Snake-Eye Resistant and Robust PKE from (Ring-)LWE With Small Secrets

Amit Deo and Benoît Libert

Zama, France

Abstract. Snake-eye resistance is a property of public-key encryption schemes that captures the infeasibility of generating a ciphertext which decrypts to the same plaintext under two honestly generated public keys. This property is non-trivial to achieve when ciphertexts should additionally hide the identity of the receiver (which rules out the trivial solution of appending the public key to the ciphertext). It is motivated by applications in *oblivious message retrieval* (OMR) protocols, where a server obliviously recovers messages from a public repository on behalf of clients without learning which client recovers which message. Snake-eye resistance was identified by Liu and Tromer (Crypto'22) as a key property allowing to hedge against denial-of-service attacks in OMR protocols. Recently (Eurocrypt'25), Liu *et al.* proved a conjecture saying that the LWE-based cryptosystem of Peikert-Vaikuntanathan-Waters (Crypto'08) is snake-eye resistant. They also showed that widely used (R)LWE-based public-key encryption schemes with small-norm secret keys fall short of providing the same property. They left open the question of realizing efficient RLWE-based PKE schemes that simultaneously feature snake-eye resistance and a simple decryption circuit which can be evaluated under FHE. In this work, we first show that a simple tweak in LWE-based PKE schemes with binary secret keys suffices to provide snake-eye resistance. As a second contribution, we show how to achieve snake-eye resistance (and in fact the stronger *robustness* property) under the RLWE assumption. Our constructions are the first examples of robust encryption from the RLWE assumption that do not use the KEM-DEM paradigm. Finally, we discuss different RLWE-based constructions that may offer better performance when used to build OMR protocols.

Keywords. Snake-eye resistance, robustness, (Ring-)LWE, oblivious message retrieval, denial-of-service attacks.

1 Introduction

User anonymity is a central requirement of many end-to-end encryption applications such as private messaging [19] and privacy-preserving cryptocurrencies [11,46]. Without it, sensitive meta-data may be leaked by means of traffic analysis, thus damaging sender/recipient anonymity. A typical situation is that of a bulletin board or public ledger containing a list of ciphertexts. To preserve anonymity, ciphertexts should not leak the identity of the sender or intended recipient. On the other hand, utility of the bulletin board should allow recipients

to identify ciphertexts corresponding to them.

Bellare *et al.* [10] put forth a property of *key-private* asymmetric encryption schemes whereby ciphertexts appear independent of their public keys and hide the receiver's identity. While this notion addresses the anonymity problem, it does not rule out mis-communication where a recipient mistakenly decrypts a ciphertext that was not intended for them. The latter problem can be avoided using *robust* encryption [1]. The notion of robustness captures the infeasibility of finding a ciphertext that decrypts to valid messages for two distinct key pairs. This ensures that, whenever users attempt to decrypt a ciphertext not intended for them, the decryption algorithm outputs a reject symbol without revealing anything about the actual recipient. In anonymous message delivery systems, it thus provides a mechanism allowing a sender to privately inform receivers that specific data items are pertinent to them, by appending an encryption (sometimes called a "clue" or "signal") of some fixed message to these items.

In cases where a bulletin board is prohibitively large, users may not have the computational power to download and store it locally, nor to perform a decryption attempt on every item's signal. Recent solutions [9,43,35] to this issue introduce an untrusted server to which clients can outsource the task of scanning the entire bulletin board in order to find the relevant messages. All of these solutions aim to hide the list of messages intended for a recipient (i.e., the list of *pertinent* messages) from the server. Among these solutions, *fuzzy message detection* [9] tolerates a server that can incorrectly identify a non-pertinent message as pertinent with some fairly large pre-defined false positive rate. As a result, it achieves a weak notion of privacy where truly pertinent messages are contained within some decoy set known to the server. As an improvement, *private signaling* [43] and *oblivious message retrieval* (OMR) [35] protocols satisfy a stronger notion of recipient privacy.

OMR and DoS Attacks. Practical OMR protocols [35,37,30,31] use fully homomorphic encryption (FHE) with bulletin board entries consisting of two parts: a payload and a clue (a.k.a. signal). The payload carries the (encrypted) message and the clue is used to decide pertinency. In particular, the clue takes the form of a relatively lightweight public key encryption (PKE) ciphertext that encrypts some pre-defined value. A bulletin board entry is marked as pertinent for a recipient if the clue decrypts to 1 under the recipient's "clue secret key" csk, where csk is the secret key underlying the public key used to produce clues.

In order to hide pertinency, the recipient provides the server with an FHE encryption of csk, which allows it to homomorphically decrypt clues. This allows the server to obtain FHE ciphertexts encrypting a list of binary values indicating pertinency. Next, the server can homomorphically compute carefully chosen inner-products of payload vectors and pertinency vectors so as to transfer pertinent payloads in a highly compressed "digest" to their receiver. Importantly, the size of the digest should be proportional to the maximal number of pertinent payloads and grow at most poly-logarithmically in the size of the bulletin board.

To ensure that pertinent messages are successfully detected, the clue in [35] takes the form of a PVW [49] encryption of $\boldsymbol{x} = (0, \ldots, 0)$ with plaintext space

$\{0,1\}^\ell$. The server first homomorphically runs the PVW decryption algorithm and then (still homomorphically) computes $\prod_{i=1}^{\ell}(1-x_i)$. Setting ℓ appropriately ensures that it is very unlikely that an *honestly* produced PVW ciphertext gets marked as pertinent for an unintended recipient.

As pointed out in [35], an active attacker willing to mount a denial of service (DoS) attack may be able to *maliciously* craft clues that trigger false positives for multiple recipients. To prevent such attacks, Liu and Tromer [35] relied on the so-called *snake-eye resistance conjecture* on PVW, which states it is infeasible to find an element of the ciphertext space that decrypts to $(0,\ldots,0)$ under two distinct keys. In a follow-up work [37], they replaced PVW by a more efficient PKE scheme based on the Ring Learning With Errors (RLWE) assumption so as to obtain better OMR performance and smaller public keys.

More recently, Liu *et al.* [34] removed the unproven conjecture from [35]. They showed under the Learning-With-Errors (LWE) assumption that PVW as well as a new, more efficient LWE-based encryption scheme (dubbed LWEmongrass) are both snake-eye resistant. They also argue that the decryption circuit of the latter can be efficiently evaluated using FHE. As a result, DoS-resistant OMR protocols can now rely on a standard LWE assumption. Along the way, they showed an attack against the snake-eye resistance of an LWE-based version of the efficient encryption of [37], where secret-keys consists of small-norm vectors. They argued that the same attack carries over to RLWE-based PKE schemes (e.g., [41]) with short secret keys.

To date, it remains an open question to build PKE schemes that are both snake-eye resistant and FHE-friendly under the RLWE assumption. As previously alluded to, in order to be usable in OMR protocols, a PKE should have a relatively simple decryption circuit that can be homomorphically evaluated in a reasonably efficient way. This prevents us from directly using robust schemes based on the KEM-DEM hybrid encryption paradigm.

Robustness and Snake-Eye Resistance. Robustness and snake-eye resistance are closely related. If a PKE scheme is strongly robust [1], it is trivially snake-eye resistant. In the converse direction, [34] gave a black-box transformation that builds a strongly robust PKE from a snake-eye resistant one.[1]

While several post-quantum NIST candidates were shown [22,54] to be simultaneously key-private and robust, they all rely on the KEM-DEM paradigm, which makes it much harder to homomorphically evaluate their decryption circuit in OMR protocols. For the time being, the only known robust or snake-eye resistant PKE schemes that avoid the KEM-DEM paradigm [22] are the LWE-based schemes highlighted in [34] (namely, PVW and LWEmongrass).

The main difficulty that arises when trying to adapt these solutions to the ring setting is that the attack of [34] seemingly poses a threat to any scheme where secret keys consist of small-norm vectors or ring elements. The reason is that these schemes involve ciphertexts $(\mathbf{a}, b) \in \mathbb{Z}_q^{n+1}$ where the decryption algorithm computes an inner product $b - \langle \mathbf{a}, \mathbf{s} \rangle \bmod q$. This allows a malicious

[1] It appends a snake-eye resistant encryption of 0 to the message-carrying ciphertext.

attacker to defeat snake-eye resistance by outputting a small-norm $(\mathbf{a}, b)$, which decrypts to 0 under any sufficiently small key $\mathbf{s} \in \mathbb{Z}^n$.

1.1 Contributions

As a first contribution, we show (perhaps surprisingly in light of the observations made by Liu *et al.* [34]) that the shortLWE construction discussed in [34] can be made snake-eye resistant under a standard assumption at essentially no cost.

Recall that [34] provides an attack on the snake-eye resistance of shortLWE and argue that any LWE-based PKE with similarly short secret keys is likely to be vulnerable to this attack. In Section 3, we show that the simple introduction of an extra norm check upon decryption is sufficient to provide snake-eye resistance. Specifically, we simply modify the decryption procedure to ensure that the first part $\mathbf{a} \in \mathbb{Z}_q^n$ of a ciphertext $(\mathbf{a}, \mathbf{b}) \in \mathbb{Z}_q^{n+\ell}$ has an infinity norm larger than some bound $4r$, where r is an upper bound on the decryption error term.

We note that our additional check can be done publicly and does not introduce any additional cost when it comes to homomorphically evaluating its decryption circuit in an OMR application.

Similarly to [34], we prove snake-eye resistance assuming the hardness of a correlation LWE problem (which is at least as hard as a standard LWE problem, as shown in [34]). We first prove that our augmented shortLWE scheme is snake-eye resistant when its secret keys are binary. In a second step, we show (in Section 3.2) that it retains security and snake-eye resistance when replacing binary secret keys by secret keys sampled from a binomial distribution. For both binary and binomial secret keys, we can then obtain strong robustness [1] by applying a transformation from [34, Section 9.2].

As a second contribution, we provide two constructions of strongly robust (and thus snake-eye resistant) PKE from the RLWE assumption. To our knowledge, these constructions are the first RLWE-based constructions of robust PKE schemes that do not rely on hybrid encryption nor random oracles. Prior to our results, the difficulty of porting existing robust LWE-based constructions to RLWE was discussed in [34].[2]

These two RLWE-based constructions can be seen as variants of the randomness-recovering scheme of Hoffmann *et al.* [24] that cover two different parameter settings. The construction in Section 4 is proven robust under the Ring Short Integer Solution (Ring SIS) assumption [39] and considers the usual NTT-friendly setting, where the ring $R_q = \mathbb{Z}_q[X]/(X^n+1)$ is isomorphic to $\mathbb{Z}_q^n$. In the full version, we present a variant that is proven robust in the statistical sense (without any assumption), but relies on partially splitting rings which are less NTT-friendly. The former allows us to pack more $\mathbb{Z}_q$-elements into a single plaintext. This is particularly advantageous when exploiting the SIMD capability of BFV/BGV [20,15] for the required range checks during homomorphic decryption. However, it turns out that the number of BFV/BGV multiplications and multiplication levels required for homomorphic decryption is fairly large. So, in our quest for OMR-compatible schemes, we investigate alternatives departing

[2] See Remark 5.5 of the full version.

from the randomness-recovering PKE scheme of [24].

In Section 5.1, we propose a KEM-DEM-like PKE scheme that leads to a statistically snake-eye private signaling scheme (each private signal consisting of an encryption of 0). This result uses the private signaling syntax of [30]. The construction relies on both RLWE and binary-secret LWE (the reason for using both assumptions will be explained shortly). It is compatible with NTT-friendly rings while providing statistical snake-eye resistance. In order to aid efficiency, we truncate ring elements of the ciphertext to decrease the number of homomorphic linear transformations. The removal of a modular reduction in the decryption circuit significantly reduces the number of multiplications in the homomorphic decryption circuit compared to the NTT-friendly scheme discussed above.

Despite this improvement, the homomorphic decryption performance using BFV remains significantly higher than in DoS-resistant schemes [34] based on the LWE assumption. On the other hand, we benefit from shorter public keys comprised of $\approx n \cdot \log q$ bits while [34] requires public keys in $\mathbb{Z}_q^{\ell \times n}$ for some repetition parameter $\ell > 1$ allowing for a false positive probability $\leq ((4r+1)/q)^{-\ell}$ where r is a bound on the decryption error term and q is the ciphertext modulus.

1.2 Technical Overview

Snake-eye resistance of shortLWE. The shortLWE PKE [34] was initially introduced as an adaptation of [42] using ideas described in [37]. It is essentially a variant of the Lindner-Peikert scheme [33] that uses a *binary* secret key $\mathbf{S} \in \{0,1\}^{n \times \ell}$ instead of a Gaussian key. Its public key is of the form $(\mathbf{A}, \mathbf{U}^\top = \mathbf{A}^\top \mathbf{S} + \mathbf{E} \bmod q)$ where $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ is uniformly random and $\mathbf{E} \in \mathbb{Z}^{m \times \ell}$ has small entries. A shortLWE encryption of $\boldsymbol{\mu} \in \mathbb{Z}_p^\ell$ takes the usual form $(\boldsymbol{a}, \boldsymbol{b}) = (\mathbf{A}\mathbf{r} + \mathbf{e}_0, \mathbf{U}\mathbf{r} + \mathbf{e}_1 + \Delta \cdot \boldsymbol{\mu})$ where $\Delta := \lfloor q/p \rceil$. Decryption *begins* with the usual computation of $\boldsymbol{\mu}^* := \lfloor (\mathbf{b} - \mathbf{S}^\top \mathbf{a})/\Delta \rceil$, but is augmented with a range check on the ciphertext error. In particular, the decryption outputs $\perp$ if the proposed ciphertext error $\mathbf{z} = \mathbf{b} - \mathbf{S}^\top \mathbf{a} - \Delta \cdot \boldsymbol{\mu}^*$ satisfies $\|\mathbf{z}\|_\infty > r$, and outputs $\boldsymbol{\mu}^*$ otherwise.

In order to make shortLWE snake-eye resistant, we simply introduce another test in the decryption algorithm and output $\perp$ whenever $\|\mathbf{a}\|_\infty < 4r$. The argument relies on a correlated LWE assumption in dimension n which has been shown [34] to follow from the LWE assumption in dimension roughly $n/2$ for binary secrets. Suppose a ciphertext $(\mathbf{a}, \mathbf{b})$ breaks snake-eye resistance i.e., it decrypts to the same non-$\perp$ plaintext for two distinct key pairs $(\mathbf{S}_j, (\mathbf{A}_j, \mathbf{U}_j))_{j \in \{1,2\}}$. Then, it follows that $\|\mathbf{a}\|_\infty \geq 4r$ and, for $j \in \{1,2\}$,

$$\mathbf{z}_j = \mathbf{b} - \mathbf{S}_j^\top \mathbf{a} - \Delta \cdot \boldsymbol{\mu}^* \bmod q$$

satisfies $\|\mathbf{z}_j\|_\infty \leq r$ for some plaintext $\boldsymbol{\mu}^*$. Subtracting the two equations, we obtain $(\mathbf{S}_1 - \mathbf{S}_2) \cdot \mathbf{a} \bmod q = \mathbf{z}_2 - \mathbf{z_1} + q\boldsymbol{\theta}$ for some $\boldsymbol{\theta} \in \mathbb{Z}^\ell$. However, since $r < q/4$, it must be that $\boldsymbol{\theta} = \mathbf{0}$ since the left-hand-side member has entries in $(-q/2, q/2]$ and we have $\|\mathbf{z}_1 - \mathbf{z}_2\|_\infty \leq 2r$. This implies $\|(\mathbf{S}_1 - \mathbf{S}_2) \cdot \mathbf{a} \bmod q\|_\infty \leq 2r$.

Now, the correlation LWE assumption says that one cannot distinguish $(\mathbf{A}_1, \mathbf{U}_1, \mathbf{A}_2, \mathbf{U}_2, \mathbf{S}_2 - \mathbf{S}_1)$ where $(\mathbf{U}_1, \mathbf{U}_2)$ is uniform from $(\mathbf{A}_1, \mathbf{S}_1^\top \mathbf{A}_1 + \mathbf{E}_1^\top, \mathbf{A}_2, \mathbf{S}_2^\top \mathbf{A}_2 + \mathbf{E}_2^\top, \mathbf{S}_2 - \mathbf{S}_1)$, for uniform $\mathbf{S}_1, \mathbf{S}_2 \in \{0,1\}^{n \times \ell}$ and Gaussian $\mathbf{E}_1, \mathbf{E}_2 \in \mathbb{Z}^{m \times \ell}$. In the former case, $\mathbf{U}_1$ and $\mathbf{U}_2$ are uniform over $\mathbb{Z}_q^{\ell \times n}$ and the adversary's view in the snake-eye resistance experiment is independent of $(\mathbf{S}_2, \mathbf{S}_1)$ if $\mathbf{U}_1$ and $\mathbf{U}_2$ are used to produce the two public keys. A statistical argument then shows that the probability that the adversary outputs an $\mathbf{a} \in \mathbb{Z}_q^n$ such that $\|\mathbf{a}\|_\infty \geq 4r$ and $\|(\mathbf{S}_1 - \mathbf{S}_2) \cdot \mathbf{a} \bmod q\|_\infty \leq 2r$ is at most $2^{-\ell}$. On the other hand, if the public keys use $\mathbf{U}_j = \mathbf{S}_j^\top \mathbf{A}_j + \mathbf{E}_j^\top$ for $j \in \{1,2\}$, the probability that the adversary outputs such an $\mathbf{a}$ is exactly its advantage in the experiment. Therefore, an adversary breaking the snake-eye resistance with advantage noticeably larger than $2^{-\ell}$ provides a distinguisher for the correlated LWE problem.

Strongly Robust RLWE-based PKE. Our RLWE-based robust schemes make use of power-of-two cyclotomic rings $R = \mathbb{Z}[X]/(X^n + 1)$. They achieve robustness by means of a mechanism allowing the receiver to retrieve the sender's random coins in order to perform a sanity check on the ciphertext. Both variants are adapted from the randomness-recovering hybrid encryption system of Hoffmann *et al.* [24]. The scheme uses three moduli q, p and t, where q is the ciphertext modulus, p is the randomness-recovery modulus and t is the plaintext modulus. We define $R_q = R/qR$ and $R_t = R/tR$.

The scheme has secret keys $s \in R$, where s has Gaussian coefficients, and a public key the form $\mathsf{pk} = (a, b = p \cdot (a \cdot s + e) \bmod q) \in R_q \times R_q^\times$ A plaintext $\mu \in R_t$ is encrypted as $(c_1, c_2) = (a \cdot r + e_1, b \cdot r + e_2 + \Delta\mu) \in R_q \times R_q$ in the standard way [41], using randomness $r, e_1, e_2 \in R$ with Gaussian coefficients. The decryption procedure first recovers a candidate plaintext $\mu^* = \lfloor (c_2 - p \cdot c_1 \cdot s \bmod q)/\Delta \rfloor$. Next, it subtracts $\Delta\mu^*$ from c_2 and then recovers (r, e_1, e_2) as in the decryption algorithm of [24]. The latter guarantees that, if a ciphertext (c_1, c_2) is not rejected, there exist sufficiently small $r, e_1 \in R$ such that $c_1 = a \cdot r + e_1$. In our proof of robustness, we leverage this property to show that, for two honestly generated public keys $\mathsf{pk} = (a, b)$, $\mathsf{pk}' = (a', b')$, no PPT adversary can output a ciphertext $(c_1, c_2) \in R_q$ and two colliding message-randomness pairs that explain (c_1, c_2) as valid ciphertexts under pk and pk', respectively.

In the case of fully splitting rings with $q = 1 \bmod 2n$, we prove this under the Ring SIS assumption [39]. If an adversary can find a ciphertext (c_1, c_2) that decrypts to valid $\mu, \mu' \in R_t$ under key pairs $(\mathsf{sk}, \mathsf{pk}) = (s, (a, b = p \cdot (a \cdot s + \widetilde{e})))$ and $(\mathsf{sk}', \mathsf{pk}') = (s', (a', b' = p \cdot (a' \cdot s' + \widetilde{e}')))$, the decryption algorithm recovers short (r, r', e_1, e_1') such that $c_1 = a \cdot r + e_1 = a' \cdot r' + e_1'$. This implies

$$a \cdot r - a' \cdot r' + (e_1 - e_1') = 0 \bmod q. \tag{1}$$

Letting $\boldsymbol{v} = (r, r', e_1 - e_1')$, we argue that $\boldsymbol{v} \neq 0$ via a public check that rejects (c_1, c_2) when $\|c_1\|_\infty$ is too small. Then, we obtain a short non-zero $\boldsymbol{v} \in R^3$ such that $\langle (a, a', 1), \boldsymbol{v} \rangle = 0 \bmod q$. For some $k > 0$, this yields a non-zero vector of short ring elements $\boldsymbol{w} = (w_1, \ldots, w_k) \in R^k$, such that $\sum_{i=1}^k a_i \cdot w_i = 0 \bmod q$, which solves a standard Ring SIS instance. This is achieved by defining a and a' appropriately with respect to $(a_1, \ldots a_k)$ and choosing k large enough so that

one of the a_i's are invertible.

The above construction can only be proven *computationally* robust (albeit under a standard assumption). In partially splitting rings, we can prove statistical robustness since, with high probability over the random choice of $a, a' \in R_q$, there exist no small $r, r', e_1, e_1' \in R$ satisfying (1). A proof of this fact follows from an infinity norm variant of [7, Lemma 5].

Snake-eye resistance in OMR protocols. When the robust RLWE-based PKE schemes are turned into private signaling schemes by encrypting a zero plaintext, recovering the encryption randomness is costly as it requires a homomorphic reduction modulo p to obtain an encryption of e_2, which is expensive. Using the interpolation function of [25, Proposition 1] with the BFV leveled FHE [20], it requires $\approx \sqrt{q}$ homomorphic multiplications if q is the ciphertext modulus of the PKE scheme (and the plaintext modulus of BFV).

To decrease the number of homomorphic multiplications in OMR protocols, we consider a different method that avoids modular reduction. We replace the randomness recovery method from [24] by a KEM-DEM-like construction. That is, instead of recovering the randomness r as above, we encapsulate it using the RLWE-based PKE scheme of Lyubashevsky, Peikert and Regev [41] and then use r as an (R)LWE secret key to encrypt the message μ. When r has binary coefficients, homomorphically decrypting an RLWE encryption of r is doable via two batched range checks using interpolation functions from [25]. This significantly reduces the number of homomorphic multiplications compared to the solution based homomorphic reductions mod p.

In Section 5.1, we provide an optimized version of this construction as a private-signaling scheme, which relies on both RLWE and LWE with binary secrets. If d is the binary-secret LWE dimension, we encrypt a random d-bit message $\mathbf{r} \in \{0,1\}^d$ using a truncated version of the LPR cryptosystem [41]. Then, we use $\mathbf{r}$ as an LWE secret to compute a vector $\mathbf{c} = \mathbf{A}^\top \mathbf{r} + \mathsf{noise}$, which is included in the private signal. A counting argument from [21] then ensures statistical snake-eye resistance since $\mathbf{c}$ is a statistically binding commitment to $\mathbf{r}$. The reason for using LWE instead of RLWE in $\mathbf{c}$ is two-fold. First, the latter enables a statistical argument of snake-eye resistance (as in our statistically robust RLWE-based PKE). Second, LWE offers more flexibility in the choice of d, which can be smaller than the RLWE dimension n (in particular, it does not have to be a power of 2). In practice, the uniform matrix $\mathbf{A}$ can be derived from a random oracle so as to not increase the public-key size of the private signaling scheme.

1.3 Related Work

There are multiple works on outsourcing message retrieval from a public bulletin board in a privacy-preserving manner. For example, [9] uses decoy sets (whose size can be tuned) to obscure which messages belong to a recipient. However, the server learns a super-set of the pertinent messages.

Stronger security is achieved using private signaling [43,26] and oblivious message retrieval [35,36,37,27,34] protocols. A drawback of the private signaling approaches in [43,26] is that trusted execution environments such as Intel SGX

are required. The study of OMR protocols was initiated by Liu and Tromer [35] using an FHE-based approach and LWE-based encryption. This same approach was optimized using a RLWE-based encryption in [37]. Subsequently, Lee and Yeo [30] further improved on this line of work in terms of run-time and digest-size using a Newton's identity-based pertinency vector compression method [16]. A limitation of these works is that, at each retrieval, a bound on the number of pertinent messages must be known. If this bound is violated, the retrieval fails. The need for such a bound is removed by [27]. The group setting where a single message can have multiple intended recipients is studied in [36].

Liu and Tromer [35] discussed DoS attacks whereby a malicious adversary can carefully craft malformed signals/clues that are falsely identified as being pertinent to many users. Clearly, this can affect the utility of OMR and even ruin it in the case where a bound on the number of pertinent messages is required. To hedge against such attacks, Liu and Tromer [35] made a *snake-eye resistance conjecture* on the LWE-based PKE scheme [49] used by their protocol. This conjecture was proven in [34] where the snake-eye resistance of certain LWE-based schemes was established. Liu *et al.* [34] further described attacks on the snake-eye resistance of (R)LWE-based PKE schemes with short keys.

Robust encryption was first formalized by Abdalla *et al.* [1] and motivated by applications in private auction protocols [51,1] and searchable public-key encryption [12]. A form of robustness is also useful in anonymous broadcast encryption schemes [6,32]. Liu *et al.* [34] presented a generic transformation from a snake-eye resistant PKE into a strongly robust one [1]. Robust PKE schemes from lattices may be built in the KEM-DEM paradigm [22,54], but these are ill-suited to FHE-based OMR protocols which require efficient homomorphic decryption circuits. We also note that our RLWE-based schemes achieve robustness without relying on a random oracle and while preserving their additive homomorphism (up to a small growth in bound checks).

2 Background

We denote by $\mathbb{Z}$, $\mathbb{N}$ and $\mathbb{R}$ the sets of integers, natural numbers and reals, respectively. For any $a \in \mathbb{N}$, $[a]$ will stand for the set $\{1, \ldots, a\}$. For $b, c \in \mathbb{R}$, we define the sets $[b, c] = \{x \in \mathbb{Z} : b \leq x \leq c\}$ as well as $(b, c] = \{x \in \mathbb{Z} : b < x \leq c\}$.

Vectors will be denoted in bold e.g., $\mathbf{x}$ and $\mathbf{x}[i]$ will denote the i-th entry of $\mathbf{x}$. The inner product of two vectors $\mathbf{x}$ and $\mathbf{y}$ will be written as $\langle \mathbf{x}, \mathbf{y} \rangle$ and the component-wise product will be written as $\mathbf{x} \circ \mathbf{y}$, where $\mathbf{x} \circ \mathbf{y}[i] = \mathbf{x}[i] \cdot \mathbf{y}[i]$. Matrices will be written as both upper-case and bold e.g., $\mathbf{A}$.

For distribution $\mathcal{D}$, $x \hookleftarrow \mathcal{D}$ denotes the act of sampling x from $\mathcal{D}$. For a finite set X, $U(X)$ denotes the uniform distribution over X. The *centered* binomial distribution with 2η trials and a success probability of $p = 1/2$ in each trial will be denoted $\mathcal{B}_\eta$. To sample from $\mathcal{B}_\eta$, we may simply sample $(c_i, d_i) \hookleftarrow \{0, 1\}^2$ uniformly for $i \in \{1, \ldots, \eta\}$ and then output $\sum_{i=1}^{\eta}(c_i - d_i)$. We use the standard asymptotic notation o, O, Θ, Ω. As usual, a function $\rho(\lambda)$ is negligible in the security parameter λ if for any positive integer c, there exists an integer λ_c such

that $|\rho(\lambda)| < \lambda^{-c}$ for any $\lambda > \lambda_c$. The set of negligible functions is denoted $\mathsf{negl}(\lambda)$. The set of polynomially bounded functions in λ (i.e., functions that are $O(n^c)$ for some constant c) will be denoted as $\mathsf{poly}(\lambda)$.

2.1 Lattices and 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 $\mathbf{B} = \{\mathbf{b}_1, \ldots, \mathbf{b}_n\} \subseteq \mathbb{R}^n$.

Discrete Gaussian Distributions. Let $\mathbf{\Sigma} \in \mathbb{R}^{n \times n}$ be a symmetric positive-definite matrix, and let $\mathbf{c} \in \mathbb{R}^n$. We define the Gaussian function on $\mathbf{x} \in \mathbb{R}^n$ by $\rho_{\mathbf{\Sigma},\mathbf{c}}(\mathbf{x}) = \exp(-\pi(\mathbf{x}-\mathbf{c})^\top \mathbf{\Sigma}^{-1}(\mathbf{x}-\mathbf{c}))$. In the case $\mathbf{\Sigma} = \sigma^2 \cdot \mathbf{I}_n$ for a real value $\sigma > 0$ and $\mathbf{c} = \mathbf{0}$, we denote by ρ_σ the Gaussian function $\rho_\sigma(\mathbf{x}) = \exp(-\pi\|\mathbf{x}\|^2/\sigma^2)$ with standard deviation σ. For any lattice $\Lambda \subset \mathbb{R}^n$, the discrete Gaussian distribution $D_{\Lambda,\sigma}$ has probability mass $\Pr_{X \sim D_{\Lambda,\sigma}}[X = \mathbf{x}] = \frac{\rho_\sigma(\mathbf{x})}{\rho_\sigma(\Lambda)}$ for any $\mathbf{x} \in \Lambda$, where we define $\rho_\sigma(\Lambda) = \sum_{\mathbf{y} \in \Lambda} \rho_\sigma(\mathbf{y})$.

Lemma 1 ([5, Lemma 1.5],[38, Lemma 4.4]).

1. *For any $k > 0$, $\Pr[|z| > k \cdot \sigma \mid z \hookleftarrow D_{\mathbb{Z},\sigma}] \leq 2 \cdot \exp(-k^2/2)$.*
2. *For any $k > 1$, $\Pr[\|\mathbf{z}\| > k \cdot \sigma\sqrt{m} \mid z \hookleftarrow D_{\mathbb{Z}^m,\sigma}] \leq k^m \cdot \exp(\frac{m}{2}(1-k^2))$.*

Definition 1. *A distribution χ supported over $\mathbb{Z}$ is (B,ϵ)-bounded if we have $\Pr_{x \hookleftarrow \chi}[|x| > B] < \epsilon$.*

We write that χ is B-bounded when ϵ is bounded by a negligible function of the security parameter. Lemma 1 implies that $D_{\mathbb{Z},\sigma}$ is B-bounded for $B = \sigma\sqrt{2\lambda}$.

The LWE Assumption. We now recall the Learning With Errors problem [50].

Definition 2. *Let $\lambda \in \mathbb{N}$ be a security parameter and take integers $n = n(\lambda)$, $m = m(\lambda)$, $q = q(\lambda)$. Let $\chi_s = \chi_s(\lambda)$, $\chi_e = \chi_e(\lambda)$ be efficiently sampleable distributions over $\mathbb{Z}_q$ and $\mathbb{Z}$ respectively. The $\mathsf{LWE}_{n,m,q,\chi_s,\chi_e}$ assumption posits that the following distance is a negligible function for any PPT algorithm $\mathcal{A}$:*

$$\begin{aligned}\mathbf{Adv}^{\mathcal{A},\mathsf{LWE}}_{n,m,q,\chi_s,\chi_e}(\lambda) := \big|&\Pr[\mathcal{A}(1^\lambda, \mathbf{A}, \mathbf{u}) = 1 \mid \mathbf{A} \hookleftarrow U(\mathbb{Z}_q^{n \times m}), \mathbf{u} \hookleftarrow U(\mathbb{Z}_q^m)] \\ &- \Pr[\mathcal{A}(1^\lambda, \mathbf{A}, \mathbf{A}^\top \mathbf{s} + \mathbf{e} \bmod q) = 1 \mid \mathbf{A} \hookleftarrow U(\mathbb{Z}_q^{n \times m}), \mathbf{s} \hookleftarrow \chi_s^n, \mathbf{e} \hookleftarrow \chi_e^m]\big|.\end{aligned}$$

A common distribution for $\mathbf{s}$ is the uniform distribution $\chi_s = U(\mathbb{Z}_q)$. Small secrets from distributions such as the Gaussian distribution, the uniform binary distribution $\chi_s = U(\{0,1\})$ or the uniform ternary distribution $U(\{-1,0,1\})$ are also popular as they often lead to more efficient constructions. A typical choice for χ_e is the discrete Gaussian distribution $D_{\mathbb{Z},\alpha q}$, for some parameter $\alpha \in (\sqrt{n}/q, 1)$. In this case, choosing $\alpha q > 2\sqrt{n}$ allows for quantum reductions from standard lattice problems with approximation factor $\gamma = \widetilde{O}(n/\alpha) = \widetilde{O}(\sqrt{n}q)$ to LWE (see, e.g., [50,14]). The best lattice algorithms for approximation factor γ run in time at least $2^{\widetilde{\Omega}(n/\log\gamma)}$ [52]. Yet another practical choice for the secret distribution is the centered binomial distribution [13].

We also use the following LWE variant defined in [34].

Definition 3. *Let $\lambda \in \mathbb{N}$ be a security parameter and take integers $n = n(\lambda)$, $m = m(\lambda)$, $q = q(\lambda)$. Let $\chi_s = \chi_s(\lambda)$, $\chi_e = \chi_e(\lambda)$ be efficiently sampleable distributions over $\mathbb{Z}_q$ and $\mathbb{Z}$ respectively. The* LWE*-with-correlation assumption* ($\mathsf{CorLWE}_{n,m,q,\chi_s,\chi_e}$) *says that* $\mathbf{Adv}^{\mathcal{A},\mathsf{CorLWE}}_{n,m,q,\chi_s,\chi_e}(\lambda) := \big|p_0 - p_1\big|$*, is a negligible function of λ for any PPT algorithm $\mathcal{A}$, where*

$$p_0 := \Pr\big[\mathcal{A}(1^\lambda, \mathbf{A}_1, \mathbf{A}_2, \mathbf{u}_1, \mathbf{u}_2, \bar{\mathbf{s}}) = 1 \mid \mathbf{A}_1, \mathbf{A}_2 \hookleftarrow U(\mathbb{Z}_q^{n\times m}), \\ \mathbf{s}_1, \mathbf{s}_2 \hookleftarrow \chi_s^n, \mathbf{u}_1, \mathbf{u}_2 \hookleftarrow U(\mathbb{Z}_q^m), \bar{\mathbf{s}} = \mathbf{s}_1 - \mathbf{s}_2 \bmod q\big],$$

$$p_1 := \Pr\big[\mathcal{A}(1^\lambda, \mathbf{A}_1, \mathbf{A}_2, \mathbf{A}_1^\top \mathbf{s}_1 + \mathbf{e}_1 \bmod q, \mathbf{A}_2^\top \mathbf{s}_2 + \mathbf{e}_2 \bmod q, \bar{\mathbf{s}}) = 1 \\ \mid \mathbf{A}_1, \mathbf{A}_2 \hookleftarrow U(\mathbb{Z}_q^{n\times m}), \mathbf{s}_1, \mathbf{s}_2 \hookleftarrow \chi_s^n, \mathbf{e}_1, \mathbf{e}_2 \hookleftarrow \chi_e^m, \bar{\mathbf{s}} = \mathbf{s}_1 - \mathbf{s}_2 \bmod q\big]$$

Liu *et al.* [34] gave reductions from LWE to CorLWE in the cases of uniform, binary and ternary secrets. They showed that binary-secret CorLWE (i.e., with $\chi_s = U(\{0,1\})$ and $\chi_e = D_{\mathbb{Z},\alpha q}$ for some parameter $\alpha \in (0,1)$) is at least as hard has LWE with binary secrets in smaller dimension $n' = cn/2$ for some $c \in (0,1)$.

2.2 Rings and Ideal Lattices

Let n be a power of 2 and a prime modulus q. Let the cyclotomic ring $R = \mathbb{Z}[X]/(X^n + 1)$ and its quotient ring $R_q = R/qR$. Each element of R is a polynomial of degree $n-1$ in $\mathbb{Z}[X]$ and can be interpreted as an element of $\mathbb{Z}[X]$ via the natural coefficient embedding that maps $a = \sum_{i=0}^{n-1} a_i X^i \in R$ to $(a_0, a_1, \ldots, a_{n-1}) \in \mathbb{Z}^n$. An element of R_q can be viewed as a degree-$(n-1)$ polynomial over $\mathbb{Z}_q[X]$ and represented as a vector over $\mathbb{Z}^n$ with coefficients in $\{-\frac{q-1}{2}, \ldots, \frac{q-1}{2}\}$. For any $r \in R_q$, $\phi(r) \in \mathbb{Z}_q^n$ denotes its coefficient vector.

We define $\mathsf{rot} : R \to \mathbb{Z}^{n\times n}$ as the ring homomorphism that sends $a \in R$ to a matrix in $\mathbb{Z}^{n\times n}$ of which the i-th row is $\phi\big(a \cdot X^{i-1} \bmod (X^n+1)\big) \in \mathbb{Z}^n$. Note that $\phi(a \cdot s) = \mathsf{rot}(a) \cdot \phi(s)$.

As in [29], for any lattice Λ, $D^{\mathsf{coeff}}_{\Lambda,\sigma}$ 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 $D_{\Lambda,\sigma}$.

The Euclidean/infinity norm of an element $a \in R$ and a vector $\boldsymbol{v} \in R^k$ is defined by viewing elements of R as elements of $\mathbb{Z}^n$ via the coefficient embedding.

If q is a prime and n is a power of 2 such that $q = 1 \bmod 2n$, then there is a ring isomorphism from R_q to $\mathbb{Z}_q^n$ (where multiplication in the latter is performed component-wise). In such a case, we say that the ring fully splits. The mapping that realizes this isomorphism is known as the number theoretic transform (NTT). In particular, if we have $\mathsf{NTT}(r) = (\hat{r}_1, \ldots, \hat{r}_n) \in \mathbb{Z}_q^n$ and $\mathsf{NTT}(s) = (\hat{s}_1, \ldots, \hat{s}_n) \in \mathbb{Z}_q^n$, then $\mathsf{NTT}(r+s) = (\hat{r}_1 + \hat{s}_1, \ldots, \hat{r}_n + \hat{s}_n)$ and $\mathsf{NTT}(r \cdot s) = (\hat{r}_1 \cdot \hat{s}_1, \ldots, \hat{r}_n \cdot \hat{s}_n)$. Furthermore, for any $r \in R_q$, $\mathsf{NTT}(r) = M \cdot \phi(r)$ for some matrix $M \in \mathbb{Z}_q^{n\times n}$ that depends only on (n, q).

We recall the ring Learning-With-Errors (RLWE) problem in Hermite normal form. Note that we make the number of samples k explicit in our definition.

Definition 4. *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 $\chi = \chi(n)$ be an error distribution over*

R_q. *The* $\mathsf{RLWE}_{n,k,q,\chi}$ *assumption says that the following distance is a negligible function for any PPT algorithm* $\mathcal{A}$*:*

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

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 [41]. We also sometimes consider the case where the secret s in the definition above comes from a distribution χ_s that does not coincide with χ. When this is the case, we denote the assumption as $\mathsf{RLWE}_{n,k,q,\chi_s,\chi}$.

We also use the Ring SIS assumption [44] in the proof of robustness of some of our ring-based constructions.

Definition 5. *Let integers* $n = n(\lambda)$*,* $k = k(\lambda)$*, and a prime* $q = q(n) > 2$*, for a security parameter* λ*. Let the ring* $R_q = R/(qR)$ *with* $R = \mathbb{Z}[X]/(X^n + 1)$*. The* ***Ring Short Integer Solution*** *problem (*$\mathsf{RSIS}_{n,k,q,\beta}$*) is, given a random tuple* $(a_1, \ldots, a_k) \sim U(R_q^k)$*, to find a non-zero* $(e_1, \ldots, e_k) \in R^k \setminus \{\mathbf{0}^k\}$ *such that* $\|e_i\|_\infty < \beta$ *for each* $i \in [k]$*.*

For $k > \frac{\log q}{\log(2\beta)}$, $\gamma = 16\beta \cdot kn \log^2 n$ and $q > \frac{\gamma\sqrt{n}}{4 \log n}$, [39,48] showed that solving $\mathsf{RSIS}_{n,k,q,\beta}$ is as hard as solving $\mathrm{SVP}^\infty_\gamma$ in any ideal of the ring $\mathbb{Z}[X]/(X^n + 1)$.

2.3 Snake-Eye Resistant and Robust Public-Key Encryption

The notion of δ-snake-eye resistance [35,34] ensures that it is infeasible to compute a single ciphertext that decrypts to the same plaintext under two honestly generated key pairs, except with some small probability δ.

Definition 6 (Snake-eye resistance). *A PKE scheme is* δ*-snake-eye resistant if, for any* $\mathsf{pp} \leftarrow \mathsf{ParGen}(1^\lambda)$*, any* $(\mathsf{pk}_1, \mathsf{sk}_1) \leftarrow \mathsf{Keygen}(\mathsf{pp})$*,* $(\mathsf{pk}_2, \mathsf{sk}_2) \leftarrow \mathsf{Keygen}(\mathsf{pp})$ *and any PPT adversary* $\mathcal{A}$*, we have*

$$\Pr[\mathsf{Decrypt}(\mathsf{pp}, \mathsf{sk}_1, \mathsf{ct}) = \mathsf{Decrypt}(\mathsf{pp}, \mathsf{sk}_2, \mathsf{ct}) \neq \perp | \ \mathsf{ct} \leftarrow \mathcal{A}(\mathsf{pp}, \mathsf{pk}_1, \mathsf{pk}_2)] \leq \delta.$$

If $\delta \in \mathsf{negl}(\lambda)$*, the scheme is simply said to be snake-eye resistant.*

Snake-eye resistance was previously considered under the name of *collision-freeness* in [45]. A stronger notion called *robustness* [1] requires the infeasibility of outputting a ciphertext that decrypts to valid (possibly distinct) plaintexts under two distinct honestly generated keys.

Definition 7 (Strong robustness). *A PKE scheme is* δ*-strongly robust if, for any* $\mathsf{pp} \leftarrow \mathsf{ParGen}(1^\lambda)$*, any* $(\mathsf{pk}_1, \mathsf{sk}_1) \leftarrow \mathsf{Keygen}(\mathsf{pp})$*,* $(\mathsf{pk}_2, \mathsf{sk}_2) \leftarrow \mathsf{Keygen}(\mathsf{pp})$ *and any PPT adversary* $\mathcal{A}$*, we have*

$$\Pr[\mathsf{Decrypt}(\mathsf{pp}, \mathsf{sk}_1, \mathsf{ct}) \neq \perp \ \wedge \ \mathsf{Decrypt}(\mathsf{pp}, \mathsf{sk}_2, \mathsf{ct}) \neq \perp | \ \mathsf{ct} \leftarrow \mathcal{A}(\mathsf{pp}, \mathsf{pk}_1, \mathsf{pk}_2)] \leq \delta.$$

If $\delta \in \mathsf{negl}(\lambda)$*, the scheme is simply said to be strongly robust.*

These notions are mostly useful in the context of PKE schemes that provide anonymity, meaning that ciphertexts do not betray the public key under which they are encrypted. As pointed out in [1], without anonymity, robustness can be trivially achieved by simply appending the public key to the ciphertext.

2.4 Private Signaling

We now recall the definition of private signaling [43]. We use the syntax and the model of [30], which incorporates a notion of key-unlinkability.

Definition 8 ([30]). *A Private Signaling (PS) scheme is a tuple* $(\mathsf{ParGen}, \mathsf{Keygen}, \mathsf{Signal}, \mathsf{Detect})$ *of PPT algorithms with the following specification:*

- ParGen *is a PPT algorithm that takes as input a security parameter* $\lambda \in \mathbb{N}$ *and outputs public parameters* pp.
- Keygen *is a PPT algorithm that inputs* pp *and outputs a key pair* $(\mathsf{sk}, \mathsf{pk})$.
- Signal *is a PPT algorithm that takes in* pk *and returns a signal* sig.
- Detect *is a deterministic detection algorithm that inputs* sk *and a candidate signal* sig. *It outputs* 1 *(meaning that* sig *is a valid signal for* pk*) or* 0.

We now recall the correctness and security definitions of [30] and augment them with the notion of snake-eye resistance, which was not considered in [30]. We note that the notion of snake-eye resistance implies that of soundness since it accounts for maliciously generated signals.

Definition 9 (Completeness and security of private signaling).

- *A PS scheme is complete if, for any* $\mathsf{pp} \leftarrow \mathsf{ParGen}(1^\lambda)$,

$$\Pr[\mathsf{Detect}(\mathsf{pp}, \mathsf{sk}, \mathsf{sig}) = 1 \mid (\mathsf{sk}, \mathsf{pk}) \leftarrow \mathsf{Keygen}(\mathsf{pp}); \\ \mathsf{sig} \leftarrow \mathsf{Signal}(\mathsf{pk})] \geq 1 - \mathsf{negl}(\lambda).$$

- *A PS scheme is sound if, for any* $\mathsf{pp} \leftarrow \mathsf{ParGen}(1^\lambda)$,

$$\Pr[\ \mathsf{Detect}(\mathsf{pp}, \mathsf{sk}', \mathsf{sig}) = 1 \mid\ (\mathsf{sk}, \mathsf{pk}) \leftarrow \mathsf{Keygen}(\mathsf{pp}); \\ (\mathsf{sk}', \mathsf{pk}') \leftarrow \mathsf{Keygen}(\mathsf{pp});\ \mathsf{sig} \leftarrow \mathsf{Signal}(\mathsf{pk})\ ] \in \mathsf{negl}(\lambda).$$

- *A PS scheme is receiver-private if, for any* $\mathsf{pp} \leftarrow \mathsf{ParGen}(1^\lambda)$*, and any PPT adversary* $\mathcal{A}$*, we have*

$$\begin{aligned}\mathbf{Adv}_{\mathcal{A}}(\lambda) := \big|\, &\Pr[1 \leftarrow \mathcal{A}(\mathsf{pk}_0, \mathsf{pk}_1, \mathsf{sig}) \mid\ (\mathsf{sk}_0, \mathsf{pk}_0) \leftarrow \mathsf{Keygen}(\mathsf{pp}); \\ &\qquad (\mathsf{sk}_1, \mathsf{pk}_1) \leftarrow \mathsf{Keygen}(\mathsf{pp});\ \mathsf{sig} \leftarrow \mathsf{Signal}(\mathsf{pk}_0)] \\ &- \Pr[1 \leftarrow \mathcal{A}(\mathsf{pk}_0, \mathsf{pk}_1, \mathsf{sig}) \mid\ (\mathsf{sk}_0, \mathsf{pk}_0) \leftarrow \mathsf{Keygen}(\mathsf{pp}); \\ &\qquad (\mathsf{sk}_1, \mathsf{pk}_1) \leftarrow \mathsf{Keygen}(\mathsf{pp});\ \mathsf{sig} \leftarrow \mathsf{Signal}(\mathsf{pk}_1)]\,\big| \in \mathsf{negl}(\lambda).\end{aligned}$$

- *A PS scheme is* δ*-snake-eye resistant if, for any* $\mathsf{pp} \leftarrow \mathsf{ParGen}(1^\lambda)$*, any* $(\mathsf{pk}_0, \mathsf{sk}_0) \leftarrow \mathsf{Keygen}(\mathsf{pp})$*,* $(\mathsf{pk}_1, \mathsf{sk}_1) \leftarrow \mathsf{Keygen}(\mathsf{pp})$ *and any PPT* $\mathcal{A}$*,*

$$\Pr[\mathsf{Detect}(\mathsf{pp}, \mathsf{sk}_0, \mathsf{sig}) = \mathsf{Detect}(\mathsf{pp}, \mathsf{sk}_1, \mathsf{sig}) = 1 \mid \mathsf{sig} \leftarrow \mathcal{A}(\mathsf{pp}, \mathsf{pk}_0, \mathsf{pk}_1)] \leq \delta.$$

If $\delta \in \mathsf{negl}(\lambda)$*, the scheme is just said to be snake-eye resistant.*

2.5 Useful Lemmas

We rely on Lemma 2 with $\mathbf{u} = 0$, in which case a proof was already given in [8, Lemma 2.3] when $\mathbf{R}$ is sampled from the uniform binary distribution. It was extended in [40] to distributions centered on 0.

Lemma 2 ([40, Lemma 2.5]). *Fix integers $\ell, q \in \mathbb{Z}$ and any vectors $\mathbf{u} \in [\pm q/2]^\ell$ and $\mathbf{w} \in [\pm q/2]^n$. Let $\mathcal{D}$ be a distribution with support $\{-1, 0, 1\}$ such that $\mathcal{D}[0] = 1/2$ and $\mathcal{D}[\pm 1] = 1/4$. Then, over the random choice of $\mathbf{R} \hookleftarrow \mathcal{D}^{\ell \times n}$,*

$$\Pr_{\mathbf{R}}\left[\left\|\mathbf{u} + \mathbf{R} \cdot \mathbf{w} \bmod q\right\|_\infty < \frac{1}{2} \cdot \left\|\mathbf{w}\right\|_\infty\right] < 2^{-\ell}$$

2.6 BGV/BFV Leveled Homomorphic Encryption

Typical applications of our constructions involve homomorphic evaluation of decryption or detection circuits using leveled BGV/BFV homomorphic encryption [15,20]. We overview the main characteristics of BGV/BFV here. For simplicity we refer to BFV, but all observations hold similarly for BGV. We note that one may use the BFV plaintext space $\mathbb{Z}_t^N$ which is isomorphic to $R_t = \mathbb{Z}_t[X]/(X^N + 1)$ for power-of-two N and prime $t = 1 \bmod 2N$ [53] via the NTT. Each copy of $\mathbb{Z}_t$ in the plaintext space is referred to as a "slot". BFV natively supports slot-wise homomorphic additions and multiplications. That is, ciphertexts encrypting $(x_1, \ldots, x_N)$ and $(y_1, \ldots, y_N)$ may be homomorphically added to obtain a new ciphertext encrypting $(x_1+y_1, \ldots, x_N+y_N)$, or homomorphically multiplied to obtain an encryption of $(x_1 \cdot y_1, \ldots, x_N \cdot y_N)$.

Homomorphic operations introduce extra "noise" to a ciphertext. As a result, one must set parameters to ensure that the noise does not grow too large for correct decryption. Multiplication is the main contributor to noise growth, and the crucial quantity influencing parameter selection is the multiplicative depth of the computation. The higher the multiplicative depth, the larger the ciphertext modulus and ring dimension in the BFV parameters (when fixing the noise distribution as suggested in the Homomorphic Encryption Standard [3]). This leads to larger public key/ciphertext sizes and slower computation. Additionally, homomorphic multiplication is much slower than addition. Therefore, to gauge overall performance of a BFV application, one can calculate the multiplicative depth and total number of homomorphic multiplications required.

A further operation of interest is multiplying an encrypted vector $\boldsymbol{v} \in \mathbb{Z}_t^N$ by a plaintext matrix $M \in \mathbb{Z}_t^{N \times N}$. In order to achieve this, one can use homomorphic rotations to rotate the entries of $\boldsymbol{v}$ while expressing $M \cdot \boldsymbol{v}$ as a sum of Hadamard products between different rotations of M and rotations of $\boldsymbol{v}$. The homomorphic rotations are performed by applying automorphisms followed by a key-switching operation. Assuming that $\boldsymbol{v}$ fits into a single ciphertext, only $O(\sqrt{N})$ rotations are required [23]. If k' ciphertexts are required to encrypt $\boldsymbol{v}$, the number of rotations increases by a factor of k' [30]. Overall, the noise increase due to plaintext matrix multiplication is asymptotically smaller than that of a ciphertext-ciphertext multiplication.

3 Simple Snake-Eye Resistant PKE from Short-key LWE

We show that the shortLWE scheme considered in [37,34] (and inspired by the Lindner-Peikert [33] scheme) can be made snake-eye resistant by the simple action of testing whether the first ciphertext component is a vector of sufficiently large infinity norm. We emphasize that this test can be done publicly and does not increase the cost of homomorphically evaluating the decryption circuit.

We first show that a variant of the LWE-based cryptosystem of [33] is snake-eye resistant when using binary secret keys. In Section 3.2, we generalize the result to binomial secret keys. A benefit of binary secret keys is that they make it easier to homomorphically decrypt using TFHE [17] (which also uses binary keys). In particular, one may simply encode the range check on $\mathbf{z}$ using a specially crafted look-up table and apply TFHE programmable bootstrapping [18].

3.1 The Case of Binary Secret Keys

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

1. Choose dimensions $n, m \in \mathsf{poly}(\lambda)$, a ciphertext modulus $q \in \mathsf{poly}(\lambda)$, a plaintext modulus $p < q$, and an integer $\ell \in \mathsf{poly}(\lambda)$ denoting the number of message components.
2. Choose a noise parameter $\alpha \in (0,1)$ and take the discrete Gaussian noise distribution $\chi_e = D_{\mathbb{Z},\alpha q}$. Let $r > 0$ be the smallest integer such that $\Pr_{\mathbf{e} \hookleftarrow \chi_e^{m+n+1}}[|\langle \mathbf{e}, \mathbf{1}^{m+n+1} \rangle| > r] \in \mathsf{negl}(\lambda)$.

Return public parameters $\mathsf{pp} := (n,\ m,\ q,\ p,\ \alpha,\ r)$.

Keygen(pp)**:** On input of pp, generate a key pair as follows.

1. Choose a random matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$.
2. Sample $\mathbf{S} \hookleftarrow U(\{0,1\}^{n \times \ell})$, $\mathbf{E} \hookleftarrow D_{\mathbb{Z}^m,\alpha q}^{\ell}$ and compute

$$\mathbf{U}^\top = \mathbf{A}^\top \mathbf{S} + \mathbf{E} \in \mathbb{Z}_q^{m \times \ell}$$

Return $\mathsf{sk} = \mathbf{S} \in \{0,1\}^{n \times \ell}$ and $\mathsf{pk} = (\mathsf{pp}, \mathbf{A}, \mathbf{U})$.

Encrypt(pk, μ)**:** Given a public key pk and a message $\boldsymbol{\mu} \in \mathbb{Z}_p^\ell$, Sample $\mathbf{r} \hookleftarrow U(\{0,1\}^m)$, $\mathbf{e}_0 \hookleftarrow D_{\mathbb{Z}^n,\alpha q}$, $\mathbf{e}_1 \hookleftarrow D_{\mathbb{Z}^\ell,\alpha q}$. Then, compute

$$\mathbf{a} = \mathbf{A} \cdot \mathbf{r} + \mathbf{e}_0 \bmod q, \qquad \mathbf{b} = \mathbf{U} \cdot \mathbf{r} + \mathbf{e}_1 + \Delta \cdot \boldsymbol{\mu} \bmod q, \tag{2}$$

where $\Delta = \lfloor q/p \rceil$. Output the ciphertext $\mathsf{ct} = (\mathbf{a}, \mathbf{b}) \in \mathbb{Z}_q^n \times \mathbb{Z}_q^\ell$.

Decrypt$(\mathsf{sk}, \mathsf{ct})$**:** Given $\mathsf{ct} = (\mathbf{a}, \mathbf{b}) \in \mathbb{Z}_q^{(n+\ell)}$ and $\mathsf{sk} = \mathbf{S} \in \{0,1\}^{n \times \ell}$,

1. Return $\perp$ if $\|\mathbf{a}\|_\infty < 4r$, when $\mathbf{a}$ is seen as a vector in $\mathbb{Z}^n \cap (-q/2, q/2]^n$.
2. Compute the plaintext $\boldsymbol{\mu} = \lfloor (\mathbf{b} - \mathbf{S}^\top \mathbf{a})/\Delta \rceil \bmod p$ and the corresponding decryption error $\mathbf{z} = \mathbf{b} - \mathbf{S}^\top \mathbf{a} - \Delta \cdot \boldsymbol{\mu} \bmod q$. If $\|\mathbf{z}\|_\infty > r$, return $\perp$. Otherwise, return $\boldsymbol{\mu} \in \mathbb{Z}_p^\ell$.

CORRECTNESS. The scheme requires $r < \Delta/2 \leq \frac{q}{2p}$ to enable correct decryption. A small correctness error is induced by the fact that honestly generated ciphertexts may be rejected at step 1 in the unlikely event that $\|\mathbf{a}\|_\infty < 4r$. We show that this extra correctness error is negligible under the LWE assumption.

Lemma 3. *Assume $n > \lambda/\log(q/8r)$. Under the $\mathsf{LWE}_{n,m,q,\chi_s,\chi_e}$ assumption with $\chi_s = U(\{0,1\})$ and $\chi_e = D_{\mathbb{Z},\alpha q}$, the probability that an honestly generated ciphertext is rejected at step 1 of $\mathsf{Decrypt}$ is negligible.*

Proof. For a truly random vector $\mathbf{a} \sim U(\mathbb{Z}_q^n)$, we have $\|\mathbf{a}\|_\infty < 4r$ with probability $(8r)^n/q^n$, which is exponentially small as a function of λ if $n > \lambda/\log(q/8r)$. Under the LWE assumption with binary secrets, a vector $\mathbf{a}$ produced by $\mathsf{Encrypt}$ is computationally indistinguishable from a random vector over $\mathbb{Z}_q^n$.

More precisely, assuming that honestly generated ciphertexts get rejected with non-negligible probability ϵ at step 1 of $\mathsf{Decrypt}$, we can build an LWE distinguisher $\mathcal{B}$ with advantage $\epsilon - 2^{-\lambda}$. On input of an LWE instance $(\mathbf{A}, \mathbf{a})$ where $\mathbf{a}$ is either a truly random vector $\mathbf{a} \sim U(\mathbb{Z}_q^n)$ or a pseudorandom vector of the form $\mathbf{a} = \mathbf{A} \cdot \mathbf{r} + \mathbf{e}_0$ with $\mathbf{r} \sim U(\{0,1\}^m)$ and $\mathbf{e}_0 \sim D_{\mathbb{Z}^n,\alpha q}$, the distinguisher $\mathcal{B}$ returns[3] 1 if $\|\mathbf{a}\|_\infty < 4r$ and 0 otherwise. □

PARAMETERS. Since the discrete Gaussian distribution $\chi_e = D_{\mathbb{Z},\alpha q}$ is B-bounded (i.e., $\Pr_{x \hookleftarrow \chi_e}[|x| > B] \leq \mathsf{negl}(\lambda)$) for $B = \alpha q\sqrt{2\lambda}$ by Lemma 1, we can set $r = B((m+n)+1)$. Here, the matrix $\mathbf{A}$ does not have to be very wide and we can set $m = n$ (rather than $m = \Omega(n \cdot \log q)$ if we were to use the Leftover Hash Lemma to make $\mathbf{Ar}$ statistically uniform in the ciphertext).

Theorem 1. *The above scheme is δ-snake-eye resistant under the CorLWE assumption if $\delta \leq 2^{-\ell} + \mathsf{negl}(\lambda)$. Concretely, we have $\delta \leq 2^{-\ell} + \mathbf{Adv}^{\mathcal{A},\mathsf{CorLWE}}_{n,m,q,\chi_s,\chi_e}(\lambda)$ with $\chi_s = U(\{0,1\})$ and $\chi_e = D_{\mathbb{Z},\alpha q}$.*

Proof. Assuming that an adversary $\mathcal{A}$ can break the snake-eye resistance of the scheme with non-negligible advantage δ, we build a distinguisher $\mathcal{B}$ with advantage $\geq \delta - 2^{-\ell}$ for the CorLWE problem with binary secrets.

Algorithm $\mathcal{B}$ is given a CorLWE instance $(\mathbf{A}_1, \mathbf{U}_1, \mathbf{A}_2, \mathbf{U}_2, \bar{\mathbf{S}})$, where $\mathbf{A}_1, \mathbf{A}_2 \sim U(\mathbb{Z}_q^{n\times m})$, $\bar{\mathbf{S}} \in \{-1,0,1\}^{n\times\ell}$, and $\mathbf{U}_1, \mathbf{U}_2 \in \mathbb{Z}_q^{\ell\times m}$. The goal of $\mathcal{B}$ is to decide if $\mathbf{U}_1, \mathbf{U}_2 \sim U(\mathbb{Z}_q^{\ell\times m})$ or if

$$\mathbf{U}_j^\top = \mathbf{A}_j^\top \mathbf{S}_j + \mathbf{E}_j \qquad \forall j \in \{1,2\}, \tag{3}$$

for some noise matrices $\mathbf{E}_1, \mathbf{E}_2 \sim D^\ell_{\mathbb{Z}^m,\alpha q}$ and random binary $\mathbf{S}_1, \mathbf{S}_2 \in \{0,1\}^{n\times\ell}$ such that $\bar{\mathbf{S}} = \mathbf{S}_1 - \mathbf{S}_2$. Note that, in both cases, we have $\bar{\mathbf{S}} \sim \mathcal{D}^{n\times\ell}$, where $\mathcal{D}$ is the distribution over $\{-1,0,1\}$ from which 0 is sampled with probability $1/2$ and $1, -1$ are sampled with probability $1/4$ each.

To do this, $\mathcal{B}$ runs $\mathcal{A}$ on input of public keys $\mathsf{pk}_1 = (\mathbf{A}_1, \mathbf{U}_1)$ and $\mathsf{pk}_2 =$

[3] Here, $\mathbf{B}$ does not need the secret key since $\mathbf{S}$ is not used until step 2 of $\mathsf{Decrypt}$. The distinguisher does not even have to build an entire ciphertext $(\mathbf{a}, \mathbf{b})$.

$(\mathbf{A}_2, \mathbf{U}_2)$. Then, $\mathcal{A}$ is expected to output a ciphertext $\mathsf{ct} = (\mathbf{a}, \mathbf{b})$ that breaks the snake-eye resistance of the scheme. At this point, $\mathcal{B}$ computes $\mathbf{w} = \bar{\mathbf{S}}^\top \cdot \mathbf{a} \bmod q$. If $\|\mathbf{a}\|_\infty \geq 4r$ and $\|\mathbf{w}\|_\infty \leq 2r$, $\mathcal{B}$ outputs 1 to indicate that $\mathbf{U}_1$ and $\mathbf{U}_2$ are pseudorandom matrices of the form (3). Otherwise, it outputs 0, betting that $\mathbf{U}_1$ and $\mathbf{U}_2$ are uniformly random matrices. We claim that $\mathcal{B}$'s advantage as a CorLWE distinguisher is at least $\delta - 2^{-\ell}$.

Indeed, if $\mathbf{U}_1$ and $\mathbf{U}_2$ are pseudorandom matrices of the form (3), then pk_1 and pk_2 are distributed as in the real scheme so $\mathcal{A}$ succeeds with probability δ. When it succeeds, it outputs a ciphertext $(\mathbf{a}, \mathbf{b}) \in \mathbb{Z}_q^n \times \mathbb{Z}_q^\ell$ satisfying the conditions $\|\mathbf{a}\|_\infty \geq 4r$ and $\mathbf{b} - \mathbf{S}_j^\top \cdot \mathbf{a} \bmod q = \Delta\boldsymbol{\mu} + \mathbf{z}_j$ for each $j \in \{1, 2\}$, for some plaintext $\boldsymbol{\mu} \in \mathbb{Z}_p^\ell$ and decryption errors $\mathbf{z}_1, \mathbf{z}_2 \in \mathbb{Z}^\ell$ such that $\|\mathbf{z}_1\|_\infty, \|\mathbf{z}_2\|_\infty \leq r$. We note that the latter conditions imply $\|\bar{\mathbf{S}}^\top \mathbf{a} \bmod q\|_\infty \leq 2r$. Indeed, we have

$$\begin{aligned}
\|\bar{\mathbf{S}}^\top \mathbf{a} \bmod q\|_\infty &= \|(\mathbf{S}_1^\top - \mathbf{S}_2^\top) \cdot \mathbf{a} \bmod q\|_\infty \\
&= \|(\mathbf{b} - \mathbf{S}_2^\top \mathbf{a}) \bmod q) - (\mathbf{b} - \mathbf{S}_1^\top \mathbf{a}) \bmod q) \bmod q\|_\infty \\
&= \|(\mathbf{b} - \mathbf{S}_2^\top \mathbf{a}) \bmod q) - (\mathbf{b} - \mathbf{S}_1^\top \mathbf{a}) \bmod q) + q \cdot \boldsymbol{\theta}\|_\infty \\
&= \|(\Delta\boldsymbol{\mu} + \mathbf{z}_2) - (\Delta\boldsymbol{\mu} + \mathbf{z}_1) + q \cdot \boldsymbol{\theta}\|_\infty = \|(\mathbf{z}_2 - \mathbf{z}_1) + q \cdot \boldsymbol{\theta}\|_\infty
\end{aligned}$$

with $\boldsymbol{\theta} \in \{-1, 0, 1\}^\ell$. Since $\|\mathbf{z}_2 - \mathbf{z}_1\|_\infty \leq 2r \leq q/2$ and $\bar{\mathbf{S}}^\top \mathbf{a} \bmod q \in (q/2, q/2]^\ell$, we can only have $\boldsymbol{\theta} = \mathbf{0}^\ell$ and thus $\|\bar{\mathbf{S}}^\top \mathbf{a} \bmod q\|_\infty \leq 2r$. Hence, $\mathcal{B}$ outputs 1 whenever $\mathcal{A}$ succeeds.

In contrast, if $\mathbf{U}_1, \mathbf{U}_2 \sim U(\mathbb{Z}_q^{\ell \times m})$, the matrix $\bar{\mathbf{S}}^\top \sim \mathcal{D}^{\ell \times n}$ is completely independent of $\mathcal{A}$'s view and the distribution of $\bar{\mathbf{S}}^\top \mathbf{a} \bmod q$ is the same as if $\bar{\mathbf{S}}$ was chosen *after* $\mathcal{A}$ outputs $\mathbf{a}$. By applying Lemma 2 with $\mathbf{u} = \mathbf{0}$, the probability to have $\|\bar{\mathbf{S}}^\top \mathbf{a}\|_\infty \leq \frac{1}{2}\|\mathbf{a}\|_\infty$ is at most $2^{-\ell}$. So, if $\|\mathbf{a}\|_\infty \geq 4r$, the probability that $\mathcal{B}$ outputs 1 (i.e., that $\|\bar{\mathbf{S}}^\top \mathbf{a}\|_\infty \leq 2r$) is bounded by $2^{-\ell}$. This yields the stated lower bound on $\mathcal{B}$'s distinguishing advantage against CorLWE. □

Difficulty of a Ring Version. The proof of Theorem 1 does not extend to a ring version of the scheme as it requires the rows of $\bar{\mathbf{S}}^\top$ to be independent in order to apply Lemma 2 (in the ring setting, $\mathbf{S}_1^\top$ and $\mathbf{S}_2^\top$ would become structured matrices with correlated rows). Also, the reduction [34] from LWE to CorLWE is not known to carry over to the ring setting.

Tradeoff Between Efficiency and False Positive Rate. In the statement of Theorem 1, the bound on δ contains a term $2^{-\ell}$ that can be seen as the false positive rate in the resulting private signaling scheme. To obtain a better efficiency, we may want to tolerate a relatively small false positive rate like 2^{-30} and set $\ell = 30$ rather than $\ell \approx 128$. Then, Theorem 1 ensures that the effective false positive rate cannot be significantly higher than 2^{-30} as long as the CorLWE assumption holds.

Achieving Strong Robustness. While Theorem 1 only shows snake-eye resistance, the scheme can be made strongly robust via the technique suggested in [34, Section 9.2] for PVW and LWEmongrass. The idea is to enlarge the message space by λ extra components, which are required to be zeroes in legitimate

plaintexts. A plaintext $\boldsymbol{\mu} \in \mathbb{Z}_p^\ell$ is encoded as $\hat{\boldsymbol{\mu}} = (\boldsymbol{\mu} \mid \mathbf{0}^\lambda) \in \mathbb{Z}_p^{\ell+\lambda}$ and encrypted into $(\mathbf{a}, \mathbf{b}) = (\mathbf{A}\cdot\mathbf{r} + \mathbf{e}_0, \mathbf{U}\cdot\mathbf{r} + \mathbf{e}_1 + \Delta\hat{\boldsymbol{\mu}}) \in \mathbb{Z}_q^n \times \mathbb{Z}_q^{\ell+\lambda}$. The decryption algorithm remains unchanged except that it returns $\perp$ if the rightmost λ components of $\hat{\boldsymbol{\mu}}$ differ from $\mathbf{0}^\lambda$. As in [34], this technique of achieving robustness increases the size of $\mathbf{U}$ by λm additional elements of $\mathbb{Z}_q$.

3.2 Binomially Distributed LWE Secrets

This section generalizes the result of Section 3.1 to use a binomially distributed secret key. Such keys are notably used in the Kyber construction [13]. Firstly, we prove Lemma 2 using a centered binomial distribution and then generalize existing results on correlation LWE.

Lemma 4 ([40, Lemma 2.5]). *Fix positive integers $\eta, n, \ell, q \in \mathbb{Z}$ and any vector $\mathbf{w} \in [\pm q/2]^n$. Then, we have $\Pr_{\mathbf{R}}\left[\left\|\mathbf{R}\cdot\mathbf{w} \bmod q\right\|_\infty < \frac{1}{2}\cdot\left\|\mathbf{w}\right\|_\infty\right] < 2^{-\ell}$ over the random choice of $\mathbf{R} \hookleftarrow \mathcal{B}_\eta^{\ell\times n}$.*

Proof. For $\eta = 1$, Lemma 2 with $\boldsymbol{u} = \mathbf{0}$ implies the result and we thus assume that $\eta > 1$. We use the fact that, if $x \hookleftarrow \mathcal{B}_{\eta-1}$ and $y \hookleftarrow \mathcal{B}_1$, then the distribution of $x + y$ is $\mathcal{B}_\eta$. From this, it follows that

$$\begin{aligned}\Pr_{\mathbf{R}\sim\mathcal{B}_\eta^{\ell\times n}}&\left[\left\|\mathbf{R}\mathbf{w} \bmod q\right\|_\infty < \frac{1}{2}\left\|\mathbf{w}\right\|_\infty\right]\\ &= \Pr_{\substack{\mathbf{R}_1\sim\mathcal{B}_1^{\ell\times n}\\ \mathbf{R}_2\sim\mathcal{B}_{\eta-1}^{\ell\times n}}}\left[\left\|\mathbf{R}_1\mathbf{w} + \mathbf{R}_2\mathbf{w} \bmod q\right\|_\infty < \frac{1}{2}\left\|\mathbf{w}\right\|_\infty\right] < 2^{-\ell}\end{aligned}$$

where the inequality follows from applying Lemma 2 with $\mathbf{u} = \mathbf{R}_2\mathbf{w} \bmod q$ and using the fact that $\mathbf{R}_1$ is sampled independently of $\mathbf{u}$. □

We then show that CorLWE with binomially distributed secrets is as hard as CorLWE with binary secrets. Using [34, Lemma 2], this shows that the former is at least as hard as binary-secret LWE. One could reduce to CorLWE with secret distribution $\mathcal{B}_1$, but this would lead to an LWE distribution where different entries follow different distributions (which is not the case for Short-key LWE).

We note that the reduction essentially relates a correlated LWE problem with non-binary secret to an LWE problem with a binary secret. It does not show that larger secrets in the correlated LWE setting lead to a harder problem.

Lemma 5. *For any positive integers n, m, q, η and any error distribution χ, $\mathsf{CorLWE}_{n,m,q,\mathcal{B}_\eta,\chi}$ is at least as hard as $\mathsf{CorLWE}_{n,m,q,U(\{0,1\}),\chi}$.*

Proof. Starting with a $\mathsf{CorLWE}_{n,m,q,U(\{0,1\}),\chi}$ instance $(\mathbf{A}_1, \mathbf{A}_2, \mathbf{b}_1, \mathbf{b}_2, \bar{\mathbf{s}})$, we perform a simple mapping. Firstly, we sample $(\mathbf{s}_i^{(1)}, \mathbf{s}_i^{(2)})_{i=2}^\eta, (\mathbf{t}_i^{(1)}, \mathbf{t}_i^{(2)})_{i=1}^\eta$ as uniform binary vectors. Next, we compute $\mathbf{s}_1' = -\mathbf{t}_1^{(1)} + \sum_{i=2}^\eta(\mathbf{s}_i^{(1)} - \mathbf{t}_i^{(1)})$ and $\mathbf{s}_2' = -\mathbf{t}_1^{(2)} + \sum_{i=2}^\eta(\mathbf{s}_i^{(2)} - \mathbf{t}_i^{(2)})$. Finally, we output

$$\left(\mathbf{A}_1, \mathbf{A}_2, \mathbf{b}_1 + \mathbf{A}_1\cdot\mathbf{s}_1', \mathbf{b}_2 + \mathbf{A}_2\cdot\mathbf{s}_2', \bar{\mathbf{s}} + \mathbf{s}_1' - \mathbf{s}_2'\right),$$

which is a valid pseudorandom instance of $\mathsf{CorLWE}_{n,m,q,\mathcal{B}_\eta,\chi}$ if $\mathbf{b}_i = \mathbf{A}_i^\top \cdot \mathbf{s}_i + \mathbf{e}_i$ for $\mathbf{s}_i \sim U(\{0,1\}^n)$ and $\mathbf{e}_i \sim \chi^m$ for each $i \in \{1,2\}$. If $\mathbf{b}_1, \mathbf{b}_2$ are uniformly random, so are $\mathbf{b}_1 + \mathbf{A}_1 \cdot \mathbf{s}'_1$ and $\mathbf{b}_2 + \mathbf{A}_2 \cdot \mathbf{s}'_2$ and $\bar{\mathbf{s}} + \mathbf{s}'_1 - \mathbf{s}'_2$ is the difference between two samples of $\mathcal{B}_\eta^n$, as required. □

By [34, Lemma 2], we have the following corollary.

Corollary 1. *For any positive integers n, m, q, η and error distribution χ, the $\mathsf{CorLWE}_{n,m,q,\mathcal{B}_\eta,\chi_e}$ problem is at least as hard as $\mathsf{LWE}_{n',2m,q,U(\{0,1\}),\chi}$ in dimension $n' = 2n/\epsilon$ for any constant $0 < \epsilon < 1$.*

Lemma 6. *The* shortLWE *construction with secret distribution $\mathcal{B}_\eta$ is δ-snake-eye resistant under the $\mathsf{CorLWE}_{n,m,q,\mathcal{B}_\eta,\chi_e}$ assumption for $\delta < 2^{-\ell} + \mathsf{negl}$.*

Proof. We build a CorLWE distinguisher $\mathcal{B}$ from a snake-eye resistance adversary $\mathcal{A}$. The description of $\mathcal{B}$ is *identical* to the reduction in Theorem 1. We use the same notation and write $\mathcal{B}$'s input as $(\mathbf{A}_1, \mathbf{U}_1, \mathbf{A}_2, \mathbf{U}_2, \bar{\mathbf{S}})$.

If $\mathbf{U}_1$ and $\mathbf{U}_2$ are pseudorandom matrices of the form (3), and $\mathcal{A}$ succeeds outputting $(\mathbf{a}, \mathbf{b}) \in \mathbb{Z}_q^n \times \mathbb{Z}_q^\ell$, we find that $\|\bar{\mathbf{S}}^\top \mathbf{a} \bmod q\| \leq 2r$ by the same argument as in Theorem 1. This implies that $\mathcal{B}$ outputs 1 whenever $\mathcal{A}$ succeeds.

On the other hand, if $\mathbf{U}_1$ and $\mathbf{U}_2$ are uniform matrices, $\bar{\mathbf{S}} \sim \mathcal{B}_{2\eta}$ is completely independent of $\mathcal{A}$'s view. By Lemma 4 the probability that $\|\bar{\mathbf{S}}^\top \mathbf{a}\|_\infty \leq \frac{1}{2}\|\mathbf{a}\|_\infty$ is at most $2^{-\ell}$. So, if $\|\mathbf{a}\|_\infty \geq 4r$, the probability that $\mathcal{B}$ outputs 1 (i.e., that $\|\bar{\mathbf{S}}^\top \mathbf{a}\|_\infty \leq 2r$) is bounded by $2^{-\ell}$. Hence, $\mathcal{B}$'s advantage is at least $\delta - 2^{-\ell}$. □

4 A Strongly Robust PKE Scheme from Ring LWE

In this section, we build a strongly robust public-key encryption scheme based on the RLWE assumption. Strong robustness is achieved by means of a decryption algorithm that recovers the sender's random coins on which it performs a smallness check before deciding to reject the ciphertext or not. Under the RLWE assumption, we then prove the infeasibility of finding valid randomness leading to an encryption collision for two honestly generated public keys.

We consider a variant of the randomness-recovering RLWE-based encryption scheme proposed by Hoffmann *et al.* [24], which is itself a randomness-recovering variant of the LPR cryptosystem [41]. As described in [24], the original scheme is a hybrid KEM-DEM construction with IND-CCA2 security. Here, we remove its symmetric components as they are ill-suited to our purposes. As a result, we obtain a scheme which is additively homomorphic up to small changes in the norm check bounds[4] and strongly robust.

4.1 Description

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

[4] In order to enable ℓ_a additions on fresh ciphertexts, we can set $B = 2\ell_a \alpha q \sqrt{2\lambda}$.

1. Choose a dimension $n \in \mathsf{poly}(\lambda)$, which is a power of 2, and a prime modulus $q \in \mathsf{poly}(\lambda)$ such that $q = 1 \bmod 2n$, which define the rings $R = \mathbb{Z}[X]/(X^n + 1)$ and $R_q = R/(qR)$.
2. Choose a noise distribution χ over R. We assume that $\chi = D^{\mathsf{coeff}}_{\mathbb{Z}^n,\alpha q}$, for some noise parameter $\alpha \in (0, 1)$, so that χ is B-bounded for infinity norm bound $B = \alpha q\sqrt{2\lambda}$ by Lemma 1. Choose moduli $p, t \in \mathsf{poly}(\lambda)$.

Return the public parameters $\mathsf{pp} := (n,\ q,\ p,\ t,\ \chi,\ B)$.

Keygen(pp)**:** Given public parameters pp,

1. Sample $a \hookleftarrow U(R_q)$.
2. Sample $s \hookleftarrow D^{\mathsf{coeff}}_{\mathbb{Z}^n,\alpha q}, e \hookleftarrow D^{\mathsf{coeff}}_{\mathbb{Z}^n,\alpha q}$ and compute $b = p\cdot(a\cdot s+e)$. If $b \notin R_q^\times$, retry with fresh s, e.

Return the key pair $(\mathsf{pk}, \mathsf{sk})$ where $\mathsf{sk} := s \in R$ and

$$\mathsf{pk} := \left(\mathsf{pp},\ (a, b) \in R_q \times R_q^\times\right).$$

Encrypt(pk, μ)**:** Given a public key pk and a message $\mu \in R_t$, sample $r \hookleftarrow D^{\mathsf{coeff}}_{\mathbb{Z}^n,\alpha q}, e_1, e_2 \hookleftarrow D^{\mathsf{coeff}}_{\mathbb{Z}^n,\alpha q}$. Compute

$$c_1 = a \cdot r + e_1, \qquad c_2 = b \cdot r + e_2 + \Delta \cdot \mu, \tag{4}$$

where $\Delta = \lfloor q/t \rceil$. Output the ciphertext $\mathsf{ct} = (c_1, c_2) \in R_q^2$.

Decrypt$(\mathsf{sk}, \mathsf{ct})$**:** Given $\mathsf{sk} = s \in R$ and $\mathsf{ct} = (c_1, c_2) \in R_q^2$, return $\perp$ if $\|c_1\|_\infty < B$ when $\phi(c_1)$ is viewed as a vector in $(-q/2, q/2]^n$. Otherwise,

1. Compute $\zeta = c_2 - p \cdot c_1 \cdot s \bmod q$
2. Compute $\mu = \lfloor \zeta/\Delta \rceil \bmod t$.
3. Compute $e_2 = (\zeta - \Delta\mu \bmod q) \bmod p$. Return $\perp$ if $\|e_2\|_\infty > B$.
4. Compute $r = b^{-1} \cdot (c_2 - \Delta\mu - e_2) \in R_q$ and return $\perp$ if $\|r\|_\infty > B$.
5. Compute $e_1 = c_1 - a \cdot r \bmod q$ and return $\perp$ if $\|e_1\|_\infty > B$.
6. Return $\mu \in R_t$.

4.2 Correctness and Security

Lemma 7. *Let $\sigma = \alpha q$ and assume that the $\mathsf{RLWE}_{n,q,1,\chi}$ problem is hard. The scheme provides correctness with overwhelming probability over the randomness of* Keygen *and* Encrypt *if $p/2 > \alpha q\sqrt{2\lambda}$ and* [5] *$\Delta > 2(2p \cdot (\alpha q)^2 \cdot 2n + \alpha q \cdot \sqrt{2\lambda})$.*

Proof. We begin by bounding the probability that Decrypt aborts due to the check $\|c_1\|_\infty < B$. In particular, this probability is $\rho := \Pr[\|a \cdot r + e_1\|_\infty < B]$ where $a \hookleftarrow U(R_q), r, e_1 \hookleftarrow \chi$. Clearly, $a \cdot r + e_1$ is a $\mathsf{RLWE}_{n,1,q,\chi}$ sample. Defining $\rho' := \Pr[\|u\|_\infty < B]$ for $u \hookleftarrow R_q$, we have $\rho' = ((2B - 1)/q)^n = ((2\sigma\sqrt{2\lambda} - 1)/q)^n \leq 2^{-\lambda}$. Considering a $\mathsf{RLWE}_{n,q,1,\chi}$ adversary $\mathcal{C}$ that takes a challenge $(\bar{a}, \bar{b})$ and outputs 1 if and only if $\|\bar{b}\|_\infty < B$, we obtain $\rho \leq 2^{-\lambda} + \mathbf{Adv}^{\mathcal{C},\mathsf{RLWE}}_{n,1,q,\chi}(\lambda)$.

[5] We use Lemma 1 with $k = \sqrt{2}$ for the Euclidean bound of a discrete Gaussian but different constants may be used

Next, at step 2 of Decrypt, we have $e_2 \hookleftarrow \chi$ and

$$\zeta = c_2 - p \cdot c_1 \cdot s \bmod q = p \cdot (e \cdot r - e_1 \cdot s) + e_2 + \Delta \cdot \mu \bmod q$$

so that step 2 recovers the correct $\mu \in R_t$ if $\|p \cdot (e \cdot r - e_1 \cdot s) + e_2\|_\infty < \Delta/2$. This inequality holds by the Cauchy-Schwarz inequality and Lemma 1 with $k = \sqrt{2}$ for the Euclidean norm bound of a Gaussian vector. Then, at step 3, we have $e_2 = (\zeta - \Delta\mu \bmod q) \bmod p$ if $\|p \cdot (e \cdot r - e_1 \cdot s) + e_2\|_\infty < q/2$ and $\|e_2\|_\infty < p/2$, in which case step 4 recovers the correct $r \in R$. This completes the proof. □

Parameters. Correctness requires that $\Delta = \lfloor q/t \rceil > 2(2p \cdot (\alpha q)^2 \cdot 2n + \alpha q \cdot \sqrt{2\lambda})$ and $p/2 > \alpha q\sqrt{2\lambda}$. We can thus set $p > 2\sqrt{2\lambda}\alpha q$ and $\Delta > 2\alpha q\sqrt{2\lambda} \cdot (8(\alpha q)^2 n + 1)$. If we use a noise with standard deviation $\alpha q = \Omega(\sqrt{n})$ and a plaintext modulus $t = O(1)$, we can choose $p = \Theta(n)$ and $q = \Theta(n^{3.5})$.

The proof of IND-CPA security is almost identical to the proof of [24, Theorem 1] and can be found in the full version.

Theorem 2. *Under the* $\mathsf{RLWE}_{n,2k,q,\chi}$ *assumption with* $k = \lceil \lambda / \log(q/n) \rceil$ *and* $\chi = D^{\mathsf{coeff}}_{\mathbb{Z}^n, \alpha q}$*, the scheme provides IND-CPA security with pseudorandom ciphertexts.* (The proof is given in the full version.)

We now prove the strong robustness property (which implies snake-eye resistance). The scheme of [24] can be proven strongly robust in the same way.

Theorem 3. *The above scheme is strongly robust under the* $\mathsf{RSIS}_{n,k,q,\beta}$ *assumption with* $k = O(\lambda / \log(q/n))$ *and* $\beta = 2\alpha q\sqrt{2\lambda}$ *if* $q > n$*.*

Proof. Using the randomness-recovering property, we can prove strong robustness under the RSIS assumption. On input of two honestly generated public keys (a, b) and (a', b'), an adversary can only break strong robustness by outputting a ciphertext $(c_1, c_2) \in R_q^2$ for which there exist (r, e_1), (r', e_1') such that

$$c_1 = a \cdot r + e_1 = a' \cdot r' + e_1' \bmod q \tag{5}$$

where $\|r\|_\infty, \|r'\|_\infty, \|e\|_\infty, \|e'\|_\infty \leq B$. We show that it would contradict RSIS.

We build an algorithm $\mathcal{B}$ that uses a robustness adversary $\mathcal{A}$ to solve an RSIS instance. The reduction $\mathcal{B}$ takes as input an instance with $k \geq \lceil \lambda / \log(q/n) \rceil + 2$ ring elements $(a_1, \ldots, a_k) \in R_q^k$ with the goal of finding $\bar{\boldsymbol{e}} = (\bar{e}_1, \ldots, \bar{e}_k) \in R^k$ such that $\sum_{i=1}^k a_i \cdot \bar{e}_i = 0 \bmod q$ with $\|\bar{e}_i\|_\infty \leq \beta$ for each $i \in [k]$.

To do that, $\mathcal{B}$ determines the smallest $j \in [3, k]$ such that a_j is invertible in R_q. If none of $\{a_i\}_{i=3}^k$ is invertible, $\mathcal{B}$ aborts. Since $\{a_i\}_{i=3}^k$ are sampled independently from $U(R_q)$ and $k \geq \lceil \lambda / \log(q/n) \rceil + 2$, the probability that none of them is a unit is bounded by $(n/q)^{k-2} < 2^{-\lambda}$ (the probability is maximal in fully splitting rings where a random a is non-invertible with probability n/q).

If $\mathcal{B}$ does not abort, it generates the public keys by setting $a = a_1 \cdot a_j^{-1} \in R_q$ and $a' = a_2 \cdot a_j^{-1} \in R_q$. Next, it honestly generates $b = p \cdot (a \cdot s + e) \bmod q$ and $b' = p \cdot (a' \cdot s' + e') \bmod q$ by sampling $s, s' \hookleftarrow D^{\mathsf{coeff}}_{\mathbb{Z}^n, \alpha q}$, and $e, e' \hookleftarrow D^{\mathsf{coeff}}_{\mathbb{Z}^n, \alpha q}$. Then,

it feeds $\mathcal{A}$ with $\mathsf{pk} = (a, b)$ and $\mathsf{pk}' = (a', b')$, which are distributed exactly as in the real scheme, and keeps $\mathsf{sk} = s$ and $\mathsf{sk}' = s'$ for later use.

By hypothesis, $\mathcal{A}$ is expected to output a ciphertext $(c_1, c_2) \in R_q^2$ that decrypts to valid (i.e., non-$\perp$) messages under both sk and sk'. By running the $\mathsf{Decrypt}$ algorithm using sk and sk' (which are known to $\mathcal{B}$), $\mathcal{B}$ recovers pairs (r, e_1) and (r', e_1') for which a collision (5) occurs. This allows $\mathcal{B}$ to obtain a short vector $\boldsymbol{v} \triangleq (r, -r', e_1 - e_1')^\top \in R^3$ such that

$$\left[\begin{array}{c|c|c} a & a' & 1 \end{array}\right] \cdot \boldsymbol{v} = 0 \bmod q$$

and thus $\langle (a_1, a_2, a_j), \boldsymbol{v} \rangle = 0 \bmod q$. We cannot have $\boldsymbol{v} = \boldsymbol{0}$ since this would imply $r = r' = 0$ and $e_1 = e_1' = c_1$, so that c_1 would be rejected at the very beginning of $\mathsf{Decrypt}$. Then, $\mathcal{B}$ can solve its RSIS instance by outputting $\bar{\boldsymbol{e}} \triangleq (\bar{e}_1, \ldots, \bar{e}_k)^\top \in R^k$ such that $(\bar{e}_1, \bar{e}_2, \bar{e}_j) = (r, -r', e_1 - e_1')$ and $\bar{e}_i = 0$ for all $i \in [k] \setminus \{1, 2, j\}$. Note that $\bar{\boldsymbol{e}}$ is a valid solution since $\|r\|_\infty, \|r'\|_\infty \leq \beta$ and $\|e_1 - e_1'\|_\infty \leq \beta$. □

A Variant With Statistical Robustness. In the full version, we show that the scheme can be made statistically strongly robust. This variant is less NTT-friendly as it does not support fully splitting rings.

4.3 Snake-eye Resistant Private Signaling

In OMR protocols, a signal consists of an encryption of a zero plaintext and the detection algorithm checks that a given ciphertext actually decrypts to zero. For this application to the generation of private signals (a.k.a. clues) in OMR protocols, we can simplify the decryption algorithm and remove step 2.

The ParGen and Keygen algorithms are exactly as in Section 4.1. The signaling and detection algorithms proceed as follows. The snake-eye resistance of the scheme follows from the strong robustness of the underlying PKE scheme. Its receiver privacy property immediately follows from the pseudorandomness of ciphertexts in the PKE system.

$\mathsf{Signal}(\mathsf{pk})$: Given pk, sample $r \hookleftarrow D_{\mathbb{Z}^n, \alpha q}^{\mathsf{coeff}}, e_1, e_2 \hookleftarrow D_{\mathbb{Z}^n, \alpha q}^{\mathsf{coeff}}$. Then, compute

$$c_1 = a \cdot r + e_1, \qquad c_2 = b \cdot r + e_2, \tag{6}$$

Output the signal $\mathsf{sig} = (c_1, c_2) \in R_q^2$.

$\mathsf{Detect}(\mathsf{sk}, \mathsf{sig})$: Given $\mathsf{sk} = s \in R$ and $\mathsf{sig} = (c_1, c_2) \in R_q^2$, return 0 if $\|c_1\|_\infty < B$ when c_1 is seen as a polynomial with coefficients in $(-q/2, q/2]$. Otherwise,

1. Compute $\zeta = c_2 - p \cdot c_1 \cdot s \bmod q$
2. Compute $e_2 = \zeta \bmod p$. Return 0 if $\|e_2\|_\infty > B$.
3. Compute $r = b^{-1} \cdot (c_2 - e_2) \in R_q$ and return 0 if $\|r\|_\infty > B$.
4. Compute $e_1 = c_1 - a \cdot r \bmod q$ and return 0 if $\|e_1\|_\infty > B$.
5. Return 1.

With the above simplifications, the proof of Theorem 3 carries over and ensures that, under the RLWE assumption, no PPT adversary can output a valid signal $(c_1, c_2) \in R_q^2$ for two distinct key pairs $(\mathsf{pk}, \mathsf{sk})$ and $(\mathsf{pk}', \mathsf{sk}')$.

4.4 Homomorphic Evaluation of the Detection Circuit

For an OMR application, it is required to homomorphically evaluate the decryption circuit given a ciphertext (c_1, c_2) and an FHE encryption of the secret key s. As is typical, we use BFV with a plaintext space of $\mathbb{Z}_q^N$ to utilize its SIMD capability. This plaintext space is isomorphic to $\mathbb{Z}_q[X]/(X^N+1)$ where we assume that $q = 1 \bmod 2N$ and $N > n$ a power-of-two. This standard choice of N is widely used in practice (e.g., in the OpenFHE library [4]) and leads to efficient ring multiplication. Note that q is both the BFV plaintext modulus and the "ciphertext modulus" of the signaling scheme. The secret key s is encrypted by repeating its coefficient vector in $\mathbb{Z}_q^n$ a total of N/n times to derive a plaintext in $\mathbb{Z}_q^N$. We only give a brief overview of the homomorphic evaluation here as we improve on the current private signaling scheme in Section 5.1. In particular, we leave details on how to process a *batch* of signals together to the full version, while (unrealistically) ignoring the cost of linear operations for now. All ring multiplications in the signaling scheme will be treated as matrix-vector products.

We can begin by computing an encryption of $\zeta = c_2 - p \cdot c_1 \cdot s \bmod q$ by linearity. This is done in part by multiplying a plaintext anti-circulant matrix derived from c_1 by the encryption of s using the techniques outlined in Section 2.6. From here, we must compute an encryption of $e_2 = \zeta \bmod p$. This step is not a linear function due to the reduction modulo p, but a method costing around $\sqrt{q-3}$ multiplications and $\lceil \log(q-1) \rceil + 1$ levels exists [25, Proposition 1]. Next note that we only ever decrypt to 1 when e_2 and $r = b^{-1} \cdot (c_2 - e_2) \bmod q$ satisfy their range checks simultaneously. We can compute an encryption of this candidate r using the encryption of e_2 and linearity as b and c_2 are known. Finally, computing an encryption of e_1 from r is also possible due to linearity. What is left is to perform the three range checks on (e_1, e_2, r). Multiplying the results of these checks either gives 1 if all range checks pass, or 0 otherwise. Counting all individual coefficients, there are $3n$ range checks to combine. Therefore, combining all of the checks costs $3n-1$ multiplications and $\lceil \log(3n) \rceil$ levels.

We next overview the cost of range checks. We can define a range-check function on $\mathbb{Z}_q$ by

$$f_B(x) = \begin{cases} 1 \text{ if } x \in [-B, B] \\ 0 \text{ otherwise.} \end{cases}$$

Two methods have been proposed for homomorphic evaluation of f_B. The first is a "low-degree" method that simply computes a degree-$(q-1)$ polynomial interpolation of the above function requiring a total of $\sqrt{2 \cdot (q-1)}$ multiplications and $\lceil \log(q-1) \rceil + 1$ levels [25,47]. A second "high-degree" method [37] homomorphically computes $1 - \left(x \cdot \prod_{i=1}^{B} (x^2 - i^2) \right)^{q-1} \bmod q$. Fermat's Little Theorem guarantees that this results in the correct evaluation for prime q. Although the total degree of this polynomial is $(q-1) \cdot (2B+1)$, only $B + 1 + \lceil \log(q-1) \rceil$ multiplications are required with $1 + \lceil \log(B+1) \rceil + \lceil \log(q-1) \rceil$ levels. We focus on this latter method as it was shown to yield fewer multiplications and faster implementations [37] when $B \ll q$.

Overall, the total number of levels is

$$2\lceil\log(q-1)\rceil + 2 + \lceil\log(B+1)\rceil + \lceil\log(3n)\rceil\,.$$

When considering the batching of N signals at once using SIMD (as described in the full version), the total number of homomorphic multiplications is around $n \cdot \sqrt{q-3} + 3n \cdot (B + 1 + \lceil\log(q-1)\rceil) + 3n - 1$.

As mentioned previously, multiplicative depth places a constraint on the BFV ring degree, which in turn places a constraint on the plaintext modulus q. For example, at depths greater than 30, OpenFHE dictates a ring degree of at least $N = 2^{16}$ for security which leads to a plaintext modulus of at least $q = 786433$. This means that the depth above will always be considerably larger than $2\lceil\log(q-1)\rceil \geq 40$. Even more troubling is the $\sqrt{q-3} > 886$ number of multiplications to compute a batch of N modulo p reductions. For the simultaneous processing of N signals, the total number of homomorphic multiplications due to modulo p reduction is $n \cdot \sqrt{q-3}$. For a typical value of $n = 1024$, the number of multiplications is in the region of one million which is prohibitively expensive. To reduce this bottleneck, we present an alternative private signaling scheme in the next section.

5 Snake-eye Resistant Private Signaling Schemes with Simpler Detection Circuits

We now consider variants that do not require modular reduction or rounding in the detection circuit, and are thus more amenable to homomorphic evaluation. In the full version, we give a variant based solely on the RLWE assumption. In Section 5.1, we rely on both LWE and RLWE to decrease the number of homomorphic multiplications and levels in the detection circuit.

5.1 Optimized Scheme from RLWE and Binary-Secret LWE

For each index $i \in [n]$ and any ring element $a \in R_q$ with coefficient embedding $\phi(a) = (a_0, \ldots, a_{n-1})^\top \in \mathbb{Z}_q^n$, $\phi_{i-1}(a) = (a_i, \ldots, a_0, -a_{n-1}, \ldots, -a_{i+1})^\top$ denotes the i-th row of the matrix $\mathsf{rot}(a) \in \mathbb{Z}_q^{n\times n}$.

As in [28,30], we rely on an RLWE-based PKE scheme where each ciphertext can be seen as truncation of LPR ciphertexts [41] in the sense that only the first d components of $c_2 = b \cdot u + e_2 + \Delta \cdot r$ are included in the ciphertext. We use the truncated LPR cryptosystem to encrypt a random $\mathbf{r} \in \{0,1\}^d$. The latter is then used as a binary LWE secret to compute $\mathbf{c}_3 = \mathbf{A}^\top \cdot \mathbf{r} + \mathbf{e}_3$ for a small $\mathbf{e}_3$. Since d does not have to be a power of 2, we can choose it significantly smaller than n, which allows a shorter truncated LPR ciphertext. Importantly, the uniform matrix $\mathbf{A} \in \mathbb{Z}_q^{d\times m}$ can be derived from a random oracle (like the ring element $a \in R_q$) so as to avoid increasing its public key size. In the random-oracle-optimized variant, each public key consists of only one ring element $b \in R_q$.

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

1. Choose dimensions $d, m, n \in \mathsf{poly}(\lambda)$, where n 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 $R_q = R/(qR)$.
2. Choose noise distributions χ and χ' over R and $\mathbb{Z}$, respectively. We assume that $\chi = D^{\mathsf{coeff}}_{\mathbb{Z}^n,\alpha q}$, $\chi' = D_{\mathbb{Z}^n,\alpha' q}$ for noise parameters $\alpha, \alpha' \in (0,1)$, and that χ (resp. χ') is B-bounded (resp. B'-bounded) for infinity norm bounds $B = \alpha q\sqrt{2\lambda}$, $B' = \alpha' q\sqrt{2\lambda}$. Let B_e be an upper bound on the infinity norm of the decryption error.

Return the public parameters $\mathsf{pp} := (d,\ m,\ n,\ q,\ \chi,\ \chi',\ B,\ B',\ B_e)$.

Keygen(pp)**:** Given public parameters pp,

1. Choose a random matrix $\mathbf{A} \hookleftarrow U(\mathbb{Z}_q^{d\times m})$. Sample $a \hookleftarrow U(R_q)$.
2. Sample $s \hookleftarrow D^{\mathsf{coeff}}_{\mathbb{Z}^n,\alpha q}, e \hookleftarrow D^{\mathsf{coeff}}_{\mathbb{Z}^n,\alpha q}$ and compute $b = a \cdot s + e$.

Return the key pair $(\mathsf{pk}, \mathsf{sk})$ where $\mathsf{sk} := s \in R$ and

$$\mathsf{pk} := \big(\mathsf{pp},\ (\mathbf{A},\ (a,b)) \in \mathbb{Z}_q^{m\times d} \times R_q \times R_q\big)\,.$$

Signal(pk)**:** Choose $\mathbf{r} = (r_1, \ldots, r_d)^\top \hookleftarrow U(\{0,1\}^d)$. Then, sample $u \hookleftarrow D^{\mathsf{coeff}}_{\mathbb{Z}^n,\alpha q}$, $e_1 \hookleftarrow D^{\mathsf{coeff}}_{\mathbb{Z}^n,\alpha q}$, $\mathbf{e}_2 = (e_{2,1}, \ldots, e_{2,d})^\top \hookleftarrow D_{\mathbb{Z}^d,\alpha q}$, $\mathbf{e}_3 \hookleftarrow D_{\mathbb{Z}^m,\alpha' q}$. Compute

$$\begin{aligned} c_1 &= a \cdot u + e_1 \\ c_{2,i} &= \langle \phi_{i-1}(b), \phi(u)\rangle + e_{2,i} + \Delta \cdot r_i \qquad \forall i \in [d], \\ \mathbf{c}_3 &= \mathbf{A}^\top \cdot \mathbf{r} + \mathbf{e}_3 \end{aligned} \tag{7}$$

where $\Delta = \lfloor q/2 \rfloor$. Output the signal

$$\mathsf{sig} = (c_1, \mathbf{c}_2 = (c_{2,1}, \ldots, c_{2,d}), \mathbf{c}_3) \in R_q \times \mathbb{Z}_q^d \times \mathbb{Z}_q^m. \tag{8}$$

Detect(sk, ct)**:** Given $\mathsf{sk} = s \in R$ and $\mathsf{sig} \in R_q \times \mathbb{Z}_q^d \times \mathbb{Z}_q^m$, parse sig as in (8). Return 0 if $\|\mathbf{c}_3\|_\infty < B'$ when $\mathbf{c}_3$ is seen as a vector in $\mathbb{Z}^m \cap (-q/2, q/2]^m$.

1. Compute

$$\forall i \in [d]: \begin{array}{l} \boldsymbol{\zeta}_0[i] = c_{2,i} - \langle \phi_{i-1}(c_1), \phi(s)\rangle \bmod q \\ \boldsymbol{\zeta}_1[i] = \boldsymbol{\zeta}_0[i] - \Delta \bmod q, \end{array}$$

2. Define $\boldsymbol{\mu}_0, \boldsymbol{\mu}_1 \in \{0,1\}^d$ such that, for each $i \in [d]$,

$$\boldsymbol{\mu}_0[i] = \begin{cases} 1 & \text{if } \ \boldsymbol{\zeta}_0[i] \in [-B_e, B_e] \\ 0 & \text{otherwise} \end{cases}$$

and

$$\boldsymbol{\mu}_1[i] = \begin{cases} 1 & \text{if } \ \boldsymbol{\zeta}_1[i] \in [-B_e, B_e] \\ 0 & \text{otherwise} \end{cases}$$

Compute $\boldsymbol{\mu} = \boldsymbol{\mu}_1 - \boldsymbol{\mu}_0 \in \{-1, 0, 1\}^d$ and return 0 if $\boldsymbol{\mu} \circ \boldsymbol{\mu} \neq (1, 1, \ldots, 1)^\top$.

3. Define $\mathbf{r} \in \{0,1\}^d$ such that, for each $i \in [d]$, $\mathbf{r}[i] = 1$ if $\boldsymbol{\mu}[i] = 1$ and $\mathbf{r}[i] = 0$ if $\boldsymbol{\mu}[i] = -1$.[6]

[6] This can be done by setting $r = 2^{-1} \cdot (\phi^{-1}(\boldsymbol{\mu}) + \mathbf{1}_R) \bmod q$.

4. Compute $\mathbf{e}_3 = \mathbf{c}_3 - \mathbf{A}^\top \cdot \mathbf{r} \bmod q$ and return 0 if $\|\mathbf{e}_3\|_\infty > B'$.

In Detect, computing $\mathbf{r}$ via steps 2-3 has the same effect as computing $\mathbf{r}[i] = \lfloor (c_{2,i} - \langle \phi_{i-1}(c_1), \phi(s) \rangle \bmod q)/\Delta \rfloor \bmod 2$ for each $i \in [d]$. Since homomorphically evaluating a rounding function is somewhat expensive, the above description makes it easier to explain how we can homomorphically run it when the homomorphic rounding is replaced by two SIMD smallness tests at step 2.

Parameters. The matrix $\mathbf{A} \in \mathbb{Z}_q^{d \times m}$ does not have to be very wide. As the proof of Theorem 5 suggests, we can even have $m \leq d$ for a sufficiently large q. For example, if we take $q > 8(4B' + 1)$, we only need $m > (2d + \lambda)/3$.

5.2 Correctness and Security

Lemma 8. *The above private signaling scheme provides completeness if* $d > \lambda + 1$, $m > (\lambda + 1)/\log(\frac{q}{2B'-1})$, $B_e \leq \Delta/2$ *and* $B_e \geq 4n \cdot (\alpha q)^2 + \alpha q\sqrt{2\lambda}$.

Proof. We first argue that the likelihood of the detection algorithm outputting 0 due to $\|\mathbf{c}_3\|_\infty$ being larger than B' is negligible.

To do so, consider an arbitrary non-zero $\mathbf{r}^* \in \{0,1\}^d$ and arbitrary $\mathbf{e}_3^* \in \mathbb{Z}_q^m$. The probability that the public key component $\mathbf{A} \hookleftarrow U(\mathbb{Z}_q^{d \times m})$ happens to satisfy $\|\mathbf{A}^\top \cdot \mathbf{r}^* + \mathbf{e}_3^*\|_\infty < B'$ is $((2B' - 1)/q)^m < 2^{-(\lambda+1)}$. This follows from the fact that there is an index j such that $\mathbf{r}^*[j] = 1$, which implies that each entry of $\mathbf{A}^\top \cdot \mathbf{r}^* + \mathbf{e}_3^*$ is uniformly distributed by virtue of the j-th column of $\mathbf{A}^\top$. Therefore, considering the negligible probability that $\mathbf{r} \hookleftarrow U(\{0,1\}^d)$ could be $\mathbf{0}^d$, the probability that $\|\mathbf{c}_3\|_\infty < B'$ is at most $2^{-d} + 2^{-(\lambda+1)} \leq 2^{-\lambda}$.

Next, let $\mathbf{e_2}$ and $\mathbf{r} \in \{0,1\}^d$ denote the variables used to produce a signal $(c_1, \mathbf{c}_2, \mathbf{c}_3)$. Define $e_2 \in R_q$ and $r \in R_q$ to be the ring elements whose first d entries coincide with $\mathbf{e_2}$ and $\mathbf{r}$ respectively, but with the final $n - d$ coefficients set to 0. We then have that the vector $\mathbf{c}_2$ is a truncated version of $\phi(b \cdot u + e_2' + \Delta r')$ and that $\boldsymbol{\zeta}_0$ is simply a truncated version of

$$\begin{aligned}\phi(b \cdot u + e_2 + \Delta r - c_1 \cdot s) &= \phi\big((b - a \cdot s) \cdot u + e_2 + \Delta r - e_1 \cdot s\big) \\ &= \phi(e \cdot u - e_1 \cdot s + e_2 + \Delta r).\end{aligned}$$

Similarly, $\boldsymbol{\zeta}_1$ is a truncated version of $\phi(e \cdot u - e_1 \cdot s + e_2' + \Delta(r' - \phi^{-1}(\mathbf{1})))$. By the assumption in the lemma statement, we have the bound $\|e \cdot u + e_1 \cdot s + e_2'\|_\infty \leq B_e$ due to Lemma 1 and the Cauchy-Schwarz inequality. Now, if $\mathbf{r}_i = 0$, the detection decodes to $\boldsymbol{\mu}[i] = \boldsymbol{\mu}_1[i] - \boldsymbol{\mu}_0[i] = 0 - 1 = -1$ and recovers $\mathbf{r}_i = 0$. For $\mathbf{r}_i = 1$, the detection decodes to $\boldsymbol{\mu}[i] = 1 - 0 = 1$ and correctly recovers $\mathbf{r}_i = 1$. The range check on $\boldsymbol{e}_3$ thus passes by the bound B' and Lemma 1. □

Theorem 4. *The above private signaling scheme is receiver-private under the* $\mathsf{RLWE}_{n,2,q,\chi}$ *and* $\mathsf{LWE}_{d,m,q,\chi_s,\chi_e}$ *assumptions with* $\chi = D^{\mathsf{coeff}}_{\mathbb{Z}^n,\alpha q}$, $\chi_s = U(\{0,1\})$ *and* $\chi_e = D_{\mathbb{Z},\alpha' q}$. (The proof is included in the full version.)

Theorem 5. *The above private signaling scheme is statistically snake-eye resistant if* $m > (2d + \lambda)/\log(\frac{q}{4B'+1})$.

Proof. The proof is adapted from the statistical argument of [21, Lemma 5.3]. Given two honestly generated public keys $(\mathbf{A}, a, b)$ and $(\mathbf{A}', a', b')$, a successful adversary must come up with a signal $(c_1, \mathbf{c}_2, \mathbf{c}_3) \in R_q \times \mathbb{Z}_q^d \times \mathbb{Z}_q^m$ for which there exist $(\mathbf{r}, \mathbf{e}_3) \in \{0,1\}^d \times [-B', B']^m$ and $(\mathbf{r}', \mathbf{e}_3') \in \{0,1\}^d \times [-B', B']^m$ such that

$$\mathbf{c}_3 = \mathbf{A}^\top \mathbf{r} + \mathbf{e}_3 = \mathbf{A}'^\top \mathbf{r}' + \mathbf{e}_3' \bmod q. \tag{9}$$

Equivalently, (9) can be written

$$\underbrace{\left[\begin{array}{c|c} \mathbf{A}^\top & -\mathbf{A}'^\top \end{array}\right]}_{\triangleq\ \bar{\mathbf{A}}^\top} \cdot \underbrace{\begin{bmatrix} \mathbf{r} \\ \mathbf{r}' \end{bmatrix}}_{\triangleq\ \bar{\mathbf{r}}} = \mathbf{e}_3' - \mathbf{e}_3 \bmod q, \tag{10}$$

where $\bar{\mathbf{A}}^\top$ is uniformly distributed over $\mathbb{Z}_q^{m \times 2d}$ since $\mathbf{A}, \mathbf{A}' \sim U(\mathbb{Z}_q^{d \times m})$ are sampled independently as part of pk and pk' in the snake-eye resistance experiment.

We first note that $\bar{\mathbf{r}} \neq \mathbf{0}^{2d}$ since, otherwise, (9) would imply $\mathbf{c}_3 = \mathbf{e}_3 = \mathbf{e}_3'$ and the Detect algorithm would immediately reject sig. We now claim that, with overwhelming probability over the choice of $(\mathbf{A}, \mathbf{A}')$, there exists no vector $\bar{\mathbf{r}} \in \{0,1\}^{2d} \setminus \{\mathbf{0}^{2d}\}$ such that $\bar{\mathbf{A}}^\top \cdot \bar{\mathbf{r}}$ lives in the hypercube $\mathcal{C} \triangleq \mathbb{Z}^m \cap [-2B', 2B']^m$, meaning that (10) cannot hold for any vectors $\mathbf{e}_3, \mathbf{e}_3' \in [-B', B']^m$.

For a given fixed vector $\bar{\mathbf{r}} \in \{0,1\}^{2d} \setminus \{\mathbf{0}^{2d}\}$, the probability over the random choice of $\bar{\mathbf{A}}^\top \sim U(\mathbb{Z}_q^{m \times 2d})$ that $\bar{\mathbf{A}}^\top \cdot \bar{\mathbf{r}}$ lands in $\mathcal{C}$ is $(4B'+1)^m/q^m$. If we now take a union bound over all possible $\bar{\mathbf{r}} \in \{0,1\}^{2d} \setminus \{\mathbf{0}^{2d}\}$, the probability that one of them leads to $\bar{\mathbf{A}}^\top \cdot \bar{\mathbf{r}} \in \mathcal{C}$ is at most $2^{2d}(4B'+1)^m/q^m$, which is smaller than $2^{-\lambda}$ when $m > (2d+\lambda)/\log(\frac{q}{4B'+1})$. □

One could potentially replace the statistical argument of Theorem 5 by a SIS assumption following the recipe used for Theorem 3. The idea would be to consider $\boldsymbol{c}_3$ and a SIS matrix $\mathbf{A}_{\mathsf{SIS}} = [\mathbf{A}_1^\top \mid \mathbf{A}_2^\top \mid \mathbf{A}_3]$ where $\mathbf{A}_1^\top, \mathbf{A}_2^\top \in \mathbb{Z}_q^{m \times d}$ and $\mathbf{A}_3 \in \mathbb{Z}_q^{m \times m}$ is invertible. The invertibility of $\mathbf{A}_3$ may be guaranteed by taking a subset of the columns of a wider SIS matrix. Then, the reduction would set the public key components $\mathbf{A}^\top = \mathbf{A}_3^{-1} \cdot \mathbf{A}_1^\top$ and $\mathbf{A}^\top = -\mathbf{A}_3^{-1} \cdot \mathbf{A}_2^\top$. With the notation of Equation (10), the SIS solution $(\boldsymbol{r}, \boldsymbol{r}', \boldsymbol{e}_3 - \boldsymbol{e}_3')$ would have Euclidean norm $2 \cdot \sqrt{d} + 2B' \cdot \sqrt{m}$. Alternatively, the solution has an infinity norm bound of B'. Unfortunately, setting the parameters according to Euclidean or infinity norm SIS hardness using the lattice estimator tool [2] does not lead to improved parameters. This is because the hardness of SIS with respect to these norms does not consider the more specific geometry of the solution $(\boldsymbol{r}, \boldsymbol{r}', \boldsymbol{e}_3 - \boldsymbol{e}_3')$ where $\boldsymbol{r}, \boldsymbol{r}'$ are binary and $\|\boldsymbol{e}_3 - \boldsymbol{e}_3'\|_\infty \leq B'$. On the other hand, the statistical argument above does consider this geometry.

Remark 1. While the scheme is specified as a private signaling system, it can be turned into a *strongly robust* PKE scheme by encoding a plaintext $\boldsymbol{\mu} \in \{0,1\}^m$ into $\mathbf{c}_3 = \mathbf{A}^\top \mathbf{r} + \mathbf{e}_3 + \boldsymbol{\Delta}\boldsymbol{\mu}$. The proof of strong robustness bears little difference from that of Theorem 5. In particular, all that changes is the inclusion of an

additional term of $\Delta \cdot (\boldsymbol{\mu}' - \boldsymbol{\mu}) \in \{-\Delta, 0, \Delta\}^m$ in the right-hand-side member of (10). In turn, the proof should consider the union of hypercubes

$$\mathcal{C} \triangleq \mathbb{Z}^m \cap ([-2B', 2B']^m \cup [\Delta - 2B', \Delta]^m \cup [-\Delta, -\Delta + 2B']^m)\,,$$

so that the probability of landing in the hypercube becomes $2^{2d}(8B'+3)^m/q^m$.

While this provides an alternative strongly robust PKE based on LWE/RLWE, the resulting scheme is not additively homomorphic. Indeed, decrypting the sum of two valid ciphertexts may lead to a $\perp$ plaintext.

5.3 Implementation Strategy

In the full version of this paper, we overview concrete parameters and performance of running the Detect algorithm homomorphically. To derive relatively small bounds (B_e, B'), we utilize the previously deployed strategy of setting the RLWE secret distribution to have a ternary, fixed hamming weight coefficient vector [37]. Due to space constraints, we defer the high-level pseudocode for the BFV homomorphic detection of a batch of N signals to the full version. As discussed earlier, a high multiplicative depth greatly hinders performance. The largest contributors to the multiplicative depth are the two range checks with bound B_e and B'. In particular, the multiplicative depth can be estimated as

$$\begin{aligned} &1 + \lceil \log(B_e + 1) \rceil + \lceil \log(q-1) \rceil + \\ \max\{1 + \lceil \log(B'+1) \rceil + \lceil \log(q-1) \rceil, 1 + \lceil \log(d) \rceil\} + \\ &\lceil \log(m+1) \rceil. \end{aligned}$$

The first line corresponds to the range checks with bound B_e using the strategy of Section 4.4. The second line stems from the range checks with bound B' and checking that $\boldsymbol{\mu} \circ \boldsymbol{\mu} = 1^d$. The third line corresponds to combining the m range check results with bound B' with the $\boldsymbol{\mu}$-check from the second line. In terms of multiplications required to process a batch of N signals, we obtain approximately

$$\begin{aligned} &2d \cdot (B_e + 1 + \lceil \log(q-1) \rceil) + d + \\ &m \cdot (B' + 1 + \lceil \log(q-1) \rceil) + \\ &m + d - 1. \end{aligned}$$

The first line here corresponds to the range checks with bound B_e along with the computation of $\boldsymbol{\mu} \circ \boldsymbol{\mu}$, the second line corresponds to the range checks with bound B' and the third line corresponds to combining the range check on e_3 with the check on $\boldsymbol{\mu} \circ \boldsymbol{\mu}$.

Although this number is considerably smaller than that of Section 4 (as the modulo p reduction is removed), performance remains worse than existing LWE-based solutions (e.g., [34]). On the other hand, the public/clue key sizes are roughly $2\times$ smaller at around 2.44kB for concrete parameters. For more details on our concrete implementation, see the full version.

Our schemes can be seen as a first step at building DoS-resistant OMR schemes under structured lattice assumptions, with further improvements left to future work.

References

1. Abdalla, M., Bellare, M., Neven, G.: Robust Encryption. In: TCC (2010). https://doi.org/10.1007/978-3-642-11799-2_28
2. Albrecht, M., Player, R., Scott, S.: On the concrete hardness of Learning with Errors. Journal of Mathematical Cryptology **9**(3) (2015)
3. Albrecht, M., Chase, M., Chen, H., Ding, J., Goldwasser, S., Gorbunov, S., Halevi, S., Hoffstein, J., Laine, K., Lauter, K., Lokam, S., Micciancio, D., Moody, D., Morrison, T., Sahai, A., Vaikuntanathan, V.: Homomorphic encryption security standard. Tech. rep., HomomorphicEncryption.org, Toronto, Canada (November 2018)
4. Badawi, A.A., Alexandru, A., Bates, J., Bergamaschi, F., Cousins, D.B., Erabelli, S., Genise, N., Halevi, S., Hunt, H., Kim, A., Lee, Y., Liu, Z., Micciancio, D., Pascoe, C., Polyakov, Y., Quah, I., R.V., S., Rohloff, K., Saylor, J., Suponitsky, D., Triplett, M., Vaikuntanathan, V., Zucca, V.: OpenFHE: Open-source fully homomorphic encryption library. Cryptology ePrint Archive, Paper 2022/915 (2022), https://eprint.iacr.org/2022/915, https://eprint.iacr.org/2022/915
5. Banaszczyk, W.: New bounds in some transference theorems in the geometry of number. Mathematische Annalen **296**, 625–635 (1993)
6. Barth, A., Boneh, D., Waters, B.: Private encrypted content distribution using private broadcast encryptions. In: Financial Cryptography (2006). https://doi.org/10.1007/11889663_4
7. Baum, C., Damgård, I., Lyubashevsky, V., Oechsner, S., Peikert, C.: More efficient commitments from structured lattice assumptions. In: SCN (2018)
8. Baum, C., Lyubashevsky, V.: Simple amortized proofs of shortness for linear relations over polynomial rings. Cryptology ePrint Archive Report 2017/759, https://eprint.iacr.org/2017/759
9. Beck, G., Len, J., Miers, I., Green, M.: Fuzzy message detection. In: Kim, Y., Kim, J., Vigna, G., Shi, E. (eds.) CCS '21: 2021 ACM SIGSAC Conference on Computer and Communications Security, Virtual Event, Republic of Korea, November 15 - 19, 2021. pp. 1507–1528. ACM (2021). https://doi.org/10.1145/3460120.3484545, https://doi.org/10.1145/3460120.3484545
10. Bellare, M., Boldyreva, A., Desai, A., Pointcheval, D.: Key-privacy in public-key encryption. In: Asiacrypt (2001). https://doi.org/10.1007/3-540-45682-1_33, https://doi.org/10.1007/3-540-45682-1_33
11. Ben-Sasson, E., Chiesa, A., Garman, C., Green, M., Miers, I., Tromer, E., Virza, M.: Zerocash: Decentralized anonymous payments from bitcoin. In: 2014 IEEE Symposium on Security and Privacy, SP 2014, Berkeley, CA, USA, May 18-21, 2014. pp. 459–474. IEEE Computer Society (2014). https://doi.org/10.1109/SP.2014.36, https://doi.org/10.1109/SP.2014.36
12. Boneh, D., Di Crescenzo, G., Ostrovsky, R., Persiano, G.: Public key encryption with keyword search. In: Eurocrypt (2004). https://doi.org/10.1007/978-3-540-24676-3_30
13. Bos, J., Ducas, L., Kiltz, E., Lepoint, T., Lyubashevsky, V., Schanck, J.M., Schwabe, P., Seiler, G., Stehlé, D.: CRYSTALS - kyber: A cca-secure module-lattice-based KEM. In: EuroS&P (2018). https://doi.org/10.1109/EUROSP.2018.00032, https://doi.org/10.1109/EuroSP.2018.00032
14. Brakerski, Z., Langlois, A., Peikert, C., Regev, O., Stehlé, D.: Classical hardness of learning with errors. In: STOC (2013). https://doi.org/10.1145/2488608.2488680

15. Brakerski, Z., Gentry, C., Vaikuntanathan, V.: (leveled) fully homomorphic encryption without bootstrapping. ACM Trans. Comput. Theory **6**(3), 13:1–13:36 (2014). https://doi.org/10.1145/2633600, https://doi.org/10.1145/2633600
16. Cheon, J.H., Lee, K., Park, J.H., Yeo, Y.: SIMD-Aware Homomorphic Compression and Application to Private Database Query. CoRR **abs/2408.17063** (2024). https://doi.org/10.48550/ARXIV.2408.17063, https://doi.org/10.48550/arXiv.2408.17063
17. Chillotti, I., Gama, N., Georgieva, M., Izabachène, M.: Tfhe: fast fully homomorphic encryption over the torus. Journal of Cryptology **33** (2020)
18. Chillotti, I., Joye, M., Paillier, P.: Programmable bootstrapping enables efficient homomorphic inference of deep neural networks. In: International Symposium on Cyber Security Cryptography and Machine Learning. Springer (2021)
19. Corrigan-Gibbs, H., Boneh, D., Mazières, D.: Riposte: An anonymous messaging system handling millions of users. In: 2015 IEEE Symposium on Security and Privacy, SP (2015). https://doi.org/10.1109/SP.2015.27
20. Fan, J., Vercauteren, F.: Somewhat practical fully homomorphic encryption. IACR Cryptol. ePrint Arch. p. 144 (2012), http://eprint.iacr.org/2012/144
21. Gentry, C., Peikert, C., Vaikuntanathan, V.: Trapdoors for hard lattices and new cryptographic constructions. In: STOC (2008). https://doi.org/10.1145/1374376.1374407
22. Grubbs, P., Maram, V., Paterson, K.: Anonymous, robust post-quantum public key encryption. In: Eurocrypt 2022 (2022). https://doi.org/10.1007/978-3-031-07082-2_15, https://doi.org/10.1007/978-3-031-07082-2_15
23. Halevi, S., Shoup, V.: Faster homomorphic linear transformations in HElib. In: Crypto (2018). https://doi.org/10.1007/978-3-319-96884-1_4, https://doi.org/10.1007/978-3-319-96884-1_4
24. Hoffmann, C., Libert, B., Momin, C., Peters, T., Standaert, F.X.: POLKA: Towards Leakage-Resistant Post-Quantum CCA-Secure Public Key Encryption. In: PKC (2023). https://doi.org/10.1007/978-3-031-31368-4_5
25. Iliashenko, I., Nègre, C., Zucca, V.: Integer functions suitable for homomorphic encryption over finite fields. In: WAHC (2021). https://doi.org/10.1145/3474366.3486925, https://doi.org/10.1145/3474366.3486925
26. Jakkamsetti, S., Liu, Z., Madathil, V.: Scalable private signaling. IACR Cryptol. ePrint Arch. p. 572 (2023), https://eprint.iacr.org/2023/572
27. Jia, Y., Madathil, V., Kate, A.: HomeRun: High-efficiency Oblivious Message Retrieval, Unrestricted. In: ACM-CCS (2024). https://doi.org/10.1145/3658644.3670381, https://doi.org/10.1145/3658644.3670381
28. Joye, M.: TFHE public key encryption revisited. In: CT-RSA (2024). https://doi.org/10.1007/978-3-031-58868-6_11
29. Katsumata, S., Yamada, S.: Partitioning via non-linear polynomial functions: More compact IBEs from ideal lattices and bilinear maps. In: Asiacrypt (2016). https://doi.org/10.1007/978-3-662-53890-6_23
30. Lee, K., Yeo, Y.: SophOMR: Improved oblivious message retrieval from SIMD-aware homomorphic compression. Cryptology ePrint Archive, Paper 2024/1814 (2024), https://eprint.iacr.org/2024/1814
31. Liang, H., Liu, Z., Tromer, E., Xie, X., Yu, Y.: InstantOMR: Oblivious Message Retrieval with Low Latency and Optimal Parallelizability. In: USENIX Security (2026)
32. Libert, B., Paterson, K., Quaglia, E.: Anonymous broadcast encryption: Adaptive security and efficient constructions in the standard model. In: PKC (2012). https://doi.org/10.1007/978-3-642-30057-8_13

33. Lindner, R., Peikert, C.: Better key sizes (and attacks) for LWE-based encryption. In: CT-RSA (2011). https://doi.org/10.1007/978-3-642-19074-2_21, https://doi.org/10.1007/978-3-642-19074-2_21
34. Liu, Z., Sotiraki, K., Tromer, E., Wang, Y.: Snake-eye Resistant PKE from LWE for Oblivious Message Retrieval and Robust Encryption. In: Eurocrypt (2025). https://doi.org/10.1007/978-3-031-91131-6_5
35. Liu, Z., Tromer, E.: Oblivious message retrieval. In: Crypto (2022). https://doi.org/10.1007/978-3-031-15802-5_26
36. Liu, Z., Tromer, E., Wang, Y.: Group oblivious message retrieval. In: IEEE Security and Privacy (2024). https://doi.org/10.1109/SP54263.2024.00115
37. Liu, Z., Tromer, E., Wang, Y.: PerfOMR: Oblivious Message Retrieval with Reduced Communication and Computation. In: USENIX (2024)
38. Lyubashevsky, V.: Lattice signatures without trapdoors. In: Pointcheval, D., Johansson, T. (eds.) Eurocrypt (2012). https://doi.org/10.1007/978-3-642-29011-4_43
39. Lyubashevsky, V., Micciancio, D.: Generalized compact knapsacks are collision resistant. In: ICALP. pp. 144–155. Springer (2006)
40. Lyubashevsky, V., Nguyen, N.K., Seiler, G.: Shorter lattice-based zero-knowledge proofs via one-time commitments. In: PKC (2021). https://doi.org/10.1007/978-3-030-75245-3_9
41. Lyubashevsky, V., Peikert, C., Regev, O.: On ideal lattices and learning with errors over rings. In: Eurocrypt (2010). https://doi.org/10.1007/978-3-642-13190-5_1
42. Lyubashevsky, V., Peikert, C., Regev, O.: On ideal lattices and learning with errors over rings. J. ACM **60**(6), 43 (2013)
43. Madathil, V., Scafuro, A., Seres, I.A., Shlomovits, O., Varlakov, D.: Private signaling. In: 31st USENIX Security Symposium, (2022), https://www.usenix.org/conference/usenixsecurity22/presentation/madathil
44. Micciancio, D.: Generalized compact knapsacks, cyclic lattices, and efficient one-way functions from worst-case complexity assumptions. In: FOCS (2002). https://doi.org/10.1109/SFCS.2002.1181960
45. Mohassel, P.: A closer look at anonymity and robustness in encryption schemes. In: Asiacrypt (2010). https://doi.org/10.1007/978-3-642-17373-8_29
46. Noether, S.: Ring signature confidential transactions for monero. IACR Cryptol. ePrint Arch. p. 1098 (2015), http://eprint.iacr.org/2015/1098
47. Paterson, M., Stockmeyer, L.J.: On the number of nonscalar multiplications necessary to evaluate polynomials. SIAM J. Comput. **2**(1), 60–66 (1973). https://doi.org/10.1137/0202007, https://doi.org/10.1137/0202007
48. Peikert, C., Rosen, A.: Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In: TCC (2006). https://doi.org/10.1007/11681878_8
49. Peikert, C., Vaikuntanathan, V., Waters, B.: A framework for efficient and composable oblivious transfer. In: Crypto (2008). https://doi.org/10.1007/978-3-540-85174-5_30
50. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. In: STOC (2005). https://doi.org/10.1145/1060590.1060603
51. Sako, K.: An auction protocol which hides bids of losers. In: PKC (2000). https://doi.org/10.1007/978-3-540-46588-1_28
52. Schnorr, C.P.: A hierarchy of polynomial time lattice basis reduction algorithms. Theor. Comput. Sci. **53**(2-3) (1987). https://doi.org/10.1016/0304-3975(87)90064-8

53. Smart, N.P., Vercauteren, F.: Fully homomorphic SIMD operations. Designs, Codes & Cryptography **71**(1) (2014). https://doi.org/10.1007/S10623-012-9720-4, https://doi.org/10.1007/s10623-012-9720-4
54. Xagawa, K.: Anonymity of NIST PQC Round 3 KEMs. In: Eurocrypt (2022). https://doi.org/10.1007/978-3-031-07082-2_20