations. Furthermore, we use an explicit rejection mechanism (returning an error symbol for invalid ciphertexts) to avoid the additional leakage caused by implicit rejection. As a result, all the operations of POLKA can be protected against leakage in a much cheaper way than state-of-the-art designs, opening the way towards schemes that are both quantum-safe and leakage-resistant.

1 Introduction

Recent research efforts showed that designing post-quantum chosen-ciphertext-secure public-key encryption (PKE) schemes that allow efficient implementations offering side-channel security guarantees is extremely challenging with existing techniques. One well-documented issue arises from the Fujisaki-Okamoto (FO) transform that is frequently used for building key encapsulation mechanisms (KEMs) with chosen-ciphertext (IND-CCA) security from PKE schemes or KEMs that only provide weak security notions like one-wayness under passive attacks (OW-CPA security) [42, 43]. The FO transformation and its variants are, for example, used in the NIST post-quantum finalists KYBER [5, 21] and SABER [11, 31], where the CCA-secure KEM is combined with a secret-key (authenticated) encryption scheme into a hybrid PKE system.

Recall that a KEM system ($\mathsf{Keygen}, \mathsf{Encaps}, \mathsf{Decaps}$) is a PKE scheme that does not take any plaintext as input, but rather computes an encryption of a random symmetric key K. To encrypt a plaintext M via the hybrid KEM/DEM framework [74], the Encaps algorithm often samples a random m, which is used to derive a symmetric key K and random coins r from a random oracle $(K, r) \leftarrow H(m)$ before deterministically encapsulating K as $c_{kem} = \mathsf{Encaps}_{pk}(m, r)$. Next, a secret-key scheme (E, D) (a.k.a. data encapsulation mechanism, or DEM) is used to compute $c_{sym} = \mathsf{E}_K(M)$ in order to obtain a hybrid PKE ciphertext $c = (c_{kem}, c_{sym})$. The receiver can then recover $m = \mathsf{Decaps}_{sk}(c_{kem})$ and $(K, r) \leftarrow H(m)$ before obtaining $M = \mathsf{D}_K(c_{sym})$. It is known that the hybrid construction provides IND-CCA security if the underlying KEM is itself IND-CCA-secure and if the DEM satisfies a similar security notion in the secret-key setting [74]. In order to secure the KEM part against chosen-ciphertext attacks, the FO transform usually checks the validity of the incoming c_{kem} by testing if $c_{kem} = \mathsf{Encaps}_{pk}(m, r)$ (a step known as "re-encryption") after having recovered the random coins r from $(K, r) \leftarrow H(m)$ upon decryption.

In the FO transform, the first computation during a decryption attempt is $\mathsf{Decaps}_{sk}(c_{kem})$, where Decaps is the underlying decapsulation of the OW-CPA secure KEM. While this has no impact in a black-box security analysis, in the context of side-channel chosen-ciphertext attacks the adversary remains able to target this component using many c_{kem} values [63,68,77] of its choice, leaving an important source of vulnerabilities. Indeed, the adversary is free to adaptively feed Decaps_{sk} with (invalid) ciphertexts and craft c_{kem} in such a way that an internal message m with only few unknown bits is re-encrypted via the FO transform. This allows side-channel attacks to directly exploit the leakage of these bits obtained during the re-encryption test $c_{kem} \overset{?}{=} \mathsf{Encaps}_{pk}(m, r)$ to infer information about sk. This task is surprisingly easy since all the leakage samples of the deterministic re-encryption can be exploited for this purpose (i.e., much more than the few rounds of leakage that are typically exploited in divide-and-conquer side-channel attacks against symmetric encryption schemes) [59].

In parallel, several pieces of work started to analyze masked implementations of KYBER and SABER [12,18,22,41]. These works typically indicate large overheads when high security levels are required, which can be directly connected to a large amount of leaking intermediate computations [6]. In particular, these implementations all consider a uniform protection level for all their operations, that is in contrast with the situation of symmetric cryptography where so-called leveled implementations, in which different (more or less sensitive) parts of a mode of operation are protected with different (more or less expensive) side-channel countermeasures, can lead to important performance gains [15].

In this paper, we therefore initiate the study of quantum-safe CCA-secure public-key encryption schemes that have good features for leakage-resistant (LR) implementations. The seed ingredients we propose for this purpose are three-fold. First, we leverage the *rigidity* property introduced by Bernstein and Persichetti [16], as it allows building CCA-secure encryption schemes without relying on re-encryption nor on the FO transform. Despite removing an important

source of leakage, getting rid of the FO transform is not yet sufficient to enable leveled implementations for KYBER (or SABER), since the rest of their operations remains expensive to protect [6]. Therefore, we also propose to randomize the decryption process by incorporating a "dummy ciphertext". It brings the direct benefit of removing the adversary's control on all intermediate computations that are dummied, while making these computations ephemeral, which is in general helpful against leakage. This second step already allows an interesting leveling between computations that require security against simple power analysis (SPA) and differential power analysis (DPA) attacks.[1] Eventually, we observe that the structure of the KEM's remaining DPA target shares similarities with the key-homomorphic re-keying schemes used in symmetric cryptography to prevent side-channel attacks [35, 39, 61]. Building on this observation, we propose to implement this DPA target such that only its key-homomorphic parts are (efficiently) protected thanks to masking, by relying on the recently introduced Learning With Physical Rounding (LWPR) assumption [38]. In short, the LWPR assumption is a physical version of the crypto dark matter introduced by Boneh *et al.* [20]. The latter assumes that low-complexity PRFs can be obtained by mixing linear mappings over different small moduli. LWPR further leverages the possibility that one of these mappings is computed by a leakage function.

We additionally observe that by carefully instantiating the symmetric authenticated encryption scheme of the DEM as an Encrypt-then-MAC one with a one-time key-homomorphic MAC, the overheads due to the side-channel countermeasures can be reduced to linear in the number of shares used for masking for this part of the computation. And we finally combine these different ingredients into a new efficient post-quantum CCA-secure public-key encryption scheme, called POLKA (standing for POst-quantum Leakage-resistant public Key encryption Algorithm), that *simultaneously* provides excellent features against leakage and a proof of IND-CCA security (in the sense of the standard definition without leakage) under the standard RLWE assumption.

Without leakage, we show that POLKA provides CCA security in the quantum random oracle model (QROM) [19]. Our construction is a hybrid KEM-DEM encryption scheme built upon a variant of a public-key encryption scheme due to Lyubashevsky, Peikert and Regev (LPR) [58], which is well-known to provide IND-CPA security under the ring learning-with-errors (RLWE) assumption. In order to obtain a KEM, we modify the LPR system so as to recover the sender's random coins upon decryption. In contrast with the FO transformation and its variants, this is achieved without derandomizing an IND-CPA system, by deriving the sender's random coins. Instead of encrypting a random message m to derive our symmetric key K, we always "encrypt" 0 and hash the random coins consisting of a tuple $(r, e_1, e_2) \in R$ of small-norm ring elements sampled from the noise distribution. These elements (r, e_1, e_2) are then encoded into a pair $c_{kem} = (a \cdot r + e_1, b \cdot r + e_2) \in R_q^2$, where $a, b \in R_q$ are random-looking elements included

---

[1] Informally, SPAs are side-channel attacks where the adversary can only observe the leakage of a few inputs to the target operation for a given secret. DPAs are attacks where the adversary can observe the leakage of many such inputs.

in the public key. Using its secret key, the decryptor can extract (r, e_1, e_2) from c_{kem} and check their smallness. This verification/extraction step is designed in such a way that decapsulation natively provides rigidity [16] without relying on re-encryption. Namely, due to the way to recover (r, e_1, e_2) from c_{kem}, we are guaranteed that deterministically re-computing $c_{kem} = (a \cdot r + e_1, b \cdot r + e_2)$ would yield the incoming ciphertext. This allows dispensing with the need to explicitly re-compute c_{kem} in the real scheme, thus eliminating an important source of side-channel vulnerability that affects KYBER and SABER.

In a black-box security analysis, our KEM can be seen as an injective trapdoor function that maps $(r, e_1, e_2) \in R^3$ to $(a \cdot r + e_1, b \cdot r + e_2)$. As long as we sample (r, e_1, e_2) from a suitable distribution, c_{kem} is pseudorandom under the RLWE assumption. However, to ease the use of efficient side-channel countermeasures upon decryption, we also leverage the fact that our injection satisfies a (bounded) form of additive homomorphism for appropriate parameters. That is, if we generate what we call a *dummy* ciphertext c'_{kem} by having the decryptor honestly run the basic encapsulation step using its own random coins, the decapsulation of $\bar{c}_{kem} = c_{kem} + c'_{kem}$ should give the sum of the random coins chosen by the sender and the receiver. Then, we can easily remove the additional *dummy* random coins after some additional tests. Introducing $\bar{c}_{kem}$ in the decryption process removes the adversary's freedom of forcing the computation of $\mathsf{Decaps}_{sk}(c_{kem})$ to take place on a c_{kem} under its control, which helps us protecting the secret key. Moreover, the underlying coins of c_{kem} are now split into two shares upon decryption and they are only recombined in a step where we can safely derive K. To implement this idea, we prove the CCA security of our scheme in its variant endowed with a probabilistic decryption algorithm.

With leakage, we argue that POLKA offers a natural path towards efficient leveled implementations secured against side-channel attacks. For this purpose, we first use a methodology inspired from [15] to identify the level of security required for all its intermediate computations. We then focus on how to secure the polynomial multiplication used in POLKA against DPA, by combining masking for its key-homomorphic parts and a variant of the aforementioned LWPR assumption after the shares recombination. Our contributions in this respect are twofold. On the one hand, we define the LWPR variant on which POLKA relies and discuss its difference from the original one. Given that LWPR is an admittedly recent assumption and in view of the important performance gains it can lead to, we additionally specify instances to serve as cryptanalysis targets. On the other hand, we describe a hardware architecture for these masked operations, which confirms these excellent features (e.g., simplicity to implement them securely, performance overheads that are linear in the number of shares). Overall, protecting the long-term secret of our prototype implementation only needs to combine the masking of key-homomorphic computations (which has linear overheads in the number of shares) with SPA security for other computations, that is quite directly/cheaply obtained thanks to parallelism in hardware. Protecting the message confidentiality additionally requires protecting its symmetric cryptographic components (i.e., hash function and authenticated encryption).

We insist that our leakage analysis is not (yet) about sufficient security conditions but about necessary conditions that implementers must fulfill to avoid critical attack paths. Such an analysis has been shown sufficiently informative to compare designs in the symmetric setting. We see no reason why things would significantly differ here. Formally proving the leakage-resistance of POLKA is a natural next step and an important open problem. So our goal is to show that (i) by considering the need for side-channel countermeasures as a design criterion, we can considerably limit the attack vectors, and (ii) by combining design tweaks (e.g., rigidity, dummy operations, LWPR, key-homomorphism) we can make the cost of countermeasures significantly lower than, say KYBER or SABER.

As a last result to confirm the generality of our findings, we also show that they can naturally apply to an LR variant of the NTRU cryptosystem [27], which already satisfies the rigidity property, can be enhanced with a dummy mechanism and has internal computations that also generate LWPR samples.

2 Technical Overview of POLKA & Related Works

TECHNICAL OVERVIEW. Our construction can be seen as a rigid and randomness-recovering version of the RLWE-based encryption scheme described in [58]. By "randomness-recovering," we mean that the decryption procedure recovers the message *and* the sender's random coins. A randomness-recovering encryption scheme is rigid [16] if, when the decryptor obtains a message m and randomness r, running the encryption algorithm on input of (m, r) necessarily yields the incoming ciphertext. While rigidity can always be achieved by adding a re-encryption step (as pointed out in [48]), this generally introduces one-more place of potential side-channel vulnerabilities, which is precisely exploited in [63,68,77]. In order to eliminate the need for an explicit re-encryption step, it is thus desirable to have a decryption algorithm which is natively injective (when seen as a deterministic function). The first difficulty is thus to build a rigid, randomness-recovering PKE/KEM under the standard RLWE assumption. Our goal is to achieve this without sacrificing the efficiency of the original LPR system while remaining reasonably competitive with NIST finalists.

The LPR cryptosystem is not randomness-recovering. In a cyclotomic ring $R = \mathbb{Z}[X]/(X^n + 1)$, it involves a public key containing a pair $(a, b = a \cdot s + e)$, where $a \in R/(qR)$ is uniform, $s \in R$ is the secret key and e is a noise. To encrypt $m \in R/(pR)$ (for some moduli $p < q$), the sender chooses small-norm randomness $r, e_1, e_2 \in R$ and computes $(c_1, c_2) = (a \cdot r + e_1, b \cdot r + e_2 + m \cdot \lfloor q/p \rceil)$, so that the receiver can obtain $c_2 - c_1 \cdot s \bmod q = m \cdot \lfloor q/p \rceil + \mathsf{small}$. While m is then computable, there is no way to recover (r, e_1, e_2) from the "decryption error" term small. To address this problem, a folklore solution is to introduce distinct powers of p. Suppose we want to build a randomness-recovering encryption of 0 (which is sufficient to build a KEM). The sender can then compute $(c_1, c_2) = (p^2 \cdot a \cdot r + p \cdot e_1, p^2 \cdot b \cdot r + e_2)$, which allows the receiver to obtain $\mu = c_2 - c_1 \cdot s \bmod q = p^2 er - pe_1 s + e_2$. Since the right-hand-side member is small, the receiver can efficiently decode $(r, e_1, e_2) \in R^3$ from μ.

Unfortunately, the latter construction is not rigid. Suppose that an adversary can somehow compute a non-trivial pair $(u, u \cdot s) \in R^2$ given $(a, a \cdot s + e)$. It can then faithfully compute $(c_1, c_2) = (p^2 \cdot a \cdot r + p \cdot e_1, p^2 \cdot b \cdot r + e_2)$ and turn it into $(c_1', c_2') = (c_1 + u, c_2 + u \cdot s)$, which yields a "decryption collision" $\mu = c_2 - c_1 \cdot s = c_2' - c_1' \cdot s$. Besides, as shown in [32], computing a pair $(u, u \cdot s)$ (for an arbitrary, possibly non-invertible $u \neq 0$) given $(a, a \cdot s + e)$ can only be hard in rings $R/(qR) \cong \mathbb{Z}_q[X]/(\Phi_1(X)) \times \cdots \times \mathbb{Z}_q[X]/(\Phi_t(X))$ that have no small-degree factors, which rules out NTT-friendly rings. Even for rings where $\Phi(X) = X^n + 1$ splits into degree-$n/2$ factors, the problem (called SIP-LWE in [32]) is non-standard and its hardness is not known to be implied by RLWE.[2] Here, we take a different approach since we aim at rigidity without relying on stronger assumptions than RLWE and without forbidding fully splitting rings.

We modify the original LPR system in the following way. The public key contains a random $a \in R/(qR)$ and a pseudorandom $b \in R/(qR)$, which is now of the form $b = p \cdot (a \cdot s + e)$, for small secrets $s, e \in R$ and a public integer p such that $\|e\|_\infty < p/2$. We also require b to be invertible over $R/(qR)$, so that the key generation phase must be repeated with new candidates (s, e) until b is a unit. To compute an encapsulation, we sample Gaussian ring elements $r, e_1, e_2 \in R$ and compute $c_{kem} = (c_1, c_2) = (a \cdot r + e_1, b \cdot r + e_2)$, where $K = H(r, e_1, e_2)$ is the encapsulated key. Decapsulation is performed by using $s \in R$ to compute $\mu = c_2 - p \cdot c_1 \cdot s \bmod q$, which is a small-norm element $\mu = e_2 + p \cdot \mathsf{small} \in R$ that reveals $e_2 = \mu \bmod p$. Given e_2, the receiver then obtains $r = (c_2 - e_2) \cdot b^{-1}$ and $e_1 = c_1 - a \cdot r$, and checks the smallness of (r, e_1, e_2). The decapsulation phase is natively *rigid* (without re-encryption) as it outputs small-norm $(r, e_1, e_2) \in R^3$ *if and only if* $(c_1, c_2) = (a \cdot r + e_1, b \cdot r + e_2)$.

Our hybrid encryption scheme builds on a variant of the above KEM with *explicit rejection*, where the decapsulation phase returns an error symbol $\bot$ on input of an invalid c_{kem}. It thus departs from NIST finalists that all rely on KEMs with implicit rejection, where the decapsulation algorithm never outputs $\bot$, but rather handles invalid encapsulations c_{kem} by outputting a random key $K' \leftarrow H(z, c_{kem})$ derived from an independent long-term secret z.[3] While our scheme could have relied on implicit rejection in a similar way, we chose to avoid additional computations involving extra secret key components z. The reason is that, if we were to introduce additional key material z, it should also be DPA-protected with possibly heavy side-channel countermeasures.

When it comes to proving security in the QROM, the use of an explicit-rejection KEM introduces some difficulty as it is not clear how to deal with invalid ciphertexts. While the classical ROM allows inspecting all random oracle queries and determining if one of them explains a given ciphertext, we cannot

---

[2] D'Anvers *et al.* [32] defined a homogeneous variant of SIP-LWE which is unconditionally hard, even in fully splitting rings. Still, relying on this variant incurs a partial re-encryption to enforce the equality $c_2 = c_2'$.

[3] When the hybrid KEM-DEM framework is instantiated with an implicit rejection KEM, invalid ciphertexts are usually rejected during the symmetric decryption step as decrypting c_{sym} with a random key K' yields $\bot$.

use this approach in the QROM because RO-queries are made on superpositions of inputs. Our solution is to use an implicit-rejection KEM only in the security proof. In a sequence of games, we first modify the decryption oracle so as to make the rejection process implicit. Then, we argue that, as long as the DEM component is realized using an authenticated symmetric encryption scheme, the modified decryption oracle is indistinguishable from the real one. After having modified the decryption oracle, we can adapt ideas from Saito *et al.* [73] in order to tightly relate the security of the hybrid scheme to the RLWE assumption.

As mentioned earlier, avoiding re-encryption does not suffice to ensure side-channel resistance. As a first countermeasure, we modify the decapsulation step and add a dummy ciphertext $(c_1', c_2') = (a \cdot r' + e_1', b \cdot r' + e_2')$ for fresh receiver-chosen randomness r', e_1', e_2' to (c_1, c_2) before proceeding with the decapsulation of $(\bar{c}_1, \bar{c}_2) = (c_1 + c_1', c_2 + c_2')$. This simple trick prevents the adversary from controlling the ring elements that multiply the secret key s at the only step where it is involved. We even show in Section 5 how this computation can be protected against DPA with minimum overheads by combination the masking countermeasure and a LWPR assumption. Additionally, the choice of (c_1', c_2') as an honestly generated encapsulation allows continuing the decryption process as if $(\bar{c}_1, \bar{c}_2)$ was the ciphertext computed from the (still) small-norm coins $(\bar{r}, \bar{e}_1, \bar{e}_2) = (r + r', e_1 + e_1', e_2 + e_2')$. That is, we do not have to remove the noise terms as we can retrieve $\bar{r}$, $\bar{e}_1$ and $\bar{e}_2$ and test their smallness. Since r', e_1' and e_2' do have small norm, if the decryption succeeds until this step, then r, e_1 and e_2 must be small as well (with a small constant slackness factor 3). Therefore, the dummy ciphertext/KEM makes it possible to eliminate an exponential amount of invalid ciphertexts without having ever tried to re-compute the correct (r, e_1, e_2). In case of an early rejection, and because the secret key s is now protected with a hidden and pseudorandom $(\bar{c}_1, \bar{c}_2)$, the leakage only provides limited information related to the ephemeral values in (r', e_1', e_2') which were sampled independently of the adversary's view. If no rejection occurs, (r', e_1', e_2') has components of (small but) sufficiently large norm to hide (at least most of the bits of) (r, e_1, e_2) if the adversary gets the full leakage of $(\bar{r}, \bar{e}_1, \bar{e}_2)$. At that time, we can safely recover (r, e_1, e_2) and check their norm (to eliminate the slackness) for technical reasons. This computation can only be repeated through many decryption queries on fixed inputs, and therefore only require SPA security (with averaging), which is cheaper to ensure than DPA security. As for the DEM, the general solutions outlined in [15] are a natural option. But we show an even cheaper one that leverages a key-homomorphic MAC.

RELATED WORK. As a KEM variant of LPR, POLKA bears high-level resemblance with NewHope [4] and Peikert's KEM [65] in that its encapsulation procedure computes $(c_1, c_2) = (a \cdot r + e_1, b \cdot r + e_2)$. The main differences are that its public key is of the form $b = p \cdot (as + e)$ (instead of $b = as + e$ in [4,65]) and c_2 is sent along with c_1, whereas [4,65] send a compressed version of c_2 from which a key is extractable using s (and a "reconciliation" technique from [34,65]). Also, decapsulation first recovers e_2 (instead of a noisy version of ars in [4,65]) in order to enforce rigidity and randomness recovery. In [4,65], the compression of

7

c_2 makes it hard to combine both properties without using Fujisaki-Okamoto.

The NTRU cryptosystem [47] can be made rigid, as observed in [27, Section 2.3] where a re-encryption-free variant of an NTRU-KEM due to Saito *et al.* [73] was proposed. While the latter construction admits highly efficient instantiations (as well as a tight proof in the QROM [73]), these rely on a less standard assumption than RLWE. In a polynomial ring $R = \mathbb{Z}[X]/(\Phi(x))$, NTRU involves a public key of the form $h = p \cdot g/f \in R/(qR)$, for some moduli $p < q$ and small-norm polynomials $f, g \in R$. Its most efficient variants rely on the assumption (often called "NTRU assumption") that $h = p \cdot g/f$ is indistinguishable from a random invertible element of $R/(qR)$, even when f and g are sampled from a narrow distribution over R. In fact, some NTRU-based NIST candidates even sample f and g from a distribution of polynomials over $\mathbb{Z}[X]/(\Phi(X))$ with ternary coefficients in $\{-1, 0, 1\}$. Other works [56] sample them from a wider, Gaussian distribution with standard deviation $\sigma < q^{1/2}$.

Despite recent progress [66], the NTRU problem is still less understood than RLWE. It is in fact asymptotically easier for small-magnitude f, g: Kirchner and Fouque [54] gave a heuristic slightly sub-exponential algorithm in the setting where $q < \deg(\Phi)$ and $\|f\|_\infty, \|g\|_\infty = O(1)$. There is even a parameter regime [2] where NTRU is easy and RLWE remains hard when q is larger than the secrets by a sub-exponential factor in $n = \deg(\Phi)$.

Stehlé and Steinfeld [75] showed that, when f and g are sampled from a discrete Gaussian with standard deviation $\sigma > \mathsf{poly}(n) \cdot q^{1/2}$, the NTRU problem is information-theoretically hard as the distribution of g/f is statistically uniform over the units of $R/(qR)$. Using this property, Stehlé and Steinfeld gave an IND-CPA-secure variant of NTRU that solely relies on the RLWE assumption. On the downside, sampling f and g from such a wide Gaussian requires a significantly larger modulus q to ensure correctness. We show in Supplementary Material A that the re-encryption-free hybrid NTRU candidate of [27, Section 2.3] can be made side-channel-resistant by applying the same countermeasures as in POLKA. Nevertheless, if we tune it so as to only rely on the RLWE assumption by applying the result of [75], it is significantly less efficient than POLKA. Asymptotically, it requires a modulus $q = \widetilde{O}(n^7)$ while POLKA can make do with $q = O(n^3)$ when the noise is sampled from a Gaussian of standard deviation $\alpha q = \Omega(\sqrt{n})$.

Recently, Duman *et al.* [37] gave a rigorous analysis of the security of optimized NTRU-based KEMs. While some of their constructions are rigid, no discussion on side-channel resistance was given in [37]. In fact, only their third construction seems compatible with our dummy-ciphertext technique. If we similarly set the parameters of POLKA according to the concrete hardness of RLWE against known attacks [3], our scheme competes fairly well with the most efficient NTRU candidates of [37]. In a fully optimized instantiation, ciphertexts and keys fit within 3Kb. Moreover, in contrast with [37], we can avoid the use of assumptions that posit the pseudorandomness of small-polynomial ratios.

D'Anvers *et al.* [32] introduced a technique that obviates the need for re-encryption when Fujisaki-Okamoto is applied to specific KEMs based on Module LWE [55]. They somehow relaxed the rigidity property of [16] by showing that,

as long as the SIP-LWE problem is hard, a cheaper test on error terms allows dispensing with the need to re-compute the entire ciphertext. Although their error-term-checking technique offers noticeable computational savings in some NIST candidates [4,46,57], it does not (and was not claimed to) protect against side-channel leakage, as observed in [77]. POLKA departs from their approach as it relies neither on the FO transform, nor on SIP-LWE.

The ideas we use regarding polynomial operations build on the early observation that they have good features for efficient masking: see [71,72] for examples of key randomization (which we leverage for the parts of POLKA that require DPA security) and [70] for an example of message randomization (which we leverage for the parts of POLKA that require SPA security), respectively. As also observed in these previous works, it is generally the additional operations besides the polynomial multiplication that make post-quantum schemes challenging to protect against leakage, and it is precisely this issue that POLKA aims to contribute to, by taking advantage of the concept of leveled implementations.

Eventually, the interest of an explicit rejection mechanism was also mentioned in the recent independent work of Hövelmanns et al. [50].

3 Background

3.1 Lattices and Discrete Gaussian Distributions

An n-dimensional lattice $\Lambda \subseteq \mathbb{R}^n$ is the set $\Lambda = \{\sum_{i=1}^n z_i \cdot \mathbf{b}_i \mid \mathbf{z} \in \mathbb{Z}^n\}$ of all integer linear combinations of a set of linearly independent basis vectors $\mathbf{B} = \{\mathbf{b}_1, \ldots, \mathbf{b}_n\} \subseteq \mathbb{R}^n$. Let $\mathbf{\Sigma} \in \mathbb{R}^{n \times n}$ be a symmetric positive definite matrix, and $\mathbf{c} \in \mathbb{R}^n$. The n-dimensional Gaussian function on $\mathbb{R}^n$ is defined as $\rho_{\mathbf{\Sigma}, \mathbf{c}}(\mathbf{x}) = \exp(-\pi(\mathbf{x} - \mathbf{c})^\top \mathbf{\Sigma}^{-1}(\mathbf{x} - \mathbf{c}))$. In the special case where $\mathbf{\Sigma} = \sigma^2 \cdot \mathbf{I}_n$ and $\mathbf{c} = \mathbf{0}$, we denote it by ρ_σ. For any lattice $\Lambda \subset \mathbb{R}^n$, the discrete Gaussian distribution $D_{\Lambda, \mathbf{\Sigma}, \mathbf{c}}$ has probability mass $\Pr_{X \sim D_{\Lambda, \mathbf{\Sigma}, \mathbf{c}}}[X = \mathbf{x}] = \frac{\rho_{\mathbf{\Sigma}, \mathbf{c}}(\mathbf{x})}{\rho_{\mathbf{\Sigma}, \mathbf{c}}(\Lambda)}$ for any $\mathbf{x} \in \Lambda$. When $\mathbf{c} = \mathbf{0}$ and $\mathbf{\Sigma} = \sigma^2 \cdot \mathbf{I}_n$ we denote it by $D_{\Lambda, \sigma}$.

Lemma 1 ([62, Lemma 4.4]). *For $\sigma = \omega(\sqrt{\log n})$ there is a negligible function $\varepsilon = \varepsilon(n)$ such that $\Pr_{\mathbf{x} \sim D_{\mathbb{Z}^n, \sigma}}[\|\mathbf{x}\| > \sigma\sqrt{n}] \leq \frac{1+\varepsilon}{1-\varepsilon} \cdot 2^{-n}$.*

3.2 Rings and Ideal Lattices.

Let n a power of 2 and define the rings $R = \mathbb{Z}[X]/(X^n+1)$ and $R_q = R/qR$. Each element of R is a $(n-1)$-degree polynomial in $\mathbb{Z}[X]$ and can be interpreted as an element of $\mathbb{Z}[X]$ via the natural coefficient embedding that maps the polynomial $a = \sum_{i=0}^{n-1} a_i X^i \in R$ to $(a_0, a_1, \ldots, a_{n-1}) \in \mathbb{Z}^n$. An element of R_q can similarly be viewed as a degree-$(n-1)$ polynomial over $\mathbb{Z}_q[X]$ and represented as an n-dimensional vector with coefficients in the range $\{-(q-1)/2, \ldots, (q-1)/2\}$.

The Euclidean and infinity norms of an element of $a \in R$ are defined by viewing elements of R as elements of $\mathbb{Z}^n$ via the coefficient embedding.

The ring R can also be identified as the subring of anti-circulant matrices in

$\mathbb{Z}^{n \times n}$ by viewing each $a \in R$ as a linear transformation $r \to a \cdot r$. This implies that, for any $a, b \in R$, $\|a \cdot b\|_\infty \le \|a\| \cdot \|b\|$ by the Cauchy-Schwartz inequality.

As in [52], for any lattice Λ, $D_{\Lambda,\sigma}^{\mathsf{coeff}}$ denotes the distribution of a ring element $a = \sum_{i=0}^{n-1} a_i X^i \in R$ of which the coefficient vector $(a_0, \ldots, a_{n-1})^\top \in \mathbb{Z}^n$ is sampled from the discrete Gaussian distribution $D_{\Lambda,\sigma}$.

We now recall the ring variant of the Learning-With-Errors assumption [69]. The ring LWE (RLWE) problem is to distinguish between a polynomial number of pairs of the form $(a_i, a_i \cdot s + e_i)$, where $a_i \sim U(R_q)$ and $s, e_i \in R$ are sampled from some distribution χ of bounded-magnitude ring elements, and random pairs $(a_i, b_i) \sim U(R_q^2)$. In Definition 1, the number of samples k is made explicit.

Definition 1. *Let $\lambda \in \mathbb{N}$ a security parameter. Let positive integers $n = n(\lambda)$, $k = k(\lambda)$, and a prime $q = q(n) > 2$. Let an error distribution $\chi = \chi(n)$ over R. The $\mathsf{RLWE}_{n,k,q,\chi}$ assumption says that the following distance is a negligible function for any PPT algorithm $\mathcal{A}$,*

$$\mathbf{Adv}_{n,k,q,\chi}^{\mathcal{A},\mathsf{RLWE}}(\lambda) := \big| \Pr[\mathcal{A}(1^\lambda, \{(a_i, v_i)\}_{i=1}^k) = 1]$$
$$- \Pr[\mathcal{A}(1^\lambda, \{(a_i, a_i s + e_i)\}_{i=1}^k) = 1]\big|,$$

where $a_1, \ldots, a_i, v_1, \ldots, v_k \hookleftarrow U(R_q)$, $s \hookleftarrow \chi$, $e_1, \ldots, e_i \hookleftarrow \chi$.

For suitable parameters, the RLWE assumption is implied by the hardness of worst-case instances of the approximate shortest vector problem in ideal lattices.

Lemma 2 ([58]). *Let n a power of 2. Let $\Phi_m(X) = X^n + 1$ the m-th cyclotomic polynomial where $m = 2n$, and $R = \mathbb{Z}[X]/(\Phi_m(X))$. Let $q = 1 \bmod 2n$. Let also $r = \omega(\sqrt{\log n})$ Then, there is a randomized reduction from $2^{\omega(\log n)} \cdot (q/r)$-approximate R-SVP to $\mathsf{RLWE}_{n,\mathsf{poly}(n),q,\chi}$ where $\chi = D_{\mathbb{Z}^n,r}^{\mathsf{coeff}}$.*

4 POLKA: Rationale and Specifications

Our starting point is a variant of the LPR cryptosystem [58], which builds on a rigid randomness-recovering KEM. As in [58], the public key contains a random ring element $a \in R_q$ and a pseudorandom $b \in R_q$. Here, b is of the form $b = p \cdot (a \cdot s + e)$ (instead of $b = a \cdot s + e$ as in [58]), for secret $s, e \in R$ sampled from the noise distribution and where p is an integer such that $\|e\|_\infty < p$. Another difference with [58] is that decryption requires b to be invertible over R_q.

The encryptor samples ring elements $r, e_1, e_2 \in R$ from a Gaussian distribution and uses them to derive a symmetric key $K = H(r, e_1, e_2)$. The latter is then encapsulated by computing a pair $(c_1, c_2) = (a \cdot r + e_1, b \cdot r + e_2)$. The decryption algorithm uses $s \in R$ to compute $\mu = c_2 - p \cdot c_1 \cdot s \in R_q$, which is a small-norm ring element $\mu = e_2 + p \cdot (er - e_1 s) \in R$. This allows recovering $e_2 = \mu \bmod p$, which in turn reveals $r = (c_2 - e_2) \cdot b^{-1} \in R_q$ and $e_1 = c_1 - a \cdot r \in R_q$. After having checked the smallness of (r, e_1, e_2), the decryption procedure obtains $K = H(r, e_1, e_2)$. The scheme provides the *rigidity* property of [16] as the decryptor obtains $(r, e_1, e_2) \in R^3$ such that $\|r\|, \|e_1\|, \|e_2\| \le B$, for some norm

bound B, if and only if $(c_1, c_2) = (a \cdot r + e_1, b \cdot r + e_2)$. This ensures that no re-encryption is necessary to check the validity of the input pair (c_1, c_2).

In this section, our hybrid encryption scheme builds on a KEM with *explicit rejection* (in the terminology of [48,67]), meaning that invalid encapsulations are rejected as soon as they are noticed in the decryption algorithm. In the security proof, we will switch to an implicit rejection mechanism (as defined [67, Section 5.3]), where the decapsulation algorithm outputs a random key on input of an invalid encapsulation. The rejection of malformed encapsulations is then deferred to the symmetric decryption step.

4.1 The Scheme With an Additive Mask

We now describe a version of the scheme that has good features for side-channel resistant implementation, where the decryption algorithm first adds a "dummy ciphertext" to (c_1, c_2) before proceeding with the actual decryption.

Keygen(1^λ): Given a security parameter $\lambda \in \mathbb{N}$,

1. Choose a dimension $n \in \mathbb{N}$, a prime modulus $q = 1 \bmod 2n$. Let the rings $R = \mathbb{Z}[X]/(X^n + 1)$ and $R_q = R/(qR)$ such that $\Phi(X) = X^n + 1$ splits into linear factors over R_q. Let $R_q^\times$ the set of units in R_q.
2. Choose a noise parameter $\alpha \in (0, 1)$, and let a norm bound $B = \alpha q \sqrt{n}$. Choose an integer $p \in \mathbb{N}$ such that $4B < p < \frac{q}{8(B^2+1)}$.
3. Sample $a \hookleftarrow U(R_q)$ and $s, e \hookleftarrow D_{\mathbb{Z}^n, \alpha q}^{\mathsf{coeff}}$ and compute $b = p \cdot (a \cdot s + e)$. If $b \notin R_q^\times$, restart step 3.
4. Choose an authenticated symmetric encryption scheme $\Pi^{sym} = (\mathsf{K}, \mathsf{E}, \mathsf{D})$ with key length $\kappa \in \mathsf{poly}(\lambda)$ and message space $\{0,1\}^{\ell_m}$.
5. Let a domain $D_E := \{(r, e_1, e_2) \in R^3 : \|r\|, \|e_1\|, \|e_2\| \le B\}$. Choose a hash function $H : D_E \to \{0,1\}^\kappa$ modeled as a random oracle.

Return the key pair (PK, SK) where

$$PK := \big(n, \ q, \ p, \ \alpha, \ a \in R_q, \ b \in R_q^\times, \ \Pi^{sym}, \ H, \ B\big) \text{ and } SK := s \in R.$$

Encrypt(PK, M): Given a public key PK and a message $M \in \{0,1\}^{\ell_m}$:

1. Sample $r, e_1, e_2 \hookleftarrow D_{\mathbb{Z}^n, \alpha q}^{\mathsf{coeff}}$ and compute

$$c_1 = a \cdot r + e_1 \in R_q, \qquad c_2 = b \cdot r + e_2 \in R_q$$

together with $K = H(r, e_1, e_2) \in \{0,1\}^\kappa$.
2. Compute $c_0 = \mathsf{E}_K(M)$.

Output the ciphertext $C = (c_0, c_1, c_2)$.

Decrypt(SK, C): Given $SK = s \in R$ and $C = (c_0, c_1, c_2)$, do the following:

1. Sample $r', e_1', e_2' \hookleftarrow D_{\mathbb{Z}^n, \alpha q}^{\mathsf{coeff}}$ and return $\bot$ if $\|r'\| > B$, or $\|e_1'\| > B$, or $\|e_2'\| > B$. Otherwise, compute $c_1' = a \cdot r' + e_1'$ and $c_2' = b \cdot r' + e_2'$.

11

2. Compute $\bar{c}_1 = c_1 + c_1'$ and $\bar{c}_2 = c_2 + c_2'$.
3. Compute $\bar{\mu} = \bar{c}_2 - p \cdot \bar{c}_1 \cdot s$ over R_q.
4. Compute $\bar{e}_2 = \bar{\mu} \bmod p$. If $\|\bar{e}_2\| > 2B$, return $\perp$.
5. Compute $\bar{r} = (\bar{c}_2 - \bar{e}_2) \cdot b^{-1} \in R_q$. If $\|\bar{r}\| > 2B$, return $\perp$.
6. Compute $\bar{e}_1 = \bar{c}_1 - a \cdot \bar{r} \in R_q$. If $\|\bar{e}_1\| > 2B$, return $\perp$.
7. Compute $r = \bar{r} - r'$, $e_1 = \bar{e}_1 - e_1'$ and $e_2 = \bar{e}_2 - e_2'$. If $\|r\| > B$, or $\|e_1\| > B$, or $\|e_2\| > B$, then return $\perp$.
8. Compute $K = H(r, e_1, e_2) \in \{0,1\}^\kappa$ and return

$$M = \mathsf{D}_K(c_0) \in \{0,1\}^{\ell_m} \cup \{\perp\}.$$

The use of fully splitting rings may require multiple attempts to find an invertible $b \in R_q^\times$ as step 3 of Keygen. The proof of Lemma 7 shows that, unless the RLWE assumption is false, a suitable b can be found after at most $\lceil \lambda / \log n \rceil$ iterations, except with negligible probability $2^{-\lambda}$. In practice, a small number of attempts suffices since a random ring element is invertible with probability $1 - n/q$, which is larger than $1 - 1/n$ with our choice of parameters.

CORRECTNESS. Let $\bar{r} = r + r'$, $\bar{e}_1 = e_1 + e_1'$ and $\bar{e}_2 = e_2 + e_2'$ over R. At step 2, Decrypt computes $\bar{c}_1 = a \cdot \bar{r} + \bar{e}_1$, $\bar{c}_2 = b \cdot \bar{r} + \bar{e}_2$ over R_q, where $\|\bar{r}\|, \|\bar{e}_1\|, \|\bar{e}_2\| \leq 2B$ with probability $1 - 2^{-\Omega(n)}$ over the randomness of Encrypt and Decrypt (by Lemma 1). At step 3, the decryptor obtains

$$\begin{aligned}
\bar{\mu} &= \bar{c}_2 - p \cdot \bar{c}_1 \cdot s \bmod q \\
&= (b \cdot \bar{r} + \bar{e}_2) - p \cdot (a \cdot \bar{r} + \bar{e}_1) \cdot s \bmod q \\
&= p \cdot (as + e) \cdot \bar{r} + \bar{e}_2 - p \cdot a\bar{r}s - p \cdot \bar{e}_1 s \bmod q \\
&= \bar{e}_2 + p \cdot e\bar{r} - p \cdot \bar{e}_1 s,
\end{aligned}$$

where the last equality holds over R with overwhelming probability over the randomness of Keygen, Encrypt and Decrypt. Indeed, Lemma 1 implies that $\|s\|, \|e\| \leq \alpha q \sqrt{n}$ with probability $1 - 2^{-\Omega(n)}$ over the randomness of Keygen. With probability $1 - 2^{-\Omega(n)}$ over the randomness of Encrypt and Decrypt, we also have $\|\bar{r}\|, \|\bar{e}_1\|, \|\bar{e}_2\| \leq 2\alpha q \sqrt{n}$. Then, the Cauchy-Schwartz inequality implies

$$\|\bar{e}_2 + p \cdot e\bar{r} - p \cdot \bar{e}_1 s\|_\infty < 2\alpha q \sqrt{n} + 4p \cdot (\alpha q)^2 n \qquad (1)$$
$$< 4p(B^2 + 1) < q/2.$$

Since $p/2 > 2\alpha q \sqrt{n}$, step 4 recovers $\bar{e}_2$ with overwhelming probability. Since $b \in R_q^\times$, Decrypt obtains $\bar{r}$ at step 5 and $\bar{e}_1$ at step 6. Therefore, it also recovers (r, e_1, e_2) at step 7 and the correct symmetric key $K = H(r, e_1, e_2)$ at step 8. Correctness thus follows from the correctness of Π^{sym}.

Remark 1. We note that correctness is guaranteed whenever $\|s\|, \|e\| \leq B$ and $\|\bar{r}\|, \|\bar{e}_1\|, \|\bar{e}_2\| \leq 2B$, as it is a sufficient condition to have inequalities (1). This will be used in the security proof.

ON DECRYPTION FAILURES. Due to the rigidity and randomness recovery properties of the scheme, the probability of decryption failure does not depend on the

specific secret key s in use as long as $\|s\| \leq B$. If $(r, e_1, e_2) \in D_E$ and $\|s\| \leq B$, we always have $\|\bar{\mu}\|_\infty \leq q/2$, where $\bar{\mu} = \bar{c}_2 - p \cdot \bar{c}_1 \cdot s \bmod q = p \cdot (e \cdot \bar{r} - \bar{e}_1 \cdot s) + \bar{e}_2$ unless the ciphertext is rejected at step 1 (which does not depend on s). If $(r, e_1, e_2) \notin D_E$, then either: (i) We still have $\|\bar{\mu}\|_\infty \leq q/2$ and Decrypt obtains (r, e_1, e_2), which are necessarily rejected; or (ii) $\|\bar{\mu}\|_\infty > q/2$ but the extracted $(r^\dagger, e_1^\dagger, e_2^\dagger)$ cannot land in D_E since, otherwise, the rigidity property would imply $(c_1, c_2) = (a \cdot r^\dagger + e_1^\dagger, b \cdot r^\dagger + e_2^\dagger)$, in which case we would have $\|\bar{\mu}\|_\infty \leq q/2$ unless the ciphertext is rejected at step 1. Hence, if Decrypt does not return $\bot$ at step 1, it computes $(r, e_1, e_2) \in D_E$ if and only if $(c_1, c_2) = (a \cdot r + e_1, b \cdot r + e_2)$ no matter which s of norm $\|s\| \leq B$ is used at step 2.

In contrast, when $m = 0$ is encrypted in the LPR cryptosystem, we have $\mu = c_2 - c_1 \cdot s \bmod q = e \cdot r - e_1 \cdot s + e_2$. An adversary can then fix small (r, e_2) and play with many e_1's until it triggers a decryption failure when $\|\mu\|_\infty > q/2$. The probability that this happens depends on the secret s (as a different s' may not cause rejection for a fixed (r, e_1, e_2)). In KYBER and SABER, the FO transform allows restricting the adversary's control over e_1 (which is derived from a random message m using a random oracle and re-computed for verification upon decryption) so as to make such attacks impractical. The FO transform is thus crucial to offer a sufficient security margin against attacks like [30, 33].

4.2   Black-Box Security Analysis

Our security proof uses ideas from Saito *et al.* [73, Section 4] to prove (tight) security in the QROM. Their approach exploits the implicit rejection mechanism of their KEM. Namely, when the incoming encapsulation (c_1, c_2) is found invalid upon decryption in [73], the decapsulated symmetric key K is replaced by a random-looking $K = H'(u, (c_1, c_2))$, where u is a random string included in the secret key and H' is an independent random oracle.

Here, in order to simplify the analysis of side-channel leakages in the real scheme, it is desirable to minimize the amount of secret key operations in the decryption algorithm and the amount of key material to protect against leakage. Therefore we refrain from introducing an additional secret key component u. Instead, our security proof will first switch (in Game_2) to a modified decryption algorithm where the rejection mechanism goes implicit and the decapsulation procedure computes K as a random function of (c_1, c_2). At this point, we will be able to apply the techniques from [73].

Since the implicit/explicit decapsulation mechanisms are used as part of a hybrid encryption system, we can argue that they are indistinguishable by relying on the ciphertext integrity of the symmetric encryption scheme. This is the reason why we are considering the CCA security of the hybrid combination as a whole, rather than that of its KEM component.[4] We note that similar ideas were

---

[4] The underlying explicit rejection KEM can be proven CCA-secure secure in the ROM but we do not prove it CCA-secure in the QROM as we only consider the CCA security of the hybrid PKE scheme.

previously used in the security proofs of hybrid PKE schemes [1, 49], but usually in the opposite direction (to go from implicit rejection to explicit rejection).

For the rest, the proof in the ROM carries over to the QROM since it avoids ROM techniques that do not work in the QROM: we do not rely on the extraction of encryption randomness by inspecting the list of RO queries to answer decryption queries, which is not possible when queries are made on superpositions of inputs, and the RO is programmed identically for all queries.

Theorem 1. *If Π^{sym} is a symmetric authenticated encryption scheme, the construction of Section 4.1 provides IND-CCA security in the QROM under the* RLWE *assumption.*

Proof. The proof considers a sequence of hybrid games, which is similar to that of [73, Theorem 4.2] from Game_3 to Game_5. For each i, we denote by W_i the event that the adversary wins (i.e., $d' = d$) in Game_i. We also denote by $\mathsf{Encaps}(PK, (r, e_1, e_2))$ the deterministic algorithm that takes as inputs PK and explicit randomness $(r, e_1, e_2) \in R^3$, and outputs $(c_1, c_2) = (a \cdot r + e_1, b \cdot r + e_2)$.

Game_0: This is the real IND-CCA game. The challenger faithfully answers (quantum) random oracle queries. All (classical) decryption queries are answered by running the real decryption algorithm. Note that a decryption query triggers a random oracle query at step 8 of Decrypt. In the challenge phase, the adversary $\mathcal{A}$ outputs messages M_0, M_1 and obtains a challenge $C^\star = (c_0^\star, c_1^\star, c_2^\star)$, where $c_1^\star = a \cdot r^\star + e_1^\star$, $c_2^\star = b \cdot r^\star + e_2^\star$, with $r^\star, e_1^\star, e_2^\star \hookleftarrow D_{\mathbb{Z}^n, \alpha q}^{\mathsf{coeff}}$, and $c_0^\star = \mathsf{E}_{K^\star}(M_d)$ for some $d \hookleftarrow U(\{0, 1\})$. Eventually, $\mathcal{A}$ outputs $b' \in \{0, 1\}$ and its advantage is $\mathbf{Adv}(\mathcal{A}) := |\Pr[W_0] - 1/2|$.

Game_1: In this game, the challenger aborts and replaces $\mathcal{A}$'s output by a random bit $d' \in \{0, 1\}$ in the event that $\|s\| > B$ or $\|e\| > B$ at step 3 of Keygen. By Lemma 1, we have $|\Pr[W_1] - \Pr[W_0]| \le 2^{-\Omega(n)}$.

Game_2: We modify the decryption algorithm. Throughout the game, the challenger uses an independent random oracle $H_Q : R_q^2 \to \{0, 1\}^\kappa$ that is only accessible to $\mathcal{A}$ via decryption queries (i.e., $\mathcal{A}$ has no direct access to H_Q). This random oracle is used to run the following decryption algorithm.

$\mathsf{Decrypt}_2$: Given $SK = s$ and $C = (c_0, c_1, c_2)$, initialize a Boolean variable flag $= 0$. Then, do the following.

1. Sample $r', e_1', e_2' \hookleftarrow D_{\mathbb{Z}^n, \alpha q}^{\mathsf{coeff}}$. If $\|r'\| > B$, or $\|e_1'\| > B$, or $\|e_2'\| > B$, then set flag $= 1$ and return $\perp$.[5] Otherwise, compute $c_1' = a \cdot r' + e_1'$ and $c_2' = b \cdot r' + e_2'$.
2. Compute $\bar{c}_1 = c_1 + c_1'$ and $\bar{c}_2 = c_2 + c_2'$.
3. Compute $\bar{\mu} = \bar{c}_2 - p \cdot \bar{c}_1 \cdot s$ over R_q.
4. Compute $\bar{e}_2 = \bar{\mu} \bmod p$. If $\|\bar{e}_2\| > 2B$, set flag $= 1$.

---

[5] $\mathsf{Decrypt}_2$ still uses explicit rejection at step 1 because the secret key is not needed at this step and the goal of implicit rejection is to handle validity checks that depend on the secret key and the ciphertext.

5. Compute $\bar{r} = (\bar{c}_2 - \bar{e}_2) \cdot b^{-1} \in R_q$. If $\|\bar{r}\| > 2B$, set flag $= 1$.
6. Compute $\bar{e}_1 = \bar{c}_1 - a \cdot \bar{r} \in R_q$. If $\|\bar{e}_1\| > 2B$, set flag $= 1$.
7. Compute $r = \bar{r} - r'$, $e_1 = \bar{e}_1 - e_1'$ and $e_2 = \bar{e}_2 - e_2'$. If $\|r\| > B$, or $\|e_1\| > B$, or $\|e_2\| > B$, then set flag $= 1$.
8. If flag $= 0$, compute $K = H(r, e_1, e_2) \in \{0,1\}^\kappa$. Otherwise, compute $K = H_Q(c_1, c_2)$.
9. Compute and return $M = \mathsf{D}_K(c_0) \in \{0,1\}^{\ell_m} \cup \{\bot\}$.

Lemma 3 shows that, if the adversary can distinguish Game_2 from Game_1, we can turn it into an adversary against the ciphertext integrity of Π^{sym} (of which the definition is recalled in Supplementary Material B).

Game_3: We now simulate the random oracle[6] $H : D_E \to \{0,1\}^\kappa$ as

$$H(r, e_1, e_2) = H_Q'\big(\mathsf{Encaps}(PK, (r, e_1, e_2))\big) \qquad (2)$$

where $H_Q' : R_q^2 \to \{0,1\}^\kappa$ is another random oracle to which $\mathcal{A}$ has no direct access. At each decryption query, $\mathsf{Decrypt}_2$ consistently computes K as per (2) when flag $= 0$. In the computation of $C^\star = (c_0^\star, c_1^\star, c_2^\star)$, the symmetric key $K^\star$ is similarly obtained as $K^\star = H_Q'\big(c_1^\star, c_2^\star\big)$, where $(c_1^\star, c_2^\star) = \mathsf{Encaps}(PK, (r^\star, e_1^\star, e_2^\star))$. Lemma 4 shows that, from $\mathcal{A}$'s view, Game_3 is identical to Game_2, so that we have $\Pr[W_3] = \Pr[W_2]$.

Game_4: This game is like Game_3 except that the random oracle H is now simulated as $H(r, e_1, e_2) = H_Q\big(\mathsf{Encaps}(PK, (r, e_1, e_2))\big)$, where $H_Q : R_q^2 \to \{0,1\}^\kappa$ is the random oracle introduced in Game_2. In the computation of the challenge ciphertext $C^\star = (c_0^\star, c_1^\star, c_2^\star)$, the symmetric key $K^\star$ is similarly obtained as $K^\star = H_Q\big(c_1^\star, c_2^\star\big)$, where $(c_1^\star, c_2^\star) = \mathsf{Encaps}(PK, (r^\star, e_1^\star, e_2^\star))$, and K is computed in the same way when flag $= 0$ at step 8 of $\mathsf{Decrypt}_2$. That is, Game_4 is identical to Game_3 except that H_Q' has been replaced by H_Q in the simulation of H. Lemma 5 shows that $\Pr[W_4] = \Pr[W_3]$ as the two games are perfectly indistinguishable.

Game_5: This game is like Game_4 except that we modify the decryption oracle. At each query $C = (c_0, c_1, c_2)$, if flag $= 0$ at the end of step 1, then the decryption oracle computes $K = H_Q(c_1, c_2)$ and returns $M = \mathsf{D}_K(c_0) \in \{0,1\}^{\ell_m} \cup \{\bot\}$ (i.e., it ignores steps 2-7 of $\mathsf{Decrypt}_2$ and jumps to step 8 after having set flag $= 1$). Lemma 6 shows that $\Pr[W_5] = \Pr[W_4]$.

Game_6: We now remove the change introduced in Game_1. Namely, Game_6 is like Game_5, but we no longer replace $\mathcal{A}$'s output by a random bit if $\|s\| > B$ or $\|e\| > B$ at the end of Keygen. By Lemma 1, $|\Pr[W_6] - \Pr[W_5]| \leq 2^{-\Omega(n)}$.

In Game_6, we note that the decryption oracle does not use the secret s anymore.

Game_7: We modify the generation of PK. The challenger initially samples $a_1, \ldots, a_k \hookleftarrow U(R_q)$, $e_1, \ldots, e_k \hookleftarrow D_{\mathbb{Z}^n, \alpha q}^{\mathsf{coeff}}$, where $k = \lceil \lambda / \log n \rceil$, and computes $b_i = a_i \cdot s + e_i$ for each $i \in [k]$. If none of the obtained $\{b_i\}_{i=1}^k$ is

---

[6] We may assume that H outputs $\bot$ on input of a triple $(r, e_1, e_2) \notin D_E$. A hash function can always check domain membership before any computation.

invertible, the challenger aborts and replaces $\mathcal{B}$'s output by a random bit. Otherwise, it determines the first index $i \in [k]$ such that $b_i \in R_q^\times$ and defines the public key by setting $a = a_i$ and $b = p \cdot b_i$. Lemma 7 shows that, under the RLWE assumption, this modified key generation procedure does not affect $\mathcal{A}$'s view and we have $|\Pr[W_7] - \Pr[W_6]| \leq \mathbf{Adv}^{\mathsf{RLWE}}(\lambda) + 2^{-\lambda}$.

Game$_8$: We change again the generation of the public key. We replace the pseudorandom ring elements $\{b_i = a_i \cdot s + e_i\}_{i=1}^k$ of Game$_7$ by truly random $b_1, \ldots, b_k \hookleftarrow U(R_q)$ at the beginning of the game. Under the RLWE assumption, this change goes unnoticed and a straightforward reduction shows that $|\Pr[W_8] - \Pr[W_7]| \leq \mathbf{Adv}^{\mathsf{RLWE}}(\lambda)$. As a result, since $\gcd(p, q) = 1$, the public key is now distributed so that $a \sim U(R_q)$ and $b \sim U(R_q^\times)$.

Game$_9$: We change the generation of the challenge $C^\star = (c_0^\star, c_1^\star, c_2^\star)$. In this game, instead of computing $c_1^\star = a \cdot r^\star + e_1^\star$, $c_2^\star = b \cdot r^\star + e_2^\star$ with $r^\star, e_1^\star, e_2^\star \hookleftarrow D_{\mathbb{Z}^n, \alpha q}^{\mathsf{coeff}}$, we now sample $c_1^\star, c_2^\star \hookleftarrow U(R_q)$ uniformly. Then, we compute $c_0^\star$ as a symmetric encryption of M_d under the key $K^\star = H_Q(c_1^\star, c_2^\star)$. Lemma 8 shows that Game$_9$ is indistinguishable from Game$_8$ under the RLWE assumption.

In Game$_9$, $\mathcal{A}$ can no longer query H on short ring elements $(r^\star, e_1^\star, e_2^\star)$ that underlie $(c_1^\star, c_2^\star)$ (in which case we would have $H_Q(c_1^\star, c_2^\star) = H(r^\star, e_1^\star, e_2^\star)$). With overwhelming probability $1 - 2^{-\Omega(n)}$, there exist no $r^\star, e_1^\star, e_2^\star \in R$ of norm $\leq B$ such that $c_1^\star = a \cdot r^\star + e_1^\star$ and $c_2^\star = b \cdot r^\star + e_2^\star$. Since $\mathcal{A}$ has no direct access to $H_Q(\cdot)$, this means that $H_Q(c_1^\star, c_2^\star)$ is now independent of $\mathcal{A}$'s view.

Game$_{10}$: In this game, we modify the decryption oracle that now rejects all ciphertexts of the form $C = (c_0, c_1^\star, c_2^\star)$ with $c_0 \neq c_0^\star$ after the challenge phase. Game$_{10}$ is identical to Game$_9$ until the event E_{10} that $\mathcal{A}$ queries the decryption of a ciphertext $C = (c_0, c_1^\star, c_2^\star)$ that would not have been rejected in Game$_9$. Since $c_0^\star = \mathsf{E}_{K^\star}(M_d)$ is encrypted under a random key $K^\star = H_Q(c_1^\star, c_2^\star)$ that is independent of $\mathcal{A}$'s view, E_{10} would imply an attack against the ciphertext integrity of Π^{sym} (as defined in Supplementary Material B)). We have $|\Pr[W_{10}] - \Pr[W_9]| \leq \Pr[E_{10}] \leq 2^{-\Omega(n)} + \mathbf{Adv}^{\mathsf{AE-INT}}(\lambda)$, where Q is the number of decryption queries.

In Game$_{10}$, the challenge $C^\star = (c_0^\star, c_1^\star, c_2^\star)$ is obtained by encrypting $c_0^\star$ under a random key $K^\star$ which is never used anywhere but in the computation of $c_0^\star = \mathsf{E}_{K^\star}(M_d)$. At this point, the adversary is essentially an adversary against the indistinguishability (under passive attacks) of the authenticated encryption scheme Π^{sym}. We have $|\Pr[W_{10}] - 1/2| \leq \mathbf{Adv}^{\mathsf{AE-IND}}(\lambda)$.

Putting the above altogether, we can bound the advantage of an IND-CCA adversary as

$$\mathbf{Adv}^{\mathrm{cca}}(\mathcal{A}) \leq \frac{3}{1 - 2^{-\lambda}} \cdot \mathbf{Adv}^{\mathcal{B}, \mathsf{RLWE}}_{n, \lceil \lambda / \log n \rceil, q, \chi}(\lambda) + Q(Q+1) \cdot \mathbf{Adv}^{\mathsf{AE-INT}}(\lambda) \qquad (3)$$

$$+ \mathbf{Adv}^{\mathsf{AE-IND}}(\lambda) + \frac{1}{2^{\Omega(n)}},$$

where Q is the number of decryption queries. $\qquad \square$

Lemma 3. Game_2 *is indistinguishable from* Game_1 *as long as the authenticated encryption scheme* Π^{sym} *provides ciphertext integrity. Concretely, we have the inequality* $|\Pr[W_2] - \Pr[W_1]| \leq \frac{Q \cdot (Q+1)}{2} \cdot \mathbf{Adv}^{\mathsf{AE\text{-}INT}}(\lambda)$.

Proof. The distinguishing advantage of $\mathcal{A}$ between the two games can be bounded by the difference between the probabilities that a ciphertext gets rejected in Game_2 and Game_1.

The difference between the two decryption oracles is that, when an invalid pair (c_1, c_2) is detected at steps 4-7, $\mathsf{Decrypt}_2$ sets $\mathsf{flag} = 1$ and keeps going when the decryption oracle of Game_1 would stop and return $\bot$.

In Game_2, let us assume that $\mathcal{A}$ queries a ciphertext (c_0, c_1, c_2) for which $\mathsf{Decrypt}_2$ can be distinguished from $\mathsf{Decrypt}$ because, after step 1, it returns $\bot$ with a significantly different probability. This means that, in Game_2, we have $\mathsf{flag} = 1$ at step 8, so that $\mathsf{Decrypt}_2$ computes $K = H_Q(c_1, c_2)$. Since $\mathcal{A}$ has no direct access to H_Q, the key K is independent of $\mathcal{A}$'s view and uniformly random over $\{0, 1\}^\kappa$. Hence, if $\mathsf{Decrypt}_2$ does not return $\bot$ at step 9, it implies $\mathsf{D}_K(c_0) \neq \bot$, meaning that $\mathcal{A}$ managed to forge a valid ciphertext for a completely random key K. Concretely, the distinguishing advantage of $\mathcal{A}$ between the two games can be bounded by $\frac{Q \cdot (Q+1)}{2} \cdot \mathbf{Adv}^{\mathsf{AE\text{-}INT}}(\lambda)$, with Q the number of decryption queries and $\mathbf{Adv}^{\mathsf{AE\text{-}INT}}(\lambda)$ the reduction's advantage against the ciphertext integrity of the authenticated encryption scheme Π^{sym}. To see this, we consider a sub-sequence of hybrid games that bridges between Game_1 and Game_2.

$\mathsf{Game}_{1.i}$ $(0 \leq i \leq Q)$: In this game, the challenger answers the first i decryption queries by running $\mathsf{Decrypt}_2$. In the last $Q - i$ decryption queries, it simulates the decryption oracle by running $\mathsf{Decrypt}$ as in Game_1.

$\mathsf{Game}_{1.0}$ is thus identical to Game_1 while $\mathsf{Game}_{1.Q}$ corresponds to Game_2. We show that the distinguishing advantage between $\mathsf{Game}_{1.i}$ and $\mathsf{Game}_{1.(i-1)}$ is at most $i \cdot \mathbf{Adv}^{\mathsf{AE\text{-}INT}}(\lambda)$.

Suppose that $\mathcal{A}$ can distinguish between $\mathsf{Game}_{1.i}$ and $\mathsf{Game}_{1.(i-1)}$ with advantage $\varepsilon = |\Pr[W_{1.i}] - \Pr[W_{1.(i-1)}]|$. This can only happen if the i-th decryption query gets rejected with significantly different probabilities in the two games. Then, we build an adversary $\mathcal{B}^{\mathsf{AE}}$ with advantage ε/i against the ciphertext integrity (as defined in Supplementary Material B) of Π^{sym}. It first guesses an indexes $j \hookleftarrow U([1, i])$ that will be correct with probability $1/i$. Namely, with probability $1/i$, $j \in [1, i]$ will be the index of the first decryption query of the form $(\cdot, c_1^\diamond, c_2^\diamond)$, where $(c_0^\diamond, c_1^\diamond, c_2^\diamond)$ denotes the i-th decryption query.

During the IND-CCA game, $\mathcal{B}^{\mathsf{AE}}$ faithfully answers all decryption queries by simulating $H_Q(\cdot, \cdot)$ on its own until the j-th query. At the j-th query $(c_0, c_1^\diamond, c_2^\diamond)$, it implicitly defines $H_Q(c_1^\diamond, c_2^\diamond)$ to be the challenge key $K^\diamond$ in the ciphertext integrity game. To this end, it submits c_0 to its own decryption oracle and returns whatever the latter returns. At each query of the form $(\cdot, c_1^\diamond, c_2^\diamond)$ after the j-th query and before the i-th query, $\mathcal{B}^{\mathsf{AE}}$ proceeds exactly as in the j-th query. For all queries $(\cdot, c_1, c_2)$ where $(c_1, c_2) \neq (c_1^\diamond, c_2^\diamond)$, it faithfully runs $\mathsf{Decrypt}_2$ and simulates $H_Q(\cdot, \cdot)$. At the i-th query $(c_0^\diamond, c_1^\diamond, c_2^\diamond)$, $\mathcal{B}^{\mathsf{AE}}$ halts and outputs $c_0^\diamond$ as a fake

ciphertext in the ciphertext integrity game.

If $\mathcal{B}^{\mathsf{AE}}$ successfully guesses $j \in [1, i]$, the simulation of $\mathsf{Game}_{1.i}$ is perfect until the i-th query, at which point $\mathcal{B}^{\mathsf{AE}}$ wins against its challenger. □

Lemma 4. *If $q > 8p(\alpha q)^2 n$ and $p > 4\alpha q \sqrt{n}$, Game_3 is perfectly indistinguishable from Game_2.*

Proof. We show that, from the adversary's view, the random oracle H of Game_3 is identical to that of Game_2. To this end, we consider the function $\mathsf{Encaps} : D_E \to R_q^2 : (r, e_1, e_2) \to \mathsf{Encaps}(PK, (r, e_1, e_2))$ and note that it is injective over its domain $D_E := \{(r, e_1, e_2) \in R^3 : \|r\|, \|e_1\|, \|e_2\| \leq B\}$. Indeed, there cannot exist colliding $(r, e_1, e_2), (\tilde{r}, \tilde{e}_1, \tilde{e}_2) \in D_E$. The equality $b \cdot r + e_2 = b \cdot \tilde{r} + \tilde{e}_2$ over R_q would imply

$$p \cdot (a \cdot s + e)(r - \tilde{r}) \bmod q = \tilde{e}_2 - e_2.$$

Combining this with $a \cdot r + e_1 = a\tilde{r} + \tilde{e}_1$ (over R_q) would yield the equality

$$p \left(s \cdot (\tilde{e}_1 - e_1) + e(r - \tilde{r}) \right) = \tilde{e}_2 - e_2,$$

which would hold over R as the left-hand-side member is a polynomial with coefficients smaller than $4p(\alpha q)^2 n < q/2$ in absolute value. However, this impossible if $\tilde{e}_2 \neq e_2$ since $p > 4B > \|\tilde{e}_2 - e_2\|_\infty$. If $\tilde{e}_2 = e_2$, it implies $\tilde{r} = r$ (since $b \in R_q^\times$) and then $\tilde{e}_1 = e_1$.

Since Encaps is injective over the domain D_E and H_Q is a random function, so is $H_Q\big(\mathsf{Encaps}(PK, (\cdot, \cdot, \cdot))\big)$ over D_E. The two games are thus perfectly indistinguishable from $\mathcal{A}$'s view since H behaves as a random function either way. □

Lemma 5. *Game_4 is perfectly indistinguishable from Game_3.*

Proof. We assume that the event considered in Game_1 does not occur (so that $\|s\|, \|e\| \leq B$) since both games have the same output distribution otherwise.

We call a ciphertext (c_0, c_1, c_2) *good* if (c_1, c_2) satisfy

$$c_1 = a \cdot r + e_1 \in R_q, \qquad c_2 = b \cdot r + e_2 \in R_q$$

for some triple $(r, e_1, e_2) \in R^3$ such that $\|r\|, \|e_1\|, \|e_2\| \leq B$. For a good ciphertext, if the decryption oracle of Game_3 sets $\mathsf{flag} = 1$ at all, the ciphertext gets rejected at step 1 of $\mathsf{Decrypt}_2$. Indeed, if we have $\|e_2 + e_2'\| > 2B$ at step 4 of $\mathsf{Decrypt}_2$, $\|e_2\| \leq B$ implies $\|e_2'\| > B$ and similar implications hold for r' and e_1'. Then, if a good ciphertext is not rejected at step 1 and $\|s\|, \|e\| \leq B$, Remark 1 implies that $\mathsf{flag} = 0$ at step 8.

Consequently, in Game_3, the random oracle H_Q is never evaluated on a good ciphertext since, when $\mathsf{Decrypt}_2$ sets $\mathsf{flag} = 1$ for such a ciphertext, it rejects it at step 1. At the same time, if we have $\mathsf{flag} = 0$ at step 8, we know that the ciphertext is good. This means that (either via a query to $H(\cdot, \cdot, \cdot)$ or a decryption query) H_Q' is only evaluated on good ciphertexts in Game_3. Since H_Q and H_Q' are evaluated on disjoint sub-domains in Game_3, $\mathcal{A}$'s view is exactly the same as if they were emulated using a single random oracle, as in Game_4. □

Lemma 6. Game_5 *is perfectly indistinguishable from* Game_4.

Proof. Game_5 and Game_4 are identical from $\mathcal{A}$'s view until the two decryption oracles give different outputs for some decryption query (c_0, c_1, c_2).

We note that, for any decryption query $C = (c_0, c_1, c_2)$ such that the decryption oracle of Game_5 sets $\mathsf{flag} = 1$ after step 1, it always computes $K = H_Q(c_1, c_2)$ at step 8 of $\mathsf{Decrypt}_2$. So, we only need to worry about decryption queries for which the decryption oracle of Game_4 never sets flag to 1.

For such ciphertexts $C = (c_0, c_1, c_2)$ leading to $\mathsf{flag} = 0$, the rigidity of the encapsulation mechanism ensures that step 7 computes $(r, e_1, e_2) \in D_E$ such that $c_1 = a \cdot r + e_1$, $c_2 = b \cdot r + e_1$, meaning that $(c_1, c_2) = \mathsf{Encaps}(PK, (r, e_1, e_2))$. In this case, at step 8 of $\mathsf{Decrypt}_2$, the decryption oracle of Game_4 computes $K = H(r, e_1, e_2) = H_Q(c_1, c_2)$, exactly as the decryption oracle of Game_5 does.

In both cases $\mathsf{flag} \in \{0, 1\}$, Game_5 is thus indistinguishable from Game_4 as the two decryption oracles always output the same result. $\square$

Lemma 7. *Under the* $\mathsf{RLWE}_{n,k,q,\chi}$ *assumption where* $\chi = D_{\mathbb{Z}^n,\alpha q}^{\mathsf{coeff}}$ *and* $k = \lceil \lambda / \log n \rceil$, Game_7 *is indistinguishable from* Game_6 *if* $q > n^2$. *Concretely, there is a PPT algorithm* $\mathcal{B}$ *such that* $|\Pr[W_7] - \Pr[W_6]| \leq \mathbf{Adv}_{n,k,q,\chi}^{\mathcal{B},\mathsf{RLWE}}(\lambda) + 2^{-\lambda}$.

Proof. Game_7 can only be distinguished from Game_6 in the event F_7 that the challenger $\mathcal{B}$ fails to sample a pair $(a_i, b_i = a_i \cdot s + e_i) \in R_q \times R_q$ such that $b_i \in R_q^\times$ after $k = \lceil \lambda / \log n \rceil$ attempts. If F_7 occurs with non-negligible probability, we can build an RLWE distinguisher $\mathcal{B}$ as follows.

Algorithm $\mathcal{B}$ is given as input an RLWE instance $\{(a_i, b_i) \in R_q^2\}_{i=1}^k$ with $k = \lceil \lambda / \log n \rceil$ samples and must decide if $b_i \sim U(R_q)$ for each $i \in [k]$ or $b_i = a_i \cdot s + e_i$ with $e_i \sim D_{\mathbb{Z}^n,\alpha q}^{\mathsf{coeff}}$. If none of the ring elements $\{b_i\}_{i=1}^k$ is invertible over R_q, the reduction $\mathcal{B}$ returns 1 (meaning that $b_i = a_i \cdot s + e_i$ for all $i \in [k]$). Otherwise, it returns 0 (meaning that $b_i \sim U(R_q)$ for each $i \in [k]$).

We claim that $\Pr[F_7] \leq \mathbf{Adv}^{\mathcal{B},\mathsf{RLWE}}(\lambda) + 2^{-\lambda}$. Indeed, a random $b_i \sim U(R_q)$ is non-invertible with probability $n/q < 1/n$ since Φ splits into degree-1 factors over R_q. Hence, if we sample $k = \lceil \lambda / \log n \rceil$ independent $b_1, \ldots, b_k \hookleftarrow U(R_q)$, the probability of not obtaining any unit in R_q is smaller than $(1/n)^k \leq 2^{-\lambda}$. Hence, if $b_i \sim U(R_q)$ for all $i \in [k]$, $\mathcal{B}$ outputs 1 with probability at most $2^{-\lambda}$. Now, if $b_i = a_i \cdot s + e_i$ for each $i \in [k]$, $\mathcal{B}$ outputs 1 with probability $\Pr[F_7]$ by hypothesis. This shows that $|\Pr[W_7] - \Pr[W_6]| \leq \Pr[F_7] \leq \mathbf{Adv}^{\mathcal{B},\mathsf{RLWE}}(\lambda) + 2^{-\lambda}$. $\square$

Lemma 8. *Under the* RLWE *assumption,* Game_9 *is indistinguishable from* Game_8. *We have* $|\Pr[W_9] - \Pr[W_8]| \leq (1 - 2^{-\lambda})^{-1} \cdot \mathbf{Adv}_{n,k,q,\chi}^{\mathcal{B},\mathsf{RLWE}}(\lambda)$, *where* $\chi = D_{\mathbb{Z}^n,\alpha q}^{\mathsf{coeff}}$ *and* $k = 2\lceil \lambda / \log n \rceil$ *is the number of samples.*

Proof. Assuming that the adversary can distinguish between Game_9 and Game_8, we build an RLWE distinguisher $\mathcal{B}$.

The reduction $\mathcal{B}$ is given $k' = \lceil \lambda / \log n \rceil$ tuples $(a_i, b_i, c_{i,1}, c_{i,2}) \in R_q \times R_q$ where $(a_i, b_i) \sim U(R_q \times R_q)$ and each pair $(c_{i,1}, c_{i,2})$ is either uniform over $R_q \times R_q$ or of the form $(c_{i,1}, c_{i,2}) = (a_i \cdot r + e_{i,1}, b_i \cdot r + e_{i,2})$ for some $e_{i,1}, e_{i,2} \sim D_{\mathbb{Z}^n,\alpha q}^{\mathsf{coeff}}$. To

define the public key, $\mathcal{B}$ first checks if there exists an index $i \in [k']$ such that b_i is invertible over R_q. If no such index is found, $\mathcal{B}$ aborts and outputs a random bit. Otherwise, it picks the first $i \in [k']$ such that $b_i \in R_q^\times$ and defines PK by setting $a = a_i$ and $b = b_i$. Then, it simulates $\mathcal{A}$'s view as in Game_8. In the challenge phase, $\mathcal{A}$ outputs M_0, M_1. At this point, $\mathcal{B}$ constructs the challenge ciphertext by setting $c_1^\star = c_{i,1}$, $c_2^\star = c_{i,2}$ and $c_0^\star = \mathsf{E}_{K^\star}(M_d)$, where $K^\star = H_Q(c_1^\star, c_2^\star)$. After the challenge phase, $\mathcal{B}$ answers all queries as in the first stage and eventually outputs whatever $\mathcal{A}$ outputs.

Let Fail the event that none of the $\{b_i\}_{i=1}^{k'}$ is invertible. We have $\Pr[\neg\mathsf{Fail}] \geq 1 - 2^{-\lambda}$ by the same arguments as in the proof of Lemma 7. Conditionally on $\neg\mathsf{Fail}$, if $(c_{i,1}, c_{i,2}) = (a_i \cdot r + e_{i,1}, b_i \cdot r + e_{i,2})$, then $\mathcal{A}$'s view is exactly as in Game_8. Then, when $(c_{i,1}, c_{i,2}) = (a_i \cdot r + e_{i,1}, b_i \cdot r + e_{i,2})$ for each $i \in [k']$, $\mathcal{B}$ outputs 1 with probability $\frac{1}{2} \cdot \Pr[\mathsf{Fail}] + \Pr[W_8] \cdot \Pr[\neg\mathsf{Fail}]$.

Conditionally on event $\neg\mathsf{Fail}$, if $(c_{i,1}, c_{i,2}) \sim U(R_q \times R_q)$, $\mathcal{A}$'s view is identical to that of Game_9. When $(c_{i,1}, c_{i,2}) \sim U(R_q \times R_q)$, $\mathcal{B}$ thus outputs 1 with probability $\frac{1}{2} \cdot \Pr[\mathsf{Fail}] + \Pr[W_9] \cdot \Pr[\neg\mathsf{Fail}]$. Then, $\mathcal{B}$'s advantage as an RLWE distinguisher is at least $|\Pr[W_9] - \Pr[W_8]| \cdot \Pr[\neg\mathsf{Fail}]$. This yields the claimed inequality $|\Pr[W_9] - \Pr[W_8]| \leq (1 - 2^{-\lambda})^{-1} \cdot \mathbf{Adv}^{\mathcal{B},\mathsf{RLWE}}(\lambda)$. $\qquad\square$

We note that bound (3) tightly relates the security of the scheme to the RLWE assumption. On the other hand, it loses a quadratic factor $O(Q^2)$ with respect to the ciphertext integrity of the symmetric authenticated encryption scheme. However, the term $Q(Q + 1) \cdot \mathbf{Adv}^{\mathsf{AE\text{-}INT}}(\lambda)$ becomes statistically negligible if Π^{sym} is realized using an information-theoretically secure one-time MAC, as we discuss in the following instantiation section.

4.3 Parameters and Instantiations

PARAMETERS. In an instantiation in fully splitting rings R_q (which allows faster multiplications using the NTT in $\mathsf{Encrypt}$), we use $\Phi(X) = X^n + 1$, where n is a power of 2, with a modulus $q = 1 \bmod 2n$.

For correctness, we need to choose $\alpha \in (0, 1)$, q and p such that $p/2 > 2B$ and $4p(B^2 + 1) < q/2$, which satisfy the requirements of Lemma 4. To apply Lemma 2, we can set $\alpha \in (0, 1)$ so that $\alpha q = \Omega(\sqrt{n})$. To satisfy all conditions, we may thus set $p = \Theta(n)$, $q = \Theta(n^3)$ and $\alpha^{-1} = \Theta(n^{2.5})$.

CONCRETE PROPOSALS. Around the 128-bit security level, n has to be somewhere between 512 and 1024. When r, e_1, e_2 are sampled from $D_{\mathbb{Z}^n, \alpha q}^{\mathsf{coeff}}$, we relax the constraint $\alpha q = \Omega(\sqrt{n})$ but still choose $\alpha q \approx \lambda^{1/2}$ in order to apply Lemma 1 and we take into account the constraints of Lemma 4 and Lemma 7. If we set $n = 1024$ so as to have a power of 2, we can use $q = 1813307393$, $\alpha q = 12$, and $p = 1537$, so that the pair (c_1, c_2) fits within 7.75Kb.

In order to obtain a more optimized instantiation, we can further sample s, e, r, e_1, e_2 from a centered binomial distributions as suggested in [4, 21, 37] and set the parameters according to the concrete hardness against known attacks via the LWE estimator [3]. As in [37], we use the distribution $\bar{\psi}_2^n$ obtained by reducing $\psi_2^n \bmod 3$, where ψ_k^n is defined over $\mathbb{Z}^n$ as the distribution

$\{\sum_{i=1}^{k}(a_i - b_i) \mid a_i, b_i \leftarrow U(\{0,1\}^n)\}$. This modification only entails minor modifications in the security proof, which are detailed in Supplementary Material C. We then obtain ciphertexts of 4Kb by setting $n = 1024$, $p = 5$ and $q = 59393$.

In order to push optimizations even further, we could use rings where n does not have to be a power of 2, as suggested in [37]. For example, the cyclotomic polynomial $\Phi_{3n} = X^n - X^{n/2} + 1$ with $q = 1 \bmod 3n$ and $n = 2^i 3^j$ still gives a fully splitting $R_q = \mathbb{Z}_q[X]/(\Phi_{3n})$. This allows choosing $n = 768$, $p = 5$ (so that $\|\bar{e}_2\|_\infty \leq (p-1)/2$), and $q = 28 \cdot 3n + 1 = 64513$. Correctness is ensured since we have $\|a \cdot b\|_\infty \leq 2\|a\|\|b\|$ for any $a, b \in R = \mathbb{Z}[X]/(\Phi_{3n})$, so that $\|\bar{e}_2 + p \cdot (e \cdot \bar{r} - \bar{e}_1 \cdot s)\|_\infty \leq (p-1)/2 + 8 \cdot p \cdot n \leq (q-1)/2$. We would then obtain a ciphertext (c_1, c_2) that only takes 3Kb to represent.

In all instantiations, the public key can compressed down to roughly 50% of the ciphertext size if we derive the random $a \in R_q$ from a hash function modeled as a random oracle (as considered by many NIST candidates).

INSTANTIATING THE DEM COMPONENT. The symmetric authenticated encryption scheme Π^{sym} can in general be instantiated with a a leakage-resistant Enc-then-MAC mode of operation. Candidates for this purpose that rely on a masked block cipher or permutation can be found in [15]. Yet, POLKA encourages the following more efficient solution based on a key-homomorphic one-time MAC. Specifically, if ℓ_m is the message length, we can use a key length $\kappa = \lambda + 2\ell_m$ and a pseudorandom generator $G : \{0,1\}^\lambda \to \{0,1\}^{\ell_m}$. To encrypt $M \in \{0,1\}^{\ell_m}$, we parse $K = H(r, e_1, e_2) \in \{0,1\}^\kappa$ as a triple $K = (K_0, K_1, K_2) \in \{0,1\}^\lambda \times (\{0,1\}^{\ell_m})^2$. Then, we compute a ciphertext $c = (\bar{c}, \tau) = (M \oplus G(K_0), K_1 \cdot \bar{c} + K_2)$, where the one-time MAC $\tau = K_1 \cdot \bar{c} + K_2$ is computed over $\text{GF}(2^{\ell_m})$.

This specific MAC allows "annihilating" the quadratic term $O(Q^2)$ in the security bound (3). In the ciphertext integrity experiment (defined in Supplementary Material B), the adversary's advantage can then be bounded as $(Q+1)/2^{\ell_m}$ if Q is the number of decryption queries. To make the term $Q(Q+1) \cdot \mathbf{Adv}^{\text{AE-INT}}(\lambda)$ statistically negligible in (3), we can assume that $\ell_m \geq \lambda + 3\log^2 \lambda$ in order to have $(Q+1)^2 Q/2^{\ell_m} < 2^{3\log^2 \lambda}/2^{\ell_m} < 2^{-\lambda}$. Concretely, if we set $\lambda = 128$ and assume $Q < 2^{60}$, we can choose $\ell_m \geq 308$.

In terms of leakage, the computation $K_1 \cdot \bar{c} + K_2$ is linear in the key and can therefore be masked with overheads that are linear in the number of shares (rather than quadratic for a block cipher or permutation). The constraint on the message length could be relaxed by hashing the message at the cost of an additional idealized assumption, which we leave as a scope for further research.

5 Side-Channel Security Analysis

We now discuss the leakage properties of POLKA. In Section 5.1 we introduce the general ideas supporting its leveled implementation and explain how its security requirements can be efficiently fulfilled. In Section 5.2, we focus on its most novel part, namely the variant of the LWPR assumption on which this implementation relies. We also provide cryptanalysis challenges to motivate further research on

hard physical learning problems. In Section 5.3 we describe a hardware architecture for the most sensitive DPA target of POLKA. Our descriptions borrow the terminology introduced in [45] for symmetric cryptography. Namely, we denote as leakage-resilient implementations of which confidentiality guarantees may vanish in the presence of leakage, but are restored once leakage is removed from the adversary's view; and we denote as leakage-resistant implementations that preserve confidentiality against leakage even for the challenge encryption.

5.1 Leveled Implementation and Design Goals

The high-level idea behind leveled implementations is that it may not be necessary to protect all the parts of implementation with equally strong (and therefore expensive) side-channel countermeasures. In the following, we describe how POLKA could be implemented in such a leveled manner. For this purpose, and as a first step, we follow the heuristic methodology introduced in [15] and identify its SPA and DPA targets in decryption. The resulting leveled implementation of POLKA is represented in Figure 1. The lighter green-colored (dummied) operations need to be protected against SPA. The darker green-colored operations need to be protected against SPA with averaging (avg-SPA), which is a SPA where the adversary can repeat the measurement of a fixed target intermediate computation in order to remove the leakage noise. The lighter blue-colored operations need to be protected against DPA with unknown (dummied) inputs (UP-DPA). The darker blue-colored operations must be protected against DPA. Operations become generically more difficult to protect against side-channel attacks when moving from the left to the right of the figure. Eventually, securing the first four steps in the figure is needed to protect the long-term secret of POLKA, and therefore to ensure leakage-resilience. By contrast, securing the fifth step is only needed to ensure leakage-resistance (these values can leak in full in case only leakage-resilience is required). Next, we first explain how these security requirements can be efficiently satisfied by hardware designers. We then discuss the advantages of this implementation over a uniformly protected one.

SPA and DPA protections. All the operations requiring SPA protection (with or without averaging) can be efficiently implemented thanks to parallelism in hardware. Typically, we expect that an implementation manipulating 128 bits or more in parallel is currently difficult to attack via SPA, even when leveraging advanced analytical strategies [78].[7] A bit more concretely, this reference shows that single-trace attacks are possible for Signal-to-Noise Ratios (SNRs) higher than one. Adversaries targeting a 128-bit secret based on 8-bit (resp., 32-bit) hypotheses would face an SNR of $\frac{1}{16}$ (resp., $\frac{1}{4}$). Securing the computations in steps 1, 3 and 4 of Figure 1 against side-channel attacks should therefore lead to limited overheads. Note that the dependency on a dummy ciphertext in step 3 prevents the adversary to control the intermediate computations (and for example to try canceling the algorithmic noise for those sensitive operations).

---

[7] As will be clear in conclusions, software implementations are left as an interesting open problem. In this case, the typical option to obtain security against SPA would be to emulate parallelism thanks to the shuffling countermeasure [79].
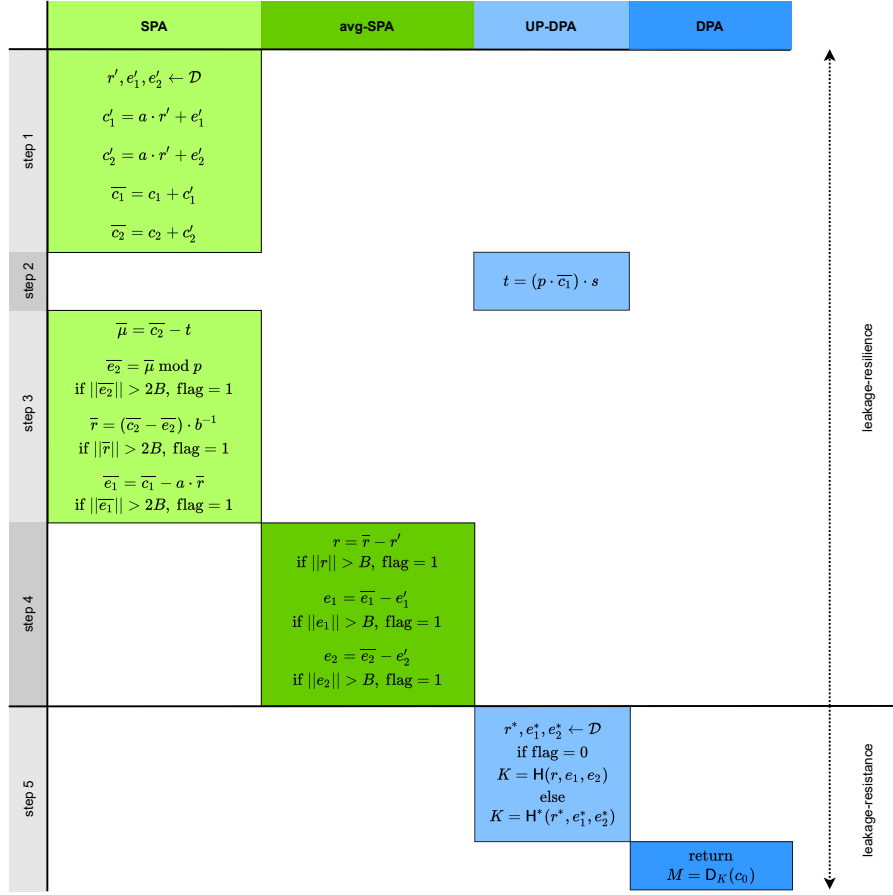
Fig. 1: Leveled implementation of `POLKA`.

Security against DPA is in general expected to be significantly more expensive to reach. The standard approach for this purpose is to mask all the operations that can be targeted, which leads to (roughly) quadratic performance overheads [51]. Furthermore, implementing masking securely is a sensitive process, which requires dealing with composition issues [9,29], physical defaults such as glitches [60,64] or transitions [7,28] or even their combination [25,26].

The main observation we leverage in `POLKA` is that its most critical DPA sensitive operation shares similarities with the key-homomorphic re-keying schemes used in symmetric cryptography to prevent side-channel attacks [35,38,39,61]. Namely, the operation $t = (p \cdot \overline{c_1}) \cdot s$ in **step** 2 of Figure 1 can indeed be viewed as the product between a long-term secret s and an ephemeral (secret) value $(p \cdot \overline{c_1})$. As a result, it can be directly computed as $t = \sum_{i=1}^{d} (p \cdot \overline{c_1}) \cdot s^i$, where $s = s^1 + s^2 + \ldots s^d$ and the s^i's are the additive shares of the long-term secret

s. Besides the linear (rather than quadratic) overheads that such a solution enables, key-homomorphic primitives have two important advantages for masking. First, their long-term secret can be refreshed with linear randomness requirements [10]. Second, their natural implementation offers strong immunity against composition issues and physical defaults [24]. On top of this, the fact that the variable input of $(p \cdot \overline{c_1}) \cdot s^i$ is dummied (hence unknown) implies that it will need one less share than in a known input setting [17].

Even more importantly, and as discussed in the aforementioned papers on fresh re-keying, it is then possible to re-combine the shares and to perform the rest of the computations on unshared values, hence extending the interest of a leveled approach. Various models have been introduced for this purpose in the literature, depending on the type of multiplication to perform. The first re-keying schemes considered multiplications in binary fields that require a sufficient level of noise to be secure [13,14]. Dziembowski et al. proposed a (more expensive) wPRF-based re-keying that is secure even if its output is leaked in full [39]. Duval et al. proposed an intermediate solution that only requires the (possibly noise-free) leakage function to be surjective and "incompatible" with the field multiplication: they for example show that this happens when combining multiplications in prime fields with the Hamming weight leakage function, which they formalized as the LWPR assumption [38]. Given that the multiplication of POLKA is based on prime moduli, we next focus on this last model, which provides a nice intermediate between efficiency and weak physical assumptions.

As for the operations of step 5 of Figure 1, we first observe that despite the inputs of H being ephemeral, it is possible that an adversary obtains a certain level of control over them by incrementally increasing c_1 or c_2. This explains why it must be secure against DPA (with unknown plaintexts since r, e_1 and e_2 are unknown as long as steps 3 and 4 are secure against SPA). Finally, the protection of the authenticated encryption is somewhat orthogonal to POLKA since it is needed for any DEM. The standard option for this purpose would be to use a leakage-resistant mode of operation that ensures side-channel security with decryption leakage. As discussed in [15], state-of-the-art modes allow the authenticated encryption scheme to be leveled (i.e., to mix SPA-secure operations with DPA-secure ones), like the rest of POLKA. But as mentioned in Section 4.3, an even more efficient solution is to use an Enc-then-MAC scheme with a one-time key-homomorphic MAC that is linear in the key and therefore easy to mask.

Discussion. The main advantage of POLKA is that its structure allows avoiding the costly implementation of uniformly protected operations based on masking. In this respect, it is worth recalling that: (i) the removal of the dummy ciphertext takes place as late as possible in the process (i.e., just before the hashing and symmetric decryption), and (ii) if only the long-term secret s must be protected (i.e., if only leakage-resilience is required), step 5 of Figure 1 does not need countermeasures. Overall, these design tweaks strongly limit the side-channel attack surface and the need to mask non-linear operations compared to algorithms like KYBER or SABER, at the cost of an admittedly provocative LWPR assumption.

We note that the explicit rejection process allows avoiding the need to manipulate additional long-term secret material. This is in contrast with an implicit rejection mechanism where a pseudorandom "garbage key" is generated in case of invalid ciphertext, the generation of which must be protected against DPA. Yet, ensuring the leakage-resistance of POLKA against timing attacks still requires running a dummy hash function when flag $= 1$ in order to ensure constant time. Otherwise, the same granular increase of c_1 or c_2 as mentioned to justify the DPA security requirements of H could leak information on e_1, e_2 and r if a timing channel allowed detecting whether a $\perp$ message is generated during step 4 or step 5 (by the authenticated encryption scheme). This also means that it should be hard to distinguish whether the flag is 0 or 1 with SPA to ensure leakage-resistance. We conjecture the latter is simpler/cheaper than protecting another long-term secret against DPA (as required with implicit rejection), but as mentioned in introduction, POLKA could be adapted with an implicit rejection mechanism as well (in which case, it should also be hard to distinguish whether the key used to decrypt is a garbage one or not thanks to SPA).

5.2 Learning With Physical Rounding Assumption

We now move to the main assumption that allows POLKA to be masked with linear overheads. Namely, we study the security of step 2 in Figure 1 after the recombinations of the shares. In other words, we study the security of the long-term secret s assuming that the adversary can observe the leakage of the (unmasked) output t.[8] We start by recalling the LWPR problem introduced at CHES 2021 [38], then discuss its adaptation to polynomial multiplications needed for POLKA. We finally propose parameters & cryptanalysis challenges.

A. The original LWPR problem can be viewed as an adaptation of the crypto dark matter proposed by Boneh et al. in [20], which showed that low-complexity PRFs can be obtained by mixing linear functions over different small moduli. Duval et al. observed that letting one of these functions being implicitly computed by a leakage function can lead to strong benefits for masking against side-channel attacks. Intuitively, it implies that a designer only has to implement a key-homomorphic function securely (i.e., the first crypto dark matter mapping), since the second (physical) mapping never has to be explicitly computed: it is rather the leakage function that provides its output to the adversary. The formal definition of the resulting LWPR problem is given next.

Definition 2 (Learning with physical rounding [38]). *Let $q, x, y \in \mathbb{N}^*$, q prime, for a secret $\kappa \in \mathbb{F}_q^{x \times y}$. The $\mathsf{LWPR}_{\mathsf{L_g}, q}^{x,y}$ sample distribution is given by:*

$$\mathcal{D}_{\mathsf{LWPR}_{\mathsf{L_g}, q}^{x,y}} := (r, \mathsf{L_g}(\kappa \cdot r)) \text{ for } r \in \mathbb{F}_q^y \text{ uniformly random,}$$

---

[8] As mentioned in subsection 5.1, the security of the internal computations of $t = \sum_{i=1}^{d} (p \cdot \overline{c_1}) \cdot s^i$ is obtained thanks to masking. So here, we only need to argue that the leakage of the recombined t does not lead to strong attacks.

where $\mathsf{L_g} : \mathbb{F}_q^x \to \mathbb{R}^d$ *is the physical rounding function. Given query access to* $\mathcal{D}_{\mathsf{LWPR}_{\mathsf{L_g},q}^{x,y}}$ *for a uniformly random* κ*, the* $\mathsf{LWPR}_{\mathsf{L_g},q}^{x,y}$ *problem is* $(\chi, \tau, \mu, \epsilon)$*-hard to solve if after the observation of* χ LWPR *samples, no adversary can recover the key* κ *with time complexity* τ *, memory complexity* μ *and probability* $\geq \epsilon$.

Concretely, the LWPR problem consists in trying to retrieve a secret key matrix κ using the information leakage emitted during its product with a random vector r. It corresponds to a learning problem similar to LWR [8], with the rounding function instantiated with a leakage function. Its security depends on the dimensions (q, x, y) and the leakage function considered. In [38], it is argued that this problem is hard in the case of the Hamming weight leakage function that is frequently encountered in practice (with a regular binary representation). This problem can then be used as the basis of a fresh re-keying mechanism, producing an ephemeral key $\in \mathbb{F}_q^x$ thanks to the aforementioned product. It can be implemented serially (by setting the x parameter to 1), in which case words of $\log_2(q)$ bits of ephemeral key will be produced one by one, or in a parallel manner by increasing x, therefore producing $x \times \log_2(q)$ bits of ephemeral key at once.

The security analysis of Duval et al. first shows that the complexity of various (algebraic and statistical) attacks against such a fresh re-keying scheme grows exponentially with the (main) security parameter y. Due to the inevitably heuristic nature of their cryptanalysis, they next use the parallelism parameter x as a way to obtain security margins. Denoting the words of the ephemeral key as z_i, a serial implementation will leak independent $\mathsf{HW}(\mathsf{g}(z_i))$ values to the adversary, while a parallel one will leak $\sum_{i=1}^{x} \mathsf{HW}(\mathsf{g}(z_i))$ values, with $\mathsf{g}(z_i)$ denoting the binary representation of z_i. As a result, the amount of mutual information per word that such Hamming weight leakages provide rapidly decreases as $\frac{\log_2(x \log_2(q))}{x \log_2(q)}$, which is expected to make attacks more challenging.

The instance proposed by Duval et al. uses a 31-bit prime modulus $p = 2^{31} - 1$ with parameters $x = 4$ and $y = 4$ (i.e., it assumes that four $\log_2(p)$-bit multiplications can be performed in parallel). It claims 128 bits of security.

As can be seen in Figure 1, `step 2` of `POLKA` shares strong similarities with the aforementioned fresh-re-keying scheme based on LWPR, by simply viewing the intermediate value t as an ephemeral key. We next discuss the differences between the original LWPR assumption and the one needed for `POLKA`.

B. Ring-LWPR. Leveraging the fact that ring variants of learning problems are common [58], we now describe a ring version of the LWPR problem. Let $r := p \cdot \overline{c_1}$. Seeing s as a long-term secret (similar to κ in the original LWPR problem), the t value can be re-written as $t = r \cdot s$. Further denoting s_i (resp., r_i) the coefficients of s (resp., r), we can write:

$$r \cdot s = \left(\sum_{i=0}^{n-1} s_i X^i \right) \cdot \left(\sum_{i=0}^{n-1} r_i X^i \right) = \sum_{i=0}^{2n-2} \left(\sum_{j=max(0,i-n+1)}^{min(i,n-1)} s_j r_{i-j} \right) X^i,$$

$$= \sum_{i=0}^{n-2} \left(\sum_{j=0}^{i} s_j r_{i-j} \right) X^i + \left(\sum_{j=0}^{n-1} s_j r_{i-j} \right) X^{n-1} + \sum_{i=n}^{2n-2} \left(\sum_{j=i-n+1}^{n-1} s_j r_{i-j} \right) X^i,$$

$$= \sum_{i=0}^{n-2} \left(\sum_{j=0}^{i} s_j r_{i-j} \right) X^i + \left(\sum_{j=0}^{n-1} s_j r_{i-j} \right) X^{n-1} + \sum_{i=0}^{n-2} \left(\sum_{j=i+1}^{n-1} s_j r_{n+i-j} \right) X^{n+i},$$

$$= \sum_{i=0}^{n-2} \left(\sum_{j=i+1}^{n-1} s_j r_{i-j} - \sum_{j=0}^{i} s_j r_{n+i-j} \right) X^i + \left(\sum_{j=0}^{n-1} s_j r_{i-j} \right) X^{n-1}.$$

The above equation highlights the matrix representation of the polynomial multiplication carried out in POLKA. If we represent polynomials as n-dimension vectors, where the i-th coefficient is the polynomial's i-th coefficient, the product $r \cdot s$ can be represented as the following matrix-vector product:

$$\begin{pmatrix} r_0 & -r_{n-1} & \dots & -r_2 & -r_1 \\ r_1 & r_0 & \ddots & & -r_2 \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ r_{n-2} & & \ddots & \ddots & -r_{n-1} \\ r_{n-1} & r_{n-2} & \dots & r_1 & r_0 \end{pmatrix} \cdot \begin{pmatrix} s_0 \\ s_1 \\ \vdots \\ s_{n-2} \\ s_{n-1} \end{pmatrix},$$

The key is represented as the vector (rather than the matrix) in order to optimize memory usage when splitting it into shares. This product can therefore be seen as a large LWPR instance, with two significant differences. First, a circulant matrix is used instead of one having independent coefficients (which we will discuss when selecting parameters in the next subsection). Second, the size of the matrix is (much) larger than the one in the original LWPR. Concretely, this second difference implies that in practice, these products are unlikely to be performed in one step: they will rather be decomposed into several submatrix-subvector products. For this purpose, let $x \in \mathbb{N}$ be a divider of n, the s matrix can then be split in $\frac{n}{x}$ $(x \times n)$-submatrices, denoted $(B_u)_{0 \le u < \frac{n}{x}}$. The product can then be decomposed into $\frac{n}{x}$ subproducts, as illustrated by the following box:

with x serving as a parameter to adapt the security vs. performance tradeoff, as in the original LWPR. For a given k, one can explicitly obtain the coefficient i, j of B_u. For a proposition $\mathcal{P}$, denote $\mathbb{1}_{\mathcal{P}} := \begin{cases} 1 \text{ if } \mathcal{P} \\ 0 \text{ else} \end{cases}$. Then, $B_u^{i,j} = (2\mathbb{1}_{(ux+i-j)<0} -$

$1) \cdot r_{ux+i-j \pmod{n}}$ and the $(x \times n)$-submatrices are therefore Toeplitz, determined by their first line and first column, each other value being equal to their top-left neighbor. Concretely, the x parameter sets the number of coefficients that are computed in parallel, so that an adversary will be granted access to $\frac{n}{x}$ samples given by the leakage function applied to $x \log_2(q)$ bit-values.

Definition 3 (Ring learning with Physical rounding). *Let $q, x, n \in \mathbb{N}$, q prime, for a secret $s \in R_q$. The $\mathsf{RLWPR}_{\mathsf{L_g},q}^{n,x}(s)$ sample distribution is given as:*

$$\mathcal{D}_{\mathsf{RLWPR}_{\mathsf{L_g},q}^{n,x}(s)} := \left(r, \left(\mathsf{L_g}\left(B_u \cdot s \right) \right)_{0 \le u < \frac{n}{x}} \right),$$

where $\mathsf{L_g}$ is the physical rounding function and the (B_u) are submatrices made of elements of r as defined above. Given query access to $\mathcal{D}_{\mathsf{RLWPR}_{\mathsf{L_g},q}^{n,x}(s)}$ for a uniformly random s, the $\mathsf{RLWPR}_{\mathsf{L_g},q}^{n,x}(s)$ problem is $(\chi, \tau, \mu, \epsilon)$-hard to solve if after the observation of χ RLWPR samples, no adversary can recover the key s with time complexity τ, memory complexity μ and probability higher than ϵ.

Note that an implementer can also split each $(x \times n)$ submatrice into $\frac{n}{y}$ pieces (e.g., to further trade circuit size for cycles in hardware) but this has no impact on the security of the RLWPR assumption, since the internal computations are assumed to be secure thanks to masking, as per Footnote 8. By contrast, more parallel implementations (reflected by a large y) may increase the level of noise in the measurements and therefore the security of the masked computations [36]. So overall, the security of the above RLWPR problem only depends on n and x. For a similar reason, the polynomial multiplication can be implemented naively or in the NTT domain, as long as the inverse NTT is applied on every share before recombination. A more efficient solution for the NTT case would be to recombine the shares in the NTT domain (so that the inverse NTT is computed only once). This would provide the adversary with leakages having a slightly different structure than in the above RLWPR problem. We leave the security analysis of this variant as an interesting scope for further research.

C. Choice of parameters and cryptanalysis challenges. Applying the security analysis of LWPR described in Part A of this subsection to RLWPR, we could choose instances based on the main security parameter n using the parallelism parameter x to obtain security margins (a necessary condition to reach λ bits of security is that $(n+1)\log_2 q + 3\log_2 n \ge \lambda$). However, as mentioned in Part B of this subsection, the RLWPR problem is not exactly the same as the LWPR one. Negatively, the $(x \times n)$ submatrices are Toeplitz and they are not independent. While using structured matrices is not unusual in the context of hard learning problems (see for example [58] for RLWE or [44, 53] for LPN variants) and we could not identify parts of the analyses in [38] that become significantly easier in this case, the corresponding problems are less studied. It is therefore natural to consider additional security margins. As for the non-independence issue, considering that the security of the full RLWPR is at least as strong as the security of one of its subproducts, we can conservatively assume

28

that one RLWPR sample will generate at most $\frac{n}{x}$ leakages about this subsecret. So parameters' choices covering that the data complexity of attacks against RLWPR is reduced by this factor compared to attacks against LWPR should be safe. Positively, the secret s in the leveled implementation of POLKA is not multiplied with a public r since this r value is dummied. So concretely, the side-channel adversary will only be provided with the leakage of this ephemeral value.

Putting things together, and considering the instances proposed in subsection 4.3, we propose the following sets of parameters as interesting targets for cryptanalysis with a time complexity of less than 2^{128} and at most 2^{64} queries to a RLWPR-b$_{\mathsf{HW_g},q}^{n,x,y}(s)$, assuming a Hamming weight leakage function.

	$\log_2(q)$	n	x
Set 1	31	1024	8
Set 2	31	1024	4
Set 3	16	1024	8

We leave the investigation of more aggressive parameters (especially reducing the level of parallelism x) as interesting directions for further research.

5.3 Hardware Performance Evaluation

We finally complete our results with an FPGA prototype for the masked computation of t (i.e., step 2 in Figure 1), which is the most sensitive operation in POLKA. As a first proof-of-concept, we consider a naive implementation of the multiplication for this purpose (i.e., without NTT), which is easier to design in hardware. As mentioned in Section 5.2, the computation of t is decomposed in the different subproducts $(B_u \cdot s)$. To perform them, the hardware module dynamically generates submatrices of coefficients together with their corresponding signs, computes the modular multiplications and accumulates the intermediate results. Leveraging on the key-homomorphism of the operations, masking the complete computation boils down to operating with a shared representation of s (with d shares) and to sequentially process each share s^i with the same vector r. In order to avoid excessive circuit size, each submatrix B_u as well as the long term key shares $(s^i)_{0 \leq i < d}$ are split in $\frac{n}{y}$ smaller pieces, respectively denoted as $(B_{u,v})_{0 \leq v < \frac{n}{y}}$ and $(s_v^i)_{0 \leq v < \frac{n}{y}}$, that are then processed sequentially.

The global architecture of our hardware module is shown in Figure 2. It takes as input the vector r (circled in blue), the subvector share s_v^i (circled in red) and control signals (circled in green), and it outputs the value of u (circled in black). The blocks $B_{u,v}$ are generated by the BlockGenerator module. To avoid relying on costly large MUXes, they are generated iteratively using a shift register. The shift register holds $P_{u,v} = \left[P_{u,v}^0, P_{u,v}^1, \ldots, P_{u,v}^{n-1}\right]$, a permuted version of the vector r in such a way that the first row of a given $B_{u,v}$ is obtained by keeping the first y coefficients stored in the registers (i.e., $\left(B_{u,v}^{0,j} = P_{u,v}^j\right)_{0 \leq j < y}$). The $x-1$ remaining rows of the block are similarly generated in parallel by keeping the

29

first y coefficients of a rotated version of $P_{u,v}$. Such a rotation operation has the advantage of being easily and virtually freely implemented in hardware by applying appropriate wiring to the rotated state. To perform a full computation over a share, the different $B_{u,v}$ are sequentially generated by updating the value of $P_{u,v}$. In practice, when v only is incremented, the $P_{u,v}$ is rotated by y coefficients to the left (i.e., $P^i_{u,v+1} = P^{(i+y) \bmod n}_{u,v}$). When iterating over u only, it is rotated by x coefficients to the right (i.e., $P^i_{u+1,v} = P^{(n-x-i)\bmod n}_{u,v}$). When both u and v are iterated (which occurs when $v = n/y - 1$), the direction and the amount of the rotation depends on the instance parameters and are trivially derived by combining rotations occurring when only u or v are incremented.
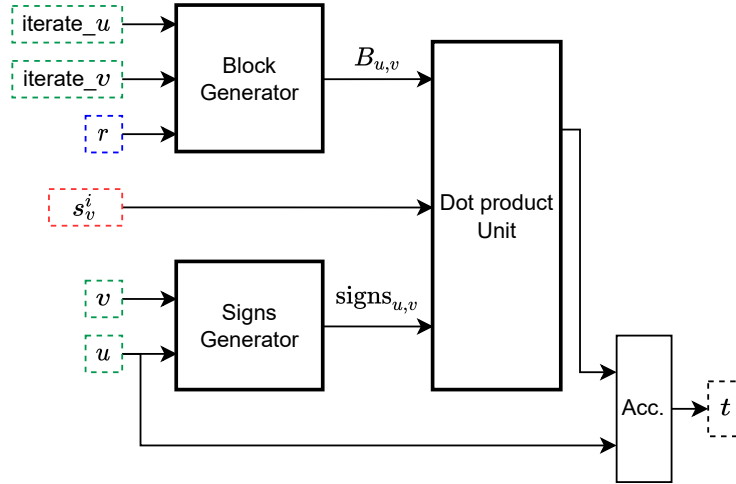

Fig. 2: Global architecture of our RLWPR prototype.

The signs of the submatrix coefficients $\text{signs}^{i,j}_{u,v}$ are generated independently by the `SignsGenerator` module. The latter is a simple combinatorial logic block that naively computes each sign based on the indices u and v. These are control signals that are forwarded together with $B_{u,v}$ to the computation core `DotProductUnit`. The latter is the main computation core of the design and computes the subproduct between $B_{u,v}$ and s^i_v in 1 clock cycle (hence matching the RLWPR assumption). Fully implemented with combinatorial logic, it is composed of $x \times y$ parallel multiplication cores followed by adder trees used to sum up the y multiplication results computed for a single row. IIn practice, the design relies on the Digital Signal Processors (DSPs) resources, nowadays embedded in FPGAs to reduce the area cost of the multiplication circuitry. Finally, the x intermediate results obtained are used to update the corresponding accumulators. A MUX is used to select the values of the accumulators that need to be updated (denoted $t^{0 \le i < x-1}_u$ in Figure 3). These are then routed back to be added modulo

q with the values outputted by the adder trees. Overall, the processing of a share is performed in n^2/xy clock cycles, leading to dn^2/xy cycles of latency for the computation part when all the shares are considered.
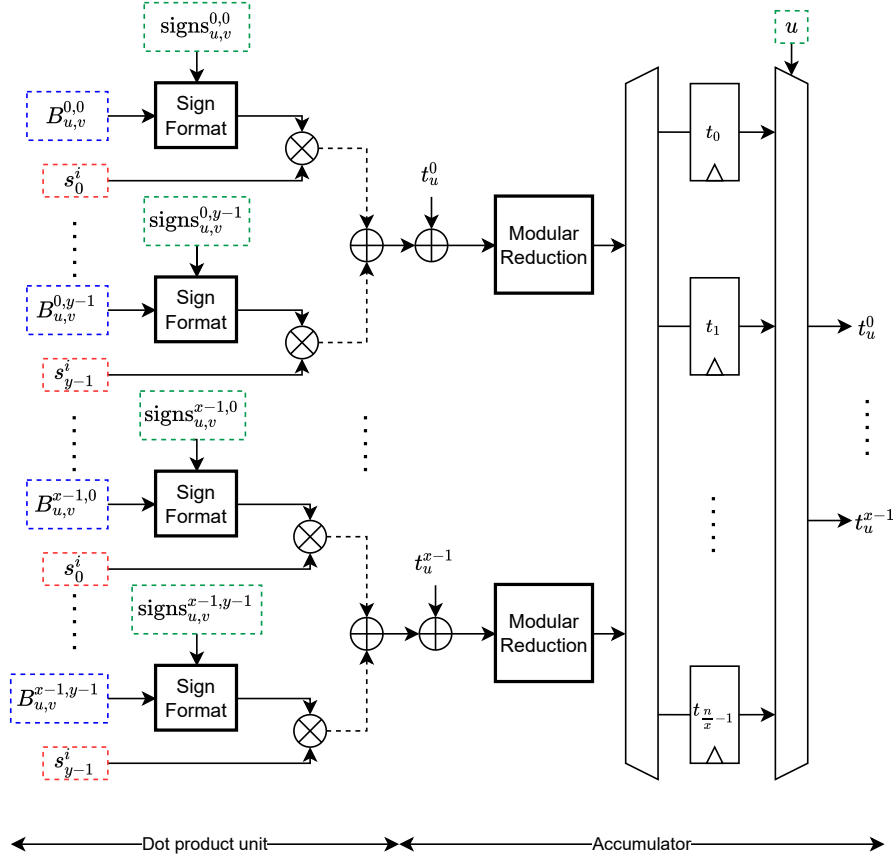

Fig. 3: Subproduct computation core and accumulator mechanism.

As depicted in Figure 4, the long-term secret shares s_v^i are stored in independent memory elements. A MUX is used in order to properly route the value s_v^i from the different memory blocks' instances. Overall, d memory blocks with a capacity of $n\log_2(q)$ bits are required to store the sharing of s. In practice, these RAMs are implemented with BRAM block resources, which are dedicated memory blocks embedded on nowadays FPGA. In order to reduce the risk of transition-based shares recombinations, the output of each RAM is forced to zero when no valid read operation has been issued. Besides, a dead cycle (i.e., a cycle without any read operation) occurs when switching from one share to another. Combined together, these mechanisms ensure that no physical recom-

bination of the share can occur when transiting from one share to another at the cost of limited area overhead and d additional clock cycles.
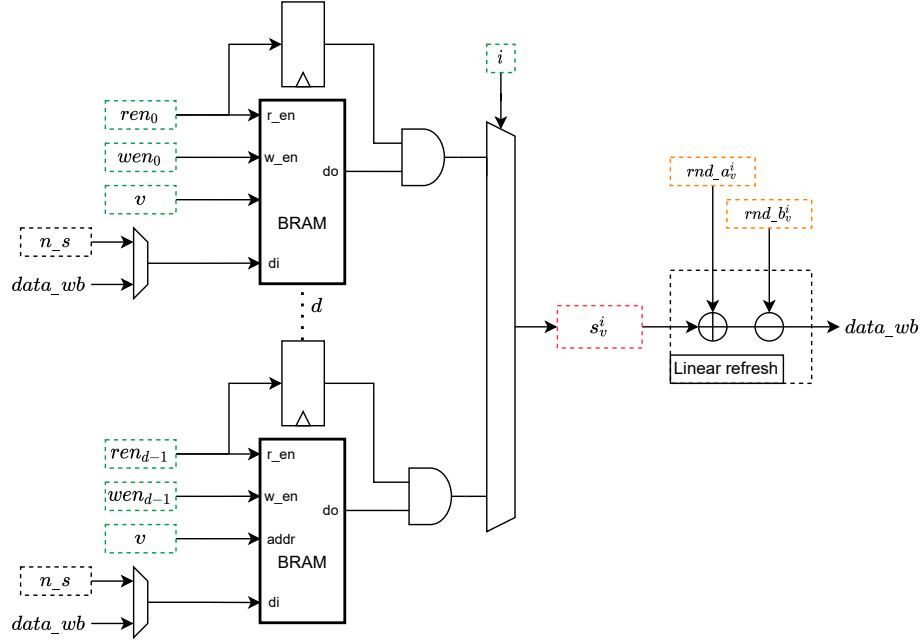

Fig. 4: Memory organization for the s_v^i storage.

After each utilization, the key shares have to be refreshed. This is done in parallel to the other computations by a dedicated module that implements the same linear refresh mechanism as the one used in [24] and [38]. Considering such a refresh mechanism, a total of dn^2/x fresh random elements over $\mathbb{F}_p$ are required when a single computation of u is considered. In the context of this work, we assumed that the $2y$ random elements required to refresh s_v^i in 1 clock cycle are provided to the core by external means through the two y elements wide buses $rnd\_a_v^i$ and $rnd\_b_v^i$. The refreshed data (denoted $data\_wb$ in Figure 4) are then routed back to the BRAMs to be written back to the memory. Properly setting the enable signal of each share BRAM write interface (denoted wen_i) ensures that only the memory block that has just been read is updated.

The FPGA implementation results of our architecture are shown in Table 1 for the parameters sets of Section 5.2. They were obtained with the Vivado v2020.1 toolset for a Xilinx Kintex7 FPGA and confirmed overheads that are linear in the number of shares d. Since based on similar or larger levels of parallelism as [38], these implementations are expected to provide similar or larger levels of security against higher-order DPA (and the security of the re-combined shares should hold for up to 2^{64} queries, as mentioned in Part C of Section 5.2).

Instance	y	d	LUTs [kLUTs]	Regs [kRegs]	B36	B18	DSPs	Latency [kcycles]	Freq. [MHz]
Set 1		2	57.18	63.74	4	0	128	65.538	27.4
Set 1	4	3	57.33	63.75	6	0	128	98.307	26.7
Set 1		4	57.60	63.76	8	0	128	131.076	26.5
Set 2		2	55.21	63.86	6	2	128	65.538	25.1
Set 2	8	3	55.60	63.87	9	3	128	98.307	25.1
Set 2		4	56.00	63.98	12	4	128	131.076	25.1
Set 3		2	27.90	33.05	4	0	64	49.154	34.23
Set 3	8	3	28.05	33.06	6	0	64	73.731	32.29
Set 3		4	28.19	33.07	8	0	64	98.308	28.47

Table 1: XC7K410T FPGA implementation results of the masked u computation. B36, B18 and DSPs correspond to physical resources RAMB36E1, RAMB18E1 and DSP48E1 that are embedded on Xilinx Serie7 FPGAs.

6 Conclusions

The uniform protection of all the operations in recent post-quantum CCA-secure public key encryption schemes against side-channel attacks is known to be very expensive. To the best of our knowledge, POLKA is the first scheme for which a protected implementation can be leveled, mixing operations that only require SPA security with a few operations that require DPA security but can be efficiently masked. We reach this goal by mixing various ideas which we believe of independent interest.[9] We also believe these techniques are generic and could be exploited for other schemes. For example, a leakage-resistant variant of the NTRU cryptosystem is discussed in Supplementary Material A.

Our results also lead to a number of interesting open problems. First, the RLWPR assumption on which a part of POLKA's physical security relies is an admittedly recent one. So further cryptanalysis (e.g., generalized to wide classes of realistic leakage functions) is an important scope for further research. The investigation of such hard physical learning problems in increasingly serial implementations is another promising direction, as it could lead to their exploitation in a software context. The same holds for a NTT-LWPR variant of RLWPR that would allow re-combining shares in the NTT domain, therefore leading to more efficient multiplications and, in general, for efforts towards a more unified / less specialized view of hard physical learning problems. More related to the high-level design of POLKA, it would be interesting to study options to further improve its potential for leveling (e.g., by removing the possibility of DPA against the hash function of step 5). From a theoretical viewpoint, evaluating whether

---

[9] For example, even a POLKA design that does not rely on the RLWPR assumption to unmask some computations would remain interesting compared to the NIST post-quantum finalists, thanks to the significant reduction of masked symmetric cryptographic primitives that dominate their performance figures [23].

post-quantum and leakage-resistant schemes could take advantage of ciphertext compression would be relevant as well. Eventually, the first leakage analysis we provide in this work is based on the heuristic (attack-based) approach of [15]. So formalizing and proving the leakage security of POLKA with an appropriate set of physical assumptions is a necessary long-term goal.

References

1. M. Abe, R. Gennaro, K. Kurosawa, and V. Shoup. Tag-KEM/DEM: A new framework for hybrid encryption and a new analysis of Kurosawa-Desmedt KEM. In *Eurocrypt*, 2005.
2. M. Albrecht, S. Bai, and L. Ducas. A subfield lattice attack on overstretched NTRU assumptions: Cryptanalysis of some FHE and graded encoding schemes. In *Crypto*, 2016.
3. M. Albrecht, R. Player, and S. Scott. On the concrete hardness of Learning with Errors. *Journal of Mathematical Cryptology*, 9(3), 2015.
4. E. Alkim, L. Ducas, T. Pöppelmann, and P. Schwabe. Post-quantum key exchange – a new hope. In *USENIX Security Symposium*, 2016.
5. R. Avanzi, J. Bos, L. Ducas, E. Kiltz, T. Lepoint, V. Lyubashevsky, J. Schanck, P. Schwabe, G. Seiler, and D. Damien Stehlé. CRYSTALS-KYBER algorithm specifications and supporting documentation. *NIST PQC Round*, 3:42, 2020.
6. M. Azouaoui, O. Bronchain, C. Hoffmann, Y. Kuzovkova, T. Schneider, and F. Standaert. Systematic study of decryption and re-encryption leakage: The case of kyber. In *COSADE*, 2022.
7. J. Balasch, B. Gierlichs, V. Grosso, O. Reparaz, and F. Standaert. On the cost of lazy engineering for masked software implementations. In *CARDIS*, 2014.
8. A. Banerjee, C. Peikert, and A. Rosen. Pseudorandom functions and lattices. In *EUROCRYPT*, 2012.
9. G. Barthe, S. Belaïd, F. Dupressoir, P. Fouque, B. Grégoire, P. Strub, and R. Zucchini. Strong non-interference and type-directed higher-order masking. In *CCS*, 2016.
10. G. Barthe, F. Dupressoir, S. Faust, B. Grégoire, F. Standaert, and P. Strub. Parallel implementations of masking schemes and the bounded moment leakage model. In *Eurocrypt (1)*, 2017.
11. A. Basso, J. M. B. Mera, J.-P. D'Anvers, A. Karmakar, S. S. Roy, M. V. Beirendonck, and F. Vercauteren. SABER algorithm specifications and supporting documentation. *NIST PQC Round*, 3:44, 2020.
12. M. V. Beirendonck, J. D'Anvers, A. Karmakar, J. Balasch, and I. Verbauwhede. A side-channel-resistant implementation of SABER. *ACM J. Emerg. Technol. Comput. Syst.*, 17(2), 2021.

13. S. Belaïd, J. Coron, P. Fouque, B. Gérard, J. Kammerer, and E. Prouff. Improved side-channel analysis of finite-field multiplication. In *CHES*, 2015.

14. S. Belaïd, P. Fouque, and B. Gérard. Side-channel analysis of multiplications in GF(2128) - application to AES-GCM. In *ASIACRYPT (2)*, 2014.

15. D. Bellizia, O. Bronchain, G. Cassiers, V. Grosso, C. Guo, C. Momin, O. Pereira, T. Peters, and F. Standaert. Mode-level vs. implementation-level physical security in symmetric cryptography - A practical guide through the leakage-resistance jungle. In *Crypto (1)*, 2020.

16. D. J. Bernstein and E. Persichetti. Towards KEM unification. Cryptology ePrint Archive, Report 2018/526, 2018.

17. F. Berti, S. Bhasin, J. Breier, X. Hou, R. Poussier, F. Standaert, and B. Udvarhelyi. A finer-grain analysis of the leakage (non) resilience of OCB. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2022(1), 2022.

18. S. Bhasin, J. D'Anvers, D. Heinz, T. Pöppelmann, and M. V. Beirendonck. Attacking and defending masked polynomial comparison for lattice-based cryptography. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2021(3), 2021.

19. D. Boneh, O. Dagdelen, M. Fischlin, A. Lehmann, C. Schaffner, and M. Zhandry. Random oracles in a quantum world. In *Asiacrypt*, 2011.

20. D. Boneh, Y. Ishai, A. Passelègue, A. Sahai, and D. J. Wu. Exploring crypto dark matter: - new simple PRF candidates and their applications. In *TCC (2)*, 2018.

21. J. Bos, L. Ducas, E. Kiltz, T. Lepoint, V. Lyubashevsky, J. Schanck, P. Schwabe, G. Seiler, and D. Stehlé. CRYSTALS - Kyber: A CCA-Secure Module-Lattice-Based KEM. In *IEEE European Symposium on Security and Privacy, EuroS&P*, 2018.

22. J. W. Bos, M. Gourjon, J. Renes, T. Schneider, and C. van Vredendaal. Masking KYBER: First- and higher-order implementations. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2021(4), 2021.

23. O. Bronchain and G. Cassiers. Bitslicing arithmetic/boolean masking conversions for fun and profit with application to lattice-based kems. Cryptology ePrint Archive, Report 2022/158, 2022.

24. O. Bronchain, T. Schneider, and F. Standaert. Reducing risks through simplicity: high side-channel security for lazy engineers. *J. Cryptogr. Eng.*, 11(1):39–55, 2021.

25. G. Cassiers, B. Grégoire, I. Levi, and F. Standaert. Hardware private circuits: From trivial composition to full verification. *IEEE Trans. Computers*, 70(10), 2021.

26. G. Cassiers and F. Standaert. Provably secure hardware masking in the transition- and glitch-robust probing model: Better safe than sorry. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2021(2):136–158, 2021.

27. C. Chen, O. Danba, J. Hoffstein, A. Hülsing, J. Rijneveld, J. Schanck, P. Schwabe, W. Whyte, Z. Zhang, T. Saito, T. Yamakawa, and K. Xagawa. NTRU algorithm specifications and supporting documentation. *NIST PQC Round*, 3:41, 2020.

28. J. Coron, C. Giraud, E. Prouff, S. Renner, M. Rivain, and P. K. Vadnala. Conversion of security proofs from one leakage model to another: A new issue. In *COSADE*, 2012.

29. J. Coron, E. Prouff, M. Rivain, and T. Roche. Higher-order side channel security and mask refreshing. In *FSE*, 2013.

30. J.-P. D'Anvers, Q. Guo, T. Johansson, A. Nilsson, F. Vercauteren, and I. Verbauwhede. Decryption failure attacks on IND-CCA secure lattice-based schemes. In *PKC*, 2019.

31. J.-P. D'Anvers, A. Karmakar, S.-S. Roy, and F. Vercauteren. Saber: Module-LWR based key exchange, CPA-secure encryption and CCA-secure KEM. In *Africacrypt*, 2018.

32. J.-P. D'Anvers, E. Orsini, and F. Vercauteren. Error term checking: Towards chosen ciphertext security without re-encryption. In *AsiaPKC*, 2021.

33. J.-P. D'Anvers, M. Rossi, and F. Virdia. (One) failure is not an option: Bootstrapping the search for failures in lattice-based encryption schemes. In *Eurocrypt*, 2020.

34. J. Ding, X. Xie, and X. Lin. A simple provably secure key exchange scheme based on the Learning With Errors problem. Cryptology ePrint Archive: Report 2012/688, 2014.

35. C. Dobraunig, F. Koeune, S. Mangard, F. Mendel, and F. Standaert. Towards fresh and hybrid re-keying schemes with beyond birthday security. In *CARDIS*, 2015.

36. A. Duc, S. Faust, and F. Standaert. Making masking security proofs concrete - or how to evaluate the security of any leaking device. In *EUROCRYPT (1)*, 2015.

37. J. Duman, K. Hövelmanns, E. Kiltz, V. Lyubashevsky, G. Seiler, and D. Unruh. A thorough treatment of highly-efficient NTRU instantiations. Cryptology ePrint Archive: Report 2021/1352.

38. S. Duval, P. Méaux, C. Momin, and F. Standaert. Exploring crypto-physical dark matter and learning with physical rounding towards secure and efficient fresh re-keying. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2021(1), 2021.

39. S. Dziembowski, S. Faust, G. Herold, A. Journault, D. Masny, and F. Standaert. Towards sound fresh re-keying with hard (physical) learning problems. In *Crypto (2)*, 2016.

40. P.-A. Fouque, P. Kirchner, T. Pornin, and Y. Yu. BAT: Small and fast KEM over NTRU lattices. In *CHES*, 2022.

41. T. Fritzmann, M. V. Beirendonck, D. B. Roy, P. Karl, T. Schamberger, I. Verbauwhede, and G. Sigl. Masked accelerators and instruction set extensions for post-quantum cryptography. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2022(1), 2022.

42. E. Fujisaki and T. Okamoto. Secure integration of asymmetric and symmetric encryption schemes. In *Crypto*, 1999.

43. E. Fujisaki and T. Okamoto. Secure integration of asymmetric and symmetric encryption schemes. *Jo. of Cryptology*, 226(21), 2013.

44. H. Gilbert, M. J. B. Robshaw, and Y. Seurin. Hb$^{\#}$: Increasing the security and efficiency of hb^{+}. In *EUROCRYPT*, 2008.

45. C. Guo, O. Pereira, T. Peters, and F. Standaert. Authenticated encryption with nonce misuse and physical leakage: Definitions, separation results and first construction - (extended abstract). In *LATINCRYPT*, 2019.

46. M. Hamburg. Three Bears. Technical report, National Institute of Standards and Technology, 2019.

47. J. Hoffstein, J. Pipher, and J. H. Silverman. NTRU: A ring-based public key cryptosystem. In *ANTS*, 1998.

48. D. Hofheinz, K. Hövelmanns, and E. Kiltz. A modular analysis of the Fujisaki-Okamoto transformation. In *TCC*, 2017.

49. D. Hofheinz and E. Kiltz. Secure hybrid encryption from weakened key encapsulation. In *Crypto*, 2007.

50. K. Hövelmanns, A. Hülsing, and C. Majenz. Failing gracefully: Decryption failures and the fujisaki-okamoto transform. *IACR Cryptol. ePrint Arch.*, page 365, 2022.

51. Y. Ishai, A. Sahai, and D. A. Wagner. Private circuits: Securing hardware against probing attacks. In *CRYPTO*, 2003.

52. S. Katsumata and S. Yamada. Partitioning via non-linear polynomial functions: More compact IBEs from ideal lattices and bilinear maps. In *Asiacrypt*, 2016.

53. E. Kiltz, K. Pietrzak, D. Venturi, D. Cash, and A. Jain. Efficient authentication from hard learning problems. *J. Cryptol.*, 30(4):1238–1275, 2017.

54. P. Kirchner and P.-A. Fouque. An improved BKW algorithm for LWE with applications to cryptography and lattices. In *Crypto*, 2015.

55. A. Langlois and D. Stehlé. Worst-case to average-case reductions for module lattices. *Designs, Code & Cryptography*, 2014.

56. A. Lopez-Alt, E. Tromer, and V. Vaikuntanathan. On-the-fly multiparty computation on the cloud via multikey fully homomorphic encryption. In *STOC*, 2012.

57. X. Lu, Y. Liu, D. Jia, H. Xue, J. He, Z. Zhang, Z. Liu, H. Yang, B. Li, and K. Wang. LAC. Technical report, National Institute of Standards and Technology, 2019.

58. V. Lyubashevsky, C. Peikert, and O. Regev. On ideal lattices and learning with errors over rings. In *Eurocrypt*, 2010.

59. S. Mangard, E. Oswald, and T. Popp. *Power analysis attacks - revealing the secrets of smart cards*. Springer, 2007.

60. S. Mangard, T. Popp, and B. M. Gammel. Side-channel leakage of masked CMOS gates. In *CT-RSA*, 2005.

61. M. Medwed, F. Standaert, J. Großschädl, and F. Regazzoni. Fresh re-keying: Security against side-channel and fault attacks for low-cost devices. In *Africacrypt*, 2010.

62. D. Micciancio and O. Regev. Worst-case to average-case reductions based on Gaussian measures. *SIAMJC*, 37(1):267–302, 2007.

63. K. Ngo, E. Dubrova, Q. Guo, and T. Johansson. A side-channel attack on a masked IND-CCA secure SABER KEM implementation. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2021(4), 2021.

64. S. Nikova, V. Rijmen, and M. Schläffer. Secure hardware implementation of nonlinear functions in the presence of glitches. *J. Cryptol.*, 24(2):292–321, 2011.

65. C. Peikert. Lattice cryptography for the Internet. In *PQCrypto*, 2014.

66. A. Pellet-Mary and S. D. On the hardness of the NTRU problem. In *Asiacrypt*, 2021.

67. E. Persichetti. *Improving the efficiency of code-based cryptography*. PhD thesis, Univ. of Auckland, 2012.

68. P. Ravi, S. S. Roy, A. Chattopadhyay, and S. Bhasin. Generic side-channel attacks on CCA-secure lattice-based PKE and KEMs. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2020(3), 2020.

69. O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In *STOC*, 2005.

70. O. Reparaz, R. de Clercq, S. S. Roy, F. Vercauteren, and I. Verbauwhede. Additively homomorphic ring-lwe masking. In T. Takagi, editor, *PQCrypto*, 2016.

71. O. Reparaz, S. S. Roy, R. de Clercq, F. Vercauteren, and I. Verbauwhede. Masking ring-lwe. *J. Cryptogr. Eng.*, 6(2):139–153, 2016.

72. O. Reparaz, S. S. Roy, F. Vercauteren, and I. Verbauwhede. A masked ring-lwe implementation. In T. Güneysu and H. Handschuh, editors, *CHES*, 2015.

73. T. Saito, K. Xagawa, and Y. Yamakawa. Tightly-secure key-encapsulation mechanism in the quantum random oracle model. In *Eurocrypt*, 2018.

74. V. Shoup. A proposal for an ISO standard for public key encryption. Manuscript, Dec. 2001.

75. D. Stehlé and R. Steinfeld. Making NTRU as secure as worst-case problems over ideal lattices. In *Eurocrypt*, 2011.

76. R. Steinfeld, S. Ling, J. Pieprzyk, C. Tartary, and H. Wang. NTRUCCA: How to strengthen NTRUEncrypt to chosen-ciphertext security in the standard model. In *PKC*, 2012.

77. R. Ueno, K. Xagawa, Y. Tanaka, A. Ito, J. Takahashi, and N. Homma. Curse of re-encryption: A generic power/EM analysis on post-quantum KEMs. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2022(1), 2022.

78. N. Veyrat-Charvillon, B. Gérard, and F. Standaert. Soft analytical side-channel attacks. In *ASIACRYPT (1)*, 2014.

79. N. Veyrat-Charvillon, M. Medwed, S. Kerckhof, and F. Standaert. Shuffling against side-channel attacks: A comprehensive study with cautionary note. In *ASIACRYPT*, 2012.

Supplementary Material

A   LR-NTRU: a Leakage-Resistant NTRU Variant

With slight changes, the re-encryption-free hybrid encryption scheme of [27, Section 2.3] can be made side-channel-resistant in the same way as our scheme of Section 4. However, if we set the parameters to rely on the standard RLWE assumption (and avoid the stronger Decision Small Polynomial Ratio assumption [56]) by applying the result of [75], it turns out to be significantly less efficient.

A.1   The Scheme With an Additive Mask

Keygen(1^λ)**:** Given a security parameter $\lambda \in \mathbb{N}$,

1. Choose a dimension $n \in \mathbb{N}$, which is a power of 2, and a prime modulus q such that $q = 1 \bmod 2n$, which define the rings $R = \mathbb{Z}[x]/(x^n + 1)$ and $R_q = R/(qR)$. We denote by $R_q^\times$ the set of units in R_q.
2. Choose a standard deviation $\sigma \in \mathbb{R}$ such that $\sigma \geq 2n\sqrt{\ln(8nq)}q^{1/2+2\varepsilon}$, a noise parameter $\alpha \in (0,1)$, and a modulus $p \in \mathbb{N}$ such that $p > 4\alpha q\sqrt{n}$ and $q > 8p(\alpha qn + 1)$. Let a norm bound $B = \alpha q\sqrt{n}$.
3. Sample $g \hookleftarrow D_{\mathbb{Z}^n,\sigma}^{\mathsf{coeff}}$. If $(g \bmod q) \notin R_q^\times$, restart step 3.
4. Sample $f' \hookleftarrow D_{\mathbb{Z}^n,\sigma}^{\mathsf{coeff}}$ and compute $f = pf'+1$. If $(f \bmod q) \notin R_q^\times$, restart step 4.
5. Choose an authenticated symmetric encryption scheme $\Pi^{sym} = (\mathsf{K}, \mathsf{E}, \mathsf{D})$ with key length $\kappa \in \mathsf{poly}(\lambda)$ and message space $\{0,1\}^{\ell_m}$.
6. Let a domain $D_E := \{(s,e) \in R^2 : \|s\|, \|e\| \leq B\}$. Choose a hash function $H : D_E \to \{0,1\}^\kappa$ that will be modeled a a random oracle.

   Return the key pair (PK, SK) where

   $$PK := \big(n,\ q,\ p,\ \alpha,\ h = p \cdot g/f \in R_q^\times,\ \Pi^{sym},\ H,\ B\big)$$

   and $SK := f \in R^\times$.

Encrypt(PK, M)**:** Given a public key PK and a message $M \in \{0,1\}^{\ell_m}$, do the following:

1. Sample $s, e \hookleftarrow D_{\mathbb{Z}^n,\alpha q}^{\mathsf{coeff}}$.
2. Compute $c_0 = h \cdot s + e \in R_q$ and $K = H(s,e) \in \{0,1\}^\kappa$.
3. Compute $c_1 = \mathsf{E}_K(M)$.

   Output the ciphertext $C = (c_0, c_1)$.

Decrypt(SK, C)**:** Given $SK = f \in R^\times$ and $C = (c_0, c_1)$, do the following.

1. Choose $s', e' \hookleftarrow D_{\mathbb{Z}^n,\alpha q}^{\mathsf{coeff}}$ and compute $d_0 = h \cdot s' + e' \in R_q$. Compute $\mu' = (c_0 + d_0) \cdot f \in R_q$. If $\|s'\| > B$ or $\|e'\| > B$, return $\perp$.
2. Compute $\bar{e} = \mu' \bmod p$ and $e = \bar{e} - e'$. Then, return $\perp$ if $\|\bar{e}\| > 2B$ or $\|e\| > B$.
3. Compute $s = (c_0 - e) \cdot h^{-1} \in R_q$ and return $\perp$ if $\|s\| > B$.
4. Compute $K = H(s,e) \in \{0,1\}^\kappa$.
5. Compute and output $M = \mathsf{D}_K(c_1) \in \{0,1\}^{\ell_m} \cup \{\perp\}$.

CORRECTNESS. In honestly generated ciphertexts, we have $c_0 = hs + e \in R_q$ for some $s, e \sim D_{\mathbb{Z}^n, \alpha q}^{\mathsf{coeff}}$. At step 1, the decryptor computes $d_0 = hs' + e' \in R_q$ and obtains

$$\mu' = p \cdot g(s + s') + (e + e')f = p \cdot \big(g(s + s') + (e + e')f'\big) + (e + e')$$

where the two equalities hold over R since, with probability $1 - 2^{-\Omega(n)}$ over the randomness of Keygen, Encrypt and Decrypt,

$$\|p \cdot \big(g(s + s') + (e + e')f'\big) + (e + e')\|_\infty < 4p\alpha q \sigma n + 2\alpha q\sqrt{n} \tag{4}$$

$$< 4p\alpha\sigma n + \frac{p}{2} < 4p(\alpha q \sigma n + 1) < q/2$$

Hence, step 2 recovers e with overwhelming probability since $p/2 > 2\alpha q\sqrt{n} \geq \|e + e'\|_\infty$. If c_0 is well-formed, Decrypt obtains s at step 3.

PARAMETERS. For correctness, we need to choose $\alpha \in (0, 1)$, q and p such that $p/2 > 2\alpha q\sqrt{n}$ and $q > 8p(\alpha q \sigma n + 1)$. In order to only rely on RLWE, we also need $\sigma \geq 2n\sqrt{\ln(8nq)}q^{1/2+2\varepsilon}$ with $\varepsilon > 4/\log q$ (in order to make the statistical distance (5) exponentially small). If we choose $\alpha q > \sqrt{n}$, we can set $q = \Theta(n^7 \log n)$, $\alpha^{-1} = \Theta(n^{6.5} \log n)$ and $\sigma = \Theta(n^{9/2} \log n)$.

DISCUSSION ON OTHER NTRU-BASED SCHEMES. Fouque *et al.* [40] recently suggested a CCA-secure KEM based on a randomness-recovering variant of NTRU where (analogously to the NTRU-CCA encryption scheme [76]) the use of a trapdoor basis allows getting rid of the masking modulus p, which in turn allows for a smaller q. While their construction can be made rigid and thus avoid the re-encryption step upon decapsulation, it would still remain less efficient than POLKA in a parameter regime allowing to rely on the standard RLWE assumption only. If we were to use a trapdoor basis to "invert" the NTRU function and recover $(s, e) \in R^2$ from $c_0 = (g/f) \cdot s + e$, applying the result of Stehlé and Steinfeld [75] would still require a modulus $q = \tilde{O}(n^6)$. Moreover, it requires significantly larger secret keys consisting of a full matrix over $R^{2\times 2}$. Their decryption algorithm also seems harder to protect against key-leakage as the secret matrix trapdoor is used on multiple occasions in the decryption algorithm [40, Section 3.2].

A.2 Security

In order to prove security under the standard RLWE assumption, we rely on the following lemma.

Lemma 9 ([75, Theorem 3]). *Let $n \geq 8$ a power of 2 and a prime $q \geq 5$ such that $\Phi(X) = X^n + 1$ splits into n linear factors over R_q. Let $0 < \varepsilon < 1$. Let $y_i \in R_q$, $z_i = -y_i p^{-1} \bmod q$ for $i \in \{1, 2\}$. If $\sigma \geq 2n\sqrt{\ln(8nq)} \cdot q^{1/2+2\varepsilon}$, then*

$$\Delta\left(\frac{y_i + p \cdot D_{\sigma, z_1}^\times}{y_2 + p \cdot D_{\sigma, z_2}^\times} \bmod q; U(R_q^\times)\right) \leq 2^{3n} \cdot q^{-\lfloor \varepsilon n\rfloor}. \tag{5}$$

The security proof is essentially identical to the proof of [73, Theorem 4.2], but we need to introduce slight changes since the underlying KEM uses explicit rejection in its decapsulation algorithm. Hence, we have to analyze the security of the hybrid PKE system as a whole (instead of focusing on the CCA security of the KEM) and first move to a game where the decryption algorithm proceeds with implicit rejection.

Theorem 2. *If Π^{sym} is a symmetric authenticated encryption scheme, the above construction provides IND-CCA security in the QROM under the* RLWE *assumption.*

Proof. The proof considers a sequence of games. In each game, we call W_i the event that $b' = b$. We also denote by $\mathsf{Encaps}(PK, (s, e))$ the deterministic algorithm that inputs PK and $(s, e) \in R^2$, and outputs $c_0 = h \cdot s + e$.

Game$_0$: This is the real IND-CCA experiment. All decryption queries are answered by running the real decryption algorithm. Note that a decryption query triggers a random oracle query at step 4 of Decrypt. In the challenge phase, the adversary $\mathcal{A}$ outputs two message M_0, M_1 and obtains a challenge ciphertext $C^\star = (c_0^\star, c_1^\star)$, where $c_0^\star = h \cdot s^\star + e^\star$, where $s^\star, e^\star \hookleftarrow D_{\mathbb{Z}^n, \alpha q}^{\mathsf{coeff}}$, and $c_1^\star = \mathsf{E}_{K^\star}(M_b)$ for some $b \hookleftarrow U(\{0, 1\})$. At the end of the game, the adversary $\mathcal{A}$ outputs $b' \in \{0, 1\}$. The adversary's advantage is defined as $\mathbf{Adv}(\mathcal{A}) := |\Pr[W_0] - 1/2|$.

Game$_1$: In this game, the challenger aborts and replaces $\mathcal{A}$'s output by a random bit in the event that $\|f\| > \sigma\sqrt{n}$ or $\|g\| > \sigma\sqrt{n}$ at step 4 of Keygen. By Lemma 1, we have $|\Pr[W_1] - \Pr[W_0]| \leq 2^{-\Omega(n)}$.

Game$_2$: In this game, we modify the decryption algorithm and use an implicit rejection mechanism in the KEM component. To this end, the challenger uses an independent random oracle $H_Q : R_q \rightarrow \{0, 1\}^\kappa$ that is only accessible to $\mathcal{A}$ via decryption queries (i.e., no direct access to H_Q is given to $\mathcal{A}$). This additional random oracle is used to run the following decryption algorithm.

> Decrypt$_2$: Given $SK = f$ and $C = (c_0, c_1)$, initialize a Boolean variable flag $= 0$. Then, do the following.
>
> 1. Choose $s', e' \hookleftarrow D_{\mathbb{Z}^n, \alpha q}^{\mathsf{coeff}}$. If $\|s'\| > B$ or $\|e'\| > B$, then set flag $= 1$ and return $\perp$. Otherwise, compute $d_0 = h \cdot s' + e' \in R_q$.
> 2. Compute $\mu' = (c_0 + d_0) \cdot f \in R_q$.
> 3. Compute $\bar{e} = \mu' \bmod p$ and $e = \bar{e} - e'$. Then, set flag $= 1$ if $\|\bar{e}\| > 2B$ or $\|e\| > B$.
> 4. Compute $s = (c_0 - e) \cdot h^{-1} \in R_q$ and set flag $= 1$ if $\|s\| > B$.
> 5. If flag $= 1$, compute $K = H_Q(c_0)$. Otherwise (i.e., if flag $= 0$), compute $K = H(s, e) \in \{0, 1\}^\kappa$.
> 6. Compute and output $M = \mathsf{D}_K(c_1) \in \{0, 1\}^{\ell_m} \cup \{\perp\}$.

Lemma 10 shows that, if the adversary can distinguish Game$_2$ from Game$_1$, we can turn it into an adversary against the ciphertext integrity of Π^{sym}.

Game₃: We now simulate $H : D_E \to \{0,1\}^\kappa$ as

$$H(s,e) = H'_Q\big(\mathsf{Encaps}(PK,(s,e))\big),$$

where $H'_Q : R_q \to \{0,1\}^\kappa$ is yet another random oracle that is not directly accessible to $\mathcal{A}$. In the computation of $C^\star = (c_0^\star, c_1^\star)$, the symmetric key $K^\star$ is similarly obtained as $K^\star = H'_Q\big(c_0^\star\big)$, where $c_0^\star = \mathsf{Encaps}(PK,(s^\star,e^\star))$. Lemma 11 shows that, from $\mathcal{A}$'s view, Game_3 is identical to Game_2, so that we have $\Pr[W_3] = \Pr[W_2]$.

Game₄: This game is identical to Game_3 except that the random oracle $H : D_E \to \{0,1\}^\kappa$ is now simulated as $H(s,e) = H_Q\big(\mathsf{Encaps}(PK,(s,e))\big)$, where $H_Q : R_q \to \{0,1\}^\kappa$ is the random oracle introduced in Game_2. In the challenge ciphertext $C^\star = (c_0^\star, c_1^\star)$, the symmetric key $K^\star$ is derived as $K^\star = H_Q\big(c_0^\star\big)$, where $c_0^\star = \mathsf{Encaps}(PK,(s^\star,e^\star))$. Lemma 12 shows that the two games are perfectly indistinguishable and $\Pr[W_4] = \Pr[W_3]$.

Game₅: This game is identical to Game_4 except that we modify the decryption oracle. At each query $C = (c_0, c_1)$, the oracle computes $K = H_Q(c_0)$ and returns $M = \mathsf{D}_K(c_1) \in \{0,1\}^{\ell_m} \cup \{\bot\}$ (i.e., it ignores steps 2-4 of $\mathsf{Decrypt}_2$ and jumps from step 1 to step 5 after having set $\mathsf{flag} = 1$). Lemma 13 shows that $\Pr[W_5] = \Pr[W_4]$.

Game₆: This game is like Game_5 but we remove the change introduced in Game_1. That is, we do not replace $\mathcal{A}$'s output by a random bit if $\|f\| > \sigma\sqrt{n}$ or $\|g\| > \sigma\sqrt{n}$ in Keygen. By Lemma 1, $|\Pr[W_6] - \Pr[W_5]| \le 2^{-\Omega(n)}$.

Game₇: We change the generation of PK. In Game_6, the decryption oracle does not use the secret f at any time. We thus replace $h = pg/f$ by a uniformly random $h \hookleftarrow U(R_q^\times)$ at the beginning of the game. By Lemma 9, the distribution of h is statistically close to that of Game_6. We have the inequality $|\Pr[W_7] - \Pr[W_6]| \le 2^{-\lambda}$.

Game₈: We change the generation of the challenge ciphertext $C^\star = (c_0^\star, c_1^\star)$ and sample $c_0^\star \hookleftarrow U(R_q)$ uniformly instead of computing $c_0^\star = hs^\star + e^\star$. It is straightforward that Game_8 is indistinguishable from Game_7 under the $\mathsf{RLWE}^\times$ assumption.[10]

Game₉: We modify the decryption oracle that now rejects all ciphertexts of the form $C = (c_0^\star, c_1)$ with $c_1 \ne c_1^\star$ after the challenge phase. Game_9 is identical to Game_8 until the event E_9 that $\mathcal{A}$ queries the decryption of a ciphertext $C = (c_0^\star, c_1)$ that would not have been rejected in Game_8. Since $c_1 = \mathsf{E}_{K^\star}(M_d)$ is encrypted under a random key $K^\star = H_Q(c_0^\star)$ that is independent of $\mathcal{A}$'s view (unless the random $c_0^\star \sim U(R_q)$ is accidentally of the form $c_0^\star = hs + e$ for small $s, e \in R$), event E_9 would imply an attack

---

[10] This assumption says that, for a fixed secret $s \hookleftarrow \chi$ sampled from a distribution $\chi = D_{\mathbb{Z}^n,\alpha q}^{\mathsf{coeff}}$, the distribution of $(a, a{\cdot}s{+}e)$ is indistinguishable from $U(R_q^\times \times R_q)$ when a is sampled from $U(R_q^\times)$ (instead of $U(R_q)$) and $e \hookleftarrow \chi$. Since a random element of R_q is invertible with non-negligible probability $1 - n/q$ when $q = 1 \bmod 2n$, this assumption is equivalent to the standard RLWE where $a \sim U(R_q)$.

against the ciphertext integrity of Π^{sym}. We have $|\Pr[W_9] - \Pr[W_8]| \leq \Pr[E_9] \leq 2^{-\Omega(n)} + Q \cdot \mathbf{Adv}^{\mathsf{AE\text{-}INT}}(\lambda)$, where Q is the number of decryption queries.

In Game_9, the challenge $C^\star = (c_0^\star, c_1^\star)$ is obtained by encrypting $c_1^\star$ under a random key K which is never used anywhere, except in the computation of $c_1^\star = \mathsf{E}_{K^\star}(M_d)$. At this point, the adversary is essentially an adversary against the indistinguishability (under passive attacks) of the authenticated encryption scheme Π^{sym}. We have $|\Pr[W_9] - 1/2| \leq \mathbf{Adv}^{\mathsf{AE\text{-}IND}}(\lambda)$. $\square$

Lemma 10. Game_2 *is indistinguishable from* Game_1 *if* Π^{sym} *is a secure authenticated encryption scheme. We have* $|\Pr[W_2] - \Pr[W_1]| \leq \frac{Q \cdot (Q+1)}{2} \cdot \mathbf{Adv}^{\mathsf{AE\text{-}INT}}(\lambda)$.

Proof. The distinguishing advantage of $\mathcal{A}$ between the two games can be bounded by the difference between the probabilities that a ciphertext gets rejected in Game_2 and Game_1.

The difference between the two decryption oracles is that, when an invalid pair c_0 is detected at steps 3-4, $\mathsf{Decrypt}_2$ sets $\mathsf{flag} = 1$ and proceeds until step 5 whereas the decryption oracle of Game_1 would directly return $\bot$.

In Game_2, let us assume that $\mathcal{A}$ queries a ciphertext (c_0, c_1) for which $\mathsf{Decrypt}_2$ significantly deviates from the decryption oracle of Game_1 in its probability to reject. This means that, in Game_2, we have $\mathsf{flag} = 1$ at step 5, so that $\mathsf{Decrypt}_2$ computes $K = H_Q(c_0)$. Since $\mathcal{A}$ has no direct access to H_Q, K is independent of $\mathcal{A}$'s view and uniformly distributed over $\{0, 1\}^\kappa$. If $\mathsf{Decrypt}_2$ does not return $\bot$ at step 6, we have $\mathsf{D}_K(c_1) \neq \bot$, which means that $\mathcal{A}$ was able to forge a valid ciphertext for a random key K independent of its view. Then, the distance $|\Pr[W_2] - \Pr[W_1]|$ can be bounded by $\frac{Q(Q+1)}{2} \cdot \mathbf{Adv}^{\mathsf{AE\text{-}INT}}(\lambda)$, where Q is the number of decryption queries and $\mathbf{Adv}^{\mathsf{AE\text{-}INT}}(\lambda)$ is the reduction's advantage against the ciphertext integrity of Π^{sym}. The reduction $\mathcal{B}^{\mathsf{AE}}$ is identical to that of Lemma 3. $\square$

Lemma 11. *If* $p > 4\alpha q \sqrt{n}$ *and* $q > 8p(\alpha \sigma n + 1)$, Game_3 *is perfectly indistinguishable from* Game_2.

Proof. We show that, from the adversary's view, the random oracle H of Game_3 behaves identically to that of Game_2.

Indeed, the function $\mathsf{Encaps} : D_E \rightarrow R_q : (s, e) \rightarrow \mathsf{Encaps}(PK, (s, e))$ is injective over $D_E := \{(s, e) \in R^2 : \|s\|, \|e\| \leq B\}$ as there cannot exist colliding pairs $(s_0, e_0), (s_1, e_1) \in D_E$. The equality $hs_0 + e_0 = hs_1 + e_1$ over R_q implies

$$p \cdot \big(gs_0 + e_0 f'\big) + e_0 = p \cdot \big(gs_1 + e_1 f'\big) + e_1,$$

which holds over R since both members are polynomials with coefficients smaller than $q/2$ in magnitude. This yields $e_0 = e_1$ (and thus $s_0 = s_1$ since $h \in R_q^\times$) after reduction modulo p since $p/2 > 2B$ and $\|e_0\|_\infty, \|e_1\|_\infty \leq B$.

Since Encaps is injective over D_E and H_Q is a random function, the composed function $H_Q\big(\mathsf{Encaps}(PK, (\cdot, \cdot, \cdot))\big)$ is itself random over D_E. The two games are thus identical from $\mathcal{A}$'s view. $\square$

Lemma 12. Game_4 *is perfectly indistinguishable from* Game_3.

Proof. We assume that the event introduced in Game_1 does not occur (and thus $\|g\|, \|f\| \leq \sigma\sqrt{n}$) since both games have identical output distributions otherwise.

We call a ciphertext (c_0, c_1) *good* if $c_0 = h \cdot s + e \in R_q$ for some $(s, e) \in R^3$ such that $\|s\|, \|e\| \leq B$. For a good ciphertext, if the decryption oracle of Game_3 sets $\mathsf{flag} = 1$ at all, the ciphertext is rejected at step 1 of $\mathsf{Decrypt}_2$: Indeed, if (c_0, c_1) is good, $\mathsf{Decrypt}_2$ cannot set $\mathsf{flag} = 1$ for the first time at step 3, nor at step 4. Hence, if a ciphertext is good and $\|s + s'\| > 2B$ or $\|e + e'\| > 2B$, it gets rejected as step 1. This implies that, if a good ciphertext is not rejected at step 1, we have $\mathsf{flag} = 0$ at step 5 because inequalities (4) hold whenever $\|f\|, \|g\| \leq \sigma\sqrt{n}$ and $\|\bar{s}\|, \|\bar{e}\| \leq 2B$.

As a result, in Game_3, H_Q is never evaluated on a good ciphertext. Also, if $\mathsf{flag} = 0$ at step 5, we are guaranteed that (c_0, c_1) is good (by the rigidity of the scheme), meaning that H'_Q is only evaluated on good ciphertexts. Since H_Q and H'_Q are evaluated on disjoint sub-domains in Game_3, $\mathcal{A}$'s view is the same as in Game_4, where they were simulated using a single random oracle. $\qquad\square$

Lemma 13. *We have* $\Pr[W_5] = \Pr[W_4]$ *as* Game_5 *is perfectly indistinguishable from* Game_4.

Proof. The two games are identical from $\mathcal{A}$'s view until the decryption oracle of Game_5 deviates from its counterpart in Game_4.

For any decryption query $C = (c_0, c_1)$ such that the decryption oracle of Game_4 sets $\mathsf{flag} = 1$, it always computes $K = H_Q(c_0)$ at step 5 of $\mathsf{Decrypt}_2$.

For ciphertexts $C = (c_0, c_1)$ such that $\mathsf{flag} = 0$ at the end of step 4 in Game_4, the rigidity of the encapsulation mechanism ensures that $\mathsf{Decrypt}_2$ computes $(s, e) \in D_E$ satisfying $c_0 = h \cdot s + e$, meaning that $c_0 = \mathsf{Encaps}(PK, (s, e))$. In this case, at step 5 of $\mathsf{Decrypt}_2$, the decryption oracle of Game_4 computes $K = H(s, e) = H_Q(c_0)$, which would also be the output of the decryption oracle in Game_5.

In both cases $\mathsf{flag} \in \{0, 1\}$, the two decryption oracles always output the same result. $\qquad\square$

B   Authenticated Symmetric Encryption

A symmetric encryption scheme is specified by a pair $(\mathsf{K}, \mathsf{E}, \mathsf{D})$, where E is the encryption algorithm; D is the decryption procedure; and K is the key generation algorithm. The security of authenticated symmetric encryption is defined by means of two games that capture the ciphertext indistinguishability and ciphertext (one-time) integrity properties.

Definition 4. *A symmetric encryption scheme is secure in the sense of authenticated encryption if it provides both indistinguishability and ciphertext integrity, as defined hereunder.*

1. **Indistinguishability.** No PPT adversary $\mathcal{A}$ has non-negligible advantage in the following game, where $\lambda \in \mathbb{N}$ is a security parameter:

$$\boxed{\mathbf{Game}_{\mathcal{A}}^{\mathsf{AE\text{-}IND}}(\lambda)}$$
$$K \leftarrow \mathsf{K}(1^\lambda)$$
$$(m_0, m_1, st) \leftarrow \mathcal{A}(\mathsf{find}, 1^\lambda)$$
$$d^\star \leftarrow U(\{0,1\})$$
$$c^\star \leftarrow \mathsf{E}_K(m_{d^\star})$$
$$d \leftarrow \mathcal{A}(\mathsf{guess}, st, c^\star)$$
$$\texttt{return 1 if } d = d^\star \texttt{ and 0 otherwise.}$$

$\mathcal{A}$'s advantage is $\mathbf{Adv}_{\mathcal{A}}^{\mathsf{AE\text{-}IND}}(\lambda) = |\Pr[\mathbf{Game}_{\mathcal{A}}^{\mathsf{AE\text{-}IND}} = 1] - 1/2|$.

2. **Ciphertext integrity.** No PPT adversary $\mathcal{A}$ has non-negligible advantage in the following game:

$$\boxed{\mathbf{Game}_{\mathcal{A}}^{\mathsf{AE\text{-}INT}}(\lambda)}$$
$$K \leftarrow \mathsf{K}(1^\lambda)$$
$$(m, st) \leftarrow \mathcal{A}(\mathsf{find}, 1^\lambda)$$
$$c \leftarrow \mathsf{E}_K(m)$$
$$c' \leftarrow \mathcal{A}^{\mathsf{D}(K, \cdot)}(\mathsf{create}, st, c)$$
$$\texttt{return 1 if } c' \neq c \texttt{ and } \mathsf{D}_K(c') \neq \perp$$
$$\texttt{0 otherwise.}$$

Here, $\mathsf{D}(K, \cdot)$ is a decryption oracle that takes as input a candidate ciphertext c and returns $\mathsf{D}_K(c) \in \{0,1\}^{\ell_m} \cup \{\perp\}$.

$\mathcal{A}$'s advantage is now defined as $\mathbf{Adv}_{\mathcal{A}}^{\mathsf{AE\text{-}INT}}(\lambda) = \Pr[\mathbf{Game}_{\mathcal{A}}^{\mathsf{AE\text{-}INT}} = 1]$.

We note that the above definition captures *one-time* flavors of indistinguishability and ciphertext integrity as the adversary is given a single ciphertext in both cases. Both security notions can thus be achieved efficiently from a pseudorandom generator.

The ciphertext integrity property can even be achieved without any assumption if Π^{sym} is instantiated using an information-theoretically secure one-time MAC via the encrypt-then-MAC technique, as suggested in Section 4.3.

C   Adapting the Scheme to Binomial Errors

In this section, we consider an optimized instantiation of the scheme where the discrete Gaussian distribution is replaced by an easier-to-sample binomial distribution, as suggested in [4, 21, 37]. Namely, let ψ_k^n the distribution over $\mathbb{Z}^n$ defined as $\{\sum_{i=1}^{k}(a_i - b_i) \mid a_i, b_i \leftarrow U(\{0,1\}^n)\}$. Let $\bar{\psi}_2^n$ the distribution obtained by reducing $\psi_2^n \bmod 3$, where ψ_k^n is defined over $\mathbb{Z}^n$ as the distribution $\{\sum_{i=1}^{k}(a_i - b_i) \mid a_i, b_i \leftarrow U(\{0,1\}^n)\}$.

C.1 Description

Keygen(1^λ)**:** Given a security parameter $\lambda \in \mathbb{N}$,

1. Choose a dimension $n \in \mathbb{N}$, a prime modulus $q = 1 \bmod 2n$. Let the rings $R = \mathbb{Z}[X]/(X^n + 1)$ and $R_q = R/(qR)$ such that $\Phi(X) = X^n + 1$ splits into linear factors over R_q. Let $R_q^\times$ the set of units in R_q.
2. Choose an integer $p \in \mathbb{N}$ such that $4 < p < q/8(n+1)$.
3. Sample $a \hookleftarrow U(R_q)$ and $s, e \hookleftarrow \bar{\psi}_2^n$ and compute $b = p \cdot (a \cdot s + e)$. If $b \notin R_q^\times$, restart step 3.
4. Choose an authenticated symmetric encryption scheme $\Pi^{sym} = (\mathsf{K}, \mathsf{E}, \mathsf{D})$ with key length $\kappa \in \mathsf{poly}(\lambda)$ and message space $\{0,1\}^{\ell_m}$.
5. Let a domain $D_E := \{(r, e_1, e_2) \in R^3 : \|r\|_\infty, \|e_1\|_\infty, \|e_2\|_\infty \leq 1\}$. Choose a hash function $H : D_E \to \{0,1\}^\kappa$.

Return the key pair (PK, SK) where

$$PK := \left(n, \; q, \; p, \; a \in R_q, \; b \in R_q^\times, \; \Pi^{sym}, \; H\right)$$

and $SK := s \in R$.

Encrypt(PK, M)**:** Given a public key PK and a message $M \in \{0,1\}^{\ell_m}$:

1. Sample $r, e_1, e_2 \hookleftarrow \bar{\psi}_2^n$ and compute

$$c_1 = a \cdot r + e_1 \in R_q, \qquad c_2 = b \cdot r + e_2 \in R_q$$

together with $K = H(r, e_1, e_2) \in \{0,1\}^\kappa$.
2. Compute $c_0 = \mathsf{E}_K(M)$.

Output the ciphertext $C = (c_0, c_1, c_2)$.

Decrypt(SK, C)**:** Given $SK = s \in R$ and $C = (c_0, c_1, c_2)$, do the following.

1. Sample $r', e_1', e_2' \hookleftarrow \bar{\psi}_2^n$. Compute $c_1' = a \cdot r' + e_1'$ and $c_2' = b \cdot r' + e_2'$.
2. Compute $\bar{c}_1 = c_1 + c_1'$ and $\bar{c}_2 = c_2 + c_2'$.
3. Compute $\bar{\mu} = \bar{c}_2 - p \cdot \bar{c}_1 \cdot s$ over R_q.
4. Compute $\bar{e}_2 = \bar{\mu} \bmod p$. If $\|\bar{e}_2\|_\infty > 2$, return $\bot$.
5. Compute $\bar{r} = (\bar{c}_2 - \bar{e}_2) \cdot b^{-1} \in R_q$. If $\|\bar{r}\|_\infty > 2$, return $\bot$.
6. Compute $\bar{e}_1 = \bar{c}_1 - a \cdot \bar{r} \in R_q$. If $\|\bar{e}_1\|_\infty > 2$, return $\bot$.
7. Compute $r = \bar{r} - r'$, $e_1 = \bar{e}_1 - e_1'$ and $e_2 = \bar{e}_2 - e_2'$. If $\|r\|_\infty > 1$, or $\|e_1\|_\infty > 1$, or $\|e_2\|_\infty > 1$, then return $\bot$.
8. Compute $K = H(r, e_1, e_2) \in \{0,1\}^\kappa$ and return

$$M = \mathsf{D}_K(c_0) \in \{0,1\}^{\ell_m} \cup \{\bot\}.$$

As a side effect of the noise distributions $\chi = \bar{\psi}_2^n$, we do not have decryption failures anymore.

CORRECTNESS. Let $\bar{r} = r + r'$, $\bar{e}_1 = e_1 + e_1'$ and $\bar{e}_2 = e_2 + e_2'$ over R. At step 2, Decrypt computes $\bar{c}_1 = a \cdot \bar{r} + \bar{e}_1$, $\bar{c}_2 = b \cdot \bar{r} + \bar{e}_2$ over R_q, where $\|\bar{r}\|_\infty, \|\bar{e}_1\|_\infty, \|\bar{e}_2\|_\infty \leq 2$. At step 3, the receiver obtains

$$
\begin{aligned}
\bar{\mu} &= \bar{c}_2 - p \cdot \bar{c}_1 \cdot s \bmod q \\
&= (b \cdot \bar{r} + \bar{e}_2) - p \cdot (a \cdot \bar{r} + \bar{e}_1) \cdot s \bmod q \\
&= p \cdot (as + e) \cdot \bar{r} + \bar{e}_2 - p \cdot a\bar{r}s - p \cdot \bar{e}_1 s \bmod q \\
&= \bar{e}_2 + p \cdot e\bar{r} - p \cdot \bar{e}_1 s
\end{aligned}
$$

where the last equality holds over R with probability 1. Indeed, $\|s\|, \|e\| \leq \sqrt{n}$ with probability 1. We also have $\|\bar{r}\|, \|\bar{e}_1\|, \|\bar{e}_2\| \leq 2\sqrt{n}$. Then, the Cauchy-Schwartz inequality implies

$$
\|\bar{e}_2 + p \cdot e\bar{r} - p \cdot \bar{e}_1 s\|_\infty < 2 + 4np < q/2
$$

Since $p > 4$, step 4 recovers $\bar{e}_2$ with overwhelming probability. Since $b \in R_q^\times$, Decrypt obtains $\bar{r}$ at step 5 and $\bar{e}_1$ at step 6. Therefore, it also recovers (r, e_1, e_2) at step 7 and the correct symmetric key $K = H(r, e_1, e_2)$ at step 8. Correctness thus follows from the correctness of Π^{sym}.

CONCRETE INSTANTIATIONS. To ensure correctness and satisfy the conditions of Lemma 14 and Lemma 17, we need $p > 4$, $q > \max(8p(n+1), n^{3/2})$.

As a concrete instantiation validated by the LWE estimator [3], we may choose $n = 1024$, $p = 5$ and $q = 59393$, so that the size of (c_1, c_2) amounts to 4Kb.

Using the rings suggested in [37], we can relax the constraint of n being a power of 2. If we consider $R = \mathbb{Z}[X]/(\Phi_{3n})$ for the cyclotomic polynomial $\Phi_{3n} = X^n - X^{n/2} + 1$ with $q = 1 \bmod 3n$ and $n = 2^i 3^j$, Φ_{3n} fully splits over $R_q = R/(qR)$. Then, if we set $n = 768$, $p = 5$ (so that $2 \cdot \|e_2\|_\infty \leq (p-1)/2$), and $q = 28 \cdot 3n + 1 = 64513$, correctness (perfectly) holds since $\|\bar{e}_2 + p \cdot (e \cdot \bar{r} - \bar{e}_1 \cdot s)\|_\infty \leq (p-1)/2 + 8 \cdot p \cdot B^2 \leq (q-1)/2$. The size of (c_1, c_2) then drops to 3Kb.

C.2 Security Proof

The security proof is slightly shorter than the proof of Theorem 1 since we do not have to take imperfect correctness into account.

Theorem 3. *If Π^{sym} is a symmetric authenticated encryption scheme, the above construction provides IND-CCA security in the QROM under the* RLWE$_{\bar{\psi}_2^n}$ *assumption.*

Proof. The proof is very close to the proof of Theorem 1 and we only highlight the changes.

Game$_0$: This is the real game at the end of which the adversary $\mathcal{A}$ ouputs $d' \in \{0, 1\}$ and wins if $d' = d$. We call W_0 this event, so that the adversary has advantage $\mathbf{Adv}(\mathcal{A}) := |\Pr[W_0] - 1/2|$.

Game$_1$: We modify the decryption algorithm. Throughout the game, the challenger uses an independent random oracle $H_Q : R_q^2 \to \{0,1\}^\kappa$ that is only accessible to $\mathcal{A}$ via decryption queries (i.e., $\mathcal{A}$ has no direct access to H_Q). This random oracle is used to run the following decryption algorithm.

Decrypt$_2$: Given $SK = s$ and $C = (c_0, c_1, c_2)$, initialize a Boolean variable flag $= 0$. Then, do the following.

1. Sample $r', e_1', e_2' \hookleftarrow \bar{\psi}_2^n$. Compute $c_1' = a \cdot r' + e_1'$ and $c_2' = b \cdot r' + e_2'$.
2. Compute $\bar{c}_1 = c_1 + c_1'$ and $\bar{c}_2 = c_2 + c_2'$.
3. Compute $\bar{\mu} = \bar{c}_2 - p \cdot \bar{c}_1 \cdot s$ over R_q.
4. Compute $\bar{e}_2 = \bar{\mu} \bmod p$. If $\|\bar{e}_2\|_\infty > 2$, set flag $= 1$.
5. Compute $\bar{r} = (\bar{c}_2 - \bar{e}_2) \cdot b^{-1} \in R_q$. If $\|\bar{r}\|_\infty > 2$, set flag $= 1$.
6. Compute $\bar{e}_1 = \bar{c}_1 - a \cdot \bar{r} \in R_q$. If $\|\bar{e}_1\|_\infty > 2$, set flag $= 1$.
7. Compute $r = \bar{r} - r'$, $e_1 = \bar{e}_1 - e_1'$ and $e_2 = \bar{e}_2 - e_2'$. If $\|r\|_\infty > 1$, or $\|e_1\|_\infty > 1$, or $\|e_2\|_\infty > 1$, then set flag $= 1$.
8. If flag $= 0$, compute $K = H(r, e_1, e_2) \in \{0,1\}^\kappa$. Otherwise, compute $K = H_Q(c_1, c_2)$.
9. Compute and return $M = \mathsf{D}_K(c_0) \in \{0,1\}^{\ell_m} \cup \{\bot\}$.

By repeating exactly the same proof as in Lemma 3, we can show that, if the adversary can distinguish Game$_1$ from Game$_0$, we can turn it into an adversary against the ciphertext integrity of Π^{sym} (as defined in Supplementary Material B).

Game$_2$: We now simulate the random oracle[11] $H : D_E \to \{0,1\}^\kappa$ as

$$H(r, e_1, e_2) = H_Q'\big(\mathsf{Encaps}(PK, (r, e_1, e_2))\big) \qquad (6)$$

where $H_Q' : R_q^2 \to \{0,1\}^\kappa$ is an independent random oracle to which $\mathcal{A}$ is *not* given access. At each decryption query, Decrypt$_2$ computes K as per (6) when flag $= 0$. In the generation of $C^\star = (c_0^\star, c_1^\star, c_2^\star)$, $K^\star$ is similarly obtained as $K^\star = H_Q'\big(c_1^\star, c_2^\star\big)$, where $(c_1^\star, c_2^\star) = \mathsf{Encaps}(PK, (r^\star, e_1^\star, e_2^\star))$. Lemma 14 shows that, from $\mathcal{A}$'s view, Game$_3$ is identical to Game$_2$, so that $\Pr[W_3] = \Pr[W_2]$.

Game$_3$: This game is like Game$_2$ but the random oralce H is now simulated as $H(r, e_1, e_2) = H_Q\big(\mathsf{Encaps}(PK, (r, e_1, e_2))\big)$, where $H_Q : R_q^2 \to \{0,1\}^\kappa$ is the random oracle introduced in Game$_1$. In the computation of the challenge $C^\star = (c_0^\star, c_1^\star, c_2^\star)$, the key $K^\star$ is now obtained as $K^\star = H_Q\big(c_1^\star, c_2^\star\big)$, where $(c_1^\star, c_2^\star) = \mathsf{Encaps}(PK, (r^\star, e_1^\star, e_2^\star))$, and K is computed in the same way when flag $= 0$ at step 8 of Decrypt$_2$. In short, Game$_3$ is identical to Game$_2$ except that H_Q' has been replaced by H_Q in the simulation of H. Lemma 15 shows that $\Pr[W_3] = \Pr[W_2]$.

Game$_4$: This game is like Game$_3$ except that we modify the decryption oracle. At each query $C = (c_0, c_1, c_2)$, if flag $= 0$ at the end of step 1, then the decryption oracle computes $K = H_Q(c_1, c_2)$ and returns $M = \mathsf{D}_K(c_0) \in \{0,1\}^{\ell_m} \cup \{\bot\}$ (i.e., it ignores steps 2-7 of Decrypt$_2$ and directly moves to step 8). Lemma 16 shows that $\Pr[W_4] = \Pr[W_3]$.

---

[11] We may assume that H outputs $\bot$ on input of a triple $(r, e_1, e_2) \notin D_E$.

In Game_6, we note that the decryption oracle does not use the secret s anymore.

Game_5: We modify the generation of PK. The challenger initially samples $a_1, \ldots, a_k \hookleftarrow U(R_q)$, $e_1, \ldots, e_k \hookleftarrow \bar{\psi}_2^n$, where $k = \lceil \lambda/\log n \rceil$, and computes $b_i = a_i \cdot s + e_i$ for each $i \in [k]$. If none of the obtained $\{b_i\}_{i=1}^k$ is in $R_q^\times$, the challenger aborts and replaces $\mathcal{B}$'s output by a random bit. Otherwise, it determines the first index $i \in [k]$ such that $b_i \in R_q^\times$ and defines PK by setting $a = a_i$ and $b = p \cdot b_i$. Lemma 17 shows that, under the RLWE assumption, this change does not affect $\mathcal{A}$'s view and $|\Pr[W_5] - \Pr[W_4]| \le \mathbf{Adv}^{\mathsf{RLWE}}(\lambda) + 2^{-\lambda}$.

Game_6: We change again the generation of the public key. We replace the pseudorandom ring elements $\{b_i = a_i \cdot s + e_i\}_{i=1}^k$ of Game_5 by truly random $b_1, \ldots, b_k \hookleftarrow U(R_q)$ at the beginning of the game. Under the RLWE assumption, this change goes unnoticed and a straightforward reduction shows that $|\Pr[W_6] - \Pr[W_5]| \le \mathbf{Adv}^{\mathsf{RLWE}}(\lambda)$. As a result, since $\gcd(p, q) = 1$, the public key is now distributed so that $a \sim U(R_q)$ and $b \sim U(R_q^\times)$.

Game_7: We change the challenge ciphertext $C^\star = (c_0^\star, c_1^\star, c_2^\star)$. Instead of computing $c_1^\star = a \cdot r^\star + e_1^\star$, $c_2^\star = b \cdot r^\star + e_2^\star$ with $r^\star, e_1^\star, e_2^\star \hookleftarrow \bar{\psi}_2^n$, we now sample $c_1^\star, c_2^\star \hookleftarrow U(R_q)$ uniformly. Then, $c_0^\star$ is computed as a symmetric encryption of M_d under the key $K^\star = H_Q(c_1^\star, c_2^\star)$. Lemma 18 shows that Game_7 is indistinguishable from Game_6 under the RLWE assumption.

In Game_7, $\mathcal{A}$ can no longer query H on short $(r^\star, e_1^\star, e_2^\star)$ that underlie $(c_1^\star, c_2^\star)$. With probability $1 - 2^{-\Omega(n)}$, there exist no $r^\star, e_1^\star, e_2^\star \in R$ of infinity norm ≤ 1 such that $c_1^\star = a \cdot r^\star + e_1^\star$ and $c_2^\star = b \cdot r^\star + e_2^\star$. Since $\mathcal{A}$ has no direct access to $H_Q(\cdot)$, $H_Q(c_1^\star, c_2^\star)$ is henceforth independent of $\mathcal{A}$'s view.

Game_8: The decryption oracle now rejects all post-callenge queries of the form $C = (c_0, c_1^\star, c_2^\star)$ with $c_0 \ne c_0^\star$. Game_8 is identical to Game_7 until the event E_8 that $\mathcal{A}$ queries the decryption of a ciphertext $C = (c_0, c_1^\star, c_2^\star)$ that would not have been rejected in Game_7. Since $c_0^\star = \mathsf{E}_{K^\star}(M_d)$ is encrypted under a random key $K^\star = H_Q(c_1^\star, c_2^\star)$ independent of $\mathcal{A}$'s view, E_8 would break the ciphertext integrity of Π^{sym}. We have $|\Pr[W_8] - \Pr[W_7]| \le \Pr[E_8] \le 2^{-\Omega(n)} + \mathbf{Adv}^{\mathsf{AE\text{-}INT}}(\lambda)$, where Q is the number of decryption queries.

In Game_8, the challenge $C^\star = (c_0^\star, c_1^\star, c_2^\star)$ is obtained by encrypting $c_0^\star$ under a random key $K^\star$ which never shows up anywhere, except in the computation of $c_0^\star = \mathsf{E}_{K^\star}(M_d)$. At this point, the adversary is essentially an adversary against the indistinguishability (under passive attacks) of the authenticated encryption scheme Π^{sym}. We have $|\Pr[W_8] - 1/2| \le \mathbf{Adv}^{\mathsf{AE\text{-}IND}}(\lambda)$.

By collecting probabilities, we can bound the adversary's advantage as

$$\mathbf{Adv}^{\mathrm{cca}}(\mathcal{A}) \le \frac{3}{1 - 2^{-\lambda}} \cdot \mathbf{Adv}_{n, \lceil \lambda/\log n \rceil, q, \chi}^{\mathcal{B}, \mathsf{RLWE}}(\lambda) + Q(Q+1) \cdot \mathbf{Adv}^{\mathsf{AE\text{-}INT}}(\lambda) \qquad (7)$$

$$+ \mathbf{Adv}^{\mathsf{AE\text{-}IND}}(\lambda) + \frac{1}{2^{\Omega(n)}}.$$

$\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad$ $\square$

Lemma 14. *If $q > 8p(n + 1)$ and $p > 4$, Game_2 is perfectly indistinguishable from Game_1.*

Proof. We show that, from $\mathcal{A}$'s view, the random oracle H of Game_2 is identical to that of Game_1. This follows from the injecivity of $\mathsf{Encaps} : D_E \to R_q^2 : (r, e_1, e_2) \to \mathsf{Encaps}(PK, (r, e_1, e_2))$ over the domain $D_E := \{(r, e_1, e_2) \in R^3 : \|r\|_\infty, \|e_1\|_\infty, \|e_2\|_\infty \leq 1\}$. Indeed, suppose that $(r, e_1, e_2) \neq (\tilde{r}, \tilde{e}_1, \tilde{e}_2) \in D_E$ collide. The equality $b \cdot r + e_2 = b \cdot \tilde{r} + \tilde{e}_2$ over R_q would imply

$$p \cdot (a \cdot s + e)(r - \tilde{r}) \bmod q = \tilde{e}_2 - e_2.$$

Combining this with $a \cdot r + e_1 = a\tilde{r} + \tilde{e}_1$ (over R_q) would yield the equality

$$p \left(s \cdot (\tilde{e}_1 - e_1) + e(r - \tilde{r}) \right) = \tilde{e}_2 - e_2,$$

which would hold over R since $\|p \left(s \cdot (\tilde{e}_1 - e_1) + e(r - \tilde{r}) \right)\|_\infty \leq 4pn < q/2$. However, this impossible if $\tilde{e}_2 \neq e_2$ because $p > 4 > \|\tilde{e}_2 - e_2\|_\infty$. If $\tilde{e}_2 = e_2$, then it implies $\tilde{r} = r$ (since $b \in R_q^\times$) and $\tilde{e}_1 = e_1$.

Since Encaps is injective over D_E and H_Q is a random function, so is the composed function $H_Q\big(\mathsf{Encaps}(PK, (\cdot, \cdot, \cdot))\big)$. $\square$

Lemma 15. *Game_3 is perfectly indistinguishable from Game_2.*

Proof. We call a ciphertext (c_0, c_1, c_2) *good* if (c_1, c_2) satisfy

$$c_1 = a \cdot r + e_1 \in R_q, \qquad c_2 = b \cdot r + e_2 \in R_q$$

for some triple $(r, e_1, e_2) \in R^3$ such that $\|r\|_\infty, \|e_1\|_\infty, \|e_2\|_\infty \leq 1$. For a good ciphertext, the decryption oracle of Game_2 never sets $\mathsf{flag} = 1$. Indeed, if $\|e_2 + e_2'\|_\infty > 2$ at step 4 of $\mathsf{Decrypt}_2$ and (c_1, c_2) is good, we must have $\|e_2'\|_\infty > 1$, which is impossible.

Consequently, in Game_2, the random oracle H_Q is never evaluated on a good ciphertext. At the same time, if we have $\mathsf{flag} = 0$ at step 8, we know that the ciphertext is good by the rigidity property. Hence, H_Q' is only evaluated (either via a query to $H(\cdot, \cdot, \cdot)$ or a decryption query) on good ciphertexts in Game_2. Since H_Q and H_Q' are evaluated on disjoint sub-domains in Game_2, $\mathcal{A}$'s view remains the same if we simulate them as in Game_3, using a single random oracle. $\square$

Lemma 16. *Game_4 is perfectly indistinguishable from Game_3.*

Proof. Game_4 and Game_3 are identical from $\mathcal{A}$'s view until the two decryption oracles give different outputs for some decryption query (c_0, c_1, c_2).

We note that, for any decryption query $C = (c_0, c_1, c_2)$ such that the decryption oracle of Game_4 sets $\mathsf{flag} = 1$, it always computes $K = H_Q(c_1, c_2)$ at step 8 of $\mathsf{Decrypt}_2$. So, we focus on decryption queries for which the decryption oracle of Game_3 never sets flag to 1.

For such ciphertexts $C = (c_0, c_1, c_2)$, the rigidity of the encapsulation mechanism ensures that step 7 computes $(r, e_1, e_2) \in D_E$ such that $c_1 = a \cdot r + e_1$,

$c_2 = b \cdot r + e_1$, meaning that $(c_1, c_2) = \mathsf{Encaps}(PK, (r, e_1, e_2))$. Then, at step 8 of $\mathsf{Decrypt}_2$, the decryption oracle of Game_3 computes $K = H(r, e_1, e_2) = H_Q(c_1, c_2)$, just like the decryption oracle of Game_4.

In both cases $\mathsf{flag} \in \{0, 1\}$, Game_4 is perfectly indistinguishable from Game_3 as the two decryption oracles never disagree. □

Lemma 17. *If $q > n^{3/2}$, then Game_5 is indistinguishable from Game_4 under the $\mathsf{RLWE}_{n,k,q,\chi}$ assumption where $\chi = \bar{\psi}_2^n$ and $k = 2\lceil \lambda / \log n \rceil$. There is a PPT algorithm $\mathcal{B}$ such that $|\Pr[W_5] - \Pr[W_4]| \leq \mathbf{Adv}_{n,k,q,\chi}^{\mathcal{B}, \mathsf{RLWE}}(\lambda) + 2^{-\lambda}$.*

Proof. The proof is identical to that of Lemma 7, except that we set $k = 2\lceil \lambda / \log n \rceil$ to make sure that $(n/q)^k < 2^{-\lambda}$. □

Lemma 18. *Under the $\mathsf{RLWE}_{n,k,q,\chi}$ assumption with $\chi = \bar{\psi}_2^n$, Game_7 is indistinguishable from Game_6. Namely, $|\Pr[W_7] - \Pr[W_6]| \leq (1 - 2^{-\lambda})^{-1} \cdot \mathbf{Adv}_{n,k,q,\chi}^{\mathcal{B}, \mathsf{RLWE}}(\lambda)$, where $k = 4\lceil \lambda / \log n \rceil$.*

Proof. The proof and the reduction are identical to those of Lemma 8 □