

FALCON⁺⁺: Shorter Signatures without NTRU Smoothing Estimates

Hao Yan and Nicholas Zhao

Imperial College London, London, United Kingdom
{h.yan22,n.zhao22}@imperial.ac.uk

Abstract. FALCON is a digital signature scheme based on NTRU lattices, known for its compact signatures. Its weak-smoothness variant, FALCON^{ws} (ASIACRYPT 2025), reduces signature sizes by allowing narrower Gaussian sampling distributions. Further reductions require efficient sampling at narrower widths and a security analysis that accounts for the resulting distributions. We analyze the distribution of accepted preimages directly, without estimating the smoothing parameter of the NTRU lattice. We also modify the rejection-corrected Klein–GPV sampler to reduce sampling cost while controlling the resulting distributional error. We combine these results to construct FALCON⁺⁺, a variant of FALCON, and give UF-CMA and SUF-CMA bounds in the classical random-oracle model. For ring degrees 512 and 1024, our estimates give signature sizes of 419 and 921 bytes, respectively. These are 37.09% and 28.05% smaller than FALCON signatures, and 11.60% and 5.92% smaller than the estimates for FALCON^{ws} at their respective parameter sets.

1 Introduction

Compact signatures are a central objective of lattice-based hash-and-sign constructions. The GPV framework [17] realizes this paradigm through trapdoor preimage sampling. NTRU-based schemes include FALCON [13] and Mitaka [10], while gadget trapdoors [25] support another family of hash-and-sign constructions, including compact-gadget signatures [32]. FALCON combines compact NTRU trapdoors with fast Fourier sampling [8], yielding particularly small signatures. With the dimension and other parameters fixed, reducing the Gaussian width offers a natural route to shorter signatures. Doing so, however, requires both efficient sampling and a security analysis that accounts for the resulting distributions.

These requirements have motivated several approaches to lattice Gaussian sampling, such as the Klein–GPV sampler [22,17], Peikert’s parallel sampler [26], and hybrid sampling [27]. Other approaches seek to sample at smaller widths by changing the sampling mechanism itself [30,31,11,6]. For FALCON-type signatures, [33,7] uses rejection correction to obtain discrete Gaussian outputs at smaller widths. Such correction removes sampling bias, but its repetition cost can increase as the width decreases.

A complementary line of work focuses on the security analysis. The use of Rényi divergence instead of statistical distance permits larger sampling errors within a prescribed security loss [4,28,29]. Fouque et al. [12] apply this approach to FALCON-type signatures and introduce FALCON⁺, while Zhang et al. [33] investigate weak smoothness and obtain FALCON^{ws} with smaller Gaussian widths.

The smoothing parameter remains a common foundation of these analyses, even as the required accuracy and the measures of distributional error are refined. The Rényi smoothing parameter [23] further connects smoothing conditions with divergence-based security analysis. Turning these conditions into concrete parameters, however, requires bounds on the relevant smoothing parameters. For NTRU lattices, tight estimates are difficult to obtain, and concrete sampling widths are typically selected using sufficient conditions of the form $s \geq Q(T) \eta_{\varepsilon_0}(\mathbb{Z})$, where $Q(T)$ is the sampler-dependent trapdoor quality. These conditions may require larger Gaussian widths than the underlying distributional comparisons actually need, limiting the scope for further reductions in signature size. Refined smoothing estimates can help improve these guarantees [11,19].

This motivates a different approach: directly analyzing the marginal distribution of the accepted preimage and its relation to the discrete Gaussian, without relying on an estimate of the NTRU smoothing parameter. Such an analysis must also account for sampling efficiency, since rejection correction can become costly at narrow Gaussian widths. This leads to two questions:

Can direct analysis of the accepted preimage distribution avoid estimating the NTRU smoothing parameter?

Can this analysis yield better parameter sets than FALCON^{ws}/FALCON⁺, while controlling sampling cost and security loss?

1.1 Our Contributions

We develop a sampling method and an analysis that address these questions, and apply them to a variant of FALCON. Our contributions are threefold.

Analysis without NTRU smoothing parameters. We develop a framework for analyzing preimage distributions without estimating the smoothing parameter of the NTRU lattice. For a sampler that draws a fresh syndrome after each rejection, we directly characterize the marginal distribution of the accepted preimage and its relation to the discrete Gaussian on $\mathbb{Z}^d$. We further derive computable upper bounds on the distributional deviation from the full sequence of Gram–Schmidt lengths of the trapdoor basis (Sections 3–4).

A sampler with a tunable rejection multiplier. We propose a variant of the rejection-corrected Klein–GPV sampler [7,33] with a tunable rejection multiplier, allowing fewer sampling trials at the cost of a controlled deviation from the target Gaussian. The multiplier can be selected to meet a prescribed

sampler-loss budget, providing flexibility in the choice of sampling accuracy and efficiency at narrow Gaussian widths. At our selected parameters, the analytical comparison indicates a reduction in the expected iteration count by factors of approximately 17.8 and 12.0 relative to the adapted $\text{FALCON}^{\text{ws}}$ estimates, with a sampler-loss budget of at most two bits (Table 1).

Table 1. Expected sampling iterations at matched narrow-width parameters.

Parameter or expected iteration count	$n = 512$	$n = 1024$
Modulus q	509	1949
Trapdoor-quality parameter γ	1.245	1.170
Signature standard deviation σ_{sig}	14.3747	28.2544
$\text{FALCON}^{\text{ws}}$: expected correction iterations	33.77	20.94
FALCON^{++} (ours): expected total iterations	1.90	1.74

The $\text{FALCON}^{\text{ws}}$ estimates apply [33, Lem. 18] to the parameters above. The FALCON^{++} profile-derived upper estimates include norm rejection and use a sampler-loss budget of at most two bits (Section 6).

FALCON^{++} : security and compact signatures. Combining our analysis of preimage distributions with the proposed sampler, we construct FALCON^{++} , a FALCON -type signature scheme building on FALCON^+ and $\text{FALCON}^{\text{ws}}$ and designed for narrower Gaussian widths. We give UF-CMA and SUF-CMA reductions in the classical random-oracle model, accounting for the modified sampling distribution and signing retries (Theorems 3 and 4). We propose two parameter sets targeting NIST security Categories I and V, with ring degrees 512 and 1024, respectively. Their signature size are reduced by 11.60% and 5.92% relative to $\text{FALCON}^{\text{ws}}$. The public-key size is also reduced by 10% at degree 512 and remains unchanged at degree 1024. Table 2 compares the resulting sizes and heuristic attack costs with $\text{FALCON}/\text{FALCON}^+$ and $\text{FALCON}^{\text{ws}}$.

Table 2. Public-key and signature sizes (bytes) and heuristic attack costs at each scheme’s own parameters.

Scheme	$n = 512$			$n = 1024$		
	$ pk $	$ sig $	Attack (C/Q)	$ pk $	$ sig $	Attack (C/Q)
$\text{FALCON} / \text{FALCON}^+$	897	666	120/108	1793	1280	273/248
$\text{FALCON}^{\text{ws}}$	640	474	142/129	1408	979	284/258
FALCON^{++} (ours)	576	419	136/124	1408	921	273/248

C/Q gives heuristic classical/quantum attack-cost exponents.

1.2 Our Techniques

Preimage distributions over the integers. As in FALCON^+ [12], each signing attempt uses a fresh salt. In our scheme, a failed correction test also starts a new attempt. For a previously undefined random-oracle input, the resulting syndrome $\mathbf{c}$ is uniform over the quotient. The candidate preimage $\mathbf{x}$ lies in $\Lambda + \mathbf{c}$, so averaging over $\mathbf{c}$ gives a distribution on the whole integer lattice $\mathbb{Z}^d$. We compute this marginal distribution and account for the rejection step to compare the accepted preimage with the discrete Gaussian on $\mathbb{Z}^d$. This comparison does not require an estimate of the NTRU lattice’s smoothing parameter.

Our bounds on distributional error and sampling cost depend on the full sequence of Gram–Schmidt lengths of the trapdoor basis. To evaluate these bounds, we refine Prest’s random-projection heuristic [27, Sec. 6.5.1] to describe the two blocks of the NTRU basis. The model captures the decrease in GSO lengths within each block and uses the lattice determinant to relate their scales. We compare its predictions with the GSO lengths of independently generated keys (Figure 2), and use the model to estimate distributional error and sampling cost at narrow Gaussian widths (Section 4.3).

Amplification and clipping. At narrow widths, the rejection-corrected Klein–GPV sampler [7] may require many trials before accepting a sample. We increase the acceptance probability by multiplying it by a common factor and capping the result at one. Without clipping, this scaling would increase the acceptance rate without changing the output distribution. The deviation from the target Gaussian therefore comes from clipping. We measure the weight removed by clipping, normalize it, and use the resulting *clipping deficit* to control both the expected number of iterations and the Rényi divergence of the output distribution (Theorem 1).

We bound this deficit through higher moments of the normalized correction factor. Using Poisson summation and root-lattice theta series, we obtain moment bounds that depend on the full sequence of GSO lengths. Finite Fourier sums let us compute upper bounds on the theta factors. These bounds then tell us how much we can increase the multiplier while keeping the cumulative sampler loss within a prescribed budget (Corollary 2).

From the sampler to FALCON^{++} . Our proof follows the framework of FALCON^+ [12]. We use the preceding Rényi divergence bounds to replace the sampler’s output with a discrete Gaussian over $\mathbb{Z}^{2n}$ in the signing simulation. The resulting loss accounts for all candidates, including those rejected by the norm test.

Under the assumptions stated in Section 5.2, UF-CMA reduces to multi-target R-ISIS. For SUF-CMA, we also need to handle forgeries that reuse the message and salt of a returned signature. Here we invoke the $\chi\text{-BDD}^{\mathcal{O}_s}$ assumption with a preimage oracle. It allows us to replace the syndrome induced by a Gaussian preimage with a uniform syndrome, while sampling the preimage

from the corresponding conditional Gaussian distribution. The two joint distributions are computationally indistinguishable under this assumption. Without the preimage oracle, this assumption is analogous to the decisional decoding assumptions considered in rank-metric identity-based encryption [14] and HQC [2]. The remaining task is to find a second short preimage, as captured by R-SPISIS.

1.3 Related Work

Gaussian sampling and smoothing. Existing approaches include Markov-chain methods [30,31] and samplers based on lattice extensions, quotients, or code structure [11,6]. For ideal lattices with a norm-separating relative basis, theta-series factorization gives exact fixed-coset identities without smoothing conditions [24]. We focus on preimage sampling for NTRU-based signatures, where exact rejection correction [7] can require many trials at narrow Gaussian widths. Our sampler allows a controlled deviation from the target Gaussian to reduce this cost.

FALCON-type signatures. FALCON⁺ sharpens FALCON’s concrete security analysis [12], and FALCON^{ws} supports narrower Gaussian widths under weak smoothness [33]. Related analyses concern Rényi smoothing [23], discrete Gaussian leftover hash lemmas [20], and security estimates based on the secret basis [19]. Building on FALCON⁺ and FALCON^{ws}, we combine our preimage-distribution analysis with a sampler that permits controlled Gaussian error to construct FALCON⁺⁺. This combination enables us to explore narrower Gaussian widths for shorter signatures while accounting for both sampling cost and the resulting security loss.

Discrete Gaussian sampling for lattice algorithms. Building on earlier DGS-based SVP algorithms [1], recent work further explores discrete Gaussian sampling as a way to find shortest lattice vectors [16,21,15]. Our sampler is designed for trapdoor-based signature generation over NTRU lattices, rather than for solving SVP or mounting lattice attacks.

1.4 Organization

Section 2 introduces the preliminaries. Sections 3 and 4 develop the sampler, moment bounds, and multiplier selection. Section 5 presents FALCON⁺⁺ and its security bounds, and Section 6 evaluates the selected parameters.

2 Preliminaries

We use column vectors and the Euclidean norm $\|\cdot\|$. For an integer $t \geq 1$, write $[t] = \{1, \dots, t\}$. The notation $x \leftarrow D$ denotes sampling from a distribution D , and $\text{Unif}(S)$ is the uniform distribution on a finite set S . For $u \in \mathbb{R}$, let $(u)_+ := \max\{u, 0\}$. We use $\log_2$ for binary logarithms and $\ln$ for natural logarithms.

2.1 Lattices and Gaussian notation

A full-rank lattice with ordered basis $\mathbf{B} = [\mathbf{b}_0 | \cdots | \mathbf{b}_{d-1}] \in \mathbb{R}^{d \times d}$ is $\Lambda = \mathcal{L}(\mathbf{B}) = \mathbf{B}\mathbb{Z}^d$. Write $\tilde{\mathbf{b}}_0, \dots, \tilde{\mathbf{b}}_{d-1}$ for its Gram-Schmidt vectors. Its covolume satisfies $\det(\Lambda) = |\det(\mathbf{B})| = \prod_{i=0}^{d-1} \|\tilde{\mathbf{b}}_i\|$. For $\Lambda \subseteq \mathbb{Z}^d$, the quotient $\mathbb{Z}^d/\Lambda$ consists of $\det(\Lambda)$ cosets. We distinguish a coset from a chosen representative $\mathbf{c} \in \mathbb{Z}^d$, writing the coset as $\Lambda + \mathbf{c}$.

For a width $s > 0$ and center $\mathbf{c} \in \mathbb{R}^d$, define

$$\rho_{s,\mathbf{c}}(\mathbf{x}) = e^{-\pi\|\mathbf{x}-\mathbf{c}\|^2/s^2}, \quad \rho_{s,\mathbf{c}}(S) = \sum_{\mathbf{x} \in S} \rho_{s,\mathbf{c}}(\mathbf{x}).$$

For a discrete set S with $0 < \rho_{s,\mathbf{c}}(S) < \infty$, the discrete Gaussian on S is

$$\mathsf{D}_{S,s,\mathbf{c}}(\mathbf{x}) = \frac{\rho_{s,\mathbf{c}}(\mathbf{x})}{\rho_{s,\mathbf{c}}(S)}, \quad \mathbf{x} \in S.$$

We omit the center subscript when it is zero. In particular, $\mathsf{D}_{\Lambda+\mathbf{t},s}$ is supported on the shifted lattice $\Lambda + \mathbf{t}$ and centered at the origin. We convert between the cryptographic width s and the standard-deviation parameter σ by

$$s = \sqrt{2\pi} \sigma. \tag{1}$$

Here σ is the standard deviation of the corresponding continuous Gaussian, rather than necessarily that of its discrete counterpart. For $x > 0$, define

$$\vartheta_3(x) := \sum_{k \in \mathbb{Z}} e^{-\pi x k^2}.$$

Poisson summation gives $\rho_u(\mathbb{Z}) = u\vartheta_3(u^2)$ for $u > 0$.

For the NTRU setting, let n be a power of two, $q \geq 2$ an integer, $\mathcal{R} = \mathbb{Z}[X]/(X^n+1)$, and $\mathcal{R}_q = \mathcal{R}/q\mathcal{R}$. We identify ring elements with their coefficient vectors, so $\mathcal{R}^2$ is identified with $\mathbb{Z}^{2n}$ and all norms use this coefficient embedding. A public key $h \in \mathcal{R}_q$ defines

$$\pi_h : \mathcal{R}^2 \longrightarrow \mathcal{R}_q, \quad \pi_h(u, v) = u + hv \bmod q, \quad \Lambda_h = \ker(\pi_h).$$

The map π_h is surjective. Thus Λ_h is a full-rank integer lattice of dimension $d = 2n$, with

$$\mathcal{R}^2/\Lambda_h \simeq \mathcal{R}_q, \quad \det(\Lambda_h) = |\mathcal{R}_q| = q^n.$$

For each syndrome $c \in \mathcal{R}_q$, fix a representative $\boldsymbol{\tau}_c \in \pi_h^{-1}(c)$. Its preimages form the coset $\pi_h^{-1}(c) = \Lambda_h + \boldsymbol{\tau}_c$. A secret NTRU basis serves as a trapdoor for sampling short preimages [13]; the key-generation procedure is specified in Section 5.1.

2.2 The Klein proposal

Algorithm 1 gives the Klein–GPV sampler [22,17]. We call its output law the Klein proposal to distinguish it from the rejection-corrected distribution. At an arbitrary width s , this law need not equal $D_{\Lambda,s,\mathbf{c}}$. Throughout the analysis, each one-dimensional Gaussian draw in the algorithm is exact.

Algorithm 1: Klein proposal [22,17]

Input: Ordered basis $\mathbf{B}$ of Λ , center $\mathbf{c}$, and width $s > 0$.
Output: A lattice point $\mathbf{z} \in \Lambda$.
1 Compute $\mathbf{B} = \mathbf{Q}\mathbf{R}$ with $\mathbf{Q}$ orthogonal and $\mathbf{R} = (r_{ij})$ upper triangular with positive diagonal; set $\bar{\mathbf{c}} \leftarrow \mathbf{Q}^T \mathbf{c}$;
2 **for** $i = d - 1, \dots, 0$ **do**
3 Set $\xi_i \leftarrow \frac{\bar{c}_i - \sum_{j=i+1}^{d-1} r_{ij} a_j}{r_{ii}}$;
4 Draw $a_i \leftarrow D_{\mathbb{Z}, s/\|\tilde{\mathbf{b}}_i\|, \xi_i}$;
5 **end**
6 Set $\mathbf{z} \leftarrow \mathbf{B}\mathbf{a}$ and **return** $\mathbf{z}$;

For an output $\mathbf{z} = \mathbf{B}\mathbf{a}$, let $\xi_i(\mathbf{z})$ be the conditional center at step i . Since $r_{ii} = \|\tilde{\mathbf{b}}_i\|$, multiplying the one-dimensional probabilities gives [31, Eq. (10)]

$$\begin{aligned} p_{\text{KL}, \mathbf{B}, s, \mathbf{c}}(\mathbf{z}) &:= \prod_{i=0}^{d-1} D_{\mathbb{Z}, s/\|\tilde{\mathbf{b}}_i\|, \xi_i(\mathbf{z})}(a_i) \\ &= \frac{\rho_{s, \mathbf{c}}(\mathbf{z})}{\prod_{i=0}^{d-1} \rho_{s/\|\tilde{\mathbf{b}}_i\|, \xi_i(\mathbf{z})}(\mathbb{Z})}, \quad \mathbf{z} \in \Lambda. \end{aligned} \quad (2)$$

When $\mathbf{B}$ and s are fixed, we abbreviate this law to $p_{\text{KL}, \mathbf{c}}$. Its denominator generally depends on the output, unlike the normalizing mass $\rho_{s, \mathbf{c}}(\Lambda)$ of the target Gaussian. To obtain a preimage in $\Lambda + \mathbf{c}$, we use the representative $\mathbf{c}$ as the sampling center and set $\mathbf{x} = \mathbf{c} - \mathbf{z}$. Section 3 analyzes this preimage when the coset is chosen uniformly and a rejection test is applied.

2.3 Multi-target R-ISIS and second-preimage R-SPISIS

In the classical random-oracle model, an adversary given the public key makes at most Q_s signing and Q_H hash queries. UF-CMA requires a valid forgery on an unsigned message; SUF-CMA requires a valid message–signature pair never returned by the signer. Advantages are the corresponding success probabilities [12].

Following [12, Sec. 2.5], let $t \geq 1$, $s, \beta > 0$, and let PK denote the possibly nonuniform public-key distribution on $\mathcal{R}_q$ induced by key generation.

Multi-target R-ISIS. The problem $t\text{-R-ISIS}_{q, \text{PK}, \beta}$ gives $h \leftarrow \text{PK}$ and t independent uniform targets $u_i \in \mathcal{R}_q$, independent of h . It asks for $(j, \mathbf{x})$ with $\mathbf{x} \in \mathcal{R}^2$ such that

$$j \in [t], \quad \pi_h(\mathbf{x}) = u_j, \quad \|\mathbf{x}\| \leq \beta.$$

Second-preimage R-SPISIS. The problem t -R-SPISIS $_{q,\text{PK},s,\beta}$ gives $h \leftarrow \text{PK}$ and t pairs generated independently conditioned on h by drawing

$$c_i \leftarrow \text{Unif}(\mathcal{R}_q), \quad X_i \leftarrow D_{A_h + \tau_{c_i}, s, \mathbf{0}},$$

redrawing both values until $\|X_i\| \leq \beta$. The task is to find $(j, X') \in [t] \times \mathcal{R}^2$ satisfying

$$j \in [t], \quad \pi_h(X') = c_j, \quad \|X'\| \leq \beta, \quad X' \neq X_j.$$

Write $\text{Adv}_{\mathcal{A}}^{\mathcal{P}}$ for the success probability in either problem $\mathcal{P}$, over all experiment randomness.

2.4 χ -BDD $^{\mathcal{O}_s}$ with an exact preimage oracle

Following [33, Def. 6], the experiment $\chi\text{-BDD}_{\text{PK}}^{\mathcal{O}_s}[C]$ samples $h \leftarrow \text{PK}$ and, for a common hidden bit b , generates $C \geq 1$ challenges independently conditioned on (h, b) :

$$\begin{cases} T_i \leftarrow \text{Unif}(\mathcal{R}_q), & b = 0, \\ X_i \leftarrow D_{\mathcal{R}^2, s, \mathbf{0}}, \quad T_i = \pi_h(X_i), & b = 1. \end{cases}$$

Given $(h, T_1, \dots, T_C)$, the adversary may adaptively query each challenge index $i \in [C]$ at most once, receiving an independent sample

$$\mathcal{O}_s(i) : \quad Y_i \leftarrow D_{A_h + \tau_{T_i}, s, \mathbf{0}}.$$

The advantage is $\text{Adv}_{\mathcal{A}}^{\chi\text{-BDD}_{\text{PK}}^{\mathcal{O}_s}[C]} := |p_1 - p_0|$, where p_b is the probability of outputting 1 in branch b ; the assumption requires this advantage to be small within the specified resource bounds. The exact oracle is part of the assumption, accepts only challenge indices, and supplies no trapdoor. No norm test is applied in this experiment. We write $\chi\text{-BDD}_{\text{PK}}[C]$ for the same experiment without access to $\mathcal{O}_s$.

2.5 Rényi divergence

Let $\alpha > 1$ and let $P \ll Q$, meaning that $Q(x) = 0$ implies $P(x) = 0$. The multiplicative, unlogged Rényi divergence of order α is

$$\mathcal{R}_{\alpha}(P\|Q) = \left(\sum_{x \in \Omega} P(x)^{\alpha} Q(x)^{1-\alpha} \right)^{1/(\alpha-1)},$$

with the sum replaced by an integral for densities. If $P \not\ll Q$, its value is $+\infty$. Thus $\mathcal{R}_{\alpha}(P\|Q) \geq 1$, with equality precisely when $P = Q$ up to null sets. The divergence is directional and must not be treated as a metric. The limiting multiplicative divergence is

$$\mathcal{R}_{\infty}(P\|Q) := \text{ess sup}_{x \sim Q} \frac{dP}{dQ}(x).$$

For discrete distributions the essential supremum is the supremum over x with $Q(x) > 0$.

We repeatedly use the following properties from the lattice-cryptography literature [4,28]:

- **Data processing:** for every randomized channel K ,

$$\mathcal{R}_\alpha(PK\|QK) \leq \mathcal{R}_\alpha(P\|Q).$$

- **Multiplicativity:** $\mathcal{R}_\alpha(\prod_i P_i\|\prod_i Q_i) = \prod_i \mathcal{R}_\alpha(P_i\|Q_i)$. More generally, if every conditional step of an adaptive experiment is bounded by r_i for every history, then the joint divergence is at most $\prod_i r_i$.
- **Probability preservation:** for every event E ,

$$P(E) \leq (\mathcal{R}_\alpha(P\|Q)Q(E))^{(\alpha-1)/\alpha}.$$

At order infinity this becomes the direct domination

$$P(E) \leq \mathcal{R}_\infty(P\|Q)Q(E).$$

3 Fresh-Syndrome Klein–GPV Sampling with Clipped Rejection

We present a variant of the rejection-corrected Klein–GPV sampler that draws a fresh syndrome after each correction rejection. We first give the algorithm, then derive its output distribution and exact-sampling baseline. The cost of this baseline motivates amplifying and clipping the correction probability. We finally express the acceptance rate and the divergence from the Gaussian target through a clipping deficit.

3.1 The Fresh-Syndrome Sampler and Its Output Distribution

Fix an ordered basis $\mathbf{B}$ of a full-rank integer lattice $\Lambda \subseteq \mathbb{Z}^d$ and a Gaussian width $s > 0$. Algorithm 2 uses a multiplier $\Gamma_{\mathbf{B},s} > 0$ to control its correction acceptance probability. Every rejection restarts both the syndrome selection and the Klein proposal. The role of the multiplier will be explained later.

Algorithm 2: Fresh-syndrome Klein–GPV sampling with clipped rejection

Input: Basis B of $\Lambda \subseteq \mathbb{Z}^d$, width $s > 0$, multiplier $\Gamma_{B,s} > 0$, and an oracle returning independent uniform syndrome classes.
Output: (c, x) with $x \in \Lambda + c$.

- 1 Precompute $\rho_{s/\|\tilde{b}_i\|}(\mathbb{Z})$ for $i = 0, \dots, d-1$;
- 2 **repeat**
- 3 $c \leftarrow \text{Unif}(\mathbb{Z}^d/\Lambda)$; take its fixed representative $c \in \mathfrak{c}$;
- 4 $\Delta_B \leftarrow 1$;
- 5 **for** $i = d-1, \dots, 0$ **do**
- 6 $\xi_i \leftarrow \frac{\langle c - \sum_{j>i} z_j b_j, \tilde{b}_i \rangle}{\|\tilde{b}_i\|^2}$;
- 7 $z_i \leftarrow D_{\mathbb{Z}, s/\|\tilde{b}_i\|, \xi_i}$;
- 8 $\gamma_i \leftarrow \xi_i - z_i$;
- 9 $\Delta_B \leftarrow \Delta_B \cdot \frac{\rho_{s/\|\tilde{b}_i\|, \gamma_i}(\mathbb{Z})}{\rho_{s/\|\tilde{b}_i\|}(\mathbb{Z})}$;
- 10 **end**
- 11 $x \leftarrow c - \sum_{i=0}^{d-1} z_i b_i$;
- 12 $u \leftarrow [0, 1]$;
- 13 **until** $u \leq \min\{1, \Gamma_{B,s} \Delta_B\}$;
- 14 **return** (c, x) ;

Algorithm 2 retains Klein’s sequential sampling procedure and the correction factor in [7,33], with two modifications. First, after each rejection, it draws a fresh uniform syndrome class and uses its representative as the target center for the next iteration. In the signature scheme, this corresponds to sampling a fresh salt and recomputing the syndrome, see Section 5.1. Second, it accepts with probability $\min\{1, \Gamma_{B,s} \Delta_B\}$, where the multiplier $\Gamma_{B,s}$ controls amplification. The correction reuses the centers computed in this procedure (Algorithm 1), with no additional Gaussian draws.

A trial means one iteration of the outer repeat–until loop, from drawing a syndrome class through the acceptance test. Trials use independent uniform classes and independent randomness.

Let $Q_{B,s}$ denote the distribution of the preimage generated in a single trial of Algorithm 2, before the acceptance test. This distribution includes the randomness of both the uniform syndrome class and the Klein sampler.

Lemma 1. For $x \in \mathbb{Z}^d$, let $\gamma_i(x) = \frac{\langle x, \tilde{b}_i \rangle}{\|\tilde{b}_i\|^2}$. In a trial generating x , the accumulated correction factor is

$$\Delta_B(x) = \prod_{i=0}^{d-1} \frac{\rho_{s/\|\tilde{b}_i\|, \gamma_i(x)}(\mathbb{Z})}{\rho_{s/\|\tilde{b}_i\|}(\mathbb{Z})} \in (0, 1]. \quad (3)$$

The probability that a single trial generates $\mathbf{x}$ before the acceptance test is

$$Q_{\mathbf{B},s}(\mathbf{x}) = \frac{\rho_s(\mathbf{x})}{\det(\Lambda) \prod_{i=0}^{d-1} \rho_{s/\|\tilde{\mathbf{b}}_i\|, \gamma_i(\mathbf{x})}(\mathbb{Z})}. \quad (4)$$

Proof. Producing a fixed $\mathbf{x}$ uniquely determines the sampled class to be $\mathbf{x} + \Lambda$, which has probability $1/\det(\Lambda)$. For any representative $\mathbf{c}$ of that class, write the required lattice output as $\mathbf{z} = \mathbf{c} - \mathbf{x} = \sum_i z_i \mathbf{b}_i$. Since $\langle \mathbf{b}_j, \tilde{\mathbf{b}}_i \rangle = 0$ for $j < i$, the Klein centers satisfy

$$\begin{aligned} \gamma_i = \xi_i - z_i &= \frac{\langle \mathbf{c} - \sum_{j \geq i} z_j \mathbf{b}_j, \tilde{\mathbf{b}}_i \rangle}{\|\tilde{\mathbf{b}}_i\|^2} \\ &= \frac{\langle \mathbf{x}, \tilde{\mathbf{b}}_i \rangle}{\|\tilde{\mathbf{b}}_i\|^2} = \gamma_i(\mathbf{x}). \end{aligned}$$

Substituting into the accumulator gives Eq. (3). Since $z_i \in \mathbb{Z}$, integer-shift invariance gives

$$\rho_{s/\|\tilde{\mathbf{b}}_i\|, \gamma_i(\mathbf{x})}(\mathbb{Z}) = \rho_{s/\|\tilde{\mathbf{b}}_i\|, \xi_i}(\mathbb{Z}).$$

Together with $\rho_{s,\mathbf{c}}(\mathbf{z}) = \rho_s(\mathbf{x})$, Eq. (2) and the class probability yield Eq. (4). Neither expression depends on the representative. The integer-shift equality also identifies Eq. (3) with the rejection-correction factor of the Klein-GPV sampler expressed through the recursive centers ξ_i . $\square$

Lem. 1 describes the preimage generated in one trial before the acceptance test. We now incorporate this test and the repetition after rejection to characterize the final output distribution and the expected number of trials. Define

$$\delta_{\mathbf{B},s} := \frac{\rho_s(\mathbb{Z}^d)}{\det(\Lambda) \prod_{i=0}^{d-1} \rho_{s/\|\tilde{\mathbf{b}}_i\|}(\mathbb{Z})}. \quad (5)$$

The following lemma gives both quantities.

Lemma 2 (Output distribution). *The preimage output by Algorithm 2 has distribution $\nu_{\mathbf{B},s}$, given by*

$$\begin{aligned} \nu_{\mathbf{B},s}(\mathbf{x}) &= \frac{\Gamma_{\mathbf{B},s} \delta_{\mathbf{B},s}}{\mathbb{E}_{X \leftarrow Q_{\mathbf{B},s}}[\min\{1, \Gamma_{\mathbf{B},s} \Delta_{\mathbf{B}}(X)\}]} \\ &\quad \times D_{\mathbb{Z}^d, s, \mathbf{0}}(\mathbf{x}) \min \left\{ 1, \frac{1}{\Gamma_{\mathbf{B},s} \Delta_{\mathbf{B}}(\mathbf{x})} \right\}. \end{aligned} \quad (6)$$

for every $\mathbf{x} \in \mathbb{Z}^d$. Let N_{Kl} denote the number of iterations of the outer repeat-until loop in Algorithm 2, including the final accepted iteration. Then

$$\mathbb{E}[N_{\text{Kl}}] = \frac{1}{\mathbb{E}_{X \leftarrow Q_{\mathbf{B},s}}[\min\{1, \Gamma_{\mathbf{B},s} \Delta_{\mathbf{B}}(X)\}]}.$$

Proof. Define the single-trial acceptance probability

$$p_{\text{acc}} = \mathbb{E}_{X \leftarrow Q_{\mathbf{B},s}}[\min\{1, \Gamma_{\mathbf{B},s} \Delta_{\mathbf{B}}(X)\}] \in (0, 1].$$

Since $Q_{\mathbf{B},s}$ and $\Delta_{\mathbf{B}}$ are positive on $\mathbb{Z}^d$, $p_{\text{acc}} > 0$. First acceptance at trial t with preimage $\mathbf{x}$ requires rejection in the preceding $t - 1$ trials, followed by proposing and accepting $\mathbf{x}$. By independence, this event has probability

$$(1 - p_{\text{acc}})^{t-1} Q_{\mathbf{B},s}(\mathbf{x}) \min\{1, \Gamma_{\mathbf{B},s} \Delta_{\mathbf{B}}(\mathbf{x})\}.$$

Summing over the possible first-acceptance times gives

$$\begin{aligned} \nu_{\mathbf{B},s}(\mathbf{x}) &= \sum_{t=1}^{\infty} (1 - p_{\text{acc}})^{t-1} Q_{\mathbf{B},s}(\mathbf{x}) \min\{1, \Gamma_{\mathbf{B},s} \Delta_{\mathbf{B}}(\mathbf{x})\} \\ &= \frac{Q_{\mathbf{B},s}(\mathbf{x}) \min\{1, \Gamma_{\mathbf{B},s} \Delta_{\mathbf{B}}(\mathbf{x})\}}{p_{\text{acc}}}. \end{aligned}$$

Using $\min\{1, y\} = y \min\{1, 1/y\}$ for $y > 0$, rewrite the numerator as

$$\Gamma_{\mathbf{B},s} Q_{\mathbf{B},s}(\mathbf{x}) \Delta_{\mathbf{B}}(\mathbf{x}) \min\left\{1, \frac{1}{\Gamma_{\mathbf{B},s} \Delta_{\mathbf{B}}(\mathbf{x})}\right\}.$$

Substituting Lem. 1 cancels the shifted one-dimensional masses. Equation (5) and $D_{\mathbb{Z}^d, s, \mathbf{0}}(\mathbf{x}) = \rho_s(\mathbf{x})/\rho_s(\mathbb{Z}^d)$ then give Eq. (6) after expanding p_{acc} . Summing the first-acceptance probability over $\mathbf{x}$ instead gives

$$\Pr[N_{\text{Kl}} = t] = (1 - p_{\text{acc}})^{t-1} p_{\text{acc}}, \quad t \geq 1.$$

Thus N_{Kl} is geometric with mean $1/p_{\text{acc}}$. □

Note that the theorem holds without a smoothing assumption. The class $\mathbf{x} + \Lambda$ determines its fixed representative $\mathbf{c}$, so this marginal also determines the joint output $(\mathbf{c}, \mathbf{x})$.

Taking $\Gamma_{\mathbf{B},s} = 1$, the acceptance probability is $\Delta_{\mathbf{B}}(\mathbf{x})$. Substituting Lem. 1 into its expectation gives $p_{\text{acc}} = \delta_{\mathbf{B},s} \in (0, 1]$. Lem. 2 gives

$$\mathbf{x} \sim D_{\mathbb{Z}^d, s, \mathbf{0}}, \quad p_{\text{acc}} = \delta_{\mathbf{B},s}, \quad \mathbb{E}[N_{\text{Kl}}] = \frac{1}{\delta_{\mathbf{B},s}}.$$

When $\delta_{\mathbf{B},s}$ is small, it requires many trials. Figure 1 illustrates how the modeled cost grows as the width decreases. This cost motivates taking $\Gamma_{\mathbf{B},s} > 1$ to increase acceptance. Section 3.2 analyzes the resulting clipping and its effect on the output distribution.

Exact-sampling cost at narrow widths

$n = 512$, $d = 1024$, $q = 953$, $\gamma = 1.17$, $\Gamma_{B,s} = 1$




Fig. 1. Modeled mean trials for $\Gamma_{B,s} = 1$, with $n = 512$, $q = 953$, and $\gamma = 1.17$. The curve plots the reciprocal of the normalization in Eq. (5), evaluated using modeled GSO lengths. The first block follows Eq. (15), while all second-block lengths are scaled by a common factor so that the product of all $2n$ lengths equals q^n .

3.2 Clipping Deficit, Sampling Cost, and Divergence

To analyze the clipped output distribution, define the normalized parameter

$$\tau = \Gamma_{B,s} \delta_{B,s}.$$

When $\Gamma_{B,s} = 1$, $\tau = \delta_{B,s}$. The value $\tau = 1$ instead corresponds to $\Gamma_{B,s} = 1/\delta_{B,s}$. To quantify the change from exact sampling, define the clipping deficit

$$d_{B,s}(\tau) := \mathbb{E}_{X \leftarrow Q_{B,s}} \left[\left(\frac{\Delta_B(X)}{\delta_{B,s}} - \frac{1}{\tau} \right)_+ \right]. \quad (7)$$

For a candidate $\mathbf{x}$, clipping caps the weight $\Gamma_{B,s} \Delta_B(\mathbf{x})$ at 1. The removed weight is therefore $(\Gamma_{B,s} \Delta_B(\mathbf{x}) - 1)_+$. Thus Eq. (7) can equivalently be written as

$$d_{B,s}(\tau) = \frac{1}{\tau} \mathbb{E}_{X \leftarrow Q_{B,s}} [(\Gamma_{B,s} \Delta_B(X) - 1)_+] = \frac{\mathbb{E}_{X \leftarrow Q_{B,s}} [(\Gamma_{B,s} \Delta_B(X) - 1)_+]}{\mathbb{E}_{X \leftarrow Q_{B,s}} [\Gamma_{B,s} \Delta_B(X)]}.$$

Thus $d_{B,s}(\tau)$ is the expected weight removed by clipping, normalized by the mean untruncated weight τ . It lets us rewrite the output law in a form that directly controls both sampling cost and divergence.

Theorem 1. For every $\tau > 0$, Algorithm 2, with $\Gamma_{\mathbf{B},s} = \tau/\delta_{\mathbf{B},s}$, satisfies $0 \leq d_{\mathbf{B},s}(\tau) < 1$. Its output distribution is

$$\nu_{\mathbf{B},s}(\mathbf{x}) = \frac{D_{\mathbb{Z}^d,s,\mathbf{0}}(\mathbf{x})}{1 - d_{\mathbf{B},s}(\tau)} \min \left\{ 1, \frac{\delta_{\mathbf{B},s}}{\tau \Delta_{\mathbf{B}}(\mathbf{x})} \right\}, \quad \mathbf{x} \in \mathbb{Z}^d.$$

Moreover,

$$\begin{aligned} \mathbb{E}[N_{\text{Kl}}] &= \frac{1}{\tau(1 - d_{\mathbf{B},s}(\tau))}, \\ \mathcal{R}_{\infty}(\nu_{\mathbf{B},s} \| D_{\mathbb{Z}^d,s,\mathbf{0}}) &\leq \frac{1}{1 - d_{\mathbf{B},s}(\tau)}. \end{aligned}$$

Proof. Both $Q_{\mathbf{B},s}$ and $D_{\mathbb{Z}^d,s,\mathbf{0}}$ are positive on $\mathbb{Z}^d$, and $\Delta_{\mathbf{B}}(\mathbf{x}), \delta_{\mathbf{B},s} > 0$, so all ratios below are well defined. Substituting Lem. 1 and Eq. (5) gives

$$\mathbb{E}_{Q_{\mathbf{B},s}} \left[\frac{\Delta_{\mathbf{B}}(X)}{\delta_{\mathbf{B},s}} \right] = \sum_{\mathbf{x} \in \mathbb{Z}^d} D_{\mathbb{Z}^d,s,\mathbf{0}}(\mathbf{x}) = 1.$$

Using $\Gamma_{\mathbf{B},s} = \tau/\delta_{\mathbf{B},s}$ in the acceptance rule, the single-trial acceptance probability p_{acc} satisfies

$$\begin{aligned} p_{\text{acc}} &= \mathbb{E}_{Q_{\mathbf{B},s}} \left[\min \left\{ 1, \tau \frac{\Delta_{\mathbf{B}}(X)}{\delta_{\mathbf{B},s}} \right\} \right] \\ &= \tau \mathbb{E}_{Q_{\mathbf{B},s}} \left[\min \left\{ \frac{\Delta_{\mathbf{B}}(X)}{\delta_{\mathbf{B},s}}, \frac{1}{\tau} \right\} \right] \\ &= \tau \left(\mathbb{E}_{Q_{\mathbf{B},s}} \left[\frac{\Delta_{\mathbf{B}}(X)}{\delta_{\mathbf{B},s}} \right] - \mathbb{E}_{Q_{\mathbf{B},s}} \left[\left(\frac{\Delta_{\mathbf{B}}(X)}{\delta_{\mathbf{B},s}} - \frac{1}{\tau} \right)_+ \right] \right) \\ &= \tau(1 - d_{\mathbf{B},s}(\tau)). \end{aligned}$$

The penultimate step uses $\min\{y, a\} = y - (y - a)_+$. The deficit is nonnegative by definition, and $p_{\text{acc}} > 0$ implies $d_{\mathbf{B},s}(\tau) < 1$.

Rewriting Lem. 2 with $\Gamma_{\mathbf{B},s} = \tau/\delta_{\mathbf{B},s}$ and $p_{\text{acc}} = \tau(1 - d_{\mathbf{B},s}(\tau))$ gives

$$\nu_{\mathbf{B},s}(\mathbf{x}) = \frac{D_{\mathbb{Z}^d,s,\mathbf{0}}(\mathbf{x})}{1 - d_{\mathbf{B},s}(\tau)} \min \left\{ 1, \frac{\delta_{\mathbf{B},s}}{\tau \Delta_{\mathbf{B}}(\mathbf{x})} \right\}.$$

The geometric waiting time in Lem. 2 then gives

$$\mathbb{E}[N_{\text{Kl}}] = \frac{1}{p_{\text{acc}}} = \frac{1}{\tau(1 - d_{\mathbf{B},s}(\tau))}.$$

Finally, the likelihood ratio is

$$\frac{\nu_{\mathbf{B},s}(\mathbf{x})}{D_{\mathbb{Z}^d,s,\mathbf{0}}(\mathbf{x})} = \frac{1}{1 - d_{\mathbf{B},s}(\tau)} \min \left\{ 1, \frac{\delta_{\mathbf{B},s}}{\tau \Delta_{\mathbf{B}}(\mathbf{x})} \right\} \leq \frac{1}{1 - d_{\mathbf{B},s}(\tau)}.$$

Taking the supremum over $\mathbf{x} \in \mathbb{Z}^d$ proves the claimed $\mathcal{R}_{\infty}$ bound, since the target is positive everywhere. $\square$

Theorem 1 relates the sampling cost and the divergence from the target discrete Gaussian to the clipping deficit $d_{\mathbf{B},s}(\tau)$. In Section 4, we derive computable upper bounds on this deficit and use them to select τ , aiming for a small expected number of trials while controlling $\mathcal{R}_\infty(\nu_{\mathbf{B},s} \parallel \mathbf{D}_{\mathbb{Z}^d,s,\mathbf{0}})$. This divergence bound is then used to quantify the loss in the security analysis of Section 5.2.

4 Moment Bounds and Heuristic Cost Model

This section bounds the clipping deficit from a trapdoor's Gram-Schmidt profile and derives parameter choices for a target loss. It then presents the profile heuristic used in the concrete evaluation.

4.1 Key-Specific Moment Bounds

Theorem 1 expresses both the sampling cost and the divergence bound in terms of the clipping deficit $d_{\mathbf{B},s}(\tau)$. We now derive a computable upper bound on $d_{\mathbf{B},s}(\tau)$ from the Gram-Schmidt lengths of $\mathbf{B}$. The argument first reduces the deficit to a moment, then bounds that moment using the basis profile.

For $y, a > 0$ and every integer $k \geq 2$, we have the pointwise inequality $(y - a)_+ \leq \frac{(k-1)^{k-1}}{k^k} a^{1-k} y^k$. The inequality follows by maximizing $(y - a)/y^k$ over $y > a$, whose maximum is attained at $y = ka/(k-1)$. For $y \leq a$, it is immediate. Applying this inequality pointwise with $y = \Delta_{\mathbf{B}}/\delta_{\mathbf{B},s}$ and $a = 1/\tau$ gives

$$d_{\mathbf{B},s}(\tau) \leq \frac{(k-1)^{k-1}}{k^k} \tau^{k-1} \mathbb{E}_{X \leftarrow Q_{\mathbf{B},s}} \left[\left(\frac{\Delta_{\mathbf{B}}(X)}{\delta_{\mathbf{B},s}} \right)^k \right]. \quad (8)$$

The moment on the right depends on $\mathbf{B}, s$ and the order k , but not on τ . We can therefore bound it before choosing the amplification parameter. The following theorem gives such a bound using only the Gaussian width and the Gram-Schmidt lengths. For $u > 0$, define

$$g_{u,\mathbb{Z}}(t) = \frac{\rho_{u,t}(\mathbb{Z})}{u} = \sum_{m \in \mathbb{Z}} e^{-\pi u^2 m^2} e^{2\pi i m t}, \quad \int_0^1 g_{u,\mathbb{Z}}(t) dt = 1.$$

Theorem 2 (Fresh-syndrome correction moments). *For every integer $k \geq 2$,*

$$\mathbb{E}_{X \leftarrow Q_{\mathbf{B},s}} \left[\left(\frac{\Delta_{\mathbf{B}}(X)}{\delta_{\mathbf{B},s}} \right)^k \right] \leq \mathcal{M}_k(\mathbf{B}, s), \quad (9)$$

where

$$\mathcal{M}_k(\mathbf{B}, s) := \frac{\vartheta_3(s^2/k)^d}{\vartheta_3(s^2)^{dk}} \prod_{i=0}^{d-1} \int_0^1 g_{s/\|\tilde{\mathbf{b}}_i\|,\mathbb{Z}}(t)^k dt. \quad (10)$$

Proof. See Appendix D.1. □

Substituting Theorem 2 into Eq. (8) bounds the clipping deficit. To express the resulting bound explicitly in terms of the GSO lengths, we first rewrite the integral factors as root-lattice theta series. Let

$$A_{k-1} = \{(m_1, \dots, m_k) \in \mathbb{Z}^k : \sum_i m_i = 0\}$$

and write $\Theta_L(t) := \sum_{v \in L} e^{-\pi t \|v\|^2}$.

Lemma 3. *For every integer $k \geq 2$ and $u > 0$,*

$$\int_0^1 g_{u, \mathbb{Z}}(t)^k dt = \Theta_{A_{k-1}}(u^2).$$

Proof. Expand the k Fourier series $g_{u, \mathbb{Z}}(t) = \sum_{m \in \mathbb{Z}} e^{-\pi u^2 m^2} e^{2\pi i m t}$. Integration over one period keeps exactly the tuples satisfying $m_1 + \dots + m_k = 0$, which are the vectors of A_{k-1} . $\square$

Combining this identity with the moment bound gives the following explicit upper bound on the clipping deficit.

Corollary 1. *Fix a basis $\mathbf{B}$, a width $s > 0$, and a nonempty finite set $\mathcal{O} \subseteq \{2, 3, \dots\}$ of integer orders. For every $\tau > 0$,*

$$d_{\mathbf{B}, s}(\tau) \leq \min_{k \in \mathcal{O}} \left\{ \frac{(k-1)^{k-1}}{k^k} \tau^{k-1} \frac{\vartheta_3(s^2/k)^d}{\vartheta_3(s^2)^{dk}} \prod_{i=0}^{d-1} \Theta_{A_{k-1}} \left(\frac{s^2}{\|\tilde{\mathbf{b}}_i\|^2} \right) \right\}. \quad (11)$$

Proof. Combine Eq. (8), Theorem 2, and Lem. 3, then take the minimum over $k \in \mathcal{O}$. $\square$

For a fixed order k , amplification enters this bound through τ^{k-1} . The basis enters through the relative widths $s/\|\tilde{\mathbf{b}}_i\|$ in the theta product. Each theta factor decreases toward one as its relative width increases, and the product combines the contributions of all dimensions. Different orders give different bounds, so we choose the minimum over $k \in \mathcal{O}$.

To evaluate the bound in Eq. (11), the following lemma replaces each root-lattice theta factor by an explicit finite-sum upper bound.

Lemma 4 (Finite evaluation of $\Theta_{A_{k-1}}(u^2)$). *Let $u > 0$, and let $k \geq 2$, $M \geq 0$, and $R \geq 1$ be integers. Set*

$$\varepsilon_{k, M}(u) = 2k \exp \left(-\pi u^2 \frac{k}{k-1} (M+1)^2 \right).$$

If $\varepsilon_{k, M}(u) < 1$, then

$$\Theta_{A_{k-1}}(u^2) \leq \frac{1}{R(1 - \varepsilon_{k, M}(u))} \sum_{j=0}^{R-1} \left(1 + 2 \sum_{m=1}^M e^{-\pi u^2 m^2} \cos \frac{2\pi m j}{R} \right)^k. \quad (12)$$

Proof. See Appendix D.2. $\square$

Together with the theta representation above, this finite-sum bound makes $\mathcal{M}_k(\mathbf{B}, s)$ computable from the full GSO profile. For evaluation across different profiles, its theta factors can also be bounded between tabulation nodes using log-convexity.

Set $\phi_k(x) = \log \Theta_{A_{k-1}}(x)$, and let $\Pr_x[V = v] \propto e^{-\pi x \|v\|^2}$ on A_{k-1} . Then

$$\phi'_k(x) = -\pi \mathbb{E}_x[\|V\|^2], \quad \phi''_k(x) = \pi^2 \text{Var}_x(\|V\|^2) > 0. \quad (13)$$

For $x \in [x_j, x_{j+1}]$, put $\lambda = (x - x_j)/(x_{j+1} - x_j)$. Convexity gives $\phi_k(x) \leq (1 - \lambda)\phi_k(x_j) + \lambda\phi_k(x_{j+1})$. Thus upper bounds at the tabulation nodes extend to each interval by chord interpolation in $x = u^2$.

4.2 Parameter Selection from the Moment Bound

In our parameter selection, we use the preceding deficit bound to choose Γ under a prescribed security-loss budget. We restrict the choice to $0 < \tau \leq 1$ and select the largest value allowed by the moment bounds over a public nonempty finite set of integer orders $\mathcal{O} \subseteq \{2, 3, \dots\}$. This minimizes the resulting upper bound on the expected number of correction trials within this choice rule.

Let $C \geq 1$ bound the number of correction-accepted candidates in the security experiment, including candidates later rejected by the norm test, and allocate $b > 0$ bits of security loss to the sampler. Define

$$d_\star := 1 - 2^{-b/C}.$$

By Theorem 1, the condition $d_{\mathbf{B},s}(\tau) \leq d_\star$ ensures that the C -fold multiplicative Rényi loss is at most $(1 - d_\star)^{-C} = 2^b$. We therefore invert the deficit bound of Corollary 1 to obtain an admissible value of τ for each order, then choose the largest of these values.

Corollary 2 (Parameter selection for a target security loss). *For each $k \in \mathcal{O}$, define*

$$\tau_k := \min \left\{ 1, \left(\frac{k^k d_\star}{(k-1)^{k-1} \mathcal{M}_k(\mathbf{B}, s)} \right)^{1/(k-1)} \right\}. \quad (14)$$

Choose

$$\tau = \max_{k \in \mathcal{O}} \tau_k, \quad \Gamma_{\mathbf{B},s} = \frac{\tau}{\delta_{\mathbf{B},s}}.$$

Then the corresponding output distribution $\nu_{\mathbf{B},s}$ satisfies

$$d_{\mathbf{B},s}(\tau) \leq d_\star, \quad \mathcal{R}_\infty(\nu_{\mathbf{B},s} \| \mathbb{D}_{\mathbb{Z}^d, s, \mathbf{0}})^C \leq 2^b,$$

and the expected number of outer-loop iterations in Algorithm 2 satisfies

$$\mathbb{E}[N_{\text{KI}}] \leq \frac{1}{\tau(1 - d_\star)}.$$




Fig. 2. Empirical fFLDL length distributions under four parameter settings, compared with the heuristic predictions of Eq. (15). No Gram–Schmidt quality rejection is applied to the sampled keys.

Proof. For each $k \in \mathcal{O}$, Corollary 1 and Eq. (14) give

$$d_{\mathbf{B},s}(\tau_k) \leq \frac{(k-1)^{k-1}}{k^k} \tau_k^{k-1} \mathcal{M}_k(\mathbf{B}, s) \leq d_\star.$$

The last inequality also holds when the minimum in Eq. (14) selects 1. Since $\mathcal{O}$ is nonempty and finite, the chosen τ equals τ_k for some $k \in \mathcal{O}$, so $d_{\mathbf{B},s}(\tau) \leq d_\star$. Theorem 1 now yields

$$\mathcal{R}_\infty(\nu_{\mathbf{B},s} \| D_{\mathbb{Z}^d,s,\mathbf{0}})^C \leq (1 - d_{\mathbf{B},s}(\tau))^{-C} \leq (1 - d_\star)^{-C} = 2^b$$

and

$$\mathbb{E}[N_{\text{K1}}] = \frac{1}{\tau(1 - d_{\mathbf{B},s}(\tau))} \leq \frac{1}{\tau(1 - d_\star)}.$$

For fixed $d_\star$, this cost bound decreases with τ , which justifies selecting the largest τ_k . $\square$




Fig. 3. Parameter selection from the moment bounds. The factor $\delta_{B,s}$ is determined by B, s and is used only to convert the selected τ into the sampler multiplier $\Gamma_{B,s}$.

4.3 Heuristic Estimate

In the experiments of Section 6, we use the following heuristics for $0 \leq i < n/2$:

– *Heuristic 1:*

$$\begin{aligned} \|\tilde{\mathbf{b}}_{2i}\| &= \|\tilde{\mathbf{b}}_{2i+1}\| = \|\tilde{\mathbf{b}}_0\| \sqrt{\frac{n-i}{n}}, \\ \|\tilde{\mathbf{b}}_0\| &\approx \gamma \sqrt{q}. \end{aligned} \tag{15a}$$

– *Heuristic 2:*

$$\begin{aligned} \|\tilde{\mathbf{b}}_{n+2i}\| &= \|\tilde{\mathbf{b}}_{n+2i+1}\| = \|\tilde{\mathbf{b}}_n\| \sqrt{\frac{n-i}{n}}, \\ \|\tilde{\mathbf{b}}_n\| &\approx \frac{qe}{2\|\tilde{\mathbf{b}}_0\|} \approx \frac{e}{2\gamma} \sqrt{q}. \end{aligned} \tag{15b}$$

Here the dimensionless secret-sampling multiplier is $\gamma = \sigma_{f,g} \sqrt{2n/q}$. In the scheme below, the same current parameter value is used in both KeyGen norm filters; it is not a secret-key encoding width. For FALCON, $\gamma = 1.17$ [13], and hence $\|\tilde{\mathbf{b}}_0\| \approx 1.17\sqrt{q}$, $\|\tilde{\mathbf{b}}_n\| \approx 1.17\sqrt{q}$, which is heuristically used by Zhang et al. [33, Sec. 5.1.1].

We derive these heuristics following Prest’s random-projection model [27, Sec. 6.5.1]:

$$\mathbb{E}[\|\tilde{\mathbf{b}}_0\|^2] \approx 2n\sigma_{f,g}^2 = \gamma^2 q,$$

so $\|\tilde{\mathbf{b}}_0\| \approx \gamma \sqrt{q}$. After j preceding directions, $\tilde{\mathbf{b}}_j$ lies in a $(2n-j)$ -dimensional subspace, hence

$$\mathbb{E}[\|\tilde{\mathbf{b}}_j\|^2] \approx \frac{2n-j}{2n} \mathbb{E}[\|\tilde{\mathbf{b}}_0\|^2], \quad 0 \leq j < n.$$

Applying the same decay to the second block, Stirling’s formula gives

$$\left(\prod_{j=0}^{n-1} \frac{2n-j}{2n}\right)^{1/n} \approx \frac{2}{e}.$$

Combining this estimate with $\det(\mathbf{B}) = q^n$ yields

$$\|\tilde{\mathbf{b}}_n\| \approx \frac{qe}{2\|\tilde{\mathbf{b}}_0\|} \approx \frac{e}{2\gamma} \sqrt{q}.$$

The equality within each pair has a different origin. In the bit-reversal order of fast Fourier orthogonalization, $\mathbf{b}_{2i+1} = S\mathbf{b}_{2i}$, where S is the isometric negacyclic shift by $n/2$ positions. The preceding span is S -invariant, so Gram–Schmidt gives $\tilde{\mathbf{b}}_{2i+1} = S\tilde{\mathbf{b}}_{2i}$ and hence equal norms [8]. Taking $j = 2i$ above then gives $\sqrt{(n-i)/n}$ and yields Heuristics 1 and 2. Exact balancing occurs at $\gamma = \sqrt{e/2} \approx 1.1658$, where both initial scales equal $\sqrt{qe/2}$.

Figure 2 compares the heuristic predictions with empirical GSO profiles from four independently generated sets of valid NTRU secret pairs, without Gram–Schmidt quality rejection. MARD and P95 report the mean and 95th percentile, respectively, of the coordinatewise relative errors between the empirical median and the prediction, expressed as percentages.

5 Signature Scheme and Security

We present FALCON⁺⁺ as a variant of FALCON and isolate the changes to its key-generation distribution, signing distribution, and retry control flow. Figure 4 gives the defining algorithms. Section 5.2 states the UF-CMA and SUF-CMA bounds in the classical ROM.

5.1 FALCON⁺⁺: a variant of FALCON

FALCON⁺⁺ uses FALCON’s mathematical public-key form and verification equation, while modifying key generation, signing, and the signature codec. Fix the **KeyGen**⁺⁺ parameters $(\sigma_{f,g}, \gamma)$, signing width s , norm bound β , and a global correction multiplier $\hat{\Gamma} > 0$. **Sign**⁺⁺ instantiates the fresh-syndrome sampler of Section 3 by drawing a fresh salt r and setting $c = \mathcal{H}(h, r, M)$ in each trial. The correction-accepted preimage is then subjected to the norm test $\|\mathbf{x}\| \leq \beta$, and sampling restarts upon rejection. **KeyGen**⁺⁺ follows the NTRU trapdoor-generation procedure of FALCON and FALCON^{ws} [13,33], using secret width $\sigma_{f,g}$ and quality parameter $\gamma = \sigma_{f,g} \sqrt{2n/q}$. For $a \in \mathcal{R}$, write $a^* = a(x^{-1})$ and set $D = ff^* + gg^*$. The sampled pair must satisfy

$$\|(f, g)\|_2^2 \leq \lfloor \gamma^2 q \rfloor, \quad \left\| \left(\frac{qg^*}{D}, -\frac{qf^*}{D} \right) \right\|_2^2 \leq \gamma^2 q. \quad (16)$$

The solver NTRUGen performs the required invertibility and solvability checks and returns (f, g, F, G) satisfying $fG - gF = q$.

For $h = f^{-1}g \bmod q$, Figure 4 constructs the trapdoor basis $\mathbf{B}$ of Λ_h , where $\mathcal{A}(a) \in \mathbb{Z}^{n \times n}$ denotes the negacyclic convolution matrix of a . We use $\mathbf{c} = (c, 0) \in \mathcal{R}^2$ for the fixed coefficientwise lift of $c \in \mathcal{R}_q$, so $\pi_h(\mathbf{c}) = c$, and write $[\cdot]_q$ for coefficientwise centered reduction modulo q . We use canonical rANS encoding as in [33] for signature compression.

KeyGen⁺⁺($n, q; \sigma_{f,g}, \gamma$) $(f, g, F, G) \xleftarrow{\\$} \text{NTRUGen}(n, q; \sigma_{f,g}, \gamma)$ $h \leftarrow f^{-1}g \bmod q$ $\mathbf{B} \leftarrow \left(\begin{array}{c c} \mathcal{A}(g) & \mathcal{A}(G) \\ \hline -\mathcal{A}(f) & -\mathcal{A}(F) \end{array} \right) \in \mathbb{Z}^{2n \times 2n}$ return $(sk_{\text{alg}}, pk) = ((f, g, F, G), h)$ Verify($pk = h, M, \text{sig} = (r, e)$) if $r \notin \mathcal{R}_{\text{salt}}$, return 0 $s_2 \leftarrow \text{Decode}(e)$ if $s_2 = \perp$ or $\text{Encode}(s_2) \neq e$, return 0 $c \leftarrow \mathcal{H}(h, r, M)$ $s_1 \leftarrow [c - hs_2]_q$ return 1 $[\ (s_1, s_2) \ \leq \beta]$	Sign⁺⁺($sk_{\text{alg}} = (f, g, F, G), pk = h, M$) $\mathbf{B} \leftarrow \text{Expand}(f, g, F, G)$ (or cached basis) repeat repeat $r \xleftarrow{\\$} \mathcal{R}_{\text{salt}}$ $c \leftarrow \mathcal{H}(h, r, M) \in \mathcal{R}_q$ $\mathbf{c} \leftarrow (c, 0)$ $\mathbf{z} \xleftarrow{\\$} \text{Klein}(\mathbf{B}, \mathbf{c}, s)$ $\mathbf{x} \leftarrow \mathbf{c} - \mathbf{z} = (s_1, s_2)$ $u \xleftarrow{\\$} [0, 1]$ until $u \leq \min\{1, \hat{\Gamma} \Delta_{\mathbf{B}}(\mathbf{x})\}$ until $\ (s_1, s_2) \ \leq \beta$ $e \leftarrow \text{Encode}(s_2)$ return $\text{sig} = (r, e)$
--	---

Fig. 4. Construction of FALCON⁺⁺. The parameterized-norm-cap KeyGen⁺⁺ and unchanged verification algorithm are shown on the left; the fresh-syndrome clipped-rejection Sign⁺⁺ is shown on the right.

5.2 Unforgeability Bounds

In this subsection, we state conditional UF-CMA and SUF-CMA bounds for FALCON⁺⁺ in the classical random-oracle model, as formalized in Theorems 3 and 4, respectively. Our security reductions build on the proof framework of [12], adapted to our proposed sampler. Both statements explicitly assume the stopping-time and adaptive-composition transitions, as well as the programming and retention of the random oracle used in their appendices.

Let $\mathcal{K}$ denote the support of $(\mathbf{B}, h)$ induced by KeyGen⁺⁺($n, q; \sigma_{f,g}, \gamma$), where $\mathbf{B} = \text{Expand}(f, g, F, G)$.

Theorem 3 (Unforgeability). *Let $\mathcal{A}$ be any adversary against the UF-CMA security of FALCON⁺⁺, running in time $t_{\mathcal{A}}$ and making at most Q_s signing queries and Q_H random-oracle queries. Assume $0 < \tau_{\mathbf{B}} = \hat{\Gamma} \delta_{\mathbf{B},s} \leq 1$ for every*

$(\mathbf{B}, h) \in \mathcal{K}$, $\alpha \in [2, \infty]$. Let $C_s, L_s \in \mathbb{Z}_{>0}$ satisfy $Q_s \leq C_s \leq L_s$. Then there exists an adversary $\mathcal{B}$ against ordinary-key $(Q_H + 1)$ -target R-ISIS, running in time $t_{\mathcal{B}} \approx t_{\mathcal{A}}$, such that

$$\begin{aligned} \text{Adv}_{\mathcal{A}}^{\text{UF-CMA}} &\leq \sum_{i=0}^{Q_s-1} \binom{C_s}{i} p_{\text{norm}}^i (1 - p_{\text{norm}})^{C_s-i} \\ &\quad + \sum_{i=0}^{C_s-1} \binom{L_s}{i} p_{\text{corr}}^i (1 - p_{\text{corr}})^{L_s-i} \\ &\quad + \frac{L_s(L_s + Q_H)}{2^\kappa} + \left[r_\alpha^{C_s} \text{Adv}_{\mathcal{B}}^{(Q_H+1)\text{-R-ISIS}_{q,\text{PK},\beta}} \right]^{(\alpha-1)/\alpha}, \quad (17) \end{aligned}$$

where $r_\alpha = \sup_{\mathcal{K}} \mathcal{R}_\alpha(\nu_{\mathbf{B},s} \| \mathcal{D}_{\mathcal{R}^2,s,\mathbf{0}})$, $p_{\text{corr}} = \inf_{\mathcal{K}} \tau_{\mathbf{B}}(1 - d_{\mathbf{B},s}(\tau_{\mathbf{B}}))$, and $p_{\text{norm}} = \inf_{\mathcal{K}} \Pr_{X \leftarrow \nu_{\mathbf{B},s}}[\|X\| \leq \beta]$.

In Corollary 2, we set $C = C_s$. The conditional derivation is given in Appendix A.

Theorem 4 (Strong Unforgeability). *Let $\mathcal{A}$ be any adversary against the **SUF-CMA** security of FALCON^{++} , running in time $t_{\mathcal{A}}$ and making at most Q_s signing queries and Q_H random-oracle queries. Assume $0 < \tau_{\mathbf{B}} = \hat{\Gamma} \delta_{\mathbf{B},s} \leq 1$ for every $(\mathbf{B}, h) \in \mathcal{K}$. Let $\alpha \in [2, \infty]$ and $C_s \in \mathbb{Z}_{>0}$ satisfy $Q_s \leq C_s$. Then there exist adversaries $\mathcal{B}_1$, $\mathcal{D}$, and $\mathcal{B}_{\text{SP}}$ against ordinary-key $(Q_H + 1)$ -target R-ISIS, χ -BDD $_{\text{PK}}^{\mathcal{O}_s}[C_s]$, and Q_s -target R-SPISIS, respectively, running in time $t_{\mathcal{B}_1}, t_{\mathcal{D}}, t_{\mathcal{B}_{\text{SP}}} \approx t_{\mathcal{A}}$, such that*

$$\begin{aligned} \text{Adv}_{\mathcal{A}}^{\text{SUF-CMA}} &\leq \text{Adv}_{\mathcal{A}}^{\text{UF-CMA}} + Q_s 2^{-\kappa} \left(\frac{Q_s + 1}{2p_{\text{sig}}^2} + \frac{2(Q_H + 1)}{p_{\text{sig}}} \right) \\ &\quad + \sum_{i=0}^{Q_s-1} \binom{C_s}{i} p_{\text{norm}}^i (1 - p_{\text{norm}})^{C_s-i} \\ &\quad + \left[r_\alpha^{C_s} \left(\text{Adv}_{\mathcal{B}_1}^{(Q_H+1)\text{-R-ISIS}_{q,\text{PK},\beta}} + \text{Adv}_{\mathcal{D}}^{\chi\text{-BDD}_{\text{PK}}^{\mathcal{O}_s}[C_s]} \right. \right. \\ &\quad \quad \left. \left. + \text{Adv}_{\mathcal{B}_{\text{SP}}}^{Q_s\text{-R-SPISIS}_{q,\text{PK},s,\beta}} \right) \right]^{(\alpha-1)/\alpha}, \quad (18) \end{aligned}$$

where r_α , p_{corr} , and p_{norm} are the same as in Theorem 3, and $p_{\text{sig}} = p_{\text{corr}} p_{\text{norm}} > 0$.

The conditional derivation is given in Appendix B.

6 Parameter Selection and Concrete Evaluation

We present two parameter sets, $\text{FALCON}^{++}\text{-512}$ and $\text{FALCON}^{++}\text{-1024}$, with ring degrees $n = 512$ and $n = 1024$, respectively, under the query budget

Table 3. Selected parameters for FALCON⁺⁺-512 and FALCON⁺⁺-1024.

Parameter	FALCON ⁺⁺ -512	FALCON ⁺⁺ -1024
Target security λ	128	256
Modulus q	509	1949
Secret-width / KeyGen cap γ	1.245	1.170
Signature standard deviation σ_{sig}	14.3747	28.2544
Signature norm cap β	506	1407
Sampler-loss budget	2 bits	2 bits
Salt wire length / entropy κ	41 bytes / 325 bits	41 bytes / 325 bits

Table 4. Parameter comparison of FALCON⁺, FALCON^{ws}, and FALCON⁺⁺.

Scheme	q	$ pk $	$ sig $
FALCON ⁺ -512 [13,12]	12289	897	666
FALCON ^{ws} -512 [33]	953	640	474
FALCON ⁺⁺ -512	509	576	419
FALCON ⁺ -1024 [13,12]	12289	1793	1280
FALCON ^{ws} -1024 [33]	1949	1408	979
FALCON ⁺⁺ -1024	1949	1408	921

$(Q_H, Q_s) = (2^{96}, 2^{64})$. Table 3 lists the main parameters, and Fig. 5 summarizes the parameter-selection process. The FALCON⁺⁺ reference implementation and code for reproducing the numerical tables are available online.¹ Both sets use a uniformly sampled 325-bit salt, encoded in $\lceil 325/8 \rceil = 41$ bytes with the three unused high bits set to zero. We compare public-key and signature sizes. For FALCON⁺⁺ and FALCON^{ws}, fixed-width coefficient encoding gives $|pk| = |h| = \lceil n \lceil \log_2 q \rceil / 8 \rceil$. We use entropy estimates of signature size, since rANS coding can attain lengths close to the entropy limit [33, Sec. 7]. With $H(\sigma) = \log_2(\sigma\sqrt{2\pi e})$, we estimate the FALCON⁺⁺ signature size as $|sig| = 41 + \lceil nH(\sigma_{\text{sig}})/8 \rceil$.² Table 4 compares the resulting sizes. Relative to FALCON^{ws}’s reported estimates, the signature estimates decrease by 11.60% and 5.92% for FALCON⁺⁺-512 and FALCON⁺⁺-1024, respectively.

6.1 Acceptance bounds and sampling cost

We set

$$\beta = \left\lceil 1.1 \sigma_{\text{sig}} \sqrt{2n} \right\rceil. \quad (19)$$

¹ <https://github.com/FanPUYen/FALCONpp>

² The accepted law P of s_2 differs from $G = D_{\mathbb{Z}^n, s, 0}$ but satisfies $H(P) \leq H(G) + \sqrt{(r_\infty/p_{\text{norm}} - 1) \text{Var}_G[-\log_2 G]}$. Numerical results give corrections of about 0.0202/0.000186 bytes for $n = 512/1024$, leaving the rounded estimates 419/921 bytes unchanged.




Fig. 5. Parameter dependencies with moment-based multiplier selection. The dashed arrow shows the dependence of d_* on C_s . The concrete evaluation presets sufficient caps and checks them against the resulting acceptance bounds. For a fixed global $\hat{\Gamma}$, evaluation uses $\tau_{\mathbf{B}} = \hat{\Gamma} \delta_{\mathbf{B},s}$. Acceptance, divergence, and trial counts are bounded rather than evaluated exactly.

Table 5. Sampling parameters and heuristic-profile bounds for the selected sets.

Quantity	FALCON ⁺⁺ -512	FALCON ⁺⁺ -1024
Moment order k	130	145
Global $\hat{\Gamma}$ (display only)	3.762128	9.067426
r_∞ upper estimate	$1 + 7.52 \times 10^{-20}$	$1 + 7.43 \times 10^{-20}$
p_{corr} lower estimate	0.526341	0.575072
p_{norm} lower estimate	$1 - e^{-9.9275}$	$1 - e^{-19.9940}$
Our sampler total trials	1.90	1.74
Sampler trials [33, Lem. 18]	33.77	20.94

The correction-acceptance and divergence bounds combine Theorem 1, Corollary 1, and Lem. 4. Following Fig. 5, we first use the heuristic GSO profile and $s = \sqrt{2\pi}\sigma_{\text{sig}}$ to evaluate $\delta_{\mathbf{B},s}$ and the moment bounds $\mathcal{M}_k(\mathbf{B}, s)$. With a sampler-loss budget of $b = 2$ bits and the candidate caps C_s preset in Section 6.2, we set $d_\star = 1 - 2^{-b/C_s}$ and apply Corollary 2 to select τ . The conversion $\Gamma_{\mathbf{B},s} = \tau/\delta_{\mathbf{B},s}$ is evaluated on the modeled profile to choose a global operational multiplier $\hat{\Gamma}$, reported in Table 5. For each operational value, we substitute $\tau = \hat{\Gamma}\delta_{\mathbf{B},s}$ back into the deficit bound to evaluate r_∞ and p_{corr} .

Lemma 5 (Computable p_{norm} bound). *For integer $\beta \geq 0$ with $u = (\beta^2 + 1)/(2n\sigma_{\text{sig}}^2) > 1$, we have*

$$p_{\text{norm}} \geq \max \left\{ 0, 1 - r_\infty e^{-n(u-1-\ln u)} \right\}. \quad (20)$$

The proof is in Appendix C.1. For positive acceptance bounds, the mean number of raw trials per signature satisfies $T_{\text{total}} \leq 1/(p_{\text{corr}}p_{\text{norm}})$.

For the comparison with FALCON^{ws} sampler at the same parameters, we use the adapted estimate from [33, Lem. 18]:

$$T_{\text{WS}}(n, \epsilon) = \exp \left(8.4n \int_1^2 \frac{(\epsilon/2)^t}{t^2} dt \right), \quad 0 < \epsilon \leq 0.047. \quad (21)$$

Here $u_\star = \sqrt{2\pi}\sigma_{\text{sig}}/(\gamma\sqrt{q})$ and $\epsilon = \vartheta_3(u_\star^2) - 1$. This uses the adapted approximation $\max_i \|\tilde{\mathbf{b}}_i\| \approx \gamma\sqrt{q}$. It compares analytical repetition estimates at the same $(n, q, \gamma, \sigma_{\text{sig}})$, rather than measured running times. At FALCON^{ws}'s original parameters [33, Table 7], the same calculation with $\epsilon = 1/(2n)$ gives 1.25/1.23 trials, which is a different comparison.

6.2 Trial caps and sampler loss

For the fixed query cap $Q_s = 2^{64}$, we preset sufficient C_s and L_s and check them against the selected profile bounds. This verifies a feasible budget rather than computing a joint optimum over the caps and multiplier.

Lemma 6 (Sufficient candidate caps). *For FALCON⁺⁺-512 with $\lambda = 128$ and FALCON⁺⁺-1024 with $\lambda = 256$, respectively, under the profile lower bounds of Table 5,*

$$\min \left\{ C_s \in \mathbb{Z}_{\geq Q_s} : \sum_{i=0}^{Q_s-1} \binom{C_s}{i} p_{\text{norm}}^i (1 - p_{\text{norm}})^{C_s-i} < 2^{-\lambda} \right\} < \begin{cases} 2^{64} + 2^{50}, & \text{FALCON}^{++}\text{-512}, \\ 2^{64} + 2^{36}, & \text{FALCON}^{++}\text{-1024}. \end{cases}$$

The proof is in Appendix C.2.

Lemma 7 (Sufficient raw-trial caps). *Fix $C_s = 2^{64} + 2^{50}$ for FALCON⁺⁺-512 and $C_s = 2^{64} + 2^{36}$ for FALCON⁺⁺-1024. Under the same profile bounds and target values of λ ,*

$$\min \left\{ L_s \in \mathbb{Z}_{\geq C_s} : \sum_{i=0}^{C_s-1} \binom{L_s}{i} p_{\text{corr}}^i (1 - p_{\text{corr}})^{L_s-i} < 2^{-\lambda} \right\} < 2^{65}.$$

The proof is in Appendix C.3.

Using the unrounded profile-specific divergence bounds, the sampler loss is

$$b_{\text{samp}} = C_s \log_2 r_{\infty} \leq \begin{cases} 2.00 \text{ bits}, & \text{FALCON}^{++}\text{-512}, \\ 1.98 \text{ bits}, & \text{FALCON}^{++}\text{-1024}. \end{cases} \quad (22)$$

These sampler-loss bounds are below the respective classical attack-cost margins of 8 and 17 bits reported in Section 6.3.

6.3 Heuristic attack-cost estimates

We estimate key-recovery and forgery costs using the full-dimensional BKZ success conditions of [13, Sec. 2.5.1, Eqs. (2.3)–(2.4)], with $\sigma_{f,g} = \gamma \sqrt{q/(2n)}$ and norm cap β . For BKZ block size b , we use the Core-SVP cost convention [3]

$$c_{\text{cl}}(b) = 0.292b, \quad c_{\text{qu}}(b) = 0.265b. \quad (23)$$

Reported work exponents are rounded down to integers.

Table 6 gives classical/quantum key-recovery costs of 136/124 bits for FALCON⁺⁺-512 and 273/248 bits for FALCON⁺⁺-1024. Forgery costs 169/153 and 372/338 bits, respectively, so key recovery is less costly for both parameter sets. The classical exponents exceed the target values 128 and 256 by 8 and 17 bits.

Table 6. Heuristic attack costs and χ -BDD model diagnostics.

Quantity	FALCON ⁺⁺ -512	FALCON ⁺⁺ -1024
Key recovery: BKZ block size	468	936
Key recovery: cost (bits)	136/124	273/248
Forgery: BKZ block size	581	1277
Forgery: cost (bits)	169/153	372/338
χ -BDD (single character, no $\mathcal{O}_s$)		
BKZ block size b	1024	2048
Classical reduction (bits)	299	598
Samples ($\log_2 N_{\text{suff}}$)	1401	2813

To assess the χ -BDD $^{\mathcal{O}_s}$ assumption, we follow FALCON^{ws} [33, Sec. 3.3.1] and consider lattice dual attacks against ordinary χ -BDD [9], heuristically disregarding the preimage oracle $\mathcal{O}_s$. For our generic single-character model, Table 6 reports classical lattice-reduction cost exponents of 299 and 598 for FALCON⁺⁺-512 and FALCON⁺⁺-1024, respectively, while the sufficient sample budgets exceed the available challenge counts.

References

1. Aggarwal, D., Dadush, D., Regev, O., Stephens-Davidowitz, N.: Solving the shortest vector problem in 2^n time via discrete Gaussian sampling. In: Proceedings of the 47th Annual ACM Symposium on Theory of Computing. pp. 733–742. ACM (2015). <https://doi.org/10.1145/2746539.2746606>
2. Aguilar Melchor, C., Blazy, O., Deneuville, J.C., Gaborit, P., Zémor, G.: Efficient encryption from random quasi-cyclic codes. IEEE Transactions on Information Theory **64**(5), 3927–3943 (2018). <https://doi.org/10.1109/TIT.2018.2804444>
3. Alkim, E., Ducas, L., Pöppelmann, T., Schwabe, P.: Post-quantum key exchange—a new hope. In: 25th USENIX Security Symposium (USENIX Security 16). pp. 327–343. USENIX Association, Austin, TX (Aug 2016), <https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/alkim>
4. Bai, S., Langlois, A., Lepoint, T., Stehlé, D., Steinfeld, R.: Improved security proofs in lattice-based cryptography: Using the Rényi divergence rather than the statistical distance. In: Iwata, T., Cheon, J.H. (eds.) Advances in Cryptology – ASIACRYPT 2015, Part I. LNCS, vol. 9452, pp. 3–24. Springer, Berlin, Heidelberg (2015). https://doi.org/10.1007/978-3-662-48797-6_1
5. Barbosa, M., Barthe, G., Doczkal, C., Don, J., Fehr, S., Grégoire, B., Huang, Y.H., Hülsing, A., Lee, Y., Wu, X.: Fixing and mechanizing the security proof of Fiat–Shamir with aborts and Dilithium. In: Handschuh, H., Lysyanskaya, A. (eds.) Advances in Cryptology – CRYPTO 2023, Part V. LNCS, vol. 14085, pp. 358–389. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-38554-4_12
6. Bollauf, M.F., Lie, M., Ling, C.: On Gaussian sampling for q -ary lattices and linear codes with Lee weight. In: Tauman Kalai, Y., Kamara, S.F. (eds.) Advances in Cryptology – CRYPTO 2025, Part I. LNCS, vol. 16000, pp. 321–352. Springer, Cham (2025). https://doi.org/10.1007/978-3-032-01855-7_11

7. Brakerski, Z., Langlois, A., Peikert, C., Regev, O., Stehlé, D.: Classical hardness of learning with errors. In: Proceedings of the 45th Annual ACM Symposium on Theory of Computing. pp. 575–584. ACM (2013). <https://doi.org/10.1145/2488608.2488680>
8. Ducas, L., Prest, T.: Fast Fourier orthogonalization. In: Proceedings of the 2016 ACM International Symposium on Symbolic and Algebraic Computation. pp. 191–198. ACM (2016). <https://doi.org/10.1145/2930889.2930923>
9. Ducas, L., Pulles, L.N.: Accurate score prediction for dual-sieve attacks. Cryptology ePrint Archive, Paper 2023/1850 (2023), <https://eprint.iacr.org/2023/1850>
10. Espitau, T., Fouque, P.A., Gérard, F., Rossi, M., Takahashi, A., Tibouchi, M., Wallet, A., Yu, Y.: Mitaka: A simpler, parallelizable, maskable variant of Falcon. In: Dunkelman, O., Dziembowski, S. (eds.) Advances in Cryptology – EUROCRYPT 2022, Part III. LNCS, vol. 13277, pp. 222–253. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-07082-2_9
11. Espitau, T., Wallet, A., Yu, Y.: On Gaussian sampling, smoothing parameter and application to signatures. In: Guo, J., Steinfeld, R. (eds.) Advances in Cryptology – ASIACRYPT 2023, Part VII. LNCS, vol. 14444, pp. 65–97. Springer, Singapore (2023). https://doi.org/10.1007/978-981-99-8739-9_3
12. Fouque, P.A., Gajland, P., de Groote, H., Janneck, J., Kiltz, E.: A closer look at FALCON. In: Daemen, J., Thomé, E. (eds.) Advances in Cryptology – EUROCRYPT 2026, Part IV. LNCS, vol. 16544, pp. 3–31. Springer, Cham (2026). https://doi.org/10.1007/978-3-032-25327-9_1
13. Fouque, P.A., Hoffstein, J., Kirchner, P., Lyubashevsky, V., Pornin, T., Prest, T., Ricosset, T., Seiler, G., Whyte, W., Zhang, Z.: FALCON: Fast-Fourier lattice-based compact signatures over NTRU. Specification, version 1.2 (Oct 2020), <https://falcon-sign.info/falcon.pdf>, accessed 29 July 2026
14. Gaborit, P., Hauteville, A., Phan, D.H., Tillich, J.P.: Identity-based encryption from codes with rank metric. In: Katz, J., Shacham, H. (eds.) Advances in Cryptology – CRYPTO 2017, Part III. LNCS, vol. 10403, pp. 194–224. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-63697-9_7
15. Gao, Y., Feng, Y., Hu, H.: Discrete Gaussian sampling meets BDGL decoding: Solving the shortest vector problem in $2^{0.5596n+o(n)}$ time. Cryptology ePrint Archive, Paper 2026/1844 (2026), <https://eprint.iacr.org/2026/1844>
16. Gao, Y., Feng, Y., Hu, H.: Solving the shortest vector problem in $2^{0.7314n+o(n)}$ time via discrete Gaussian sampling on superlattices. Cryptology ePrint Archive, Paper 2026/1587 (2026), <https://eprint.iacr.org/2026/1587>
17. Gentry, C., Peikert, C., Vaikuntanathan, V.: Trapdoors for hard lattices and new cryptographic constructions. In: Proceedings of the 40th Annual ACM Symposium on Theory of Computing. pp. 197–206. ACM (2008). <https://doi.org/10.1145/1374376.1374407>
18. Hoeffding, W.: Probability inequalities for sums of bounded random variables. Journal of the American Statistical Association **58**(301), 13–30 (1963). <https://doi.org/10.1080/01621459.1963.10500830>
19. Jia, H., Zhang, S., Yu, Y., Tang, C.: Revisiting the concrete security of FALCON-type signatures. In: Bai, S., Persichetti, E. (eds.) Public-Key Cryptography – PKC 2026, Part I. LNCS, vol. 16551, pp. 334–363. Springer, Cham (2026). https://doi.org/10.1007/978-3-032-26731-3_11
20. Jin, H., Liu, F.H., Wang, Z., Gu, D.: Discrete Gaussians modulo sub-lattices: New leftover hash lemmas for discrete Gaussians. In: Jager, T., Pan, J. (eds.) Public-Key Cryptography – PKC 2025, Part II. LNCS, vol. 15675, pp. 301–330. Springer, Cham (2025). https://doi.org/10.1007/978-3-031-91823-0_10

21. Kim, J.: One discrete Gaussian sample in $2^{n/2+o(n)}$ time. Cryptology ePrint Archive, Paper 2026/1599 (2026), <https://eprint.iacr.org/2026/1599>
22. Klein, P.N.: Finding the closest lattice vector when it's unusually close. In: Proceedings of the Eleventh Annual ACM-SIAM Symposium on Discrete Algorithms. pp. 937–941. Society for Industrial and Applied Mathematics (2000)
23. Ling, C., Luzzi, L., Yan, H.: The Rényi smoothing parameter and its applications in lattice-based cryptography. Cryptology ePrint Archive, Paper 2025/986 (2025), <https://eprint.iacr.org/2025/986>
24. Liu, R., Liang, Z., Ding, D.: Dimension reduction for theta series of ideal lattices. Cryptology ePrint Archive, Paper 2026/2028 (2026), <https://eprint.iacr.org/2026/2028>
25. Micciancio, D., Peikert, C.: Trapdoors for lattices: Simpler, tighter, faster, smaller. In: Pointcheval, D., Johansson, T. (eds.) Advances in Cryptology – EUROCRYPT 2012. LNCS, vol. 7237, pp. 700–718. Springer, Berlin, Heidelberg (2012). https://doi.org/10.1007/978-3-642-29011-4_41
26. Peikert, C.: An efficient and parallel Gaussian sampler for lattices. In: Rabin, T. (ed.) Advances in Cryptology – CRYPTO 2010. LNCS, vol. 6223, pp. 80–97. Springer, Berlin, Heidelberg (2010). https://doi.org/10.1007/978-3-642-14623-7_5
27. Prest, T.: Gaussian Sampling in Lattice-Based Cryptography. Ph.D. thesis, École Normale Supérieure, Paris, France (2015), <https://tprest.github.io/pdf/pub/thesis-thomas-prest.pdf>
28. Prest, T.: Sharper bounds in lattice-based cryptography using the Rényi divergence. In: Takagi, T., Peyrin, T. (eds.) Advances in Cryptology – ASIACRYPT 2017, Part I. LNCS, vol. 10624, pp. 347–374. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-70694-8_13
29. Takashima, K., Takayasu, A.: Tighter security for efficient lattice cryptography via the Rényi divergence of optimized orders. In: Au, M.H., Miyaji, A. (eds.) Provable Security – 9th International Conference, ProvSec 2015. LNCS, vol. 9451, pp. 412–431. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-26059-4_23
30. Wang, Z., Ling, C.: On the geometric ergodicity of Metropolis–Hastings algorithms for lattice Gaussian sampling. IEEE Transactions on Information Theory **64**(2), 738–751 (Feb 2018). <https://doi.org/10.1109/TIT.2017.2742509>
31. Wang, Z., Ling, C.: Lattice Gaussian sampling by Markov Chain Monte Carlo: Bounded distance decoding and trapdoor sampling. IEEE Transactions on Information Theory **65**(6), 3630–3645 (Jun 2019). <https://doi.org/10.1109/TIT.2019.2901497>
32. Yu, Y., Jia, H., Wang, X.: Compact lattice gadget and its applications to hash-and-sign signatures. In: Handschuh, H., Lysyanskaya, A. (eds.) Advances in Cryptology – CRYPTO 2023, Part V. LNCS, vol. 14085, pp. 390–420. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-38554-4_13
33. Zhang, S., Jia, H., Ran, D., Yu, Y., Yu, Y., Wang, X.: GPV preimage sampling with weak smoothness and its applications to lattice signatures. In: Hanaoka, G., Yang, B.Y. (eds.) Advances in Cryptology – ASIACRYPT 2025, Part III. LNCS, vol. 16247, pp. 233–264. Springer, Singapore (December 2025). https://doi.org/10.1007/978-981-95-5099-9_8

A Proof of Theorem 3

This appendix derives the displayed inequality under the assumptions on stopping times, adaptive composition, and random-oracle programming and retention stated in Theorem 3.

The parameters C_s and L_s are proof-only upper bounds and are not part of Sign^{++} . They bound, respectively, the total numbers of correction-accepted candidates and raw Klein trials over all signing queries.

Game G_0 . This is the plain UF-CMA game. Hence

$$\Pr[G_0^{\mathcal{A}} \Rightarrow 1] = \text{Adv}_{\mathcal{A}}^{\text{UF-CMA}}.$$

Game G_1 . This game is identical to G_0 , except that it aborts on a bad salt/point collision involving one of the first L_s raw signer trials. A bad collision occurs if (h, r, M) was already defined when the trial begins, or if an external random-oracle query later hits an internal signer point that did not produce a returned signature.

Claim 1.

$$|\Pr[G_0^{\mathcal{A}} \Rightarrow 1] - \Pr[G_1^{\mathcal{A}} \Rightarrow 1]| \leq \frac{L_s(L_s + Q_H)}{2^\kappa}. \quad (24)$$

Proof. The new abort can arise only from a point collision involving one of the at most L_s raw signer trials under consideration. Such a collision is covered by two cases: a collision among the signer trials, with probability at most $\binom{L_s}{2}/2^\kappa$, or a collision with one of the at most Q_H external random-oracle queries or the final forgery-point query, with probability at most $L_s(Q_H + 1)/2^\kappa$. Therefore, a union bound gives

$$\frac{\binom{L_s}{2} + L_s(Q_H + 1)}{2^\kappa} \leq \frac{L_s(L_s + Q_H)}{2^\kappa},$$

which holds for $L_s \geq 1$. $\square$

Game G_2 . This game is identical to G_1 , except that it aborts if more than C_s correction-accepted candidates are needed to answer the Q_s signing queries within the first L_s raw trials. In particular, the norm tests of the first C_s candidates have all been evaluated before this abort occurs, regardless of whether they pass or fail.

Claim 2.

$$|\Pr[G_1^{\mathcal{A}} \Rightarrow 1] - \Pr[G_2^{\mathcal{A}} \Rightarrow 1]| \leq \sum_{i=0}^{Q_s-1} \binom{C_s}{i} p_{\text{norm}}^i (1 - p_{\text{norm}})^{C_s-i}. \quad (25)$$

Proof. For each correction-accepted candidate $j \leq C_s$ within the first L_s raw trials, let

$$Y_j = \begin{cases} 1, & \text{if the } j\text{-th candidate passes the norm test,} \\ 0, & \text{if the } j\text{-th candidate fails the norm test.} \end{cases}$$

If fewer than C_s such candidates are produced, set the remaining Y_j 's to 1. A new abort in G_2 implies $\sum_{j=1}^{C_s} Y_j < Q_s$. Lem. 2 gives each candidate conditional law $\nu_{\mathbf{B},s}$ and, conditioned on the full history $\mathcal{F}_{j-1}$, we have $\Pr[Y_j = 1 \mid \mathcal{F}_{j-1}] \geq p_{\text{norm}}$. Since each Y_j , conditioned on the entire preceding history, succeeds with probability at least p_{norm} , the probability of obtaining fewer than Q_s successes among the C_s slots is no larger than in C_s independent Bernoulli trials with success probability p_{norm} . The probability of the new abort is therefore at most

$$\Pr \left[\sum_{j=1}^{C_s} Y_j < Q_s \right] \leq \Pr[\text{Bin}(C_s, p_{\text{norm}}) < Q_s],$$

which is exactly the displayed lower tail. $\square$

Game G_3 . This game is identical to G_2 , except that it aborts if answering the signing queries would require more than L_s raw Klein trials.

Claim 3.

$$|\Pr[G_2^{\mathcal{A}} \Rightarrow 1] - \Pr[G_3^{\mathcal{A}} \Rightarrow 1]| \leq \sum_{i=0}^{C_s-1} \binom{L_s}{i} p_{\text{corr}}^i (1 - p_{\text{corr}})^{L_s-i}. \quad (26)$$

Proof. A new abort in G_3 can occur only if fewer than C_s of the first L_s raw Klein trials are correction-accepted. Beside, every raw trial considered before this new abort has correction-acceptance probability

$$q_{\mathbf{B}} = \tau_{\mathbf{B}}(1 - d_{\mathbf{B},s}(\tau_{\mathbf{B}})) \geq p_{\text{corr}}.$$

Thus the probability that fewer than C_s correction acceptances occur among the first L_s raw trials is no larger than for L_s independent Bernoulli trials with success probability p_{corr} . Consequently,

$$\Pr[G_3 \text{ triggers its new abort}] \leq \Pr[\text{Bin}(L_s, p_{\text{corr}}) < C_s],$$

which is exactly the displayed lower tail. $\square$

Game G_4 . Replace the raw Klein trials up to the next correction acceptance by direct sampling $X \leftarrow \nu_{\mathbf{B},s}$. A fresh salt r is sampled from the undefined points, $C = \pi_h(X)$ is programmed at (h, r, M) , and the same C_s cap and norm test are then applied, as shown in Figure 6.

Game G_3	Game G_4
repeat $N_{\text{raw}} \leftarrow N_{\text{raw}} + 1$ abort if $N_{\text{raw}} > L_s$ $r \xleftarrow{\$} \mathcal{R}_{\text{salt}}$ abort if (h, r, M) was already defined or used $C \leftarrow \mathcal{H}(h, r, M)$ $Z \leftarrow \text{Klein}(\mathbf{B}, C, s)$ $X \leftarrow C - Z$ until $\text{CorrAccept}(X)$	repeat $r \xleftarrow{\$} \mathcal{R}_{\text{salt}}$ until (h, r, M) is absent from the retained $\mathcal{H}$ table $X = (s_1, s_2) \leftarrow \nu_{\mathbf{B}, s}$ $C \leftarrow \pi_h(X)$ $\mathcal{H}(h, r, M) \leftarrow C$

Fig. 6. Comparison of the steps that differ between Games G_3 and G_4 .

Claim 4.

$$\Pr[G_3^{\mathcal{A}} \Rightarrow 1] \leq \Pr[G_4^{\mathcal{A}} \Rightarrow 1].$$

Proof. We first observe that the points created by correction- or norm-rejected trials are invisible on every non-aborting execution of G_3 . Indeed, if a later signer or adversarial random-oracle query hits such a point, G_3 aborts by definition. We may therefore erase these points after rejection, resample a signer salt whenever its point is already defined, and remove the L_s cap. These changes only continue executions on which G_3 would otherwise abort, and hence cannot decrease the winning probability.

It remains to compare the generation of one correction-accepted candidate. Let

$$\mathcal{S}(M) = \{r \in \mathcal{R}_{\text{salt}} : (h, r, M) \text{ is undefined}\}.$$

In the modified G_3 , rejected trials leave no retained state. Hence every retry again chooses $r \leftarrow U(\mathcal{S}(M))$ and uses a fresh uniform syndrome. By Lem. 2, at the first correction acceptance the joint distribution is

$$\Pr[R = r, X = x, C = c] = \frac{\mathbf{1}[r \in \mathcal{S}(M)]}{|\mathcal{S}(M)|} \nu_{\mathbf{B}, s}(x) \mathbf{1}[c = \pi_h(x)].$$

This is exactly the distribution obtained in G_4 by first sampling $r \leftarrow U(\mathcal{S}(M))$ and $X \leftarrow \nu_{\mathbf{B}, s}$, then setting $C = \pi_h(X)$ and programming $\mathcal{H}(h, r, M) = C$. The candidate cap and norm test are unchanged. Therefore the resulting public view is distributed exactly as in G_4 , after changes that only remove aborts, which proves the claim. $\square$

Note that the leftover-hash lemma is not needed here, as the syndrome is never replaced by an independent uniform value. G_4 preserves the syndrome $C = \pi_h(X)$ induced by the accepted preimage.

Game G_5 . This game is identical to G_4 , except that $X \leftarrow \nu_{\mathbf{B},s}$ before the C_s cap is replaced by $X \leftarrow \mathsf{D}_{\mathcal{R}^2,s,\mathbf{0}}$. It still sets $C = \pi_h(X)$ and uses the same retained-table salt rule, C_s cap, programming rule, and norm test.

Claim 5. For every $\alpha \in [2, \infty]$,

$$\Pr[G_4^{\mathcal{A}} \Rightarrow 1] \leq (r_\alpha^{C_s} \Pr[G_5^{\mathcal{A}} \Rightarrow 1])^{(\alpha-1)/\alpha}.$$

Proof. Fix the sampled key $(\mathbf{B}, h)$ and condition on the complete history before the next candidate is generated. The only difference is therefore $X \leftarrow \nu_{\mathbf{B},s}$ and $X \leftarrow \mathsf{D}_{\mathcal{R}^2,s,\mathbf{0}}$. Both games then set $C = \pi_h(X)$. Hence the conditional Rényi divergence of the next candidate transition is

$$\mathcal{R}_\alpha((\nu_{\mathbf{B},s}, \pi_h(\nu_{\mathbf{B},s})) \parallel (\mathsf{D}_{\mathcal{R}^2,s,\mathbf{0}}, \pi_h(\mathsf{D}_{\mathcal{R}^2,s,\mathbf{0}}))) \leq r_\alpha.$$

The common salt and all subsequent processing, including random-oracle programming and the norm test, do not increase this divergence.

Since the C_s cap allows at most C_s candidate transitions, adaptive composition gives $\mathcal{R}_\alpha(\text{View}_{G_4} \parallel \text{View}_{G_5}) \leq r_\alpha^{C_s}$. Applying the probability-preservation property to the winning event gives $\Pr[G_4^{\mathcal{A}} \Rightarrow 1] \leq (r_\alpha^{C_s} \Pr[G_5^{\mathcal{A}} \Rightarrow 1])^{(\alpha-1)/\alpha}$. $\square$

Direct reduction from G_5 to R-ISIS.

Claim 6. There exists an adversary $\mathcal{B}$ against ordinary-key $(Q_H + 1)$ -R-ISIS such that

$$\Pr[G_5^{\mathcal{A}} \Rightarrow 1] \leq \text{Adv}_{\mathcal{B}}^{(Q_H+1)\text{-R-ISIS}_{q,\text{PK},\beta}}.$$

Proof. Adversary $\mathcal{B}$ receives $h, u_1, \dots, u_{Q_H+1}$, where the u_i 's are independent and uniform in $\mathcal{R}_q$. It simulates G_5 by answering signing queries with the same C_s -capped sampling process and norm test, choosing the retained salt only after acceptance and programming the corresponding point to $\pi_h(X)$. For each fresh adversarial RO query with public key h , it returns the next target u_ℓ and writes in the table. This perfectly simulates G_5 . At most Q_H targets are used during the attack, and one additional target may be used when the forgery point is evaluated.

Suppose that $\mathcal{A}$ outputs a valid fresh forgery $(M^*, (r^*, e^*))$. Since M^* was not signed, the point (h, r^*, M^*) was not programmed by the signing simulation. Hence $c^* = \mathcal{H}(h, r^*, M^*) = u_{j^*}$ for a recorded index j^* . Set $s_2^* = \text{Decode}(e^*)$ and $s_1^* = [u_{j^*} - h s_2^*]_q$. Validity implies $s_1^* + h s_2^* = u_{j^*} \bmod q$ and $\|(s_1^*, s_2^*)\| \leq \beta$, so (j^*, s_1^*, s_2^*) is a valid R-ISIS solution. $\square$

Combining Claims 1–6 gives Eq. (17) under the theorem's explicit transition assumptions. This establishes the conditional statement, not an unconditional reduction.

B Proof of Theorem 4

As in Appendix A, C_s is a proof-only bound and is not part of Sign^{++} . No raw-trial bound L_s is needed.

Game G_0 . This is the real SUF-CMA experiment, and hence

$$\Pr[G_0^{\mathcal{A}} \Rightarrow 1] = \text{Adv}_{\mathcal{A}}^{\text{SUF-CMA}}.$$

Game G_1 . Let $\mathcal{Q}_M$ be the set of messages on which $\mathcal{A}$ queried the signing oracle. If the message in the final forgery (M^*, σ^*) does not belong to $\mathcal{Q}_M$, then it constitutes a UF-CMA forgery. This game is identical to G_0 , except that $\mathcal{A}$ wins only if $M^* \in \mathcal{Q}_M$.

$$\Pr[G_0^{\mathcal{A}} \Rightarrow 1] \leq \text{Adv}_{\mathcal{A}}^{\text{UF-CMA}} + \Pr[G_1^{\mathcal{A}} \Rightarrow 1]. \quad (27)$$

Game G_2 . This game is identical to G_1 , except that its signing oracle is modified using the **Sign** $\rightarrow$ **Prog** $\rightarrow$ **Trans** transformation. Only the random-oracle point corresponding to the signature eventually returned by the signing oracle is programmed. The signing oracle aborts if this point was already defined.

Claim 1.

$$|\Pr[G_1^{\mathcal{A}} \Rightarrow 1] - \Pr[G_2^{\mathcal{A}} \Rightarrow 1]| \leq \left(\frac{Q_s + 1}{2p_{\text{sig}}^2} + \frac{2(Q_H + 1)}{p_{\text{sig}}} \right) Q_s 2^{-\kappa}. \quad (28)$$

Proof. The claim follows from the **Sign** $\rightarrow$ **Prog** $\rightarrow$ **Trans** argument of [5, Thm. 3], as used in [12, Claim 7]. Applying it with acceptance probability p_{sig} , point probability $2^{-\kappa}$, and $Q_H + 1$ random-oracle queries gives Eq. (28). $\square$

Game G_3 . This game is identical to G_2 , except that the raw trials up to the next correction acceptance are replaced by one direct draw $X = (s_1, s_2) \leftarrow \nu_{\mathbf{B},s}$ with $C = \pi_h(X)$. The accepted salt and its existing-point flag, the norm test, and the programming rule remain unchanged.

Claim 2.

$$\Pr[G_2^{\mathcal{A}} \Rightarrow 1] = \Pr[G_3^{\mathcal{A}} \Rightarrow 1]. \quad (29)$$

Proof. The argument is similar to the transition from G_3 to G_4 in Appendix A. $\square$

Game G_4 . This game is identical to G_3 , except that it aborts before drawing more than C_s candidates from $\nu_{\mathbf{B},s}$.

Claim 3.

$$|\Pr[G_3^{\mathcal{A}} \Rightarrow 1] - \Pr[G_4^{\mathcal{A}} \Rightarrow 1]| \leq \sum_{i=0}^{Q_s-1} \binom{C_s}{i} p_{\text{norm}}^i (1 - p_{\text{norm}})^{C_s-i}. \quad (30)$$

Proof. The proof is analogous to that of Claim 2 in Appendix A. $\square$

Game G_5 . This game is identical to G_4 , except that every draw $X \leftarrow \nu_{B,s}$ before the C_s cap is replaced by $X \leftarrow D_{\mathcal{R}^2,s,\mathbf{0}}$. It retains $C = \pi_h(X)$ and all common post-processing.

Claim 4. For every $\alpha \in [2, \infty]$,

$$\Pr[G_4^A \Rightarrow 1] \leq (r_\alpha^{C_s} \Pr[G_5^A \Rightarrow 1])^{(\alpha-1)/\alpha}. \quad (31)$$

Proof. The proof is analogous to that of Claim 5 in Appendix A. $\square$

Let $\mathcal{P}_{\text{ret}}$ be the RO points (r, M) of returned signatures. Note that $M^* \in \mathcal{Q}_M$, but the RO point (r^*, M^*) can be fresh or previously returned. Thus we define two events that partition the winning event of G_5 :

$$E_{\text{point}} = \{(r^*, M^*) \notin \mathcal{P}_{\text{ret}}\}, \quad E_{\text{reuse}} = \{(r^*, M^*) \in \mathcal{P}_{\text{ret}}\}.$$

Hence, on E_{point} , the forgery point was not created by the signing oracle. For E_{point} , $\mathcal{B}_I$ uses the reduction of Claim 6 in Appendix A. It simulates G_5 and assigns the next R-ISIS target to every new RO point not created by the signing oracle. There are at most $Q_H + 1$ such points. A successful forgery gives the short preimage $([c^* - hs_2^*]_q, s_2^*)$ of one of these targets. Hence

$$\Pr[G_5^A \Rightarrow 1 \wedge E_{\text{point}}] \leq \text{Adv}_{\mathcal{B}_I}^{(Q_H+1)\text{-R-ISIS}_{q,\text{PK},\beta}}. \quad (32)$$

Game G_6 . This game is identical to G_5 , except that every pair (C, X) used by the signing oracle is generated as $C \leftarrow \text{Unif}(\mathcal{R}^2/\Lambda_h)$, $X \leftarrow D_{\Lambda_h + \tilde{C}, s, \mathbf{0}}$.

Claim 5.

$$|\Pr[G_5^A \Rightarrow 1 \wedge E_{\text{reuse}}] - \Pr[G_6^A \Rightarrow 1 \wedge E_{\text{reuse}}]| \leq \text{Adv}_{\mathcal{D}}^{\chi\text{-BDD}_{\text{PK}}^{\mathcal{O}_s}[C_s]}. \quad (33)$$

Proof. The distinguisher uses at most C_s instances of $\chi\text{-BDD}_{\text{PK}}^{\mathcal{O}_s}[C_s]$. For each instance, it sets $C = T_i$ and obtains $X = Y_i$ from the exact preimage oracle. When the hidden bit is $b = 1$, exact conditional sampling implies that (C, X) has the joint law $X \leftarrow D_{\mathcal{R}^2,s,\mathbf{0}}$ and $C = \pi_h(X)$, as in G_5 . When $b = 0$, C is uniform in $\mathcal{R}^2/\Lambda_h$ and $X \leftarrow D_{\Lambda_h + \tilde{C}, s, \mathbf{0}}$, as in G_6 . All remaining steps are identical. $\square$

Game G_7 . This game is identical to G_6 , except that the candidate cap C_s is removed. The collision rule is unchanged. Removing the cap can only increase the probability of a successful execution, so $\Pr[G_6^A \Rightarrow 1 \wedge E_{\text{reuse}}] \leq \Pr[G_7^A \Rightarrow 1 \wedge E_{\text{reuse}}]$.

Claim 6.

$$\Pr[G_7^A \Rightarrow 1 \wedge E_{\text{reuse}}] \leq \text{Adv}_{\mathcal{B}_{\text{SP}}}^{Q_s\text{-R-SPISIS}_{q,\text{PK},s,\beta}}. \quad (34)$$

Proof. The R-SPISIS challenger supplies Q_s norm-passing pairs (c_j, X_j) . $\mathcal{B}_{\text{SP}}$ uses the j th pair for the j th signing query, samples one salt r_j , and aborts if (h, r_j, M_j) is already defined. Otherwise it programs this point to c_j and returns the signature obtained from X_j . This simulates G_7 . If the forgery reuses the j th returned point, verification gives $X^* = ([c_j - hs_2^*]_q, s_2^*)$. Since $r^* = r_j$ and $M^* = M_j$, we have $X^* \neq X_j$. Thus (j, X^*) solves the R-SPISIS instance. $\square$

By the two events and Claims 5 and 6,

$$\begin{aligned} \Pr[G_5^{\mathcal{A}} \Rightarrow 1] &\leq \text{Adv}_{\mathcal{B}_I}^{(Q_H+1)\text{-R-ISIS}_{q,\text{PK},\beta}} + \text{Adv}_{\mathcal{D}}^{\chi\text{-BDD}_{\text{PK}}^{\mathcal{O}_s}[C_s]} \\ &\quad + \text{Adv}_{\mathcal{B}_{\text{SP}}}^{Q_s\text{-R-SPISIS}_{q,\text{PK},s,\beta}}. \end{aligned}$$

Combining this inequality with Claims 1–4 gives Eq. (18) under the theorem's explicit transition assumptions. This establishes the conditional statement only. $\square$

C Proofs of the Parameter-Selection Results

This appendix proves the norm-acceptance lemma and the trial-budget results used in Theorems 3 and 4.

C.1 Proof of Lem. 5

Proof (Lem. 5). Let $Z_i \leftarrow \mathcal{D}_{\mathbb{Z},s,0}$ independently for $1 \leq i \leq 2n$, where $s = \sqrt{2\pi} \sigma_{\text{sig}}$. For $0 < t < 1/(2\sigma_{\text{sig}}^2)$, let $s_t^{-2} = s^{-2} - t/\pi$. By Poisson summation, we have

$$\begin{aligned} \mathbb{E}\left[e^{t \sum_{i=1}^{2n} Z_i^2}\right] &= \prod_{i=1}^{2n} \mathbb{E}[e^{t Z_i^2}] = \left(\frac{1}{\rho_s(\mathbb{Z})} \sum_{z \in \mathbb{Z}} e^{t z^2 - \pi z^2 / s^2}\right)^{2n} \\ &= \left(\frac{1}{\rho_s(\mathbb{Z})} \sum_{z \in \mathbb{Z}} e^{-\pi(s^{-2} - t/\pi) z^2}\right)^{2n} \\ &= \left(\frac{\rho_{s_t}(\mathbb{Z})}{\rho_s(\mathbb{Z})}\right)^{2n} = \left(\frac{s_t}{s} \frac{\rho_{1/s_t}(\mathbb{Z})}{\rho_{1/s}(\mathbb{Z})}\right)^{2n} \\ &\leq \left(\frac{s_t}{s}\right)^{2n} = (1 - 2\sigma_{\text{sig}}^2 t)^{-n}. \end{aligned}$$

The inequality follows from $s_t > s$, since $\rho_{1/s_t}(\mathbb{Z}) = \sum_{z \in \mathbb{Z}} e^{-\pi s_t^2 z^2}$ is a decreasing function of s_t . By the Chernoff inequality,

$$\Pr\left[\sum_{i=1}^{2n} Z_i^2 \geq \beta^2 + 1\right] \leq e^{-t(\beta^2+1)} (1 - 2\sigma_{\text{sig}}^2 t)^{-n}.$$

For $u = (\beta^2 + 1)/(2n\sigma_{\text{sig}}^2) > 1$, we choose $t = (1 - u^{-1})/(2\sigma_{\text{sig}}^2)$ to minimize the upper bound and have

$$\Pr\left[\sum_{i=1}^{2n} Z_i^2 \geq \beta^2 + 1\right] \leq \exp[-n(u - 1 - \ln u)].$$

Since $\nu_{\mathbf{B},s}(x) \leq r_{\infty} \mathcal{D}_{\mathcal{R}^2,s,0}(x)$, we have

$$\Pr_{\nu_{\mathbf{B},s}}[\|X\| > \beta] \leq r_{\infty} \Pr\left[\sum_{i=1}^{2n} Z_i^2 \geq \beta^2 + 1\right] \leq r_{\infty} e^{-n(u-1-\ln u)}.$$

$\square$

C.2 Proof of Lem. 6

Proof (Lem. 6). For fixed $p_{\text{norm}} > 0$, the lower tail in Lem. 6 is nonincreasing in C_s and tends to zero. This follows by coupling successive binomials through one additional independent Bernoulli trial, and by the law of large numbers. Thus the exact minimum exists.

For $0 \leq (Q_s - 1)/C_s < p_{\text{norm}} < 1$, the lower-tail Chernoff–Hoeffding inequality gives

$$\begin{aligned} & \sum_{i=0}^{Q_s-1} \binom{C_s}{i} p_{\text{norm}}^i (1 - p_{\text{norm}})^{C_s-i} \\ &= \Pr[\text{Bin}(C_s, p_{\text{norm}}) \leq Q_s - 1] \\ &\leq 2^{-C_s \text{KL}((Q_s-1)/C_s \| p_{\text{norm}})}. \end{aligned}$$

One obtains this form by applying Hoeffding’s upper-tail inequality to the complementary Bernoulli variables [18, Thm. 1, Eq. (2.1)]. Here $\text{KL}(x \| p_{\text{norm}}) = x \log_2(x/p_{\text{norm}}) + (1-x) \log_2((1-x)/(1-p_{\text{norm}}))$. A sufficient strict condition is therefore

$$C_s \text{KL}((Q_s - 1)/C_s \| p_{\text{norm}}) > \lambda, \quad (Q_s - 1)/C_s < p_{\text{norm}}. \quad (35)$$

The minimum passing this sufficient test can exceed the exact binomial minimum. For $p_{\text{norm}} = 1$, $C_s = Q_s$ makes the tail zero. A certified lower bound on p_{norm} may be substituted by monotonicity. When that bound changes with C_s , each proposed cap must be re-evaluated; the fixed-probability coupling does not prove monotonicity of an arbitrary parameter-selection routine. For the selected sets, use the displayed lower bounds in Table 5 and evaluate the KL test at $C_s = 2^{64} + 2^{50} - 1$ for FALCON⁺⁺-512 and $C_s = 2^{64} + 2^{36} - 1$ for FALCON⁺⁺-1024. The exponents exceed 3.76×10^{13} and 1.41×10^{10} , respectively, both greater than λ . This proves the strict upper bounds on the minima. $\square$

C.3 Proof of Lem. 7

Proof (Lem. 7). The same coupling argument, with required successes C_s , shows existence of the exact minimum in Lem. 7. For $0 \leq (C_s - 1)/L_s < p_{\text{corr}} < 1$,

$$\begin{aligned} & \sum_{i=0}^{C_s-1} \binom{L_s}{i} p_{\text{corr}}^i (1 - p_{\text{corr}})^{L_s-i} \\ &= \Pr[\text{Bin}(L_s, p_{\text{corr}}) \leq C_s - 1] \\ &\leq 2^{-L_s \text{KL}((C_s-1)/L_s \| p_{\text{corr}})}. \end{aligned}$$

The same Chernoff–Hoeffding theorem applies with p_{corr} , so

$$L_s \text{KL}((C_s - 1)/L_s \| p_{\text{corr}}) > \lambda, \quad (C_s - 1)/L_s < p_{\text{corr}} \quad (36)$$

is sufficient. At $p_{\text{corr}} = 1$, the minimum is $L_s = C_s$. Strict inequalities are required throughout; equality with the allocated loss is reported as failure. Powers

displayed for readability are rounded upward and rechecked using the full KL expression, not its looser quadratic relaxation. For the selected sets, keep the candidate caps in Lem. 7 fixed and evaluate the test at $L_s = 2^{65} - 1$, using the displayed lower bounds in Table 5. The exponents exceed 7.37×10^{16} for FALCON⁺⁺-512 and 6.06×10^{17} for FALCON⁺⁺-1024, both greater than λ . This proves the strict upper bounds on the minima. $\square$

D Proofs for the Fresh-Syndrome Sampler

D.1 Proof of Theorem 2

Proof (Theorem 2). Set $u_i := s/\|\tilde{\mathbf{b}}_i\|$. Combining Eqs. (4) and (3), and substituting Eq. (5), gives directly

$$\mathbb{E}_{Q_{\mathbf{B},s}} \left[\left(\frac{\Delta_{\mathbf{B}}}{\delta_{\mathbf{B},s}} \right)^k \right] = \frac{\det(\Lambda)^{k-1}}{\rho_s(\mathbb{Z}^d)^k} \sum_{\mathbf{x} \in \mathbb{Z}^d} F_k(\mathbf{x}), \quad (37)$$

where $F_k(\mathbf{x}) := \rho_s(\mathbf{x}) \prod_{i=0}^{d-1} \rho_{u_i, \gamma_i(\mathbf{x})}(\mathbb{Z})^{k-1}$.

It therefore suffices to control this lattice sum. For $u > 0$, define

$$\phi_{u,k}(t) := e^{-\pi t^2/u^2} \rho_{u,t}(\mathbb{Z})^{k-1}.$$

Let $\mathbf{B} = \mathbf{Q}\mathbf{R}$ be the QR decomposition, so the column of $\mathbf{Q}$ indexed by i is $\tilde{\mathbf{b}}_i/\|\tilde{\mathbf{b}}_i\|$, and put $\mathbf{y} = \mathbf{Q}^\top \mathbf{x}$. Then $y_i = \langle \mathbf{x}, \tilde{\mathbf{b}}_i \rangle / \|\tilde{\mathbf{b}}_i\|$ and $\gamma_i(\mathbf{x}) = y_i / \|\tilde{\mathbf{b}}_i\|$. Moreover, $\|\mathbf{x}\|^2 = \sum_i y_i^2$ because $\mathbf{Q}$ is orthogonal. Together with $s = u_i \|\tilde{\mathbf{b}}_i\|$, we have

$$\begin{aligned} F_k(\mathbf{x}) &= \rho_s(\mathbf{x}) \prod_{i=0}^{d-1} \rho_{u_i, \gamma_i(\mathbf{x})}(\mathbb{Z})^{k-1} = \prod_{i=0}^{d-1} e^{-\pi y_i^2/s^2} \rho_{u_i, y_i/\|\tilde{\mathbf{b}}_i\|}(\mathbb{Z})^{k-1} \\ &= \prod_{i=0}^{d-1} e^{-\pi (y_i/\|\tilde{\mathbf{b}}_i\|)^2/u_i^2} \rho_{u_i, y_i/\|\tilde{\mathbf{b}}_i\|}(\mathbb{Z})^{k-1} = \prod_{i=0}^{d-1} \phi_{u_i,k} \left(\frac{y_i}{\|\tilde{\mathbf{b}}_i\|} \right). \end{aligned}$$

We next establish the two one-dimensional facts used below. For $\mathbf{m} \in \mathbb{Z}^{k-1}$, write $\mathbf{m} = (m_1, \dots, m_{k-1})$. Completing the square gives

$$t^2 + \sum_{j=1}^{k-1} (m_j - t)^2 = k \left(t - \frac{\langle \mathbf{1}, \mathbf{m} \rangle}{k} \right)^2 + \|\mathbf{m}\|^2 - \frac{\langle \mathbf{1}, \mathbf{m} \rangle^2}{k}.$$

Expanding the shifted Gaussian masses and substituting $z = t - \langle \mathbf{1}, \mathbf{m} \rangle/k$ in the Fourier transform therefore gives

$$\begin{aligned} \hat{\phi}_{u,k}(\omega) &= \sum_{\mathbf{m} \in \mathbb{Z}^{k-1}} \int_{\mathbb{R}} \exp \left[-\frac{\pi}{u^2} \left(t^2 + \sum_{j=1}^{k-1} (m_j - t)^2 \right) \right] e^{-2\pi i \omega t} dt \\ &= \frac{u}{\sqrt{k}} e^{-\pi u^2 \omega^2/k} \sum_{\mathbf{m} \in \mathbb{Z}^{k-1}} \exp \left[-\frac{\pi}{u^2} \left(\|\mathbf{m}\|^2 - \frac{\langle \mathbf{1}, \mathbf{m} \rangle^2}{k} \right) \right] e^{-2\pi i \omega \langle \mathbf{1}, \mathbf{m} \rangle/k}. \end{aligned}$$

The coefficients in the final sum are nonnegative. Hence

$$|\hat{\phi}_{u,k}(\omega)| \leq \hat{\phi}_{u,k}(0) e^{-\pi u^2 \omega^2 / k}. \quad (38)$$

We also have

$$\hat{\phi}_{u,k}(0) = \sum_{n \in \mathbb{Z}} \int_0^1 e^{-\pi(v+n)^2/u^2} \rho_{u,v+n}(\mathbb{Z})^{k-1} dv = \int_0^1 \rho_{u,v}(\mathbb{Z})^k dv. \quad (39)$$

We now lift these facts to d dimensions as follows. Noting $\mathbf{x} = \mathbf{Q}\mathbf{y}$, we have

$$\begin{aligned} \hat{F}_k(\mathbf{w}) &= \int_{\mathbb{R}^d} \prod_{i=0}^{d-1} \phi_{u_i,k} \left(\frac{y_i}{\|\tilde{\mathbf{b}}_i\|} \right) e^{-2\pi i \langle \mathbf{Q}^\top \mathbf{w}, \mathbf{y} \rangle} d\mathbf{y} \\ &= \prod_{i=0}^{d-1} \int_{\mathbb{R}} \phi_{u_i,k} \left(\frac{y_i}{\|\tilde{\mathbf{b}}_i\|} \right) e^{-2\pi i (\mathbf{Q}^\top \mathbf{w})_i y_i} dy_i \\ &= \prod_{i=0}^{d-1} \|\tilde{\mathbf{b}}_i\| \hat{\phi}_{u_i,k} \left(\|\tilde{\mathbf{b}}_i\| (\mathbf{Q}^\top \mathbf{w})_i \right). \end{aligned}$$

In particular, Eq. (39) and $\prod_i \|\tilde{\mathbf{b}}_i\| = \det(\Lambda)$ imply

$$\hat{F}_k(\mathbf{0}) = \det(\Lambda) \prod_{i=0}^{d-1} \int_0^1 \rho_{u_i,t}(\mathbb{Z})^k dt. \quad (40)$$

Applying Eq. (38) to every factor and using $u_i \|\tilde{\mathbf{b}}_i\| = s$ gives

$$\begin{aligned} |\hat{F}_k(\mathbf{w})| &\leq \hat{F}_k(\mathbf{0}) \exp \left(-\frac{\pi s^2}{k} \sum_{i=0}^{d-1} (\mathbf{Q}^\top \mathbf{w})_i^2 \right) \\ &= \hat{F}_k(\mathbf{0}) e^{-\pi s^2 \|\mathbf{w}\|^2 / k}, \end{aligned}$$

where the last equality again uses the orthogonality of $\mathbf{Q}$. Since F_k is a Schwartz function, Poisson summation over $\mathbb{Z}^d$ and the preceding bound give

$$\begin{aligned} \sum_{\mathbf{x} \in \mathbb{Z}^d} F_k(\mathbf{x}) &= \sum_{\mathbf{w} \in \mathbb{Z}^d} \hat{F}_k(\mathbf{w}) \leq \sum_{\mathbf{w} \in \mathbb{Z}^d} |\hat{F}_k(\mathbf{w})| \\ &\leq \hat{F}_k(\mathbf{0}) \sum_{\mathbf{w} \in \mathbb{Z}^d} e^{-\pi s^2 \|\mathbf{w}\|^2 / k} = \hat{F}_k(\mathbf{0}) \vartheta_3(s^2/k)^d. \end{aligned} \quad (41)$$

Finally, substitute Eqs. (40) and (41) into Eq. (37). Using $\rho_s(\mathbb{Z}^d) = s^d \vartheta_3(s^2)^d$, $\det(\Lambda) \prod_i u_i = s^d$, and $\rho_{u_i,t}(\mathbb{Z}) = u_i g_{u_i,\mathbb{Z}}(t)$, we obtain

$$\begin{aligned} \mathbb{E}_{Q_{\mathbf{B},s}} \left[\left(\frac{\Delta_{\mathbf{B}}}{\delta_{\mathbf{B},s}} \right)^k \right] &\leq \frac{\vartheta_3(s^2/k)^d \det(\Lambda)^k}{\rho_s(\mathbb{Z}^d)^k} \prod_{i=0}^{d-1} \int_0^1 \rho_{u_i,t}(\mathbb{Z})^k dt \\ &= \frac{\vartheta_3(s^2/k)^d}{\vartheta_3(s^2)^{dk}} \prod_{i=0}^{d-1} \int_0^1 g_{u_i,\mathbb{Z}}(t)^k dt \\ &= \mathcal{M}_k(\mathbf{B}, s). \end{aligned}$$

□

D.2 Proof of Lem. 4

Proof (Lem. 4). For convenience, denote the finite average in the statement by

$$\mathcal{S}_{R,k,M}(u) := \frac{1}{R} \sum_{j=0}^{R-1} \left(1 + 2 \sum_{m=1}^M e^{-\pi u^2 m^2} \cos \frac{2\pi m j}{R} \right)^k.$$

Put

$$\begin{aligned} \mathcal{G}_u(z) &:= \sum_{m \in \mathbb{Z}} e^{-\pi u^2 m^2} z^m, \\ \mathcal{G}_{u,M}(z) &:= \sum_{|m| \leq M} e^{-\pi u^2 m^2} z^m, \\ \Theta_{A_{k-1}}^{(\leq M)}(u^2) &:= \sum_{\substack{\mathbf{m} \in A_{k-1} \\ \|\mathbf{m}\|_\infty \leq M}} e^{-\pi u^2 \|\mathbf{m}\|^2} = [z^0] \mathcal{G}_{u,M}(z)^k. \end{aligned}$$

Here $\mathcal{G}_u(z)$ is a Laurent series, and $\mathcal{G}_{u,M}(z)$ is its finite truncation. For a Laurent series $H(z)$, the notation $[z^0]H(z)$ is its constant coefficient.

It suffices to prove

$$(1 - \varepsilon_{k,M}(u)) \Theta_{A_{k-1}}(u^2) \leq \Theta_{A_{k-1}}^{(\leq M)}(u^2) \leq \mathcal{S}_{R,k,M}(u).$$

We first prove the right-hand inequality. Let $\zeta_R = e^{2\pi i/R}$. For every $0 \leq j < R$,

$$\begin{aligned} \mathcal{G}_{u,M}(\zeta_R^j) &= \sum_{m=-M}^M e^{-\pi u^2 m^2} (\zeta_R^j)^m \\ &= 1 + \sum_{m=1}^M e^{-\pi u^2 m^2} \left(e^{2\pi i m j/R} + e^{-2\pi i m j/R} \right) \\ &= 1 + 2 \sum_{m=1}^M e^{-\pi u^2 m^2} \cos \frac{2\pi m j}{R}. \end{aligned}$$

Thus $\mathcal{G}_{u,M}(\zeta_R^j)$ is exactly the cosine sum defining $\mathcal{S}_{R,k,M}(u)$. Moreover, for every $n \in \mathbb{Z}$, the root-of-unity filter satisfies

$$\frac{1}{R} \sum_{j=0}^{R-1} \zeta_R^{jn} = \begin{cases} 1, & R \mid n, \\ 0, & R \nmid n. \end{cases}$$

Write

$$\mathcal{G}_{u,M}(z)^k = \sum_{n=-kM}^{kM} c_n z^n, \quad c_n := [z^n] \mathcal{G}_{u,M}(z)^k.$$

Every c_n is nonnegative. Applying the root-of-unity filter coefficientwise therefore gives

$$\begin{aligned}
\mathcal{S}_{R,k,M}(u) &= \frac{1}{R} \sum_{j=0}^{R-1} \mathcal{G}_{u,M}(\zeta_R^j)^k = \frac{1}{R} \sum_{j=0}^{R-1} \sum_{n=-kM}^{kM} c_n \zeta_R^{jn} \\
&= \sum_{n=-kM}^{kM} c_n \left(\frac{1}{R} \sum_{j=0}^{R-1} \zeta_R^{jn} \right) = \sum_{\substack{-kM \leq n \leq kM \\ R|n}} c_n \\
&= \sum_{\ell \in \mathbb{Z}} [z^{\ell R}] \mathcal{G}_{u,M}(z)^k \geq [z^0] \mathcal{G}_{u,M}(z)^k = \Theta_{A_{k-1}}^{(\leq M)}(u^2).
\end{aligned}$$

For the left-hand inequality, consider the probability distribution on A_{k-1} defined by

$$\mu_{u,k}(\mathbf{m}) := \frac{e^{-\pi u^2 \|\mathbf{m}\|^2}}{\Theta_{A_{k-1}}(u^2)}, \quad \mathbf{m} \in A_{k-1}.$$

By definition, $\Theta_{A_{k-1}}(u^2)$ is the normalizing constant, while its truncated sum gives

$$\frac{\Theta_{A_{k-1}}^{(\leq M)}(u^2)}{\Theta_{A_{k-1}}(u^2)} = \Pr_{\mathbf{m} \leftarrow \mu_{u,k}} [\|\mathbf{m}\|_\infty \leq M].$$

To bound the complementary event, let

$$H := \text{span}(A_{k-1}) = \left\{ \mathbf{x} \in \mathbb{R}^k : \sum_{j=1}^k x_j = 0 \right\}, \quad \mathbf{p}_i := \mathbf{e}_i - \frac{1}{k} \mathbf{1}.$$

Since $\mu_{u,k} = D_{A_{k-1}, 1/u}$ is 0-subgaussian with parameter $1/u$ [25, Sec. 2.4 and Lem. 2.8], every nonzero $\mathbf{v} \in H$ satisfies

$$\Pr_{\mathbf{m} \leftarrow \mu_{u,k}} [|\langle \mathbf{m}, \mathbf{v} \rangle| \geq t] \leq 2 \exp\left(-\frac{\pi u^2 t^2}{\|\mathbf{v}\|^2}\right).$$

Because $\mathbf{m} \in H$, we have $m_i = \langle \mathbf{m}, \mathbf{p}_i \rangle$ and $\|\mathbf{p}_i\|^2 = (k-1)/k$. Hence, for every $t \geq 0$,

$$\Pr_{\mathbf{m} \leftarrow \mu_{u,k}} [|m_i| \geq t] \leq 2 \exp\left(-\pi u^2 \frac{k}{k-1} t^2\right).$$

Because each m_i is integral, a union bound over the k coordinates now gives

$$\begin{aligned}
1 - \frac{\Theta_{A_{k-1}}^{(\leq M)}(u^2)}{\Theta_{A_{k-1}}(u^2)} &= \Pr_{\mathbf{m} \leftarrow \mu_{u,k}} [\|\mathbf{m}\|_\infty > M] \\
&\leq \sum_{i=1}^k \Pr_{\mathbf{m} \leftarrow \mu_{u,k}} [|m_i| \geq M+1] \\
&\leq 2k \exp\left(-\pi u^2 \frac{k}{k-1} (M+1)^2\right) \\
&= \varepsilon_{k,M}(u),
\end{aligned}$$

Thus $(1 - \varepsilon_{k,M}(u))\Theta_{A_{k-1}}(u^2) \leq \Theta_{A_{k-1}}^{(\leq M)}(u^2)$. Combining this with the root-of-unity bound and dividing by $1 - \varepsilon_{k,M}(u) > 0$ proves Eq. (12). $\square$