

On the eCK Security of MQV-Style LWE-Based Authenticated Key Exchange

Ruoyu Ding¹, Tianrun Yu¹, Yanbin Pan², Chi Cheng^{3,4✉}, and Jian Weng^{3✉}

¹ College of Cyber Security, Jinan University, Guangzhou, China

² State Key Laboratory of Mathematical Sciences, Academy of Mathematics and Systems Science, Chinese Academy of Sciences, Beijing, China

³ Cyberspace Institute of Advanced Technology, Pei Dingyi Cryptography Research Institute, Guangzhou University, Guangzhou 510006, China

⁴ Faculty of Data Science, City University of Macau, Macau, China
chengchizz@qq.com, cryptjweng@gmail.com

Abstract. Designing Authenticated Key Exchange (AKE) directly from the Learning with Errors (LWE) assumption, without relying on generic constructions, remains a prominent and challenging goal in post-quantum cryptography. Such designs inherently require reconciliation mechanisms, which may introduce additional leakage against active adversaries under key reuse. Yet previous signal leakage attacks of Bindel et al. (LATIN-CRYPT 2021) and Qin et al. (ESORICS 2022) fail against protocols hardened with dedicated countermeasures.

In this paper, we propose new signal leakage attacks capable of compromising LWE-based (authenticated) key exchange, where prior attacks fail. First, we target the landmark MQV-style AKE protocol proposed by Zhang et al. (EUROCRYPT 2015), known as ZZDSD-AKE. Determining whether it can natively achieve security under the strong extended Canetti-Krawczyk (eCK) model has remained a long-standing open question. By exploiting the ephemeral key compromise defined within the eCK model, our new attack based on a *geometric* perspective recovers the static secret key from signal leakage, definitively resolving this open question with a negative result. Furthermore, we extend our methodology to the GDLL-KE protocol of Gao et al. (IEEE TC 2018), which is designed to be robust against key-reuse via a randomized noise countermeasure. Our attack shows that its ad-hoc countermeasure fails to eliminate secret-dependent leakage.

Our implementations demonstrate that the proposed attacks recover the secret key in roughly 1,700 queries against ZZDSD-AKE and 180 queries against GDLL-KE, respectively. These results reveal that MQV-style structures and simple randomization are insufficient for securing AKEs directly based on LWE against sophisticated active attacks, providing critical insights for future protocol design and standardization.

Keywords: Learning with errors · Key exchange protocol · Signal leakage attack · Post-quantum cryptography

1 Introduction

In the classical domain, implicit authentication is best exemplified by the Menezes-Qu-Vanstone (MQV) protocol [36,31] and its variants (e.g., HMQV [29], CMQV [46], and OAKE [50]), which directly combine static and ephemeral keys. The corresponding design paradigm is widely referred to as MQV-style AKE. However, adapting this MQV-style structure directly to the (Ring) Learning with Errors (LWE) setting [43,34] is fundamentally hindered by the inherent noise of lattice-based cryptography.

In a basic (Ring-)LWE-based protocol analogous to the Diffie-Hellman (DH) key exchange [14], Alice computes and sends $p_A = as_A + e_A$ to Bob, and Bob responds with $p_B = as_B + e_B$. Due to the error terms e_A and e_B , the parties only arrive at approximately equal shared values close to as_As_B . To ensure exact key agreement, *reconciliation* mechanisms or *signal* functions [18,38,33] are required [1,9,7]. The signal function, however, may leak information about the secret key under key reuse scenarios, leading to signal leakage attacks [15,6,41]. Another approach is to build key exchange protocols from (Ring-)LWE-based public-key encryption [34,8], and can be extended to AKE via generic constructions [21,55,12]. However, the latter is incompatible with well-known constructions of MQV-style AKE [49], and incurs extra primitives and design complexity [37]. Inspired by HMQV, Zhang et al. [53] proposed a landmark AKE protocol directly from Ring-LWE using the signal function from [18], which we call ZZDSD-AKE. This protocol inherits the simplicity of HMQV by avoiding signatures and MACs, and it is proven secure in the basic Bellare-Rogaway (BR) security model [4] with weak forward secrecy.

Nevertheless, there exists a notable security gap between ZZDSD-AKE and HMQV. Specifically, HMQV has been proven secure in stronger security models [29,3,28], such as the Canetti-Krawczyk (CK) model [10] and the extended CK (eCK) model [30]. These models capture more realistic attack scenarios, including key-compromise impersonation and ephemeral key leakage. In contrast, ZZDSD-AKE has only been proven BR-secure, and the authors explicitly noted that a proof in the CK model is not viable. The eCK model is widely regarded as a stronger model, but it is incomparable to CK (i.e., neither implies the other) [13]. Thus, an analysis for ZZDSD-AKE in the eCK model is still missing.

A rigorous eCK analysis of ZZDSD-AKE is critical, bridging the fundamental gap between theoretical MQV-style designs and their practical deployment risks. Compared with the BR model, the eCK model allows the adversary to reveal ephemeral secret keys, which corresponds to real-world threats such as side-channel attacks [23] or insecure random number generators [19]. Whether ZZDSD-AKE is secure in the eCK model remains an open question [6]. Moreover, an eCK analysis of ZZDSD-AKE helps clarify the security boundary of MQV-style LWE-based AKE constructions, offering deeper insight into the security loss incurred when transferring the classical HMQV paradigm into the LWE setting. Indeed, prior works have shown that BR security alone may be insufficient in this context. For example, Bindel et al. [6] analyzed another BR-

secure LWE-based AKE [16] that shares a similar structure with the eCK-secure CMQV, and demonstrated it to be eCK-insecure via a signal leakage attack, which was further improved by Qin et al. [41]. However, both works failed to extend their attack to ZZDSD-AKE. These observations motivate a thorough eCK analysis of ZZDSD-AKE, both to resolve this long-standing open problem and to better understand the limits of MQV-style implicit authentication in the lattice-based setting.

Beyond fully authenticated frameworks, the vulnerability of key reuse extends critically to basic unauthenticated key exchange protocols. These protocols, natively designed only for passive security settings (i.e., IND-CPA), may also be expected to remain robust against key reuse attacks [20,15,17] due to efficiency considerations. To circumvent heavy generic transformations, Gao et al. [22] proposed GDLL-KE, which injects ad-hoc session noise to achieve claimed key-reuse robustness without a formal proof. By design, this randomized perturbation is intended to smear out the sharp algebraic patterns leveraged in state-of-the-art attacks, causing existing attacks by Bindel et al. [6] and Qin et al. [41] to fail. Consequently, the resilience of GDLL-KE against broader statistical or geometric vulnerabilities remains entirely unexplored.

The persistent lack of a rigorous eCK analysis for ZZDSD-AKE, together with the unverified robustness of GDLL-KE, leaves a significant gap in our understanding of reconciliation-based key exchange in practice. This naturally poses a critical open question: *Are these reconciliation-based constructions inherently secure in strong security models, or do they still harbor unforeseen vulnerabilities?*

1.1 Our Contributions

In this work, we answer the above question negatively by presenting the first eCK attack against ZZDSD-AKE and the first complete key-recovery attack against GDLL-KE. Our results demonstrate that the signal function can introduce previously overlooked leakage, ultimately enabling full secret-key recovery. Our main contributions are summarized as follows:

- We propose a *novel* signal leakage attack on ZZDSD-AKE, resolving its open question in the negative. Unlike prior work that relies on algebraic structures to extract per-coefficient information, our attack exploits a geometric bias induced by the signal function, which is further amplified by ephemeral secret leakage in the eCK model. This bias actually implies a variance bias along the secret direction, enabling us to recover an approximate secret through spectral methods. Combined with a parity-recovery phase that corrects the remaining ambiguity, we obtain the exact static secret key. We implement the attack against all parameter sets of ZZDSD-AKE and successfully recover the secret key in 99–100% of trials with roughly 1,500–1,700 queries.
- Beyond the attack on ZZDSD-AKE, we further illustrate the generality of our spectral methodology by applying it to the GDLL-KE protocol. Despite GDLL-KE’s different design and explicit randomization countermeasure, its

leakage reduces to the same form as in ZZDSD-AKE. Our spectral method recovers the reused secret key in about 180 queries without any ephemeral leakage. This completely breaks the claimed key-reuse security of GDLL-KE, and demonstrates that our attack methodology is not confined to the eCK model or MQV-style constructions.

1.2 Related Work

Lattice-based AKEs. Various generic constructions for lattice-based AKEs have been proposed based on key encapsulation mechanisms (KEMs) [21,49,37]. These KEMs are typically built from lattice-based public-key encryption. Beyond generic constructions, KEMs can also serve as drop-in replacements for the DH key exchange in some classical AKE protocols, thereby enabling post-quantum variants. For example, this approach has been pursued in real-world AKE protocols, including TLS [44], Signal [26], FIDO2 [5], the Noise framework [2], and WireGuard [27]. Recent works have further improved the efficiency of post-quantum TLS [55,12]. Parallel to the KEM-based paradigm, MQV-style constructions [53,16] aim for compact implicit authentication without generic transformations.

Attacks on lattice-based key exchange. The cryptanalysis of (Ring-)LWE-based key exchange under key reuse was initiated by Fluhrer [20]. Subsequently, Ding et al. [15] introduced the signal leakage attack from a critical observation that the signal function leaks information about Bob’s reused secret key. Bindel et al. [6] improved the attack via sparse signal collection and extended it to AKEs. Subsequent works [41,32] achieved further improvements. Unfortunately, these attacks are inapplicable to ZZDSD-AKE and GDLL-KE. Another type of attacks against (Ring-)LWE-based key exchange is the key mismatch attack [17], which recovers Alice’s secret under key reuse by checking whether the shared keys of both parties match. A series of works [39,40,25,45] has studied key mismatch attacks against lattice-based KEMs. Moreover, Gong and Zhao [24] proposed a key-mismatch-like attack against the one-pass variant of ZZDSD-AKE, but it does not work against the two-pass ZZDSD-AKE. Finally, the eCK setting is practically well-motivated, since ephemeral key exposures are highly realistic due to modern side-channel vulnerabilities, including EM/power analysis [42,54], cache attacks [11], and CPU frequency leakage [48,51].

Spectral methods in side-channel attacks. More broadly, the spectral decomposition approach has been independently applied in side-channel attacks against lattice-based signatures [54], where the covariance of leakage-filtered signatures reveals the secret direction and ultimately enables key recovery.

2 Preliminaries

2.1 Notation

We use bold lowercase letters for vectors and bold uppercase letters for matrices, and vectors are in row form. We denote $\mathbf{b}_i$ as the i -th row of a matrix $\mathbf{B}$. The

inner product of two vectors $\mathbf{u}$ and $\mathbf{v}$ is denoted by $\langle \mathbf{u}, \mathbf{v} \rangle$. The Euclidean ℓ_2 norm (resp. ℓ_∞ norm) of $\mathbf{v}$ is denoted by $\|\mathbf{v}\|$ (resp. $\|\mathbf{v}\|_\infty$). For $x \in \mathbb{R}$, $\lfloor x \rfloor$ is the greatest integer not exceeding x , and $\lceil x \rceil$ denotes rounding to the nearest integer. For a prime $q > 2$, let $\mathcal{R}_q = \mathbb{Z}_q[x]/(x^n+1)$ be the polynomial ring, where $\mathbb{Z}_q = \{-\frac{q-1}{2}, \dots, \frac{q-1}{2}\}$. A polynomial $a = \sum_{i=0}^{n-1} a_i x^i \in \mathcal{R}_q$ can be represented as a vector $\mathbf{a} = (a_0, a_1, \dots, a_{n-1})$, and $a[i]$ (resp. $\mathbf{a}[i]$) denotes the i -th coefficient (resp. entry) of a (resp. $\mathbf{a}$). Let χ_σ denote the discrete Gaussian distribution over $\mathcal{R}_q$ with standard deviation σ . Sampling from χ_σ is written as $a \leftarrow \chi_\sigma$, where each $a[i]$ follows the discrete Gaussian distribution $D_{\mathbb{Z}, \sigma}$ over $\mathbb{Z}$ with standard deviation σ . For $a \in \mathcal{R}_q$, define the associated negacyclic matrix $\text{Rot}(a) \in \mathbb{Z}_q^{n \times n}$ by

$$\text{Rot}(a) = \begin{pmatrix} a_0 & -a_{n-1} & \cdots & -a_1 \\ a_1 & a_0 & \cdots & -a_2 \\ \vdots & \vdots & \ddots & \vdots \\ a_{n-1} & a_{n-2} & \cdots & a_0 \end{pmatrix}.$$

For $a, b \in \mathcal{R}_q$, the transpose of the coefficient vector of $a \cdot b$ equals $\text{Rot}(a) \cdot \mathbf{b}^\top$.

The Signal Function. We introduce some notations of the signal function Sig and robust extractor Mod_2 used in ZZDSD-AKE and GDLL-KE. For a prime $q > 2$, let $E = \{-\lfloor \frac{q}{4} \rfloor, \dots, \lceil \frac{q}{4} \rceil\}$. Given $x \in \mathbb{Z}_q$, the signal function Sig is defined as

$$\text{Sig}(x) = \begin{cases} 0, & x \in E, \\ 1, & x \notin E. \end{cases} \quad (1)$$

We define $\text{Mod}_2 : \mathbb{Z}_q \times \{0, 1\} \rightarrow \{0, 1\}$ as

$$\text{Mod}_2(x, w) = \left(x + \frac{q-1}{2} \cdot w \bmod q \right) \bmod 2.$$

For $a \in \mathcal{R}_q$, we define $\text{Sig}(a) = (\text{Sig}(a[0]), \dots, \text{Sig}(a[n-1]))$ and extend Mod_2 to $\mathcal{R}_q$ coefficient-wise in the same way.

2.2 Zhang et al.'s AKE Protocol

The Ring-LWE-based AKE protocol, referred to as ZZDSD-AKE, was proposed by Zhang, Zhang, Ding, Snook, and Dagdelen [53] at EUROCRYPT 2015 and is illustrated in Figure 1. In this protocol, $H_1 : \{0, 1\}^* \rightarrow \chi_\alpha$ is a hash function which outputs an invertible element sampled from χ_α , and $H_2 : \{0, 1\}^* \rightarrow \{0, 1\}^l$ is a key derivation function that outputs an l -bit session key.

The ZZDSD-AKE can be viewed as a Ring-LWE variant of HMQV [29], and it is proven to be secure in the BR security model [4] with weak forward secrecy. Alice and Bob hold static (i.e., long-term) key pairs (p_A, s_A) and (p_B, s_B) , respectively. In each session, Alice and Bob generate ephemeral key pairs (x_A, r_A) and (y_B, r_B) , and derive the session key from both the static and ephemeral keys, respectively. Zhang et al. adopt a similar idea to HMQV to combine static




Fig. 1. ZZDSD-AKE protocol [53].

and ephemeral secret keys by computing the hash $d = H_1(ID_B, ID_A, y_B, x_A)$ and $\hat{r}_B = d \cdot s_B + r_B$. The term $\hat{r}_B$ can be viewed as a signature of the transcript (ID_B, ID_A, y_B, x_A) . Actually, Zhang et al. apply *noise flooding* and *rejection sampling*⁵, both standard in lattice-based signatures, to statistically ensure that the static key s_A (or s_B) is hidden by the ephemeral secret r_A (or r_B). As a result, two discrete Gaussian distributions χ_α and χ_β are used for generating static keys and ephemeral keys, respectively, where β is much larger than α . The specific parameters of ZZDSD-AKE are presented in Table 1, where the prime q is required to satisfy $q \equiv 1 \pmod{2n}$.

The ZZDSD-AKE protocol naturally supports reuse of static keys across sessions, and it is proven secure against active attacks in the BR model. Zhang et al. noted that ZZDSD-AKE cannot be proven secure in the stronger CK model [10], which allows the adversary to reveal the session state K_B or K_A . The eCK model [30], which is also stronger than BR and even permits leakage of ephemeral secrets r_A or r_B , remains unresolved for ZZDSD-AKE, and no security guarantees are known in this model either. Previous signal leakage attacks proposed in [6, 41, 32] fail against ZZDSD-AKE, even in the eCK model. These attacks rely on coefficient-level extraction methods, whereas multiplication by d mixes the coefficients of s_B , preventing such extraction even when r_B is revealed (as elaborated in Section 3). As a result, the eCK security of ZZDSD-AKE is still not well understood. In this work, we show that it is vulnerable to signal leakage attacks in the eCK model.

2.3 Gao et al.’s Reusable-Key Protocol

The GDLL-KE protocol, proposed by Gao, Ding, Li, and Liu [22], is also based on R-LWE and designed to be secure with key reuse. Figure 2 illustrates this pro-

⁵ For simplicity, we omit the details of rejection sampling in the description of ZZDSD-AKE, as it does not affect our attacks.




Fig. 2. GDLL-KE protocol [22].

Table 1. Parameters of ZZDSD-AKE and GDLL-KE.

Protocol	Parameter set	n	α	β	q
ZZDSD-AKE	I ₁	1024	3.397	$2^{16.1}$	2^{45}
	I ₂	1024	3.397	$2^{17.1}$	2^{47}
	II ₁	2048	3.397	$2^{17.1}$	2^{47}
	II ₂	2048	3.397	$2^{18.7}$	2^{50}
GDLL-KE	—	1024	3.192	—	2^{30}

tol. Gao et al. employ a randomized countermeasure to prevent signal leakage attacks. Specifically, Bob generates a noise term e_p in each session to randomize the shared key K_B and the signal $\mathbf{w}_B$. They claim that this random noise can prevent the signal leakage attack in [15], without providing a formal proof. However, subsequent signal leakage attacks presented in [6,41,32] do not apply to GDLL-KE, since multiplication by the fresh masking polynomial e_p similarly mixes the coefficients of s_B , causing their coefficient-level recovery to fail. Table 1 summarizes the parameters of GDLL-KE.

2.4 Adversarial Model and Goals

We consider an active adversary $\mathcal{A}$ that acts as a malicious initiator (i.e., Alice) and repeatedly interacts with Bob. We define a query as one interaction with Bob, in which the adversary sends a chosen public key x_A and receives a response from Bob. In each query, $\mathcal{A}$ is able to arbitrarily select x_A (in both ZZDSD-AKE and GDLL-KE) and observes the corresponding signal $\mathbf{w}_B$ output by Bob. In both protocols, the adversary mounts adaptive chosen-input interactions, and the ultimate goal is to recover Bob’s static (or reused) secret key s_B from the leakage of $\mathbf{w}_B$.

Ephemeral key leakage in the eCK model [30]. The eCK model provides an EphemeralKeyReveal query, allowing the adversary to obtain the ephemeral secret key of a party in a given session. This is permitted provided that the same party’s

Experiment $\text{KR}_A()$ 1: Generate public parameters pp and static key pairs $(p_A, s_A), (p_B, s_B)$ honestly 2: $s'_B \leftarrow \mathcal{A}^{\mathcal{O}_B}(\text{pp}, p_A, s_A, p_B)$ 3: return 1 if $s'_B = s_B$, and 0 otherwise
Oracle $\mathcal{O}_B(x_A)$ 1: Start a new responder session π_B at Bob by sending x_A 2: Receive the response $(y_B, \mathbf{w}_B)$ 3: Obtain r_B through $\text{EphemeralKeyReveal}(\pi_B)$ 4: return $(y_B, \mathbf{w}_B, r_B)$

Fig. 3. Adversarial capabilities in our key recovery attack against ZZDSD-AKE.

static key is not compromised in the **Test** session⁶. This models the possible leakage of ephemeral secrets due to side-channel attacks or insecure randomness generation.

Adversarial capabilities against ZZDSD-AKE. In ZZDSD-AKE, the adversary acts as a malicious initiator possessing an honestly generated static key pair (p_A, s_A) , and may choose the ephemeral public key x_A arbitrarily in each query. The only extra capability required by our attack is to obtain Bob’s ephemeral secret r_B . This capability is explicitly captured by the eCK model, in which our attack against ZZDSD-AKE is carried out. For each query the adversary is additionally allowed to issue an $\text{EphemeralKeyReveal}$ query to reveal Bob’s ephemeral secret r_B . Figure 3 formalizes the capabilities used in our attack. We remark that the adversarial capabilities considered in our attack go beyond those covered by the original BR security proof of ZZDSD-AKE, and may be achievable in realistic attack scenarios.

After recovering s_B via our attack, $\mathcal{A}$ can win the eCK security game as follows. The adversary $\mathcal{A}$ impersonates Bob to another honest initiator U , and uses the recovered s_B and freshly sampled randomness to compute the response. After U completes the session, $\mathcal{A}$ selects this session as the **Test** session. Using s_B and its own randomness, $\mathcal{A}$ computes the same session key as U , and thus distinguishes the real session key from a random key returned by **Test**. The **Test** session is fresh: no reveal query is made on it or on either party’s static key. No matching session exists, since Bob is impersonated by $\mathcal{A}$.

Adversarial capabilities against GDLL-KE. In GDLL-KE, the adversary $\mathcal{A}$ can choose a public key x_A arbitrarily. Since GDLL-KE is explicitly intended to support key reuse under active attacks, such a malicious initiator is already covered by the claimed threat setting. Unlike the attack on ZZDSD-AKE, however, our attack against GDLL-KE requires *no ephemeral-secret leakage*. The adversary strictly relies on the publicly transmitted signal $\mathbf{w}_B$ and the masking term e_p to

⁶ The **Test** session is a fresh session for which the adversary is challenged to distinguish its session key from a random key in the security experiment.




Fig. 4. Geometric intuition of the signal-conditioned spectral recovery visualized in two dimensions. (a) The unconstrained row vectors $\mathbf{d}_i$ form the initial isotropic distribution, concealing the secret key direction. (b) The signal $\mathbf{w}_B$ acts as a filter, retaining only the row vectors bounded within a thin orthogonal slab (blue points). (c) The filtering suppresses the variance of the retained samples along the secret direction, as illustrated by the purple ellipse. The unit eigenvector $\mathbf{u}$ (blue arrow) corresponding to the smallest eigenvalue of the retained sample covariance matrix closely aligns with the hidden secret direction.

reconstruct the reused secret key s_B . That is, the oracle $\mathcal{O}_B(x_A)$ returns only $(e_p, \mathbf{w}_B)$.

3 Attack Idea

This section explains the intuition behind our attacks. Before delving into the detailed algebraic formulations, we first highlight the core geometric concept underlying our attack. As illustrated in Figure 4, our key insight is to shift from algebraic coefficient-level analysis to a high-dimensional geometric perspective. The signal $\mathbf{w}_B[i] = 0$ acts as a geometric filter that preferentially retains sample vectors with small projections onto the secret key $\mathbf{s}_B$ as shown in Figure 4(b). This selection bias suppresses the variance of the retained samples, specifically along the secret direction. Consequently, the secret direction can be recovered by identifying the direction of minimum variance via spectral decomposition of the sample covariance matrix. We show below how the adversary exploits this geometric bias by leveraging the leakage permitted in the eCK model.

We start from the useful decomposition of Bob’s approximate shared key K_B introduced by Bindel et al. [6]. Specifically, from $p_A = as_A + 2e_A$ and $p_B = as_B + 2e_B$, we obtain

$$\begin{aligned} K_B &= (p_A c + x_A)(ds_B + r_B) + 2cg_B \\ &= x_A \cdot d \cdot s_B + \Delta + \varepsilon, \end{aligned} \tag{2}$$

where $\Delta = dcp_B s_A + cp_A r_B + x_A r_B$ and $\varepsilon = 2cg_B + 2dc(e_A s_B - e_B s_A)$.




Fig. 5. Empirical distribution of all coefficients of ds_B for the parameter set Π_1 , over 10,000 samples of d (std: $522.22 \approx \alpha^2 \sqrt{n}$). Bars are colored according to the signal $\mathbf{w}_B[i]$: blue for $\mathbf{w}_B[i] = 0$ and beige for $\mathbf{w}_B[i] = 1$. Hatched bars around $\pm q/(4k)$ indicate the transition regions, where $\mathbf{w}_B[i]$ is ambiguous due to ε (i.e., can be either 0 or 1).

The adversary knows (p_B, p_A, s_A, x_A) and is able to compute the hash values c and d . The obstacle is Δ , which depends on Bob's ephemeral secret r_B and, without r_B , is indistinguishable from a uniform ring element. In the eCK model, the adversary can obtain r_B using `EphemeralKeyReveal` and then compute Δ exactly. Then, the adversary is able to eliminate the effect of the known Δ by retaining those signals $\mathbf{w}_B[i] = \text{Sig}(K_B[i])$ with small $\Delta[i]$. To illustrate our core idea, we defer the detailed treatment of Δ to Section 4.1 and simplify the expression of K_B to

$$K_B^* = x_A \cdot d \cdot s_B + \varepsilon. \quad (3)$$

Since the noise term ε is small relative to $x_A ds_B$, the signal $\mathbf{w}_B = \text{Sig}(K_B^*)$ is governed primarily by the term $x_A ds_B$.

Note that removing Δ , however, does not directly reveal s_B from the signal $\mathbf{w}_B$. A fundamental barrier stands in the way. The ring multiplication between d and s_B mixes all coefficients of s_B , so each signal $\mathbf{w}_B[i]$ depends on all coefficients of s_B , preventing per-coefficient extraction. This is why prior attacks [6, 41, 32] that rely on coefficient-level extraction methods fail. Instead, we use a geometric perspective to analyze the information of $d \cdot s_B$ carried by the signal $\mathbf{w}_B$, thereby overcoming the limitation of coefficient-level extraction methods.

Signal-Conditioned Bias. The adversary chooses x_A as a constant polynomial $x_A = k > 0$. According to Equation (3), each coefficient satisfies

$$K_B^*[i] = k \cdot (ds_B)[i] + \varepsilon[i].$$

The distribution of ds_B is characterized by Lemma 1. Then, by Lemma 1, each coefficient of ds_B is approximately Gaussian with standard deviation $\alpha^2\sqrt{n}$. The adversary sets k such that most values of $k \cdot (ds_B)[i]$ lie within $(-q/2, q/2)$ to avoid modular wrap-around. Recall from Equation (1) that $\text{Sig}(x) = 0$ if x lies in the central interval E . Thus, the signal $\mathbf{w}_B[i] = 0$ is equivalent to

$$|k \cdot (ds_B)[i] + \varepsilon[i]| < q/4,$$

which simplifies to $|(ds_B)[i]| < \frac{q}{4k}$ approximately. A suitable choice of k makes $\mathbf{w}_B[i] = 0$ a reliable indicator that the value of $|(ds_B)[i]|$ is smaller than $q/(4k)$, as illustrated in Figure 5. Since $|\varepsilon[i]| \ll q$, values of $|(ds_B)[i]|$ slightly above this threshold may still yield $\mathbf{w}_B[i] = 0$, while values slightly below it may be flipped to $\mathbf{w}_B[i] = 1$. We call such regions *the transition regions* (hatched bars in Figure 5).

Lemma 1 ([6]). *For $a, b \leftarrow \chi_\sigma$, let $c = a \cdot b \in \mathcal{R}_q$. Then for each coefficient $c[i]$ ($0 \leq i < n$), we have $\mathbb{E}[c[i]] = 0$ and $\text{Var}(c[i]) = n\sigma^4$. When n is sufficiently large, $c[i]$ approximately follows a Gaussian distribution with mean 0 and standard deviation $\sigma^2\sqrt{n}$ (by the central limit theorem).*

Now write $d = d_0 + d_1x + \dots + d_{n-1}x^{n-1}$ and $\mathbf{s}_B = (s_B[0], \dots, s_B[n-1])$. Multiplication by d in $\mathcal{R}_q$ corresponds to the negacyclic matrix

$$\mathbf{Rot}(d) = \begin{pmatrix} d_0 & -d_{n-1} & \dots & -d_1 \\ d_1 & d_0 & \dots & -d_2 \\ \vdots & \vdots & \ddots & \vdots \\ d_{n-1} & d_{n-2} & \dots & d_0 \end{pmatrix} \in \mathbb{Z}_q^{n \times n}. \quad (4)$$

Hence $(ds_B)[i] = \langle \mathbf{d}_i, \mathbf{s}_B \rangle$, where $\mathbf{d}_i$ is the i -th row of $\mathbf{Rot}(d)$, and

$$\mathbf{w}_B[i] = \text{Sig}(K_B^*[i]) = \text{Sig}(k \cdot \langle \mathbf{d}_i, \mathbf{s}_B \rangle + \varepsilon[i]).$$

The signal $\mathbf{w}_B[i] = 0$ holds when $|\langle \mathbf{d}_i, \mathbf{s}_B \rangle| < b$, where $b \approx \frac{q}{4k}$, and b becomes sufficiently small given a sufficiently large k . Geometrically, it means that these row vectors $\mathbf{d}_i$ are biased toward the slab $\{\mathbf{d} \in \mathbb{R}^n \mid |\langle \mathbf{d}, \mathbf{s}_B \rangle| < b\}$, which forms a thin neighborhood of the hyperplane orthogonal to $\mathbf{s}_B$, as visualized in Figure 4(b). For comparison, Figure 4(a) illustrates the isotropic Gaussian distribution of $\mathbf{d}_i$ in the absence of the signal filter. After multiple queries, the adversary collects many row vectors $\mathbf{d}_i$ corresponding to $\mathbf{w}_B[i] = 0$. Let the collected vectors be the set $\mathcal{D}$ defined as

$$\mathcal{D} = \{\mathbf{d}_i \mid \mathbf{d}_i \text{ is the } i\text{-th row of } \mathbf{Rot}(d) \text{ and } \mathbf{w}_B[i] = 0\}.$$

As illustrated in Figure 5, the inner products $\langle \mathbf{d}_i, \mathbf{s}_B \rangle$ for $\mathbf{d}_i \in \mathcal{D}$ follow a centrally truncated Gaussian distribution. In other words, vectors with large projections onto $\mathbf{s}_B$ are suppressed, while those with small projections are retained. This induced selection bias is precisely the leakage exploited by our attack.

Spectral Recovery of the Secret Direction. Since the bias is symmetric with respect to $\pm \mathbf{s}_B$, the first moment (e.g., the mean) carries little information. Instead, the second moment (e.g., the variance) reveals the secret direction. Concretely, arrange vectors in $\mathcal{D}$ as rows of a matrix $\mathbf{D} \in \mathbb{R}^{M \times n}$, where $M = |\mathcal{D}|$, and form the (unnormalized) covariance matrix

$$\mathbf{\Sigma}_M = \mathbf{D}^\top \mathbf{D}. \quad (5)$$

Let the spectral decomposition of $\mathbf{\Sigma}_M$ be

$$\mathbf{\Sigma}_M = \mathbf{Q} \begin{bmatrix} \lambda_1 & & & \\ & \lambda_2 & & \\ & & \ddots & \\ & & & \lambda_n \end{bmatrix} \mathbf{Q}^\top, \quad (6)$$

where $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$. By filtering with $\mathbf{w}_B[i] = 0$, all collected row vectors $\mathbf{d}_i$ admit small projections onto the secret key $\mathbf{s}_B$. As a result, the variance of their projections onto the direction of $\mathbf{s}_B$ is suppressed, while the variance in every orthogonal direction remains essentially unchanged. Figure 4(c) illustrates this phenomenon: the direction orthogonal to $\mathbf{s}_B$ (along the major axis of the purple ellipse) exhibits larger variance, while the direction aligned with $\mathbf{s}_B$ has smaller variance. Consequently, $\mathbf{s}_B$ coincides with the minimum-variance direction in $\mathbb{R}^n$.

The following lemma shows that the direction of minimum variance is exactly the eigenvector associated with the smallest eigenvalue of the covariance matrix.

Lemma 2 (Minimum-variance direction [47]). *Let X be a random vector in $\mathbb{R}^n$ with covariance matrix $\mathbf{\Sigma}$, and let the eigenvalues of $\mathbf{\Sigma}$ satisfy $0 \leq \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$ with corresponding unit eigenvectors $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_n$. Then, over all unit vectors $\mathbf{v} \in \mathbb{R}^n$,*

$$\lambda_1 = \min_{\|\mathbf{v}\|=1} \text{Var}(\langle X, \mathbf{v} \rangle),$$

and the minimum value is attained at $\mathbf{u}_1$.

Applying Lemma 2 to the covariance matrix $\mathbf{\Sigma}_M$, this minimum variance is attained at the unit eigenvector $\mathbf{u} \in \mathbb{R}^n$ corresponding to the smallest eigenvalue λ_1 of $\mathbf{\Sigma}_M$. Therefore, $\mathbf{u}$ closely aligns with the normalized secret direction, satisfying

$$\mathbf{u} \approx \pm \frac{\mathbf{s}_B}{\|\mathbf{s}_B\|}. \quad (7)$$

To convert $\mathbf{u}$ into an approximate secret vector $\mathbf{s}'$, we must estimate the Euclidean norm $\|\mathbf{s}_B\|$. Recall that coefficients of $\mathbf{s}_B$ are sampled independently from a discrete Gaussian distribution with standard deviation α . We exploit the norm concentration behavior of high-dimensional random vectors, presented as the following lemma.

Lemma 3 ([35]). Let $\mathbf{v} = (v_1, v_2, \dots, v_n) \in \mathbb{R}^n$ be a random vector whose entries are independent and identically distributed (i.i.d.) random variables with zero mean and standard deviation σ . Then we have

$$\mathbb{E}[\|\mathbf{v}\|] \leq \sigma\sqrt{n},$$

and asymptotically $\mathbb{E}[\|\mathbf{v}\|] \sim \sigma\sqrt{n}$ as $n \rightarrow \infty$.

Applying Lemma 3, we estimate that $\|\mathbf{s}_B\| \approx \alpha\sqrt{n}$. Hence the adversary can obtain an approximate secret

$$\mathbf{s}' = \alpha\sqrt{n} \cdot \mathbf{u}, \quad (8)$$

which is very close to $\mathbf{s}_B$ or $-\mathbf{s}_B$. When the number M of collected vectors is large enough, $\|\mathbf{s}' - \mathbf{s}_B\|_\infty$ or $\|\mathbf{s}' + \mathbf{s}_B\|_\infty$ is less than $1/2$. The adversary can refine $\mathbf{s}'$ into the exact secret key by a simple rounding. We give a theoretical analysis of the effectiveness of the spectral method in Section 4.1, which is similar to that in [54]. In fact, using the parity recovery and post-processing described in Section 4.2, $\|\mathbf{s}' \pm \mathbf{s}_B\|_\infty < 1$ is enough to recover $\mathbf{s}_B$, which is easier to achieve in practice.

Remark 1. Conditioning on $\mathbf{w}_B[i] = 1$ (the outer signal zone) instead of 0 also leaks the secret direction. In that case the retained vectors with large $|\langle \mathbf{d}_i, \mathbf{s}_B \rangle|$ are preferentially aligned with $\pm \mathbf{s}_B$, amplifying the variance in the secret direction. As a result, $\mathbf{s}_B$ is then revealed by the eigenvector corresponding to the largest eigenvalue λ_n of Σ_M . However, as Figure 5 shows, the event $\mathbf{w}_B[i] = 1$ is much rarer than $\mathbf{w}_B[i] = 0$, since it corresponds to the tails of the distribution of $(ds_B)[i]$. An attack conditioning on $\mathbf{w}_B[i] = 1$ would require significantly more queries to collect enough vectors for a robust spectral recovery. Thus, we condition on $\mathbf{w}_B[i] = 0$ in our attacks.

Extension to GDLL-KE. The above analysis extends to GDLL-KE almost verbatim. In GDLL-KE, Bob's approximate shared key $K_B = x_A e_p s_B + \varepsilon'$, where $\varepsilon' = 2e_p^2 + 2e'_B$ and $e_p \leftarrow \chi_\alpha$ is freshly sampled in each session. In ZZSD-AKE, the adversary first neutralizes the extra term Δ through the eCK leakage and then works with $x_A ds_B$. However, in GDLL-KE the situation is simpler, since the signal leakage is directly from the term $x_A e_p s_B$.

Similarly, choosing $x_A = k$ as a suitable constant again makes the signal $\mathbf{w}_B[i] = 0$ an approximate test for whether $(e_p s_B)[i]$ is close to 0. Consequently, the corresponding row vectors of $\mathbf{Rot}(e_p)$ for which $\mathbf{w}_B[i] = 0$ are biased toward vectors whose absolute inner products with $\mathbf{s}_B$ are small, causing them to concentrate near the hyperplane orthogonal to $\mathbf{s}_B$. Therefore, the same spectral method applies, and the eigenvector corresponding to the smallest eigenvalue of the sample covariance matrix again approximates $\pm \mathbf{s}_B / \|\mathbf{s}_B\|$.

In short, although GDLL-KE is designed to support key reuse and is claimed to resist signal leakage attacks, our attack shows that the random masking term e_p does not prevent the signal from leaking information about s_B .

4 eCK Key Recovery Attack against ZZDSD-AKE

In this section, we describe our full attack against ZZDSD-AKE in detail. The attack consists of two main phases. First, the adversary collects row vectors $\mathbf{d}_i$ for which $\mathbf{w}_B[i] = 0$, and recovers an approximate secret $\mathbf{s}'$ using the spectral method. Second, the adversary retrieves the parity (i.e., least significant bit) of each $\mathbf{s}_B[i]$, and combines it with the approximate secret $\mathbf{s}'$ to reconstruct $\mathbf{s}_B$.

4.1 Approximate Secret Recovery via Spectral Decomposition

Parameter Selection. The adversary sets the ephemeral public key x_A as a constant polynomial $x_A = k \in \mathbb{Z}_q$ with $k > 0$. Recall from Section 3 that the signal $\mathbf{w}_B[i] = 0$ is approximately equivalent to $|(ds_B)[i]| < q/(4k)$, so the choice of k determines the selection window for $(ds_B)[i]$. A larger k narrows the selection window and enhances the variance suppression in the secret direction, which spectral recovery exploits. However, if k is excessively large, $k \cdot (ds_B)[i]$ exceeds $\pm q/2$ for a large $(ds_B)[i]$, leading to modular wrap-around. This breaks the signal-dependent bias: $\mathbf{w}_B[i] = 0$ no longer indicates $|(ds_B)[i]| < q/(4k)$, as large values of $|(ds_B)[i]|$ can also yield $\mathbf{w}_B[i] = 0$ after wrap-around.

We therefore choose k such that $q/(4k) \approx \alpha^2 \sqrt{n}$. For example, for parameter set Π_1 with $\alpha^2 \sqrt{n} \approx 522$, we set $k = \lfloor q/(4 \cdot 530) \rfloor$. By this setting, the selection window aligns with one standard deviation of $(ds_B)[i]$. With this setting, since $(ds_B)[i]$ approximately follows a Gaussian distribution with standard deviation $\alpha^2 \sqrt{n}$, we have

$$\Pr(|k \cdot (ds_B)[i]| < q/2) = \Pr(|(ds_B)[i]| < 2\alpha^2 \sqrt{n}) \approx 0.9545$$

by standard properties of the Gaussian distribution. In other words, more than 95% of the coefficients satisfy $|k \cdot (ds_B)[i]| < q/2$ and avoid modular wrap-around. The remaining few cases, where $|(ds_B)[i]| > 2\alpha^2 \sqrt{n}$, cause the modular wrap-around errors. In practice, the spectral recovery can tolerate such a small number of errors with a sufficiently large sample size.

Signal-Conditioned Collection. In each query, the adversary (acting as Alice) sends $x_A = k$ to Bob and receives $(y_B, \mathbf{w}_B)$. The adversary then computes $c = H_1(ID_A, ID_B, x_A)$, $d = H_1(ID_B, ID_A, y_B, x_A)$, and using the revealed ephemeral secret r_B , evaluates

$$\Delta = d \cdot c \cdot p_B \cdot s_A + c \cdot p_A \cdot r_B + x_A \cdot r_B.$$

Note that the ephemeral secret r_B is obtained by an `EphemeralKeyReveal` query defined in the eCK model. Following the heuristic analyses in [6,41], the products $dcp_B s_A$ and $cp_A r_B$ in Δ can be heuristically modeled as pseudorandom elements of $\mathcal{R}_q$, motivated by the fresh r_B and hash d . Hence, even when $x_A = k$ is fixed, the sum Δ is heuristically approximated by a uniform distribution over $\mathcal{R}_q$, and in particular over $\mathbb{Z}_q$ coefficient-wise.

For each $i \in \{0, \dots, n-1\}$, $\mathbf{d}_i$ is the i -th row of the negacyclic matrix $\mathbf{Rot}(d)$. According to Equation (2), the signal $\mathbf{w}_B[i]$ satisfies

$$\mathbf{w}_B[i] = \text{Sig}(k \cdot \langle \mathbf{d}_i, \mathbf{s}_B \rangle + \Delta[i] + \varepsilon[i]). \quad (9)$$

Although $\Delta[i]$ is known to the adversary, a large $\Delta[i]$ affects the signal $\mathbf{w}_B[i]$ regardless of $|\langle \mathbf{d}_i, \mathbf{s}_B \rangle|$. Hence, $\Delta[i]$ enlarges the transition regions (hatched bars) in Figure 5 and renders $\mathbf{w}_B[i] = 0$ a less reliable indicator of small $|\langle \mathbf{d}_i, \mathbf{s}_B \rangle|$. To mitigate this effect, the adversary collects $\mathbf{d}_i$ if and only if

1. $\mathbf{w}_B[i] = 0$, and
2. $|\Delta[i]| < \delta_1$, where $\delta_1 = \lfloor q/16 \rfloor$.

The first condition realizes the signal-conditioned selection described in Section 3. The second condition filters out cases where the known $\Delta[i]$ is too large, and ensures that $\mathbf{w}_B[i] = 0$ reliably implies $|\langle \mathbf{d}_i, \mathbf{s}_B \rangle| < \frac{q/4+\delta_1}{k}$, or $|\langle \mathbf{d}_i, \mathbf{s}_B \rangle| < 1.25\alpha^2\sqrt{n}$ with $k = q/(4\alpha^2\sqrt{n})$.

The choice $\delta_1 = \lfloor q/16 \rfloor$ offers a practical trade-off between the quality of spectral recovery and the efficiency of our attack. A smaller threshold yields higher-quality spectral recovery, since the known $\Delta[i]$ perturbs the signal less, but it also reduces the number of collected rows per query. However, a larger threshold would increase the yield but broaden the transition regions and weaken the bias induced by $\mathbf{w}_B[i] = 0$. With $\delta_1 = \lfloor q/16 \rfloor$, the probability that $|\Delta[i]| < \delta_1$ is about $2\delta_1/q = 1/8$ under the above heuristic. Thus, it retains a constant fraction of the rows for which $\mathbf{w}_B[i] = 0$, striking a balance between the two extremes. We omit ε in the above discussion, as its standard deviation $\sigma_\varepsilon \approx \sqrt{4n\alpha^2\beta^2 + 8n^3\alpha^8}$ is much smaller than $q/16$. For instance, for the Π_1 parameters, $\sigma_\varepsilon \approx 2^{25}$ while $q/16 \approx 2^{43}$.

The adversary repeats the above process until M row vectors are collected. Let $\mathcal{D} = (\mathbf{d}^{(1)}, \dots, \mathbf{d}^{(M)})$ denote the resulting set of collected vectors, where each $\mathbf{d}^{(j)}$ corresponds to a row vector $\mathbf{d}_i$ satisfying $\mathbf{w}_B[i] = 0$ and $|\Delta[i]| < \delta_1$. By construction, $\mathcal{D}$ is a signal-conditioned sample set, corresponding to the central truncation of the (approximate Gaussian) distribution of $\langle \mathbf{d}_i, \mathbf{s}_B \rangle$, i.e., samples with relatively small $|\langle \mathbf{d}_i, \mathbf{s}_B \rangle|$ are retained. This truncation removes large projections onto $\mathbf{s}_B$, thereby significantly reducing the variance along the secret direction. The inherent bias is exactly what our subsequent steps require.

Spectral Recovery. The adversary arranges all vectors in $\mathcal{D}$ as the rows of a matrix $\mathbf{D} \in \mathbb{R}^{M \times n}$, and forms the unnormalized covariance matrix $\mathbf{\Sigma}_M \in \mathbb{R}^{n \times n}$ according to Equation (5). Then, the adversary performs the spectral decomposition of $\mathbf{\Sigma}_M$ to obtain

$$\mathbf{\Sigma}_M = \mathbf{Q}\mathbf{\Lambda}\mathbf{Q}^\top,$$

where $\mathbf{\Lambda} = \text{diag}(\lambda_1, \dots, \lambda_n)$ with $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$. As discussed in Section 3, the unit eigenvector $\mathbf{u} \in \mathbb{R}^n$ corresponding to the smallest eigenvalue λ_1 is an approximation of $\pm \mathbf{s}_B / \|\mathbf{s}_B\|$. The adversary does not know the Euclidean norm

of $\mathbf{s}_B$, but it can be well estimated from Lemma 3 as $\|\mathbf{s}_B\| \approx \alpha\sqrt{n}$ in practice. Therefore, by scaling $\mathbf{u}$ with $\alpha\sqrt{n}$, the adversary obtains an initial estimate $\mathbf{s}' = \alpha\sqrt{n} \cdot \mathbf{u}$.

To assess the effectiveness of spectral recovery, the norm $\|\mathbf{s}' - \mathbf{s}_B\|_\infty$ is particularly suitable, as it measures the maximum deviation across all coefficients. If every coefficient error is strictly smaller than $1/2$, then simple coefficient-wise rounding recovers the exact coefficients of $\mathbf{s}_B$ ⁷. Inspired by the sample complexity analysis of spectral recovery in [54], we give an approximate theoretical estimate of $\|\mathbf{s}' - \mathbf{s}_B\|_\infty$ in the following.

We assume that the sign of $\mathbf{s}'$ is correct. To simplify the theoretical analysis, we idealize the retained rows as independent samples from a subgaussian distribution.⁸ Let $\mathbf{u}_B = \mathbf{s}_B/\|\mathbf{s}_B\|$ be the normalized unit direction. Let $\Sigma = \mathbb{E}[\mathbf{d}_i^\top \mathbf{d}_i \mid \mathbf{w}_B[i] = 0, |\Delta[i]| < \delta_1]$ be the population covariance matrix where the unit eigenvector corresponding to its smallest eigenvalue is $\mathbf{u}_B$, and $\Sigma'_M = \frac{1}{M} \mathbf{D}^\top \mathbf{D}$ be the sample covariance matrix measured over M samples. Note that scaling by $1/M$ does not change the eigenvectors. For the eigenvalues $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$ of Σ , there exists a gap $\delta = \lambda_2 - \lambda_1$ separating the secret direction from the orthogonal subspace.

Lemma 4 (Davis-Kahan Inequality [47]). *Let $\Sigma, \Sigma'_M \in \mathbb{R}^{n \times n}$ be symmetric matrices. Given their respective unit eigenvectors corresponding to the smallest eigenvalues, $\mathbf{u}_B$ and $\mathbf{u}$, we have*

$$\|\mathbf{u}_B - \theta \mathbf{u}\| \leq \frac{2\sqrt{2}\|\Sigma - \Sigma'_M\|}{\delta},$$

where $\theta \in \{-1, 1\}$.

Lemma 5 (Following Remark 4.7.3 in [47]). *Let Σ'_M be the sample covariance matrix of M independent subgaussian random vectors in $\mathbb{R}^n$ with population covariance Σ . There exist $K > 1$ and C_0 such that for any $t \geq 0$, we have*

$$\|\Sigma'_M - \Sigma\| \leq C_0 K^2 \lambda_n \left(\sqrt{\frac{n+t}{M}} + \frac{n+t}{M} \right),$$

with probability at least $1 - 2e^{-t}$.

Lemma 6. *Under the idealized sampling and covariance assumptions above, let $\mathbf{s}' = \alpha\sqrt{n} \cdot \mathbf{u}$ be the unrounded spectral estimate of the secret key. For a sufficiently large sample size $M > n$, there exists a constant $C_1 > 0$ such that with high probability, the ℓ_2 error of the estimate is bounded by*

$$\|\mathbf{s}' - \mathbf{s}_B\| \leq C_1 \frac{\alpha \lambda_n n}{\delta \sqrt{M}} + |\alpha\sqrt{n} - \|\mathbf{s}_B\||.$$

⁷ The sign of $\mathbf{s}'$ can be easily determined. After rounding, $\mathbf{s}'$ is equivalent to either $\mathbf{s}_B$ or $-\mathbf{s}_B$. We can recover the true sign by selecting the candidate with smaller coefficients of $p_B - a\mathbf{s}'$.

⁸ Rows within each matrix $\mathbf{Rot}(d)$ are correlated, but rows across different queries are independent as d is fresh in each session.




Fig. 6. $\|s' - s_B\|_\infty$ over 50 spectral recovery runs for the parameter set Π_1 .

Proof. By the triangle inequality, we have

$$\|s' - s_B\| = \|\alpha\sqrt{n}\mathbf{u} - \|s_B\|\mathbf{u}_B\| \leq \alpha\sqrt{n}\|\mathbf{u} - \mathbf{u}_B\| + |\alpha\sqrt{n} - \|s_B\||.$$

Combining Lemma 4 and Lemma 5,

$$\|\mathbf{u} - \mathbf{u}_B\| \leq \frac{2\sqrt{2}C_0K^2\lambda_n}{\delta} \left(\sqrt{\frac{n+t}{M}} + \frac{n+t}{M} \right).$$

Taking $t = n$, the above bound on $\|\mathbf{u} - \mathbf{u}_B\|$ holds with probability at least $1 - 2e^{-n}$. Since $M > n$, we have $\sqrt{(n+t)/M} + (n+t)/M \leq (\sqrt{2} + 2)\sqrt{n/M}$. Multiplying by $\alpha\sqrt{n}$ and absorbing the constants into C_1 , we obtain the desired bound.

Under the above idealized setting, allowing nonzero $\Delta[i]$ with $|\Delta[i]| < \delta_1$ relaxes the conditioning on $|\langle \mathbf{d}_i, \mathbf{s}_B \rangle|$ compared with the $\Delta[i] = 0$ case. This affects the secret-direction eigenvalue and the constants δ and C_1 in Lemma 6, but preserves the $1/\sqrt{M}$ dependence of the sampling-error term. We now give a heuristic analysis of $\|s' - s_B\|_\infty$ based on Lemma 6. Concretely, assuming that all entries of $\mathbf{e} = s' - s_B$ independently follow a Gaussian distribution with standard deviation σ , then we have $\mathbb{E}[\|\mathbf{e}\|] \approx \sigma\sqrt{n}$ by Lemma 3 and $\mathbb{E}[\|\mathbf{e}\|_\infty] \approx \sigma\sqrt{2\ln n}$ (see Remark 2.7.7 in [47]) for sufficiently large n . Consequently,

$$\mathbb{E}[\|\mathbf{e}\|_\infty] \approx \sqrt{\frac{2\ln n}{n}} \mathbb{E}[\|\mathbf{e}\|].$$

By Lemma 3, we have $|\alpha\sqrt{n} - \|s_B\|| = O(1)$ in expectation. Substituting this heuristic scaling relation into the ℓ_2 bound in Lemma 6 yields the heuristic estimate for the expected ℓ_∞ error as

$$\mathbb{E}[\|s' - s_B\|_\infty] = \mathcal{O} \left(\frac{\alpha\lambda_n\sqrt{n\ln n}}{\delta\sqrt{M}} + \sqrt{\frac{\ln n}{n}} \right). \quad (10)$$




Fig. 7. The number h of coefficients with $|\mathbf{s}'[i] - \mathbf{s}_B[i]| > 0.5$ over 50 spectral recovery runs for the parameter set Π_1 .

Experimental results. This theoretical estimate aligns well with our experimental results in Figure 6, which illustrates the evolution of $\|\mathbf{s}' - \mathbf{s}_B\|_\infty$ as the number M of collected vectors increases for the parameter set Π_1 under the filter $|\Delta[i]| < \delta_1$. It is clear that the error decays proportionally to $1/\sqrt{M}$ rather than linearly in M . As a consequence, driving the maximum coefficient error below the rounding threshold (i.e., $\|\mathbf{s}' - \mathbf{s}_B\|_\infty < 0.5$) requires a substantial increase in M . For example, reducing the error from 1.0 to 0.5 requires roughly a fourfold increase in the sample size. This clarifies why achieving exact coefficient recovery purely through spectral decomposition demands a prohibitively large M (e.g., $M > 10^6$ for the parameter set Π_1).

On the other hand, if very few coefficients of $\mathbf{s}' - \mathbf{s}_B$ deviate by $1/2$, the adversary can still correct them by enumerating. We use h to denote the number of coefficients of $\mathbf{s}'$ satisfying $|\mathbf{s}'[i] - \mathbf{s}_B[i]| > 1/2$ for $i \in \{0, \dots, n-1\}$. Figure 7 illustrates the evolution of this metric with M . Crucially, h decays much faster than $\|\mathbf{s}' - \mathbf{s}_B\|_\infty$. For the parameter set Π_1 , it drops precipitously from over 300 at $M = 150,000$ to near zero by $M = 950,000$. This observation confirms that a moderate M is sufficient for an adversary to recover the secret key through efficient enumeration.

Query Complexity. The expected number of queries required to accumulate M row vectors is determined by two filtering conditions, parameterized by k and δ_1 .

By setting $k \approx q/(4\alpha^2\sqrt{n})$, the condition for small $|(ds_B)[i]|$ to fall into the central interval E translates to $|(ds_B)[i]| < \alpha^2\sqrt{n}$. Since this threshold corresponds exactly to one standard deviation of $(ds_B)[i]$, the probability of observing the signal $\mathbf{w}_B[i] = 0$ is approximately 0.68 by standard Gaussian properties.

Simultaneously, the noise filter $|\Delta[i]| < \delta_1$ with $\delta_1 = \lfloor q/16 \rfloor$ accepts a coefficient with probability approximately $2\delta_1/q = 1/8 = 0.125$. In fact, $\Delta[i]$ and $(ds_B)[i]$ are correlated through the ephemeral polynomial d . To obtain an ide-

alized complexity estimate, we heuristically treat them as independent random variables. The probability of satisfying both conditions is then estimated as

$$\Pr(\mathbf{w}_B[i] = 0 \text{ and } |\Delta[i]| < \delta_1) \approx 0.68 \times 0.125 = 0.085.$$

Since each query provides n candidate row vectors, the adversary expects to collect approximately $0.085n$ valid row vectors per query. Consequently, given a target sample size M , the overall query complexity of the spectral recovery phase is roughly estimated to be $M/(0.085n)$.

4.2 Parity Recovery and Post-Processing

In this subsection, we consider refining an approximate secret $\mathbf{s}' = \alpha\sqrt{n} \cdot \mathbf{u}$ recovered by a moderate M into the exact secret key. First, we propose a parity recovery method to extract the parity of each coefficient of $\mathbf{s}_B$. It ensures that the adversary can correct errors by rounding to the nearest integer of matching parity when $\|\mathbf{s}' - \mathbf{s}_B\|_\infty < 1$. Moreover, we introduce a post-processing method that combines the outputs of spectral recovery and parity recovery with an improved norm estimate of $\mathbf{s}_B$ and a small-scale enumeration, thereby improving the accuracy and efficiency of the attack.

Parity Recovery. The adversary acts as before, except that the ephemeral public key is set to $x_A = (q-1)/2$, and tries to extract the parity of each $\mathbf{s}_B[i]$ based on the below fact.

Lemma 7. *For any prime $q > 5$ and any $v \in \mathbb{Z}$ with $|v| < q/4$, it holds that*

$$\text{Sig}\left(\frac{q-1}{2} \cdot v \bmod q\right) \equiv v \pmod{2}.$$

Proof. Let $x = \frac{q-1}{2} \cdot v \bmod q$. We analyze x by considering the parity of v . We first consider the case that v is even. Let $v = 2y$ for some integer y . Then we have

$$x \equiv \frac{q-1}{2} \cdot 2y = y(q-1) \equiv -y = -\frac{v}{2} \pmod{q}.$$

Since $|v|$ is small relative to q (e.g., $|v| < q/4$), the value $x \equiv -v/2$ lies strictly within the central interval $E = \{-\lfloor \frac{q}{4} \rfloor, \dots, \lceil \frac{q}{4} \rceil\}$. By the definition of the signal function in Equation (1), $\text{Sig}(x) = 0$. Thus, $\text{Sig}(x) \equiv v \pmod{2}$ holds.

Then, we deal with the case that v is odd. Let $v = 2y + 1$ for some integer y . Then we have

$$x \equiv \frac{q-1}{2}(2y+1) = y(q-1) + \frac{q-1}{2} \equiv -y + \frac{q-1}{2} \pmod{q}.$$

Substituting $y = \frac{v-1}{2}$, we obtain:

$$x \equiv -\frac{v-1}{2} + \frac{q-1}{2} = \frac{q-v}{2} \pmod{q}.$$

Since $|v|$ is small, the value $\frac{q-v}{2}$ is clustered closely around $\frac{q}{2}$. Specifically, for $|v| < q/4$, it is easy to verify that $\frac{q-v}{2}$ strictly exceeds $\lceil \frac{q}{4} \rceil$ and is significantly far from the interval E . Consequently, $x \notin E$, which yields $\text{Sig}(x) = 1$. When v is odd, $\text{Sig}(x) \equiv v \pmod{2}$ also holds. In both cases, the signal function correctly extracts the parity of v , concluding the proof.

Recall that in Equation (9), the signal $\mathbf{w}_B[i]$ is generated from the noisy term as

$$\mathbf{w}_B[i] = \text{Sig} \left(\frac{q-1}{2} \cdot \langle \mathbf{d}_i, \mathbf{s}_B \rangle + \Delta[i] + \varepsilon[i] \bmod q \right). \quad (11)$$

Let $v = \langle \mathbf{d}_i, \mathbf{s}_B \rangle$. Since the coefficients of $\mathbf{s}_B$ and $\mathbf{d}_i$ are Gaussian, v has a standard deviation of approximately $\alpha^2 \sqrt{n}$. With $\alpha = 3.397$ and $n \leq 2048$, the magnitude $|v|$ is typically just a few hundred. Compared to the modulus boundaries (e.g., $q/4 \approx 2^{43}$ for parameter set I₁), the value $|v|/2$ is exceptionally negligible.

According to the proof of Lemma 7, the unperturbed base value $\frac{q-1}{2} \cdot v \bmod q$ is tightly centered around 0 (if v is even) or $q/2$ (if v is odd). To avoid flipping the final output signal $\mathbf{w}_B[i]$, the cumulative perturbation $\Delta[i] + \varepsilon[i]$ must not push the value across the boundaries of the interval E , which are located at a distance of $q/4$ from the respective centers. Since $|v|/2$ and $\varepsilon[i]$ are too small, the available error tolerance for $\Delta[i]$ is overwhelmingly large (nearly $q/4$).

Crucially, $\Delta[i]$ is fully computable by the adversary in the eCK model. By imposing a filter such as $|\Delta[i]| < \delta_2$ with $\delta_2 = \lfloor 7q/32 \rfloor$, the adversary can eliminate the risk of signal corruption by selectively collecting only those signal bits where $\Delta[i]$ is safely bounded away from the $q/4$ threshold. The choice of $\delta_2 = \lfloor 7q/32 \rfloor$ reserves a $q/32$ gap to tolerate the small error $\varepsilon[i]$. As a result, for all collected $\mathbf{w}_B[i]$ and $\mathbf{d}_i$, the condition of Lemma 7 holds despite the noise, guaranteeing that

$$\mathbf{w}_B[i] \equiv \langle \mathbf{d}_i \bmod 2, \mathbf{s}_B \bmod 2 \rangle \pmod{2}.$$

In other words, each retained pair $(\mathbf{d}_i, \mathbf{w}_B[i])$ yields a linear equation over $\mathbb{F}_2$ in the unknown parity bits of $\mathbf{s}_B$. Note that, unlike the spectral recovery, the adversary retains the coefficients regardless of the value of $\mathbf{w}_B[i]$, since both 0 and 1 are informative for the linear system.

Suppose that after a few queries, the adversary has collected M' valid pairs satisfying the filtering condition $|\Delta[i]| < \delta_2$, where $M' > n$. Let

$$\mathcal{P} = \{(\bar{\mathbf{d}}_j, w_j)\}_{j=1}^{M'}$$

denote this collection of retained pairs, where each $\bar{\mathbf{d}}_j \in \mathbb{F}_2^n$ is the collected $\mathbf{d}_i \bmod 2$, and $w_j \in \mathbb{F}_2$ is its corresponding signal $\mathbf{w}_B[i]$. The adversary randomly selects a subset of n pairs from $\mathcal{P}$. Without loss of generality, let this subset be indexed as $\{(\bar{\mathbf{d}}_k, w_k)\}_{k=1}^n$. The adversary then forms the matrix $\mathbf{A} \in \mathbb{F}_2^{n \times n}$ whose k -th row is $\bar{\mathbf{d}}_k$, and the column vector $\mathbf{b} = (w_1, \dots, w_n)^\top \in \mathbb{F}_2^n$. If $\mathbf{A}$ has full rank over $\mathbb{F}_2$, the adversary solves the linear system $\mathbf{A}\mathbf{x}^\top = \mathbf{b}$ over $\mathbb{F}_2$ to

directly obtain the parity row vector $\mathbf{x} = \mathbf{s}_B \bmod 2 \in \mathbb{F}_2^n$. If $\mathbf{A}$ is rank-deficient, the adversary simply draws another random subset of n pairs from $\mathcal{P}$ and retries.

To successfully recover the target secret, the adversary needs to construct a full-rank $n \times n$ matrix $\mathbf{A}$ by selecting n vectors from the gathered pool $\mathcal{P}$. To analyze the feasibility of this step and bound the expected number of trials, we first recall a classical result concerning the rank distribution of random binary matrices.

Lemma 8 ([52]). *Let $\mathbf{A} = (a(i, j))$ be an $n \times n$ random binary matrix whose entries are independent and satisfy $\Pr(a(i, j) = 0) = \Pr(a(i, j) = 1) = \frac{1}{2}$. Then*

$$\Pr(\text{rank}(\mathbf{A}) = n) = \prod_{i=1}^n \left(1 - \frac{1}{2^i}\right).$$

Although the reduced vectors in $\mathcal{P}$ are not strictly uniformly random, they are approximately uniform over $\mathbb{F}_2^n$ in practice under the concrete parameters n and α . Empirically, over 5000 collected vectors of dimension 2048, the overall proportion of 1s is 0.5032, while the per-vector proportions range from 0.4956 to 0.5171 with standard deviation 0.0077, which is highly consistent with the uniform distribution over $\mathbb{F}_2$. Based on Lemma 8 and using a heuristic estimate, a randomly selected binary matrix $\mathbf{A} \in \mathbb{F}_2^{n \times n}$ is full rank with probability

$$\prod_{i=1}^n (1 - 2^{-i}) \approx 0.288,$$

when n is sufficiently large (e.g., $n \geq 1024$). Thus, a small number of retries (typically 3–5) suffices to find a full-rank $\mathbf{A}$ with reasonable probability. If a full-rank matrix still cannot be found after these retries, the adversary can issue a few additional queries to expand the pool $\mathcal{P}$ and repeat the process, incurring negligible overhead to the overall attack. Concretely, under the uniformity heuristic for $\Delta[i]$, the filter $|\Delta[i]| < \delta_2$ accepts each coefficient with probability approximately $2\delta_2/q \approx 7/16$, yielding about $7n/16$ valid pairs per query. Thus, collecting $M' > n$ valid pairs requires approximately $16M'/(7n)$ queries. For example, taking a reasonable pool size $M' = 1.5n$ requires only about 3.43 queries.

By obtaining $\mathbf{s}_B \bmod 2$, the adversary can refine the approximate secret $\mathbf{s}'$ by rounding each coordinate to the nearest integer that respects the known parity. Specifically, for each i , the adversary chooses $\lfloor \mathbf{s}'[i] \rfloor$ if $\lfloor \mathbf{s}'[i] \rfloor \equiv \mathbf{s}_B[i] \pmod{2}$, and $\lceil \mathbf{s}'[i] \rceil$ otherwise.

Experimental results. Figure 8 presents the average error $\|\mathbf{s}' - \mathbf{s}_B\|_\infty$ over 50 independent trials for the parameter set Π_1 , adopting spectral recovery and parity correction under moderate values of M . Compared with pure spectral recovery, the required M for $\|\mathbf{s}' - \mathbf{s}_B\|_\infty$ to approach 0 is reduced by nearly one order of magnitude after introducing parity correction, i.e., from 10^6 to 10^5 . The average h is shown in Figure 9. As can be observed, while driving the error




Fig. 8. Average of $\|\mathbf{s}' - \mathbf{s}_B\|_\infty$ over 50 runs (spectral recovery and parity correction per run) for the parameter set Π_1 .




Fig. 9. Average of the number h of coefficients with $|\mathbf{s}'[i] - \mathbf{s}_B[i]| > 0.5$ over 50 runs (spectral recovery and parity correction per run) for the parameter set Π_1 .

$\|\mathbf{s}' - \mathbf{s}_B\|_\infty$ close to 0 demands $M = 350,000$, the adversary can further relax the required M for key recovery to 250,000, since an enumeration-based error correction is highly affordable when $h \leq 2$. The experimental results also show that the average number of queries needed to recover the parity is 3.66 over 50 parity recovery runs, which closely matches the above heuristic estimate.

Post-Processing. As shown in Equation (10), the heuristic ℓ_∞ error bound contains a tail term bounded by $\mathcal{O}(\sqrt{(\ln n)/n})$. This term arises directly from the norm estimation error, i.e., approximating the true key length $\|\mathbf{s}_B\|$ with its statistical expectation $\alpha\sqrt{n}$. When the sample size M is relatively small, this norm estimation gap noticeably impacts the coefficient-wise accuracy of $\mathbf{s}'$.

To mitigate the impact of this tail term and better approximate the true scale of $\mathbf{s}_B$, the adversary performs a guided search for the scale factor C . Specifically,

Algorithm 1 Post-Processing for Full Key Recovery

Input: Unit eigenvector $\mathbf{u} \in \mathbb{R}^n$, parity vector $\mathbf{x} \in \mathbb{F}_2^n$ **Output:** Static secret key $\mathbf{s}_B$ or **Failure**

```
1: for  $C \in \{\alpha\sqrt{n} - 10, \alpha\sqrt{n} - 9.9, \dots, \alpha\sqrt{n} + 10\}$  do  $\triangleright$  Search over scale parameter
2:   for  $\text{sign} \in \{1, -1\}$  do  $\triangleright$  Resolve sign ambiguity
3:      $\mathbf{s}' \leftarrow \text{sign} \cdot C \cdot \mathbf{u}$ 
4:      $\mathbf{s}'_B \leftarrow \text{ParityCorrection}(\mathbf{x}, \mathbf{s}')$   $\triangleright$  Correct coefficients using parity
5:     if  $(p_B - a \cdot \mathbf{s}'_B)/2$  is consistent with  $\chi_\alpha$  then  $\triangleright$  Verify candidate key
6:       return  $\mathbf{s}'_B$ 
7:     end if
8:   for  $i \in \{0, \dots, n-1\}$  do  $\triangleright$  Perform single-coefficient enumeration
9:     for  $p \in \{2, -2\}$  do
10:       $\mathbf{s}''_B \leftarrow \mathbf{s}'_B$ 
11:       $\mathbf{s}''_B[i] \leftarrow \mathbf{s}'_B[i] + p$ 
12:      if  $(p_B - a \cdot \mathbf{s}''_B)/2$  is consistent with  $\chi_\alpha$  then
13:        return  $\mathbf{s}''_B$ 
14:      end if
15:    end for
16:  end for
17: end for
18: end for
19: return Failure
```

the adversary enumerates candidate scales C within the interval $[\alpha\sqrt{n}-P, \alpha\sqrt{n}+P]$ with a step size of 0.1. In practice, setting $P = 10$ is sufficient to cover the true norm. For each candidate scale C , the adversary considers both possible directional signs, generating the continuous estimates $\mathbf{s}' = C \cdot \mathbf{u} \in \mathbb{R}^n$ and $-\mathbf{s}'$. For each candidate continuous vector, the adversary applies the parity correction using the recovered parity vector to correct each coefficient, as discussed above. The resulting candidate integer vector is then verified against Bob's static public key $p_B = a \cdot \mathbf{s}_B + 2e_B$.

Despite the above refinements, a very small number of coefficients may still remain incorrect after the parity correction. In practice, an additional single-coefficient enumeration is enough to recover the final secret. Specifically, for each index $i \in \{0, \dots, n-1\}$, the adversary tests the two candidates obtained by adjusting the i -th coefficient by ± 2 , and verifies the resulting key against p_B . Once a consistent candidate is found, the exact static secret key $\mathbf{s}_B$ is recovered. The complete post-processing procedure is summarized in Algorithm 1.

5 Key Recovery Attack against GDLL-KE

In this section, we present a signal leakage attack against GDLL-KE that fully recovers Bob's reused secret key, thereby invalidating the claimed key-reuse security. The overall attack follows the same two-phase approach as the attack on ZZDSD-AKE, but admits two important simplifications. First, the term Δ does not appear in GDLL-KE, thus the filtering conditions required in Section 4.1 are

unnecessary. Second, the role of the hash output d in ZZDSD-AKE is effectively replaced by the random masking term e_p , which is freshly sampled by Bob in each session and publicly transmitted together with the signal $\mathbf{w}_B$.

In GDLL-KE, the adversary does not need to compromise any ephemeral secrets. The adversary simply impersonates Alice by choosing an arbitrary public key x_A and sending it to Bob. The responder Bob, holding the reused secret key s_B , computes

$$K_B = x_A \cdot s_B \cdot e_p + 2e_p^2 + 2e'_B, \quad (12)$$

where $e_p, e'_B \leftarrow \chi_\alpha$ are freshly sampled in each session, and outputs $\mathbf{w}_B = \text{Sig}(K_B)$ together with e_p .

Approximate Secret Recovery. The adversary sets $x_A = k \in \mathbb{Z}_q$ as a constant polynomial with $k \approx q/(4\alpha^2\sqrt{n})$, following the same parameter selection as in the eCK attack against ZZDSD-AKE in Section 4.1. Each coefficient of K_B satisfies

$$K_B[i] = k \cdot (e_p s_B)[i] + \varepsilon[i],$$

where $\varepsilon = 2e_p^2 + 2e'_B$ is small relative to q . Then each signal bit is

$$\mathbf{w}_B[i] = \text{Sig}(K_B[i]) = \text{Sig}(k \cdot (e_p s_B)[i] + \varepsilon[i]).$$

As in the ZZDSD-AKE case, $\mathbf{w}_B[i] = 0$ occurs approximately when $|(e_p s_B)[i]| < q/(4k)$. Write the ring multiplication as $(e_p s_B)[i] = \langle \mathbf{r}_i, \mathbf{s}_B \rangle$, where $\mathbf{r}_i$ is the i -th row of the negacyclic matrix $\mathbf{Rot}(e_p)$. The adversary collects all rows $\mathbf{r}_i$ for which $\mathbf{w}_B[i] = 0$. Since GDLL-KE has no Δ term, no further filtering is needed. The adversary repeats this process over multiple queries until M row vectors are accumulated.

The adversary arranges the M vectors as rows of $\mathbf{D} \in \mathbb{R}^{M \times n}$ and forms the unnormalized covariance matrix $\Sigma_M = \mathbf{D}^\top \mathbf{D}$. After computing the spectral decomposition $\Sigma_M = \mathbf{Q} \Lambda \mathbf{Q}^\top$ with $\lambda_1 \leq \dots \leq \lambda_n$, the eigenvector $\mathbf{u} \in \mathbb{R}^n$ corresponding to the smallest eigenvalue λ_1 approximates $\pm \mathbf{s}_B / \|\mathbf{s}_B\|$, as established in Section 3. Using the norm estimate $\|\mathbf{s}_B\| \approx \alpha\sqrt{n}$ by Lemma 3, the adversary obtains the approximate secret $\mathbf{s}' = \alpha\sqrt{n} \cdot \mathbf{u}$. Following the theoretical analysis in Section 4.1, the population covariance is conditioned only on $\mathbf{w}_B[i] = 0$, corresponding to the $\Delta[i] = 0$ case. Consequently, the ℓ_∞ error $\|\mathbf{s}' - \mathbf{s}_B\|_\infty$ of the spectral estimate also decays as $1/\sqrt{M}$. Similar to the analysis in Section 4.1, we estimate the query complexity of spectral recovery without filtering Δ to be $M/(0.68n)$ for a target sample size M .

Parity Recovery. The adversary switches the public key to $x_A = \frac{q-1}{2}$. Then each signal bit $\mathbf{w}_B[i]$ satisfies

$$\mathbf{w}_B[i] = \text{Sig}\left(\frac{q-1}{2} \cdot \langle \mathbf{r}_i, \mathbf{s}_B \rangle + \varepsilon[i]\right).$$

The noise term $\varepsilon[i] = 2e_p^2[i] + 2e'_B[i]$ is even and its magnitude is far below $q/4$ for the parameters under consideration. Hence it does not affect the signal

function output, and Lemma 7 applies directly, giving

$$\mathbf{w}_B[i] \equiv \langle \bar{\mathbf{r}}_i, \bar{\mathbf{s}}_B \rangle \pmod{2},$$

where $\bar{\mathbf{r}}_i = \mathbf{r}_i \bmod 2$ and $\bar{\mathbf{s}}_B = \mathbf{s}_B \bmod 2$. Since GDLL-KE has no Δ term, no noise filter is required.

Due to the fact that no coefficients need to be filtered out, every single signal bit contributes a valid linear equation over $\mathbb{F}_2$. The adversary can thus readily assemble the binary linear system and solve for the parity vector $\bar{\mathbf{s}}_B$ using the exact same procedure detailed in Section 4.2. Notably, this absence of a noise filter drastically improves the query efficiency, as a single query may be sufficient to construct a full-rank matrix over $\mathbb{F}_2$. Concretely, for $M' = 1.5n$, the cost of parity recovery is estimated to be 2 queries, which suffices to find a full-rank matrix over $\mathbb{F}_2$ with reasonable probability based on Lemma 8.

Post-Processing. Given the spectral direction estimate $\mathbf{u}$ and the parity vector $\bar{\mathbf{s}}_B$, the adversary recovers the exact secret key $\mathbf{s}_B$ through the same post-processing procedure described for ZZDSD-AKE in Section 4.2 and summarized in Algorithm 1.

Remark on the Random Masking Countermeasure. Gao et al. [22] introduced the random noise term e_p as a countermeasure against signal leakage attacks, arguing that the randomization of K_B and $\mathbf{w}_B$ in each session prevents the adversary from accumulating useful statistical information. However, our attack reveals that relying solely on the randomness of individual signal values is fundamentally insufficient. While e_p does randomize the signal values in each session, it fails to eliminate secret-dependent correlations in the joint distribution of $(e_p, \mathbf{w}_B)$. In particular, conditioning on $\mathbf{w}_B[i] = 0$ favors rows of $\mathbf{Rot}(e_p)$ with small projections onto $\mathbf{s}_B$. These signal-conditioned rows exhibit variance suppression along the secret direction, similar to the rows retained in ZZDSD-AKE. This conditional bias can be accumulated and exploited by the spectral method despite the fresh masking in each session.

6 Experiments

We implement our attacks against ZZDSD-AKE and GDLL-KE in SageMath 10.8 and Python 3.14. All experiments are carried out on a desktop with an AMD Ryzen 5 3600 CPU (12 logical cores) and 16 GB RAM. In our implementation of the attack against ZZDSD-AKE, the signal-conditioned collection phase is parallelized across 10 worker processes. The source code is available at https://github.com/frostry/eCK_attack.

6.1 Attack against ZZDSD-AKE

We evaluate the attack on all four parameter sets of ZZDSD-AKE listed in Table 1, spanning $n \in \{1024, 2048\}$, $\log_2 \beta \in \{16.1, 17.1, 18.7\}$ and $\log_2 q \in$

Table 2. Success rate, query complexity, and running time of the attack against ZZDSD-AKE on parameter sets I_1 , I_2 , Π_1 , and Π_2 , over 100 trials per setting. Query counts and running times are reported as mean $\pm$ standard deviation.

Parameter set	M	Success rate	Queries	Time (s)
I_1	110,000	90/100	1260.99 ± 18.08	11.08 ± 1.37
	130,000	100/100	1483.74 ± 19.37	12.52 ± 1.12
I_2	110,000	90/100	1259.24 ± 16.80	11.55 ± 1.84
	130,000	99/100	1487.64 ± 20.98	12.70 ± 1.34
Π_1	250,000	94/100	1429.04 ± 15.35	40.01 ± 6.58
	290,000	100/100	1655.05 ± 16.12	46.14 ± 5.98
Π_2	250,000	97/100	1429.18 ± 15.84	47.99 ± 8.12
	290,000	100/100	1657.09 ± 17.67	53.37 ± 5.02

$\{45, 47, 50\}$. Here, the modulus q is chosen to satisfy $q \equiv 1 \pmod{2n}$, as in [53]. Following the analysis in Section 4.1, the adversary sets $x_A = \lfloor q/(4k') \rfloor$ in the spectral phase with threshold $\delta_1 = \lfloor q/16 \rfloor$, where $k' \approx \alpha^2 \sqrt{n}$ (i.e., $k' = 375$ for $n = 1024$ and $k' = 530$ for $n = 2048$). In the parity recovery phase, $x_A = (q-1)/2$ with threshold $\delta_2 = \lfloor 7q/32 \rfloor$.

For each parameter set we test two different choices of the sample size M and run 100 independent trials per setting. A trial is counted as successful if the post-processing in Algorithm 1 outputs $\mathbf{s}_B$ exactly. The results are summarized in Table 2. At the smaller sample size the attack already succeeds in roughly 90–97% of trials, and increasing M by 20,000 or 40,000 samples raises the success rate to 99–100% across all four parameter sets, at the cost of roughly 220–230 additional queries. For $n = 1024$ each full attack completes in about 12 seconds, and for $n = 2048$ the running time grows to roughly 50 seconds, dominated by the larger collection size. The few failed runs at smaller values of M correspond to cases where at least two coefficients of $\mathbf{s}'$ still differ from the corresponding coefficients of $\mathbf{s}_B$ by ± 2 after parity correction. The average number of queries, dominated by the spectral phase, is close to the estimate $M/(0.085n)$ from Section 4.1. For example, the average number of queries is 1655.05, which closely matches $M/(0.085n) = 1665.9$ for $M = 290,000$ under parameter set Π_1 .

6.2 Attack against GDLL-KE

For GDLL-KE, we use the parameter set in Table 1 ($n = 1024$, $\alpha = 3.192$, $q \approx 2^{30}$). We set $x_A = \lfloor q/(4 \cdot 335.462) \rfloor$ for the approximate secret recovery and $x_A = (q-1)/2$ for the parity recovery. Since GDLL-KE has no Δ term, no $|\Delta[i]|$ filter is applied, and every coefficient with $\mathbf{w}_B[i] = 0$ is retained.

We test at $M = 110,000$ and $M = 130,000$, again over 100 independent trials. The results are reported in Table 3. At $M = 110,000$ the attack succeeds in 99/100 trials with an average of 156.44 queries, and at $M = 130,000$ the success

Table 3. Success rate, query complexity, and running time of the attack against GDLL-KE, over 100 trials. Query counts and running times are reported as mean $\pm$ standard deviation.

M	Success rate	Queries	Time (s)
110,000	99/100	156.44 ± 3.17	8.66 ± 0.56
130,000	100/100	184.23 ± 3.01	10.19 ± 0.47

rate reaches 100% with 184.23 queries on average. Each full attack completes in about 9–10 seconds on our setup.

Compared with ZZDSD-AKE at the same dimension $n = 1024$, the absence of the Δ filter reduces the average number of required queries from 10^3 to 10^2 , which represents a substantial reduction for the same sample size M .

7 Conclusion

We presented a new signal leakage attack against Ring-LWE-based key exchange protocols whose static or reused keys had been considered safe, either through MQV-style designs or through ad-hoc randomization countermeasures. Our first attack answers the open question on the eCK security of ZZDSD-AKE in the negative. By leveraging the ephemeral key compromise allowed in the eCK model, an active adversary recovers the static secret key $\mathbf{s}_B$ in roughly 1,500–1,700 queries across all four parameter sets. Note that our attack against ZZDSD-AKE does not invalidate its claimed BR security, since the attack relies on capabilities that lie outside the BR model. Our second attack breaks the claimed key-reuse security of GDLL-KE with about 180 queries, showing that injecting per-session noise as a stand-alone countermeasure does not defeat signal leakage exploitation.

Our results demonstrate that neither MQV-style implicit authentication nor additional noise randomization alone suffices to protect static (or reused) secret keys in practical settings. In particular, our attacks show that security analyses restricted to coefficient-wise leakage may overlook vulnerabilities arising from the global geometric structure of the leakage distribution. We therefore recommend that future lattice-based AKE designs either adopt formally proven generic transformations, such as Fujisaki–Okamoto-style constructions, or provide rigorous guarantees against geometric leakage attacks in advanced models such as eCK. It remains open whether a lightweight modification of ZZDSD-AKE or GDLL-KE can provably resist our attacks while retaining the efficiency of AKE protocols directly built from the LWE problem.

Acknowledgments. We sincerely thank the anonymous reviewers for their valuable comments. We used ChatGPT solely for spell checking and basic language editing. This work was supported by the National Natural Science Foundation of China (Grant Nos. 62332007, U22B2028, 62672136, and 62372445), the Shenzhen Major Science and Technology Special Project (No. KJZD20240903101108011), and the Guangdong Hong Kong Joint Laboratory for Data Security and Privacy Protection.

References

1. Alkim, E., Ducas, L., Pöppelmann, T., Schwabe, P.: Post-quantum key exchange: a new hope. In: Proceedings of the 25th USENIX Conference on Security Symposium. p. 327–343. SEC’16, USENIX Association, USA (2016)
2. Angel, Y., Dowling, B., Hülsing, A., Schwabe, P., Weber, F.: Post quantum noise. In: Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security. p. 97–109. CCS ’22, Association for Computing Machinery, New York, NY, USA (2022)
3. Barthe, G., Crespo, J.M., Lakhnech, Y., Schmidt, B.: Mind the gap: Modular machine-checked proofs of one-round key exchange protocols. In: Oswald, E., Fischlin, M. (eds.) Advances in Cryptology - EUROCRYPT 2015. pp. 689–718. Springer Berlin Heidelberg, Berlin, Heidelberg (2015)
4. Bellare, M., Rogaway, P.: Entity authentication and key distribution. In: Stinson, D.R. (ed.) Advances in Cryptology — CRYPTO’ 93. pp. 232–249. Springer Berlin Heidelberg, Berlin, Heidelberg (1994)
5. Bindel, N., Cremers, C., Zhao, M.: Fido2, ctap 2.1, and webauthn 2: Provable security and post-quantum instantiation. In: 2023 IEEE Symposium on Security and Privacy (SP). pp. 1471–1490 (2023). <https://doi.org/10.1109/SP46215.2023.10179454>
6. Bindel, N., Stebila, D., Veitch, S.: Improved attacks against key reuse in learning with errors key exchange. In: Longa, P., Ràfols, C. (eds.) Progress in Cryptology – LATINCRYPT 2021. pp. 168–188. Springer International Publishing, Cham (2021)
7. Bos, J., Costello, C., Ducas, L., Mironov, I., Naehrig, M., Nikolaenko, V., Raghunathan, A., Stebila, D.: Frodo: Take off the ring! practical, quantum-secure key exchange from lwe. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security. pp. 1006–1018. CCS ’16, Association for Computing Machinery, New York, NY, USA (2016)
8. Bos, J., Ducas, L., Kiltz, E., Lepoint, T., Lyubashevsky, V., Schanck, J.M., Schwabe, P., Seiler, G., Stehle, D.: Crystals - kyber: A cca-secure module-lattice-based kem. In: 2018 IEEE European Symposium on Security and Privacy (EuroS&P). pp. 353–367 (2018)
9. Bos, J.W., Costello, C., Naehrig, M., Stebila, D.: Post-quantum key exchange for the tls protocol from the ring learning with errors problem. In: 2015 IEEE Symposium on Security and Privacy. pp. 553–570 (2015). <https://doi.org/10.1109/SP.2015.40>
10. Canetti, R., Krawczyk, H.: Analysis of key-exchange protocols and their use for building secure channels. In: Pfitzmann, B. (ed.) Advances in Cryptology — EUROCRYPT 2001. pp. 453–474. Springer Berlin Heidelberg, Berlin, Heidelberg (2001)
11. Chen, B., Wang, Y., Shome, P., Fletcher, C., Kohlbrenner, D., Paccagnella, R., Genkin, D.: GoFetch: Breaking Constant-Time cryptographic implementations using data Memory-Dependent prefetchers. In: 33rd USENIX Security Symposium (USENIX Security 24). pp. 1117–1134. USENIX Association, Philadelphia, PA (2024)
12. Chen, J., Zhou, B., Chen, R., Jiang, H., Wang, Y., Huang, X., Zhao, Y., Yung, M.: Post-quantum tls 1.3 handshake from cpa-secure kems with tighter reductions. In: Daemen, J., Thomé, E. (eds.) Advances in Cryptology – EUROCRYPT 2026. pp. 499–529. Springer Nature Switzerland, Cham (2026)

13. Cremers, C.J.F.: Session-state reveal is stronger than ephemeral key reveal: Attacking the naxos authenticated key exchange protocol. In: Abdalla, M., Pointcheval, D., Fouque, P.A., Vergnaud, D. (eds.) *Applied Cryptography and Network Security*. pp. 20–33. Springer Berlin Heidelberg, Berlin, Heidelberg (2009)
14. Diffie, W., Hellman, M.: New directions in cryptography. *IEEE Transactions on Information Theory* **22**(6), 644–654 (1976). <https://doi.org/10.1109/TIT.1976.1055638>
15. Ding, J., Alsayigh, S., Saraswathy, R.V., Fluhrer, S., Lin, X.: Leakage of signal function with reused keys in rlwe key exchange. In: 2017 IEEE International Conference on Communications (ICC). pp. 1–6 (2017)
16. Ding, J., Branco, P., Schmitt, K.: Key exchange and authenticated key exchange with reusable keys based on RLWE assumption. *Cryptology ePrint Archive*, Paper 2019/665 (2019), <https://eprint.iacr.org/2019/665>
17. Ding, J., Fluhrer, S., Rv, S.: Complete attack on rlwe key exchange with reused keys, without signal leakage. In: Susilo, W., Yang, G. (eds.) *Information Security and Privacy*. pp. 467–486. Springer International Publishing, Cham (2018)
18. Ding, J., Xie, X., Lin, X.: A simple provably secure key exchange scheme based on the learning with errors problem. *Cryptology ePrint Archive*, Paper 2012/688 (2012), <https://eprint.iacr.org/2012/688>
19. Feltz, M., Cremers, C.: Strengthening the security of authenticated key exchange against bad randomness. *Designs, Codes and Cryptography* **86**(3), 481–516 (2018). <https://doi.org/10.1007/s10623-017-0337-5>
20. Fluhrer, S.: Cryptanalysis of ring-LWE based key exchange with key share reuse. *Cryptology ePrint Archive*, Paper 2016/085 (2016), <https://eprint.iacr.org/2016/085>
21. Fujioka, A., Suzuki, K., Xagawa, K., Yoneyama, K.: Strongly secure authenticated key exchange from factoring, codes, and lattices. In: Fischlin, M., Buchmann, J., Manulis, M. (eds.) *Public Key Cryptography – PKC 2012*. pp. 467–484. Springer Berlin Heidelberg, Berlin, Heidelberg (2012)
22. Gao, X., Ding, J., Li, L., Liu, J.: Practical randomized rlwe-based key exchange against signal leakage attack. *IEEE Transactions on Computers* **67**(11), 1584–1593 (2018)
23. Genêt, A., de Guertechin, N.L., Kaluđerović, N.: Full key recovery side-channel attack against ephemeral sike on the cortex-m4. In: Bhasin, S., De Santis, F. (eds.) *Constructive Side-Channel Analysis and Secure Design*. pp. 228–254. Springer International Publishing, Cham (2021)
24. Gong, B., Zhao, Y.: Cryptanalysis of rlwe-based one-pass authenticated key exchange. In: Lange, T., Takagi, T. (eds.) *Post-Quantum Cryptography*. pp. 163–183. Springer International Publishing, Cham (2017)
25. Guo, Q., Mårtensson, E., Åström, A.: The perils of limited key reuse: Adaptive and parallel mismatch attacks with post-processing against kyber. *IACR Communications in Cryptology* **1**(3) (2024)
26. Hashimoto, K., Katsumata, S., Kwiatkowski, K., Prest, T.: An efficient and generic construction for signal’s handshake (x3dh): Post-quantum, state leakage secure, and deniable. *Journal of Cryptology* **35**(3), 17 (2022). <https://doi.org/10.1007/s00145-022-09427-1>
27. Hülsing, A., Ning, K.C., Schwabe, P., Weber, F.J., Zimmermann, P.R.: Post-quantum wireguard. In: 2021 IEEE Symposium on Security and Privacy (SP). pp. 304–321 (2021)

28. Kiltz, E., Pan, J., Riepel, D., Ringerud, M.: Multi-user cdh problems and the concrete security of naxos and hmqv. In: Rosulek, M. (ed.) *Topics in Cryptology – CT-RSA 2023*. pp. 645–671. Springer International Publishing, Cham (2023)
29. Krawczyk, H.: Hmqv: A high-performance secure diffie-hellman protocol. In: Shoup, V. (ed.) *Advances in Cryptology – CRYPTO 2005*. pp. 546–566. Springer Berlin Heidelberg, Berlin, Heidelberg (2005)
30. LaMacchia, B., Lauter, K., Mityagin, A.: Stronger security of authenticated key exchange. In: Susilo, W., Liu, J.K., Mu, Y. (eds.) *Provable Security*. pp. 1–16. Springer Berlin Heidelberg, Berlin, Heidelberg (2007)
31. Law, L., Menezes, A., Qu, M., Solinas, J., Vanstone, S.: An efficient protocol for authenticated key agreement. *Designs, Codes and Cryptography* **28**(2), 119–134 (2003). <https://doi.org/10.1023/A:1022595222606>
32. Li, Z., Xu, J., Hu, L.: An improved signal leakage attack against dxl key exchange protocol. In: Lin, D., Wang, M., Yung, M. (eds.) *Information Security and Cryptology*. pp. 193–210. Springer Nature Singapore, Singapore (2025)
33. Liu, K., Gu, X., Ren, P., Nie, X., Lv, Y.: Speedy error reconciliation. In: Chen, R., Deng, R.H., Yung, M. (eds.) *Information Security and Cryptology*. pp. 166–186. Springer Nature Singapore, Singapore (2026)
34. Lyubashevsky, V., Peikert, C., Regev, O.: On ideal lattices and learning with errors over rings. In: Gilbert, H. (ed.) *Advances in Cryptology – EUROCRYPT 2010*. pp. 1–23. Springer Berlin Heidelberg, Berlin, Heidelberg (2010)
35. May, A., Nowakowski, J.: Too many hints – when LLL breaks LWE. In: Guo, J., Steinfeld, R. (eds.) *Advances in Cryptology – ASIACRYPT 2023*. pp. 106–137. Springer Nature Singapore, Singapore (2023)
36. Menezes, A., Qu, M., Vanstone, S.: Some new key agreement protocols providing implicit authentication. In: *Workshop on Selected Areas in Cryptography (SAC)’95*. pp. 22–32 (1995)
37. Pan, J., Wagner, B., Zeng, R.: Lattice-based authenticated key exchange with tight security. In: Handschuh, H., Lysyanskaya, A. (eds.) *Advances in Cryptology – CRYPTO 2023*. pp. 616–647. Springer Nature Switzerland, Cham (2023)
38. Peikert, C.: Lattice cryptography for the internet. In: Mosca, M. (ed.) *Post-Quantum Cryptography*. pp. 197–219. Springer International Publishing, Cham (2014)
39. Qin, Y., Cheng, C., Ding, J.: A complete and optimized key mismatch attack on nist candidate newhope. In: Sako, K., Schneider, S., Ryan, P.Y.A. (eds.) *Computer Security – ESORICS 2019*. pp. 504–520. Springer International Publishing, Cham (2019)
40. Qin, Y., Cheng, C., Zhang, X., Pan, Y., Hu, L., Ding, J.: A systematic approach and analysis of key mismatch attacks on lattice-based nist candidate kems. In: Tibouchi, M., Wang, H. (eds.) *Advances in Cryptology – ASIACRYPT 2021*. pp. 92–121. Springer International Publishing, Cham (2021)
41. Qin, Y., Ding, R., Cheng, C., Bindel, N., Pan, Y., Ding, J.: Light the signal: Optimization of signal leakage attacks against lwe-based key exchange. In: Atluri, V., Di Pietro, R., Jensen, C.D., Meng, W. (eds.) *Computer Security – ESORICS 2022*. pp. 677–697. Springer International Publishing, Cham (2022)
42. Rajendran, G., Ravi, P., D’Anvers, J.P., Bhasin, S., Chattopadhyay, A.: Pushing the limits of generic side-channel attacks on lwe-based kems - parallel pc oracle attacks on kyber kem and beyond. *IACR Transactions on Cryptographic Hardware and Embedded Systems* **2023**(2), 418–446 (2023)

43. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. In: Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing. p. 84–93. STOC '05, Association for Computing Machinery, New York, NY, USA (2005)
44. Schwabe, P., Stebila, D., Wiggers, T.: Post-quantum tls without handshake signatures. In: Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security. p. 1461–1480. CCS '20, Association for Computing Machinery, New York, NY, USA (2020). <https://doi.org/10.1145/3372297.3423350>
45. Shao, M., Liu, Y., Zhou, Y.: Pairwise and parallel: Enhancing the key mismatch attacks on kyber and beyond. In: Proceedings of the 19th ACM Asia Conference on Computer and Communications Security. p. 548–559. ASIA CCS '24, Association for Computing Machinery, New York, NY, USA (2024)
46. Ustaoglu, B.: Obtaining a secure and efficient key agreement protocol from (h)mqv and naxos. *Designs, Codes and Cryptography* **46**(3), 329–342 (2008)
47. Vershynin, R.: *High-Dimensional Probability: An Introduction with Applications in Data Science*. Cambridge Series in Statistical and Probabilistic Mathematics, Cambridge University Press, 2nd edn. (2026)
48. Wang, Y., Paccagnella, R., Wandke, A., Gang, Z., Garrett-Grossman, G., Fletcher, C.W., Kohlbrenner, D., Shacham, H.: Dvfs frequently leaks secrets: Hertzbleed attacks beyond sike, cryptography, and cpu-only data. In: 2023 IEEE Symposium on Security and Privacy (SP). pp. 2306–2320 (2023)
49. Xue, H., Lu, X., Li, B., Liang, B., He, J.: Understanding and constructing ake via double-key key encapsulation mechanism. In: Peyrin, T., Galbraith, S. (eds.) *Advances in Cryptology – ASIACRYPT 2018*. pp. 158–189. Springer International Publishing, Cham (2018)
50. Yao, A.C.C., Zhao, Y.: Oake: a new family of implicitly authenticated diffie-hellman protocols. In: Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security. pp. 1113–1128. CCS '13, Association for Computing Machinery, New York, NY, USA (2013)
51. Yu, T., Cheng, C., Yang, Z., Wang, Y., Pan, Y., Weng, J.: Hints from hertz: Dynamic frequency scaling side-channel analysis of number theoretic transform in lattice-based kems. *IACR Transactions on Cryptographic Hardware and Embedded Systems* **2024**(3), 200–223 (2024)
52. Zeng, K., Yang, C.H., Rao, T.R.N.: On the linear consistency test (lct) in cryptanalysis with applications. In: Brassard, G. (ed.) *Advances in Cryptology — CRYPTO'89 Proceedings*. pp. 164–174. Springer New York, New York, NY (1990)
53. Zhang, J., Zhang, Z., Ding, J., Snook, M., Dagdelen, Ö.: Authenticated key exchange from ideal lattices. In: Oswald, E., Fischlin, M. (eds.) *Advances in Cryptology - EUROCRYPT 2015*. pp. 719–751. Springer Berlin Heidelberg, Berlin, Heidelberg (2015)
54. Zhang, S., Lin, X., Yu, Y., Wang, W.: Improved power analysis attacks on falcon. In: Hazay, C., Stam, M. (eds.) *Advances in Cryptology – EUROCRYPT 2023*. pp. 565–595. Springer Nature Switzerland, Cham (2023)
55. Zhou, B., Jiang, H., Zhao, Y.: Cpa-secure kems are also sufficient for post-quantum tls 1.3. In: Chung, K.M., Sasaki, Y. (eds.) *Advances in Cryptology – ASIACRYPT 2024*. pp. 433–464. Springer Nature Singapore, Singapore (2025)