On the Inapplicability of Residue-Confined Error Attacks to CRT-FHE with Non-Degenerate Error Distributions

Anil Kumar Pradhan[1] and Abhraneel Dutta[2]

[1]Vaultree Ltd., Ireland
anil@vaultree.com
[2]University of Miami, USA
abdutta@miami.edu

Abstract

This paper responds to the cryptanalysis presented in *"Key Recovery from Residue-Confined Errors in the Pradhan CRT-RLWE Construction"* (Olejnik and Naskrecki, 2026), which claims key recovery and plaintext recovery against the CRT-FHE scheme introduced in *"A New CRT-based Fully Homomorphic Encryption"* (Pradhan et al., 2026).

We demonstrate that the reported recovery attacks rely entirely on a deliberately constructed error distribution whose support lies in the kernel of the CRT encoding function. Consequently, the effective encryption error is identically zero, reducing the attacked construction to a noiseless RLWE instance. The subsequent key-recovery and plaintext-recovery results therefore follow from the deliberate elimination of the encryption noise, rather than from any cryptanalytic method capable of recovering secrets from properly noise-bearing CRT-FHE ciphertexts.

Crucially, the proposed attack neither removes nor estimates errors from normally generated CRT-FHE ciphertexts. Instead, it assumes an artificial error distribution specifically chosen to guarantee their annihilation during encoding. This condition does not hold for the intended centered-binomial error distribution, under which the probability that an entire error polynomial belongs to the encoder's kernel is negligible at the reported parameters. Thus, the presented recovery procedures do not constitute attacks against the concrete CRT-FHE instantiation with its intended non-degenerate error distribution.

We conclude that the reported key-recovery and plaintext-recovery results concern a deliberately degenerate, effectively zero-error instantiation and do not demonstrate a practical security compromise of CRT-FHE under its intended error sampling.

We further show that the distributional mismatch identified in the claimed RLWE hardness reduction does not constitute a cryptanalytic attack. In fact, the CRT-encoded error distribution under the intended parameters admits an exact, invertible transformation to ordinary RLWE with an explicitly characterized non-degenerate ternary error law.

1 Introduction

The CRT-FHE scheme proposed by Pradhan *et al.* [2] employs the Chinese Remainder Theorem (CRT) to encode plaintexts and error terms before incorporating them into RLWE-based ciphertexts. Recently, Olejnik and Naskrecki [3] presented key-recovery and plaintext-recovery attacks against this construction, based on a specially chosen error distribution. In this work, we clarify a fundamental limitation of these attacks: they rely on deliberately eliminating the encryption error before it enters the ciphertext, rather than recovering secrets in the presence of the nonzero errors used in the intended CRT-FHE instantiation.

The central assumption of the proposed attacks is that the sampled errors belong to p_2R, where p_2 is one of the CRT moduli. Specifically, by selecting errors of the form $e = p_2e'$, the

authors ensure that

$$\mathrm{CRT}_{p_1,p_2}(0, e) = 0.$$

Thus, although the initially sampled errors may be nonzero, their contributions are completely annihilated by the CRT encoding *before* key generation and encryption. Consequently, the resulting public-key and ciphertext equations contain no effective encryption error and reduce to a noiseless RLWE instance. The reported recovery procedures therefore exploit a deliberately constructed zero-error setting rather than overcoming the noise that provides security in RLWE-based encryption.

More importantly, this assumption does not hold for the regular small-error distributions used in CRT-FHE. In particular, when the error coefficients satisfy

$$|e_i| < p_2,$$

the condition $e_i \equiv 0 \pmod{p_2}$ holds if and only if $e_i = 0$. Hence, every nonzero coefficient in this range remains nonzero modulo p_2, and the CRT encoder cannot annihilate a nonzero error polynomial. For the centered-binomial distribution B_2 and the reported choice $p_2 = 3$, all sampled coefficients satisfy $|e_i| < p_2$. Therefore, complete error annihilation occurs only when the entire error polynomial is zero, an event of probability

$$\Pr[e = 0] = \left(\frac{3}{8}\right)^n,$$

assuming independently sampled coefficients. This probability is negligible for the reported ring dimensions.

These observations demonstrate that the reported recovery procedures do not apply to the intended CRT-FHE instantiation with its non-degenerate centered-binomial error distribution. The attacks neither remove nor recover the nonzero encoded errors of normally generated ciphertexts; instead, they impose a specially constructed error distribution that makes those errors vanish before encryption. The distinction between this degenerate instantiation and the intended construction is essential when interpreting the claimed security implications.

The remainder of this paper formalizes these observations using the original CRT encoding, key-generation, and encryption equations. We establish the conditions under which error annihilation occurs and demonstrate why the residue-confined recovery procedures do not extend to the regular small-error setting.

Beyond the residue-confined recovery attacks, Section 4 of [3] questions the RLWE hardness reduction by demonstrating that direct error scaling and reduced CRT encoding produce different error distributions. Although this discrepancy identifies an issue in the original same-distribution reduction, it does not establish a security compromise of the resulting CRT-RLWE problem. We show that the actual CRT-RLWE public-key distribution admits an exact, invertible transformation to ordinary RLWE with the induced error law. In particular, under the reported parameters and centered-binomial sampling, this law reduces to a bounded three-point distribution, which remains non-degenerate. Consequently, the statistical-distance arguments of [3] demonstrate a mismatch between two error representations, not an efficient recovery or distinguishing attack against the intended CRT-FHE instantiation.

2 Construction and Mathematical Analysis

We examine the key-recovery and plaintext-recovery arguments of [3] using the CRT encoding and encryption equations of the original CRT-FHE construction [2]. We first show that the proposed attacks deliberately eliminate the encoded error before encryption. We then establish that this error-annihilation mechanism does not apply to regular nonzero errors whose coefficients have magnitude smaller than the auxiliary CRT modulus p_2.

2.1 CRT Encoding and Error Annihilation

Let

$$\mathcal{R} = \mathbb{Z}[X]/(X^n + 1), \qquad \mathcal{R}_q = \mathcal{R}/q\mathcal{R},$$

and let p_1, p_2 be coprime CRT moduli satisfying $p_1 p_2 < q/2$. Define

$$\Delta_1 = p_2[p_2^{-1}]_{p_1}, \qquad \Delta_2 = p_1[p_1^{-1}]_{p_2}.$$

The CRT encoding is defined as

$$\mathrm{CRT}_{p_1,p_2}(m, e) = [\Delta_1 m + \Delta_2 e]_{p_1 p_2}. \tag{1}$$

Its coefficients are represented canonically before embedding the resulting polynomial into $\mathcal{R}_q$. Since $p_1 p_2 < q/2$, the canonical representative embedding is injective.

In particular, the encoded error is determined by

$$\mathrm{CRT}_{p_1,p_2}(0, e) = [\Delta_2 e]_{p_1 p_2}.$$

As Δ_2 is divisible by p_1 and invertible in the second CRT component, its kernel is precisely

$$\mathrm{CRT}_{p_1,p_2}(0, e) = 0 \quad \Longleftrightarrow \quad e \equiv 0 \pmod{p_2\mathcal{R}}. \tag{2}$$

Error distribution assumed by the attack. The cryptanalysis of [3] considers an error distribution supported entirely on multiples of p_2, namely

$$e = p_2 e', \qquad e' \leftarrow B_1.$$

Substituting this choice into the CRT encoding gives

$$\begin{aligned} \mathrm{CRT}_{p_1,p_2}(0, e) &= [\Delta_2 p_2 e']_{p_1 p_2} \\ &= [p_1 p_2 [p_1^{-1}]_{p_2} e']_{p_1 p_2} \\ &= 0. \end{aligned} \tag{3}$$

Thus, although the sampled polynomial e may be nonzero, its encoded contribution is identically zero. This cancellation occurs during CRT encoding, *before* the error is incorporated into the public key or ciphertext modulo q.

Consequently, the attack does not eliminate an error from an existing noisy ciphertext. Instead, it selects an error distribution that guarantees the complete absence of effective error in the resulting encryption equations.

2.2 Inapplicability to Regular Small Errors

The preceding cancellation relies entirely on the special condition $e \in p_2\mathcal{R}$. We now show that this condition cannot hold for a nonzero error polynomial whose coefficients have magnitude strictly smaller than p_2.

Proposition 1. *Let $e = \sum_{i=0}^{n-1} e_i X^i \in \mathcal{R}$ satisfy*

$$\|e\|_\infty < p_2.$$

Then

$$\mathrm{CRT}_{p_1,p_2}(0, e) = 0 \quad \Longleftrightarrow \quad e = 0.$$

In particular, every nonzero error polynomial satisfying $\|e\|_\infty < p_2$ has a nonzero CRT-encoded error.

Proof. By (2), the CRT-encoded error vanishes if and only if

$$e_i \equiv 0 \pmod{p_2}$$

for every coefficient e_i.

Since $|e_i| < p_2$, the only integer in the interval $(-p_2, p_2)$ divisible by p_2 is zero. Therefore,

$$e_i \equiv 0 \pmod{p_2} \quad \Longleftrightarrow \quad e_i = 0.$$

It follows that all coefficients must vanish, and hence $e = 0$. □

Proposition 1 establishes a direct mathematical limitation of the proposed attacks. For regular nonzero errors satisfying $\|e\|_\infty < p_2$, the CRT encoder does not eliminate the error. In particular,

$$e \neq 0, \quad \|e\|_\infty < p_2 \quad \Longrightarrow \quad \mathrm{CRT}_{p_1,p_2}(0, e) \neq 0. \tag{4}$$

This conclusion is deterministic and does not depend on the probability distribution of the error, provided the stated coefficient bound holds.

The centered-binomial distribution. The reported CRT-FHE parameters use $p_2 = 3$, while the centered-binomial distribution B_2 generates coefficients in

$$\{-2, -1, 0, 1, 2\}.$$

Consequently, every sampled error polynomial satisfies

$$\|e\|_\infty \leq 2 < p_2.$$

By Proposition 1, any nonzero error sampled from this distribution necessarily produces a nonzero CRT-encoded error.

For completeness, the coefficient probabilities under B_2 are

$$\Pr[e_i = j] = \begin{cases} 1/16, & j = \pm 2, \\ 4/16, & j = \pm 1, \\ 6/16, & j = 0. \end{cases}$$

Hence, for independently sampled coefficients,

$$\Pr[e \in p_2\mathcal{R}] = \Pr[e = 0] = \left(\frac{3}{8}\right)^n. \tag{5}$$

At the reported ring dimension $n = 8192$, this probability is negligible.

Therefore, unlike the specially constructed distribution $p_2 B_1$, for which complete error annihilation occurs with probability one, the intended B_2 distribution produces nonzero encoded errors except with negligible probability.

2.3 Consequences for the Recovery Attacks

We now examine how the preceding distinction affects the reported key-recovery and plaintext-recovery procedures.

Key recovery. The CRT-FHE public key satisfies

$$b = as + \mathrm{CRT}_{p_1,p_2}(0, e_{\mathrm{pk}}) \pmod{q\mathcal{R}}. \tag{6}$$

Under the specially chosen attack distribution, equation (3) gives

$$b = as \pmod{q\mathcal{R}}.$$

If a is invertible in $\mathcal{R}_q$, the secret key can then be obtained as

$$s = a^{-1}b \pmod{q\mathcal{R}}.$$

This recovery is a direct consequence of the absence of an encoded error in the public key. For a regular nonzero error satisfying $\|e_{\mathrm{pk}}\|_\infty < p_2$, however, Proposition 1 gives

$$\varepsilon_{\mathrm{pk}} := \mathrm{CRT}_{p_1,p_2}(0, e_{\mathrm{pk}}) \neq 0.$$

The public-key equation remains

$$b = as + \varepsilon_{\mathrm{pk}} \pmod{q\mathcal{R}},$$

and the same inversion yields

$$a^{-1}b = s + a^{-1}\varepsilon_{\mathrm{pk}} \not\equiv s \pmod{q\mathcal{R}}. \tag{7}$$

Here the inequality follows from the nonzero encoded error and invertibility of a.

Thus, the reported single-inversion recovery formula does not recover the secret key in the presence of a nonzero encoded error.

Plaintext recovery. Encryption in CRT-FHE is given by

$$\begin{aligned} c_0 &= [bu + \mathrm{CRT}_{p_1,p_2}(m, e_0)]_q, \\ c_1 &= [-au + \mathrm{CRT}_{p_1,p_2}(0, e_1)]_q. \end{aligned} \tag{8}$$

Under the attack distribution, all the relevant CRT error contributions vanish:

$$\mathrm{CRT}_{p_1,p_2}(0, e_{\mathrm{pk}}) = \mathrm{CRT}_{p_1,p_2}(0, e_1) = 0,$$

and

$$\mathrm{CRT}_{p_1,p_2}(m, e_0) = \mathrm{CRT}_{p_1,p_2}(m, 0).$$

Writing $M(m) = \mathrm{CRT}_{p_1,p_2}(m, 0)$, the ciphertext equations reduce to

$$\boxed{\begin{aligned} b &= as, \\ c_0 &= asu + M(m), \pmod{q\mathcal{R}}. \\ c_1 &= -au \end{aligned}} \tag{9}$$

If a is invertible, these equations immediately give

$$c_0 + ba^{-1}c_1 = M(m) \pmod{q\mathcal{R}},$$

without requiring knowledge of the secret key. The no-secret-key recovery mechanism considered in [3] similarly exploits the resulting noiseless structure.

In contrast, for regular nonzero encryption errors, define

$$M_e = \mathrm{CRT}_{p_1,p_2}(m, e_0), \qquad \varepsilon_1 = \mathrm{CRT}_{p_1,p_2}(0, e_1).$$

Then, whenever a is invertible, equation (8) gives

$$c_0 + ba^{-1}c_1 = M_e + ba^{-1}\varepsilon_1 \pmod{q\mathcal{R}}. \tag{10}$$

Unlike (9), this expression contains the encoded encryption errors. In particular, the simplification to $M(m)$ is no longer justified by the residue-confinement argument.

The proposed recovery analysis therefore does not establish plaintext recovery from regular CRT-FHE ciphertexts containing nonzero encoded errors.

2.4 Summary of the Mathematical Limitation

The attack and the intended error sampling lead to fundamentally different encoded error distributions:

$$\begin{aligned} &\text{Attack distribution:} && e \leftarrow p_2 B_1 && \Longrightarrow \mathrm{CRT}_{p_1,p_2}(0,e) = 0, \\ &\text{Regular small error:} && 0 < \|e\|_\infty < p_2 && \Longrightarrow \mathrm{CRT}_{p_1,p_2}(0,e) \neq 0. \end{aligned}$$

The reported recovery procedures rely on the first case, where the effective error is deliberately annihilated before encryption. They do not provide a corresponding method for eliminating the nonzero encoded errors present in the second case.

Consequently, the residue-confined recovery arguments do not demonstrate key recovery or plaintext recovery against the concrete CRT-FHE instantiation using its intended non-degenerate centered-binomial error distribution.

2.5 Inapplicability to Errors After Homomorphic Evaluation

The residue-confined error argument also does not extend to errors that become multiples of p_2 during homomorphic evaluation. The attack relies on applying the CRT encoding to an error $e \in p_2 R$ *before* encryption, such that

$$\mathrm{CRT}_{p_1,p_2}(0,e) = 0.$$

However, once encryption has been performed, ciphertexts and their subsequent homomorphic evaluations operate in R_q, not in $R_{p_1 p_2}$.

Suppose an evaluated ciphertext has phase

$$c_0 + c_1 s \equiv M + p_2 e' \pmod{qR},$$

where M denotes the encoded message contribution. Even though the resulting error is divisible by p_2 as an integer polynomial, this does not imply

$$c_0 + c_1 s \equiv M \pmod{qR}.$$

Indeed, since $\gcd(q, p_2) = 1$, multiplication by p_2 is invertible in R_q. Consequently,

$$p_2 e' \equiv 0 \pmod{qR} \quad \Longleftrightarrow \quad e' \equiv 0 \pmod{qR}.$$

Thus, divisibility by p_2 does not eliminate the error in the ciphertext modulus.

Moreover, coefficient-wise reduction modulo p_2 is not a well-defined ring homomorphism from R_q to R_{p_2} when $p_2 \nmid q$. Therefore, the CRT-annihilation property cannot be applied retroactively to ciphertexts following homomorphic operations.

Hence, even if homomorphic evaluation produces an error that is a multiple of p_2, the error remains present modulo q unless it is itself zero modulo q. The proposed residue-confined recovery mechanism therefore does not extend to errors arising after encryption or during homomorphic evaluation.

3 A Distributional Mismatch Is Not a Security Break: Exact RLWE Representation of CRT-RLWE

Section 4 of [3] argues that the RLWE-to-CRT-RLWE transformation in [2] does not preserve the original error distribution. While the reduced CRT encoding and direct error scaling are indeed distinct operations, this observation does not establish a cryptanalytic attack or imply that the resulting CRT-RLWE distribution is insecure.

We establish an exact transformation from CRT-RLWE with its actual encoded error distribution to an ordinary RLWE instance with a well-defined, nonzero error distribution. In particular, for the reported parameters and centered-binomial sampling, the transformed error follows a bounded three-point law. This provides a direct characterization of the underlying RLWE problem without requiring the distribution-preservation identity disputed in [3].

3.1 Distinguishing Error Scaling from CRT Encoding

Let $p = p_1 p_2$ and

$$\Delta_2 = p_1[p_1^{-1}]_{p_2}.$$

The two error transformations considered in [3] are

$$T(e) = [\Delta_2 e]_q$$

and

$$C(e) = \mathrm{CRT}_{p_1,p_2}(0, e),$$

where the latter is interpreted using a fixed coefficientwise representative before embedding into $\mathcal{R}_q$.

These transformations need not coincide:

$$T(e) \neq C(e) \pmod{q\mathcal{R}}$$

for some $e \in \mathcal{R}$. Accordingly, their induced error distributions may differ.

However, the inequality of these distributions does not imply that the CRT-RLWE distribution is computationally distinguishable from uniform. In particular, statistical distance between the isolated error distributions is not the distinguishing advantage against RLWE samples, where the error is hidden by the secret-dependent term as.

The relevant question is therefore whether the actual encoded error distribution admits an appropriate RLWE hardness characterization, rather than whether it is identical to the original sampled distribution.

3.2 Exact Reduction to Ordinary RLWE

We now establish an exact equivalence for the reduced CRT encoding itself, without replacing it by an unreduced integer lift.

Proposition 2. *Let* $\gcd(p_1, p_2) = 1$ *and* $\gcd(p_1, q) = 1$. *Assume the reduced CRT encoder uses either the least-nonnegative or the centered coefficientwise representative. Let*

$$b = as + \mathrm{CRT}_{p_1,p_2}(0, e) \pmod{q\mathcal{R}},$$

where a is uniformly sampled from $\mathcal{R}_q$ and e follows an arbitrary integer-polynomial distribution χ.

Then the resulting CRT-RLWE distribution is exactly and efficiently equivalent to an ordinary RLWE distribution with the same secret and an explicitly induced error law.

Proof. By the defining CRT congruences,

$$\mathrm{CRT}_{p_1,p_2}(0, e) \equiv 0 \pmod{p_1\mathcal{R}}.$$

Under either specified representative, each coefficient of the encoded error is therefore an integer multiple of p_1.

Consequently, there exists an integer polynomial $z \in \mathcal{R}$ such that

$$\mathrm{CRT}_{p_1,p_2}(0, e) = p_1 z.$$

More precisely, writing

$$k = [p_1^{-1}]_{p_2},$$

the coefficients of z are the chosen representatives of

$$z_i \equiv k e_i \pmod{p_2}.$$

The public-key equation becomes

$$b = as + p_1 z \pmod{q\mathcal{R}}.$$

Since $\gcd(p_1, q) = 1$, multiplication by p_1 is invertible in $\mathcal{R}_q$. Apply the public transformation

$$a' = p_1^{-1} a, \qquad b' = p_1^{-1} b \pmod{q\mathcal{R}}.$$

We obtain

$$\boxed{b' = a's + z \pmod{q\mathcal{R}}.} \tag{11}$$

Furthermore, a' remains uniformly distributed because multiplication by p_1^{-1} is a bijection on $\mathcal{R}_q$.

Thus (a', b') is an ordinary RLWE sample with error law

$$\chi_{\text{eff}} = (e \mapsto z)_{\#}\chi.$$

The transformation is invertible:

$$(a', b') \longmapsto (p_1 a', p_1 b').$$

Both directions preserve the uniform comparison distribution and can be applied samplewise, including to multiple samples sharing the same secret.

Hence the two decisional problems are exactly equivalent. □

Proposition 2 shows that the reduced CRT encoding does not place the public-key distribution outside the ordinary RLWE framework. Rather, it induces a different, explicitly defined RLWE error distribution.

The equivalence is exact and does not require the incorrect identification $\text{CRT}(0, e) = \Delta_2 e$ modulo q.

3.3 Exact Error Distribution at the Reported Parameters

We specialize the preceding result to the reported parameters

$$p_1 = 65537, \qquad p_2 = 3,$$

for which

$$k = [p_1^{-1}]_3 = 2.$$

Let each coefficient of e be independently sampled from the centered-binomial distribution B_2:

$$\Pr[e_i = j] = \begin{cases} 1/16, & j = \pm 2, \\ 4/16, & j = \pm 1, \\ 6/16, & j = 0. \end{cases}$$

Centered CRT representative. For the centered representative, the normalized error coefficient is

$$z_i = \langle 2e_i \rangle_3,$$

where $\langle \cdot \rangle_3$ denotes the centered representative modulo 3.

Direct enumeration gives

$$\boxed{\Pr[z_i = j] = \begin{cases} 5/16, & j = -1, \\ 6/16, & j = 0, \\ 5/16, & j = 1. \end{cases}} \tag{12}$$

Therefore, the centered CRT-RLWE public-key distribution is exactly equivalent to ordinary RLWE with a symmetric, bounded ternary error law.

In particular,

$$\Pr[z_i \neq 0] = \frac{5}{8}, \qquad \mathbb{E}[z_i] = 0, \qquad \mathbb{E}[z_i^2] = \frac{5}{8}.$$

Least-nonnegative CRT representative. For the least-nonnegative representative, the normalized coefficient is

$$z_i = [2e_i]_3 \in \{0, 1, 2\},$$

with distribution

$$\Pr[z_i = j] = \begin{cases} 6/16, & j = 0, \\ 5/16, & j = 1, \\ 5/16, & j = 2. \end{cases} \tag{13}$$

This is likewise an ordinary RLWE error distribution supported on three bounded integers. If desired, the public shift $z_i \mapsto z_i - 1$ gives support $\{-1, 0, 1\}$, with probabilities $6/16, 5/16, 5/16$, respectively.

Hence both representative conventions admit exact ordinary-RLWE formulations with non-degenerate bounded errors.

3.4 Interpretation of the Statistical-Distance Claims

Example 4.5 and Corollary 4.6 of [3] establish a large statistical distance between the errors obtained by directly scaling B_2 with Δ_2 and those obtained through the reduced CRT encoder.

These results demonstrate that the two error transformations are not distribution-preserving.

However, the distance

$$\mathrm{TV}(T_{\#}\chi, C_{\#}\chi)$$

compares two error distributions, not the computational distinguishability of either corresponding RLWE distribution from uniform.

In particular, even a statistical distance approaching one between the isolated error laws does not establish an algorithm for recovering the secret or distinguishing the actual CRT-RLWE samples.

By Proposition 2, the actual distribution can instead be written exactly as

$$(a', a's + z), \qquad z \leftarrow \chi_{\mathrm{eff}}.$$

The appropriate security question is therefore the hardness of ordinary RLWE with χ_{eff}, not the statistical distance between χ_{eff} and an unrelated scaled-error distribution.

For the intended B_2 parameters, χ_{eff} is the explicit three-point distribution derived above. Its non-degeneracy does not itself constitute a hardness proof, but neither does the distributional discrepancy reported in [3] constitute evidence of an efficient attack against this RLWE instance.

3.5 Scope of the Security Implications

Theorem 4.8 of [3] correctly observes that invertible scaling of ordinary RLWE samples preserves the corresponding decisional problem. However, this fact does not turn the authors' residue-confined instantiation into an attack against the intended centered-binomial instantiation.

For the specially chosen law $\chi = p_2 B_1$, the reduced CRT encoding produces

$$\chi_{\text{eff}} = \delta_0,$$

which is a noiseless RLWE instance.

For the intended B_2 distribution, the induced law is instead the non-degenerate three-point distribution in (12) or (13), depending on the representative convention.

These are fundamentally different security assumptions.

Our result also clarifies the scope of the original security reduction. The identification of the reduced CRT error with $\Delta_2 e$ does not establish hardness under the unchanged B_2 error law. The appropriate formulation is instead ordinary RLWE hardness for the explicitly induced law χ_{eff}.

This gives an exact reduction for the CRT-RLWE public-key distribution. A complete IND-CPA security proof for the encryption scheme must additionally account for the encryption algorithm and its ciphertext-error distributions.

Consequently, Section 4 of [3] identifies a distributional mismatch in the original same-law reduction, but does not establish that the actual CRT-RLWE distribution is computationally easy. The reported statistical-distance bounds do not constitute key-recovery, plaintext-recovery, or distinguishing attacks against the intended CRT-FHE instantiation.

3.6 Hardness Preservation Under Invertible Scalar Transformations

The use of scalar transformations in Ring-LWE security reductions is well established. In particular, Lyubashevsky, Peikert, and Regev [1, Lemma 2.23] establish an efficient transformation that maps Ring-LWE samples with error distribution ψ to samples with a suitably scaled and discretized error distribution χ. Their transformation preserves uniform samples and therefore transfers decisional Ring-LWE hardness between the corresponding problems, subject to the stated discretization and coprimality conditions.

In the present setting, a simpler exact equivalence follows directly from invertible scalar multiplication in $\mathcal{R}_q$. Let c be an integer satisfying $\gcd(c, q) = 1$, and consider the transformation

$$\Phi_c(a, b) = ([ca]_q, [cb]_q).$$

For an RLWE sample

$$b = as + e \pmod{q\mathcal{R}},$$

we obtain

$$cb = (ca)s + ce \pmod{q\mathcal{R}}.$$

Since c is invertible modulo q, Φ_c is a bijection on $\mathcal{R}_q^2$ and preserves the uniform distribution. Its inverse is $\Phi_{c^{-1}}$.

Consequently, for any error distribution χ over $\mathcal{R}_q$,

$$\text{DLWE}_{q,\chi} \equiv \text{DLWE}_{q,c_\#\chi}, \tag{14}$$

where $c_\#\chi$ denotes the distribution of $ce \bmod q$ for $e \leftarrow \chi$. The equivalence is exact and preserves the distinguishing advantage and number of samples.

Applying this observation to the actual CRT-RLWE public-key distribution, write

$$b = as + \text{CRT}_{p_1,p_2}(0, e) = as + p_1 z \pmod{q\mathcal{R}},$$

where z follows the induced distribution χ_{eff}.

Because $\gcd(p_1, q) = 1$, multiplication of both public-key components by p_1^{-1} gives

$$\Big([p_1^{-1}a]_q, [p_1^{-1}b]_q\Big) = (a', a's + z).$$

Therefore,

$$\boxed{\text{CRT-RLWE}_{q,\chi} \equiv \text{RLWE}_{q,\chi_{\text{eff}}}.} \tag{15}$$

This equivalence is independent of the incorrect identification $\text{CRT}_{p_1,p_2}(0, e) = \Delta_2 e \bmod q$. It uses the actual reduced CRT encoding and its induced error distribution.

For the reported parameters $p_1 = 65537$, $p_2 = 3$, and centered CRT representatives, sampling $e \leftarrow B_2$ yields

$$\Pr[z_i = -1] = \frac{5}{16}, \qquad \Pr[z_i = 0] = \frac{6}{16}, \qquad \Pr[z_i = 1] = \frac{5}{16}.$$

Thus, the actual CRT-RLWE distribution is exactly equivalent to ordinary RLWE with a non-degenerate ternary error distribution.

Importantly, Lemma 2.23 of [1] does not assert that arbitrary error distributions remain statistically unchanged under scaling. Rather, it establishes a hardness-preserving transformation to a specified scaled and discretized distribution.

Likewise, equation (15) establishes an exact equivalence with the induced error law, not with the original B_2 law. The hardness of the resulting ternary-error RLWE problem must be justified under its own assumptions.

Nevertheless, the distributional mismatch identified in [3] does not by itself establish an efficient attack on the actual CRT-RLWE distribution. The correct security problem is explicitly characterized by (15).

4 Discussion

The preceding analysis establishes a fundamental limitation of the recovery attacks presented in [3]. Their key-recovery and plaintext-recovery procedures rely on an artificially constructed error distribution whose support lies entirely in the kernel of the CRT encoding function. By choosing $e = p_2 e'$, the attacks guarantee that $\text{CRT}_{p_1,p_2}(0, e) = 0$ before the error is incorporated into the public key or ciphertext. Consequently, the resulting construction contains no effective encryption error, and the reported recovery equations follow directly from this noiseless setting.

This distinction is essential. The attacks do not identify a method for removing, estimating, or overcoming the nonzero errors present in normally generated CRT-FHE ciphertexts. Instead, they replace the intended error distribution with one specifically selected to ensure complete error annihilation during encoding. Thus, the resulting recovery procedures exploit the absence of effective encryption noise rather than a weakness that has been demonstrated for the intended noisy CRT-FHE instantiation.

Furthermore, as established in Proposition 1, the error-annihilation condition cannot hold for any nonzero error polynomial satisfying $\|e\|_\infty < p_2$. In particular, under the reported parameters $p_2 = 3$ and the centered-binomial distribution B_2, the encoded error is nonzero whenever the sampled error polynomial is nonzero. The probability of complete error annihilation is only $(3/8)^n$, which is negligible at the reported ring dimensions. Hence, the central condition required by the proposed recovery procedures is not satisfied by the intended error sampling, except with negligible probability.

It is also important to distinguish an attack against a deliberately selected, effectively zero-error instantiation from a cryptanalytic attack against a specified encryption scheme operating with its intended error distribution. The former does not establish the latter. Although the constructed distribution illustrates the consequences of choosing errors from the kernel of the CRT

encoder, it provides no evidence that normally generated CRT-FHE ciphertexts are susceptible to the same recovery procedures.

Accordingly, the reported recovery results should not be interpreted as demonstrating key recovery or plaintext recovery against CRT-FHE under its intended centered-binomial error distribution. Any such claim would require a recovery method that operates in the presence of the nonzero encoded errors generated by that distribution, which is not established by the residue-confined error analysis.

5 Conclusion

We have examined the residue-confined error attacks of [3] against the CRT-FHE construction of [2]. Our analysis demonstrates that the reported recovery procedures depend entirely on a specially chosen error distribution that is annihilated by CRT encoding before encryption. The attacked instance therefore contains zero effective error, making the resulting key-recovery and plaintext-recovery equations a consequence of this deliberately noiseless construction.

In contrast, regular nonzero errors satisfying $\|e\|_\infty < p_2$ cannot be annihilated by the CRT encoder. Consequently, the proposed recovery mechanism does not apply to the intended centered-binomial instantiation of CRT-FHE, except on a negligible complete-error-annihilation event.

We conclude that the reported recovery procedures do not demonstrate a security compromise of CRT-FHE under its intended error distribution. Rather than recovering secrets from noisy RLWE ciphertexts, they exploit an artificially constructed instance in which the effective encryption noise has already been eliminated.

References

[1] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. A Toolkit for Ring-LWE Cryptography. In Thomas Johansson and Phong Q. Nguyen, editors, *Advances in Cryptology – EUROCRYPT 2013*, volume 7881 of *Lecture Notes in Computer Science*, pages 35–54. Springer, 2013.

[2] A. K. Pradhan, A. Dutta, H. Jangir, and D. Das, *A New CRT-based Fully Homomorphic Encryption*, IACR Communications in Cryptology (CiC), 2026. `https://doi.org`.

[3] L. Olejnik and B. Naskrecki, *Key Recovery from Residue-Confined Errors in the Pradhan CRT-RLWE Construction*, IACR Cryptology ePrint Archive, Paper 2026/1734 (2026). `https://eprint.iacr.org/2026/1734`.