# Single-Hop HRA-Secure Owner-Encrypted Lattice-Based Proxy Re-Encryption without Statistical Noise Flooding

Haotian Yin[1][0000−0002−5088−6116], Jie Zhang[1][0000−0001−7415−7869], Yuji Dong[2][0000−0001−6715−7588], Dominik Wojtczak[3][0000−0001−5560−0546], and Eng Gee Lim[1][0000−0003−0199−7386]

[1] School of Advanced Technology,
Xi'an Jiaotong-Liverpool University, Suzhou, 215123, China
haotian.yin23@student.xjtlu.edu.cn, jie.zhang01@xjtlu.edu.cn, enggee.lim@xjtlu.edu.cn
[2] School of Internet of Things,
Xi'an Jiaotong-Liverpool University, Suzhou, 215400, China
yuji.dong02@xjtlu.edu.cn
[3] Department of Computer Science,
University of Liverpool, Liverpool, L69 3BX, UK
d.wojtczak@liverpool.ac.uk

**Abstract.** Existing lattice-based proxy re-encryption (PRE) schemes that achieve security against honest re-encryption attacks (HRA) typically rely on statistical noise flooding to hide dependencies between the source ciphertext, the re-encryption key, and the re-encrypted ciphertext. While effective, this technique causes large noise growth and significant parameter expansion.

We present an owner-encrypted proxy re-encryption (OE-PRE) construction that achieves single-hop HRA security without statistical noise flooding. In the owner-encrypted setting, ciphertexts are generated directly using the data owner's secret key, yielding a clean inherited-noise term during re-encryption.

To formalise the leakage argument, we use noisy error-leakage LWE (NEL-LWE), which captures LWE hardness in the presence of a noisy observation of the error and follows from a special case of the Leaky-LWE framework of Lai, Swarnakar, and Woo (CiC 2025). We prove single-hop HRA security from standard LWE/RLWE-style pseudorandomness assumptions, the NEL-LWE consequence of Leaky-LWE, and PRF security via a direct HRA hybrid proof. Our parameter discussion shows that the HRA-specific computational padding remains at the same scale as the encryption error, in contrast to the large flooding widths required in prior HRA-secure lattice PRE schemes. The proposed scheme targets single-hop owner-generated access delegation, such as encrypted file sharing and publish/subscribe delivery.

**Keywords:** Proxy re-encryption · Honest re-encryption security · Lattice cryptography · Error leakage

# 1 Introduction

Proxy re-encryption (PRE), introduced by Blaze, Bleumer, and Strauss [BBS98], enables a semi-trusted proxy to transform a ciphertext encrypted for one user into a ciphertext decryptable by another user, without revealing the underlying plaintext.

Lattice-based PRE has attracted significant attention due to its post-quantum security and its natural compatibility with key-switching techniques. In such schemes, re-encryption typically transforms a ciphertext by approximately evaluating the source-key decryption relation and re-encoding the result under the target key. As a consequence, the re-encrypted ciphertext may retain residual noise inherited from the source ciphertext, together with additional key-switching noise. This residual noise can become visible through honest re-encryption transcripts and may leak information correlated with the source secret key, even when the underlying encryption scheme is IND-CPA-secure.

Cohen [Coh19] formalised this phenomenon by showing that IND-CPA security alone is insufficient for PRE. An adversary may obtain honestly re-encrypted ciphertexts and exploit information leaked through the re-encryption procedure. Existing lattice-based HRA-secure PRE constructions address this issue using statistical *noise flooding*: the proxy adds a large fresh Gaussian error during re-encryption, so that the resulting ciphertext becomes statistically independent of the original re-encryption randomness. This statistical hiding is then used to establish re-encryption simulatability and, consequently, HRA security.

While effective, this approach comes at a high cost. The added masking noise must statistically dominate the pre-existing error terms, often requiring a variance substantially larger than that of the underlying encryption noise. As a result, re-encryption consumes a significant portion of the available noise budget and forces larger correctness parameters. Although FHE-style noise-management tools such as modulus switching or bootstrapping can be used to support further re-encryptions, they introduce additional parameters and computational costs. In the absence of such tools, statistical flooding directly reduces the number of feasible re-encryption hops. This dependence on statistical domination raises a natural question:

> Can HRA security be achieved through computational masking of the inherited ciphertext noise?

## 1.1 Our Results

In this work, we answer this question affirmatively in the owner-encrypted single-hop setting. We show that, when ciphertexts are generated directly by the data owner, the inherited encryption noise has a clean enough form to be protected by computational leakage masking rather than by statistical domination.

The proposed scheme targets the single-hop owner-generated access-delegation layer that appears in encrypted file-sharing/storage systems and publish/subscribe

delivery. In these settings, a data owner or publisher generates an encrypted object and a semi-trusted storage server or broker later transforms it once for an authorised recipient [AFGH06, PRSV17]. We give a detailed discussion of the application scope and limitations in Section 5.

Our contributions are as follows.

- We propose a BGV-style owner-encrypted proxy re-encryption (OE-PRE) construction. Re-encryption remains unidirectional and uses Alice's secret key together with Bob's public key to generate the re-encryption key. The owner-encrypted form yields a simple inherited-noise term, which is the basis of our leakage-masking argument.
- We use noisy error-leakage LWE (NEL-LWE) to capture the hardness of distinguishing LWE samples in the presence of noisy error leakage. The NEL-LWE condition needed in our proof follows from the recent Leaky-LWE framework of Lai, Swarnakar, and Woo [LSW25] in the error-only leakage setting.
- We prove single-hop HRA security directly inside the HRA game. The proof preserves the algebraic consistency of visible honest-to-honest re-encryption transcripts, while showing that honest-to-corrupted re-encryption responses reveal only computationally harmless noisy error leakage.

Our model is a stateful-proxy OE-PRE model: the proxy keeps a private PRF key $\kappa_{\mathsf{pad}}$ used only to derive deterministic per-ciphertext padding. The HRA adversary may obtain admissible re-encryption keys and re-encryption transcripts, but it is not given the proxy's private padding state.

Overall, our work shows that the single-hop owner-generated access-delegation layer of lattice-based PRE can be protected by computational leakage masking while remaining reducible to standard LWE-style assumptions.

### 1.2 Technical Overview

The key idea is to replace statistical domination of the inherited ciphertext noise by computational leakage masking. In flooding-based HRA proofs, the re-encryption noise is chosen large enough to statistically erase the dependence on the incoming ciphertext and re-encryption randomness. In our construction, the proxy instead exposes only a padded version of the owner-encryption error and proves that this leakage is computationally harmless. The tool for this step is NEL-LWE, which follows from the Leaky-LWE framework under the parameter conditions stated in Corollary 1.

The owner-encrypted setting is crucial. A standard public-key LWE ciphertext contains inherited noise of the form $e_{\mathsf{pk}}v + se_1 + e_0$, which is a correlated convolution of several random sources and does not directly match the error-leakage form needed in our proof. By contrast, an owner-generated ciphertext under Alice's secret key $s_A$ satisfies

$$c_0 + c_1 s_A = \mathsf{m} + pe_{\mathsf{ct}} \pmod{q},$$

where $e_{\mathsf{ct}}$ is a fresh encryption error. This clean inherited error is the quantity protected by our leakage argument.

During re-encryption, the proxy performs BGV-style key switching from $s_A$ to Bob's key and adds a deterministic padding value $e_{\mathsf{pad}}(\mathsf{ct})$ derived from a proxy-held PRF key and the source ciphertext encoding. The padding is fixed for each source ciphertext, a method used in recent threshold encryption [MS25, OT26]. This prevents an adversary from obtaining many independent masked versions of the same error and averaging them to recover information about $e_{\mathsf{ct}}$. For honest-to-corrupted re-encryption queries, the exposed legacy term has the form

$$e_{\mathsf{ct}} + e_{\mathsf{pad}}(\mathsf{ct}),$$

which is exactly the noisy error-leakage pattern that NEL-LWE handles.

Fixed padding must also be protected in visible, honest-to-honest transcripts. If it were added in the clear, a visible re-encryption key and its matching output would allow public subtraction of $pe_{\mathsf{pad}}(\mathsf{ct})$. We therefore add a fresh target-key zero-encryption shield $C_{\mathsf{fr}}$ to every re-encryption. After subtracting the visible key-switching contribution, the adversary sees a BGV encryption-form sample with offset $pe_{\mathsf{pad}}$ under an honest target public key, rather than the padding itself. Standard BGV pseudorandomness hides this fixed offset.

Our HRA proof is direct. Instead of first proving a stand-alone re-encryption simulatability notion, we work inside the HRA game and preserve the public consistency between visible honest-to-honest re-encryption keys and their outputs. The theorem is restricted to one re-encryption hop. After an honest-to-honest hop, the inherited effective noise also contains key-switching and zero-encryption-shield errors depending on the honest intermediate secret. If such a ciphertext is later re-encrypted to a corrupted target, this effective noise may itself become part of the leakage. Handling this would require an additional leakage theorem for BGV effective errors or a different refresh mechanism; we return to this limitation in Section 4.

### 1.3 Related Work

Early lattice-based PRE constructions focused mainly on efficient delegation and IND-CPA security. Polyakov, Rohloff, Sahu, and Vaikuntanathan [PRSV17] used key-switching techniques to build efficient lattice-based PRE, while Cohen [Coh19] later showed that IND-CPA security alone is insufficient for PRE because honest re-encryption transcripts may leak information about the re-encryption process. This motivated the study of security against honest re-encryption attacks.

Existing HRA-secure lattice PRE constructions typically obtain re-encryption simulatability through statistical smudging or flooding. Cohen, Cousins, Genise, Kline, Polyakov, and Rv [CCG+25] give a tightly analysed RLWE-based HRA-secure PRE construction using calibrated flooding. Related large-noise masking mechanisms also appear in broader PRE settings, including the source-hiding approach of Fuchsbauer, Kamath, Klein, and Pietrzak [FKKP19], the strong blurring construction of Chandran, Chase, Liu, Nishimaki, and Xagawa [CCL+14],

the unbounded multi-hop construction of Wan, Wang, Xue, and Wang [WWXW25], and the fine-grained multi-hop schemes of Zhou, Liu, and Han [ZLH24]. These works differ in functionality, but their HRA mechanisms rely on statistical hiding rather than on computational error-leakage masking.

Our work is also related to leakage-resilient variants of LWE. Lai, Swarnakar, and Woo [LSW25] introduced the Leaky-LWE framework, which studies LWE hardness in the presence of auxiliary leakage on the secret and error terms. The NEL-LWE condition used in this paper is the error-only leakage case, instantiated with leakage $\boldsymbol{\ell} = \mathbf{e} + \mathbf{e}_{\mathsf{pad}}$.

### 1.4 Organisation

The remainder of this paper is organised as follows. Section 2 reviews the necessary preliminaries, notation, and security notions. Section 3 presents the proposed OE-PRE construction, introduces the NEL-LWE-based security framework, and establishes correctness and single-hop HRA security by a direct HRA hybrid. Section 3.4 compares the HRA-specific noise requirements with flooding-based approaches. Section 4 gives a bounded multi-hop correctness analysis and discusses the additional obstacle to multi-hop HRA security. Section 5 discusses the application scope and future parameter-design directions. Section 6 concludes the paper.

## 2 Preliminaries

Let $\lambda \in \mathbb{N}$ denote the security parameter. For an integer $N \in \mathbb{N}$, define $[N] := \{1, 2, \ldots, N\}$. A probabilistic polynomial-time (PPT) algorithm is an algorithm whose running time is polynomial in the security parameter $\lambda$. We write $\mathsf{poly}(\lambda)$ for an unspecified polynomial and $\mathsf{negl}(\lambda)$ for an unspecified negligible function. An event occurs with overwhelming probability if it occurs with probability at least $1 - \mathsf{negl}(\lambda)$. We use $\perp$ to denote an abort symbol.

We use bold lowercase letters, e.g., $\mathbf{x}$, for vectors and bold uppercase letters, e.g., $\mathbf{A}$, for matrices.

For a distribution $D$, we write $x \xleftarrow{\$} D$ to denote that $x$ is sampled according to $D$. For a finite set $S$, $U(S)$ denotes the uniform distribution over $S$, and $x \xleftarrow{\$} U(S)$ denotes uniform sampling from $S$. For two distribution ensembles $X = \{X_\lambda\}_\lambda$ and $Y = \{Y_\lambda\}_\lambda$, we write $X \approx_s Y$ for statistical indistinguishability and $X \approx_c Y$ for computational indistinguishability.

For a structured object $X$, $\langle X \rangle$ denotes its canonical encoding used as an index in simulator tables and PRF evaluations. We use $m$ for the number of LWE samples and $\mathsf{m}$ for plaintext messages.

Unless stated otherwise, all probabilities are taken over the randomness of all algorithms and experiments involved.

### 2.1 Rings, Embeddings, and Discrete Gaussians

Let $R = \mathbb{Z}[X]/(X^d + 1)$, where $d = d(\lambda)$ is a power of two, and let $R_q := R/qR \cong \mathbb{Z}_q[X]/(X^d + 1)$. Every element $a \in R$ is represented as $a = z_0 + z_1 X + \cdots + z_{d-1} X^{d-1}$, where $z_i \in \mathbb{Z}$. Throughout the paper, we take $q$ to be a prime modulus satisfying $q \equiv 1 \pmod{2d}$, so that $X^d + 1$ splits completely over $\mathbb{F}_q$. Consequently, $R_q \cong \prod_{k=0}^{d-1} \mathbb{F}_q$ via the Number Theoretic Transform (NTT).

We use the coefficient embedding $\tau : R \to \mathbb{Z}^d, \tau\left(\sum_{i=0}^{d-1} z_i X^i\right) := (z_0, \ldots, z_{d-1})$. We extend $\tau$ component-wise to $R^m$, thereby identifying $R^m$ with $\mathbb{Z}^{dm}$. For $x \in R_q$, let $[x]_q \in R$ denote its coefficient-wise centred representative, with coefficients in $(-q/2, q/2]$. Arithmetic modulo $q$ is performed in $R_q$, whereas all norm bounds on elements of $R_q$ are taken after applying $[\cdot]_q$.

For $a \in R$, we set $\|a\|_2 := \|\tau(a)\|_2, \|a\|_\infty := \|\tau(a)\|_\infty$, and for $x \in R_q$, we set $\|x\|_2 := \|[x]_q\|_2, \|x\|_\infty := \|[x]_q\|_\infty$. The notation is extended component-wise to vectors, and $\|a\|$ denotes $\|a\|_2$. All Gaussian parameters are measured with respect to the coefficient embedding.

The following expansion factor bound is due to [BGV14, Proposition 2.1].

**Proposition 1 (Expansion factor of $R$).** *Let $\gamma_R := \sup_{a,b \in R \setminus \{0\}} \frac{\|ab\|}{\|a\|\|b\|}$. Then $\gamma_R \le \sqrt{d}$.*

We identify each plaintext $\mathsf{m} \in R_p$ with its coefficient-wise centred representative $[\mathsf{m}]_p \in R$, and then view this representative modulo $q$ as an element of $R_q$. Throughout the paper, plaintexts satisfy $\|[\mathsf{m}]_p\|_\infty \le B_{\mathsf{m}} < p/2$.

Let $B_g \ge 2$ be the gadget base and let $w = \lceil \log_{B_g} q \rceil$. Define the gadget vector

$$\mathbf{g} = (g_0, g_1, \ldots, g_{w-1}) = (1, B_g, B_g^2, \ldots, B_g^{w-1}) \in R_q^w.$$

For $c \in R_q$, the gadget decomposition algorithm $\mathsf{Decomp}_{B_g} : R_q \to R^w$ outputs $\mathbf{d} = (d_0, \ldots, d_{w-1}) \in R^w$ such that $c = \sum_{i=0}^{w-1} g_i d_i \pmod q$. For simplicity, we write $\mathsf{Decomp}$ for $\mathsf{Decomp}_{B_g}$ when the gadget base is clear from context. The digits are represented using centred representatives and therefore satisfy $\|d_i\|_\infty \le B_g/2$ for all $i = 0, \ldots, w-1$.

Let $\Lambda \subseteq \mathbb{R}^k$ be a full-rank lattice. We write $D_{\Lambda,\sigma,\mathbf{c}}$ for the discrete Gaussian distribution over $\Lambda$ with parameter $\sigma$ and centre $\mathbf{c}$, and write $D_{\Lambda,\sigma}$ when $\mathbf{c} = \mathbf{0}$. Using the coefficient embedding, $D_{R^m,\sigma}$ denotes the distribution over $R^m$ induced by $D_{\mathbb{Z}^{dm},\sigma}$ centred at $\mathbf{0}$.

**Definition 1 (Smoothing parameter by Micciancio and Regev [MR07]).** *Let $\Lambda \subseteq \mathbb{R}^k$ be a full-rank lattice and let $\Lambda^*$ denote its dual lattice. For $\epsilon > 0$, the $\epsilon$-smoothing parameter of $\Lambda$ is*

$$\eta_\epsilon(\Lambda) := \inf\left\{ s > 0 : \rho_{1/s}(\Lambda^* \setminus \{\mathbf{0}\}) \le \epsilon \right\}.$$

By abuse of notation, $\eta_\epsilon(R^m)$ denotes $\eta_\epsilon(\tau(R^m)) = \eta_\epsilon(\mathbb{Z}^{dm})$.

**Lemma 1 ( [MR07]).** *For every dimension $k$, $\eta_\epsilon(\mathbb{Z}^k) \le \sqrt{\frac{\ln(2k(1+1/\epsilon))}{\pi}}$.*

### 2.2 LWE and Noisy Error-Leakage LWE

We use the term LWE in a generalised sense over the ring $R$.

**Definition 2 (LWE over $R$).** *Let $D_{\mathbf{s}}$ be a distribution over $R^n$, and let $D_{\mathbf{e}}$ be a distribution over $R^m$. Define the real LWE distribution as*

$$\mathcal{D}^0_{\mathsf{LWE}} := \left\{ (\mathbf{A}, \mathbf{b}) \;\middle|\; \begin{array}{l} \mathbf{A} \xleftarrow{\$} U(R_q^{n\times m}), \\ \mathbf{s} \xleftarrow{\$} D_{\mathbf{s}}, \\ \mathbf{e} \xleftarrow{\$} D_{\mathbf{e}}, \\ \mathbf{b}^T = \mathbf{s}^T\mathbf{A} + \mathbf{e}^T \pmod q \end{array} \right\}$$

*and the uniform LWE distribution as*

$$\mathcal{D}^1_{\mathsf{LWE}} := U(R_q^{n\times m}) \times U(R_q^m).$$

*The $\mathsf{LWE}_{R,n,m,q,D_{\mathbf{s}},D_{\mathbf{e}}}$ assumption states that $\mathcal{D}^0_{\mathsf{LWE}} \approx_c \mathcal{D}^1_{\mathsf{LWE}}$.*

When $D_{\mathbf{e}} = D_{R^m,\sigma}$, we write $\mathsf{LWE}_{R,n,m,q,D_{\mathbf{s}},\sigma}$. When $n = 1$, the above corresponds to Ring-LWE; when $n > 1$, it corresponds to Module-LWE; and when $R = \mathbb{Z}$, it corresponds to standard LWE.

To model the leakage pattern induced by our re-encryption procedure, we introduce the following variant of LWE.

**Definition 3 (Noisy Error-Leakage LWE (NEL-LWE)).** *Let $D_{\mathbf{s}}$ be a distribution over $R^n$, and let $D_{R^m,\sigma}$ and $D_{R^m,\sigma_{\mathsf{pad}}}$ be discrete Gaussian distributions over $R^m$. The leakage variable $\boldsymbol{\ell}$ models a noisy observation of the LWE error.*

*Define the real NEL-LWE distribution as*

$$\mathcal{D}^0_{\mathsf{NEL}} = \left\{ (\mathbf{A}, \mathbf{b}, \boldsymbol{\ell}) \;\middle|\; \begin{array}{l} \mathbf{A} \xleftarrow{\$} U(R_q^{n\times m}), \\ \mathbf{s} \xleftarrow{\$} D_{\mathbf{s}}, \\ \mathbf{e} \xleftarrow{\$} D_{R^m,\sigma}, \\ \mathbf{e}_{\mathsf{pad}} \xleftarrow{\$} D_{R^m,\sigma_{\mathsf{pad}}}, \\ \mathbf{b}^T = \mathbf{s}^T\mathbf{A} + \mathbf{e}^T \pmod q, \\ \boldsymbol{\ell} = \mathbf{e} + \mathbf{e}_{\mathsf{pad}} \end{array} \right\}$$

*and the uniform NEL-LWE distribution as*

$$\mathcal{D}^1_{\mathsf{NEL}} = \left\{ (\mathbf{A}, \mathbf{u}, \boldsymbol{\ell}) \;\middle|\; \begin{array}{l} \mathbf{A} \xleftarrow{\$} U(R_q^{n\times m}), \\ \mathbf{u} \xleftarrow{\$} U(R_q^m), \\ \mathbf{e} \xleftarrow{\$} D_{R^m,\sigma}, \\ \mathbf{e}_{\mathsf{pad}} \xleftarrow{\$} D_{R^m,\sigma_{\mathsf{pad}}}, \\ \boldsymbol{\ell} = \mathbf{e} + \mathbf{e}_{\mathsf{pad}} \end{array} \right\}.$$

*The $\mathsf{NEL\text{-}LWE}_{R,n,m,q,D_{\mathbf{s}},\sigma,\sigma_{\mathsf{pad}}}$ assumption states that $\mathcal{D}^0_{\mathsf{NEL}} \approx_c \mathcal{D}^1_{\mathsf{NEL}}$.*

The NEL-LWE assumption can be viewed as a special case of the Leaky-LWE framework of Lai et al. [LSW25] with the leakage template $\boldsymbol{\ell} = (L_{\mathbf{s}}^T, L_{\mathbf{e}}^T)\begin{pmatrix}\mathbf{s}\\ \mathbf{e}\end{pmatrix} + \mathbf{f}$. Specifically, the leakage $\boldsymbol{\ell} = \mathbf{e} + \mathbf{e}_{\mathsf{pad}}$ corresponds to the error-only setting $L_{\mathbf{s}} = \mathbf{0}$ and $L_{\mathbf{e}} = \mathbf{I}_m$, with $\mathbf{f} = \mathbf{e}_{\mathsf{pad}}$.

**Corollary 1 (NEL-LWE from Leaky-LWE).** *Let $\epsilon = \epsilon(\lambda)$ be negligible. Suppose that the standard $\mathsf{LWE}_{R,n,m,q,D_{\mathbf{s}},\sigma^\star}$ assumption holds, where $\sigma^\star \geq \sqrt{2}\eta_\epsilon(R^m)$. Define $A := (\sigma^\star)^2 + 2\eta_\epsilon(R^m)^2$. Let $\sigma^2 > A$.*

*If*

$$\sigma_{\mathsf{pad}}^2 \geq \frac{\sigma^2 A}{\sigma^2 - A},$$

*then $\mathsf{NEL\text{-}LWE}_{R,n,m,q,D_{\mathbf{s}},\sigma,\sigma_{\mathsf{pad}}}$ holds, i.e., $\mathcal{D}^0_{\mathsf{NEL}} \approx_c \mathcal{D}^1_{\mathsf{NEL}}$.*

*Proof.* This claim follows from Theorem 3 of Lai et al. [LSW25] by setting $L_{\mathbf{s}} = \mathbf{0}$, $L_{\mathbf{e}} = \mathbf{I}_m$, and $\mathbf{f} = \mathbf{e}_{\mathsf{pad}}$. The operator norm of the error-leakage matrix is 1, and the leakage noise is sampled independently from $D_{R^m,\sigma_{\mathsf{pad}}}$. The stated lower bound on $\sigma_{\mathsf{pad}}$ is exactly the admissibility condition obtained from this instantiation. □

*Remark 1 (Ternary secrets).* In BGV-style instantiation, $\chi_s$ may be a ternary distribution. This is consistent with the NEL-LWE reduction, since the leakage is error-only and the underlying LWE assumption allows an arbitrary fixed secret distribution.

### 2.3 Owner-Encrypted Proxy Re-Encryption

Unlike standard public-key PRE, owner-encrypted proxy re-encryption (OE-PRE) does not allow arbitrary third parties to encrypt messages under the delegator's public key. Instead, the delegator generates the initial ciphertext herself using her secret key and later delegates decryption rights through a re-encryption key.

**Definition 4 (Owner-Encrypted Proxy Re-Encryption).** *An owner-encrypted proxy re-encryption scheme* OE-PRE *consists of the following PPT algorithms.*

- $\mathsf{Setup}(1^\lambda) \to \mathsf{pp}$. *The setup algorithm outputs public parameters* $\mathsf{pp}$.
- $\mathsf{ProxyInit}(\mathsf{pp}) \to \mathsf{st}_\mathsf{P} = \kappa_{\mathsf{pad}}$. *The proxy initialises its inner state.*
- $\mathsf{KeyGen}(\mathsf{pp}) \to (\mathsf{pk}, \mathsf{sk})$. *The key-generation algorithm outputs a public key and a secret key.*
- $\mathsf{ReKeyGen}(\mathsf{pp}, \mathsf{pk}_A, \mathsf{sk}_A, \mathsf{pk}_B) \to \mathsf{rk}_{A\to B}$. *Given the delegator's public and secret keys $(\mathsf{pk}_A, \mathsf{sk}_A)$ and the delegatee's public key $\mathsf{pk}_B$, the algorithm outputs a re-encryption key $\mathsf{rk}_{A\to B}$. The delegatee's secret key is not required.*
- $\mathsf{Enc}(\mathsf{pp}, \mathsf{sk}_A, \mathsf{m}) \to \mathsf{ct}_A$. *The owner-encryption algorithm takes the delegator's secret key $\mathsf{sk}_A$ and a message $\mathsf{m} \in \mathcal{M}$, and outputs an owner-generated ciphertext $\mathsf{ct}_A$.*

- $\mathsf{ReEnc}(\mathsf{pp}, \mathsf{rk}_{A\to B}, \mathsf{ct}_A, \mathsf{pk}_A, \mathsf{pk}_B; \mathsf{st}_\mathsf{P}) \to \mathsf{ct}_B$. *The re-encryption algorithm takes a re-encryption key and a ciphertext under the delegator's key domain, and outputs a ciphertext under the delegatee's key domain.*
- $\mathsf{Dec}(\mathsf{pp}, \mathsf{sk}, \mathsf{ct}) \to \mathsf{m}$ *or* $\perp$. *The decryption algorithm takes a secret key and a ciphertext, and outputs either a message* $\mathsf{m} \in \mathcal{M}$ *or* $\perp$.

*Remark 2 (Proxy initialisation).* The proxy may maintain local state across re-encryption queries. While PRE servers typically store re-encryption keys and access-control metadata, our construction also uses proxy-local state explicitly to ensure deterministic reuse of the padding value for each source ciphertext.

*Correctness.* The single-hop correctness notion requires that decrypting a ciphertext after one application of $\mathsf{ReEnc}$ recovers the original message with overwhelming probability. Since the security theorem proved in this work is single-hop, the detailed iterative $L$-hop correctness definition is deferred to Section 4, where we analyse bounded iteration only at the level of correctness.

### 2.4 Security for OE-PRE

We define single-hop IND-CPA and HRA security for OE-PRE by adapting the standard HRA game [AFGH06,Coh19,FKKP19] to the owner-encrypted setting. The game maintains a set $\mathsf{Deriv}$ that contains the challenge ciphertext and all ciphertexts obtained from it via admissible re-encryption. The main difference is that honest encryption uses the owner's secret key.

We use a selective-key formulation. All honest and corrupted users are created before the challenge phase, and the corruption status of each user is fixed throughout the experiment. Adaptive post-challenge key generation and adaptive corruptions are left outside the present model.

We use the following shorthand throughout the security proofs. For compactness, define

$$\mathsf{H2H}(i,j) \Leftrightarrow i \in \mathsf{Hon} \wedge j \in \mathsf{Hon}, \quad \mathsf{H2C}(i,j) \Leftrightarrow i \in \mathsf{Hon} \wedge j \in \mathsf{Cor},$$

and define $\mathsf{C2H}(i,j)$ and $\mathsf{C2C}(i,j)$ analogously.

**Definition 5 (Security for OE-PRE).** *Let* $\mathsf{ATK} \in \{\mathsf{OE\text{-}CPA}, \mathsf{OE\text{-}HRA}\}$. *An OE-PRE scheme* $\mathsf{OE\text{-}PRE}$ *is* $\mathsf{ATK}$*-secure if for every PPT adversary* $\mathcal{A}$*,*

$$\mathsf{Adv}^{\mathsf{ATK}}_{\mathsf{OE\text{-}PRE}}(\mathcal{A}) = \left| \Pr\left[\mathsf{ATK}^{b}_{\mathsf{OE\text{-}PRE},\mathcal{A}}(1^\lambda) = b\right] - \frac{1}{2} \right|$$

*is negligible, where the probability is over* $b \in \{0,1\}$*, the randomness of the experiment defined in Fig. 1, and the adversary. We call the cases* $\mathsf{ATK} = \mathsf{OE\text{-}CPA}$ *and* $\mathsf{ATK} = \mathsf{OE\text{-}HRA}$ *IND-CPA security and HRA security, respectively.*

*Remark 3 (Why we prove HRA directly).* Prior HRA analyses often use re-encryption simulatability as a modular tool. For our construction, a direct HRA

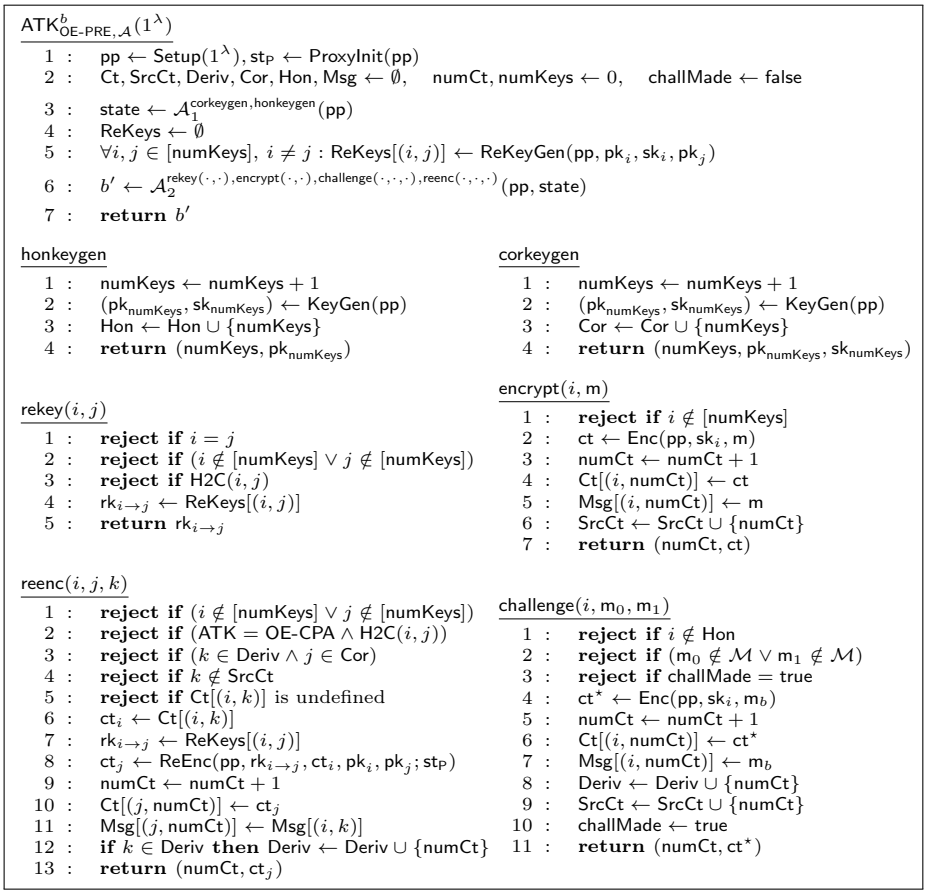


**Fig. 1.** HRA and IND-CPA games for OE-PRE. Here, $\mathsf{ATK} \in \{\mathsf{OE\text{-}HRA}, \mathsf{OE\text{-}CPA}\}$, and "**reject**" means returning $\perp$. The set SrcCt contains the identifiers of ciphertexts generated directly by encrypt or challenge, and re-encryption queries are accepted only for such source ciphertexts. The set Deriv contains the identifiers of the challenge ciphertext and all ciphertexts derived from it through re-encryption. The games differ only in how H2C re-encryption access is handled. In the OE-CPA game, such access is forbidden, whereas the OE-HRA game allows it except for ciphertexts derived from the challenge.

proof is more precise. Computational RES, as formulated by Döttling and Nishimaki [DN18, Appendix A.2], asks the simulator to handle adversarially supplied source ciphertexts, whereas our HRA game only re-encrypts honestly generated owner ciphertexts. Statistical RES is also stronger than needed, since it aims to erase the relevant dependence at the distributional level. We therefore prove HRA security directly inside the game, where the simulator only has to preserve visible honest-to-honest consistency and handle the honest-source noisy leakage appearing in admissible honest-to-corrupted responses.

## 3 BGV-Style Owner-Encrypted PRE

In this section, we present our BGV-style OE-PRE construction over the ring $R = \mathbb{Z}[X]/(X^d + 1)$.

### 3.1 Construction

All ciphertext components are elements of $R_q$, and messages are elements of $R_p$, represented in $R_q$ by their centred representatives. Let $\chi_s$ denote the secret-key distribution, $\chi_e$ the encryption-error distribution, and $\chi_{\mathsf{pad}}$ the padding-error distribution. Throughout the paper, $\chi_e$ and $\chi_{\mathsf{pad}}$ are instantiated as discrete Gaussian distributions. Recall that $\mathbf{g} = (1, B_g, \ldots, B_g^{w-1})$ denotes the gadget vector introduced in Section 2.

*Deterministic padding.* Let $F : \{0,1\}^\lambda \times \{0,1\}^* \to \{0,1\}^{\ell(\lambda)}$ be a secure PRF, and let $\mathsf{Sample}_{\chi_{\mathsf{pad}}} : \{0,1\}^{\ell(\lambda)} \to R$ be a deterministic sampler such that

$$\mathsf{Sample}_{\chi_{\mathsf{pad}}}(U_{\ell(\lambda)}) \approx_s \chi_{\mathsf{pad}}.$$

The proxy stores $\mathsf{st}_{\mathsf{P}} = \kappa_{\mathsf{pad}}$, where $\kappa_{\mathsf{pad}} \xleftarrow{\$} \{0,1\}^\lambda$. For a source ciphertext $\mathsf{ct}$, define

$$e_{\mathsf{pad}}(\mathsf{ct}) = \mathsf{Sample}_{\chi_{\mathsf{pad}}}\big(F_{\kappa_{\mathsf{pad}}}(\langle \mathsf{ct} \rangle)\big).$$

For fixed $\kappa_{\mathsf{pad}}$, this value is deterministic in the canonical ciphertext encoding $\langle \mathsf{ct} \rangle$ and is reused for repeated re-encryptions of the same source ciphertext.

*Algorithms.* The setup, key generation, encryption, and decryption algorithms are shown in Fig. 2, which follows the owner-encrypted variant of the BGV-style encryption framework. The proxy initialisation, re-encryption key generation, and re-encryption algorithms are shown in Fig. 3. Let the single-hop OE-PRE scheme be

$$\mathsf{OE\text{-}PRE} = (\mathsf{Setup}, \mathsf{ProxyInit}, \mathsf{KeyGen}, \mathsf{ReKeyGen}, \mathsf{Enc}, \mathsf{ReEnc}, \mathsf{Dec}).$$

$\underline{\mathsf{Setup}(1^\lambda)}$

Choose ring $R = \mathbb{Z}[X]/(X^d + 1)$
Choose moduli $p, q$ with $q$ prime, $q \equiv 1 \pmod{2d}$, $\gcd(p, q) = 1$
Choose the distributions $\chi_s$, $\chi_e := D_{R,\sigma_e}$, $\chi_{\mathsf{pad}} := D_{R,\sigma_{\mathsf{pad}}}$ over $R$
Choose gadget base $B_g \geq 2$ and set $w = \lceil \log_{B_g} q \rceil$
Choose a PRF $F$ and a deterministic sampler $\mathsf{Sample}_{\chi_{\mathsf{pad}}}$
$\mathsf{pp} \leftarrow (R, R_q, R_p, q, p, \chi_s, \chi_e, \chi_{\mathsf{pad}}, B_g, w, \mathbf{g}, F, \mathsf{Sample}_{\chi_{\mathsf{pad}}})$
**return** pp

$\underline{\mathsf{KeyGen}(\mathsf{pp})}$

$s \xleftarrow{\$} \chi_s$
$a \xleftarrow{\$} U(R_q), e \xleftarrow{\$} \chi_e$
$\mathsf{pk}_0 \leftarrow as + pe$
$\mathsf{pk}_1 \leftarrow -a$
$\mathsf{pk} \leftarrow (\mathsf{pk}_0, \mathsf{pk}_1), \mathsf{sk} \leftarrow s$
**return** $(\mathsf{pk}, \mathsf{sk})$

$\underline{\mathsf{Enc}(\mathsf{pp}, s, \mathsf{m})}$

**reject if** $\mathsf{m} \notin R_p$
$a \xleftarrow{\$} U(R_q), \quad e_{\mathsf{ct}} \xleftarrow{\$} \chi_e$
$c_0 \leftarrow as + pe_{\mathsf{ct}} + \mathsf{m}$
$c_1 \leftarrow -a$
$\mathsf{ct} \leftarrow (c_0, c_1)$
**return** ct

$\underline{\mathsf{Dec}(\mathsf{pp}, s, \mathsf{ct})}$

Parse $\mathsf{ct} = (c_0, c_1)$
$\mathsf{m}' \leftarrow c_0 + c_1 s \in R_q$
**return** $[\mathsf{m}']_q \bmod p$

**Fig. 2.** Setup, key generation, encryption, and decryption algorithms.

$\underline{\mathsf{ProxyInit}(\mathsf{pp})}$

$\kappa_{\mathsf{pad}} \xleftarrow{\$} \{0,1\}^\lambda$
$\mathsf{st}_{\mathsf{P}} \leftarrow \kappa_{\mathsf{pad}}$
**return** $\mathsf{st}_{\mathsf{P}}$

$\underline{\mathsf{ReKeyGen}(\mathsf{pp}, \mathsf{pk}_A, s_A, \mathsf{pk}_B)}$

Parse $\mathsf{pk}_B = (\mathsf{pk}_{B,0}, \mathsf{pk}_{B,1})$
**for** $i = 0, \ldots, w-1$ **do**
  $v_i \xleftarrow{\$} \chi_s, \ e_{i,0} \xleftarrow{\$} \chi_e, \ e_{i,1} \xleftarrow{\$} \chi_e$
  $\mathsf{rk}_{i,0} \leftarrow \mathsf{pk}_{B,0} v_i + pe_{i,0} + g_i s_A$
  $\mathsf{rk}_{i,1} \leftarrow \mathsf{pk}_{B,1} v_i + pe_{i,1}$
**end for**
$\mathsf{rk}_{A\to B} \leftarrow \{(\mathsf{rk}_{i,0}, \mathsf{rk}_{i,1})\}_{i=0}^{w-1}$
**return** $\mathsf{rk}_{A\to B}$

$\underline{\mathsf{ReEnc}(\mathsf{pp}, \mathsf{rk}_{A\to B}, \mathsf{ct}_A, \mathsf{pk}_A, \mathsf{pk}_B; \mathsf{st}_{\mathsf{P}})}$

Parse $\mathsf{st}_{\mathsf{P}} = \kappa_{\mathsf{pad}}$
Parse $\mathsf{ct}_A = (c_0, c_1)$
Parse $\mathsf{pk}_B = (\mathsf{pk}_{B,0}, \mathsf{pk}_{B,1})$
Parse $\mathsf{rk}_{A\to B} = \{(\mathsf{rk}_{i,0}, \mathsf{rk}_{i,1})\}_{i=0}^{w-1}$
$\mathbf{d} = (d_0, \ldots, d_{w-1})^T \leftarrow \mathsf{Decomp}(c_1)$
$\rho_{\mathsf{pad}} \leftarrow F_{\kappa_{\mathsf{pad}}}(\langle \mathsf{ct}_A \rangle)$
$e_{\mathsf{pad}} \leftarrow \mathsf{Sample}_{\chi_{\mathsf{pad}}}(\rho_{\mathsf{pad}})$
$r_{\mathsf{fr}} \xleftarrow{\$} \chi_s, \qquad e_{\mathsf{fr},0}, e_{\mathsf{fr},1} \xleftarrow{\$} \chi_e$
$C_{\mathsf{fr},0} \leftarrow \mathsf{pk}_{B,0} r_{\mathsf{fr}} + pe_{\mathsf{fr},0}$
$C_{\mathsf{fr},1} \leftarrow \mathsf{pk}_{B,1} r_{\mathsf{fr}} + pe_{\mathsf{fr},1}$
$c_0' \leftarrow c_0 + \sum_{i=0}^{w-1} d_i \mathsf{rk}_{i,0} + pe_{\mathsf{pad}} + C_{\mathsf{fr},0}$
$c_1' \leftarrow \sum_{i=0}^{w-1} d_i \mathsf{rk}_{i,1} + C_{\mathsf{fr},1}$
$\mathsf{ct}_B \leftarrow (c_0', c_1')$
**return** $\mathsf{ct}_B$

**Fig. 3.** Proxy initialisation, re-encryption key generation, and re-encryption algorithms.

*Remark 4 (Why fixed padding must be shielded).* Fixed padding cannot be added in the clear when a re-encryption key to an honest target is visible. Without the shield, a public transcript of a $\mathsf{H2H}$ re-encryption would satisfy

$$c_0' - c_0 - \sum_{i=0}^{w-1} d_i \mathsf{rk}_{i,0} = p e_{\mathsf{pad}}(\mathsf{ct}_A) \pmod q,$$

where $\mathbf{d} = \mathsf{Decomp}(c_1)$. Since $\gcd(p, q) = 1$, this would reveal the padding modulo $q$. The fresh target-key zero-encryption shield $C_{\mathsf{fr}} = (C_{\mathsf{fr},0}, C_{\mathsf{fr},1})$ changes the publicly computable value to

$$\left( c_0' - c_0 - \sum_i d_i \mathsf{rk}_{i,0},\ c_1' - \sum_i d_i \mathsf{rk}_{i,1} \right) = (p e_{\mathsf{pad}} + C_{\mathsf{fr},0},\ C_{\mathsf{fr},1}).$$

For an honest target (in particular, for a visible $\mathsf{H2H}$ re-encryption), this is a BGV encryption-form sample with offset $pe_{\mathsf{pad}}$, and the offset is hidden under the target secret. For a corrupted target, the adversary may remove the shielded term, but the padding is fixed for the source ciphertext and therefore cannot be averaged away.

*Relation to BGV encryption.* The construction inherits the standard BGV key-switching structure. The public keys, re-encryption-key components, and target-key zero-encryption shields are BGV-style LWE objects. We use the following standard pseudorandomness facts. The first item is a direct consequence of multi-sample LWE after multiplying by $p^{-1} \in R_q$. The second item is the standard pseudorandomness of BGV encryption-form and key-switching samples: BGV public-key encryption has the form $(\mathsf{pk}_0 v + pe_0 + \mu, \mathsf{pk}_1 v + pe_1)$, and key-switching keys are generated as encryptions of gadget-scaled secret-key digits under the target public key [BGV14, Lemma 4.3, Theorem 5.7]. We state the fact with arbitrary fixed offsets because in our proof the offset may be 0, $pe_{\mathsf{pad}}$, or $g_\tau s_A$. The proof in Appendix A only recalls the standard reduction viewpoint and explains why fixed offsets do not affect pseudorandomness.

**Lemma 2 (BGV-style pseudorandomness facts used in the proof).** *Assume that* $\gcd(p, q) = 1$. *Assume also that the chosen BGV parameters satisfy the standard multi-sample RLWE pseudorandomness conditions for the polynomially many public-key, owner-ciphertext, key-switching, and zero-encryption samples used in the scheme and in the security proof. Then the following holds.*

1. Owner secret-key ciphertexts. *Let* $(\mathsf{pk}_A, s_A) \leftarrow \mathsf{KeyGen}(\mathsf{pp})$, *where* $\mathsf{pk}_A = (\mathsf{pk}_{A,0}, \mathsf{pk}_{A,1}) = (a_A s_A + pe_A, -a_A)$. *Let* $Q = Q(\lambda)$ *be polynomially bounded. For* $t = 1, \ldots, Q$, *let* $\mathsf{m}_t \in R_p$, *sample independently* $a_t \leftarrow U(R_q)$ *and* $e_t \leftarrow \chi_e$, *and define* $\mathsf{ct}_t = (c_{t,0}, c_{t,1}) = (a_t s_A + pe_t + \mathsf{m}_t, -a_t) \in R_q^2$. *Then, for any PPT distinguisher not given* $s_A$,

$$(\mathsf{pk}_A, \{\mathsf{ct}_t\}_{t=1}^Q) \approx_c (\mathsf{pk}_A, \{U_t\}_{t=1}^Q),$$

*where* $U_t \leftarrow U(R_q^2)$ *are sampled independently.*

2. BGV encryption-form samples with arbitrary offsets. *Let* $(\mathsf{pk}_B, s_B) \leftarrow \mathsf{KeyGen}(\mathsf{pp})$, *where* $\mathsf{pk}_B = (\mathsf{pk}_{B,0}, \mathsf{pk}_{B,1}) = (a_B s_B + pe_B, -a_B)$. *Let* $T = T(\lambda)$ *be polynomially bounded. For* $t = 1, \ldots, T$, *let* $\mu_t \in R_q$ *be any efficiently computable offset fixed before the fresh randomness for the* $t$*-th sample is drawn and independent of the target secret* $s_B$. *Sample* $v_t \leftarrow \chi_s$ *and* $e_{t,0}, e_{t,1} \leftarrow \chi_e$, *and define* $H_t = (\mathsf{pk}_{B,0} v_t + pe_{t,0} + \mu_t,\ \mathsf{pk}_{B,1} v_t + pe_{t,1}) \in R_q^2$. *Then, for any PPT distinguisher not given* $s_B$,

$$(\mathsf{pk}_B, \{H_t\}_{t=1}^T) \approx_c (\mathsf{pk}_B, \{U_t\}_{t=1}^T),$$

*where* $U_t \leftarrow U(R_q^2)$ *are sampled independently.*

**Lemma 3 (Independent auxiliary information).** *Let* $X_0$ *and* $X_1$ *be two computationally indistinguishable distribution ensembles, and let* $Z$ *be an auxiliary distribution sampled independently of the challenge bit and of the sample from* $X_b$. *Then* $(X_0, Z) \approx_c (X_1, Z)$.

*Proof.* Any distinguisher for $(X_0, Z)$ and $(X_1, Z)$ gives a distinguisher for $X_0$ and $X_1$ by sampling $Z$ independently and forwarding the pair to the original distinguisher. Hence the advantage remains negligible. □

### 3.2 Correctness

We prove correctness by tracking the raw decryption value.

**Lemma 4 (Noise evolution).** *Let* $\mathsf{ct}_A = (c_0, c_1) \leftarrow \mathsf{Enc}(\mathsf{pp}, s_A, \mathsf{m})$ *be an owner-encrypted ciphertext, where* $\mathsf{m} \in R_p$. *Then* $c_0 + c_1 s_A = \mathsf{m} + pe_{\mathsf{ct}} \pmod q$, *where* $e_{\mathsf{ct}} \leftarrow \chi_e$ *is the encryption error.*

*Moreover, let* $\mathsf{ct}_B = (c_0', c_1') \leftarrow \mathsf{ReEnc}(\mathsf{pp}, \mathsf{rk}_{A\to B}, \mathsf{ct}_A, \mathsf{pk}_A, \mathsf{pk}_B; \mathsf{st}_P)$. *Write* $\mathsf{pk}_{B,0} = a_B s_B + pe_B, \mathsf{pk}_{B,1} = -a_B$, *and let* $\mathbf{d} = (d_0, \ldots, d_{w-1})^T = \mathsf{Decomp}(c_1)$. *Then*

$$c_0' + c_1' s_B = \mathsf{m} + p(e_{\mathsf{ct}} + e_{\mathsf{rk}} + e_{\mathsf{pad}} + e_{\mathsf{fr}}) \pmod q,$$

*where* $e_{\mathsf{rk}} := \sum_{i=0}^{w-1} d_i(e_B v_i + e_{i,0} + e_{i,1} s_B)$, *and* $e_{\mathsf{fr}} := e_B r_{\mathsf{fr}} + e_{\mathsf{fr},0} + e_{\mathsf{fr},1} s_B$ *is the decryption noise contributed by the target-key zero-encryption shield.*

The proof is a direct expansion of the owner-encryption, key-switching, and target-key zero-encryption equations, and is deferred to Appendix B.

**Lemma 5 (Noise bound).** *Assume that* $\|e_{\mathsf{ct}}\| \le B_e, \|e_{\mathsf{pad}}\| \le B_{\mathsf{pad}}$, *and, for all* $i$, $\|e_B\|, \|e_{i,0}\|, \|e_{i,1}\|, \|e_{\mathsf{fr},0}\|, \|e_{\mathsf{fr},1}\| \le B_e$, *and* $\|v_i\|, \|r_{\mathsf{fr}}\|, \|s_B\| \le B_s$. *Assume also that the centred gadget decomposition satisfies* $\|d_i\|_\infty \le B_g/2$ *for all* $i$. *Let* $\gamma_R$ *be the expansion factor of* $R$. *Then* $\|e_{\mathsf{rk}}\| \le B_{\mathsf{rk}} := w\gamma_R \left(\frac{\sqrt{d}B_g}{2}\right)(2\gamma_R B_e B_s + B_e)$, *and* $\|e_{\mathsf{fr}}\| \le B_{\mathsf{fr}} := 2\gamma_R B_e B_s + B_e$. *Consequently,*

$$\|e_{\mathsf{ct}} + e_{\mathsf{rk}} + e_{\mathsf{pad}} + e_{\mathsf{fr}}\| \le B_e + B_{\mathsf{rk}} + B_{\mathsf{pad}} + B_{\mathsf{fr}}.$$

The proof is a direct application of Proposition 1, the centred gadget bound $\|d_i\| \le \sqrt{d}B_g/2$, and the triangle inequality, and is deferred to Appendix C.

**Theorem 1 (Correctness).** *Assume that the bounds in Lemma 5 hold, that the centred representative of every plaintext satisfies* $\|[\mathsf{m}]_p\|_\infty \le B_{\mathsf{m}}$, *and that*

$$B_{\mathsf{m}} + p(B_e + B_{\mathsf{rk}} + B_{\mathsf{pad}} + B_{\mathsf{fr}}) < q/2.$$

*Then the scheme is single-hop correct. Namely, for every* $\mathsf{m} \in R_p$,

$$\mathsf{Dec}\Big(\mathsf{pp}, s_B, \mathsf{ReEnc}(\mathsf{pp}, \mathsf{rk}_{A\to B}, \mathsf{Enc}(\mathsf{pp}, s_A, \mathsf{m}), \mathsf{pk}_A, \mathsf{pk}_B; \mathsf{st}_{\mathsf{P}})\Big) = \mathsf{m}.$$

*Proof.* By Lemma 4, $c'_0 + c'_1 s_B = \mathsf{m} + pe_{\mathsf{tot}} \pmod q$, where $e_{\mathsf{tot}} = e_{\mathsf{ct}} + e_{\mathsf{rk}} + e_{\mathsf{pad}} + e_{\mathsf{fr}}$. By Lemma 5, $\|e_{\mathsf{tot}}\|_\infty \le \|e_{\mathsf{tot}}\| \le B_e + B_{\mathsf{rk}} + B_{\mathsf{pad}} + B_{\mathsf{fr}}$. Thus $\|[\mathsf{m}]_p + pe_{\mathsf{tot}}\|_\infty \le B_{\mathsf{m}} + p(B_e + B_{\mathsf{rk}} + B_{\mathsf{pad}} + B_{\mathsf{fr}}) < q/2$. Therefore, no coefficient wraps modulo $q$, and hence $[c'_0 + c'_1 s_B]_q = [\mathsf{m}]_p + pe_{\mathsf{tot}}$. The decryption algorithm returns $[c'_0 + c'_1 s_B]_q \bmod p = ([\mathsf{m}]_p + pe_{\mathsf{tot}}) \bmod p = \mathsf{m}$. □

*Remark 5 (Probabilistic correctness).* The theorem is stated under deterministic norm bounds. For discrete Gaussian errors and bounded secret distributions, standard tail bounds imply the required bounds on $e$-terms and secret/randomness terms with overwhelming probability for suitable choices of $B_e$, $B_s$, and $B_{\mathsf{pad}}$; the derived quantities $B_{\mathsf{rk}}$ and $B_{\mathsf{fr}}$ then follow from Lemma 5.

### 3.3 HRA Security

We now prove single-hop OE-HRA security directly in the HRA game. In this subsection, "single-hop" means that re-encryption queries are applied only to ciphertext records generated by encrypt or challenge; re-encryption of a previously re-encrypted ciphertext is handled only at the level of bounded correctness in Section 4. The proof does not introduce a stand-alone re-encryption simulatability notion. Instead, we keep the honest encryption and challenge oracles real and replace the HRA-relevant re-encryption interface by a stateful hybrid oracle. The hybrid oracle maintains the fixed padding table and the hidden noisy leakage needed for H2C responses, while preserving the public consistency of visible H2H re-encryption keys and outputs.

**Theorem 2 (Single-hop HRA security).** *Assume that $F$ is a secure PRF. Assume that the standard BGV-style pseudorandomness assumptions in Lemma 2 hold for the chosen BGV parameters. Assume further that the parameter conditions of Corollary 1 hold for any polynomial number of honest-source ciphertext records. Then* OE-PRE *is single-hop* OE-HRA*-secure in the game of Fig. 1.*

*Proof.* Fix a challenge bit $b \in \{0,1\}$. Let $G_0^b$ be the real single-hop OE-HRA experiment with challenge bit $b$. We transform $G_0^b$ into a game in which the challenge bit is information-theoretically hidden from the adversary. Throughout the proof, $\mathsf{H2H}(A,B)$ and $\mathsf{H2C}(A,C)$ denote the corruption types of the source-target pair as defined in the security experiment. The proof only concerns the single-hop experiment: a re-encryption query is admissible only when it is applied to a source ciphertext record generated by the honest encryption oracle or by the challenge oracle, and challenge-derived honest-to-corrupted re-encryptions are rejected by the HRA admissibility rule.

*Game $G_1^b$: random padding table and simulated re-encryption interface.* The honest key-generation, corrupted key-generation, honest encryption, and challenge oracles are unchanged. The game first replaces the PRF-derived padding values by a lazy random padding table and then answers all HRA-relevant re-encryption queries using proof-level tables. These tables are maintained lazily: whenever an entry is first accessed, the game samples it from the distribution specified below, stores it, and reuses the same value on all later accesses.

The table $\mathsf{Tbl}_e[k]$ stores a simulated owner-encryption error $\widetilde{e}_k \xleftarrow{\$} \chi_e$ for each honest-source ciphertext record $k$ that may appear in an admissible honest-to-corrupted re-encryption query. This table is only a proof device: its entries replace the real owner-encryption errors after the NEL-LWE hybrid described below.

The table $\mathsf{Tbl}_{\mathsf{pad}}[\langle \mathsf{ct} \rangle]$ stores one fixed padding value $\widetilde{e}_{\mathsf{pad}} \xleftarrow{\$} \chi_{\mathsf{pad}}$ for each source ciphertext encoding $\langle \mathsf{ct} \rangle$. The same value is reused for all re-encryptions of the same source ciphertext encoding. The real PRF-derived padding values are replaced by this random table. This changes the adversary's view by at most a negligible amount by PRF security and the statistical closeness of the deterministic sampler to $\chi_{\mathsf{pad}}$. The table preserves the essential fixedness property of the real scheme: the same source ciphertext encoding always receives the same padding value.

The table $\mathsf{Tbl}_{\mathsf{visRK}}[(A, B)]$ stores visible simulated switching-key components for each honest-to-honest pair satisfying $\mathsf{H2H}(A, B)$. On first access, the game samples

$$\left\{\widetilde{\mathsf{rk}}_{A\to B,\tau} = (\widetilde{\mathsf{rk}}_{A\to B,\tau,0}, \widetilde{\mathsf{rk}}_{A\to B,\tau,1}) \xleftarrow{\$} U(R_q^2)\right\}_{\tau=0}^{w-1}.$$

These components are returned to the adversary in visible honest-to-honest rekey queries and are reused in all corresponding honest-to-honest re-encryption outputs.

The table $\mathsf{Tbl}_{\mathsf{hidSW}}[(A, C)]$ stores hidden switching randomness for each honest-to-corrupted pair satisfying $\mathsf{H2C}(A, C)$. On first access, the game samples

$$\left\{(\widetilde{v}_\tau, \widetilde{e}_{\tau,0}, \widetilde{e}_{\tau,1}) \xleftarrow{\$} \chi_s \times \chi_e \times \chi_e\right\}_{\tau=0}^{w-1}.$$

The corresponding honest-to-corrupted re-encryption key is never revealed by the HRA admissibility rule. Nevertheless, the same hidden switching randomness is reused across all admissible re-encryption outputs for the same ordered pair $(A, C)$, matching the correlations induced by a fixed hidden re-encryption key.

Queries whose source user is corrupted are answered exactly as in the real experiment, since they do not carry protected honest-source noise. A $\mathsf{rekey}(A, C)$ query with $\mathsf{H2C}(A, C)$ is rejected, as in the real HRA game.

For a visible $\mathsf{rekey}(A, B)$ query with $\mathsf{H2H}(A, B)$, the game retrieves

$$\left\{\widetilde{\mathsf{rk}}_{A\to B,\tau} = (\widetilde{\mathsf{rk}}_{A\to B,\tau,0}, \widetilde{\mathsf{rk}}_{A\to B,\tau,1})\right\}_{\tau=0}^{w-1} = \mathsf{Tbl}_{\mathsf{visRK}}[(A, B)]$$

and returns these components to the adversary. By the BGV key-switching pseudorandomness fact in Lemma 2, these uniform pairs are computationally indistinguishable from the real key-switching components

$$(\mathsf{pk}_{B,0}v_\tau + pe_{\tau,0} + g_\tau s_A, \mathsf{pk}_{B,1}v_\tau + pe_{\tau,1}),$$

because they are BGV encryption-form samples under the hidden target secret $s_B$ with fixed offsets $\mu_\tau = g_\tau s_A$. This replacement is applied jointly to all polynomially many visible switching-key components, and subsequent re-encryption outputs are efficient post-processings of the same stored components.

For a $\mathsf{reenc}(A, B, k)$ query with $\mathsf{H2H}(A, B)$, let

$$\mathsf{ct} = (c_0, c_1) = \mathsf{Ct}[(A, k)], \qquad \mathbf{d} = (d_0, \ldots, d_{w-1})^T = \mathsf{Decomp}(c_1).$$

The game retrieves $\widetilde{e}_{\mathsf{pad}} = \mathsf{Tbl}_{\mathsf{pad}}[\langle \mathsf{ct} \rangle]$, retrieves the stored visible switching-key components $\{\widetilde{\mathsf{rk}}_{A\to B,\tau}\}_{\tau=0}^{w-1}$, samples a fresh target-key zero-encryption shield

$$\widetilde{C}_{\mathsf{fr}} = (\mathsf{pk}_{B,0}\widetilde{r}_{\mathsf{fr}} + p\widetilde{e}_{\mathsf{fr},0}, \mathsf{pk}_{B,1}\widetilde{r}_{\mathsf{fr}} + p\widetilde{e}_{\mathsf{fr},1}),$$

and returns

$$\widetilde{c}'_0 = c_0 + \sum_{\tau=0}^{w-1} d_\tau \widetilde{\mathsf{rk}}_{A\to B,\tau,0} + p\widetilde{e}_{\mathsf{pad}} + \widetilde{C}_{\mathsf{fr},0}, \qquad \widetilde{c}'_1 = \sum_{\tau=0}^{w-1} d_\tau \widetilde{\mathsf{rk}}_{A\to B,\tau,1} + \widetilde{C}_{\mathsf{fr},1}.$$

Thus repeated visible re-encryptions of the same source ciphertext under the same visible re-encryption key remain publicly consistent with a fixed switching key and a fixed padding value, exactly as in the real experiment. In particular, even if the adversary has obtained the visible re-encryption key for the ordered pair $(A, B)$, public subtraction yields only

$$(p\widetilde{e}_{\mathsf{pad}} + \widetilde{C}_{\mathsf{fr},0}, \widetilde{C}_{\mathsf{fr},1}).$$

This is a BGV encryption-form sample under the honest target public key with fixed offset $\mu = p\widetilde{e}_{\mathsf{pad}}$. The BGV encryption-form pseudorandomness fact in Lemma 2 applies with arbitrary fixed offsets: the ordinary zero-encryption shield corresponds to the special case $\mu = 0$, while the publicly visible shielded-padding transcript corresponds to the offset $\mu = p\widetilde{e}_{\mathsf{pad}}$. Since the fresh shield randomness is sampled after the offset is fixed, the fixed padding value is computationally hidden in visible honest-to-honest transcripts.

It remains to describe honest-to-corrupted re-encryption queries. Consider a $\mathsf{reenc}(A, C, k)$ query with $\mathsf{H2C}(A, C)$. Such a query is admissible only when $k$ is an honest-source, non-derived ciphertext record and is not a challenge-derived record forbidden by the HRA admissibility rule. Let

$$\mathsf{ct} = (c_0, c_1) = \mathsf{Ct}[(A, k)], \qquad \mathsf{m} = \mathsf{Msg}[(A, k)], \qquad \mathbf{d} = \mathsf{Decomp}(c_1).$$

In the real experiment, writing the source ciphertext as

$$\mathsf{ct} = (as_A + pe_{\mathsf{ct}} + \mathsf{m}, -a),$$

the honest-to-corrupted re-encryption output can be expressed as

$$c'_0 = \mathsf{m} + p(e_{\mathsf{ct}} + e_{\mathsf{pad}}) + \sum_{\tau=0}^{w-1} d_\tau(\mathsf{pk}_{C,0} v_\tau + p e_{\tau,0}) + C_{\mathsf{fr},0},$$
$$c'_1 = \sum_{\tau=0}^{w-1} d_\tau(\mathsf{pk}_{C,1} v_\tau + p e_{\tau,1}) + C_{\mathsf{fr},1}.$$

Indeed, the terms $g_\tau s_A$ in the switching key cancel because

$$\sum_{\tau=0}^{w-1} d_\tau g_\tau = c_1 \qquad \text{and} \qquad c_0 + c_1 s_A = \mathsf{m} + p e_{\mathsf{ct}}.$$

The re-encryption key $A \to C$ is never revealed in the HRA game, so the simulator only needs to reproduce the output distribution, including the correlations caused by a fixed hidden re-encryption key.

The game retrieves

$$\widetilde{e}_k = \mathsf{Tbl}_e[k], \qquad \widetilde{e}_{\mathsf{pad}} = \mathsf{Tbl}_{\mathsf{pad}}[\langle \mathsf{ct} \rangle],$$

retrieves $\{(\widetilde{v}_\tau, \widetilde{e}_{\tau,0}, \widetilde{e}_{\tau,1})\}_{\tau=0}^{w-1} = \mathsf{Tbl}_{\mathsf{hidSW}}[(A, C)]$, samples fresh shield randomness $(\widetilde{r}_{\mathsf{fr}}, \widetilde{e}_{\mathsf{fr},0}, \widetilde{e}_{\mathsf{fr},1})$, and returns

$$\widetilde{c}'_0 = \mathsf{m} + p(\widetilde{e}_k + \widetilde{e}_{\mathsf{pad}}) + \sum_{\tau=0}^{w-1} d_\tau(\mathsf{pk}_{C,0} \widetilde{v}_\tau + p\widetilde{e}_{\tau,0}) + \mathsf{pk}_{C,0} \widetilde{r}_{\mathsf{fr}} + p\widetilde{e}_{\mathsf{fr},0},$$
$$\widetilde{c}'_1 = \sum_{\tau=0}^{w-1} d_\tau(\mathsf{pk}_{C,1} \widetilde{v}_\tau + p\widetilde{e}_{\tau,1}) + \mathsf{pk}_{C,1} \widetilde{r}_{\mathsf{fr}} + p\widetilde{e}_{\mathsf{fr},1}.$$

This output has the same switching-randomness and shield-noise structure as in the real scheme under the corrupted target public key, while using the simulated noisy leakage $\widetilde{e}_k + \widetilde{e}_{\mathsf{pad}}$.

We now justify the replacement of $e_{\mathsf{ct}} + e_{\mathsf{pad}}$ by $\widetilde{e}_k + \widetilde{e}_{\mathsf{pad}}$. Fix an honest source user $A$. Collect all honest-source ciphertext records under $A$ that may be used in admissible $\mathsf{H2C}$ re-encryption queries. These records share the same secret $s_A$, so they are treated as one multi-sample NEL-LWE instance whose sample number is bounded by the total number of oracle queries. For a record $k$ in this collection, write

$$\mathsf{ct}_k = (a_k s_A + p e_k + \mathsf{m}_k, -a_k)$$

and set $\alpha_k = p^{-1} a_k$, using $\gcd(p, q) = 1$. The tuple

$$(\alpha_k, p^{-1}(c_{0,k} - \mathsf{m}_k), e_k + e_{\mathsf{pad},k})$$

has the form

$$(\alpha_k, \alpha_k s_A + e_k, e_k + e_{\mathsf{pad},k}).$$

Applying Corollary 1 jointly to all such records under the same source secret $s_A$, we obtain computational indistinguishability from the hybrid in which these tuples are replaced by

$$(\alpha_k, U_k, e_k + e_{\mathsf{pad},k}),$$

where each $U_k \xleftarrow{\$} U(R_q)$ is independent of the leakage values.

In this uniform hybrid, the second component is independent of the leakage. Therefore the leakage variables can be resampled without changing the distribution: for each record $k$,

$$(e_{\mathsf{pad},k}, e_k + e_{\mathsf{pad},k}) \equiv (\widetilde{e}_{\mathsf{pad},k}, \widetilde{e}_k + \widetilde{e}_{\mathsf{pad},k})$$

as distributions, where $e_k, \widetilde{e}_k \xleftarrow{\$} \chi_e$ and $e_{\mathsf{pad},k}, \widetilde{e}_{\mathsf{pad},k} \xleftarrow{\$} \chi_{\mathsf{pad}}$. The lazy padding table preserves the same fixed-padding semantics as the real experiment: all re-encryptions of the same source ciphertext encoding use the same sampled padding value.

Finally, by ordinary multi-sample LWE pseudorandomness together with Lemma 3, the independent uniform components $U_k$ can be changed back to LWE-form components while keeping the simulated leakages as independent auxiliary information. Thus

$$(\alpha_k, U_k, \widetilde{e}_k + \widetilde{e}_{\mathsf{pad},k}) \approx_c (\alpha_k, \alpha_k s_A + e'_k, \widetilde{e}_k + \widetilde{e}_{\mathsf{pad},k}),$$

where $e'_k \xleftarrow{\$} \chi_e$ is the owner-encryption error used only for the public ciphertext component. Applying the inverse affine transformation to the first two components yields a game in which every public source ciphertext is still distributed as an honest owner ciphertext, but every admissible honest-to-corrupted re-encryption response uses the simulated noisy leakage $\widetilde{e}_k + \widetilde{e}_{\mathsf{pad},k}$ rather than the real leakage $e_k + e_{\mathsf{pad},k}$.

The above argument is applied independently for each honest source user $A$, and the total number of records is polynomially bounded. The transition from $G_0^b$ to $G_1^b$ is therefore computationally indistinguishable by PRF security, the BGV-style pseudorandomness facts of Lemma 2, the independent-auxiliary lemma, and the NEL-LWE consequence of Corollary 1. Hence

$$G_0^b \approx_c G_1^b$$

for every $b \in \{0, 1\}$.

*Game $G_2^b$: replace the challenge ciphertext by uniform.* Let $A$ be the challenge owner and let

$$\mathsf{ct}_b^\star = (a^\star s_A + pe^\star + \mathsf{m}_b, -a^\star)$$

be the challenge ciphertext in $G_1^b$. In $G_2^b$, the game replaces $\mathsf{ct}_b^\star$ by an independent uniform element of $R_q^2$, while leaving all other oracle responses as in $G_1^b$.

By the owner-ciphertext pseudorandomness fact in Lemma 2, owner secret-key ciphertexts are jointly pseudorandom even in the presence of polynomially

many other owner ciphertexts, public keys, visible simulated re-encryption keys, and independent auxiliary values. The simulated re-encryption interface in $G_1^b$ uses only public ciphertext components, stored plaintexts for non-challenge encryption records, simulated error values, simulated switching-key components, fixed padding-table values, hidden switching randomness, and fresh shields. Challenge-derived honest-to-corrupted re-encryption queries are rejected by admissibility, and honest-to-honest re-encryption queries involving the challenge ciphertext are efficient post-processings of the public challenge ciphertext, visible simulated switching keys, fixed padding values, and fresh honest-target shields. Therefore replacing $\mathsf{ct}_b^\star$ by uniform is computationally indistinguishable, and

$$G_1^b \approx_c G_2^b.$$

In $G_2^b$, the challenge ciphertext is uniform and independent of $b$. Any admissible challenge-derived honest-to-honest re-encryption response is computed from this uniform ciphertext, simulated visible switching-key components, a fixed padding-table value, and a fresh honest-target zero-encryption shield; all these values are independent of $b$. Challenge-derived honest-to-corrupted re-encryption responses are rejected by the HRA admissibility rule. Non-challenge honest-to-corrupted re-encryption responses use plaintexts stored by the honest encryption oracle and simulated noisy leakages independent of the challenge bit. All remaining oracle responses either involve corrupted-source data or public values independent of $b$. Hence the entire adversarial view in $G_2^b$ is independent of $b$, and therefore

$$G_2^0 \equiv G_2^1.$$

Combining the hybrid transitions gives

$$G_0^0 \approx_c G_1^0 \approx_c G_2^0 \equiv G_2^1 \approx_c G_1^1 \approx_c G_0^1.$$

Thus no probabilistic polynomial-time adversary can distinguish the challenge bit with non-negligible advantage, which proves single-hop OE-HRA security. □

### 3.4 Comparison of HRA Noise Requirements

This section compares only the HRA-specific masking scale of our construction with the flooding scale used in the HRA-secure PRE scheme of Cohen et al. [CCG+25]. It is not a complete concrete parameter selection. Such a selection must also account for the key-switching and zero-encryption noise in the correctness bound, as well as the concrete security of the underlying LWE/RLWE parameters.

Our construction has three HRA-relevant noise ingredients. The fixed padding $e_{\mathsf{pad}}(\mathsf{ct}) \leftarrow D_{R,\sigma_{\mathsf{pad}}}$ is the direct computational mask for the inherited owner-encryption error. The owner-encryption width $\sigma_e$ must be large enough to satisfy the NEL-LWE admissibility condition of Corollary 1. The target-key shield $C_{\mathsf{fr}}$ is an ordinary BGV zero-encryption component used to hide the fixed padding in visible H2H transcripts; it is not the analogue of statistical flooding.

In the single-hop HRA proof, the NEL-LWE protected leakage is $e_{\mathsf{ct}} + e_{\mathsf{pad}}$, where $e_{\mathsf{ct}} \leftarrow D_{R,\sigma_e}$ and $e_{\mathsf{pad}} \leftarrow D_{R,\sigma_{\mathsf{pad}}}$. For simplicity, estimate the smoothing parameter in the single-sample case $m = 1$. Recall that $A = (\sigma^\star)^2 + 2\eta_\epsilon(R)^2$, where $\sigma^\star$ is the error width in the underlying standard LWE assumption of Corollary 1. Corollary 1 requires $\sigma_e^2 > A$ and

$$\sigma_{\mathsf{pad}}^2 \geq \frac{\sigma_e^2 A}{\sigma_e^2 - A}.$$

Equivalently, choosing $\sigma_e^2 = \alpha A$ with $\alpha > 1$ gives $\sigma_{\mathsf{pad}}^2 \geq \alpha A/(\alpha - 1)$. The balanced choice $\alpha = 2$ yields $\sigma_e, \sigma_{\mathsf{pad}} \geq \sqrt{2A}$.

The following calculation is only an order-of-magnitude illustration of the masking scale. It does not by itself imply a concrete security level. Take $d = 2048$, $\epsilon = 2^{-128}$, and $\sigma^\star \approx 3.2$. Under the coefficient embedding, $\tau(R) \cong \mathbb{Z}^d$. Lemma 1 gives $\eta_\epsilon(R) \leq 5.56$, so

$$A \leq 3.2^2 + 2 \cdot 5.56^2 \approx 72.1.$$

With $\alpha = 2$, we get

$$\sigma_e \approx 12, \qquad \sigma_{\mathsf{pad}} \approx 12,$$

and the protected leakage scale is approximately $\sqrt{\sigma_e^2 + \sigma_{\mathsf{pad}}^2} \approx 17$ under the independent Gaussian heuristic.

By contrast, Cohen et al.'s single-hop low $\nu$ HRA-secure mode [CCG+25, Section 6.1] uses flooding noise

$$\sigma_{\mathsf{fl}} = \sqrt{12\tau}\, t\, 2^{\nu/2},$$

where $t$ bounds the width of the re-encryption-dependent noise, $\tau$ bounds the number of re-encryption queries, and $\nu$ is the statistical security parameter. For $\tau = 2^9$, $\nu \geq 20$, and $t \approx 45$, this gives $\sigma_{\mathsf{fl}} \approx 3.6 \times 10^6$. The qualitative difference is that flooding statistically dominates the relevant re-encryption-dependent noise, whereas our scheme uses computational masking of the owner-encryption error together with ordinary BGV key-switching and zero-encryption noise.

## 4 Bounded Iteration: Correctness Only and the Multi-Hop HRA Obstacle

This section records bounded iterative correctness. It does not claim multi-hop HRA security.

**Definition 6 ($L$-hop correctness).** *An OE-PRE scheme is $L$-hop correct if, for every message $\mathsf{m} \in \mathcal{M}$, after running $\mathsf{Setup}$, $\mathsf{ProxyInit}$, generating keys $(\mathsf{pk}_i, \mathsf{sk}_i)$ for $i = 0, \ldots, L$, generating re-encryption keys $\mathsf{rk}_{i \to i+1}$ for $i = 0, \ldots, L-1$, computing $\mathsf{ct}_0 \leftarrow \mathsf{Enc}(\mathsf{pp}, \mathsf{sk}_0, \mathsf{m})$, and recursively setting*

$$\mathsf{ct}_{i+1} \leftarrow \mathsf{ReEnc}(\mathsf{pp}, \mathsf{rk}_{i \to i+1}, \mathsf{ct}_i, \mathsf{pk}_i, \mathsf{pk}_{i+1}; \mathsf{st}_{\mathsf{P}}),$$

*we have $\mathsf{Dec}(\mathsf{pp}, \mathsf{sk}_i, \mathsf{ct}_i) = \mathsf{m}$ for all $i \in \{0, \ldots, L\}$ with overwhelming probability. The single-hop setting corresponds to $L = 1$.*

Consider a chain $U_0 \to U_1 \to \cdots \to U_L$, where $\mathsf{ct}^{(0)} = (c_0^{(0)}, c_1^{(0)})$ is generated by the owner $U_0$, and let $s_t$ be the secret key of $U_t$. Write the raw decryption invariant after the $t$-th hop as

$$c_0^{(t)} + c_1^{(t)} s_t = \mathsf{m} + pE_t \pmod{q}.$$

For the initial owner ciphertext, $E_0 = e_{\mathsf{ct}}$. If $\mathsf{ct}^{(t-1)}$ is re-encrypted from $U_{t-1}$ to $U_t$, then

$$E_t = E_{t-1} + e_{\mathsf{pad}}^{(t)} + e_{\mathsf{rk}}^{(t)} + e_{\mathsf{fr}}^{(t)},$$

where

$$e_{\mathsf{rk}}^{(t)} = \sum_{\tau=0}^{w-1} d_\tau^{(t-1)} \big( e_{U_t} v_\tau^{(t)} + e_{\tau,0}^{(t)} + e_{\tau,1}^{(t)} s_t \big), \qquad e_{\mathsf{fr}}^{(t)} = e_{U_t} r_{\mathsf{fr}}^{(t)} + f_0^{(t)} + f_1^{(t)} s_t.$$

Thus, after $L$ hops,

$$E_L = e_{\mathsf{ct}} + \sum_{t=1}^{L} e_{\mathsf{pad}}^{(t)} + \sum_{t=1}^{L} e_{\mathsf{rk}}^{(t)} + \sum_{t=1}^{L} e_{\mathsf{fr}}^{(t)}.$$

**Theorem 3 (Bounded multi-hop correctness).** *Let $B_{\mathsf{m}}$ be a bound on the centred plaintext representative. Suppose that, except with negligible probability, the accumulated legacy-and-padding term satisfies $\left\| e_{\mathsf{ct}} + \sum_{t=1}^{L} e_{\mathsf{pad}}^{(t)} \right\| \le B_{\mathsf{leg},L}$, and for every hop $t \in \{1, \ldots, L\}$, the key-switching and zero-encryption-shield errors satisfy $\|e_{\mathsf{rk}}^{(t)}\| \le B_{\mathsf{rk}}$ and $\|e_{\mathsf{fr}}^{(t)}\| \le B_{\mathsf{fr}}$. If*

$$B_{\mathsf{m}} + p\big(B_{\mathsf{leg},L} + LB_{\mathsf{rk}} + LB_{\mathsf{fr}}\big) < q/2,$$

*then the $L$-hop ciphertext decrypts correctly.*

*Proof.* The proof is by induction on the number of hops. For $t = 0$, owner encryption gives $c_0^{(0)} + c_1^{(0)} s_0 = \mathsf{m} + pe_{\mathsf{ct}} \pmod{q}$. Assume $c_0^{(t-1)} + c_1^{(t-1)} s_{t-1} = \mathsf{m} + pE_{t-1} \pmod{q}$. Since $\sum_\tau d_\tau^{(t-1)} g_\tau = c_1^{(t-1)}$, substituting the re-encryption-key components and the fresh zero-encryption shield into the raw decryption under $s_t$ gives

$$c_0^{(t)} + c_1^{(t)} s_t = \mathsf{m} + p(E_{t-1} + e_{\mathsf{pad}}^{(t)} + e_{\mathsf{rk}}^{(t)} + e_{\mathsf{fr}}^{(t)}) \pmod{q}.$$

Thus the recurrence for $E_t$ holds. The stated bounds imply $\|E_L\| \le B_{\mathsf{leg},L} + LB_{\mathsf{rk}} + LB_{\mathsf{fr}}$, and the modulus condition prevents wraparound of the centred raw decryption value. Reducing modulo $p$ recovers $\mathsf{m}$. □

The term $B_{\mathsf{leg},L}$ can be bounded conservatively by $B_e + \sum_{t=1}^{L} B_{\mathsf{pad}}^{(t)}$, or more tightly using Gaussian tail bounds. If $e_{\mathsf{ct}}$ has width $\sigma_e$ and each padding term has width $\sigma_{\mathsf{pad}}$, then the legacy-and-padding variance grows as

$$\sigma_{\mathsf{leg},L}^2 = \sigma_e^2 + L\sigma_{\mathsf{pad}}^2.$$

Thus $B_{\mathsf{leg},L}$ grows like $O(\sqrt{L})$ under Gaussian tail estimates, while the conservative contribution of key-switching and zero-encryption-shield errors grows linearly as $L(B_{\mathsf{rk}} + B_{\mathsf{fr}})$.

This is only a correctness statement. In a chain such as $A \xrightarrow{\mathsf{H2H}} B \xrightarrow{\mathsf{H2C}} C$, the final corrupted target may see leakage involving the first-hop effective key-switching and shield errors, which depend on the honest intermediate secret $s_B$. Handling this appears to require a joint leakage theorem for BGV effective errors, or a stronger leakage assumption such as known-covariance RLWE [MS25, Definition 12]. We therefore restrict HRA security to one hop.

## 5 Discussion: Application Scope and Future Directions

### 5.1 Application Scope

The proved construction targets single-hop access delegation for owner-generated ciphertexts. This is narrower than general public-key PRE, because only the data owner generates the initial ciphertext, and narrower than general multi-hop PRE, because our HRA theorem covers only one re-encryption hop. It nevertheless captures a common access-control pattern: a data owner encrypts an object before outsourcing it, and a storage server or broker later transforms that ciphertext once for an authorised recipient. Encrypted file sharing and storage [AFGH06] and publish/subscribe delivery [PRSV17] are natural examples.

The construction should not be read as a replacement for PRE systems that require arbitrary public encryption or HRA-secure multi-hop routing. Its purpose is to isolate the owner-generated single-hop delegation step and show that the inherited-noise leakage in this step can be handled by computational masking rather than by statistical flooding.

### 5.2 Future Parameter Design and Hybrid Multi-hop Extension

A concrete parameter design should choose $d, q, p, B_g, \sigma_e, \sigma_{\mathsf{pad}}$ jointly so that the NEL-LWE condition of Corollary 1, the correctness bound of Theorem 1, and the ordinary BGV pseudorandomness requirements hold simultaneously. It should then estimate concrete LWE/RLWE security and compare ciphertext size, re-encryption-key size, noise budget, and re-encryption cost with flooding-based HRA-secure PRE.

A second direction is a hybrid computational/statistical multi-hop design. The first owner-generated hop could use NEL-LWE padding to avoid the initial large flooding term, while later hops in a longer lineage could use calibrated flooding, possibly combined with modulus switching and more precise noise accounting [CCG+25, Ogi26]. Whether this can yield smaller later flooding widths requires a separate parameter analysis and is left for future work.

## 6 Conclusion

We presented a BGV-style owner-encrypted proxy re-encryption construction that achieves single-hop HRA security without statistical noise flooding. The construction combines deterministic per-ciphertext padding with a fresh target-key zero-encryption shield. The padding masks the inherited owner-encryption error in honest-to-corrupted re-encryption queries, while the shield prevents the fixed padding from being publicly extracted in visible honest-to-honest transcripts.

The security proof is a direct HRA hybrid. Its core step uses NEL-LWE to handle the noisy leakage $e_{\mathsf{ct}} + e_{\mathsf{pad}}$, and the required NEL-LWE condition follows from the Leaky-LWE framework under suitable parameters. The resulting security claim is based on standard LWE/RLWE-style pseudorandomness assumptions, the NEL-LWE consequence of Leaky-LWE, and PRF security.

We also analysed bounded iteration at the level of correctness and identified the obstacle to multi-hop HRA security: after an honest-to-honest hop, the inherited effective noise contains key-switching and zero-encryption-shield errors depending on the honest intermediate secret. Extending HRA security beyond one hop without statistical flooding appears to require a new leakage-resilience theorem for BGV effective errors, a stronger effective-error-leakage assumption, or a hybrid computational/statistical refresh mechanism.

## References

AFGH06. Giuseppe Ateniese, Kevin Fu, Matthew Green, and Susan Hohenberger. Improved proxy re-encryption schemes with applications to secure distributed storage. *ACM Transactions on Information and System Security*, 9(1):1–30, February 2006. https://dl.acm.org/doi/10.1145/1127345.1127346.

BBS98. Matt Blaze, Gerrit Bleumer, and Martin Strauss. Divertible protocols and atomic proxy cryptography. In Kaisa Nyberg, editor, *Advances in Cryptology — EUROCRYPT'98*, pages 127–144, Berlin, Heidelberg, 1998. Springer.

BGV14. Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (Leveled) Fully Homomorphic Encryption without Bootstrapping. *ACM Trans. Comput. Theory*, 6(3):13:1–13:36, July 2014. https://dl.acm.org/doi/10.1145/2633600.

CCG+25. Aloni Cohen, David Bruce Cousins, Nicholas Genise, Erik Kline, Yuriy Polyakov, and Saraswathy Rv. HRA-Secure Homomorphic Lattice-Based Proxy Re-Encryption with Tight Security. *IACR Communications in Cryptology*, 2(1):1–30, April 2025. https://cic.iacr.org/p/2/1/22.

CCL+14. Nishanth Chandran, Melissa Chase, Feng-Hao Liu, Ryo Nishimaki, and Keita Xagawa. Re-encryption, Functional Re-encryption, and Multi-hop Re-encryption: A Framework for Achieving Obfuscation-Based Security and Instantiations from Lattices. In Hugo Krawczyk, editor, *Public-Key Cryptography – PKC 2014*, pages 95–112, Berlin, Heidelberg, 2014. Springer.

Coh19. Aloni Cohen. What About Bob? The Inadequacy of CPA Security for Proxy Reencryption. In Dongdai Lin and Kazue Sako, editors, *Public-Key Cryptography – PKC 2019*, pages 287–316, Cham, 2019. Springer International Publishing.

DN18. Nico Döttling and Ryo Nishimaki. Universal Proxy Re-Encryption. https://eprint.iacr.org/2018/840, 2018.

FKKP19. Georg Fuchsbauer, Chethan Kamath, Karen Klein, and Krzysztof Pietrzak. Adaptively Secure Proxy Re-encryption. In Dongdai Lin and Kazue Sako, editors, *Public-Key Cryptography – PKC 2019*, pages 317–346, Cham, 2019. Springer International Publishing.

LSW25. Russell Lai, Monisha Swarnakar, and Ivy Woo. Leaky LWE: Learning with Errors with Semi-Adaptive Secret- and Error-Leakage. *IACR Communications in Cryptology*, 2(3):cc2–3–35, October 2025. https://cic.iacr.org/p/2/3/21.

MR07. Daniele Micciancio and Oded Regev. Worst-Case to Average-Case Reductions Based on Gaussian Measures. *SIAM Journal on Computing*, 37(1):267–302, January 2007. https://epubs.siam.org/doi/abs/10.1137/S0097539705447360.

MS25. Daniele Micciancio and Adam Suhl. Simulation-Secure Threshold PKE from LWE with Polynomial Modulus. *IACR Communications in Cryptology*, 1(4):cc1–3–24, January 2025. https://cic.iacr.org/p/1/4/2.

Ogi26. Tabitha Ogilvie. IND-CPA-D and KR-D Security with Reduced Noise from the HintLWE Problem. In Goichiro Hanaoka and Bo-Yin Yang, editors, *Advances in Cryptology – ASIACRYPT 2025*, pages 263–295, Singapore, 2026. Springer Nature.

OT26. Hiroki Okada and Tsuyoshi Takagi. Low Communication Threshold FHE from Standard (Module-)LWE. In Goichiro Hanaoka and Bo-Yin Yang, editors, *Advances in Cryptology – ASIACRYPT 2025*, pages 165–198, Singapore, 2026. Springer Nature.

PRSV17. Yuriy Polyakov, Kurt Rohloff, Gyana Sahu, and Vinod Vaikuntanathan. Fast Proxy Re-Encryption for Publish/Subscribe Systems. *ACM Trans. Priv. Secur.*, 20(4):14:1–14:31, September 2017. https://dl.acm.org/doi/10.1145/3128607.

WWXW25. Xiaohan Wan, Yang Wang, Haiyang Xue, and Mingqiang Wang. Unbounded Multi-hop Proxy Re-encryption with HRA Security: An LWE-Based Optimization. In Willy Susilo and Josef Pieprzyk, editors, *Information Security and Privacy*, pages 124–144, Singapore, 2025. Springer Nature.

ZLH24. Yunxiao Zhou, Shengli Liu, and Shuai Han. Multi-hop Fine-Grained Proxy Re-encryption. In Qiang Tang and Vanessa Teague, editors, *Public-Key Cryptography – PKC 2024*, pages 161–192, Cham, 2024. Springer Nature Switzerland.

## A Proof of BGV-Style Pseudorandomness Facts

*Proof (of Lemma 2).* For Item 1, define the efficiently computable bijection

$$\Phi_\nu(x_0, x_1) = (-p^{-1}x_1,\ p^{-1}(x_0 - \nu)).$$

Since $\gcd(p, q) = 1$, multiplication by $p^{-1}$ is a permutation of $R_q$. Applying $\Phi_{\mathsf{m}_t}$ to $\mathsf{ct}_t = (a_t s_A + pe_t + \mathsf{m}_t, -a_t)$ gives

$$\Phi_{\mathsf{m}_t}(\mathsf{ct}_t) = (p^{-1}a_t,\ (p^{-1}a_t)s_A + e_t).$$

Since $a_t \leftarrow U(R_q)$, also $p^{-1}a_t \leftarrow U(R_q)$. Thus the transformed ciphertexts are standard LWE samples with secret $s_A$ and error distribution $\chi_e$. The public key $\mathsf{pk}_A$, after the same scaling of its first component, is another sample of the same LWE form. By the multi-sample LWE assumption, the joint distribution is computationally indistinguishable from uniform. Applying the inverse bijection preserves indistinguishability. Strictly speaking, the multi-sample LWE hybrid first replaces the joint transcript consisting of the public key and the owner-ciphertext samples by uniform. Since the public key alone is also LWE-pseudorandom, and the replacement ciphertexts are independent uniform elements, a triangle argument together with Lemma 3 yields the stated form in which the public key is kept in its real distribution.

For Item 2, we use the standard BGV encryption-form pseudorandomness fact under the target public key. Conditioned on any fixed offset sequence $(\mu_1, \ldots, \mu_T)$ that is fixed before the fresh encryption randomness and is independent of the target secret $s_B$, the samples

$$H_t = (\mathsf{pk}_{B,0}v_t + pe_{t,0} + \mu_t,\ \mathsf{pk}_{B,1}v_t + pe_{t,1})$$

are BGV encryptions, under $\mathsf{pk}_B$, of the fixed offsets $\mu_t$. The standard BGV security proof shows that polynomially many such encryption-form samples are pseudorandom under the same LWE/RLWE parameters used for public-key encryption and key switching. Equivalently, after subtracting the fixed offset from the first component, one obtains the usual zero-message BGV encryption-form samples, and adding the fixed offset back is a deterministic translation of the first component, which preserves indistinguishability from uniform. If the offsets are sampled from some distribution independent of the target secret and then fixed before the fresh encryption randomness, the same conclusion follows by conditioning on the offsets and averaging. Taking $\mu_t = g_t s_A$ gives the re-encryption-key components, taking $\mu_t = 0$ gives a target-key zero encryption, and taking $\mu_t = pe_{\mathsf{pad}}$ gives the shielded fixed-padding transcript used in the HRA proof. □

## B Proof of Noise Evolution

*Proof (of Lemma 4).* The first identity follows directly from the owner-encryption form $c_0 = as_A + pe_{\mathsf{ct}} + \mathsf{m}, c_1 = -a$.

For the re-encryption-key component, we have

$$\begin{aligned}\mathsf{rk}_{i,0} + \mathsf{rk}_{i,1}s_B &= (\mathsf{pk}_{B,0}v_i + pe_{i,0} + g_i s_A) + (\mathsf{pk}_{B,1}v_i + pe_{i,1})s_B \\ &= g_i s_A + p(e_B v_i + e_{i,0} + e_{i,1}s_B).\end{aligned}$$

For the zero-encryption shield, by construction, $C_{\mathsf{fr},0} = \mathsf{pk}_{B,0} r_{\mathsf{fr}} + p e_{\mathsf{fr},0}, C_{\mathsf{fr},1} = \mathsf{pk}_{B,1} r_{\mathsf{fr}} + p e_{\mathsf{fr},1}$. Therefore,

$$\begin{aligned} C_{\mathsf{fr},0} + C_{\mathsf{fr},1} s_B &= (\mathsf{pk}_{B,0} + \mathsf{pk}_{B,1} s_B) r_{\mathsf{fr}} + p(e_{\mathsf{fr},0} + e_{\mathsf{fr},1} s_B) \\ &= p(e_B r_{\mathsf{fr}} + e_{\mathsf{fr},0} + e_{\mathsf{fr},1} s_B) \\ &= p e_{\mathsf{fr}}. \end{aligned}$$

Since $\sum_{i=0}^{w-1} g_i d_i = c_1 \pmod q$, the modified re-encryption algorithm gives

$$\begin{aligned} c_0' + c_1' s_B &= c_0 + \sum_{i=0}^{w-1} d_i(\mathsf{rk}_{i,0} + \mathsf{rk}_{i,1} s_B) + p e_{\mathsf{pad}} + C_{\mathsf{fr},0} + C_{\mathsf{fr},1} s_B \\ &= c_0 + c_1 s_A + p \sum_{i=0}^{w-1} d_i(e_B v_i + e_{i,0} + e_{i,1} s_B) + p e_{\mathsf{pad}} + p e_{\mathsf{fr}} \\ &= \mathsf{m} + p(e_{\mathsf{ct}} + e_{\mathsf{rk}} + e_{\mathsf{pad}} + e_{\mathsf{fr}}) \pmod q. \end{aligned}$$

The lemma is proved. □

## C Proof of Noise Bound

*Proof (of Lemma 5).* For $\varepsilon_i := e_B v_i + e_{i,0} + e_{i,1} s_B$, the expansion-factor (Proposition 1) bound gives $\|\varepsilon_i\| \le 2\gamma_R B_e B_s + B_e$. Moreover, $\|d_i\| \le \sqrt{d}\, B_g/2$. Hence

$$\begin{aligned} \|e_{\mathsf{rk}}\| &= \left\| \sum_{i=0}^{w-1} d_i \varepsilon_i \right\| \\ &\le \sum_{i=0}^{w-1} \gamma_R \|d_i\| \, \|\varepsilon_i\| \\ &\le w \gamma_R \left( \frac{\sqrt{d} B_g}{2} \right) (2\gamma_R B_e B_s + B_e). \end{aligned}$$

Similarly, for $e_{\mathsf{fr}} = e_B r_{\mathsf{fr}} + e_{\mathsf{fr},0} + e_{\mathsf{fr},1} s_B$, we have $\|e_{\mathsf{fr}}\| \le 2\gamma_R B_e B_s + B_e = B_{\mathsf{fr}}$. The final claim follows from the triangle inequality. □