

ANSA-IBS: Identity-Based Signatures from Annular NTRU Trapdoors and Bimodal Fiat-Shamir with Aborts^{*}

Zhaohui Cheng^{1,2**}, Hengfu Yang², and Kaixin Xiong¹

¹ Shenzhen Olym Info. Sec. Ltd.
chengzh@myibc.net, xiongkx@myibc.net

² Hunan First Normal University
hengfuyang@hnfnu.edu.cn

Abstract. We give a lattice identity-based signature scheme from annular NTRU trapdoors and Fiat-Shamir with aborts (FSwA). Counter-based hashing maps each identity to its first invertible NTRU syndrome a . Extraction returns a short witness to $s_0 + s_1 h = a \bmod q$, and signing normalizes this relation to a fixed-target two-response equation with public commitment recovery. Gärtners iterative rejection method fits naturally into this signing layer: processing the challenge monomials separately reduces the Gaussian mask width and expected signing attempts, yielding shorter signatures. Under heuristic direct-attack estimates, our three main parameter sets target NIST security categories 2, 3, and 5, with fixed signature sizes of 8488, 11309, and 12734 bytes, respectively; we also report a lighter profile.

Keywords: identity-based signatures · lattice cryptography · Fiat-Shamir with aborts · bimodal signatures · NTRU trapdoors

1 Introduction

Identity-based signatures (IBS) let a verifier check a signature using the master public key and the signer identity, while a private-key generator issues the corresponding identity signing key. In settings where the generator is already part of the trust model, this removes the need for a verifier to obtain and validate a separate certificate for an identity-specific public key.

Bellare, Namprempre, and Neven [3] identify two generic routes to IBS: certificate-based constructions and transformations from convertible standard

^{*} This is the full version of a paper presented at CANS 2026. It includes retuned parameters and the integer-only fixed-point realization of Setup, Hybrid-Sampler extraction, and signing (Appendix B). Sections B.1–B.4 apply equally to ANSA-IBKEM, which uses the same annular trapdoor generator and Hybrid-Sampler interface. Appendix C specifies Falcon-style key generation and extraction with unchanged user-side signing and verification.

^{**} Corresponding author.

identification or signature schemes. They also give BNN-IBS as a dedicated exception based on an extended Fiat-Shamir transform for canonical identity-based identification. In lattice settings, the certificate route can use Falcon/Antrag-based message-recovery signatures for compact IBS [30,14,17,5], while [22] mirrors BNN-IBS by concatenating Dilithium-style credential/message signatures [11].

For the convertible-identification approach, a natural lattice instantiation starts from the Ducas-Lyubashevsky-Prest (DLP) NTRU relation [12]: the master public key is an NTRU public key h , an identity is mapped to a syndrome $a \in R_q$, and extraction returns (s_0, s_1) with $s_0 + s_1 h = a \bmod q$. A direct Lyubashevsky/Dilithium-style Fiat-Shamir-with-aborts (FSwA) layer [23,24,11,26] would sign with two responses $(y_0 + cs_0, y_1 + cs_1)$, as in NTRU identity-based signatures such as [32]. The difficulty is quantitative: extraction produces wider witnesses than Dilithium-style module-signature secrets, forcing large moduli and response bounds for rejection sampling and simulation. Even with response compression, this direct route gives much larger signatures.

Our signing layer builds on bimodal rejection sampling. BLISS introduced bimodal Gaussian responses for hiding secret-dependent shifts [9]; HAETAE gives a modern FSwA analysis [6], and NTRU+Sign applies the idea to an NTRU relation with recoverable hinted commitments [31]. In its usual single-step form, bimodal rejection must accommodate the full secret-dependent shift $c(s_0, s_1)$. G rtner retains its two signed branches but applies them sequentially to the nonzero monomials of a sparse challenge, coupling each sign choice with the corresponding rejection decision [16]. In ANSA-IBS, each step therefore handles $X^j(s_0, s_1)$ rather than $c(s_0, s_1)$, permitting a narrower Gaussian mask, fewer expected signing attempts, and smaller signatures. The one-response hint of NTRU+Sign does not apply directly to this normalized two-response representation because normalization multiplies the omitted response by an identity-dependent unit. ANSA-IBS therefore retains both response polynomials and combines iterative rejection with public commitment recovery.

The main technical step is to normalize the DLP identity relation. We require the identity map to output a unit label $a \in R_q^\times$ and, with $\hat{q} = 2^{-1} \bmod q$, rewrite $s_0 + hs_1 = a \bmod q$ as $A_0 s_0 + A_1 s_1 = \hat{q} \bmod q$, where $(A_0, A_1) = (\hat{q}a^{-1}, \hat{q}ha^{-1})$. The identity dependence moves into the public coefficients, while the target becomes the fixed scalar $\hat{q}$. Unit labels are obtained by counter-based hashing: the identity map hashes (h, id, j) until the candidate is invertible, and the signature carries the counter so that verification recomputes the first valid label. If the irreducible factors of $\phi \bmod q$ have degrees f_i , the expected number of trials is $\prod_i (1 - q^{-f_i})^{-1}$; it is below 1.000245 for our parameters.

For efficient user-key extraction, ANSA-IBS follows ANSA-IBKEM [5]: it combines Prest’s Hybrid-Sampling [28] with Antrag annular NTRU trapdoor generation [14]. A user key remains a DLP-style preimage of the identity syndrome. To sign, the user samples $y = (y_0, y_1)$, commits to $[A_0 y_0 + A_1 y_1]_{d_h}$, and derives a fixed-weight binary challenge. The selected branch signs form a ternary lift c' with $c' \equiv c \bmod 2$. For the response $z = y + c'S$, where S is the user key,

and $u = A_0y_0 + A_1y_1$, this gives $A_0z_0 + A_1z_1 + c\hat{q} = u + d$ with a binary d ; a public filter on the high bits then guarantees that the verifier recovers the commitment $[u]_{d_n}$.

The security analysis separates the adaptive identity layer from the fixed-statement signing layer. The identity layer follows a witness-first path: deterministic extraction tapes are replaced by independent random tapes, Hybrid-Sampler witnesses are compared with ideal conditional witnesses through Rényi divergence, unit labels are programmed from witnesses, and the forgery identity is localized to one normalized signing statement. In the underlying public-coin protocol, the ideal iterative sampler, conditioned on success, yields an exact centered-Gaussian response independent of the witness. A maximum point-probability bound then controls the FS_WA loss [2,8], and a single no-message forgery yields the bimodal self-target RSIS (BMST-RSIS) relation used in different forms by HAE_TAE and NTRU+Sign [6,31]. A fork satisfying the additional no-wrap conditions yields an auxiliary ordinary-RSIS solution, whose concrete estimates are weaker at the selected modulus (Section 5.6).

ANSA-IBS shares ANSA-IBKEM’s annular-trapdoor design and ML-DSA modulus $q = 8,380,417$, allowing reuse of its modular arithmetic [5].

In this work, we make three contributions. First, we realize the convertible-identification route of Bellare, Namprempre, and Neven for the DLP/NTRU relation through a normalized two-response protocol with public commitment recovery and iterative bimodal rejection. Second, we give a loss-explicit adaptive-identity reduction and concrete parameters that separately account for the lattice security and challenge-space entropy of the self-target signing relation. Third, we give an integer-only fixed-point realization of Setup, Hybrid-Sampler extraction, and signing (Appendix B). The Setup and extraction comparison also applies to ANSA-IBKEM instantiated with the same generator policy and Hybrid-Sampler interface.

2 Preliminaries

We use standard annular-trapdoor, Hybrid-Sampling, and Rényi conventions, and recall only the interfaces needed here.

Throughout, $x \xleftarrow{\$} S$ denotes uniform sampling from a finite set, and $x \leftarrow D$ denotes sampling from a named distribution or randomized algorithm. An explicit randomness argument, as in $A(\cdot; rnd)$, makes the output deterministic once the inputs and rnd are fixed.

2.1 Algebraic and Probabilistic Notation

Let $R = \mathbb{Z}[X]/\langle\phi(X)\rangle$ be the degree- N NTRU ring and let $R_q = R/qR$ for an odd prime q . We work throughout with cyclotomic rings: $\phi = \Phi_M$ and $R = \mathcal{O}_{\mathcal{K}} = \mathbb{Z}[\zeta_M]$, where $\mathcal{K} = \mathbb{Q}(\zeta_M)$. The power-of-two rings have $M = 2N$; the trinomial rings have $M = 3N$ and $\Phi_M(X) = X^N - X^{N/2} + 1$. We require ϕ to be squarefree over $\mathbb{F}_q$; its factorization determines the Chinese remainder

theorem (CRT) representation and the unit probability used below. Unless stated otherwise, elements of R_q are represented coefficientwise in $[-(q-1)/2, (q-1)/2]$. We write $\text{center_mod}_q(\cdot)$ for this map, $\text{can}_q : R \rightarrow R_q$ for coefficientwise reduction followed by centered representation, and $R_q^\times$ for the unit group.

Let $\phi_1, \dots, \phi_N : \mathcal{K} \hookrightarrow \mathbb{C}$ be the complex embeddings, and let $\text{coef}(x) \in \mathbb{R}^N$ denote the power-basis coefficient vector of x , extended componentwise to tuples over R . We equip R with the normalized canonical inner product

$$\langle x, y \rangle_{\text{can}} = \frac{1}{N} \sum_{i=1}^N \phi_i(x) \overline{\phi_i(y)}, \quad \|x\|_{\text{can}}^2 = \langle x, x \rangle_{\text{can}}.$$

Let H_N be the Gram matrix of this inner product in the power basis. For a power-of-two ring, $H_N = I_N$. For a trinomial ring, write $N = 2m$ and pair coordinates $(i, i+m)$; then H_N is block diagonal with m copies of

$$G = \begin{bmatrix} 1 & 1/2 \\ 1/2 & 1 \end{bmatrix}.$$

Fix W_N such that $W_N^T W_N = H_N$ and set $K_N = H_N^{-1}$. Thus $\|x\|_{\text{can}} = \|W_N \text{coef}(x)\|_2$.

For $x \in R$, $\|x\|_2$ and $\|x\|_\infty$ are the Euclidean and maximum norms of its coefficient vector; for pairs in R^2 , norms use concatenated coefficients. For a challenge $c \in \{0, 1\}^N$, $\|c\|_0$ is its Hamming weight and $\|c\|_2 = \sqrt{\|c\|_0}$.

For $\sigma > 0$ and center c , $D_{\mathbb{Z}, \sigma, c}$ assigns to $x \in \mathbb{Z}$ probability proportional to $\rho_{\sigma, c}(x) = \exp(-(x-c)^2/(2\sigma^2))$; we omit c when it is zero. Let $D_{R, \sigma}$ be the coefficientwise product of N copies of $D_{\mathbb{Z}, \sigma}$; this is the response distribution used by the signing algorithm. Write D_σ^{2N} for $D_{R, \sigma}^2$, viewed as a distribution on coefficient vectors in $\mathbb{Z}^{2N}$. For a power-of-two ring, the ambient witness distribution is $D_{R^2, \sigma}^{\text{amb}} = D_{R, \sigma} \times D_{R, \sigma}$. For a trinomial ring, $D_{R^2, \sigma}^{\text{amb}}$ is the centered canonical discrete Gaussian on R^2 , with mass at $S = (s_0, s_1)$ proportional to $\exp(-\|S\|_{\text{can}}^2/(2\sigma^2))$, where $\|S\|_{\text{can}}^2 = \|s_0\|_{\text{can}}^2 + \|s_1\|_{\text{can}}^2$. Equivalently, $W_N^{\oplus 2} \text{coef}$ maps this distribution to the spherical discrete Gaussian of width σ on $W_N^{\oplus 2} \mathbb{Z}^{2N}$. In coefficient coordinates, each paired block of its continuous covariance $\sigma^2 K_N$ is

$$\sigma^2 G^{-1} = \frac{4}{3} \sigma^2 \begin{bmatrix} 1 & -1/2 \\ -1/2 & 1 \end{bmatrix};$$

the corresponding continuous-Gaussian coefficient marginals have standard deviation $\sqrt{4/3} \sigma$, and paired coefficients have correlation $-1/2$.

For a Euclidean lattice L , with L^* its Euclidean dual in its real span, define the standard-deviation smoothing parameter by

$$\eta_\varepsilon(L) := \inf \left\{ \sigma > 0 : \sum_{0 \neq y \in L^*} e^{-2\pi^2 \sigma^2 \|y\|_2^2} \leq \varepsilon \right\}.$$

For an algebraic lattice $\Lambda \subseteq R^r$, $\eta_\varepsilon(\Lambda)$ denotes the smoothing parameter of its normalized canonical image under $W_N^{\oplus r}$ coef. This agrees with Prest’s standard-deviation convention [28] after rescaling the trace metric and every Gaussian width by $1/\sqrt{N}$, a common rescaling that leaves the sampler theorem invariant. We use the same convention for the Micciancio–Regev smoothing bound [25] and its Gram–Schmidt refinement [15], as well as for the quotient-Gaussian and Rényi comparisons below.

2.2 NTRU Relation and Extraction Interface

The master public key is an NTRU public key $h \in R_q$. A classical NTRU trapdoor consists of short polynomials $(f, g, F, G) \in R^4$ satisfying $fG - gF = q$. When f is invertible in R_q , the public key is $h = gf^{-1} \bmod q$, and it defines the rank-two NTRU lattice $\Lambda_{h,q} := \Lambda_h = \{(u, v) \in R^2 : u + vh = 0 \bmod q\}$. The public key h itself need not be a unit. Smoothing expressions involving $\Lambda_{h,q}$ use its normalized canonical image from Section 2.1. With the basis convention used in this paper, the corresponding trapdoor basis is written

$$B = \begin{bmatrix} g & -f \\ G & -F \end{bmatrix}.$$

Changing signs or transposing this basis only changes conventions; compatibility with h means that the rows of B form an R -basis of $\Lambda_{h,q}$.

The scheme uses Antrag annular NTRU trapdoor generation [14] through $(B, h) \leftarrow \text{AnnularTrapdoorGen}(N, q, \alpha_{\text{ann}}, \xi)$, where α_{ann} is the accepted quality target and ξ selects the middle fraction of the pre-rounding annulus. For complex embeddings ϕ_i , set $s_i = |\phi_i(f)|^2 + |\phi_i(g)|^2$. Antrag enforces $q/\alpha_{\text{ann}}^2 \leq s_i \leq \alpha_{\text{ann}}^2 q$, giving $Q_{\text{hyb}}(B) := \max_i \max(\sqrt{s_i}, q/\sqrt{s_i}) \leq \alpha_{\text{ann}} \sqrt{q}$. The generator includes the finite candidate evaluation, cutoff rejection, and deterministic NTRU-solver policy of Sections B.1–B.3, together with the second-row norm test. Its public-key law conditioned on success before Galois screening is denoted $\mathcal{H}_{\text{ann}}$ (Assumption 4.1). For the smoothing and Hybrid-Sampler checks, set

$$\rho_{\text{hyb}}(N, \varepsilon) = \frac{1}{\pi} \sqrt{\frac{\ln(2N(1 + 1/\varepsilon))}{2}}, \quad \rho_u(N, \varepsilon) = \frac{1}{\pi} \sqrt{\frac{\ln(4N(1 + 1/\varepsilon))}{2}}.$$

Lemma 2.1 (Canonical NTRU smoothing specialization). For every compatible rank-two basis B and every $\varepsilon \in (0, 1)$,

$$\eta_\varepsilon(R) \leq \rho_{\text{hyb}}(N, \varepsilon), \quad \eta_\varepsilon(\Lambda_{h,q}) \leq \rho_u(N, \varepsilon) Q_{\text{hyb}}(B).$$

Consequently, for $0 < \varepsilon_p < 1/12$, the condition $\sigma_1 \geq \rho_{\text{hyb}}(N, \varepsilon_p) Q_{\text{hyb}}(B)$ suffices for Prest’s rank-two Hybrid-Sampler comparison [28, Theorem 5.10], whose restriction $\varepsilon < 1/(6r)$ specializes to $1/12$ at module rank $r = 2$. The condition $\sigma_1 \geq \rho_u(N, \varepsilon_u) Q_{\text{hyb}}(B)$ likewise suffices for quotient smoothing over $\Lambda_{h,q}$. A proof covering both ring families appears in Appendix A.1.

As in the DLP-style NTRU IBE relation [12], an identity maps to a public syndrome $a \in R_q$, here conditioned on $a \in R_q^\times$. An identity signing key is

a short witness $(s_0, s_1) \in P_h(a) := \{(s_0, s_1) \in R^2 : s_0 + s_1 h = a \bmod q\}$. Given B, h, a , Prest's Hybrid-Sampling kernel [28] deterministically produces a raw witness (s_0, s_1) from $(B, h, a, \sigma_1; rnd)$. Let $\mathcal{C}_q$ reduce every coefficient to its unique centered representative modulo q :

$$\begin{aligned} \mathcal{C}_q : R^2 &\longrightarrow R^2, \\ \mathcal{C}_q(s_0, s_1) &= (\tilde{s}_0, \tilde{s}_1), \\ \tilde{s}_{k,i} &\equiv s_{k,i} \pmod{q}, \quad \tilde{s}_{k,i} \in [-(q-1)/2, (q-1)/2] \cap \mathbb{Z}, \end{aligned}$$

for $k \in \{0, 1\}$ and $0 \leq i < N$. We write $\text{SamplePreimage}_P^{\text{exact}}$ for the raw kernel followed by $\mathcal{C}_q$. Appendix B specifies its fixed-point realization, including an additional admissibility gate. The extraction width is calibrated as

$$\sigma_1 = R_{\text{ext}} \alpha_{\text{ann}} \sqrt{q}, \quad (1)$$

where α_{ann} and R_{ext} are specified in Section 5. The parameter σ_1 is the width of the ambient witness distribution $D_{R^2, \sigma_1}^{\text{amb}}$: coefficient-spherical for the power-of-two profiles and canonical for the trinomial profiles. Let $D_{\text{hyb}}(\cdot \mid B, h, a)$ denote the raw Hybrid-Sampler distribution, and let $D_{\text{pub}}(\cdot \mid h, a)$ denote $D_{R^2, \sigma_1}^{\text{amb}}$ conditioned on $P_h(a)$. Define their centered pushforwards by

$$D_{\text{hyb}}^{\text{ctr}} := (\mathcal{C}_q)_* D_{\text{hyb}}, \quad D_{\text{pub}}^{\text{ctr}} := (\mathcal{C}_q)_* D_{\text{pub}}.$$

We model the sampler's domain-separated extendable-output function (XOF) as an ideal random function that assigns an independent lazy uniform bit tape to each distinct input transcript and reuses that tape when the transcript repeats. The exact Gaussian subroutines consume an almost surely finite prefix and induce $D_{\text{hyb}}(\cdot \mid B, h, a)$; applying $\mathcal{C}_q$ induces $D_{\text{hyb}}^{\text{ctr}}(\cdot \mid B, h, a)$.

Because a is a unit, the DLP relation normalizes to a statement whose first coefficient is a unit: with $\hat{q} = 2^{-1} \bmod q$ and $A = (A_0, A_1) = (\hat{q}a^{-1}, \hat{q}ha^{-1}) \in R_q^\times \times R_q$, the signing relation is $AS = A_0 s_0 + A_1 s_1 = \hat{q} \pmod{q}$ for $S = (s_0, s_1)^T$.

2.3 Signature Security Definitions

Definition 2.1 (Identity-based signature scheme). An identity-based signature scheme is a tuple $\text{IBS} = (\text{Setup}, \text{Extract}, \text{Sign}, \text{Verify})$.

1. $\text{Setup}(1^\lambda)$ outputs public parameters pp , a master public key mpk , and a master secret key msk , or $\perp$.
2. $\text{Extract}(\text{pp}, \text{msk}, \text{mpk}, \text{id})$ outputs an identity signing key usk_{id} , or $\perp$.
3. $\text{Sign}(\text{pp}, \text{mpk}, \text{id}, \text{msg}, \text{usk}_{\text{id}})$ outputs a signature sig , or $\perp$.
4. $\text{Verify}(\text{pp}, \text{mpk}, \text{id}, \text{msg}, \text{sig})$ outputs 1 or 0.

Definition 2.2 (Adaptive identity-based unforgeability). The adaptive identity-based existential-unforgeability experiment $\text{ID-UF-CMA}_{\text{IBS}}(\mathcal{A})$ is defined as follows. The challenger runs $\text{Setup}(1^\lambda)$, returns 0 on failure, and otherwise

gives (pp, mpk) to $\mathcal{A}$. Identities are arbitrary byte strings within the scheme's encoding length bound (Section 3.1). Concrete bounds are stated for adversaries that reach a bounded number of distinct identity-label states (Section 4.1). The adversary has access to an extraction oracle $\mathcal{O}_{\text{ext}}(\text{id}) := \text{Extract}(\text{pp}, \text{msk}, \text{mpk}, \text{id})$ that extracts and caches the key for id if necessary, and to a signing oracle $\mathcal{O}_{\text{sig}}(\text{id}, \text{msg}) := \text{Sign}(\text{pp}, \text{mpk}, \text{id}, \text{msg}, \text{usk}_{\text{id}})$. If a signing query is made for an identity whose key is not cached, the signing oracle first extracts and caches usk_{id} internally; this implicit signing-key generation does not make id non-fresh unless id was explicitly queried to $\mathcal{O}_{\text{ext}}$. An extraction failure is cached as $\perp$ for that identity; subsequent extraction or signing queries for it return $\perp$. A signing failure returns $\perp$ without changing the cached key. Explicit extraction queries and signing-query pairs are recorded even when they fail. The adversary outputs $(\text{id}^*, \text{msg}^*, \text{sig}^*)$. The experiment returns 1 iff $\text{Verify}(\text{pp}, \text{mpk}, \text{id}^*, \text{msg}^*, \text{sig}^*) = 1$, the identity id^* was never queried to $\mathcal{O}_{\text{ext}}$, and the pair $(\text{id}^*, \text{msg}^*)$ was never queried to $\mathcal{O}_{\text{sig}}$. The advantage is $\text{Adv}_{\text{IBS}}^{\text{ID-UF-CMA}}(\mathcal{A}) := \Pr[\text{ID-UF-CMA}_{\text{IBS}}(\mathcal{A}) = 1]$. In the random-oracle model, $\mathcal{A}$ also has oracle access to the scheme's random oracles, including H_{cand} , H_{ch} , and the digest hash.

Definition 2.3 (Fixed-statement and planted-witness public-key signing games). A fixed normalized statement is $A = (A_0, A_1) \in R_q^\times \times R_q$, with signing witness $S = (s_0, s_1)^T \in R^2$ satisfying $AS = \hat{q} \bmod q$. The induced signing layer signs digest messages $\mu \in \{0, 1\}^{2\lambda}$ using the Section 3 algorithms after identity-label normalization.

In $\text{PKS-UF-CMA}_{A;S}$ the adversary receives A , has signing-oracle access under S , and wins with an accepting signature on a previously unsigned digest. In PKS-UF-NMA_A it receives A but no signing oracle, and wins with any accepting digest-signature pair; this no-message game is meaningful for every public normalized statement A .

The annular planted-witness scheme induced by the signing layer is denoted $\text{ANSA-PKS}^{\text{ann}}$. Its key generation samples an accepted setup $(B, h) \leftarrow \text{AnnularTrapdoorGen}$ before Galois screening and discards B . For this fixed h , it samples $(s_0, s_1) \leftarrow D_{R^2, \sigma_1}^{\text{amb}}$ conditioned on $a = s_0 + s_1 h \bmod q$ lying in $R_q^\times$, normalizes to $A = (\hat{q}a^{-1}, \hat{q}ha^{-1})$, and uses A as the public key. It retains the sampled $S = (s_0, s_1)^T$ as the signing key without applying C_q . The uniform variant ANSA-PKS^u is identical except that it samples $h \leftarrow U(R_q)$ before the per- h conditional witness draw. The games $\text{PKS-UF-CMA}_{\text{ANSA-PKS}^{\text{ann}}}$, $\text{PKS-UF-NMA}_{\text{ANSA-PKS}^{\text{ann}}}$, $\text{PKS-UF-CMA}_{\text{ANSA-PKS}^u}$, and $\text{PKS-UF-NMA}_{\text{ANSA-PKS}^u}$ use the corresponding planted-witness key generation with the signing and verification algorithms of Section 3. The reduction also uses a centered witness-first intermediate game that preserves the ambient draw and public label but signs with $C_q(S)$. Theorem 4.2 bounds its transition to the corresponding planted-witness game by the probability that $C_q(S) \neq S$.

3 The Two-Response Bimodal ANSA-IBS Construction

3.1 Scheme Syntax and Auxiliary Maps

An ANSA-IBS instance is determined by the parameters $N, q, \lambda, \alpha_{\text{ann}}, \xi, \sigma_1, \kappa, B_{\text{shift}}, \sigma_y, d_h, Z_{\text{BOUND}},$ and $M_{\text{iter}},$ together with the identity-label map, extraction-tape rule, challenge map, high-bit map, and fixed-length signature encoding. Here, (N, q) fix the ring and modulus, $(\alpha_{\text{ann}}, \xi)$ and σ_1 are the trapdoor-generation and extraction parameters from Section 2.2, and λ is the security parameter and digest half-length. The signing parameters are the challenge weight κ , monomial-shift bound $B_{\text{shift}},$ mask width $\sigma_y,$ high-bit drop $d_h,$ public response bound $Z_{\text{BOUND}},$ and one-step rejection constant $M_{\text{iter}}.$

All hash and pseudorandom function (PRF) inputs are serialized by a fixed injective encoding: byte strings are length-delimited where stated, ring elements use canonical coefficient encodings, integers use fixed little-endian encodings of the stated width, and literal tags are leading one-byte domain separators. The symbol $\parallel$ below denotes concatenation of these encoded fields. Set $\text{pkh} := \text{Hash}('H' \parallel h, 256).$

The public candidate map $H_{\text{cand}}(h, \text{id}, j) \in R_q,$ for $0 \leq j < 2^{16},$ is defined by

$$H_{\text{cand}}(h, \text{id}, j) = \text{HashToRq}('I' \parallel \text{pkh} \parallel \text{id} \parallel \text{u16le}(j)) \in R_q.$$

The fixed lengths of pkh and the trailing counter make this input encoding injective in $(\text{pkh}, \text{id}, j).$ We write $a_{h, \text{id}, j} = H_{\text{cand}}(h, \text{id}, j).$ The map HashToRq expands its input with a domain-separated XOF and fills the ring element with coefficients sampled exactly uniformly from $\mathbb{Z}_q$ by rejection sampling.

The identity-label wrapper has output interface

$$H_{\text{id}}(h, \text{id}) \in (\{0, \dots, 2^{16} - 1\} \times R_q^\times) \cup \{\perp\}.$$

It scans the 16-bit counter range and returns the first pair (j, a) such that $a = a_{h, \text{id}, j} \in R_q^\times.$ If no such counter exists, the map returns $\perp.$ Extraction and verification apply this public rule: the returned a must be a unit and every earlier candidate $a_{h, \text{id}, i}$ with $i < j$ must be a nonunit.

The deterministic extraction tape is

$$\begin{aligned} k_{\text{ext}}(B) &:= \text{Hash}('B' \parallel B, 512), \\ G_{\text{ext}}(k_{\text{ext}}(B), \text{id}, j) &:= \text{PRF}_{k_{\text{ext}}(B)}('G' \parallel \text{id} \parallel \text{u16le}(j), 512). \end{aligned}$$

The internal extraction and signing samplers use leading tags 'E' and 'S', respectively, followed by their root or seed and a one-byte stream index (Appendix B).

For a message msg and an identity id of byte length $|\text{id}| < 2^{32},$ signing and verification compute the public digest

$$\mu = \text{Hash}('M' \parallel \text{pkh} \parallel \text{u32le}(|\text{id}|) \parallel \text{id} \parallel \text{u16le}(j) \parallel \text{msg}, 2\lambda).$$

Let ANSA-IBS^μ denote the digest-message variant in which signing and verification take (id, μ) as the message input and use the same call $H_{\text{ch}}(\mu, w).$ Let

$\text{ANSA-IBS}^{\mu, \text{own}}$ denote the wrapper-induced proof game that lazily assigns each public digest input $x = (\text{pkh}, \text{id}, j, \text{msg})$ a cached digest μ and records the proof-only owner map $\text{own}(\mu) = x$. Set

$$\alpha_h = 2^{d_h}, \quad p_h = \left\lfloor \frac{q-1}{\alpha_h} \right\rfloor, \quad r_h = (q-1) - \alpha_h p_h, \quad R_h = 2^{d_h-1} + r_h + 1. \quad (2)$$

For $x \in R_q$, let $c_i \in [-(q-1)/2, (q-1)/2] \cap \mathbb{Z}$ be the unique centered representative of its i th coefficient. The coefficientwise high-bit map $[x]_{d_h} \in \mathbb{Z}_{p_h}^N$ is

$$([x]_{d_h})_i = \left(\text{sgn}(c_i) \left\lfloor \frac{|c_i| + 2^{d_h-1}}{2^{d_h}} \right\rfloor \right) \bmod p_h, \quad \text{sgn}(0) = 0.$$

The formula rounds each c_i to the nearest multiple of α_h , resolves exact magnitude ties away from zero, and reduces the quotient to its canonical integer representative in $\{0, \dots, p_h - 1\}$. Products with α_h use that representative. Each output coefficient determines a rounded multiple modulo q and has at most $2^{d_h+1} + 2$ centered-residue preimages; this conservative bound allows for both wraparound boundaries. Consequently, each high-bit vector has at most $(2^{d_h+1} + 2)^N$ residue preimages, and every coefficient satisfies $\alpha_h[x]_{d_h} = x + e \pmod{q}$ with $|e| \leq R_h$.

The challenge space is $\mathcal{C}_\kappa = \{c \in R : c_i \in \{0, 1\}, \|c\|_0 = \kappa\}$. The map $H_{\text{ch}}(\mu, w)$ expands ' $\mathbf{C}$ ' $\parallel \mu \parallel w$ with a domain-separated XOF, where w is encoded by packing its representatives in $\{0, \dots, p_h - 1\}$ into consecutive $\lceil \log_2 p_h \rceil$ -bit fields, least significant bit first (10-bit fields for Table 1), and samples uniformly from $\mathcal{C}_\kappa$ using the binary fixed-weight challenge-sampling method of HAETAE and NTRU+Sign [6,31].

Algorithm (H_{ch} : fixed-weight challenge sampling)

Input: a digest $\mu \in \{0, 1\}^{2\lambda}$ and commitment $w \in \mathbb{Z}_{p_h}^N$.

Output: $c \in \mathcal{C}_\kappa$, with support $\ell_1 < \dots < \ell_\kappa$.

1. Form the byte stream $\text{SHAKE256}(\text{'C'} \parallel \mu \parallel w \parallel \text{u32le}(t), 256)$ for $t = 0, 1, \dots$, concatenated in order, with w encoded as above; each block is the first 256 output bytes. Read the bytes in order and each byte's bits from least to most significant.
2. Set $c_0 = \dots = c_{N-1} = 0$.
3. For $i = N - \kappa, \dots, N - 1$:
 1. Set $b_i = \lceil \log_2(i+1) \rceil$.
 2. Repeatedly read the next b_i bits $u_0, \dots, u_{b_i-1}$ and set $j = \sum_{r=0}^{b_i-1} u_r 2^r$ until $j \leq i$. If $b_i = 0$, set $j = 0$ without reading bits. Discard rejected draws and retain all unread bits for subsequent draws.
 3. Set $c_i \leftarrow c_j$, then $c_j \leftarrow 1$.
4. Return $c = \sum_{i=0}^{N-1} c_i X^i$ and enumerate its nonzero positions in increasing order.

In the ideal-XOF model, each accepted index is uniform on $\{0, \dots, i\}$. Induction on Step 3 shows that the final support is a uniform κ -subset of $\{0, \dots, N-1\}$, so c is uniform in $\mathcal{C}_\kappa$.

One-step Gaussian transport. The signing sampler uses the randomized procedure $\mathcal{R}_v$ of Gärtner [16]. Its input is a state $y \in R^2$ and a nonzero shift $v \in R^2$ with $\|v\|_2 \leq B_{\text{shift}}$. The width $\sigma_y > 0$ and normalizer $M > 1$ are public and fixed independently of y and v . Norms and inner products use the $2N$ coefficient coordinates. Set

$$\rho_{\sigma_y}(x) = \exp\left(-\frac{\|x\|_2^2}{2\sigma_y^2}\right), \quad V = \|v\|_2^2, \quad T = \langle y, v \rangle,$$

and define the absolutely convergent series

$$\begin{aligned} S_v(y) &= \sum_{k=0}^{\infty} (-1)^k \frac{\rho_{\sigma_y}(y + kv)}{\rho_{\sigma_y}(y)} \\ &= \sum_{k=0}^{\infty} (-1)^k \exp\left(-\frac{2kT + k^2V}{2\sigma_y^2}\right). \end{aligned} \tag{3}$$

The minus- and plus-branch probabilities are

$$f_v(y) = \frac{1}{M} \begin{cases} S_v(y), & T \geq V, \\ 1 - S_v(-y), & T < V, \end{cases} \quad g_v(y) = \frac{1}{M} \begin{cases} 1 - S_v(y), & T \geq -V, \\ S_v(-y), & T < -V. \end{cases}$$

For $\alpha > 0$, put

$$M_\alpha = 1 + \frac{2\alpha\sqrt{2\pi}\exp(-\pi^2\alpha^2/2)}{\exp(-1/(2\alpha^2))(1 - \exp(-2\pi^2\alpha^2))}. \tag{4}$$

The scheme fixes $M = M_{\text{iter}} \geq M_{\alpha_{\text{iter}}}$, where $\alpha_{\text{iter}} = \sigma_y/B_{\text{shift}}$. This common choice makes $f_v(y)$, $g_v(y)$, and $1 - f_v(y) - g_v(y)$ nonnegative for every admitted shift and state. With a fresh independent $U \leftarrow U([0, 1])$, the procedure returns

$$\mathcal{R}_v(y) = \begin{cases} y - v, & U < f_v(y), \\ y + v, & f_v(y) \leq U < f_v(y) + g_v(y), \\ \perp, & U \geq f_v(y) + g_v(y). \end{cases}$$

The additions and subtractions take place in the integer coefficient lattice of R^2 , without coefficient reduction modulo q . Section 4.3 establishes the validity and Gaussian output law of this procedure.

3.2 Setup and Master-Key Generation

Algorithm 3.1 (Setup)

Input: the selected parameter set.

Output: public parameters **pp**, master public key **mpk**, master secret key **msk**, or $\perp$.

1. Run $(B, h) \leftarrow \text{AnnularTrapdoorGen}(N, q, \alpha_{\text{ann}}, \xi)$, including the Galois screen of Section 5.3; return $\perp$ on failure.
2. Set **mpk** = h , **msk** = B , and let **pp** contain the selected parameter tuple, ring description, map definitions, response-bound rule, and encoding parameters.

3.3 Identity-Key Extraction

Algorithm 3.2 (Extract)

Input: pp , $\text{msk} = B$, $\text{mpk} = h$, and identity string id .

Output: identity signing key $\text{usk}_{\text{id}} = (j, a, s_0, s_1)$, or $\perp$.

1. Compute $(j, a) = H_{\text{id}}(h, \text{id})$; if the map returns $\perp$, output failure.
2. Derive $\text{seed}_G = G_{\text{ext}}(k_{\text{ext}}(B), \text{id}, j)$.
3. Sample $(s_0, s_1) \leftarrow \text{SamplePreimage}_P^{\text{exact}}(B, h, a, \sigma_1; \text{seed}_G)$ satisfying $s_0 + s_1 h = a \bmod q$.
4. Output $\text{usk}_{\text{id}} = (j, a, s_0, s_1)$.

3.4 Signing

Algorithm 3.3 (Sign)

Input: pp , $\text{mpk} = h$, identity id , message msg , and identity signing key (j, a, s_0, s_1) .

Output: signature $\text{sig} = (j, c, z_0, z_1)$.

1. Set $\hat{q} = 2^{-1} \bmod q$, $A = (A_0, A_1) = (\hat{q}a^{-1}, \hat{q}ha^{-1}) \in R_q^\times \times R_q$, and $S = (s_0, s_1)^T$. Compute μ as in Section 3.1.
2. Sample $y = (y_0, y_1) \leftarrow D_{R, \sigma_y}^2$.
3. Compute $u = A_0 y_0 + A_1 y_1 \bmod q$, $w = [u]_{d_h}$, and $c = H_{\text{ch}}(\mu, w) = \sum_{t=1}^{\kappa} X^{\ell_t}$, where $0 \leq \ell_1 < \dots < \ell_\kappa < N$.
4. Set $(z_0, z_1) = (y_0, y_1)$. For $t = 1, \dots, \kappa$, set $v_t = X^{\ell_t} S$, restart from Step 2 if $\|v_t\|_2 > B_{\text{shift}}$, and update $(z_0, z_1)^T \leftarrow \mathcal{R}_{v_t}((z_0, z_1)^T)$ with $M = M_{\text{iter}}$ as specified in Section 3.1. Restart from Step 2 if any call returns $\perp$.
5. Reject and restart from Step 2 unless $\|z_i\|_\infty \leq Z_{\text{BOUND}}$ for $i = 0, 1$.
6. Reject and restart from Step 2 unless

$$[A_0 z_0 + A_1 z_1 + c\hat{q}]_{d_h} = [A_0 z_0 + A_1 z_1 + c\hat{q} - c]_{d_h}.$$

7. Output $\text{sig} = (j, c, z_0, z_1)$.

The unbounded exact signer is a reference kernel: a nonterminating invocation yields no completed transcript and cannot produce a winning execution. The stopped games in Appendix A.5 terminate on an inadmissible shift; the shift-failure term in Theorem 4.2 accounts for the exceptional witnesses without conditioning their extraction law. The Step 4 norm checks only cause internal restarts. They always pass for a monomial-shift-bounded witness. Lemma 4.12 bounds the probability that a witness-first (ambient) draw falls outside this class; the distributional comparisons of Theorem 4.2 transfer this bound to extracted witnesses. A bounded implementation makes at most T_{sign} attempts (Section B.5) and returns $\perp$ on exhaustion. For a bounded witness, the independent-challenge reference execution exhausts with probability at most $(1 - \rho_{\text{na}})^{T_{\text{sign}}}$, where ρ_{na} is the non-abort probability of Lemma 4.13; Appendix B accounts for the random-oracle freshness and implementation errors. Lemma 4.8 gives the truncation bound needed to implement each ideal call $\mathcal{R}_{v_t}$.

3.5 Verification

Algorithm 3.4 (Verify)

Input: pp , $\text{mpk} = h$, identity id , message msg , and signature $\text{sig} = (j, c, z_0, z_1)$.
 Output: accept or reject.

1. Query $H_{\text{id}}(h, \text{id})$ and reject unless it returns a label (j, a) whose counter equals the signature's j ; use this a below.
2. Compute μ as in Section 3.1.
3. Reject unless $c \in \mathcal{C}_\kappa$, $\|z_0\|_\infty \leq Z_{\text{BOUND}}$, and $\|z_1\|_\infty \leq Z_{\text{BOUND}}$.
4. Set $\hat{q} = 2^{-1} \bmod q$ and $A = (A_0, A_1) = (\hat{q}a^{-1}, \hat{q}ha^{-1})$, and recover $w' = [A_0z_0 + A_1z_1 + c\hat{q}]_{d_h}$.
5. Accept iff $c = H_{\text{ch}}(\mu, w')$.

Correctness follows from Lemma 4.11 and the equality check in Step 6 of the signing algorithm. Although the κ hidden signs need not agree, the public predicate guarantees that the verifier recovers the original commitment $w = [u]_{d_h}$.

Appendix B gives fixed-point realizations of Setup, extraction, and signing and relates their output laws to the exact algorithms above. For extraction, Section 4 analyzes the centered, ungated Hybrid-Sampler law and its transition to the planted-witness game, while Appendix B accounts separately for the admissibility gate.

4 Security Analysis

This section proves adaptive identity-based unforgeability in the random-oracle model. It gives the reduction plan and assumptions, identity-layer tools, fixed-statement signing analysis, and full IBS reduction; proofs are in Appendix A.

4.1 Proof Plan, Games, and Assumptions

The reduction for the unscreened public-key law $\mathcal{H}_{\text{ann}}$ has three stages. First, the wrapper assigns each public input $(\text{pkh}, \text{id}, j, \text{msg})$ a cached digest $\mu \in \{0, 1\}^{2\lambda}$ and reduces to the owner-tagged game $\text{ANSA-IBS}^{\mu, \text{own}}$, with freshness on (id, μ) . Second, identity-layer hybrids replace deterministic extraction tapes and real Hybrid-Sampler witnesses by independent tapes and ideal conditional witnesses, then program unit labels witness-first. Third, the forgery identity is localized to an annular planted-witness instance; annular NTRU switches h to uniform, iterative FSwA removes the signing oracle, a unit-label hop randomizes the normalized statement, and Lemma 4.16 maps an accepting no-message forgery to BMST-RSIS.

Running times count expected total work, including exact sampling [2, Section 2]. Oracle-query budgets are worst-case and global unless stated otherwise: Q_{ID} counts distinct identities for which witness generation is attempted, including cached failures, Q_{Hid} distinct reached identity-label states, including the

label state of the forgery identity, and Q_S signing queries. Each identity has one cached label state under the master public key, so repeated calls and candidate-counter queries do not increase Q_{Hid} . Identities hashed without extraction count toward Q_{Hid} , so a bound on issued keys alone does not bound Q_{Hid} . Q_H counts explicit H_{ch} queries plus the final fresh verification point, while Q_μ counts distinct public digest inputs $(\text{pkh}, \text{id}, j, \text{msg})$. Lemma 4.17 charges signing-oracle programming through Q_S ; rejection-loop attempts enter the non-abort denominator, not another Q_H multiplier. Challenge queries made while signing for non-target identities are answered locally by the reduction (proof of Lemma 4.18) and are not charged to Q_H .

The random-oracle model exposes the candidate oracle H_{cand} and its deterministic first-unit wrapper H_{id} . On input (h, id, j) , H_{cand} lazily samples and caches independent elements $a_{h, \text{id}, i} \leftarrow U(R_q)$ for each missing $0 \leq i \leq j$, then returns $a_{h, \text{id}, j}$. The wrapper H_{id} queries $H_{\text{cand}}(h, \text{id}, j)$ for $j = 0, 1, \dots, 2^{16} - 1$ and returns the first $(j, a_{h, \text{id}, j})$ with $a_{h, \text{id}, j} \in R_q^\times$, or $\perp$ if no unit occurs.

The identity-layer reduction generates labels when candidate or wrapper state is needed, but samples witnesses only for extraction, non-simulated signing, or the hidden target slot. The first hop replaces each Section 3.1 extraction tape by an independent cached 512-bit tape when its witness is first needed. The tape key is setup-derived from the hidden basis B , which is never a public input.

To simulate G_{ext} , ExtRand^b samples an accepted setup (B, h) , exposes its public transcript, and answers the IBS state interface. When a witness is first needed, it caches a tape r_{id} , using $G_{\text{ext}}(k_{\text{ext}}(B), \text{id}, j)$ for $b = 0$ and an independent 512-bit string for $b = 1$; later requests reuse the tape, witness, counter, and label. Let $\text{Bad}_{\text{ext-coll}}$ be a collision among cached tapes, and define

$$\begin{aligned} \text{Adv}_{\text{ext-rand}}(\mathcal{F}_{G_{\text{ext}}}) &:= \left| \Pr[\text{ExtRand}^0(\mathcal{F}_{G_{\text{ext}}}) \Rightarrow 1] - \Pr[\text{ExtRand}^1(\mathcal{F}_{G_{\text{ext}}}) \Rightarrow 1] \right| \\ &\quad + \max_{b \in \{0,1\}} \Pr[\text{Bad}_{\text{ext-coll}} \text{ in } \text{ExtRand}^b]. \end{aligned}$$

In the ideal keyed-stream model for the internal PRF, the distinguishing term is zero and the collision term is at most $\binom{Q_{\text{ID}}}{2} 2^{-512}$, below 2^{-417} for $Q_{\text{ID}} = 2^{48}$.

Witness replacement and witness-first label programming incur Rényi transfers over Q_{ID} witnesses and Q_{Hid} labels. The following assumptions quantify the post-localization public-key and statement switches at the reductions' expected running times and query budgets; efficiency means expected polynomial time.

Assumption 4.1 (Decisional annular NTRU). Fix the profile and the candidate-evaluation and NTRU-solver policies of Sections B.1–B.3. Let $\mathcal{H}_{\text{ann}}$ be the master-public-key law of AnnularTrapdoorGen conditioned on successful output, including cutoff and prescribed solver rejections but omitting the Galois screen. For any efficient distinguisher $\mathcal{D}$, define

$$\text{Adv}^{\text{dec-ann-NTRU}}(\mathcal{D}) := \left| \Pr_{h \leftarrow \mathcal{H}_{\text{ann}}} [\mathcal{D}(h) = 1] - \Pr_{u \leftarrow U(R_q)} [\mathcal{D}(u) = 1] \right|.$$

We assume this advantage is negligible for efficient distinguishers, following the decisional NTRU hypotheses used in NTRU-based encryption and KEM

analyses [21,19]. The Galois screen of Algorithm 3.1 changes this public-key law by total variation distance at most $\Delta_{\text{stab},P} + \varepsilon_{\text{stab},P}$ (Section 5.3), giving the additive screening term in Corollary B.2. In the ideal-XOF realization of Appendix B, the assumption concerns the joint law of the oracle and h conditioned on successful unscreened Setup. The uniform experiment retains that oracle marginal and samples $u \leftarrow U(R_q)$ independently; the distinguisher has the same oracle interface in both experiments.

Assumption 4.2 (Unit-conditioned one-sample label distinguishing). Sample $h \leftarrow U(R_q)$; for this fixed h , sample (s_0, s_1) from the ambient witness distribution $D_{R^2, \sigma_1}^{\text{amb}}$ conditioned on $s_0 + s_1 h \in R_q^\times$; and sample $u \leftarrow U(R_q^\times)$ independently. For any efficient distinguisher $\mathcal{D}$, define

$$\text{Adv}_{\sigma_1}^{\text{u1-lab}}(\mathcal{D}) := |\Pr[\mathcal{D}(h, s_0 + s_1 h \bmod q) = 1] - \Pr[\mathcal{D}(h, u) = 1]|.$$

This assumption concerns unit syndromes of ambient Gaussian witnesses, rather than Ring-LWE with small secrets. Lemma 4.4 gives a statistical comparison under quotient smoothing, established here for annular keys; the uniform-key game retains Assumption 4.2.

After these hybrids, independent $h \leftarrow U(R_q)$ and $u \leftarrow U(R_q^\times)$ give uniform $A_0 = \hat{q}u^{-1} \in R_q^\times$. For fixed A_0 , $h \mapsto hA_0$ is a bijection of R_q , so $A_1 = hA_0$ is uniform and independent of A_0 , as Assumption 4.3 requires.

Assumption 4.3 states the BMST-RSIS relation underlying HAETAE and NTRU+Sign [6,31] in ANSA-IBS form; it is the ring analogue of Dilithium's self-target MSIS relation [11,2].

Assumption 4.3 (Bimodal self-target RSIS). Fix the parameters $(R_q, N, q, \kappa, d_h, B_{\text{ST}}, B_{\text{ST},\infty})$. Sample $A_0 \leftarrow U(R_q^\times)$ and $A_1 \leftarrow U(R_q)$. The adversary receives $A = (A_0, A_1)$ and random-oracle access to $H_{\text{ch}} : \{0, 1\}^{2\lambda} \times \mathbb{Z}_{p_h}^N \rightarrow \mathcal{C}_\kappa$. It wins if it outputs $(\mu, c, Y_0, Y_1) \in \{0, 1\}^{2\lambda} \times \mathcal{C}_\kappa \times R^2$ such that

$$H_{\text{ch}}(\mu, [A_0 Y_0 + A_1 Y_1 + c\hat{q}]_{d_h}) = c,$$

and

$$\|(Y_0, Y_1)\|_2 \leq B_{\text{ST}}, \quad \|(Y_0, Y_1)\|_\infty \leq B_{\text{ST},\infty}.$$

The corresponding random-statement advantage is denoted

$$\text{Adv}_{H_{\text{ch}}, R_q, N, q, d_h, B_{\text{ST}}, B_{\text{ST},\infty}}^{\text{BMST-RSIS}}(\mathcal{C}).$$

For fixed A , the fixed-statement version includes A in the subscript.

Lemma 4.16 maps one accepting no-message forgery to the relation in Assumption 4.3; the high-bit map contributes a bounded lift residual, and $Y = (Y_0, Y_1)$ may be zero. Ordinary RSIS instead needs two accepting forked transcripts with the same recovered commitment, distinct challenges, and the lemma's no-wrap conditions. Thus it is only a conditional fork consequence, whereas BMST-RSIS is the direct one-forgery target.

The self-target condition also permits direct challenge search: fix $c \in \mathcal{C}_\kappa$ and bounded Y (even $Y = 0$), set $w_{c,Y} = [A_0 Y_0 + A_1 Y_1 + c\hat{q}]_{d_h}$, and vary μ until

$H_{\text{ch}}(\mu, w_{c,Y}) = c$. Each trial succeeds with probability $1/|\mathcal{C}_\kappa|$, so the challenge-space entropy is $\lambda_{\text{ch}} := \log_2 |\mathcal{C}_\kappa| = \log_2 \binom{N}{\kappa}$.

Definition 4.1 (Auxiliary ordinary-RSIS fork instance). Sample $\bar{A}_0 \leftarrow U(R_q^\times)$ and $\bar{A}_1 \leftarrow U(R_q)$ and set $\bar{A} = [\bar{A}_0 \mid \bar{A}_1 \mid 1] \in R_q^{1 \times 3}$. Find nonzero $(X_0, X_1, X_2) \in R^3$ such that $\bar{A}_0 X_0 + \bar{A}_1 X_1 + X_2 = 0 \bmod q$ and, in centered representatives,

$$\|(X_0, X_1, X_2)\|_2 \leq B_{\text{RSIS}}, \quad \|(X_0, X_1, X_2)\|_\infty < (q-1)/2.$$

For the fork output of Lemma 4.16 under its ordinary-fork conditions,

$$B_{\text{RSIS}} = 2\sqrt{B_{\text{ST}}^2 + 4NR_h^2} + \sqrt{2\kappa}, \quad \|(X_0, X_1, X_2)\|_\infty \leq B_{\text{fork},\infty}. \quad (5)$$

4.2 Distributional Tools

The identity layer compares real extraction and label distributions with trapdoor-free witness-first distributions. The following Rényi tools control these replacements over adaptive queries.

Definition 4.2 (Rényi divergence factor). For distributions P, Q on a common countable support and order $\alpha > 1$, let

$$R_\alpha(P\|Q) := \left(\sum_x P(x)^\alpha Q(x)^{1-\alpha} \right)^{1/(\alpha-1)}.$$

Lemma 4.1 (Rényi event transfer; [29]). If $R_\alpha(P\|Q) \leq r$ and E is an event, then $\Pr_P[E] \leq (r \Pr_Q[E])^{(\alpha-1)/\alpha}$.

Lemma 4.2 (Sequential Rényi accumulation; [5, Lemma 5.4]). Let Law_{real} and $\text{Law}_{\text{ideal}}$ be two laws on transcript chains $T_i = (T_{i-1}, Z_i)$ with a common initial distribution. Suppose that, for every $i \leq m$, the real prefix support of T_{i-1} is contained in the ideal prefix support of T_{i-1} and, for every prefix τ in that real prefix support,

$$R_\alpha(\text{Law}_{\text{real}}(Z_i \mid T_{i-1} = \tau) \parallel \text{Law}_{\text{ideal}}(Z_i \mid T_{i-1} = \tau)) \leq r.$$

Then $R_\alpha(\text{Law}_{\text{real}}(T_m) \parallel \text{Law}_{\text{ideal}}(T_m)) \leq r^m$.

Lemma 4.3 (Canonical Hybrid-Sampler extraction comparison). Prest's Hybrid-Sampler/Rényi comparison [28,29], in the rank-two NTRU form of [5, Lemma 5.3], gives the following bound. Fix a compatible rank-two NTRU trapdoor basis B for h and syndrome a . Let $D_{\text{hyb}}(\cdot \mid B, h, a)$ be the exact Hybrid-Sampler extraction distribution and $D_{\text{pub}}(\cdot \mid h, a)$ be $D_{R^2, \sigma_1}^{\text{amb}}$ conditioned on $P_h(a)$. If the normalized canonical Hybrid-Sampler admissibility condition holds with parameter ε_p , then for every Rényi order $a_p > 1$ there is an explicit factor $r_p = r_p(a_p, \varepsilon_p)$ such that

$$R_{a_p}(D_{\text{hyb}}(\cdot \mid B, h, a) \parallel D_{\text{pub}}(\cdot \mid h, a)) \leq r_p.$$

Rényi data processing under the deterministic map $\mathcal{C}_q$ also gives

$$R_{a_p}(D_{\text{hyb}}^{\text{ctr}}(\cdot \mid B, h, a) \| D_{\text{pub}}^{\text{ctr}}(\cdot \mid h, a)) \leq r_p.$$

Both inequalities hold for every syndrome; the IBS reduction needs them only for unit labels. Lemma 2.1 establishes the normalized-canonical width condition for both ring families, and Appendix A.1 identifies the pushforward of Prest's lattice Gaussian at the appropriate affine center with $D_{\text{pub}}(\cdot \mid h, a)$. In the small- ε_p regime, use the conservative relative-error bound $\delta'_p = ((1 + \varepsilon_p)/(1 - \varepsilon_p))^6 - 1$ and assume $\delta'_p < 1$. The resulting bound is

$$r_p \leq \left(1 + \frac{(a_p/2)(a_p - 1)(\delta'_p)^2}{(1 - \delta'_p)^{a_p+1}} \right)^{1/(a_p-1)}.$$

Lemma 4.4 (Unit-label witness programming). Let Q_h be the distribution of $a = s_0 + s_1 h \bmod q$ for $S \leftarrow D_{R^2, \sigma_1}^{\text{amb}}$. Fix $0 < \varepsilon_u < 1/3$ and assume $\sigma_1 \geq \eta_{\varepsilon_u}(\Lambda_{h,q})$ in the normalized canonical metric. The quotient-Gaussian bound [5, Lemma 5.7] gives $U(R_q)(a)/Q_h(a) \in [(1 - \varepsilon_u)/(1 + \varepsilon_u), (1 + \varepsilon_u)/(1 - \varepsilon_u)]$ for every a . Set $c_u = (1 + \varepsilon_u)/(1 - \varepsilon_u)$ and $\delta_u^\times = c_u^2 - 1$, and assume $\delta_u^\times < 1$. After conditioning on $a \in R_q^\times$, $Q_h^\times(a)/U(R_q^\times)(a) \in [c_u^{-2}, c_u^2]$ for every unit a . Moreover, for Rényi order $a_u > 1$,

$$R_{a_u}(U(R_q^\times) \| Q_h^\times) \leq r_u^\times,$$

where

$$r_u^\times \leq \left(1 + \frac{(a_u/2)(a_u - 1)(\delta_u^\times)^2}{(1 - \delta_u^\times)^{a_u+1}} \right)^{1/(a_u-1)}.$$

Writing $\rho_\times = |R_q^\times|/q^N$, a witness-programmed unit label has probability at most $\Gamma_{\text{wp}}^\times \leq c_u^2/(q^N \rho_\times)$. Because $\mathcal{C}_q$ preserves the syndrome modulo q , adjoining to either unit-label law a witness drawn from the common conditional kernel $D_{\text{pub}}^{\text{ctr}}(\cdot \mid h, a)$ leaves the displayed Rényi factor unchanged.

Lemma 4.5 (Counter-based unit-label abstraction and counter exhaustion). Let the identity map scan J_{max} independent candidates from $U(R_q)$ and return the first unit. With at most Q_{Hid} identity-label states,

$$\varepsilon_{\text{unit}} \leq Q_{\text{Hid}} \nu_{\text{nonunit}}^{J_{\text{max}}}, \quad \nu_{\text{nonunit}} = \Pr_{a \leftarrow U(R_q)}[a \notin R_q^\times].$$

Conditioned on no counter exhaustion, the returned unit label is uniform in $R_q^\times$ and independent of the counter j , where

$$\Pr[j = t \mid \text{no exhaustion}] = \frac{\nu_{\text{nonunit}}^t (1 - \nu_{\text{nonunit}})}{1 - \nu_{\text{nonunit}}^{J_{\text{max}}}}, \quad 0 \leq t < J_{\text{max}}.$$

Prior candidates are simulated as independent uniform nonunits for the public first-unit check, and later H_{cand} candidates are sampled lazily from $U(R_q)$. Thus the witness-programming Rényi factor applies only to the returned unit label. If $\phi = \prod_i \phi_i$ is squarefree with $\deg \phi_i = f_i$, then $\nu_{\text{nonunit}} = 1 - \prod_i (1 - q^{-f_i})$.

4.3 Fixed-Statement Bimodal Signing Security

Fix normalized $A = (A_0, A_1) \in R_q^\times \times R_q$ and witness $S = (s_0, s_1)^T$ with $AS = \hat{q} \bmod q$. Say S is monomial-shift-bounded when $\|(X^j s_0, X^j s_1)\|_2 \leq B_{\text{shift}}$ for all $0 \leq j < N$. The efficiently checkable event **SG** requires this bound for all shifts selected by generated signing challenges, including those not reached after an earlier transport rejection. The global shift event implies **SG**.

We establish the exact response law of the transport specified in Section 3.1. On the coefficient lattice $\mathbb{Z}^{2N}$, the centered Gaussian has probability mass

$$p_{\sigma_y}(x) = \frac{\rho_{\sigma_y}(x)}{\sum_{u \in \mathbb{Z}^{2N}} \rho_{\sigma_y}(u)}.$$

Lemma 4.6 (One-step transport identity; [16]). Let Y and Z have probability masses p_Y and p_Z . If $f_v(y), g_v(y) \geq 0$, $f_v(y) + g_v(y) \leq 1$, and

$$p_Y(z+v)f_v(z+v) + p_Y(z-v)g_v(z-v) = \frac{p_Z(z)}{M} \quad (6)$$

for every z , then $\Pr[\mathcal{R}_v(Y) \neq \perp] = 1/M$ and the output conditioned on success is distributed as Z .

For the Gaussian case, let $K_v = \{y : |\langle y, v \rangle| \leq \|v\|_2^2\}$ denote the central strip.

Lemma 4.7 (Gaussian transport functions; [16]). Suppose

$$M \geq \max_{y \in K_v} \left(1 - \sum_{k \in \mathbb{Z}} (-1)^k \frac{\rho_{\sigma_y}(y + kv)}{\rho_{\sigma_y}(y)} \right). \quad (7)$$

The branch functions f_v, g_v specified in Section 3.1 satisfy Lemma 4.6 with $Y = Z = D_{\sigma_y}^{2N}$.

Lemma 4.8 (Finite evaluation of S_v ; [16, Lemma 2]). Let

$$S_v^{(J_{\text{tr}})}(y) = \sum_{k=0}^{J_{\text{tr}}} (-1)^k \frac{\rho_{\sigma_y}(y + kv)}{\rho_{\sigma_y}(y)}.$$

For $J_{\text{tr}} \geq 1$ and $\langle y, v \rangle \geq -\|v\|_2^2$,

$$\left| S_v(y) - S_v^{(J_{\text{tr}})}(y) \right| \leq \frac{\rho_{\sigma_y}(J_{\text{tr}}v)}{\rho_{\sigma_y}(v) - \rho_{\sigma_y}(2v)}. \quad (8)$$

The bound also applies to $S_v(-y)$ under its displayed condition, so interval evaluation bounds every branch probability. If one implemented step is within $\varepsilon_{\text{step}}$ of ideal in total variation, the κ -step kernel, including abort, is within $\kappa\varepsilon_{\text{step}}$. Let $\varepsilon_{\text{impl}}$ uniformly bound the statistical distance between the complete implemented and ideal signing-output distributions, including retries and the public filter.

Lemma 4.8 gives the per-output interface required by the abstract exact-signer reduction. For the fixed-point scheme, Appendix B proves a sharper likelihood/additive transfer over the entire Q_S -query transcript. After this transfer, Lemma 4.17 applies to the exact signer with $\varepsilon_{\text{impl}} = 0$.

Lemma 4.9 (Uniform one-step constant; [16]). Let M_α be defined by Eq. (4) for $\alpha > 0$. The function M_α is strictly decreasing. If $\|v\|_2 \leq B_{\text{shift}}$, $\alpha_{\text{iter}} = \sigma_y/B_{\text{shift}}$, and $M \geq M_{\alpha_{\text{iter}}}$, then the branch functions of Section 3.1 define a valid one-step sampler for v , with success probability M^{-1} . Its abort probability is pointwise at least

$$1 - \frac{M_{\alpha_{\text{iter}}}}{M}.$$

Lemma 4.10 (Iterated response distribution). Fix $c \in \mathcal{C}_\kappa$ with support $\ell_1 < \dots < \ell_\kappa$, let $v_t = X^{\ell_t}S$, and suppose $\|v_t\|_2 \leq B_{\text{shift}}$ for every t . Starting with $z^{(0)} \leftarrow D_{\sigma_y}^{2N}$, apply $z^{(t)} \leftarrow \mathcal{R}_{v_t}(z^{(t-1)})$ sequentially with a common $M_{\text{iter}} \geq M_{\alpha_{\text{iter}}}$, aborting on $\perp$. For every t , conditioned on the first t steps succeeding, $z^{(t)}$ is exactly $D_{\sigma_y}^{2N}$; hence so is $z := z^{(\kappa)}$ conditioned on success. All steps succeed with probability $M_{\text{iter}}^{-\kappa}$, and

$$z = y + c'S, \quad c' = \sum_{t=1}^{\kappa} \epsilon_t X^{\ell_t}, \quad \epsilon_t \in \{-1, 1\}, \quad c' \equiv c \pmod{2}. \quad (9)$$

Lemma 4.11 (Mixed-sign commitment recovery). Let $u = Ay$, $AS = \hat{q}$, and let z, c' be as in Lemma 4.10. Set $d = (c' + c)/2$, whose coefficients are binary and whose support is contained in that of c . Then

$$Az + c\hat{q} = u + d \pmod{q}. \quad (10)$$

If $[Az + c\hat{q}]_{d_h} = [Az + c\hat{q} - c]_{d_h}$, the recovered value $[Az + c\hat{q}]_{d_h}$ equals the original commitment $[u]_{d_h}$.

Lemma 4.12 (Monomial-shift and centering-change bounds). For fixed h , set $p_\times(h) := \Pr_{S \leftarrow D_{R^2, \sigma_1}^{\text{amb}}} [s_0 + s_1 h \in R_q^\times]$. Assume $p_\times(h) \geq p_{\times, \min} > 0$ uniformly over the accepted annular public keys, and assume, as in [5], the Lemma 4.4 smoothing condition $\sigma_1 \geq \eta_{\varepsilon_u}(\Lambda_{h,q})$. For Table 1, one may take $p_{\times, \min} = c_u^{-1} \rho_\times$. Given $\lambda_{\text{sh}} > 0$, define

$$t_2 = (\lambda_{\text{sh}} + \log_2 Q_{\text{Hid}}) \ln 2, \quad t_3 = (\lambda_{\text{sh}} + \log_2 Q_{\text{Hid}} + \log_2 N) \ln 2.$$

Set

$$B_{\text{shift}} = \begin{cases} \sigma_1 \sqrt{2N + 2\sqrt{2N}t_2 + 2t_2}, & R = \mathbb{Z}[X]/(X^N + 1), \\ \sigma_1 \sqrt{\frac{8N}{3} + 2\sqrt{\frac{40N}{9}}t_3 + 4t_3}, & R = \mathbb{Z}[X]/(X^N - X^{N/2} + 1). \end{cases}$$

Over at most Q_{Hid} ambient witness-first draws, the probability that any witness is not monomial-shift-bounded is at most $2^{-\lambda_{\text{sh}}}$; in the unit-conditioned game, $\varepsilon_{\text{sh}}^{\text{wf}} \leq 2^{-\lambda_{\text{sh}}}/p_{\times, \min}$.

To bound the effect of coefficient centering, put $m_q = (q-1)/2$ and set $\nu_R = 1$ for a power-of-two ring and $\nu_R = 4/3$ for a trinomial ring. Couple each

ambient draw S with $\mathcal{C}_q(S)$. Over at most Q_{Hid} unit-conditioned witness-first draws, the probability that any coupled pair differs is at most

$$\varepsilon_{\text{wrap}}^{\text{wf}} \leq \frac{4NQ_{\text{Hid}}}{p_{\times, \min}} \exp\left(-\frac{m_q^2}{2\nu_R\sigma_1^2}\right).$$

For every profile in Table 1 with $Q_{\text{Hid}} \leq 2^{64}$, this bound is smaller than $2^{-497000}$.

Lemma 4.13 (Scheme-level iterative rejection). For a monomial-shift-bounded fixed statement and fixed $c \in \mathcal{C}_\kappa$, the pre-filter response conditioned on iterative success is exactly $D_{\sigma_y}^{2N}$ and witness-independent. Success has probability $M_{\text{iter}}^{-\kappa}$, independent of c and the emitted response.

For $A = (A_0, A_1)$, $c \in \mathcal{C}_\kappa$, and $z = (z_0, z_1)$, set $\tau = A_0z_0 + A_1z_1 + c\hat{q}$. Let

$$E_{\text{cap}} = \{\|z_0\|_\infty \leq Z_{\text{BOUND}} \text{ and } \|z_1\|_\infty \leq Z_{\text{BOUND}}\},$$

$$E_{\text{eq}}(A, c) = \{[\tau]_{d_h} = [\tau - c]_{d_h}\},$$

and let $E_{\text{pub}}(A, c) = E_{\text{cap}} \cap E_{\text{eq}}(A, c)$. Define

$$\rho_{\text{pub}} = \min_{\substack{A \in R_q^\times \times R_q \\ c \in \mathcal{C}_\kappa}} \Pr_{z \leftarrow D_{\sigma_y}^{2N}}[E_{\text{pub}}(A, c)],$$

$$\rho_{\text{na}} = \rho_{\text{pub}} / M_{\text{iter}}^\kappa.$$

The event $E_{\text{pub}}(A, c)$ is a deterministic post-rejection filter; given A, c, z , it depends on neither hidden signs nor the witness. Set $b_h = p_h + 1 + \mathbf{1}_{p_h \text{ odd}}$. With

$$\begin{aligned} p_{q, \sigma_y} &= \max_{r \in \mathbb{Z}_q} \Pr_{x \leftarrow D_{\mathbb{Z}, \sigma_y}}[x = r \bmod q] \\ &\leq \frac{1 + 2e^{-q^2/(2\sigma_y^2)} / (1 - e^{-3q^2/(2\sigma_y^2)})}{\sqrt{2\pi}\sigma_y}, \end{aligned} \quad (11)$$

the public-event probability satisfies

$$\rho_{\text{pub}} \geq \left(1 - \operatorname{erfc}\left(\frac{Z_{\text{BOUND}}}{\sqrt{2}\sigma_y}\right)\right)^{2N} - \kappa b_h p_{q, \sigma_y}. \quad (12)$$

Lemma 4.14 (Accepting-HVZK and commitment recovery). For a monomial-shift-bounded fixed statement, the public-coin protocol samples its challenge uniformly and independently of the first message. For fixed c , a simulator trial samples $z \leftarrow D_{\sigma_y}^{2N}$, sets $\tau = Az + c\hat{q}$, and accepts on $E_{\text{pub}}(A, c)$ with $w = [\tau]_{d_h}$. The full simulator repeats complete trials, freshly sampling $c \leftarrow U(\mathcal{C}_\kappa)$ and z each time, until acceptance; it uses only (A, c) and the public filter, while verification reconstructs $w = [A_0z_0 + A_1z_1 + c\hat{q}]_{d_h}$. Exact sampling therefore gives accepting-HVZK with zero statistical error for Lemma 4.17.

Lemma 4.15 (Maximum commitment probability). For every fixed statement with $A_0 \in R_q^\times$ and arbitrary $A_1 \in R_q$, the maximum point probabilities of the raw and simulator-recovered commitment distributions are at most

$$\varepsilon_{\text{com}} = \rho_{\text{pub}}^{-1} ((2^{d_h+1} + 2)p_{q, \sigma_y})^N. \quad (13)$$

Lemma 4.16 (One-forgery self-target extraction and ordinary-RSIS fork consequence). For the ordinary-RSIS fork, define on canonical ring inputs

$$H_{\text{ch}}^{\uparrow}(\mu, \text{can}_q(2^{d_h}w)) = H_{\text{ch}}(\mu, w), \quad w \in \mathbb{Z}_{p_h}^N,$$

and answer independently outside the lifted image. Every accepting no-message forgery for fixed A gives a solution to

$$H_{\text{ch}}(\mu, [A_0 Y_0 + A_1 Y_1 + c\hat{q}]_{d_h}) = c$$

with

$$B_{\text{ST}} = \sqrt{2N} Z_{\text{BOUND}}, \quad B_{\text{ST},\infty} = Z_{\text{BOUND}}. \quad (14)$$

The high-bit lift also determines E with

$$2^{d_h}[A_0 Y_0 + A_1 Y_1 + c\hat{q}]_{d_h} = A_0 Y_0 + A_1 Y_1 + c\hat{q} + E \bmod q, \quad \|E\|_{\infty} \leq R_h.$$

For the auxiliary ordinary-RSIS fork consequence, suppose in addition that

$$B_{\text{fork},\infty} = \max\{2Z_{\text{BOUND}}, 4R_h + 1\} < (q-1)/2, \quad (15)$$

and that two accepting forks have the same recovered commitment and distinct challenges. Their lifted difference solves ordinary RSIS for $\bar{A} = [2A_0 \mid 2A_1 \mid 1]$, which has Definition 4.1's distribution since q is odd. Here $B_{\text{RSIS}} = 2\sqrt{B_{\text{ST}}^2 + 4NR_h^2} + \sqrt{2\kappa}$, and independent challenges collide with probability $1/\binom{N}{\kappa}$.

Lemma 4.17 (Commitment-recoverable FSwA transfer). For a monomial-shift-bounded fixed statement, set $\text{Rec}_A(c, z) := [Az + c\hat{q}]_{d_h}$. Lemmas 4.13–4.15 give worst-case non-abort probability ρ_{na} , accepting-HVZK error $\varepsilon_{\text{zk}} = 0$, and accepted honest and simulated commitment point probabilities at most ε_{com} . For budgets Q_S, Q_H , set

$$L_{\text{FSwA}}^{\text{BM}} = \frac{2Q_S(Q_H + Q_S + 1)\varepsilon_{\text{com}}}{\rho_{\text{na}}} + \frac{Q_S(Q_S + 1)\varepsilon_{\text{com}}}{2\rho_{\text{na}}^2} + Q_S\varepsilon_{\text{impl}}. \quad (16)$$

Then a UF-CMA forger can be converted to a no-message forger with additional loss $L_{\text{FSwA}}^{\text{BM}}$.

The tables use ideal exact samplers, for which $\varepsilon_{\text{impl}} = 0$.

For fixed A , Lemma 4.16 reaches $\text{BMST-RSIS}_{H_{\text{ch}}, A, R_q, N, q, d_h, B_{\text{ST}}, B_{\text{ST},\infty}}$; sampling $A_0 \leftarrow U(R_q^{\times})$ and $A_1 \leftarrow U(R_q)$ gives Assumption 4.3's random-statement setting.

Proposition 4.1 (Fixed-statement signing reduction). For every monomial-shift-bounded fixed statement (A, S) and every forger respecting the Q_S, Q_H budgets,

$$\text{Adv}_{A;S}^{\text{PKS-UF-CMA}}(\mathcal{A}) \leq \text{Adv}_A^{\text{BMST-RSIS}}(\mathcal{C}_A) + L_{\text{FSwA}}^{\text{BM}}.$$

Under the additional ordinary-RSIS conditions in Lemma 4.16, the transcript-difference argument also gives the auxiliary fork instance of Definition 4.1, including the challenge-collision term.

Theorem 4.1 (Unit-conditioned ANSA-PKS^u signing bound). In the planted-witness ANSA-PKS^u game of Definition 2.3, let SG be the event defined above. If $\mathcal{A}$ stays within the fixed Q_S, Q_H budgets, then there exist a unit-label distinguisher $\mathcal{D}_{\text{lab}}$ and a self-target solver $\mathcal{C}$ such that

$$\Pr[\mathcal{A} \text{ wins} \wedge \text{SG}] \leq \text{Adv}_{\sigma_1}^{\text{u1-lab}}(\mathcal{D}_{\text{lab}}) + \text{Adv}^{\text{BMST-RSIS}}(\mathcal{C}) + L_{\text{FSwA}}^{\text{BM}}.$$

4.4 Full IBS Reduction

This subsection assembles the identity-to-signing reduction. The deterministic extraction-tape hop contributes $\text{Adv}_{\text{ext-rand}}(\mathcal{F}_{G_{\text{ext}}})$. The extraction-distribution replacement affects only identity-key generations; pure identity-label queries do not generate real extraction witnesses. Witness programming affects all generated identity labels.

Let r_p be the per-key Rényi factor from Lemma 4.3, and let $r_u^\times$ and $\Gamma_{\text{wp}}^\times$ be the unit-label quantities from Lemma 4.4. Let $\varepsilon_{\text{id-coll}}$ be the label-collision probability after Lemma 4.5:

$$\varepsilon_{\text{id-coll}} \leq \binom{Q_{\text{Hid}}}{2} \frac{1}{|R_q^\times|} = \binom{Q_{\text{Hid}}}{2} \frac{1}{q^N \rho_\times}$$

for Table 1. The external-message public-digest collision term is added in Corollary 4.1.

After witness-first programming, the collision term is bounded using $\Gamma_{\text{wp}}^\times$ from Lemma 4.4:

$$\varepsilon_{\text{wf-coll}} \leq \binom{Q_{\text{Hid}}}{2} \Gamma_{\text{wp}}^\times.$$

Define

$$\Delta_{\text{PKS}}^{\text{BM}} = \text{Adv}_{H_{\text{ch}}, R_q, N, q, d_h, B_{\text{ST}}, B_{\text{ST}, \infty}}^{\text{BMST-RSIS}}(\mathcal{C}) + L_{\text{FSwA}}^{\text{BM}}.$$

Set

$$\Delta_{\text{low}}^{\text{BM}} = \Delta_{\text{ann}} + \Delta_{\text{lab}} + \Delta_{\text{PKS}}^{\text{BM}},$$

where Δ_{ann} is the advantage from Assumption 4.1 and $\Delta_{\text{lab}} = \text{Adv}_{\sigma_1}^{\text{u1-lab}}$ is the uniform-key advantage from Assumption 4.2.

Lemma 4.18 (Identity localization and lower-game bound). Let ε be the success probability in the owner-tagged uncentered witness-first game on the event that all sampled witnesses are monomial-shift-bounded. Then there exist an annular-NTRU distinguisher, a unit-label distinguisher, and a BMST-RSIS solver whose advantages are collected in $\Delta_{\text{low}}^{\text{BM}}$ such that

$$\varepsilon \leq 3Q_{\text{Hid}}\Delta_{\text{low}}^{\text{BM}} + \varepsilon_{\text{wf-coll}}.$$

The proof localizes first in the annular-key game, then switches the resulting planted-witness signing game to uniform h and applies Theorem 4.1.

Let $\varepsilon_{\text{sh}}^{\text{wf}}$ and $\varepsilon_{\text{wrap}}^{\text{wf}}$ be the monomial-shift failure and centering-change probabilities from Lemma 4.12. Define their aggregate contribution by

$$\Theta_{\text{wf}} = 3Q_{\text{Hid}}\Delta_{\text{low}}^{\text{BM}} + \varepsilon_{\text{wf-coll}} + \varepsilon_{\text{sh}}^{\text{wf}} + \varepsilon_{\text{wrap}}^{\text{wf}}.$$

Theorem 4.2 (Owner-tagged digest-message IBS reduction). Let $\mathcal{A}$ be an ID-UF-CMA adversary against $\text{ANSA-IBS}^{\mu, \text{own}}$ in the random-oracle model that respects freshness on pairs (id, μ) and the budgets above. For Rényi orders $a_p, a_u > 1$ and $\varepsilon_{\text{unit}}$ as in Lemma 4.5, there exist $\mathcal{F}_{G_{\text{ext}}}$, $\mathcal{D}_{\text{ann}}$, $\mathcal{D}_{\text{lab}}$, and $\mathcal{C}$ such that

$$\begin{aligned} \text{Adv}_{\text{ANSA-IBS}^{\mu, \text{own}}}^{\text{ID-UF-CMA, ROM}}(\mathcal{A}) &\leq \varepsilon_{\text{unit}} + \varepsilon_{\text{id-coll}} + \text{Adv}_{\text{ext-rand}}(\mathcal{F}_{G_{\text{ext}}}) \\ &\quad + \left[r_p^{Q_{\text{ID}}} \left((r_u^\times)^{Q_{\text{Hid}}} \Theta_{\text{wf}} \right)^{(a_u-1)/a_u} \right]^{(a_p-1)/a_p}. \end{aligned}$$

Reduction work. The BMST solver uses one adversary run and one localization guess, aborting if a sampled non-target witness is not monomial-shift-bounded (Lemma 4.18). Let $t_{\mathcal{A}}$ be the adversary’s expected work, τ_{pub} a uniform upper bound on the expected cost of one public-simulator trial, and τ_{wf} a corresponding bound for one ambient witness proposal, including syndrome computation and the unit test. For the selected profiles, the sampling bound in Appendix A.5 gives

$$t_{\mathcal{C}} \leq t_{\mathcal{A}} + \frac{Q_S \tau_{\text{pub}}}{\rho_{\text{pub}}} + 1.267(Q_{\text{Hid}} + 1)\tau_{\text{wf}} + t_{\text{loc}} + t_{\text{aux}},$$

where t_{loc} is the expected work for local signing under non-target identities, including all mask and transport retries, and t_{aux} is the expected cost of the remaining candidate-counter simulation, witness-bound checks, key preparation, hashing, ring arithmetic, oracle bookkeeping, and final verification. Appendix A.6 uses the resulting solver-to-adversary work ratio $r_T = t_{\mathcal{C}}/t_{\mathcal{A}}$.

Corollary 4.1 (External-message wrapper). Let $\mathcal{A}$ be an adaptive ID-UF-CMA adversary against the external-message construction with public-digest budget Q_μ . Then there is an adversary $\mathcal{A}_\mu$ in $\text{ANSA-IBS}^{\mu, \text{own}}$ with the same identity, signing, and Fiat-Shamir query counts such that

$$\text{Adv}_{\text{ANSA-IBS}}^{\text{ID-UF-CMA, ROM}}(\mathcal{A}) \leq \text{Adv}_{\text{ANSA-IBS}^{\mu, \text{own}}}^{\text{ID-UF-CMA, ROM}}(\mathcal{A}_\mu) + \varepsilon_{\mu\text{-coll}},$$

where

$$\varepsilon_{\mu\text{-coll}} \leq \binom{Q_\mu}{2} 2^{-2\lambda}.$$

Because the digest-collision event precedes the owner-tagged digest-message reduction, its probability enters the bound once. Appendix A.6 evaluates the concrete reduction losses.

5 Concrete Parameters and Security Estimates

We instantiate four parameter sets. Under the direct-attack estimates below, IBS-1/2/3 target NIST security categories 2, 3, and 5, respectively; IBS-1s provides a lighter profile without a category claim. The analysis covers master-trapdoor recovery, preimage recovery at the sampled-key radius, alternate user-

key recovery, and generic forgery against the fixed-statement BMST-RSIS relation. These heuristic estimates use the GSA and Core-SVP conventions specified below. Appendix A.6 separately quantifies the guarantees and losses of the classical-ROM reduction.

5.1 Parameter Selection and Proof Conditions

For every parameter set, ϕ is squarefree modulo q . Writing $\phi = \prod_{i=1}^t \phi_i$ over $\mathbb{F}_q$, where $\deg(\phi_i) = f_i$, the Chinese remainder theorem gives $R_q \simeq \prod_{i=1}^t \mathbb{F}_{q^{f_i}}$. A uniform output of H_{cand} is therefore invertible with probability

$$\rho_{\times} = \prod_{i=1}^t (1 - q^{-f_i}).$$

Thus H_{id} uses $1/\rho_{\times}$ trials on average, and a union bound over at most Q_{Hid} label states bounds exhaustion of its 16-bit counter by

$$\varepsilon_{\text{unit}} \leq Q_{\text{Hid}}(1 - \rho_{\times})^{2^{16}}.$$

We set $d_h = 13$ and derive p_h , r_h , and R_h from Eq. (2).

The evaluation takes the identity-label budgets $Q_{\text{Hid}} \leq 2^{24}$ for IBS-1s and $Q_{\text{Hid}} \leq 2^{48}$ for IBS-1/2/3 as premises on the adversary, like Q_H and Q_S ; they count identities queried without extraction. Since witnesses are generated only for reached label states, $Q_{\text{ID}} \leq Q_{\text{Hid}}$. All profiles allow $Q_S, Q_H, Q_{\mu} \leq 2^{64}$; these queries may reuse identities.

All four sets use $\xi = 0.9$, $\alpha_{\text{ann}} = 1.17$, and $R_{\text{ext}} = 1.29$. Equation (1) determines σ_1 , and the bound $Q_{\text{hyb}}(B) \leq \alpha_{\text{ann}}\sqrt{q}$ ensures $\sigma_1/Q_{\text{hyb}}(B) \geq R_{\text{ext}}$ in both the Hybrid-Sampler and quotient-smoothing analyses.

The choice $\alpha_{\text{ann}} = 1.17$ matches Falcon’s conventional Gram–Schmidt quality target [30] and supports an alternative private-key generator if a weakness specific to Antrag is found. Appendix C specifies Falcon-style key generation and Gaussian preimage sampling over the same (R, q) at the current witness width σ_1 . After issuing new master and identity keys, the user-side signing and verification implementations, parameters, and key and signature formats remain unchanged. Theorem C.1 establishes this substitution under decisional NTRU for the specified Gaussian key distribution, with explicit extraction-error bounds.

In the trinomial rings, σ_1 denotes a canonical width; the corresponding continuous-Gaussian coefficient marginal scale is $\sqrt{4/3}\sigma_1$. The thresholds in Table 1 use Lemma 4.12 with the conservative design allowance $Q_{\text{sh}} = 2^{64}$ and $\lambda_{\text{sh}} = (110, 163, 234, 271)$. This gives the reported ceilings B_{shift} and also covers the smaller identity-label budgets: the lifetime shift-tail bound improves by the factor $Q_{\text{Hid}}/Q_{\text{sh}}$. The remaining signing parameters satisfy

$$\alpha_{\text{iter}} = \frac{\sigma_y}{B_{\text{shift}}}, \quad \varepsilon_{\perp} = 2^{-20},$$

$$(1 - \varepsilon_{\perp})M_{\text{iter}} \geq M_{\alpha_{\text{iter}}}.$$

For each profile, M_{iter} is the least rational with its indicated denominator that satisfies this inequality; explicitly,

$$\begin{aligned} M_{\text{iter},1s} &= \frac{10174425311254520}{10^{16}}, & M_{\text{iter},1} &= \frac{100182058729041}{10^{14}}, \\ M_{\text{iter},2} &= \frac{1000917274631927}{10^{15}}, & M_{\text{iter},3} &= \frac{10125832749231644}{10^{16}}. \end{aligned}$$

Lemma 4.9 therefore gives the uniform pointwise bound $\Pr[\mathcal{R}_v(y) = \perp] \geq 2^{-20}$ for every admitted shift and input.

We set $Z_{\text{BOUND}} = \lceil 4.25\sigma_y \rceil$ for IBS-1s/1/2. For IBS-3, we choose $\sigma_y = 507600$ with $Z_{\text{BOUND}} = 2018750$ fixed, balancing transport rejection against response filtering while retaining 22-bit response fields. Using $\sigma_y = 475000$ instead would facilitate a common mask sampler across IBS-1/2/3, with mean signing-attempt bounds of approximately 3.4482 for $\kappa = 40$ and 2.9902 for $\kappa = 35$.

These choices are made as follows. Since $q - 1 = 1023 \cdot 2^{13}$, $d_h = 13$ gives $r_h = 0$ and hence $R_h = 4097$, whereas $d_h = 14$ would give $R_h = 16385$. The multiplier 4.25 balances response-cap rejection against the BMST-RSIS bound within the 22-bit fields. Each λ_{sh} exceeds the classical BMST-RSIS estimate in Table 2 by 4 bits (10 for IBS-1s), keeping the shift-failure term below that estimate. The challenge weights give 192.3 bits of entropy for IBS-1, close to ML-DSA-44's 192.8 bits, and 287.1 bits for IBS-3, about 31 bits above the 256-bit reference. Challenge entropy and lattice-work estimates are reported separately in Table 2.

The resulting upper bound on the expected number of signing attempts is $M_{\text{iter}}^\kappa / \rho_{\text{pub}}$. For $Q_S = Q_H = 2^{64}$, Appendix A.6 evaluates ρ_{pub} , the fixed-statement FSwA loss from Eqs. (11)–(16), and the remaining proof conditions. The fixed-format signature length in bytes is

$$\left\lceil \frac{16 + \kappa \lceil \log_2 N \rceil + 2N \lceil \log_2 (2Z_{\text{BOUND}} + 1) \rceil}{8} \right\rceil,$$

where the three terms encode, respectively, the identity-label counter, the challenge support, and the two centered response polynomials.

5.2 Master-Key Security

Direct annular-NTRU recovery. We estimate direct trapdoor recovery using Antrag's projection-and-sieving analysis [14, Sec. 4.2.1]. Define the normalized pre-rounding radii

$$a_{\text{low}} = \frac{\alpha_{\text{ann}} + \alpha_{\text{ann}}^{-1}}{2} - \frac{\xi}{2}(\alpha_{\text{ann}} - \alpha_{\text{ann}}^{-1}), \quad a_{\text{high}} = \frac{\alpha_{\text{ann}} + \alpha_{\text{ann}}^{-1}}{2} + \frac{\xi}{2}(\alpha_{\text{ann}} - \alpha_{\text{ann}}^{-1}),$$

and put

$$A_{\text{ann}} = \frac{N}{6q} + \frac{a_{\text{low}}^2 + a_{\text{high}}^2}{2}.$$

Table 1. Proposed parameters and repetition bounds for two-response bimodal ANSA-IBS. Attempt bounds are rounded upward.

metric	IBS-1s	IBS-1	IBS-2	IBS-3
N	1024	1536	2048	2304
ring	Φ_{2048}	Φ_{4608}	Φ_{4096}	Φ_{6912}
q	8380417	8380417	8380417	8380417
digest bit-length $ \mu $	256	256	384	512
κ	20	27	31	40
σ_1	4369.265	4369.265	4369.265	4369.265
B_{shift}	250356	362289	348139	442113
α_{iter}	1.118	1.311	1.364	1.148
M_{iter}	1.017443	1.001821	1.000917	1.012583
M_{iter}^κ	1.413	1.050	1.029	1.649
σ_y	$2.800 \cdot 10^5$	$4.750 \cdot 10^5$	$4.750 \cdot 10^5$	$5.076 \cdot 10^5$
Z_{BOUND}	1190000	2018750	2018750	2018750
response bits/coeff	22	22	22	22
d_h / R_h	13/4097	13/4097	13/4097	13/4097
expected identity-map trials	1.0001222	1.0000000	1.0002444	1.0000000
proof signing-attempt bound	1.523	1.151	1.157	2.381
fixed signature size	5659 B	8488 B	11309 B	12734 B

For a power-of-two ring, A_{ann} is the squared-norm scalar in Antrag’s $\sqrt{q}$ normalization. In either trinomial ring, the normalized canonical root determinant is $\sqrt{q}(3/4)^{1/4}$. Accordingly, define

$$A_{\text{dir}} = \begin{cases} A_{\text{ann}}, & M = 2N, \\ \frac{2}{\sqrt{3}} A_{\text{ann}}, & M = 3N. \end{cases}$$

For the trinomial profiles, the coefficient-basis moment bound of [14, Appendix B] is looser and yields a higher block size. We report the smaller, attacker-favorable canonical estimate.

For

$$\delta_\beta = \left(\frac{(\pi\beta)^{1/\beta} \beta}{2\pi e} \right)^{1/(2(\beta-1))},$$

the reported block size is the least β satisfying

$$A_{\text{dir}} \leq \frac{8N}{3\beta} \delta_\beta^{4(\beta-N)+2}. \quad (17)$$

For a rank- d lattice, the GSA convention used here assigns the i th Gram–Schmidt length $\det(\Lambda)^{1/d} \delta_\beta^{d-2i+1}$ and evaluates the projection screen at the adjacent profile point $i = d - \beta$, including its extrapolated endpoint $i = 0$. Squaring gives exponent $4\beta - 2d + 2$, hence Eq. (17) for $d = 2N$. This adjacent-vector convention is more favorable to the attacker than evaluation at $i = d - \beta + 1$. Table 2 reports the resulting block sizes and no-offset Core-SVP costs $[0.292\beta]$ and $[0.265\beta]$.

Overstretched-NTRU attacks. The attacks of Kirchner–Fouque and Ducas–van Woerden exploit the unusually dense rank- N submodule generated by the ring multiples of a conventional small NTRU secret [20,13]; Karenin and Kirshanova formulate this phenomenon in a broader dense-submodule framework [18]. For reference, writing $q = N^a$ gives $a = 2.300, 2.173, 2.091, 2.059$ for IBS-1s/1/2/3, all below the Kirchner–Fouque asymptotic crossover $a \approx 2.783$ for ternary secrets. The finite-dimensional diagnostic points in the other direction: extrapolating the Ducas–van Woerden empirical proxy $q_{\text{fat}} \approx 0.0038N^{2.484}$ places q above this proxy by factors of approximately 73, 27, 13, 10 for the four sets, respectively. The former threshold is asymptotic, whereas the latter includes an empirically fitted constant. More importantly, both assume constant-width, small-coefficient secrets and therefore do not directly apply to the Antrag distribution.

The appropriate diagnostic for annular keys is instead the covolume of this submodule. Let $\Lambda_{f,g} = \{(rf, rg) : r \in R\}$ and put $s_i = |\phi_i(f)|^2 + |\phi_i(g)|^2$. In the canonical embedding, its root covolume relative to the domain lattice R is

$$\left(\frac{\text{vol}(\Lambda_{f,g})}{\text{vol}(R)} \right)^{1/N} = \left(\prod_{i=1}^N s_i \right)^{1/(2N)} \in \left[\frac{\sqrt{q}}{\alpha_{\text{ann}}}, \alpha_{\text{ann}} \sqrt{q} \right].$$

The rank-two NTRU lattice has relative root determinant $\sqrt{q}$. The annular acceptance condition therefore keeps the secret-generated submodule within a factor $\alpha_{\text{ann}} = 1.17$ of the balanced scale, rather than at the substantially smaller scale required by the overstretched-NTRU mechanism. These attacks consequently do not supersede the direct annular-recovery estimate above. They do not, however, account for the distribution-specific information in $ff^* + gg^*$, which motivates the relative-norm analysis below.

Relative-norm recovery. Antrag’s second strategy uses $H = hh^*$ to expose an approximate relation in the maximal real subfield $\mathcal{K}^+$. Put $R^+ = R \cap \mathcal{K}^+$, $A = ff^*$, $B = gg^*$, and $W = A + B$. The lattice

$$\Lambda = \{(A', W') \in (R^+)^2 : A'(1 + H) \equiv W' \pmod{q}\}$$

contains (A, W) . Let $s = N/2$ and define the real-subfield covolume factor

$$C_+ = \text{vol}(R^+)^{2/s} = \frac{|\Delta_{\mathcal{K}^+}|^{1/s}}{s}.$$

The normalized squared root determinant of Λ is therefore

$$V^2 = \det(\Lambda)^{2/N} = C_+ q, \quad C_+ = \begin{cases} 2^{1-2/N}, & M = 2N, \\ \sqrt{3}, & M = 3N. \end{cases}$$

For the trinomial case, this follows from $|\Delta_{\mathcal{K}^+}|^2 = |\Delta_{\mathcal{K}}|$ and $|\Delta_{\mathcal{K}}|^{1/N} = N\sqrt{3}/2$. The analysis below therefore uses the exact real-subfield volume rather than extending the power-of-two factor to Φ_{4608} and Φ_{6912} .

To formulate the decoding problem, we center both public coordinates. Put $\rho_-^2 = qa_{\text{low}}^2$ and $\rho_+^2 = qa_{\text{high}}^2$, and let

$$\bar{S}_N = \frac{1}{N} \sum_{\theta} \left| \frac{\sin(N\theta)}{\sin \theta} \right|^2,$$

where the sum ranges over the primitive embeddings. Here $\bar{S}_N = 0$ for $M = 2N$ and $\bar{S}_N = 3N/4$ for $M = 3N$ with $4 \mid M$, as in both selected trinomial profiles. Under Anrag's Heuristic 1, define

$$\bar{W} = \frac{N}{6} + \frac{\rho_+^2 + \rho_-^2}{2}, \quad E = W - \bar{W}, \quad D = A - B,$$

and

$$J = \frac{N(\rho_+^2 + \rho_-^2)}{12} + \frac{N^2 + \bar{S}_N}{72}.$$

The moments induced by Anrag's Algorithm 4, expressed in the normalized canonical metric, are

$$q^2 y = \mathbb{E} \|E\|^2 = \frac{(\rho_+^2 - \rho_-^2)^2}{12} + J, \quad q^2 v = \mathbb{E} \|D\|^2 = \frac{\rho_+^4 + \rho_+^2 \rho_-^2 + \rho_-^4}{6} + J.$$

Appendix A.7 derives these moments and their component normalization, including the distinction from the total energy in [14, Lemma 5].

The public center defines the target $t = (\bar{W}/2, \bar{W})$, whose offset from (A, W) is $(A - \bar{W}/2, E)$. Exchange symmetry gives the normalized second-moment matrix

$$q^2 \begin{bmatrix} (y+v)/4 & y/2 \\ y/2 & y \end{bmatrix}.$$

Its determinant factor is $z = yv/4$, and the optimal determinant-one quadratic twist has expected squared offset $2q^2 \sqrt{z} = q^2 \sqrt{yv}$. Combining this twist with the GSA projection radius used for the direct attack gives the no-drop condition

$$q\sqrt{yv} \leq \frac{4NC_+}{3\beta} \delta_{\beta}^{4\beta-2N+2}. \quad (18)$$

The right-hand coefficient specializes to $8N2^{-2/N}/(3\beta)$ for a power-of-two ring and $4N/(\sqrt{3}\beta)$ for a trinomial ring.

The dimension-dropping construction of [14, Eq. (12)] does not directly apply to this instance. Its $k = 0$ specialization does not recover the no-drop bound, and deleting identity-type basis vectors would require corresponding coordinates of the dense element $A = ff^*$ to vanish (Appendix A.7).

We nevertheless evaluate an attacker-favorable relaxation that permits the deletion of k error coordinates, centers the target, applies the optimal correlated twist, and uses the more favorable projected covolume factor 2 in the trinomial cases. Exhaustive optimization over $0 \leq k < N/2$ and $50 \leq \beta \leq N$ selects $k = 0$ for every parameter set. Equation (18) has no solution with $\beta \leq N$; even at

its closest point, its left-hand side exceeds its right-hand side by factors 6413, 4953, 3228, and 3314 for IBS-1s/1/2/3. Relative-norm decoding is therefore well outside the predicted BDD regime, leaving direct recovery as the controlling master-key attack. This conclusion inherits Antrag’s Heuristic 1, its idealized rounding model, and the GSA; it does not rule out unrelated algebraic attacks.

5.3 Galois-Symmetry Attacks

Relative-norm descent does not require h to lie in a proper subfield; exact Galois contraction does. Define

$$L_h = \{(u, v) \in R^2 : uh \equiv v \pmod{q}\},$$

and

$$\text{Stab}(h) = \{\sigma \in \text{Gal}(\mathcal{K}/\mathbb{Q}) : \sigma(h) = h\}.$$

If $G \subseteq \text{Stab}(h)$ has order $g > 1$, then L_h is G -stable. Under the relevant unramified-discriminant condition, the contraction theorem of Cramer, van Gent, Lesavourey, and Pellet-Mary [7] realizes L_h over the fixed field without loss in the approximation factor and reduces its integer rank from $2N$ to $2N/g$. The relevant failure event is the per-key equality $\sigma(h) = h$ for some nonidentity σ , not Galois invariance of the public-key distribution.

Such a contraction would substantially reduce the direct estimates. At $g = 2$, the classical costs for IBS-1s/1/2/3 would fall to 128, 199, 263, and 303 bits; at $g = 4$, they would fall further to 61, 96, 128, and 148 bits. Independence across Fourier representatives is therefore a security requirement: KeyGen must neither generate unsampled representatives from Galois orbits nor reuse randomness within an orbit.

The specified generator independently samples one representative from each conjugate pair. Hermitian symmetry supplies the other representative solely to obtain real coefficients and imposes no additional Galois symmetry. The implementation also rejects any public key fixed by a prime-order Galois element. This test detects every nontrivial exact stabilizer because each nontrivial finite subgroup contains an element of prime order.

The ideal AnnularTrapdoorGen distribution analyzed in Section 4 is defined before this screening step. Let $D_{\text{std}} = \mathcal{H}_{\text{ann}}$ denote that public-key distribution, let E_{stab} be the event $\text{Stab}(h) = \{1\}$, and define the uniform-key failure probability and the corresponding Assumption 4.1 distinguishing advantage by

$$\begin{aligned} \varepsilon_{\text{stab}, P} &:= \Pr_{u \leftarrow U(R_q)} [\neg E_{\text{stab}}], \\ \Delta_{\text{stab}, P} &:= \text{Adv}^{\text{dec-ann-NTRU}}(\mathcal{D}_{\text{stab}, P}), \quad \mathcal{D}_{\text{stab}, P}(h) = \mathbf{1}[\neg E_{\text{stab}}]. \end{aligned}$$

Algorithm B.2 computes the predicate $\mathcal{D}_{\text{stab}, P}$ by public automorphism tests, so it is efficient. Assumption 4.1 therefore gives

$$\varepsilon_{\text{stab}, P}^{\text{ann}} := \Pr_{h \leftarrow D_{\text{std}}} [\neg E_{\text{stab}}] \leq \varepsilon_{\text{stab}, P} + \Delta_{\text{stab}, P}. \quad (19)$$

For a fixed candidate map, screening gives $D_{\text{std}} \mid E_{\text{stab}}$, at total variation distance $\varepsilon_{\text{stab},P}^{\text{ann}}$ from D_{std} . When the candidate map uses an ideal XOF, the same distance bound holds for the joint oracle/key laws by the first-success argument in Section B.3. Thus Eq. (19) contributes a single additive term, independent of the numbers of identity and signing queries.

For the uniform term, squarefreeness of Φ_M modulo q identifies $R_q \otimes_{\mathbb{F}_q} \overline{\mathbb{F}}_q$ with N coordinates indexed by the primitive M -th roots. An automorphism of prime order ℓ acts on these coordinates in N/ℓ cycles, so its fixed subspace has $\mathbb{F}_q$ -dimension N/ℓ . A uniform $u \in R_q$ is therefore fixed with probability $q^{-N(1-1/\ell)}$. For IBS-1s/1/2/3, Algorithm B.2 tests, respectively, 3, 7, 3, 7 non-identity elements of order two and 0, 2, 0, 2 elements of order three. Since every nontrivial stabilizer contains an element of prime order, a union bound gives

$$\varepsilon_{\text{stab},P} \leq \begin{cases} 3q^{-512} < 2^{-11773}, & P = \text{IBS-1s}, \\ 7q^{-768} + 2q^{-1024} < 2^{-17660}, & P = \text{IBS-1}, \\ 3q^{-1024} < 2^{-23548}, & P = \text{IBS-2}, \\ 7q^{-1152} + 2q^{-1536} < 2^{-26491}, & P = \text{IBS-3}. \end{cases} \quad (20)$$

These probabilities are negligible at the precision reported in Table 2.

5.4 Sampled-Key-Radius Preimage Diagnostic

The ambient Gaussian used in the extraction comparison provides a reference scale for affine-preimage recovery. Following the affine-CVP convention of [5], we evaluate this scale in both coefficient and normalized canonical coordinates. Put

$$\text{GH}(\beta) = \sqrt{\frac{\beta}{2\pi e}}, \quad \nu_P = \begin{cases} 1, & P \text{ is power-of-two}, \\ 4/3, & P \text{ is trinomial}, \end{cases}$$

and let $\gamma_P = 1$ for a power-of-two ring and $\gamma_P = (3/4)^{1/4}$ for a trinomial ring. In coefficient coordinates, the continuous-Gaussian RMS proxy is $\sigma_1 \sqrt{2N\nu_P}$, and its affine-CVP crossing condition at block size β is

$$\sigma_1 \sqrt{2N\nu_P} \geq \text{GH}(\beta)^{2N/(\beta-1)} \sqrt{q}.$$

In normalized canonical coordinates, the corresponding condition is

$$\sigma_1 \sqrt{2N} \geq \text{GH}(\beta)^{2N/(\beta-1)} \sqrt{q} \gamma_P.$$

For each profile, we take the least β in $50 \leq \beta \leq 2N$ satisfying each crossing condition and retain the smaller result. This gives 984, 1507, 2173, 2386 for IBS-1s/1/2/3. The two metrics coincide for the power-of-two profiles; for the trinomial profiles, the selected coefficient-coordinate values 1507 and 2386 compare with canonical-coordinate values 1538 and 2431. Applying the no-offset Core-SVP conversions $\lfloor \beta \log_2 \sqrt{3/2} \rfloor$ and $\lfloor \beta \log_2 \sqrt{13/9} \rfloor$ gives 287/261, 440/399, 635/576, and 697/632 bits, as reported in Table 2. These figures describe the sampled-key scale, not recovery of the particular extraction sample. Section 5.5 considers alternate user-key recovery with a larger witness-norm allowance.

5.5 Alternate User-Key Recovery

An affine-preimage attack seeks another witness for the targeted user's identity relation, allowing larger witness norms than the sampled-key scale of Section 5.4. Fix an identity with label (j, a) and normalized statement $A = \hat{q}a^{-1}(1, h)$. For $S_{\text{atk}} = (s'_0, s'_1)^T \in R^2$, invertibility of a and $\hat{q}$ gives

$$AS_{\text{atk}} \equiv \hat{q} \pmod{q} \iff s'_0 + hs'_1 \equiv a \pmod{q}.$$

Thus the recovered preimage satisfies the same witness relation as the issued user key. We call it an alternate user key when it supports repeated signing through the attacker's procedure; it need not equal the issued witness or satisfy the honest signer's norm-admission checks. To sign a message with digest μ , choose $Y \in R^2$ and set

$$w_Y = [AY \bmod q]_{d_h}, \quad c = H_{\text{ch}}(\mu, w_Y), \quad z = Y - cS_{\text{atk}}.$$

Since $Az + c\hat{q} \equiv AY \pmod{q}$, verification recovers w_Y and accepts whenever $\|z_0\|_\infty, \|z_1\|_\infty \leq Z_{\text{BOUND}}$. These bounds also imply the Euclidean bound of Eq. (14). The larger witness-norm allowance leaves this public response bound unchanged. Efficient use of one recovered key requires a sufficient probability of passing the response test, with fresh commitments w_Y on retries for a prescribed digest μ ; repeating $(\mu, 0)$ reuses the same oracle answer.

For the single-forgery estimate, take $Y = 0$, so $w_Y = 0$ and $z = -cS_{\text{atk}}$. Each trial decodes a new preimage and uses a fresh input $(\mu, 0)$ to H_{ch} , giving a uniform challenge in $\mathcal{C}_\kappa$. The affine-CVP model of Section 5.4 uses radius

$$L_P(\beta) = \text{GH}(\beta)^{2N/(\beta-1)} \sqrt{q} \gamma_P$$

and models independent decodes, conditional on the reused basis, by identically distributed vectors with canonical norm at most $L_P(\beta)$ and uniform spherical directions in $\mathbb{R}^{2N}$, independent of the challenges. The radius and decoder-output law are heuristic; the probabilities in this calculation condition on that basis. For $M_c : x \mapsto cx$, orthonormality gives $\|M_c\|_{\text{can} \rightarrow \infty}^2 = \kappa$ in a power-of-two ring. In a trinomial ring, write $N = 2m$ and

$$p(c) = |\{0 \leq j < m : c_j = c_{j+m} = 1\}|.$$

The exact support-dependent norm is

$$\|M_c\|_{\text{can} \rightarrow \infty}^2 = \frac{4}{3}(\kappa + p(c)).$$

Identify linear functionals with their coefficient vectors and use the dual inner product $\langle u, v \rangle_{\text{can},*} = u^T H_N^{-1} v$. For $T_{r,j}(x) = \text{coef}(X^j x)_r$, the relation $X^{2m} = X^m - 1$ gives

$$\langle T_{r,j}, T_{r,k} \rangle_{\text{can},*} = \frac{4}{3} \mathbf{1}_{j=k} + \frac{2}{3} \mathbf{1}_{|j-k|=m}.$$

Summing over the support gives every output row squared dual norm $4(\kappa + p(c))/3$, proving the formula.

A uniform weight- κ support contains p complete pairs with probability

$$\omega_{m,\kappa,p} = \frac{\binom{m}{p} \binom{m-p}{\kappa-2p} 2^{\kappa-2p}}{\binom{2m}{\kappa}}.$$

Here $\max\{0, \kappa - m\} \leq p \leq \lfloor \kappa/2 \rfloor$. Write $Z_P = Z_{\text{BOUND}}$. Set $(D_{P,0}, \omega_{P,0}) = (\sqrt{\kappa}, 1)$ in a power-of-two profile, and $D_{P,p} = \sqrt{4(\kappa + p)/3}$, $\omega_{P,p} = \omega_{m,\kappa,p}$ otherwise. Conditional on the support class p , each squared response coefficient, divided by $(D_{P,p} L_P(\beta))^2$, is stochastically dominated by $V \sim \text{Beta}(1/2, (2N - 1)/2)$. A union bound over $2N$ coefficients therefore lower-bounds the modeled success probability $p_{\beta,P}^{\text{resp}}$ by

$$\pi_{\beta,P} = \sum_p \omega_{P,p} \left[1 - 2N \Pr \left[V > \left(\frac{Z_P}{D_{P,p} L_P(\beta)} \right)^2 \right] \right]_+.$$

Here $[x]_+ = \max\{x, 0\}$; no independence among response coefficients is required. Reusing one block-reduced basis for independent randomized affine decodes gives modeled expected work

$$W_{\beta,P}^{(s)} = 2^{s\beta} + \frac{1}{p_{\beta,P}^{\text{resp}}} \leq 2^{s\beta} + \frac{1}{\pi_{\beta,P}} \quad (\pi_{\beta,P} > 0).$$

Use $s_C = \log_2 \sqrt{3/2}$ and $s_Q = \log_2 \sqrt{13/9}$, classical repetition in both columns, and unit cost for the polynomial-time trial work. Minimizing the bound over $50 \leq \beta \leq 2N$ gives block sizes 485/485, 776/776, 1114/1114, and 1332/1332, with floored costs 141/128, 226/205, 325/295, and 389/353 bits for IBS-1s/1/2/3. The corresponding success lower bounds $\pi_{\beta,P}$ are about 7.02%, 0.43%, 3.62%, and 0.06% per trial in profile order. These are sufficient-success estimates under the decoding model, not lower bounds on attack cost. A zero certificate is vacuous, without implying $p_{\beta,P}^{\text{resp}} = 0$.

The certificate bounds success averaged over decoded preimages and challenges; it does not give the same bound for every fixed preimage. A deterministic sufficient condition for reuse is

$$\|S_{\text{atk}}\|_{\text{can}} \leq \frac{Z_{\text{BOUND}}}{\max_{c \in \mathcal{C}_\kappa} \|M_c\|_{\text{can} \rightarrow \infty}}.$$

Indeed, the operator-norm bound then gives $\|cS_{\text{atk}}\|_\infty \leq Z_{\text{BOUND}}$ for every challenge, so $Y = 0$ signs every message. At the reported block sizes, $L_P(\beta)$ exceeds this threshold; the modeled norm bound alone therefore does not establish universal success.

For the alternate user-key recovery screen, we grant the attacker a usable fixed witness at expected one-time cost at most $2^{s\beta} + 1/\pi_{\beta,P}$. We further assume efficient generation of fresh oracle inputs for each prescribed digest, with

response-test success at least $\pi_{\beta,P}$ conditional on every preceding trial. This assumption includes the contribution of Y to the response bound. Obtaining an alternate key with this per-trial signing guarantee is an attacker-favorable assumption, beyond the spherical average. For Q_F prescribed messages, each allowed at most T fresh trials, define

$$\overline{W}_{\text{reuse},\beta,P}^{(s)}(Q_F) := 2^{s\beta} + \frac{Q_F + 1}{\pi_{\beta,P}} \quad (\pi_{\beta,P} > 0).$$

This upper-bounds expected total work: key recovery is paid once, and each message uses at most $\sum_{t=0}^{T-1} (1 - \pi_{\beta,P})^t \leq 1/\pi_{\beta,P}$ trials in expectation. Conditional on obtaining the granted alternate key, the signing procedure succeeds on any prescribed message with probability at least $1 - (1 - \pi_{\beta,P})^T$. High success therefore comes from retries with the same key. Over Q_F messages, the probability that at least one remains unsigned is at most $Q_F(1 - \pi_{\beta,P})^T$. These bounds require fresh oracle inputs throughout the retry budget.

5.6 BMST-RSIS Screen and Consolidated Estimates

The fixed-statement BMST-RSIS relation governs the signing reduction. For its one-message lattice screen, fix w , set $c = H_{\text{ch}}(\mu, w)$, and consider the lift

$$A_0 Y_0 + A_1 Y_1 + c\hat{q} + E \equiv 2^{d_h} w \pmod{q},$$

where $\|Y_i\|_\infty \leq Z_{\text{BOUND}}$ for $i = 0, 1$ and $\|E\|_\infty \leq R_h$. For an attacker-favorable lattice screen, we discard the Euclidean restriction in Eq. (14), assign all three blocks the common bound $\tilde{B}_{\text{ST},\infty} = Z_{\text{BOUND}} + R_h$, and apply the forgotten- q ℓ_∞ -SIS convention of Dilithium Appendix C.3. The BMST-RSIS costs use the no-offset Core-SVP conversions $\lfloor \beta \log_2 \sqrt{3/2} \rfloor$ and $\lfloor \beta \log_2 \sqrt{13/9} \rfloor$.

Under the distinct-challenge and no-wrap conditions of Lemma 4.16, a successful forgery fork yields an ordinary-RSIS solution with common bound $B_{\text{fork},\infty}$. This implication is one-way: an ordinary-RSIS solver does not by itself produce a forgery. The conventional forgotten- q screen gives 89/81, 143/130, 209/189, and 242/220 classical/quantum bits for IBS-1s/1/2/3.

For small moduli, however, q -ary lattices may retain the exploitable Z-shaped BKZ profile analyzed in [10]. Combining its moving-window profile with the closed-form ℓ_∞ tail model of [1] gives first-hit diagnostics of 100/90, 51/46, 71/64, and 81/74 bits; allowing independent retries and taking the minimum of the two screens gives 89/81, 50/46, 71/64, and 81/73 bits. Both cited analyses concern unstructured matrix ISIS, not structured homogeneous Ring-SIS. These values are therefore sensitivity diagnostics for the auxiliary fork relation, rather than direct forgery costs. Table A.1 separately verifies the coefficientwise no-wrap condition required by the fork.

The quantum columns describe lattice-algorithm costs; the reduction applies in the classical ROM.

The corresponding BMST-RSIS Z-shape screen gives higher costs than the generic-forgery row in Table 2. Across all four profiles, that row gives the lowest

Table 2. Concrete security estimates and structural screens for the parameter sets in Table 1.

estimate	IBS-1s	IBS-1	IBS-2	IBS-3
master-key recovery β	903	1395	1843	2111
master-key recovery classical/quantum	263/239	407/369	538/488	616/559
relative-norm BDD crossing for $\beta \leq N$	none	none	none	none
sampled-key-radius key recovery β	984	1507	2173	2386
sampled-key-radius key recovery classical/quantum	287/261	440/399	635/576	697/632
alternate user-key recovery β (model bound)	485	776	1114	1332
alternate user-key recovery classical/quantum (model upper bound)	141/128	226/205	325/295	389/353
generic forgery BMST-RSIS ℓ_∞ screen β	344	545	788	914
generic forgery BMST-RSIS screen classical/quantum	100/91	159/144	230/209	267/242
challenge-space entropy λ_{ch}	138.7	192.3	228.0	287.1

reported lattice-work estimate for a first forgery among the four assessed attack classes.

Larger fully split modulus. Choosing the larger 24-bit prime $q = 16588801$ fully splits all four rings and improves the auxiliary ordinary-RSIS estimates under the adapted Z-shape model, which accounts explicitly for q . Preserving the Table 1 proof-attempt bounds for IBS-1s/1/2, and taking $\sigma_y = 738545$ for IBS-3, gives Z-shape diagnostics of 101/91, 150/136, 216/195, and 233/212 classical/quantum bits, respectively. The IBS-3 retuning has a 1.600 proof-attempt bound and fork margin 2016766; IBS-1/2/3 then require 23-bit response coefficients. Complete splitting could also improve arithmetic performance for IBS-1 and IBS-3. By comparison, the nearby prime $q = 8404993$ fully splits all four rings and preserves the sizes in Table 1, but enlarges the fork margin only slightly. These auxiliary ordinary-RSIS figures are not signing-security estimates. We retain the ML-DSA modulus to reuse its modular-reduction machinery and keep signatures smaller.

6 Fixed-Point Implementation and Performance Evaluation

The implementation follows the integer realization specified and proved in Appendix B. For the Apple M5 measurements, we compile with Apple Clang 21.0.0 using `-O3 -mcpu=native, -march=armv8.2-a+crypto+sha3`, and loop unrolling. This build enables the two-lane Arm SHA3/SHAKE routines, Arm Montgomery NTT arithmetic with signed-Barrett reduction, and the profile-specific Karatsuba and Toom-3 transforms for the incompletely split rings. Setup and extraction use fixed-capacity integer arithmetic; the independent multiprecision reference uses GMP 6.3.0.

IBS-1s and IBS-2 use full radix-two NTTs of lengths 1024 and 2048, respectively, with arithmetic derived from ANSA-IBKEM [5]. IBS-1 and IBS-3 use incomplete NTTs with degree-three and degree-nine residues, respectively.

Setup. Measurements use the hardened generator of Appendix B. Candidate evaluation makes one Q128 pass, and unresolved comparisons reject the candidate. At the upper five power-of-two levels, the NTRU solver lifts with ex-

act small-prime NTT/CRT arithmetic and reduces with scheduled 32-bit fixed-point Babai rounds; below them, and throughout the trinomial towers, it uses scheduled fixed-capacity integer lifting and reduction. Setup includes candidate rejection, the Galois screen, the second-row norm check, and master-key encoding.

Extraction. Extraction starts from a prepared master state. It hashes the identity to its first unit label, derives the extraction tape, and runs the fixed-point Hybrid-Sampler. Its statistical-isochrony guarantee concerns the proposal-count and randomness-consumption transcript; complete calls retain the bounded rejection loop of Appendix B.

Signing and verification. Signing starts from a loaded identity key and uses Q128 transport and Bernoulli evaluation. Measurements cover the complete signing call, including retries, encoding, and temporary-data clearing. Hardening fixes the designated arithmetic schedules within each attempt; the outer retry loop remains variable. Uncached verification includes identity hashing and statement preparation, while cached verification reuses the transformed identity row. Both include message hashing, decoding, and signature checks.

Measurement protocol. Table 3 reports median throughput from three ten-second runs per setting on Apple M5. Signing cycles through 1000 messages of 32 bytes. Both verification modes use the ordinary optimized build and the same signature pool for each profile. Implementation source and benchmark data are provided in the artifact.

Table 3. Measurements on Apple M5. Throughput rows are medians of three runs in ops/s; attempt counts are pooled means, total attempts divided by total signatures over the three runs.

Metric	IBS-1s	IBS-1	IBS-2	IBS-3
fixed signature	5659 B	8488 B	11309 B	12734 B
Setup	58.697	6.751	18.880	3.547
isochronous Extract	4110.666	2096.642	1990.769	1300.452
variable-time fixed sign	11086.096	7648.270	6047.464	2554.131
hardened fixed sign	8692.555	5820.770	4890.023	2067.371
fixed verify	40860.051	25347.970	20173.500	16605.297
fixed verify, per-ID cache	100704.550	48051.747	48989.117	30474.018
avg. sign attempts (ordinary/hardened)	1.479/1.482	1.125/1.126	1.129/1.126	2.298/2.296

Table 4 reports measurements of the IBS-1 configuration on an STM32F407 Cortex-M4 at 168 MHz, using Keccak and ML-DSA-compatible NTT assembly. Working memory is confined to SRAM1, with CCM disabled. Setup and extraction run off-device. Cycle counts are means of 100 complete operations. The non-hardened and hardened signing batches produced 100 signatures in 116 and 107 attempts, respectively.

SDK Flash is the linked image size, excluding key objects, Setup and Extraction code, and the test harness. Peak RAM includes working state, the 8488-byte signature buffer, and the measured operation stack, but excludes application

Table 4. STM32F407 measurements for IBS-1; memory in bytes and cycles in millions.

Build	SDK Flash	Peak Sign RAM	Peak Verify RAM	Sign cycles	Verify cycles
non-hardened Q128	37780	28984	28564	52.851	3.726
hardened Q128	44776	31296	29780	116.285	3.730

state and interrupt-stack allowance; keys reside in flash. The hardened build uses fixed-work sampler, transport, challenge-support collection, and evaluator kernels within each attempt; the outer signing loop stops at the first accepted attempt.

Published lattice IBS sizes serve only as bandwidth references because their assumptions and analyses differ. At about 120-bit security, the earlier SIS- and Ring-SIS-based hash-and-sign schemes have signatures above 5 MB and 1 MB, respectively [22, Table 2]. The same table reports a 12918-byte signature at 60-bit security for the NTRU/Ring-SIS IBS of Xie et al. [32], whereas IBS-1 has an 8488-byte fixed signature and a 159-bit classical BMST-RSIS core estimate. Luo et al.’s construction concatenates two Schnorr-type lattice signatures; its 136- and 200-bit rows have 9738- and 14680-byte signatures [22, Table 1]. The fixed IBS-1 and IBS-2 signatures are 12.8% and 23.0% smaller than those two rows, respectively. Their classical BMST-RSIS core estimates are numerically 23 and 30 bits higher, respectively, although the notions and analyses differ. ANSA-IBS signs the normalized DLP/NTRU relation directly.

7 Conclusion

This work instantiates the convertible identification approach of Bellare, Namprempre, and Neven for the DLP/NTRU identity relation. The normalized two-response bimodal signature enables public commitment recovery, while Gärtner’s iterative rejection narrows the Gaussian mask and reduces signing attempts. The proof separates the adaptive identity-extraction loss from fixed-statement signing, reducing the latter to BMST-RSIS with explicit rejection-sampling, commitment-probability, and ordinary-RSIS fork conditions. The finite-extraction and signing analyses extend this reduction to the integer-only implementation.

8 Acknowledgements and Disclosure

The authors thank the anonymous CANS 2026 reviewers for their careful reading and constructive comments.

During manuscript preparation, the authors used OpenAI’s GPT-5-based Codex/ChatGPT for proof and notation audits, parameter checks, profiling scripts, proofreading, and build checks. Assistance with Appendix B included tuning the fixed-point parameters, developing the precision-error bounds and

their proofs, and implementing code to audit those proofs. The authors reviewed all AI-assisted text, proofs, code, tables, and estimates and remain solely responsible for the paper.

A Security Proofs and Auxiliary Analysis

The following table summarizes the notation used in the proofs and the fixed-point analysis.

Symbol	Meaning
$Q_{\text{ID}}, Q_{\text{Hid}}$	witness-generation states (including failure) and reached identity-label states
Q_S, Q_H, Q_μ	signing calls, challenge queries, and distinct digest inputs
$m_{\text{loc}}, L_{\text{loc}}$	localization multiplier $3Q_{\text{Hid}}$ and its $\log_2$
a_p, r_p	Rényi order and factor for Hybrid-Sampler witnesses
$a_u, r_u^\times$	Rényi order and factor for unit-label programming
SG	all shifts selected by generated signing challenges satisfy their norm bound
$L_P(\beta)$	heuristic affine-decoding radius in Section 5.5
J_{tr}	last retained index in Lemma 4.8
n_{word}	fractional-storage word count in Table B.3
κ_{acc}	discrete-sampler acceptance scale in Section B.4

A.1 Distributional and Label-Programming Proofs

Proof of Lemma 2.1. For a power-of-two ring, the power basis is orthonormal in the normalized canonical metric. For a trinomial ring, the Ramanujan-sum formula gives

$$\frac{1}{N} \text{Tr}_{\mathcal{K}/\mathbb{Q}} \left(\zeta_{3N}^{i-j} \right) = \begin{cases} 1, & i = j, \\ 1/2, & i - j = \pm N/2, \\ 0, & \text{otherwise,} \end{cases} \quad 0 \leq i, j < N.$$

Thus the normalized canonical Gram matrix has the blocks G of Section 2.1. Let $e_j = X^j$ and $m = N/2$, and order the power basis as

$$(e_0, \dots, e_{m-1}, e_m, \dots, e_{2m-1}).$$

The first m Gram-Schmidt vectors then have length 1, and the remaining m have length $\sqrt{3}/2$. The Gram-Schmidt norm of the power basis is therefore 1 in

either ring family. For a real d -dimensional lattice L with ordered basis C , the Micciancio–Regev smoothing bound [25], in its Gram–Schmidt refinement [15] and the standard-deviation convention, is

$$\eta_\varepsilon(L) \leq \frac{1}{\pi} \sqrt{\frac{\ln(2d(1+1/\varepsilon))}{2}} \|\tilde{C}\|_{\text{GS}}.$$

Here $\|\tilde{C}\|_{\text{GS}}$ is the maximum Gram–Schmidt length. Applying the bound in dimension N to the normalized canonical image of the power basis yields $\eta_\varepsilon(R) \leq \rho_{\text{hyb}}(N, \varepsilon)$.

Let $b_1 = (g, -f)$ and $b_2 = (G, -F)$ be the rows of B . Expanding this R -basis gives the ordered $\mathbb{Z}$ -basis

$$\mathcal{B}_{\text{exp}} = (e_0 b_1, \dots, e_{N-1} b_1, e_0 b_2, \dots, e_{N-1} b_2)$$

of $\Lambda_{h,q}$. Write $*$ for complex conjugation in $\mathcal{K}$ and set $D_B = f f^* + g g^*$. Gram–Schmidt orthogonalization over $\mathcal{K}$ gives

$$\tilde{b}_1 = b_1, \quad \tilde{b}_2 = b_2 - \frac{G g^* + F f^*}{D_B} b_1 = \left(\frac{q f^*}{D_B}, \frac{q g^*}{D_B} \right),$$

where the final identity follows from $fG - gF = q$. For $u, v \in \mathcal{K}^2$, write $\langle u, v \rangle_{\mathcal{K}} = u_0 v_0^* + u_1 v_1^*$. The vectors $\tilde{b}_1$ and $\tilde{b}_2$ are $\mathcal{K}$ -orthogonal, and their squared lengths at embedding ϕ_i are respectively s_i and q^2/s_i . This field-level Gram–Schmidt calculation [28, Lemma 6.7] takes place over $\mathcal{K}$ and therefore applies equally to the nonorthogonal trinomial power basis. The corresponding real canonical spans are also orthogonal: for $x, y \in \mathcal{K} \otimes_{\mathbb{Q}} \mathbb{R}$,

$$\langle x b_1, y \tilde{b}_2 \rangle_{\text{can}} = \frac{1}{N} \text{Tr}_{\mathcal{K}/\mathbb{Q}} \left(x y^* \langle b_1, \tilde{b}_2 \rangle_{\mathcal{K}} \right) = 0.$$

For $d = (d_0, d_1) \in \mathcal{K}^2$, Parseval’s identity in normalized embedding coordinates gives

$$\|x d\|_{\text{can}}^2 \leq \max_i \{ |\phi_i(d_0)|^2 + |\phi_i(d_1)|^2 \} \|x\|_{\text{can}}^2.$$

Consequently, the linear map $x \mapsto x d$ has canonical operator norm at most the square root of the displayed maximum. Applied to the first block of $\mathcal{B}_{\text{exp}}$, this bounds its Gram–Schmidt lengths by $\max_i \sqrt{s_i}$. Projecting the second block orthogonally away from the full span of $\mathcal{K} b_1$ yields $(e_j \tilde{b}_2)_j$, whose Gram–Schmidt lengths are at most $\max_i q/\sqrt{s_i}$. Hence the normalized canonical image of $\mathcal{B}_{\text{exp}}$ has Gram–Schmidt norm at most $Q_{\text{hyb}}(B)$. Applying the basis bound in dimension $2N$ gives

$$\eta_\varepsilon(\Lambda_{h,q}) \leq \rho_u(N, \varepsilon) Q_{\text{hyb}}(B).$$

Finally, Prest’s generalized $\mathcal{K}$ -Gram–Schmidt envelope has embedding lengths $\sqrt{s_i}$ and $q/\sqrt{s_i}$, all bounded by $Q_{\text{hyb}}(B)$. Together with the estimate for $\eta_\varepsilon(R)$, this proves the stated Hybrid-Sampler width implication. $\square$

Proof of Lemma 4.3. By Lemma 2.1, Prest’s Hybrid-Sampler theorem applies with the required embedding-wise width. Choose a lift $\bar{a} \in R$ of a , set $t_a =$

$(\bar{a}, 0)$, and write $\Lambda_{h,q} = \text{Span}_R(B)$. The theorem compares the exact sampler at center t_a with the canonical lattice Gaussian $D_{\Lambda_{h,q}, \sigma_1, t_a}$. Consider the affine map

$$\psi_{t_a} : \Lambda_{h,q} \longrightarrow P_h(a), \quad v \longmapsto t_a - v.$$

If $v = (x, y) \in \Lambda_{h,q}$, then the syndrome of $t_a - v = (\bar{a} - x, -y)$ is $\bar{a} - (x + yh) = a \bmod q$; conversely, every $S \in P_h(a)$ satisfies $t_a - S \in \Lambda_{h,q}$. Thus ψ_{t_a} is a bijection. Moreover,

$$\|v - t_a\|_{\text{can}} = \|\psi_{t_a}(v)\|_{\text{can}},$$

so the bijection preserves the relevant Gaussian norm. Its pushforward of $D_{\Lambda_{h,q}, \sigma_1, t_a}$ is therefore exactly $D_{\text{pub}}(\cdot \mid h, a)$. The choice of lift is immaterial: replacing $\bar{a}$ by another lift translates t_a by an element $(qr, 0) \in \Lambda_{h,q}$ and leaves the pushforward distribution unchanged.

For module rank r , Prest's likelihood-ratio argument gives the conservative bound $((1 + \varepsilon_p)/(1 - \varepsilon_p))^{3r} - 1$ on the relative error [28, Theorem 5.10 and proof]. At rank two this is δ'_p . Affine bijections preserve likelihood ratios, so the relative-error-to-Rényi conversion of [29] and [5, Lemma 5.2] yields the displayed factor r_p for the raw distributions. Rényi data processing under $\mathcal{C}_q$ yields the same factor for their centered pushforwards. $\square$

Proof of Lemma 4.4. Define the canonical images

$$L_0 = W_N^{\oplus 2} \mathbb{Z}^{2N}, \quad L_{h,q} = W_N^{\oplus 2} \text{cof}(\Lambda_{h,q}).$$

The syndrome map $(s_0, s_1) \mapsto s_0 + s_1 h \bmod q$ is surjective with kernel $\Lambda_{h,q}$. Since $W_N^{\oplus 2}$ is invertible, it induces an isomorphism

$$W_N^{\oplus 2} \text{cof}(S) + L_{h,q} \longmapsto s_0 + s_1 h \bmod q$$

from $L_0/L_{h,q}$ to R_q . Section 2.1 identifies $D_{R^2, \sigma_1}^{\text{amb}}$ with the spherical discrete Gaussian on L_0 under this map. Consequently, Q_h is precisely the quotient Gaussian on $L_0/L_{h,q}$, transported to R_q by the displayed isomorphism. Moreover, the condition $\sigma_1 \geq \eta_{\varepsilon_u}(\Lambda_{h,q})$ is exactly the smoothing condition for $L_{h,q}$. The quotient-Gaussian bound [5, Lemma 5.7] therefore gives $U(R_q)(a)/Q_h(a) \in [c_u^{-1}, c_u]$, equivalently $Q_h(a)/U(R_q)(a) \in [c_u^{-1}, c_u]$, for every a .

Conditioning both distributions on $R_q^\times$ multiplies this ratio by the inverse ratio of their probabilities of landing in $R_q^\times$, and this probability ratio also lies in $[c_u^{-1}, c_u]$. Therefore $Q_h^\times(a)/U(R_q^\times)(a) \in [c_u^{-2}, c_u^2]$ for every unit a .

Equivalently, $U(R_q^\times)(a)/Q_h^\times(a)$ lies in the same interval $[c_u^{-2}, c_u^2]$. Write $U(R_q^\times)(a) = Q_h^\times(a)(1 + \theta_a)$. With $\delta_u^\times = c_u^2 - 1 < 1$, the ratio bound implies $|\theta_a| \leq \delta_u^\times$. Applying the relative-error-to-Rényi conversion used in [5, Lemma 5.2] and [5, Proposition 5.1] gives the displayed bound for $r_u^\times$. Finally,

$$\max_a Q_h^\times(a) \leq c_u^2 / |R_q^\times|.$$

For squarefree $\phi = \prod_i \phi_i$ with $\deg \phi_i = f_i$, the CRT gives $|R_q^\times| = q^N \prod_i (1 - q^{-f_i}) = q^N \rho_\times$. $\square$

Proof of Lemma 4.5. For one (h, id) label state, let $C_0, \dots, C_{J_{\max}-1}$ be the independent $U(R_q)$ candidates underlying the lazy oracle H_{cand} , and let $T = \min\{i : C_i \in R_q^\times\}$, with $T = \perp$ if no such index exists. The counter-exhaustion event has probability $\nu_{\text{nonunit}}^{J_{\max}}$, and a union bound over at most Q_{Hid} label states gives the displayed term.

Conditioned on $T = t$, the variables $C_0, \dots, C_{t-1}$ are independent samples from $U(R_q)$ conditioned on being nonunits, C_t is an independent sample from $U(R_q^\times)$, and $C_{t+1}, \dots, C_{J_{\max}-1}$ remain independent unrestricted uniform samples from $U(R_q)$. This follows because the candidate labels are independent, exactly uniform elements of R_q and the event $T = t$ factors across the candidates. Hence, conditioned on no counter exhaustion, the visible counter has the truncated geometric distribution stated in Lemma 4.5, the returned unit is uniform in $R_q^\times$, and the returned unit is independent of the visible counter.

This decomposition can be simulated in any adversarial query order: sample T from the truncated geometric distribution, cache C_T , fill $i < T$ as independent uniform nonunits, and fill $i > T$ as independent unrestricted uniform elements. The oracle returns only requested candidates and reuses hidden cached entries if a later H_{id} query scans the state, so candidate-only queries before the wrapper query have the lazy H_{cand} distribution conditioned on no counter exhaustion.

Thus the Rényi witness-programming step in Lemma 4.4 is applied to the returned unit label only, while the prior nonunit candidates, later unrestricted candidates, and counter distribution are simulated exactly. The CRT identity $\nu_{\text{nonunit}} = 1 - \rho_\times$ gives the final claim. $\square$

A.2 Iterative-Rejection Proofs

Proof of Lemma 4.6. An output z can arise only from input $z + v$ through the minus branch or from input $z - v$ through the plus branch. Hence Eq. (6) gives

$$\Pr[\mathcal{R}_v(Y) = z] = p_Y(z + v)f_v(z + v) + p_Y(z - v)g_v(z - v) = p_Z(z)/M.$$

Summing over z gives success probability $1/M$; division by this probability leaves the law Z . $\square$

Proof of Lemma 4.7. We give the argument in the Gaussian notation of this paper. First suppose $\langle z, v \rangle \geq 0$. Then $\langle z + v, v \rangle \geq \|v\|_2^2$ and $\langle z - v, v \rangle \geq -\|v\|_2^2$, so the definitions give

$$\begin{aligned} & \rho_{\sigma_y}(z + v)f_v(z + v) + \rho_{\sigma_y}(z - v)g_v(z - v) \\ &= \frac{1}{M} (\rho_{\sigma_y}(z + v)S_v(z + v) + \rho_{\sigma_y}(z - v)(1 - S_v(z - v))) = \frac{\rho_{\sigma_y}(z)}{M}. \end{aligned}$$

The last equality follows by expanding the two alternating series and cancelling adjacent terms. The case $\langle z, v \rangle < 0$ is the same calculation after replacing (z, v) by $(-z, v)$. Normalization therefore gives Eq. (6) with $Y = Z = D_{\sigma_y}^{2N}$.

It remains to check that the displayed quantities are probabilities. If $\langle y, v \rangle \geq 0$, the terms in $S_v(y)$ decrease in magnitude, so pairing consecutive terms gives

$S_v(y) \geq 0$. If $\langle y, v \rangle \geq -\|v\|_2^2$, then

$$S_v(y) = 1 - \frac{\rho_{\sigma_y}(y+v)}{\rho_{\sigma_y}(y)} S_v(y+v) \leq 1.$$

These inequalities, and their versions with $-y$, show $f_v, g_v \geq 0$. Their sum is $1/M$ when $\langle y, v \rangle \geq \|v\|_2^2$ or $\langle y, v \rangle < -\|v\|_2^2$. For $-\|v\|_2^2 \leq \langle y, v \rangle < \|v\|_2^2$, direct substitution gives

$$f_v(y) + g_v(y) = \frac{1}{M} \left(1 - \sum_{k \in \mathbb{Z}} (-1)^k \frac{\rho_{\sigma_y}(y + kv)}{\rho_{\sigma_y}(y)} \right) \leq 1$$

by Eq. (7). □

Proof of Lemma 4.8. Expanding the Gaussian ratio gives

$$\frac{\rho_{\sigma_y}(y + kv)}{\rho_{\sigma_y}(y)} = \exp\left(-\frac{k^2\|v\|_2^2 + 2k\langle y, v \rangle}{2\sigma_y^2}\right),$$

so $S_v(y)$ depends only on $\langle y, v \rangle$ and $\|v\|_2^2$. Put

$$a = \frac{\|v\|_2^2}{2\sigma_y^2}.$$

Under $\langle y, v \rangle \geq -\|v\|_2^2$,

$$\frac{\rho_{\sigma_y}(y + kv)}{\rho_{\sigma_y}(y)} \leq \exp(-((k-1)^2 - 1)a).$$

Writing $k = J_{\text{tr}} + 1 + h$ and using $(J_{\text{tr}} + h)^2 - J_{\text{tr}}^2 = 2J_{\text{tr}}h + h^2 \geq 3h$ for $J_{\text{tr}} \geq 1$ and $h \in \mathbb{Z}_{\geq 0}$, the unsigned tail satisfies

$$\begin{aligned} \sum_{k=J_{\text{tr}}+1}^{\infty} \frac{\rho_{\sigma_y}(y + kv)}{\rho_{\sigma_y}(y)} &\leq \sum_{h=0}^{\infty} \exp(-((J_{\text{tr}} + h)^2 - 1)a) \\ &\leq \frac{\exp(-(J_{\text{tr}}^2 - 1)a)}{1 - \exp(-3a)} = \frac{\rho_{\sigma_y}(J_{\text{tr}}v)}{\rho_{\sigma_y}(v) - \rho_{\sigma_y}(2v)}. \end{aligned}$$

This proves Eq. (8). The total-variation statement follows by a standard κ -hybrid composition of the one-step kernels. □

A uniform truncation schedule requires control of $\sigma_y/\|v\|_2$. For example, at $\sigma_y = 55$, $\|v\|_2 = 1$, and $y = 0$, truncation before index 20 leaves absolute error at least $e^{-20^2/(2 \cdot 55^2)} - e^{-21^2/(2 \cdot 55^2)} > 0.00632$. This precludes the uniform twenty-term precision claimed in [16, Section 3.3] without a lower shift-norm bound, as imposed in Section B.5.

Proof of Lemma 4.9. Put $\alpha = \sigma_y/\|v\|_2$. Poisson summation on the one-dimensional coset $y + \mathbb{Z}v$ gives, uniformly for $y \in K_v$,

$$1 - \sum_{k \in \mathbb{Z}} (-1)^k \frac{\rho_{\sigma_y}(y + kv)}{\rho_{\sigma_y}(y)} \leq 1 + \frac{2\alpha\sqrt{2\pi}}{\exp(-1/(2\alpha^2))} \sum_{m=1}^{\infty} \exp\left(-\frac{\pi^2\alpha^2(2m-1)^2}{2}\right).$$

Since $(2m-1)^2 \geq 1 + 4(m-1)$, the series is at most

$$\frac{\exp(-\pi^2 \alpha^2 / 2)}{1 - \exp(-2\pi^2 \alpha^2)},$$

which is Eq. (4). The excess $M_\alpha - 1$ is strictly decreasing in α [16]. For $\|v\|_2 \leq B_{\text{shift}}$, we have $\sigma_y / \|v\|_2 \geq \alpha_{\text{iter}}$ and hence $M_{\alpha_{\text{iter}}} \geq M_{\sigma_y / \|v\|_2}$. Lemma 4.7 therefore applies for every common constant $M \geq M_{\alpha_{\text{iter}}}$.

When $\langle y, v \rangle \geq \|v\|_2^2$ or $\langle y, v \rangle < -\|v\|_2^2$, the proof of Lemma 4.7 gives $f_v(y) + g_v(y) = 1/M$. In the remaining strip $-\|v\|_2^2 \leq \langle y, v \rangle < \|v\|_2^2$, the same proof and the bound above give

$$f_v(y) + g_v(y) \leq \frac{M_{\sigma_y / \|v\|_2}}{M} \leq \frac{M_{\alpha_{\text{iter}}}}{M}.$$

Since $M_{\alpha_{\text{iter}}} \geq 1$, both regions have abort probability at least $1 - M_{\alpha_{\text{iter}}}/M$. $\square$

Proof of Lemma 4.10. Conditioned on success of the first step, Lemmas 4.6 and 4.9 give $z^{(1)} \sim D_{\sigma_y}^{2N}$ and success probability M_{iter}^{-1} . Inductively, conditioned on success through step t , the state $z^{(t)}$ has the same Gaussian law, and the probability that all t steps succeed is M_{iter}^{-t} . Taking $t = \kappa$ proves the final-distribution and success-probability claims. Each successful step adds either v_t or $-v_t$, so

$$z^{(\kappa)} = y + \sum_{t=1}^{\kappa} \epsilon_t X^{\ell_t} S = y + c' S.$$

Every nonzero coefficient of c' is odd and has the same support as c , hence $c' \equiv c \pmod{2}$. $\square$

Proof of Lemma 4.11. From $AS = \hat{q}$ and Eq. (9),

$$Az + c\hat{q} = u + (c' + c)\hat{q} = u + \frac{c' + c}{2} = u + d \pmod{q},$$

where the second equality uses $2\hat{q} = 1 \pmod{q}$. If $c_i = 0$, then $d_i = 0$ and the recovered coefficient is the corresponding coefficient of u . If $c_i = 1$, then $d_i \in \{0, 1\}$. For $d_i = 0$ recovery is again immediate; for $d_i = 1$, the public equality between $\tau = Az + c\hat{q}$ and $\tau - c$ identifies its high bits with those of u . This holds coefficientwise. $\square$

A.3 Monomial-Shift and Centering-Change Bounds

Proof of Lemma 4.12. Write $\mathbf{s} \in \mathbb{Z}^{2N}$ for the coefficient vector of S and set $\sigma = \sigma_1$. For the coefficient-spherical power-of-two row, set $\Sigma = I_{2N}$. For a trinomial row, set $\Sigma = \text{diag}(K_N, K_N)$. In either case, the exact zero-centered discrete Gaussian satisfies

$$\mathbb{E} e^{\langle u, \mathbf{s} \rangle} \leq \exp(\sigma^2 u^T \Sigma u / 2) \quad \text{for every } u.$$

For a trinomial row, each coefficient pair within either copy of R has probability proportional to $\exp(-x^T G x / (2\sigma^2))$. For a pair direction $r \in \mathbb{R}^2$, completing the

square expresses its moment-generating function as $\exp(\sigma^2 r^T G^{-1} r / 2)$ times the ratio of shifted to zero-centered G -Gaussian masses on $\mathbb{Z}^2$; Poisson summation bounds this ratio by 1. Multiplying these bounds over all pairs and both ring components gives the stated matrix Σ .

Every diagonal entry of Σ is at most ν_R . Taking $u = te_j$ in the linear moment bound and optimizing over $t > 0$ gives

$$\max\{\Pr[s_j > m_q], \Pr[s_j < -m_q]\} \leq \exp\left(-\frac{m_q^2}{2\nu_R\sigma^2}\right).$$

Because q is odd, $[-m_q, m_q]$ is the unique centered residue interval. Hence $\mathcal{C}_q(S) \neq S$ if and only if at least one coefficient lies outside this interval. A union bound over both tails, all $2N$ coefficients, and Q_{Hid} witnesses gives the factor $4NQ_{\text{Hid}}$ before unit conditioning. Conditioning on a unit label increases the bad-event probability by at most $1/p_{\times, \min}$, which proves the centering bound.

For $A \succeq 0$, set

$$B = \Sigma^{1/2} A \Sigma^{1/2}, \quad L = \|B\|_2, \quad V = \text{tr}(B^2).$$

Then, for every $\tau > 0$,

$$\Pr\left[\mathbf{s}^T A \mathbf{s} > \sigma^2 \left(\text{tr}(\Sigma A) + 2\sqrt{V\tau} + 2L\tau\right)\right] \leq e^{-\tau}.$$

The claim is trivial for $L = 0$, so assume $L > 0$. It is enough to prove the moment-generating-function domination

$$\mathbb{E} e^{\eta \mathbf{s}^T A \mathbf{s}} \leq \det(I - 2\eta\sigma^2 B)^{-1/2}, \quad 0 \leq \eta < 1/(2\sigma^2 L).$$

For $0 \leq \eta < 1/(2\sigma^2 L)$, use the Gaussian integral identity

$$e^{\eta x^T A x} = \mathbb{E}_{Y \leftarrow N(0, 2\eta A)} e^{\langle Y, x \rangle},$$

with Y supported on the image of A if A is singular. By Fubini and the linear MGF domination,

$$\mathbb{E}_{\mathbf{s}} e^{\eta \mathbf{s}^T A \mathbf{s}} \leq \mathbb{E}_Y e^{\sigma^2 Y^T \Sigma Y / 2} = \det(I - 2\eta\sigma^2 B)^{-1/2}.$$

It remains to optimize the Chernoff bound. Let λ_i be the eigenvalues of B and set $\theta = 2\eta\sigma^2$. For $0 \leq \theta < 1/L$,

$$\log \mathbb{E} e^{\eta(\mathbf{s}^T A \mathbf{s} - \sigma^2 \text{tr}(\Sigma A))} \leq \frac{1}{2} \sum_i (-\log(1 - \theta\lambda_i) - \theta\lambda_i) \leq \frac{\theta^2 V}{4(1 - \theta L)}.$$

The last inequality uses $-\log(1 - x) - x \leq x^2/(2(1 - x))$ for $0 \leq x < 1$. With

$$\theta = \frac{2\sqrt{\tau}}{\sqrt{V} + 2L\sqrt{\tau}},$$

Chernoff's inequality gives

$$\Pr\left[\mathbf{s}^\top A \mathbf{s} - \sigma^2 \operatorname{tr}(\Sigma A) > \sigma^2(2\sqrt{V\tau} + 2L\tau)\right] \leq e^{-\tau},$$

which is the claimed quadratic-form failure bound. The statement is for the exact sampler; approximate samplers add the corresponding statistical-distance term to the reduction loss.

For a monomial X^j , let T_j denote its coefficient-basis multiplication matrix and set $A_j = \operatorname{diag}(T_j^\top T_j, T_j^\top T_j)$. Then

$$\|(X^j s_0, X^j s_1)\|_2^2 = \mathbf{s}^\top A_j \mathbf{s}.$$

In the power-of-two ring, T_j is a signed permutation. Thus $A_j = I$ for every j , and the single norm event has

$$\operatorname{tr}(\Sigma A_j) = 2N, \quad V = 2N, \quad L = 1.$$

The quadratic-form bound with $\tau = t_2$ gives the power-of-two expression for B_{shift} .

For the trinomial ring, let G_N be the block-diagonal canonical Gram matrix and $K_N = G_N^{-1}$. Multiplication by a monomial preserves the canonical inner product, so $T_j^\top G_N T_j = G_N$ and

$$U_j = G_N^{1/2} T_j G_N^{-1/2}$$

is orthogonal. For one witness polynomial, the whitened quadratic-form matrix is

$$K_N^{1/2} T_j^\top T_j K_N^{1/2} = U_j^\top K_N U_j.$$

It therefore has the spectrum of K_N , namely $2/3$ and 2 , each with multiplicity $N/2$. Accounting for both witness polynomials gives

$$\operatorname{tr}(\Sigma A_j) = \frac{8N}{3}, \quad V = \frac{40N}{9}, \quad L = 2.$$

The quadratic-form bound with $\tau = t_3$ gives the stated trinomial threshold for one shift. Union-bounding over the N possible monomials accounts for the $\log_2 N$ term in t_3 ; no such factor is needed in the power-of-two case because all monomial shifts have the same Euclidean norm. A further union bound over Q_{Hid} witnesses gives unconditioned failure probability at most $2^{-\lambda_{\text{sh}}}$.

Conditioning a witness on $s_0 + s_1 h \in R_q^\times$ increases this probability by at most $1/p_{\times, \min}$. Lemma 4.4 and $|R_q^\times|/|R_q| = \rho_\times$ give $p_{\times, \min} \geq c_u^{-1} \rho_\times$. $\square$

A.4 ANSA-PKS^u Iterative FSwA Proofs

Proof of Lemma 4.13 and Eq. (12). Fix $c \in \mathcal{C}_\kappa$. On a monomial-shift-bounded statement, all κ shifts selected by c satisfy the hypothesis of Lemma 4.10. Conditioned on iterative success, the response is therefore exactly $D_{\sigma_y}^{2N}$ and

independent of the witness. The success probability is $M_{\text{iter}}^{-\kappa}$ for every fixed c , and the successful response has the same law for all such challenges.

For Eq. (11), Poisson summation gives the lower bound $\sqrt{2\pi}\sigma_y$ on the scalar Gaussian normalizer and shows that the periodized numerator is maximized at residue zero. The geometric bound $\sum_{k \geq 1} e^{-k^2 t} \leq e^{-t}/(1 - e^{-3t})$ for $t > 0$ then gives the stated residue-mass bound.

For the displayed lower bound on ρ_{pub} , a one-dimensional discrete-Gaussian coefficient satisfies

$$\Pr_{x \leftarrow D_{z, \sigma_y}}[|x| > Z_{\text{BOUND}}] \leq \text{erfc}\left(\frac{Z_{\text{BOUND}}}{\sqrt{2}\sigma_y}\right).$$

This follows by bounding the two tails by the corresponding Gaussian integral and using the same normalizer lower bound. Independence of the $2N$ response coefficients gives

$$\Pr[E_{\text{cap}}] \geq \left(1 - \text{erfc}\left(\frac{Z_{\text{BOUND}}}{\sqrt{2}\sigma_y}\right)\right)^{2N}.$$

For an index i with $c_i = 1$, the equality $[\tau_i]_{d_h} = [\tau_i - 1]_{d_h}$ can fail only if $\tau_i \bmod q$ lies immediately after a coefficientwise high-bit boundary. Enumerating the centered rounding intervals, including wraparound, gives at most $b_h = p_h + 1 + \mathbf{1}_{p_h \text{ odd}}$ such residues; the last term accounts for the two distinct centered tie endpoints when p_h is odd. Since A_0 is a unit, row i of multiplication by A_0 has a nonzero entry. Condition on z_1 and on all coefficients of z_0 except one coefficient with nonzero scalar in row i . Then $\tau_i \bmod q$ is a nonzero affine image of this remaining one-dimensional Gaussian residue, with maximum probability of any one residue at most p_{q, σ_y} . Therefore

$$\Pr[[\tau_i]_{d_h} \neq [\tau_i - c_i]_{d_h}] \leq b_h p_{q, \sigma_y}.$$

A union bound over the κ nonzero coefficients of c gives

$$\Pr[\neg E_{\text{eq}}(A, c)] \leq \kappa b_h p_{q, \sigma_y}.$$

Combining the response-cap and equality-check bounds yields the stated lower bound on ρ_{pub} . Applying $E_{\text{pub}}(A, c)$ after iterative rejection gives non-abort probability at least $\rho_{\text{pub}}/M_{\text{iter}}^{\kappa}$. $\square$

Proof of Lemma 4.14. Consider an honest trial of the public-coin identification protocol: the prover samples its first message, and the verifier then chooses an independent uniform challenge c . For a monomial-shift-bounded statement, the Step 4 norm checks always pass. By Lemma 4.13, for each fixed c the joint mass of iterative success and final response z is $M_{\text{iter}}^{-\kappa} D_{\sigma_y}^{2N}(z)$. This factor is independent of both c and z .

Lemma 4.11 shows that the equality check $[Az + c\hat{q}]_{d_h} = [Az + c\hat{q} - c]_{d_h}$ recovers the original commitment. Honest and simulated trials therefore apply the same public filter to the same (c, z) law and set $w = [Az + c\hat{q}]_{d_h}$. Resampling

(c, z) until acceptance gives the same joint law of (w, c, z) , proving accepting-HVZK for the identification protocol. $\square$

Proof of Lemma 4.15. First consider the raw first commitment. Condition on y_1 . Since A_0 is a unit, the map

$$y_0 \bmod q \mapsto A_0 y_0 + A_1 y_1 \bmod q$$

is a bijection on R_q . The product distribution of $y_0 \bmod q$ assigns probability at most p_{q, σ_y}^N to each residue vector. Every output of the high-bit map has at most $(2^{d_h+1} + 2)^N$ residue-vector preimages, so

$$\max_w \Pr[Ay]_{d_h} = w \leq ((2^{d_h+1} + 2)p_{q, \sigma_y})^N.$$

For fixed c , the unfiltered simulator commitment has the same form with $z \leftarrow D_{\sigma_y}^{2N}$ and the public shift $c\hat{q}$, so it obeys the same bound. Conditioning on $E_{\text{pub}}(A, c)$ increases the probability of any fixed value by at most ρ_{pub}^{-1} . Lemma 4.14 identifies this accepted simulator distribution with the accepted honest distribution, proving the claim. $\square$

Proof of Lemma 4.16. Let (μ, c, z_0, z_1) be an accepting no-message forgery and set $x = A_0 z_0 + A_1 z_1 + c\hat{q} \bmod q$ and $w = [x]_{d_h}$. By the high-bit lift, there is a residual $E \in R$ with $\|E\|_\infty \leq R_h$ such that $2^{d_h} w = x + E \pmod{q}$. With $Y_0 = z_0$ and $Y_1 = z_1$, the verification equation becomes

$$H_{\text{ch}}(\mu, [A_0 Y_0 + A_1 Y_1 + c\hat{q}]_{d_h}) = c.$$

The public response caps give the displayed Euclidean and infinity bounds. The same transcript has the lifted representation

$$H_{\text{ch}}^\uparrow(\mu, \text{can}_q(A_0 Y_0 + A_1 Y_1 + c\hat{q} + E)) = c,$$

which is used in the fork calculation below.

For the auxiliary ordinary-RSIS fork consequence, suppose two accepting transcripts have the same recovered high-bit commitment w and challenges $c \neq c'$. Let E, E' be their deterministic high-bit lift residuals. They satisfy

$$A_0 Y_0 + A_1 Y_1 + c\hat{q} + E = A_0 Y'_0 + A_1 Y'_1 + c'\hat{q} + E' \pmod{q}.$$

Multiplying by 2 and using $2\hat{q} = 1 \bmod q$ gives the ordinary-RSIS relation

$$[2A_0 \mid 2A_1 \mid 1] \begin{pmatrix} Y_0 - Y'_0 \\ Y_1 - Y'_1 \\ c - c' + 2E - 2E' \end{pmatrix} = 0 \pmod{q}.$$

The Euclidean norm is at most $2\sqrt{B_{\text{ST}}^2 + 4NR_h^2} + \sqrt{2\kappa}$. The scheme-level transcript bounds give

$$B_{\text{fork}, \infty} = \max\{2Z_{\text{BOUND}}, 4R_h + 1\} < (q - 1)/2$$

as the centered-integer condition needed before interpreting the fork output as an ordinary-RSIS instance. Since distinct unsigned challenges have different supports, some coefficient of $c - c'$ is ± 1 ; adding the even vector $2E - 2E'$ leaves an odd, hence nonzero, integer coefficient. The fork bound puts every coefficient of the third component in $(-(q-1)/2, (q-1)/2)$, so this nonzero coefficient remains nonzero modulo q . For a no-message forger of success ϵ making at most Q classical challenge queries, including final verification, the general forking argument [4] gives two accepting transcripts at the same query point with distinct challenges with probability at least $\epsilon(\epsilon/Q - 1/\binom{N}{\kappa})$. $\square$

Proof of Lemma 4.17. We instantiate the corrected ROM FSwA theorem of [2, Theorem 3] with the commitment-recoverable specialization of [2, Section 6.3]; [8, Theorem 10] is the bounded-abort analogue. In that theorem's notation, ϵ bounds commitment point probabilities, p bounds the abort probability, ζ_{zk} is the accepting-HVZK error, and δ bounds the probability of a key pair violating these conditions. Its additive loss is

$$\frac{2q_S(q_H + 1)\epsilon}{1 - p} + \frac{q_S(q_S + 1)\epsilon}{2(1 - p)^2} + q_S\zeta_{zk} + \delta.$$

We use the corresponding unbounded-loop denominator form. The first-message space is the high-bit commitment space, the challenge space is $\mathcal{C}_\kappa$, and the response is (z_0, z_1) . In the commitment-recoverable interface, the theorem's random-oracle-query count is $q_H = Q_H + Q_S$, accounting for the recovered-commitment queries that translate standard (w, z) signatures into compact (c, z) signatures. The recovery map is

$$\text{Rec}_A(c, z) := [Az + c\hat{q}]_{d_h}, \quad z = (z_0, z_1),$$

which is exactly the verifier's recovered commitment in each accepting transcript.

Thus the recovered value $\text{Rec}_A(c, z)$ is the canonical first message used by the FSwA theorem; Lemmas 4.14 and 4.15 give accepting-HVZK and min-entropy for the honest and simulated distributions of recovered first messages.

Lemmas 4.13–4.15 give $\epsilon = \epsilon_{\text{com}}$, $p = p_{\text{thm}} \leq 1 - \rho_{\text{na}}$, and $\zeta_{zk} = \epsilon_{zk} = 0$. These bounds hold uniformly for monomial-shift-bounded fixed statements, so $\delta = 0$. The factor ρ_{pub}^{-1} in Eq. (13) accounts for conditioning on the public filter. The signing oracle uses Lemma 4.14's full simulator to sample an accepted (c, z) and programs $H_{\text{ch}}(\mu, \text{Rec}_A(c, z)) = c$ unless the point was already programmed inconsistently. Replacing at most Q_S implemented signing outputs by ideal ones costs at most $Q_S\epsilon_{\text{impl}}$. With $q_S = Q_S$ and $q_H = Q_H + Q_S$, the corrected FSwA theorem bounds the remaining programming failures and commitment-guessing events by

$$\frac{2Q_S(Q_H + Q_S + 1)\epsilon_{\text{com}}}{1 - p_{\text{thm}}} + \frac{Q_S(Q_S + 1)\epsilon_{\text{com}}}{2(1 - p_{\text{thm}})^2} + Q_S\epsilon_{zk}.$$

Adding the implementation hybrid and substituting gives $L_{\text{FSwA}}^{\text{BM}}$. The no-message solver answers programmed points from its local table, as in [2, proof

of Theorem 3]. In particular, the additional recovery queries occur at simulated signing points and are answered locally. At most Q_H queries, including final verification, are forwarded to the external challenge oracle. $\square$

A.5 Signing and Identity-Based Reductions

Proof of Proposition 4.1. Lemma 4.17 reduces fixed-statement UF-CMA to the no-message game with loss $L_{\text{FSwA}}^{\text{BM}}$, and Lemma 4.16 maps any accepting no-message forgery for fixed A to the bimodal self-target relation. Under Lemma 4.16's additional ordinary-RSIS conditions, transcript differencing gives the auxiliary fork instance and challenge-collision obstruction. $\square$

Efficiency of unit-conditioned witness sampling. The witness-first samplers remain efficient for every $h \in R_q$, including the uniform-key game. Condition on all ambient coefficients except the constant coefficient of s_0 . Its conditional law is $D_{\mathbb{Z}, \sigma_1, c}$ for some center c in either ring family. Reduction modulo any irreducible factor of ϕ permits at most one residue of this coefficient if the syndrome is to vanish. Uniformly in c and the residue, Poisson summation bounds its probability by

$$m_{\sigma_1} := \frac{1 + 2e^{-q^2/(2\sigma_1^2)} / (1 - e^{-3q^2/(2\sigma_1^2)})}{\sqrt{2\pi}\sigma_1(1 - 2e^{-2\pi^2\sigma_1^2} / (1 - e^{-6\pi^2\sigma_1^2}))}.$$

The numerator bounds the Gaussian mass of a residue class, maximized when its center lies in that class; the denominator lower-bounds the scalar normalizer at every center. A union bound over at most N irreducible factors gives $\Pr[s_0 + hs_1 \in R_q^\times] \geq 1 - Nm_{\sigma_1} > 0.7896$ for Table 1. Independent ambient draws until a unit is obtained therefore sample the required conditional law in fewer than 1.267 draws in expectation, uniformly over h .

Proof of Theorem 4.1. Restrict each BBD hybrid to SG: as soon as a signing challenge is generated, check all shifts selected by its support and return 0 if any norm exceeds B_{shift} . Stops are absorbing, and the surviving subprobabilities are not renormalized. Coupled executions share challenges and stopping decisions unless an oracle-programming bad event occurs, as accounted for in Lemma 4.17.

Fix a witness. In each hybrid with independent challenges, put $b = M_{\text{iter}}^{-\kappa}$ and draw $c \leftarrow U(\mathcal{C}_\kappa)$ and $z \leftarrow D_{\sigma_y}^{2N}$ independently. Let p_{bad} be the probability that c fails the shift test, a_{good} the joint probability of passing both the shift test and the public filter, and a_{all} the probability of passing the public filter alone. Since $a_{\text{all}} \leq a_{\text{good}} + p_{\text{bad}}$, Lemma 4.13 gives

$$\Pr[\text{another attempt}] = 1 - p_{\text{bad}} - ba_{\text{good}} \leq 1 - ba_{\text{all}} \leq 1 - \rho_{\text{na}}.$$

For transcripts passing both tests, the stopped call's successful-output mass is the unfiltered (c, z) mass multiplied by $b/(p_{\text{bad}} + ba_{\text{good}})$, whereas the accepting public simulator uses $1/a_{\text{all}}$. The first factor is no larger than the second, so the public simulator dominates the stopped successful-output submeasure. The

BBD programming and geometric estimates give the loss $L_{\text{FSwA}}^{\text{BM}}$; replacing the stopped oracle by this simulator can only increase the winning probability and leaves a no-message forger depending only on A .

The statement has the unit-conditioned planted-label distribution $a = s_0 + s_1 h \bmod q$ of Definition 2.3. Assumption 4.2 replaces (h, a) by (h, u) for $u \leftarrow U(R_q^\times)$ at the displayed one-sample unit-label distinguishing cost. Then $A = (\hat{q}u^{-1}, \hat{q}hu^{-1})$ is uniform over $R_q^\times \times R_q$; inversion and scaling make A_0 uniform, while $A_1 = hA_0$ is uniform and independent of A_0 . Averaging the fixed-statement solver over A gives the random-statement BMST-RSIS term. $\square$

Proof of Lemma 4.18. Localize under the accepted annular key law, for which Lemmas 4.4 and 4.12 apply. Number identity-label states by order of creation. For a fresh forgery identity, its state is first created by exactly one of three external operations: a public query to the candidate map or identity-label wrapper, a signing query, or final verification. Nested calls inherit the type of the triggering operation. A state first created by an explicit extraction query cannot belong to a fresh forgery identity. Guess a slot in $\{1, \dots, Q_{\text{Hid}}\}$ and one of the three types uniformly and independently; abort if the guessed slot is reached with a different type.

Given the annular planted-witness challenge $A^* = (A_0^*, A_1^*)$, recover

$$h^* = A_1^*(A_0^*)^{-1}, \quad a^* = \hat{q}(A_0^*)^{-1},$$

and give h^* to the IBS adversary as its master public key. At the guessed state, use Lemma 4.5's counter simulation and place a^* at the first-unit position, even if that position has not yet been queried. Forward signing requests for this identity to the planted-witness signing oracle and attach the identity's cached counter to each returned signature. Abort on an explicit extraction request for this identity. For other identities, sample witnesses first and check all monomial-shift norms; if any exceeds B_{shift} , abort the simulation without resampling the witness. Otherwise simulate their labels, extraction, and signing locally. For a forgery on the guessed identity, return its digest and (c, z_0, z_1) to the challenger; freshness is preserved.

Every execution counted by ε passes the witness-bound checks and determines a unique slot/type pair. The simulator selects this pair with probability at least $1/(3Q_{\text{Hid}})$. With the correct guess and no returned-label collision, the simulation reproduces the joint transcript law; global monomial-shift boundedness then implies SG for the target, yielding a planted-witness signing win. Lemma 4.4 bounds each later programmed label's conditional point probability by $\Gamma_{\text{wp}}^\times$; a union bound gives collision probability at most $\varepsilon_{\text{wf-coll}}$.

The simulator routes challenge queries by digest owner. It forwards a point (μ, w) to the external challenge oracle of the planted-witness game unless $\text{own}(\mu)$ is an identity already known not to be the target. Points owned by such an identity are answered from the simulator's lazy local table, both when it signs for that identity with its sampled witness and when the adversary queries them explicitly, so the two see consistent answers. Points with an unowned digest, or whose owner has not yet been classified, are forwarded, and their answers are

retained if ownership is established later. Internal non-target signing occurs only after its identity and digest are fixed; outside the digest-collision event charged in Corollary 4.1, such a point cannot later become target-owned. Hence the planted-witness adversary makes at most Q_H challenge queries, including the final verification point. Non-target signing attempts add only to the simulator's running time.

Thus the annular-key game satisfies $\varepsilon \leq 3Q_{\text{Hid}} \Pr[\mathcal{A}_{\text{PKS}}^{\text{ann}} \text{ wins} \wedge \text{SG}] + \varepsilon_{\text{wf-coll}}$. Apply the annular-NTRU switch to this lower game. A distinguisher samples the planted witness and outputs the efficiently checkable success-and-SG bit, so

$$\Pr[\mathcal{A}_{\text{PKS}}^{\text{ann}} \text{ wins} \wedge \text{SG}] \leq \Delta_{\text{ann}} + \Pr[\mathcal{A}_{\text{PKS}} \text{ wins in ANSA-PKS}^u \wedge \text{SG}].$$

Theorem 4.1 bounds the last term by $\Delta_{\text{lab}} + \Delta_{\text{PKS}}^{\text{BM}}$; by the definition of $\Delta_{\text{low}}^{\text{BM}}$,

$$\varepsilon \leq 3Q_{\text{Hid}}\Delta_{\text{ann}} + 3Q_{\text{Hid}}(\Delta_{\text{lab}} + \Delta_{\text{PKS}}^{\text{BM}}) + \varepsilon_{\text{wf-coll}} \leq 3Q_{\text{Hid}}\Delta_{\text{low}}^{\text{BM}} + \varepsilon_{\text{wf-coll}}. \square$$

Proof of Theorem 4.2. Start from the owner-tagged digest-message game G_0 . Counter exhaustion and returned-unit collisions contribute $\varepsilon_{\text{unit}} + \varepsilon_{\text{id-coll}}$. These are absorbing bad events: the game stops when one occurs, and later hops act on the stopped subprobability without renormalization. Otherwise Lemma 4.5 simulates the visible counter state exactly and programs only returned unit labels. Game G_1 replaces deterministic extraction tapes by cached uniform tapes when a witness is first needed, paying $\text{Adv}_{\text{ext-rand}}(\mathcal{F}_{G_{\text{ext}}})$.

In G_1 , every raw Hybrid-Sampler output is passed through $\mathcal{C}_q$. Game G_2 replaces each resulting $D_{\text{hyb}}^{\text{ctr}}(\cdot \mid B, h, a)$ witness by $D_{\text{pub}}^{\text{ctr}}(\cdot \mid h, a)$ at the same unit label. Pure label queries generate no witness and incur no cost. Lemmas 4.1–4.3 give the accumulated factor $r_p^{Q_{\text{ID}}}$ and the outer Rényi exponent.

Game G_3 replaces uniform unit labels with centered witness-first programming. It samples the ambient S conditioned on $a = s_0 + s_1 h \bmod q$ being a unit, programs a , and caches $\mathcal{C}_q(S)$ as the signing witness. By Lemma 4.4, adjoining the common centered conditional-witness kernel preserves the label Rényi factor. Lemmas 4.1 and 4.2 then give the inner transfer over Q_{Hid} labels with factor $(r_u^\times)^{Q_{\text{Hid}}}$. The collision probability $\varepsilon_{\text{wf-coll}}$, evaluated after this transition using $\Gamma_{\text{wp}}^\times$, therefore lies inside the inner bracket.

Game G_4 uses the same ambient draws and labels as G_3 but caches S rather than $\mathcal{C}_q(S)$. The transcripts are identical unless centering changes at least one witness-first draw, an event bounded by $\varepsilon_{\text{wrap}}^{\text{wf}}$ in Lemma 4.12. Thus G_4 is the witness-first game with the unmodified ambient witness analyzed in Lemma 4.18. Together with the preceding collision and centering terms, that lemma and the monomial-shift failure bound yield Θ_{wf} . Applying event transfer at the two Rényi transitions proves the displayed inequality; the initial terms $\varepsilon_{\text{unit}} + \varepsilon_{\text{id-coll}}$ remain outside both transitions. $\square$

Proof of Corollary 4.1. The wrapper lazily samples and caches $\mu \leftarrow \{0, 1\}^{2\lambda}$ for each new public digest input $x = (\text{pkh}, \text{id}, j, \text{msg})$ and records $\text{own}(\mu) = x$. It forwards signing queries as (id, μ) and maps the final forgery likewise. The real wrapper and the owner-tagged game coincide unless two distinct inputs receive the same digest; outside that event, external freshness

implies digest-message freshness. A union bound over at most Q_μ inputs gives $\binom{Q_\mu}{2}2^{-2\lambda}$, which is added before Theorem 4.2. $\square$

A.6 Proof Conditions and Reduction Loss

Table A.1 evaluates the proof conditions and Theorem 4.2 under the budgets of Section 5.1. We take $\varepsilon_p = (e^{2\pi^2 R_{\text{ext}}^2}/(2N) - 1)^{-1}$ and $\varepsilon_u = 2\varepsilon_p/(1 - \varepsilon_p)$. All profiles satisfy $\varepsilon_p < 2^{-35} < 1/12$, $\delta'_p < 2^{-31.6} < 1$, and $\delta_u^\times = c_u^2 - 1 < 2^{-32.2} < 1$, as Lemmas 4.3 and 4.4 require. Moreover, $\rho_\times \geq (1 - q^{-1})^N > 0.9997$, hence $p_{\times, \min} = c_u^{-1}\rho_\times > 0.9997$.

Theorem 4.2 constructs a solver $\mathcal{C}$ that uses the IBS forger to solve the BMST-RSIS problem of Assumption 4.3. To isolate the BMST contribution, set the other computational advantages to zero and retain all statistical terms, including the extraction-root collision bound $\varepsilon_{\text{root}} = \binom{Q_{\text{ID}}}{2}2^{-512}$. For $\varepsilon_{\text{BMST}} = \text{Adv}^{\text{BMST-RSIS}}(\mathcal{C})$, write $m_{\text{loc}} = 3Q_{\text{Hid}}$ and

$$\begin{aligned} h_p &= (a_p - 1)/a_p, & h_u &= (a_u - 1)/a_u, & \vartheta_{\text{id}} &= h_p h_u, \\ \ell_p &= h_p Q_{\text{ID}} \log_2 r_p, & \ell_u &= \vartheta_{\text{id}} Q_{\text{Hid}} \log_2 r_u^\times, & \ell &= \ell_p + \ell_u. \end{aligned}$$

With

$$\begin{aligned} \varepsilon_{\text{out}} &= \varepsilon_{\text{unit}} + \varepsilon_{\text{id-coll}} + \varepsilon_{\text{root}} + \varepsilon_{\mu\text{-coll}}, \\ \varepsilon_{\text{in}} &= m_{\text{loc}} I_{\text{FSwA}}^{\text{BM}} + \varepsilon_{\text{wf-coll}} + \varepsilon_{\text{sh}}^{\text{wf}} + \varepsilon_{\text{wrap}}^{\text{wf}}, \end{aligned}$$

Theorem 4.2 and Corollary 4.1 bound the external-message forgery probability by

$$\varepsilon \leq \mathcal{E}_P(\varepsilon_{\text{BMST}}) := \min\left\{1, \varepsilon_{\text{out}} + 2^\ell (m_{\text{loc}} \varepsilon_{\text{BMST}} + \varepsilon_{\text{in}})^{\vartheta_{\text{id}}}\right\}. \quad (\text{A.1})$$

For expected solver work T and the classical BMST cost exponent $c_{\text{ref}, P}$ in Table 2, we use the conditional model

$$\varepsilon_{\text{BMST}} \leq \varepsilon_{\text{ref}, P}(T, Q_H) := \min\left\{1, T2^{-c_{\text{ref}, P}} + \frac{Q_H}{|\mathcal{C}_\kappa|}\right\}. \quad (\text{A.2})$$

This model assumes linear work–success scaling for lattice attacks and includes direct challenge search. Its use as an upper bound on arbitrary BMST solvers is an additional hardness hypothesis, not a consequence of the Core-SVP estimates.

Table A.1 uses $T = Q_H = 2^{64}$ and

$$\begin{aligned} \varepsilon_{\text{ref}, P} &= \varepsilon_{\text{ref}, P}(2^{64}, 2^{64}), & b_{\text{in}, P} &= -\log_2 \varepsilon_{\text{ref}, P}, \\ \Theta_P^* &= m_{\text{loc}} \varepsilon_{\text{ref}, P} + \varepsilon_{\text{in}}, & b_{\text{ref}, P} &= -\log_2 \mathcal{E}_P(\varepsilon_{\text{ref}, P}). \end{aligned}$$

Here $b_{\text{ref}, P}$ is a probability exponent; $64 + b_{\text{ref}, P}$ is a time–success exponent when adversary and solver work are both normalized to T . A ratio r_T of solver work to adversary work adds $\log_2 r_T$ to the reduction’s time–success loss.

The reference BMST time–success exponent is $64 + b_{\text{ref}, P}$, so the reduction loss is $b_{\text{in}, P} - b_{\text{ref}, P}$. With $L_{\text{loc}} = \log_2 m_{\text{loc}}$, it decomposes, up to terms below the displayed precision, as

$$b_{\text{in}, P} - b_{\text{ref}, P} = L_{\text{loc}} + (1 - \vartheta_{\text{id}})(b_{\text{in}, P} - L_{\text{loc}}) + \ell.$$

Identity localization contributes L_{loc} ; the fixed-statement reduction of Proposition 4.1 has no multiplicative query loss. Tables A.1–A.2 optimize over $a_p \in \{2, 4, 8, \dots, 512\}$ and $a_u \in \{1.100, 1.105, \dots, 200.000\}$. This loose reduction assessment is separate from the direct-attack estimates in Table 2.

The first row of Table A.1 is the lifted bound $Z_{\text{BOUND}} + R_h$ of the BMST-RSIS screen in Section 5.6.

Table A.1. Proof conditions and conditional reduction estimates at $T = Q_H = 2^{64}$ under Eq. (A.2). Values are rounded to the displayed precision.

metric	IBS-1s	IBS-1	IBS-2	IBS-3
BMST lifted ℓ_∞ bound	1194097	2022847	2022847	2022847
$B_{\text{fork}, \infty}$	2380000	4037500	4037500	4037500
$\lfloor (q-1)/2 \rfloor - B_{\text{fork}, \infty}$	1810208	152708	152708	152708
pq, σ_y upper bound	$1.425 \cdot 10^{-6}$	$8.399 \cdot 10^{-7}$	$8.399 \cdot 10^{-7}$	$7.860 \cdot 10^{-7}$
lower bound on ρ_{pub}	0.9279	0.913	0.889	0.692
$-\log_2 \varepsilon_{\text{com}}$	5550.6	9497.2	12662.9	14466.1
$-\log_2 L_{\text{FSWA}}^{\text{BM}}$	5419.8	9366.8	12532.5	14334.4
$-\log_2 \varepsilon_{\text{wrap}}^{\text{wf}}$	> 663000	> 497000	> 663000	> 497000
$-\log_2 \varepsilon_p / -\log_2 \varepsilon_u$	36.390/35.390	35.805/34.805	35.390/34.390	35.220/34.220
identity-label budget $\log_2 Q_{\text{Hid}}$	24	48	48	48
witness budget $\log_2 Q_{\text{ID}}$	24	48	48	48
classical reference cost exponent $c_{\text{ref}, P}$	100	159	230	267
BMST advantage input $b_{\text{in}, P} = -\log_2 \varepsilon_{\text{ref}, P}$	36.000	95.000	163.684	203.000
Rényi orders a_p/a_u	512/200.000	512/200.000	512/200.000	512/200.000
localization loss L_{loc}	25.585	49.585	49.585	49.585
event exponent ϑ_{id}	0.993057	0.993057	0.993057	0.993057
exponent-contraction loss	0.072	0.315	0.792	1.065
weighted Rényi losses ℓ_p/ℓ_u	$< 10^{-6} / < 10^{-6}$	0.0041/0.0007	0.0074/0.0013	0.0093/0.0016
reference inner exponent $-\log_2 \Theta_P^*$	10.4	45.4	114.1	153.4
advantage-exponent loss $b_{\text{in}, P} - b_{\text{ref}, P}$	25.657	49.905	50.386	50.661
conditional probability exponent $b_{\text{ref}, P}$	10.343	45.095	113.298	152.339
conditional time/success exponent $64 + b_{\text{ref}, P}$	74.343	109.095	177.298	216.339

Table A.2 gives advantage losses without a work–success hypothesis. Inverting Eq. (A.1), with $[t]_+ = \max\{t, 0\}$, gives

$$\varepsilon_{\text{BMST}} \geq \varepsilon_{\text{req}, P}(\varepsilon) := \frac{1}{m_{\text{loc}}} \left[\left(\frac{[\varepsilon - \varepsilon_{\text{out}}]_+}{2^\ell} \right)^{1/\vartheta_{\text{id}}} - \varepsilon_{\text{in}} \right]_+.$$

At forgery probability 2^{-x} , the ideal loss is $-\log_2 \varepsilon_{\text{req}, P}(2^{-x}) - x$ bits. The Q128 rows invert the composition with Eq. (B.26), retaining all finite-extraction and signing errors.

Table A.2. Upper bounds on the BMST reduction’s advantage loss under the stated query budgets, with other computational advantages set to zero. All statistical terms are included, and values are rounded upward in bits.

metric	IBS-1s	IBS-1	IBS-2	IBS-3
ideal loss at 1/2	25.592	49.597	49.599	49.601
Q128 loss at 1/2	25.663	49.985	49.877	50.939
Q128 loss at 2^{-64}	26.832	51.050	50.859	51.878

A.7 Relative-Norm Moment and Projection Calculations

In Antrag's radial sampler, R^2 is uniform on $[\rho_-^2, \rho_+^2]$ and a component energy is $R^2 \cos^2 \varphi$ with an independent uniform phase. Hence

$$\mathbb{E}[R^2 \cos^2 \varphi] = \frac{\rho_-^2 + \rho_+^2}{2} \cdot \frac{1}{2} = \frac{\rho_-^2 + \rho_+^2}{4}.$$

Summing the two component means gives $(\rho_-^2 + \rho_+^2)/2$. The distinction between a component mean and the total energy accounts for the factor-of-two difference from [14, Lemma 5] in the normalization used in Section 5.2.

We now derive the moments $q^2 y$ and $q^2 v$ of Section 5.2. Fix a primitive embedding θ and write the pre-rounding values as $\phi_\theta(f_0) = R \cos \varphi e^{i\psi_1}$ and $\phi_\theta(g_0) = R \sin \varphi e^{i\psi_2}$, with R^2 uniform on $[\rho_-^2, \rho_+^2]$ and independent uniform phases φ, ψ_1, ψ_2 . Under Heuristic 1, rounding adds errors e_f, e_g with independent coefficients of variance $1/12$, treated as Gaussian for fourth moments. For $z = \phi_\theta(e_f)$,

$$\mathbb{E}|z|^2 = \frac{N}{12}, \quad \mathbb{E}z^2 = \frac{1}{12} \sum_{j=0}^{N-1} e^{2ij\theta}, \quad |\mathbb{E}z^2| = \frac{1}{12} \left| \frac{\sin(N\theta)}{\sin \theta} \right|,$$

and likewise for $\phi_\theta(e_g)$. Expanding $|\phi_\theta(f)|^2 = R^2 \cos^2 \varphi + 2 \operatorname{Re}(\phi_\theta(f_0)\bar{z}) + |z|^2$, and similarly for g , gives three uncorrelated parts; every cross covariance is odd in e_f , in e_g , or in $\cos 2\varphi$. For a complex Gaussian z , $\operatorname{Var}|z|^2 = (\mathbb{E}|z|^2)^2 + |\mathbb{E}z^2|^2$, so the two quadratic error terms contribute

$$\frac{N^2 + |\sin(N\theta)/\sin \theta|^2}{72}$$

to the variance of their sum and of their difference. Averaging over ψ_1, ψ_2 , the linear terms have total variance $2(R^2 \cos^2 \varphi + R^2 \sin^2 \varphi) \mathbb{E}|z|^2 = NR^2/6$. The means give $\mathbb{E}W_\theta = \mathbb{E}R^2 + N/6 = \bar{W}$. Hence

$$\begin{aligned} \mathbb{E}|E_\theta|^2 &= \operatorname{Var}(R^2) + \frac{N \mathbb{E}R^2}{6} + \frac{N^2 + |\sin(N\theta)/\sin \theta|^2}{72}, \\ \mathbb{E}|D_\theta|^2 &= \mathbb{E}[R^4 \cos^2 2\varphi] + \frac{N \mathbb{E}R^2}{6} + \frac{N^2 + |\sin(N\theta)/\sin \theta|^2}{72}. \end{aligned}$$

Substituting $\operatorname{Var}(R^2) = (\rho_+^2 - \rho_-^2)^2/12$, $\mathbb{E}R^2 = (\rho_+^2 + \rho_-^2)/2$, $\mathbb{E}R^4 = (\rho_+^4 + \rho_+^2 \rho_-^2 + \rho_-^4)/3$, and $\mathbb{E} \cos^2 2\varphi = 1/2$, and averaging over the primitive embeddings, gives the expressions for $q^2 y$ and $q^2 v$ in Section 5.2, with $J = N(\rho_+^2 + \rho_-^2)/12 + (N^2 + \bar{S}_N)/72$. The same parity argument gives $\mathbb{E}[E_\theta D_\theta] = 0$.

For $M = 3N$ with $4 \mid M$, every primitive embedding has $\sin^2(N\theta) = 3/4$. Put $J = M/2 = 3N/2$. Reduction of primitive residues modulo M to those modulo J is two-to-one, and

$$\sum_{\substack{1 \leq k < J \\ (k, J)=1}} \csc^2(\pi k/J) = \frac{J_2(J)}{3}, \quad J_2(J) = J^2 \prod_{p \mid J} (1 - p^{-2}) = \frac{2J^2}{3}.$$

The corresponding sum over the N primitive embeddings is N^2 , giving $\bar{S}_N = 3N/4$. The condition $4 \mid M$ is satisfied by $M = 4608, 6912$.

The centered offset in the relative-norm lattice is $((E + D)/2, E)$. Exchange symmetry makes $\mathbb{E}[ED] = 0$, so its normalized second-moment matrix is

$$C = q^2 \begin{pmatrix} (y+v)/4 & y/2 \\ y/2 & y \end{pmatrix}, \quad \det C = q^4 yv/4.$$

For any determinant-one linear twist U , $\text{tr}(UCU^\top) \geq 2\sqrt{\det C}$, with equality at an isotropizing twist. This gives the offset $q^2\sqrt{yv}$ used in Eq. (18).

Removing an identity-coordinate direction from a lattice basis forces the corresponding coordinate of its represented vectors to zero. Since (A, W) has dense $A = ff^*$, the coordinate deletion of [14, Eq. (12)] need not preserve the target. A valid specialization with no coordinate deletion must also recover the determinant and offset above. Section 5.2 therefore evaluates the more favorable relaxation that permits deletion of error coordinates; the absence of a BDD crossing is a conclusion within that model.

B Fixed-Point Realization and Concrete Security

Sections B.1–B.3 specify the Setup generator, including the rejections introduced by finite arithmetic and the NTRU solver, and establish its successful-output law. Omitting only the Galois screen gives the master-key distribution assumed in Section 4; Section B.3 relates the screened and unscreened laws. Sections B.4 and B.5 specify finite extraction and signing and compare their output laws with the exact algorithms; Section B.6 treats arithmetic hardening and the extraction count transcript. Section B.7 combines these comparisons with the ideal security reduction. The comparisons exclude operational faults. The Setup and extraction analysis also applies to ANSA-IBKEM [5] with the same generator and sampler policies.

B.1 Candidate Distribution and NTRU-Solver Policy

Fix a parameter profile P and write $N = N_P$, $H = N/2$, and $M = M_P$. Define its identifier by

$$\begin{aligned} \text{PId}(\text{IBS-1s}) &= 0\text{x}0001, & \text{PId}(\text{IBS-1}) &= 0\text{x}0002, \\ \text{PId}(\text{IBS-2}) &= 0\text{x}0003, & \text{PId}(\text{IBS-3}) &= 0\text{x}0004. \end{aligned}$$

For $b \geq 1$ and $0 \leq j < 2^b$, let $\text{rev}_b(j)$ denote the integer obtained by reversing the b -bit representation of j , including leading zeros. Define

$$\begin{aligned} Q_{\text{low}} &= \frac{347696870082016657}{54756000000}, \\ Q_{\text{high}} &= \frac{611344282389613177}{54756000000}, \\ v &= e^{2\pi i/2^{66}}. \end{aligned}$$

Algorithm B.1 defines the exact annular seed-to-pair map. Algorithm B.4 evaluates this map in finite arithmetic; rejection when a finite comparison cannot be resolved is called a cutoff and forms part of the specified key law.

Algorithm B.1 ($\text{SampleFG}_P^{\text{exact}}$)

Input: a parameter profile P and $\text{seed} \in \{0, 1\}^{512}$.

Output: (f, g) or REJECT.

1. Set $N = N_P$ and $H = N/2$. Initialize the XOF

$$\mathcal{S}_P = \text{SHAKE256}(\text{'ANSA-KG'} \parallel \text{u16le}(\text{Pid}(P)) \parallel \text{seed}), \quad (\text{B.1})$$

and read exactly $2N$ words $W_0, \dots, W_{2N-1}$ by successive calls to ReadLE64, each of which consumes eight bytes in little-endian order and returns an integer in $[0, 2^{64})$. Discard the unused XOF suffix.

2. For $0 \leq j < H$, compute

$$\begin{aligned} r_j &= Q_{\text{low}} + (Q_{\text{high}} - Q_{\text{low}})W_j/2^{64}, & \rho_j &= \sqrt{r_j}, \\ z_{\varphi,j} &= v^{W_{j+H}}, & z_{f,j} &= v^{4W_{j+N}}, \\ z_{g,j} &= v^{4W_{j+3H}}, \\ \hat{f}_j &= \rho_j \Re(z_{\varphi,j} z_{f,j}), & \hat{g}_j &= \rho_j \Im(z_{\varphi,j} z_{g,j}). \end{aligned}$$

The products $4W_{j+N}$ and $4W_{j+3H}$ are integer exponents; no phase bit is discarded.

3. Define the packed spectral permutation

$$\pi_P(j) = \begin{cases} \text{rev}_{\log_2 H}(j), & P \in \{\text{IBS-1s}, \text{IBS-2}\}, \\ j, & P \in \{\text{IBS-1}, \text{IBS-3}\}. \end{cases}$$

For $x \in \{f, g\}$ and $0 \leq i < H$, set

$$u_{x,i} = \frac{\zeta_M^{-i}}{H} \sum_{j=0}^{H-1} \hat{x}_j \zeta_H^{-i\pi_P(j)}, \quad \zeta_m = e^{2\pi i/m}.$$

Thus $\hat{x}_j$ occupies the primitive root $\zeta_M^{1+(M/H)\pi_P(j)}$, while its complex conjugate occupies the conjugate root.

4. Recover a real coefficient vector $y_x \in \mathbb{R}^N$. For $M = 4H$, set

$$y_{x,i} = \Re u_{x,i}, \quad y_{x,i+H} = \Im u_{x,i}.$$

For $M = 6H$, set

$$y_{x,i+H} = \frac{2\Im u_{x,i}}{\sqrt{3}}, \quad y_{x,i} = \Re u_{x,i} - \frac{y_{x,i+H}}{2}.$$

5. Apply DecodeOdd independently to y_f and y_g . On input y , this map first sets $z_i = \text{RoundEven}(y_i)$, where RoundEven is nearest-integer rounding with half-integer ties to even, and puts $d_i = y_i - z_i$. Let i^* be the least index maximizing $|d_i|$. If $\sum_i z_i$ is even, replace z_{i^*} by $z_{i^*} + 1$ when $d_{i^*} > 0$, and by $z_{i^*} - 1$ otherwise, including when $d_{i^*} = 0$.

6. Interpret the resulting vectors as (f, g) . Return REJECT if any coefficient lies outside $[-4095, 4095]$.
7. Require the annular predicate $\text{GoodPair}_P(f, g)$: at every canonical embedding ϕ_i ,

$$\frac{q}{\alpha_{\text{ann}}^2} \leq |\phi_i(f)|^2 + |\phi_i(g)|^2 \leq q\alpha_{\text{ann}}^2.$$

Return REJECT if any exact comparison fails; otherwise return (f, g) .

For Galois screening, define the following sets of automorphism exponents. They contain every nonidentity element of order 2 and, for the trinomial rings, every element of order 3.

P	$A_{2,P}$	$A_{3,P}$
IBS-1s	$\{1023, 1025, 2047\}$	$\emptyset$
IBS-1	$\{1025, 1279, 2303, 2305, 3329, 3583, 4607\}$	$\{1537, 3073\}$
IBS-2	$\{2047, 2049, 4095\}$	$\emptyset$
IBS-3	$\{1025, 2431, 3455, 3457, 4481, 5887, 6911\}$	$\{2305, 4609\}$.

Each profile fixes a deterministic integer NTRU solver Solve_P , including the tower, precision, public reduction schedules, capacity bounds, and rejection rules of Section B.2. It forms relative norms, solves a Bézout equation at the bottom of the tower, and lifts and reduces (F, G) with $fG - gF = q$. Its rejection rules form part of the key law in Assumption 4.1.

Every basis returned by the NTRU solver must satisfy the exact second-row norm condition

$$\|(F, G)\|_{\text{can}}^2 \leq q\alpha_{\text{ann}}^2 k_N, \quad k_N = \begin{cases} N/4 + 1, & \Phi_{M_P} = X^N + 1, \\ N/6 + 1, & \Phi_{M_P} = X^N - X^{N/2} + 1. \end{cases}$$

Here $k_N = R_{\text{cov}}^2 + 1$: the squared covering radius of R is $N/4$ for the orthonormal power basis and $N/6$ for the $N/2$ hexagonal coefficient pairs. Ideal nearest-point reduction of the second-row projection bounds its contribution by $q\alpha_{\text{ann}}^2 R_{\text{cov}}^2$, while annularity bounds the orthogonal contribution by $q\alpha_{\text{ann}}^2$.

The implemented solver need not perform nearest-point reduction; its output is checked against this bound in linear time using the metric of Section 2.1. Let $\text{Valid}_P(B)$ denote the unit, exact NTRU, encoding, second-row norm, and Galois conditions of Algorithm B.2, together with determined acceptance by the Q96 annular predicate of Section B.2. Proposition B.2 proves that they guarantee extraction preparation, so KeyGen need not construct that cache. A prescribed rejection advances to the next candidate. Arithmetic-guard violations and operational failures abort the call; they are not retries.

Algorithm B.2 ($\text{FinishKeyGenCandidate}_P$)

Input: a parameter profile P and an accepted pair (f, g) returned by Algorithm B.1 or B.4.

Output: (msk, mpk) or REJECT.

1. Require $f \in R_q^\times$ and set $h = gf^{-1} \bmod q$. Return REJECT if f is not a unit.

2. For every $a \in A_{2,P} \cup A_{3,P}$, require $h(X^a) \not\equiv h(X) \pmod{(q, \Phi_{M_P})}$. Return REJECT if any equality holds.
3. Solve $fG - gF = q$ with Solve_P . Return REJECT on a prescribed solver rejection, including failure of the bottom Bézout condition, a nonpositive reduction denominator, an exceeded admission bound, schedule exhaustion, a failed word-fit or single-prime test at an upper power-of-two level, a failed final two-prime equation test, or an unrepresentable signed-32-bit basis coefficient. In particular, reject $F_i = -2^{31}$ before forming $-F_i$. Every returned pair satisfies $fG - gF = q$ exactly: the upper-level policy establishes it by its final two-prime test, and the scheduled reducer verifies it at the degree where it is entered (32 for IBS-1s, 64 for IBS-2, and N for the trinomial profiles). An equation mismatch violates an internal invariant and is a failure that ends Setup, not a rejection. Operational failures are outside the mathematical comparison. Set

$$B = \begin{pmatrix} g & -f \\ G & -F \end{pmatrix}.$$

4. Return REJECT unless the second-row norm condition holds.
5. Set $\text{msk} = \text{EncodeMSK}(P, B)$ and $\text{mpk} = \text{EncodeMPK}(P, h)$, and return (msk, mpk) .

Here $\text{EncodeMSK}(P, B)$ writes a four-byte header (key type, format version, and profile identifier), the coefficients of f and g as 13-bit two's-complement fields in $[-4095, 4095]$, and the coefficients of F as signed 32-bit little-endian integers; the decoder reconstructs G as in Section B.4, rejects a reconstruction that is not representable as signed 32-bit coefficients, and checks the integer NTRU equation. $\text{EncodeMPK}(P, h)$ writes the header and the coefficients of h in $[0, q)$ as 23-bit fields. The master-key encoder rejects f, g coefficients outside $[-4095, 4095]$, the public-key encoder reduces coefficients modulo q , and both decoders reject noncanonical encodings.

Algorithm B.3 gives the exact-pair comparison generator. Algorithm B.5 uses the same outer loop and NTRU-solver policy with the finite candidate evaluator. The latter generator, with only Step 2 of Algorithm B.2 omitted, defines $\mathcal{H}_{\text{ann}}$; Section 5.3 accounts for the Galois screen.

Algorithm B.3 ($\text{Setup}_P^{\text{exact}}$)

Input: a parameter profile P .

Output: $(\text{pp}, \text{mpk}, \text{msk})$ or $\perp$.

1. For $t = 0, \dots, 2^{16} - 1$:
 1. Draw an independent uniform seed $\text{seed}_t \in \{0, 1\}^{512}$.
 2. Run Algorithm B.1 on seed_t and, if it returns (f, g) , run Algorithm B.2.
 3. On success, let pp identify P and return $(\text{pp}, \text{mpk}, \text{msk})$.
2. Return $\perp$ if all 2^{16} candidates are rejected.

B.2 Fixed-Point Setup

Radial interpolation, integer square roots, and inverse transforms use Q128 arithmetic; integer CORDIC supplies the Q128 trigonometric values. Certified enclosures determine coefficient rounding, while directed Q96 transforms and squared-norm evaluation certify the annular test; unresolved comparisons reject the candidate. Fourier NTRU reduction uses Q32 at the upper power-of-two levels and degree-dependent precision in the scheduled reducer, with base precisions Q256 for IBS-1s/1/3 and Q384 for IBS-2. The power-of-two degree-two step uses an exact rational quotient.

Fixed-point arithmetic. An integer A in Qs represents $A2^{-s}$. Iw and Uw denote signed and unsigned w -bit integer carriers, respectively. Write RNE for nearest-integer rounding with ties to even and RNA for nearest-integer rounding with ties away from zero; both use integer quotient and remainder. Define

$$\begin{aligned} \text{Mul}_s(A, B) &= \text{RNE}(AB/2^s), \\ \text{Div}_s(A, d) &= \text{RNE}(A/d), \\ \text{SqrtN}_s(A) &= \begin{cases} t, & A2^s - t^2 < (t+1)^2 - A2^s, \\ t+1, & A2^s - t^2 > (t+1)^2 - A2^s, \\ \text{the even member of } \{t, t+1\}, & \text{otherwise} \end{cases} \end{aligned}$$

In the definition of SqrtN_s , $A \geq 0$ and $t = \lfloor \sqrt{A2^s} \rfloor$ is obtained by exact integer square root with remainder. The squared-residual rule selects an adjacent integer whose distance from $\sqrt{A2^s}$ is less than one integer unit; bisection with widened squaring gives the same result. Q128 values use signed-I256 states, and signed products are formed from exact U512 magnitudes before a single RNE shift. For complex Qs pairs, put

$$\text{CMul}_s((a, b), (c, d)) = (\text{Mul}_s(a, c) - \text{Mul}_s(b, d), \text{Mul}_s(a, d) + \text{Mul}_s(b, c)).$$

Additions and subtractions are exact within the carrier bounds proved in Section B.3. Applied coordinatewise to a complex number, $\langle x \rangle_s$ denotes the precomputed integer $\text{RNE}(2^s x)$.

The constants $\langle \zeta_3 \rangle_s$, $\langle \zeta_9 \rangle_s$, $\langle \zeta_{27} \rangle_s$, $\langle \sqrt{3}/2 \rangle_s$, and $\langle 2/\sqrt{3} \rangle_s$ are generated by exact interval arithmetic. Write $\zeta_3 = (-1/2, s_3)$, where $4s_3^2 = 3$ and $s_3 > 0$. The real part $c_9 = \Re \zeta_9$ is the unique root of $8c^3 - 6c + 1$ in $(3/4, 4/5)$, and its imaginary part is the positive root of $s_9^2 = 1 - c_9^2$. Finally, $t_{27} = 2\Re \zeta_{27}$ is the unique root in $(19/10, 2)$ of $C_9(t) + 1$, where $C_0 = 2$, $C_1 = t$, and $C_{k+1} = tC_k - C_{k-1}$; the imaginary part of ζ_{27} is the positive root of $1 - (t_{27}/2)^2$. Rational bisection and integer square root certify the correctly rounded coordinate of each stored constant.

Integer phases and transforms. A U132 phase is an integer $p \in [0, 2^{132})$ representing $p/2^{132}$ turns. Integer CORDIC uses the precomputed tables

$$G_{128} = \text{RNE} \left(2^{128} \prod_{i=0}^{130} (1 + 2^{-2i})^{-1/2} \right), \quad A_i = \text{RNE} \left(\frac{2^{132} \arctan(2^{-i})}{2\pi} \right).$$

To compute $\text{SC}_{128}(p) = (C, S)$, write $p = \nu 2^{130} + r$, where $0 \leq \nu < 4$ and $0 \leq r < 2^{130}$. If $r = 0$, set $(x, y) = (2^{128}, 0)$. Otherwise initialize $(x_0, y_0, e_0) = (G_{128}, 0, r)$ and, for $0 \leq i \leq 130$, compute

$$\epsilon_i = \begin{cases} +1, & e_i \geq 0, \\ -1, & e_i < 0, \end{cases} \quad \begin{aligned} x_{i+1} &= x_i - \epsilon_i \text{ASR}_i(y_i), \\ y_{i+1} &= y_i + \epsilon_i \text{ASR}_i(x_i), \\ e_{i+1} &= e_i - \epsilon_i A_i, \end{aligned}$$

where $\text{ASR}_i(a) = \lfloor a/2^i \rfloor$ is a signed arithmetic shift. In the nonzero- r branch, set $(x, y) = (x_{131}, y_{131})$. In both branches, apply the quadrant map

$$(C, S) = \begin{cases} (x, y), & \nu = 0, \\ (-y, x), & \nu = 1, \\ (-x, -y), & \nu = 2, \\ (y, -x), & \nu = 3. \end{cases}$$

The 131 rotations leave residual $e_{131} \in \{-1, 0\}$ for every nonzero r ; each returned Q128 coordinate has certified absolute error below 2^{-116} . The branch $r = 0$ is exact.

Every transform root is a Q128 pair, conjugation negates its second coordinate, and all complex products use CMul_{128} . For a power-of-two profile, define, for $0 \leq k < N$,

$$\mathcal{G}[k] = \text{SC}_{128} \left(\frac{2^{132}}{M} \text{rev}_{\log_2 N}(k) \right),$$

with the exact assignments $\mathcal{G}[0] = (2^{128}, 0)$ and $\mathcal{G}[1] = (0, 2^{128})$. Given spectral inputs $X_0, \dots, X_{H-1}$, initialize $V_j = X_j$ and $w = 1$. For $\ell = \log_2 N - 1, \dots, 1$, set $m = 2^\ell$ and $t = 2w$. For $0 \leq i < m/2$ and $i t \leq j < i t + w$, replace

$$(V_j, V_{j+w}) \leftarrow (V_j + V_{j+w}, \text{CMul}_{128}(V_j - V_{j+w}, \overline{\mathcal{G}[m+i]})).$$

After each stage set $w \leftarrow 2w$. The normalized output is obtained by applying $\text{Div}_{128}(\cdot, H)$ to both coordinates of every V_j . The resulting half-spectrum order is the permutation π_P of Algorithm B.1.

The mixed-radix profiles implement a natural-order length- H inverse DFT. Set $\mathcal{R}_P[0] = \mathcal{T}_P[0] = (2^{128}, 0)$ and generate the root and twist arrays by repeated CMul_{128} from

P	$\mathcal{R}_P[1]$	$\mathcal{T}_P[1]$
IBS-1	$\text{CMul}_{128}(\text{SC}_{128}(171 \cdot 2^{124}), \langle \zeta_3 \rangle_{128})$	$\text{CMul}_{128}(\text{SC}_{128}(57 \cdot 2^{123}), \langle \zeta_9 \rangle_{128})$
IBS-3	$\text{CMul}_{128}(\text{SC}_{128}(57 \cdot 2^{125}), \langle \zeta_9 \rangle_{128}^{\otimes 5})$	$\text{CMul}_{128}(\text{SC}_{128}(19 \cdot 2^{124}), \overline{\langle \zeta_{27} \rangle_{128}}^{\otimes 2})$

Here $Z^{\otimes k}$ is the left-associated k th power under CMul_{128} , not exact complex exponentiation. For $0 \leq i < H - 1$, set $\mathcal{R}_P[i+1] = \text{CMul}_{128}(\mathcal{R}_P[i], \mathcal{R}_P[1])$ and define $\mathcal{T}_P$ analogously.

The inverse transform is recursive. At a power-of-two node of length n , first replace $(V_j)_{j < n}$ by $(V_{\text{rev}_{\log_2 n}(j)})_{j < n}$. Then, for $L = 2, 4, \dots, n$, each block beginning at $b \in \{0, L, 2L, \dots, n - L\}$ and each $0 \leq i < L/2$ performs

$$\begin{aligned} X &= V_{b+i}, \\ Y &= \text{CMul}_{128}(V_{b+i+L/2}, \overline{\mathcal{R}_P[iH/L]}), \\ (V_{b+i}, V_{b+i+L/2}) &\leftarrow (X + Y, X - Y). \end{aligned}$$

At a length- $n = 3n'$ node, recursively transform the three subsequences $(V_{r+3j})_j$, $r \in \{0, 1, 2\}$. If their outputs at index $0 \leq i < n'$ are Y_0, Y_1, Y_2 , set

$$\begin{aligned} X_0 &= Y_0, \\ X_1 &= \text{CMul}_{128}(Y_1, \overline{\mathcal{R}_P[iH/n]}), \\ X_2 &= \text{CMul}_{128}(Y_2, \overline{\mathcal{R}_P[2iH/n]}), \\ \omega_1 &= \overline{\mathcal{R}_P[H/3]}, & \omega_2 &= \overline{\mathcal{R}_P[2H/3]}, \\ Z_i &= X_0 + X_1 + X_2, \\ Z_{i+n'} &= X_0 + \text{CMul}_{128}(X_1, \omega_1) + \text{CMul}_{128}(X_2, \omega_2), \\ Z_{i+2n'} &= X_0 + \text{CMul}_{128}(X_1, \omega_2) + \text{CMul}_{128}(X_2, \omega_1). \end{aligned}$$

The recursion is $768 \rightarrow 256$ for **IBS-1** and $1152 \rightarrow 384 \rightarrow 128$ for **IBS-3**. Normalize only after the complete length- H recursion: if Z_i is its output, set

$$D_i = (\text{Div}_{128}(\Re Z_i, H), \text{Div}_{128}(\Im Z_i, H)), \quad V_i = \text{CMul}_{128}(D_i, \overline{\mathcal{T}_P[i]}).$$

The Q128 coefficient integers are then

$$Y_{i+H} = \text{Mul}_{128}(\Im V_i, \langle 2/\sqrt{3} \rangle_{128}), \quad Y_i = \Re V_i - \text{Div}_{128}(Y_{i+H}, 2).$$

Candidate sampling with cutoff. Algorithm B.4 evaluates the candidate map using certified rounding and annular comparisons.

Algorithm B.4 ($\text{SampleFG}_P^{\text{fx}}$)

Input: a profile P and a seed $seed \in \{0, 1\}^{512}$.

Output: (f, g) , REJECT, or FAIL.

1. Read the $2N$ words of Algorithm B.1. Set

$$L_Q = \text{RNE}(2^{128} Q_{\text{low}}), \quad D_Q = \text{RNE}(2^{128} (Q_{\text{high}} - Q_{\text{low}})).$$

2. For $0 \leq j < H$, compute

$$\begin{aligned}
\tilde{r}_j &= L_Q + \text{RNE}(D_Q W_j / 2^{64}), \\
\tilde{\rho}_j &= \text{SqrtN}_{128}(\tilde{r}_j), \\
(C_{\varphi,j}, S_{\varphi,j}) &= \text{SC}_{128}(W_{j+H} 2^{66}), \\
(C_{f,j}, S_{f,j}) &= \text{SC}_{128}(W_{j+N} 2^{68}), \\
(C_{g,j}, S_{g,j}) &= \text{SC}_{128}(W_{j+3H} 2^{68}), \\
A_{f,j} &= \text{Mul}_{128}(\tilde{\rho}_j, C_{\varphi,j}), \\
A_{g,j} &= \text{Mul}_{128}(\tilde{\rho}_j, S_{\varphi,j}), \\
\hat{f}_j^{\text{fx}} &= (\text{Mul}_{128}(A_{f,j}, C_{f,j}), \text{Mul}_{128}(A_{f,j}, S_{f,j})), \\
\hat{g}_j^{\text{fx}} &= (\text{Mul}_{128}(A_{g,j}, C_{g,j}), \text{Mul}_{128}(A_{g,j}, S_{g,j})).
\end{aligned}$$

The U132 phase numerators retain every bit of each input word.

3. Apply the integer inverse embedding above, obtaining Q128 coefficient arrays Y_f, Y_g . For each array in the order f, g , use the radius E_P in Table B.1. Set $z_i = \text{RNE}(Y_i / 2^{128})$ and $d_i = Y_i - z_i 2^{128}$. A coefficient rounding is determined if $|d_i| + E_P < 2^{127}$. If the rounded coefficient sum is even, let i^* be the smallest index maximizing $|d_i|$ and require

$$|d_{i^*}| > |d_k| + 2E_P \quad (k \neq i^*), \quad |d_{i^*}| > E_P.$$

Apply DecodeOdd only when these decisions are determined. Any unresolved required decision rejects the candidate; there is no precision escalation. A rejection while processing f ends the candidate before processing g .

4. Return **FAIL** if any decoded coefficient lies outside $[-4095, 4095]$ or an arithmetic guard fails. The range bounds below exclude these events on the prescribed input domain.
5. Evaluate GoodPair $_P$ once with directed Q96 intervals on the exact integers (f, g) . The radix-two or radix-three forward transform uses root enclosures of radius at most one Q96 unit; directed products and sums enclose every component. For each conjugate pair, enclose $\nu_i = |\phi_i(f)|^2 + |\phi_i(g)|^2$. Compare with outward enclosures of q/α_{ann}^2 and $q\alpha_{\text{ann}}^2$. Return **REJECT** if an embedding is determined outside, or if an interval comparison is unresolved. Return **FAIL** on an arithmetic-guard violation. Otherwise return (f, g) .

The first H embedding tests suffice because conjugate embeddings have equal norms. Key loading uses the same Q96 annular predicate. The exact integer relation and second-row norm are checked separately.

NTRU-solver policies. The scheduled integer reducer handles the trinomial towers and the lower power-of-two levels, with scalar Bézout solving at the bottom. In the power-of-two profiles its reduction levels satisfy $2 \leq n \leq n_0$, where $n_0 = 32$ for IBS-1s and $n_0 = 64$ for IBS-2. The upper five levels, of degrees N down to $2n_0$, use Pornin’s NTRUGEN-style reduction [27], combining exact small-prime NTT/CRT lifting with scheduled Q32 FFT reduction. Each

policy has its own rounding and admission rules. Both policies use the same relative-norm descent; the scheduled reducer handles the scalar base case.

Write $U = q\alpha_{\text{ann}}^2$ and $s_U = \log_2 \sqrt{U}$. At tower degree n , let $b(n) = 0$ for a power-of-two ring or $n = 1$, and $b(n) = \log_2(2/\sqrt{3})$ otherwise. The coefficient bound

$$c_{fg}(n) = \left\lceil \frac{N}{n} s_U + b(n) \right\rceil$$

follows by relative-norm descent and inverse-embedding bounds. The power-of-two tower halves degrees. The trinomial towers are $1536 \rightarrow 512 \rightarrow 256 \rightarrow \dots \rightarrow 2 \rightarrow 1$ and $2304 \rightarrow 768 \rightarrow 256 \rightarrow \dots \rightarrow 2 \rightarrow 1$; the final quadratic norm uses $X \mapsto 1 - X$.

At the bottom, reject unless $\gcd(f_1, g_1) = 1$. For $g_1 \neq 0$, take the canonical representative $u = f_1^{-1} \bmod |g_1|$ in $[0, |g_1|)$, set $v = (uf_1 - 1)/g_1$, and return $(F, G) = (qv, qu)$. For $g_1 = 0$, require $f_1 = \pm 1$ and use $(F, G) = (0, qf_1)$. Odd coefficient sums imply odd resultants in the power-of-two rings, permitting binary-GCD kernels; the trinomial solver uses general extended GCD and makes no such parity inference. Relative-norm lifting then produces a solution at each parent level.

Scheduled integer reduction. At a scheduled level $n < N$, the reduced output must have coefficient magnitudes below $2^{c_{\text{red}}(n)}$, where $c_{\text{red}}(1) = c_{fg}(1) + 24$ and $c_{\text{red}}(n) = c_{fg}(n) + 64$ for $1 < n < N$. A violation rejects the candidate. At the top level of a trinomial tower, admission is the signed-32-bit representability test of Algorithm B.2, Step 3; the value $c_{\text{red}}(N) = 32$ only sizes carriers. For extension degree e and child degree n/e , define

$$\begin{aligned} c_{\text{lift}}(n) &= \left\lceil c_{\text{red}}(n/e) + \log_2(n/e) + (e-1) \frac{N}{n} s_U + b(n) \right\rceil, \\ c_{\text{cur}}(n) &= \lfloor c_{\text{lift}}(n) + \tfrac{1}{2} \log_2(2n) \rfloor + 1, \\ p_{\text{sol}}(n) &= \max\{p_0, 64 \lceil c_{fg}(n)/64 \rceil\}, \\ K_{\text{sol}}(n) &= c_{\text{lift}}(n) - c_{fg}(n) + p_{\text{sol}}(n) + 64, \end{aligned}$$

with $p_0 = 256$ for IBS-1s/1/3 and 384 for IBS-2. The strict-bit convention $\lfloor x \rfloor + 1$ in c_{cur} reserves one bit above an exact power-of-two boundary. The power-of-two degree-two reduction instead uses an exact rational quotient and the constant round bounds stated below. These bounds determine fixed integer capacities for the level operations; an exceeded admission bound rejects, whereas a carrier overflow is an arithmetic failure. For admitted inputs, convolution bounds control products, and Karatsuba middle terms and ring folding add at most four bits. The per-object capacities dominate these bounds, excluding carrier overflow. The cap-exhaustion bound of Section B.3 therefore concerns prescribed rejections only.

Each reduction round forms the Fourier quotient $(\widehat{F}\widehat{f} + \widehat{G}\widehat{g})/(\widehat{f}^2 + \widehat{g}^2)$ at the level precision and rounds its inverse transform to an integer polynomial k , using RNA at the scheduled power-of-two levels and RNE in the trinomial towers. Power-of-two reduction retains the denominator at $Q2p_{\text{sol}}$; tri-

nomial reduction rounds it back to $\mathbb{Q}p_{\text{sol}}$. Write $(K_{R,i}, K_{I,i})$ for the integer components encoding quotient slot i at $\mathbb{Q}p_{\text{sol}}(n)$. A nonpositive denominator or $\max\{|K_{R,i}|, |K_{I,i}|\} \geq 2^{K_{\text{sol}}(n)}$ in any active slot rejects. Thus the quotient bound is componentwise in encoded integer units; its real-unit threshold is $2^{K_{\text{sol}}(n)-p_{\text{sol}}(n)}$. Form $(F - kf, G - kg)$ exactly and retain it only when $k \neq 0$ and the sum of coefficient squares decreases. A zero k or a nondecreasing update stops reduction; subsequent rounds use masked dummy operands. The last round computes k without an update and rejects if the reduction remains active with $k \neq 0$.

For the scheduled power-of-two levels of degree $n > 2$, round $r \geq 1$ also bounds the entering (F, G) coefficients and encoded quotient components by

$$\begin{aligned} B_r(n) &= \max\{c_{\text{red}}(n), c_{\text{cur}}(n) - (r-1)(p_{\text{sol}}(n) - 192)\}, \\ K_r(n) &= \max\{p_{\text{sol}}(n) + 64, K_{\text{sol}}(n) - (r-1)(p_{\text{sol}}(n) - 192)\}. \end{aligned}$$

The first round uses $c_{\text{cur}}(n)$ and $K_{\text{sol}}(n)$. An active coefficient of magnitude at least $2^{B_r(n)}$, or encoded quotient component of magnitude at least $2^{K_r(n)}$, rejects. At degree two, the exact rational reducer uses constant (F, G) bounds 2^{18038} and 2^{36049} for IBS-1s and IBS-2, respectively, and rejects rounded integer quotient coefficients of magnitude at least 2^{12096} and 2^{24103} , respectively. These bounds apply in every round, without precision-dependent shrinkage. For the scheduled reducer, the public round limits are listed in descending tower order. All entries apply to the trinomial profiles; the power-of-two profiles use only the entries for $n \leq n_0$:

P	rounds per level above degree one
IBS-1s	3, 3, 3, 3, 4, 4, 4, 4, 4, 3
IBS-1	3, 3, 3, 3, 4, 4, 4, 4, 4, 4
IBS-2	3, 3, 3, 3, 4, 5, 4, 4, 4, 4, 3
IBS-3	3, 3, 3, 3, 4, 4, 4, 4, 4, 4, 4

Upper power-of-two reduction. At depth $d \in \{0, \dots, 4\}$, with degree $n = N/2^d$, the pairs (f, g) and the reduced (F, G) are held in

$$s_d = \left\lceil \frac{c_{fg}(n) + 1}{31} \right\rceil$$

signed 31-bit words per coefficient. At the hand-off depth $d = 5$, of degree n_0 , set $s_5 = \lceil (\max\{c_{fg}(n_0), c_{\text{red}}(n_0)\} + 1)/31 \rceil$. The unreduced lift of (F, G) uses

$$\ell_d = \left\lceil \frac{31s_{d+1} + c_{fg}(n) + \log_2 n + 1}{31} \right\rceil$$

words. Lifting is exact, in a residue-number representation over 31-bit NTT primes. At each level, (f, g) is converted to Q32 once, and its normalized Fourier factors are precomputed. Each reduction round converts the current (F, G) to Q32 at a scale offset fixed by the public word counts and by the bit length of

(f, g) , computes the FFT quotient using those factors, rounds its inverse transform to k by $\lfloor x + 1/2 \rfloor$ (ties toward $+\infty$), and always subtracts the correspondingly scaled multiple of $k(f, g)$. A public schedule lowers the scale offset by 11 bits per round until it reaches zero, and drops the high words of (F, G) that the remaining offset no longer requires. The level then rejects the candidate unless the reduced (F, G) fits s_d words, that is, every higher word is a sign copy, and $fG - gF = q$ holds modulo one 31-bit prime. This single-prime test is a prescribed rejection, not an exact integer check. At depth 0, $s_0 = 1$, so the output consists of signed 31-bit coefficients. The solver then rejects unless $fG - gF = q$ holds modulo two 31-bit primes; this test is exact, because every coefficient of $fG - gF$ has magnitude below 2^{55} , less than half the product of the primes. The second-row norm is checked before encoding.

Algorithm B.5 ($\text{Setup}_P^{\text{fx}}$)

Run Algorithm B.3 with Algorithm B.4 as the candidate sampler and the specified solver Solve_P . A prescribed rejection advances to the next independent seed. A failure ends Setup. The candidate cap is 2^{16} .

B.3 Correctness and Distribution of Fixed-Point Setup

We first establish agreement between accepted finite candidates and exact annular sampling, then derive the returned-key law and the effect of Galois screening. Acceptance probabilities quantify Setup availability; cutoff probabilities also bound the difference from exact candidate evaluation.

Transform correspondence. For a power-of-two profile, table entry k targets $\zeta_{M_P}^{\text{rev}_{\log_2 N}(k)}$. Induction through the conjugate-root inverse butterflies shows that spectral coordinate j contributes

$$H^{-1} \zeta_{M_P}^{-i(1+4 \text{rev}_{\log_2 H}(j))}$$

to output i . This is exactly the packed permutation π_P in Algorithm B.1. For the mixed-radix profiles, the cyclic roots and twists target

$$\begin{aligned} \zeta_{256}^{171} \zeta_3 &= \zeta_{768}, & \zeta_{512}^{57} \zeta_9^{-1} &= \zeta_{4608}, \\ \zeta_{128}^{57} \zeta_9^5 &= \zeta_{1152}, & \zeta_{256}^{19} \zeta_{27}^{-2} &= \zeta_{6912}. \end{aligned}$$

At a node of length $n = 3n'$, write the input index as $r + 3t$ and the output index as $i + bn'$, with $r, b \in \{0, 1, 2\}$. The recursive and combining phases satisfy

$$\zeta_n^{-3ti} \zeta_n^{-ri} \zeta_3^{-br} = \zeta_n^{-(r+3t)(i+bn')}.$$

Together with the radix-2 leaves, this proves that the recursion is the natural-order inverse DFT of length H . The final division by H , conjugate twist, and profile-specific unpacking give precisely the inverse embedding in Algorithm B.1.

Range and precision. Let $\epsilon = 2^{-128}$. The radial interpolation error is below 2ϵ , and the square-root rounding error is below one Q128 unit. The 131-step CORDIC has residual at most one U132 unit and coordinate error below

Table B.1. Q128 candidate-evaluation enclosure bounds. Radii are in Q128 units; states use signed I256 and raw products use U512.

profile	coefficient error	analytic radius (Q128 units)	operational E_P
IBS-1s	$< 2^{-97}$	1375737937	2^{32}
IBS-1	$< 2^{-84}$	13194206654658	2^{70}
IBS-2	$< 2^{-97}$	1509955673	2^{32}
IBS-3	$< 2^{-78}$	985162552730499	2^{90}

2^{-116} . The polar products give initial spectral error below 2^{-100} . Root-table recurrence gives component errors below 2^{-104} for IBS-1 and 2^{-100} for IBS-3. Propagating these errors through the inverse graph gives Table B.1.

For a butterfly product, a componentwise operand error e and root error u , with exact operand magnitude bounded by A , contribute at most $2Au + 2e + 2eu + 2\epsilon$; addition sums the incoming errors. For the normalized transform, including unpacking, a conservative stage envelope is

$$G(1 + cu)^s(e + sd((A + e)u + \epsilon)),$$

with $(s, G, c, d) = (10, 1, 4, 8), (12, 2, 4, 8), (11, 1, 4, 8), (14, 4, 8, 16)$ for IBS-1s/1/2/3, respectively. Applying this envelope to the input and root error bounds and rounding upward to Q128 units gives the analytic radii in Table B.1. Radial bounds and inverse-embedding inequalities also give exact coefficient magnitudes at most $(2/\sqrt{3})\sqrt{Q_{\text{high}}}$ in the trinomial rings and $\sqrt{Q_{\text{high}}}$ otherwise. Including the at-most-one change from DecodeOdd keeps every candidate coefficient below 4095 in magnitude.

Lemma B.1 (Range safety). The Q128 candidate graph fits its I256 states and U512 products. The Q96 annular predicate fits its fixed carriers on decoded candidate inputs from Algorithm B.4; its directed intervals contain the exact embedding norms.

Proof. Every initial radial factor is below 2^{12} . Triangle bounds through each transform stage bound the scaled states, while the componentwise error recurrence bounds deviations from the exact graph. The resulting quantities fit I256 and their full products fit U512. For the Q96 forward transform, each stage sums intervals enclosing bounded integer inputs and multiplies them by certified root balls. Directed multiplication, normalization, and summation preserve inclusion. The coefficient cap bounds all forward transform intermediates. Each decoded coefficient differs from its sampled real value by at most one, so the final embedding magnitude is at most $\sqrt{Q_{\text{high}}} + N$. This stronger candidate-specific bound keeps the squared-norm sums within the Q96 carriers. $\square$

Lemma B.2 (Correctness of determined candidate decisions). Whenever Algorithm B.4 accepts, Algorithm B.1 returns the same (f, g) . A determined mathematical rejection agrees with the exact predicate; an unresolved comparison is a rejection of the finite generator only.

Proof. Table B.1 encloses each exact coefficient y_i in $[Y_i - E_P, Y_i + E_P]2^{-128}$. The strict rounding test places this interval inside one rounding

cell. The residual-separation and sign tests similarly determine the correction in DecodeOdd. Q96 interval inclusion proves the exact annular inequalities whenever GoodPair accepts. Composing these determined decisions proves the claim. An unresolved comparison supplies no decision about the corresponding exact predicate, so its rejection is retained in the specified finite law. $\square$

Let $\mathcal{K}_P^{\text{fx}}$ compose Algorithms B.4 and B.2, and let $\mathcal{K}_P^{\text{ex}}$ compose Algorithms B.1 and B.2 with the same solver and integer admission tests.

Proposition B.1 (Agreement on accepted candidates). For every seed for which the finite candidate is accepted,

$$\mathcal{K}_P^{\text{fx}}(\text{seed}) = \mathcal{K}_P^{\text{ex}}(\text{seed}). \quad (\text{B.2})$$

Proof. Lemma B.2 gives the same integer pair; both maps then apply the same deterministic NTRU solver, admission, and encoding. This implication does not assert equality of their rejection sets. $\square$

Let $\mathcal{O}$ assign an independent uniform bit tape to every distinct XOF input, and let $\mathcal{K}_{P,\mathcal{O}}^{\text{fx},0}$ omit only the public Galois screen from the specified map. Its successful public-key law is $\mathcal{H}_{\text{ann}}$ in the ideal-XOF experiment.

Corollary B.1 (Specified Setup distribution). Conditional on a fixed ideal XOF $\mathcal{O}$, the seed of the first successful candidate of Algorithm B.5 is uniform over the accepted seeds, where acceptance includes the Galois screen. The returned key is the image of this seed under the specified map. The finite cap changes the return probability but not this conditional law. The joint oracle/key law conditions the entire Setup experiment on successful output.

Proof. Fix $\mathcal{O}$ and let $\Omega_P \subseteq \{0,1\}^{512}$ be the set of seeds on which $\mathcal{K}_{P,\mathcal{O}}^{\text{fx}}$ returns a key: Algorithm B.4 accepts and no step of Algorithm B.2 rejects. Write $p_P = p_P(\mathcal{O}) = |\Omega_P|2^{-512}$ and $T_K = 2^{16}$. Algorithm B.5 fails with probability $(1 - p_P)^{T_K}$. If J is the index of the first accepted seed, then, for $p_P > 0$,

$$\Pr[\text{seed}_J = s \mid J < T_K] = \frac{1}{|\Omega_P|} \quad (s \in \Omega_P). \quad (\text{B.3})$$

Indeed, before conditioning, every $s \in \Omega_P$ has the same mass

$$2^{-512} \sum_{j=0}^{T_K-1} (1 - p_P)^j.$$

The common geometric weight also holds when a seed is repeated, and normalization gives Eq. (B.3). Conditioning on successful Setup weights the oracle marginal by $1 - (1 - p_P(\mathcal{O}))^{2^{16}}$. Averaging the fixed-oracle laws with these weights gives the joint successful-output law. $\square$

Galois screening. The stabilizer of h is a subgroup of the corresponding Galois group. Each of the four group orders has no prime divisor other than 2 or 3; by Cauchy's theorem, every nontrivial stabilizer therefore contains an element of order 2 or 3. Algorithm B.2 tests every nonidentity element of order 2 and, for the trinomial rings, every element of order 3, so it rejects exactly the keys with nontrivial stabilizer.

For a fixed candidate map, Galois screening conditions D_{std} on E_{stab} . To compare the joint ideal-XOF laws, let $p_P^0(\mathcal{O})$ and $p_P(\mathcal{O}) \leq p_P^0(\mathcal{O})$ be the unscreened and screened candidate acceptance probabilities, and put $g_T(p) = \sum_{j=0}^{T-1} (1-p)^j$. Before normalization, a particular successful seed has mass proportional to $g_{T_K}(p_P^0(\mathcal{O}))$ or $g_{T_K}(p_P(\mathcal{O}))$, respectively. The latter is at least the former, while the overall screened return probability is no greater. Hence every screened seed has at least its unscreened mass after normalization. The screened law is supported on E_{stab} , so its distance from the unscreened joint law is the unscreened probability of $\neg E_{\text{stab}}$. Data processing and Eqs. (19)–(20) bound the transition by $\Delta_{\text{stab},P} + \varepsilon_{\text{stab},P}$.

Acceptance probabilities. The measured acceptance counts for the complete map $\mathcal{K}_P^{\text{fx}}$ were 992, 979, 971, and 999 out of 1000 distinct candidate seeds for IBS-1s/1/2/3. In the ideal-XOF model these distinct inputs yield independent candidate tapes. One-sided 95% Clopper–Pearson lower bounds are 0.9856, 0.9698, 0.9606, and 0.9952, respectively; each gives $(1 - p_P)^{2^{16}} < 2^{-256}$ at its stated confidence level. These are per-profile empirical confidence bounds, not worst-case acceptance guarantees.

Cutoff comparison. Let $\mathcal{A}_P$ be admission under the same NTRU-solver policy with exact candidate decisions, and let $\zeta_{\text{cut},P}$ be the probability of a cutoff conditional on $\mathcal{A}_P$. When $\Pr[\mathcal{A}_P] > 0$ and $\zeta_{\text{cut},P} < 1$, conditioning gives

$$\Pr_{\text{fx}}[E] \leq \frac{\Pr_{\text{ex}}[E]}{1 - \zeta_{\text{cut},P}}, \quad \zeta_{\text{cut},P} \leq \frac{\Pr[\text{cutoff}]}{\Pr[\mathcal{A}_P]}.$$

Here both output laws use the same choice of including or omitting the Galois screen. By Lemma B.2 and Proposition B.1, finite acceptance implies $\mathcal{A}_P$, so $\Pr[\mathcal{A}_P] \geq p_P$ in the same experiment. Among $2 \cdot 10^6$ distinct XOF inputs per profile, no cutoff was observed. In the independent-tape model, the one-sided 95% upper bound is $1 - 0.05^{1/(2 \cdot 10^6)} < 1.5 \cdot 10^{-6}$ per candidate. Combining this with the per-profile availability lower bound gives $\zeta_{\text{cut},P} < 1.6 \cdot 10^{-6}$ at at least 90% joint confidence for each profile. Assumption 4.1 concerns the specified finite-generator law and does not depend on these empirical bounds.

B.4 Fixed-Point Hybrid-Sampling Extraction

Fix a profile P and write $N = N_P$ and $R = R_P$. In Prest’s rank-two Hybrid-Sampler [28], denote the continuous perturbations by y_1, y_2 and the discrete coordinates, in sampling order, by x_2, x_1 . Directed Q96 interval arithmetic, with Q120 reciprocals and integer square roots, prepares the projection coefficients and perturbation widths. Continuous perturbations are sampled on a Q53 grid and scaled by Q62 profile and width factors. Gaussian rejection tests approximate exponential acceptance probabilities using Q64 exponents and range-reduced polynomial evaluation: Q63 for perturbation and scalar sampling, and Q64 for paired sampling. The resulting probabilities are converted to integer acceptance thresholds; integer tables supply the discrete proposals. Transforms

with Q62 and Q126 roots combine the perturbations and cached coefficients to produce Q64 sampling centers.

The distributional comparisons use fresh private uniform tapes for each materialized identity in the noncolliding-tape abstraction of Section 4.1; repeated identity requests reuse cached state. The keyed-tape and root-collision costs are included in $\text{Adv}_{\text{ext-rand}}$, and Section B.7 accounts for the joint XOF realization.

Exact Hybrid-Sampler. Write $b_1 = (g, -f)$ and $b_2 = (G, -F)$, and put $\widehat{x} = (\phi_i(x))_{i=1}^N$ for the canonical embedding of x . For each i , the NTRU relation gives

$$\|b_{1,i}^*\|^2 = \nu_i, \quad \|b_{2,i}^*\|^2 = q^2/\nu_i, \quad \nu_i = |\widehat{f}_i|^2 + |\widehat{g}_i|^2.$$

Set $r = 129/100$ and define

$$\begin{aligned} u_i &= \widehat{f}_i/q, & \beta_i &= \frac{\widehat{g}_i}{\nu_i}, \\ \gamma_i &= -\frac{\widehat{F}_i\widehat{f}_i + \widehat{G}_i\widehat{g}_i}{\nu_i}, & \tau_{1,i} &= \sqrt{\frac{\sigma_1^2}{\nu_i} - r^2}, \\ & & \tau_{2,i} &= \sqrt{\frac{\sigma_1^2\nu_i}{q^2} - r^2}. \end{aligned} \tag{B.4}$$

For a positive conjugate-symmetric width vector τ , let $\text{CG}_P(\tau; \mathcal{S})$ use tape $\mathcal{S}$ to draw a conjugate-symmetric canonical vector Y with density proportional to

$$\exp\left(-\frac{1}{2N} \sum_{i=1}^N \frac{|Y_i|^2}{\tau_i^2}\right)$$

and return the unique $y \in R \otimes \mathbb{R}$ satisfying $\widehat{y} = Y$. Each real coordinate of a conjugate pair therefore has variance $N\tau_i^2/2$. For $d \in R \otimes \mathbb{R}$, let $D_{R,r,d}^{\text{can}}$ assign to $z \in R$ mass proportional to

$$\exp\left(-\frac{\|z - d\|_{\text{can}}^2}{2r^2}\right).$$

In the power basis, this law is coefficientwise for a power-of-two ring. For a trinomial ring it factors over pairs $(i, i + N/2)$ with quadratic form $u^2 + uv + v^2$.

Extraction root and tapes. The IBS wrapper computes the first unit label

$$(j, a) = H_{\text{id}}(h, \text{id}).$$

The concrete instantiation of the keyed tape G_{ext} in Section 3 defines the 64-byte root

$$\text{root} = \text{SHAKE256}('G' \parallel k_{\text{ext}}(B) \parallel \text{id} \parallel \text{u16le}(j)) \tag{B.5}$$

by reading the first 64 output bytes. It then derives four domain-separated tapes

$$\mathcal{S}_i = \text{SHAKE256}('E' \parallel \text{root} \parallel \text{octet}(i)), \quad 0 \leq i < 4. \tag{B.6}$$

Here $\text{octet}(i)$ is the one-byte encoding of i . Each sampling call below resumes its assigned stream after an internal rejection.

Algorithm B.6 ($\text{SamplePreimage}_P^{\text{exact}}$)

Input: an annular basis B for h in profile P , a syndrome a , the profile extraction width σ_1 , and a 512-bit root root .

Output: a centered witness $S = (s_0, s_1)^T \in R^2$.

1. Prepare the quantities in Eq. (B.4), take the lift $\bar{a} \in R$ of a with coefficients in $\{0, \dots, q-1\}$, and derive $(\mathcal{S}_0, \dots, \mathcal{S}_3)$ from Eq. (B.6).
2. Independently draw $y_1 \leftarrow \text{CG}_P(\tau_1; \mathcal{S}_0)$ and $y_2 \leftarrow \text{CG}_P(\tau_2; \mathcal{S}_1)$.
3. Form the center

$$\hat{d}_{2,i} = \hat{a}_i u_i - \hat{y}_{2,i} \quad (1 \leq i \leq N)$$

and draw $x_2 \leftarrow D_{R,r,d_2}^{\text{can}}$ using $\mathcal{S}_3$.

4. Form the center

$$\hat{d}_{1,i} = \hat{a}_i \beta_i + \hat{x}_{2,i} \gamma_i - \hat{y}_{1,i} \quad (1 \leq i \leq N)$$

and draw $x_1 \leftarrow D_{R,r,d_1}^{\text{can}}$ using $\mathcal{S}_2$.

5. Compute $v = x_1 b_1 + x_2 b_2$ in the integer ring and set $S = \text{center_mod}_q(\bar{a} - v_0, -v_1)^T$.
6. Return S .

Steps 3 and 4 specialize the Gram–Schmidt projections using the NTRU identity, with d_1 depending on the preceding x_2 . Since $b_1, b_2 \in \Lambda_{h,q}$, the returned witness satisfies $s_0 + s_1 h = a \bmod q$. Its exact output law is independent of the lift: replacing $\bar{a}$ by $\bar{a} + qe$, for $e \in R$, translates the discrete draws in law by $x_2 \mapsto x_2 + ef$ and $x_1 \mapsto x_1 - eF$, hence translates v by $(qe, 0)$ and leaves S unchanged. The finite sampler and its intermediate-law comparisons below use the nonnegative lift.

Integer preparation. Loading a valid integer basis constructs its extraction cache once for reuse across identities. At precision p , represent dyadic intervals by

$$\mathcal{I}_p(L, U) = [L2^{-p}, U2^{-p}].$$

Addition and subtraction use the exact endpoint pairs $(L_1 + L_2, U_1 + U_2)$ and $(L_1 - U_2, U_1 - L_2)$. For multiplication, let m and M be the minimum and maximum of the four integer endpoint products and return

$$\mathcal{I}_p\left(\left\lfloor \frac{m}{2^p} \right\rfloor, \left\lceil \frac{M}{2^p} \right\rceil\right). \quad (\text{B.7})$$

If $0 < C \leq D$, division of $\mathcal{I}_p(L, U)$ by $\mathcal{I}_p(C, D)$ encloses the four quotients $2^p L/C$, $2^p L/D$, $2^p U/C$, and $2^p U/D$ by directed integer division. Finally,

$$\sqrt{\mathcal{I}_p(L, U)} \subseteq \mathcal{I}_p\left(\left\lfloor \sqrt{L2^p} \right\rfloor, \left\lceil \sqrt{U2^p} \right\rceil\right) \quad (\text{B.8})$$

for $0 \leq L \leq U$, using integer square root and widened square comparison. Complex arithmetic applies these rules componentwise.

Preparation evaluates Eq. (B.4) once with directed Q96 intervals, using I128 endpoint carriers and certified radix-two or radix-three roots. The constants are the exact rationals

$$\sigma_1^2 = 1909047670656633/10^8, \quad r^2 = 16641/10^4.$$

Reciprocals of q and ν_i are enclosed on a Q120 grid; multiplication by them returns a Q96 interval. The expression for γ_i divides the first-row spectra by ν_i before multiplying by the second row, limiting intermediate magnitudes. Each radicand interval is intersected with $[0, \infty)$ before square-root evaluation; an entirely negative interval is a failure.

For an enclosure $[L, U]2^{-96}$, take its floored Q96 midpoint and round to nearest-even Q74 for u_i, β_i and Q52 for γ_i . Widths first round to the binary64 value set with ties to even, then to Q62 with ties away from zero; both maps are implemented by integer significands and discarded-bit tests. No precision escalation or multiprecision fallback is used.

If the real enclosure width is w , a component stored with f fractional bits has error at most

$$2^{-f-1} + w/2 + 2^{-96}.$$

For a width, use $\sqrt{x+w} - \sqrt{x} \leq \sqrt{w}$ for $x, w \geq 0$ and add the binary64/Q62 rounding allowance $4(5/4)2^{-53} + 2^{-63}$. Uniformly over valid keys, the resulting width errors are below $2^{-39.19}$, $2^{-38.17}$, $2^{-38.57}$, and $2^{-37.50}$ for IBS-1s/1/2/3. The projection comparisons use separate bounds for τ_1 and τ_2 .

The enclosing widths for $u, \beta, \gamma, \tau_1, \tau_2$, measured in Q96 units, are bounded by the guards $2^7, 2^{11}, 2^{35}, 2^{60}, 2^{60}$. These bounds follow by propagating root-ball errors through the forward transform, sums of squares, Q120 reciprocals, and directed products in Eq. (B.4). Annularity lower-bounds every ν_i , and the second-row norm bounds every second-row spectrum. These estimates also bound γ , uniformly over every pair (F, G) satisfying the second-row norm test. Range guards require the rounded projection coefficients to fit I64 and the stored Q62 widths to lie in $[0, 5/4]$. A violated enclosure-width or range guard signals failure; Proposition B.2 shows that preparation adds no key conditioning.

The compact master-key decoder reconstructs the omitted polynomial G from the NTRU relation. The second-row norm implies $|G_i| < 2^{17} < q/2$ for every profile, so the centered lift of $gF/f \bmod q$ reconstructs G . The implementation checks the integer NTRU equation before committing and retains the general decoder as fallback.

Proposition B.2 (Preparation of every valid key). For the four profiles in Table 1, $\text{Valid}_P(B)$ implies that the specified integer preparation returns a ready extraction context, under the mathematical resource assumptions of this appendix. Preparation adds no conditioning to the returned-key law of Algorithm B.5.

Proof. In the notation of Eq. (B.4), annularity gives

$$|\widehat{f}_i/q|, |\beta_i| \leq \frac{\alpha_{\text{ann}}}{\sqrt{q}}, \quad \tau_{j,i}^2 \leq r^2(\alpha_{\text{ann}}^4 - 1).$$

Real coefficients give conjugate pairs of embeddings. Parseval and the second-row norm condition therefore imply

$$|\widehat{F}_i|^2 + |\widehat{G}_i|^2 \leq \frac{N}{2} q \alpha_{\text{ann}}^2 k_N, \quad |\gamma_i| \leq \frac{\sqrt{|\widehat{F}_i|^2 + |\widehat{G}_i|^2}}{\sqrt{\nu_i}} \leq \alpha_{\text{ann}}^2 \sqrt{\frac{N}{2}} k_N.$$

The resulting γ bounds are below 496.562, 608.162, 992.157, and 911.651 in profile order. Including the prescribed rounding errors, these bounds fit the prepared caches. The rounded perturbation widths are below $1.206 < 5/4$.

The exact perturbation widths are nonnegative by annularity. A stored width may be zero, and the sampler consumes the same records in that case. The Q96 enclosure bounds, including the square-root bound at zero, lie within the width guards for every valid key. Midpoint rounding preserves the stated output ranges. Factoring the projection as $\gamma_i = -(\widehat{F}_i \widehat{f}_i / \nu_i + \widehat{G}_i \widehat{g}_i / \nu_i)$ keeps intermediates in range. Construction of the modular transforms then returns a ready context. $\square$

Imported keys must satisfy $\text{Valid}_P(B)$ before preparation.

Continuous perturbations. Each proposal of the grid-Gaussian (crown) sampler consumes an 18-byte record containing, in order, a uniform 16-bit selector v , a uniform 55-bit coordinate w and sign bit b , and a uniform 72-bit integer U ; multibyte fields are little-endian. Set

$$(m_0, \dots, m_{24}) = (8633, 8492, 8103, 7494, 6718, 5836, 4915, 4011, 3173, 2433, \\ 1808, 1303, 910, 616, 404, 257, 158, 95, 55, 31, 17, 9, 5, 3, 2),$$

and $m_\ell = 1$ for $25 \leq \ell < 80$, so $\sum_{\ell=0}^{79} m_\ell = 2^{16}$. The block index is the unique ℓ satisfying $\sum_{j < \ell} m_j \leq v < \sum_{j \leq \ell} m_j$. Define

$$\lambda_\ell = \text{RNE} \left(2^{64} \left(\frac{\ell^2}{64} + \ln \frac{m_\ell}{8625} \right) \right), \quad Z_{\text{mag}} = \ell 2^{55} + w.$$

The Q64 acceptance exponent is

$$E = \lambda_\ell + 16\ell w + \text{RNE}(w^2/2^{52}).$$

Write $E = sL + R$, where $L = \text{RNE}(2^{64} \ln 2) = \text{0xb17217f7d1cf79ac}$ and $0 \leq R < L$. The degree-nine Q63 evaluator uses the following integers, written in hexadecimal:

$$\begin{aligned} c_0 &= 000010640465bc10, c_5 = 0555552288547964, \\ c_1 &= 0000c6bb1df37522, c_6 = 1555554ff148016c, \\ c_2 &= 00067908673fb508, c_7 = 3fffffffb57f7500, \\ c_3 &= 002d7f463a397537, c_8 = 7fffffff700bcf, \\ c_4 &= 01110ffd6c1cb838, c_9 = 8000000000000000. \end{aligned}$$

Initialize $A = c_0$ and successively set $A = c_i - \lfloor [R/2]A/2^{63} \rfloor$ for $i = 1, \dots, 9$. Accept if $U < \lfloor A2^{9-s} \rfloor$, except that $Z_{\text{mag}} = 0, b = 1$ is rejected. The accepted

Q53 integer is $Z = (2b - 1)Z_{\text{mag}}$, so zero has a unique encoding. Each rejection consumes the complete record and advances the same stream. An output permits at most 256 proposals; exhaustion terminates extraction.

With exact offsets and acceptance probabilities, the proposal masses m_ℓ cancel, leaving mass proportional to $\exp(-Z^2/2^{116})$ on $|Z| < 320 \cdot 2^{53}$. Thus the target grid has variance parameter 512 after Q53 rescaling. Its rounded continuous reference is $\text{RNE}(2^{53}G)$ for $G \sim \mathcal{N}(0, 512)$. The comparison below accounts for the finite support, exponent rounding, polynomial evaluation, and threshold quantization. Proposal-cap exhaustion is below $2^{-998.70}$ per output and below $2^{-938.5}$ over 2^{48} extractions for every profile, using at most $2N$ outputs per extraction.

The public profile normalization

$$\omega_P = \sqrt{N_P/1024} \quad (\text{B.9})$$

is part of the fixed-point map: $\omega_P G$ has variance $N_P/2$, as required by CG_P . Let $\hat{\omega}_P = \text{RNA}(2^{62}\omega_P)$ and let $\hat{\tau}_i$ be the prepared Q62 width. Each accepted coordinate is scaled by

$$Z_P = \text{RNA}(Z\hat{\omega}_P/2^{62}), \quad \tilde{Y}_i = \text{RNA}(Z_P\hat{\tau}_i/2^{51}). \quad (\text{B.10})$$

Here Z_P is Q53 and $\tilde{Y}_i$ is Q64. The first rounding is the identity for IBS-1s; otherwise both roundings are successive, with the width multiplication using the rounded Z_P . Fill real and imaginary coordinates in conjugate-symmetric half-spectrum order to approximate $\text{CG}_P(\tau)$.

Discrete ring Gaussians. Set

$$t_0 = \frac{1}{2r^2}, \quad t_1 = \frac{3}{8r^2}, \quad A_i = \lfloor 2^{64}t_i \rfloor, \quad a_i = A_i 2^{-64}, \quad T_i = \text{RNE}(2^{64}t_i).$$

The corresponding integers are

$$\begin{aligned} A_0 &= T_0 = 5542558762607280697, \\ A_1 &= 4156919071955460522, \quad T_1 = A_1 + 1. \end{aligned}$$

For $i \in \{0, 1\}$, define the half-Gaussian proposal

$$P_i(j) = e^{-a_i j^2} / Z_i, \quad Z_i = \sum_{j \geq 0} e^{-a_i j^2}.$$

Its U96 table assigns mass $M_i(j) = \lfloor 2^{96}P_i(j) \rfloor$ to each $j \geq 1$ and the remaining mass to zero. The respective supports are $0, \dots, 14$ and $0, \dots, 17$. Directed interval evaluation isolates every table entry and bounds the omitted infinite tail.

First consider the ideal proposals P_i . For a scalar center $c = k + f$, with $k = \lfloor c \rfloor$ and $0 \leq f < 1$, draw $J \leftarrow P_0$ and a fair bit V , and set $m = V + (2V - 1)J$. Accept $z = k + m$ with probability

$$\kappa_{\text{acc}} \exp(-t_0(m - f)^2 + a_0 J^2), \quad \kappa_{\text{acc}} = K 2^{-64}, \quad K = 18272768353266032322.$$

Here $K = \text{RNE}(2^{64}\sigma_{\min}/r)$ and $\sigma_{\min} = 0\text{x}147201\text{BF}1\text{F}7\text{A}75/2^{52}$. Since $|m - f| \geq J$ and $a_0 \leq t_0$, this is a valid probability. Each integer has a unique proposal representation, and its accepted mass is $\kappa_{\text{acc}} e^{-(z-c)^2/(2r^2)}/(2Z_0)$. Conditioning on acceptance therefore gives $D_{\mathbb{Z},r,c}$.

For a trinomial pair centered at (c_0, c_1) , first propose

$$z_1 = \lfloor c_1 \rfloor + V_1 + (2V_1 - 1)J_1, \quad J_1 \leftarrow P_1,$$

then form $\mu_0 = c_0 - (z_1 - c_1)/2$ and propose

$$z_0 = \lfloor \mu_0 \rfloor + V_0 + (2V_0 - 1)J_0, \quad J_0 \leftarrow P_0.$$

The signs and proposal selectors are independent. Accept jointly with probability e^{-E} , where

$$E = t_0(z_0 - \mu_0)^2 + t_1(z_1 - c_1)^2 - a_0J_0^2 - a_1J_1^2 \geq 0.$$

Completing the square shows that the accepted mass is

$$\frac{1}{4Z_0Z_1} \exp\left(-\frac{(z_0 - c_0)^2 + (z_0 - c_0)(z_1 - c_1) + (z_1 - c_1)^2}{2r^2}\right).$$

Thus conditioning gives exactly the canonical pair law. This joint test incorporates the conditional normalizer through the proposal masses.

The fixed sampler uses the U96 tables and integer exponents. At scalar Q64 center $C = k2^{64} + \varphi$, it computes

$$E_{\text{sc}} = \text{RNE}\left(\frac{T_0(m2^{64} - \varphi)^2}{2^{128}}\right) - A_0J^2. \quad (\text{B.11})$$

For Q64 pair centers (C_0, C_1) , it retains the conditional half-bit exactly:

$$\begin{aligned} J_c &= 2C_0 + C_1 - z_12^{64}, \\ \Delta_0 &= 2^{65}z_0 - J_c, \quad \Delta_1 = 2^{65}z_1 - 2C_1, \\ E_{\text{pair}} &= \text{RNE}\left(\frac{T_0\Delta_0^2 + T_1\Delta_1^2}{2^{130}}\right) - A_0J_0^2 - A_1J_1^2. \end{aligned} \quad (\text{B.12})$$

Both exponents are nonnegative Q64 integers. The pair numerator is rounded once. Removing integer center parts before multiplication bounds the scalar and pair numerators by 199 and 202 bits, respectively, so U256 arithmetic suffices.

Write $E = kL + u$, $0 \leq u < L$, with $L = 0\text{x}b17217\text{f}7\text{d}1\text{c}\text{f}79\text{ac}$. The scalar threshold uses the degree-9 Q63 subtraction-Horner graph of the perturbation sampler: for its prescribed coefficients $c_0, \dots, c_9$, initialize $A = c_0$ and successively set $A = c_i - \lfloor [u/2]A/2^{63} \rfloor$, then replace A by $\lfloor AK/2^{64} \rfloor$. The pair threshold uses $c_j = \text{RNE}(2^{64}/j!)$: initialize $A = c_{20}$ and set $A = c_j - \text{RNE}(uA/2^{64})$ for $j = 19, \dots, 0$. With $b = 63$ for the scalar graph and $b = 64$ for the pair graph, define

$$W_{96} = \lfloor A2^{96-b-k} \rfloor, \quad T_{56} = \lfloor W_{96}/2^{40} \rfloor. \quad (\text{B.13})$$

An independent U56 integer accepts precisely when it is below T_{56} . A scalar proposal consumes 20 bytes: a 12-byte selector, one sign byte, and a 7-byte acceptance word. A pair consumes 32 bytes: the J_1 and J_0 selectors, one byte whose low two bits are V_1, V_0 , and the acceptance word. Multibyte words are little-endian; unused sign bits are discarded. Every rejection consumes a complete record. Each scalar or pair call permits 1024 proposals and returns $\perp$ on exhaustion.

Uniform acceptance lower bounds are 0.756588 for the scalar call and 0.602442 for the pair call, giving mean-proposal upper bounds 1.322 and 1.660, respectively. There are $2N$ scalar calls per power-of-two extraction and N pair calls per trinomial extraction. Over 2^{48} extractions, the union bound for exhaustion is below $2^{-1303.53}$ in every profile.

The fixed-point extractor applies the signing-key admissibility predicate

$$n_{\min, P} \leq \|X^u S\|_2^2 \leq n_{\max, P} \quad (0 \leq u < N), \quad (\text{B.14})$$

with endpoints from Table B.3. One norm computation suffices in a power-of-two ring, where monomials preserve coefficient norm; a trinomial profile evaluates all N shifts.

Algorithm B.7 ($\text{SamplePreimage}_P^{\text{fx}}$)

Input: an admitted preparation context for (B, h) , a syndrome a , and a 512-bit root $root$.

Output: a centered admissible witness $S = (s_0, s_1)^T$, or $\perp$. Exhaustion of a crown or discrete sampler terminates the call.

1. Derive $(\mathcal{S}_0, \dots, \mathcal{S}_3)$ from Eq. (B.6).
2. Independently draw two grid-Gaussian vectors from $\mathcal{S}_0$ and $\mathcal{S}_1$. Apply Eqs. (B.9)–(B.10) coordinatewise with the prepared widths $(\tau_{1,i})_i$ and $(\tau_{2,i})_i$, obtaining the Q64 canonical perturbations $\tilde{Y}_1, \tilde{Y}_2$.
3. Transform the nonnegative lift $\tilde{a}$ of a , with Q29 input scaling for a power-of-two profile and Q28 scaling for a trinomial profile. Multiply pointwise by the Q74 cache $(u_i)_i$, using an RNA shift of 39 or 38 bits, respectively. Subtract $\tilde{Y}_2$ and apply the Q62 inverse transform with exact public normalization to obtain the Q64 center D_2 .
4. Sample the I32 vector x_2 at center D_2 from $\mathcal{S}_3$, coefficientwise in a power-of-two ring and by Eq. (B.12) in a trinomial ring.
5. Transform x_2 as Q64 data using Q126 roots. Multiply the syndrome transform by the Q74 cache $(\beta_i)_i$ with the same profile-dependent shift, and the x_2 transform by the Q52 cache $(\gamma_i)_i$ with a 52-bit RNA shift. Add the two Q64 products, subtract $\tilde{Y}_1$, and apply the Q126 inverse transform with exact public normalization to obtain the Q64 center D_1 .
6. Sample the I32 vector x_1 at center D_1 from $\mathcal{S}_2$, using the same scalar or paired procedure. With exact modular polynomial arithmetic, compute

$$s_1 = \text{center_mod}_q(x_1 f + x_2 F), \quad s_0 = \text{center_mod}_q(a - s_1 h).$$

7. Return $S = (s_0, s_1)^T$ if it satisfies Eq. (B.14); otherwise return $\perp$.

Consequently, every returned witness satisfies

$$s_0 + s_1 h = a \pmod{q}. \quad (\text{B.15})$$

The witness gate in Eq. (B.14) and the crown and discrete-sampler caps are executed checks; their failure terminates extraction without deriving a new root or restarting a stream. The preparation width guards are also executed checks, whose violation signals failure before the extraction context is admitted.

Distributional comparison. For at most $Q_{\text{ID}} = 2^{48}$ adaptive extractions, let $\text{Law}_{\text{ext},P}^{\text{fx,raw}}$ be Algorithm B.7 with its witness gate and both crown and discrete-sampler caps removed, and let $\text{Law}_{\text{ext},P}^{\text{hyb}}$ be Algorithm B.6 with the same output encoding.

At Rényi order $a = 512$, let $L_{q,P}$, $L_{e,P}$, and $L_{o,P}$ denote the lifetime bounds on $\log_2 R_a$ for threshold quantization, discrete-sampler evaluation, and perturbation/projection errors, respectively. Their composition bounds the raw extraction law; restoring the witness gate and sampler caps adds the extraction-failure term.

For the comparison of the x_1 -center D_1 , define the proof-only event

$$C_P = \begin{cases} 2^{17}, & P = \text{IBS-1s}, \\ 2^{18}, & \text{otherwise}, \end{cases}$$

$$G_{2,P} = \{\|x_2\|_{\text{can}} < C_P \text{ at every extraction}\}.$$

The implementation does not test $G_{2,P}$. This event restricts the intermediate laws used to compare the x_1 -center D_1 ; the finite sampler satisfies this event automatically, and the reference-tail analysis below bounds the cost of conditioning the comparison laws on it.

In order, these comparisons remove U56 threshold quantization, replace the finite discrete sampler by the exact width- r scalar or pair law, and restore continuous perturbations and exact projections. The first two retain the bounded grid perturbations and fixed projection graph. The final comparison lifts each grid value to a latent Gaussian in its rounding cell and compares the joint law with the continuous Hybrid-Sampler; its projection bounds apply on the bounded support of the first law.

Threshold quantization. Write g_x for a finite proposal mass, $p_x = W_{96}(x)2^{-96}$, $\epsilon = 2^{-56}$, and $\alpha = \sum_x g_x p_x$. Downward rounding removes at most ϵ from each acceptance probability. Its integrated moment bound is

$$\log_2 R_a(P_{56} \| P_{96}) \leq \frac{1}{a-1} \log_2 \left(1 + \frac{(a-1)G}{\alpha(1-\epsilon/\alpha)^a} \right),$$

$$G = \sum_{p_x > 0} g_x \min \left(p_x, \epsilon, \frac{a\epsilon^2}{2p_x} \right).$$

Set $\delta_x = p_x - 2^{-56} \lfloor 2^{56} p_x \rfloor$ and $t = \sum_x g_x \delta_x / \alpha$. The inequality $(1-u)^a \leq 1 - au + (a-1) \min(u, au^2/2)$, for $0 \leq u \leq 1$, bounds the numerator of the

normalized moment. Convexity gives $(1 - t)^a \geq 1 - at$, and $t \leq \epsilon/\alpha$ yields the displayed expression. This direction permits outcomes with zero rounded acceptance. Summing rigorous acceptance lower bounds for each proposal magnitude over all signs and fractional centers bounds G .

The additional U56 losses per accepted output are below $2.337 \cdot 10^{-31}$ and $6.101 \cdot 10^{-31}$ bits for scalar and paired sampling, respectively.

Discrete-sampler evaluation. Let v bound the proposal total variation, C_{prop} its pointwise excess over the ideal proposal, δ the relative acceptance-weight error, and $p_{\min}$ the ideal acceptance lower bound. Put $\ell = (v + 2^{-96})/(p_{\min}(1 - \delta))$ and $d = 2\delta/(1 - \delta)$. The certified order-512 bound is

$$\log_2 R_a \leq \frac{1}{a-1} \log_2 \left(1 + \frac{a(a-1)d^2}{2(1-d)^{a+1}} \right) + \frac{a}{a-1} \log_2 \frac{C_{\text{prop}}}{1-\ell}.$$

The first term compares normalized weights under the ideal proposal. The second follows from pointwise domination after finite-table and U96 rounding; it is applied directly to the same moment.

The evaluator-law losses per accepted output are below $2.328 \cdot 10^{-24}$ and $1.074 \cdot 10^{-27}$ bits for scalar and paired sampling, respectively.

Perturbations and projections. The perturbation sampler contributes less than $3.304 \cdot 10^{-22}$ bits per output relative to its Q53 grid law. That grid is compared pointwise with the rounded Gaussian, at less than $4.342 \cdot 10^{-36}$ bits per coordinate. These bounds are uniform over the admitted context, center, and preceding adaptive transcript. The remaining comparison controls the displacement of the discrete centers.

Center-displacement bound. Let L be the canonical image of $\mathbb{Z}$, of a trinomial coefficient pair, or of an orthogonal sum of these blocks. For $r = 1.29$, every center c satisfies $\text{Cov}(D_{L,r,c}) \preceq C_{\text{cov}} r^2 I$, with $C_{\text{cov}} = 1.000000000000713$. Consequently, for every $a > 1$ and centers c, d ,

$$\log_2 R_a(D_{L,r,c} \| D_{L,r,d}) \leq \frac{a C_{\text{cov}}}{2r^2 \ln 2} \|c - d\|_2^2.$$

Proof. For one scalar or paired block, put $S_j = \sum_{0 \neq y \in L^*} \|y\|_2^j e^{-2\pi^2 r^2 \|y\|_2^2}$ for $j = 0, 1, 2$. Differentiating the Poisson series for the Gaussian normalizer gives the uniform bound

$$\frac{\|\text{Cov}(D_{L,r,c})\|_2}{r^2} \leq 1 + 4\pi^2 r^2 \left(\frac{S_2}{1 - S_0} + \frac{S_1^2}{(1 - S_0)^2} \right) < C_{\text{cov}}.$$

For the scalar block, the dual squared lengths are j^2 ; for the paired block they are $4(u^2 - uv + v^2)/3$. Summing the scalar series through $|j| = 32$ and the pair series through $\max(|u|, |v|) = 16$, with geometric bounds on their tails, proves the displayed numerical inequality. Orthogonal sums preserve the covariance bound.

Set

$$\psi(c) = \log \sum_{x \in L} \exp(-\|x\|_2^2/(2r^2) + \langle x, c \rangle/r^2).$$

Its Hessian is $\text{Cov}(D_{L,r,c})/r^4 \preceq C_{\text{cov}}I/r^2$, and direct expansion gives

$$(a-1) \ln R_a(D_{L,r,c} \| D_{L,r,d}) = \psi(ac - (a-1)d) - a\psi(c) + (a-1)\psi(d).$$

Integrating the Hessian bound along $d + t(c-d)$ for $0 \leq t \leq a$ bounds the right-hand side by $a(a-1)C_{\text{cov}}\|c-d\|_2^2/(2r^2)$. Division by $(a-1) \ln 2$ proves the claim. $\square$

Write $z = Z/2^{53}$, $z_P = Z_P/2^{53}$, and $e_{\omega,P} = |\widehat{\omega}_P/2^{62} - \omega_P|$. Exact midpoint-square comparisons certify the prescribed RNA constants and $e_{\omega,P} \leq 2^{-63}$. Since $|z| < 320$, the common rounded-Gaussian coupling gives

$$|z_P - \omega_P G| \leq \omega_P 2^{-54} + 320e_{\omega,P} + \iota_P 2^{-54}, \quad \iota_P = \begin{cases} 0, & N_P = 1024, \\ 1, & N_P > 1024. \end{cases}$$

Here ι_P records the additional Q53 rounding in the first operation of Eq. (B.10). Combining this inequality with the profile- and width-specific preparation error from the Q96 enclosure and the final Q64 RNA gives the center-error bounds below. These estimates also show that both raw products fit signed I128 and that the normalized Q53 value fits signed I64.

The finite sampler satisfies $G_{2,P}$ by the support bounds

$$\|x_2\|_{\text{can}} < B_{2,P}, \quad (B_{2,P})_P = (114395, 170569, 160931, 208443).$$

The matched scalar displacement is at most 15; the paired squared displacement is at most $15^2 + \frac{3}{4}18^2 = 468 < 3 \cdot 19^2$. These fit the envelopes used for the displayed support bounds. Thus the concrete law has probability one on $G_{2,P}$. These bounds also establish its arithmetic range. For the intermediate exact samplers, the x_1 -center error calculation instead uses the full cap C_P , with coefficient envelope C_P in a power-of-two ring and $\sqrt{2}C_P$ in a trinomial ring.

Let $D_{2,P}$ bound the canonical norm of the computed x_2 -center D_2 under the bounded grid perturbations. Couple each grid value with a Gaussian in its rounding cell, and use the same x_2 in the finite and exact x_1 -center computations. Write $\varepsilon_{j,P}^{\text{ctr}}$ for an upper bound on the canonical distance between the computed center $\widehat{D}_j$, interpreted in real units, and its exact continuous counterpart, for $j = 1, 2$. The bound $\varepsilon_{1,P}^{\text{ctr}}$ is uniform over $\|x_2\|_{\text{can}} < C_P$.

These errors follow by propagating the cache, root, and RNA errors through Algorithm B.7. For example, a radix- t transform stage with half-spectrum length H , exact input norm A , input error e , complex root error u , and Qs states has output error at most

$$\left(\sqrt{t} + \sqrt{t(t-1)}u\right)e + \sqrt{t(t-1)}uA + (t-1)\sqrt{2H}2^{-s}.$$

The inverse normalization divides the accumulated error by H and adds at most $\sqrt{H/2}2^{-s}$; the trinomial unpacking uses its explicit linear map and directed rounding bounds. Combining these projection errors with the perturbation coupling above gives the outward x_2 -center norm bounds

$$(D_{2,P})_P = (113786.880001, 169656.320001, 160070.639541, 207325.430323)$$

in profile order, and the following center-error and lifetime bounds:

P	$\varepsilon_{2,P}^{\text{ctr}}$	$\varepsilon_{1,P}^{\text{ctr}}$	$L_{o,P}$ upper
IBS-1s	$5.658982773752 \cdot 10^{-10}$	$7.883346625600 \cdot 10^{-10}$	0.059019662
IBS-1	$1.440803860254 \cdot 10^{-9}$	$1.981438398374 \cdot 10^{-9}$	0.375234364
IBS-2	$1.201571545795 \cdot 10^{-9}$	$1.680893581869 \cdot 10^{-9}$	0.267078445
IBS-3	$2.710797633593 \cdot 10^{-9}$	$3.708391415762 \cdot 10^{-9}$	1.318592600

The displayed $L_{o,P}$ bounds exclude the separate conditioning allowance. For the lifetime outer-hop bound, put $\ell_{\text{crown}} = 3.303455901301445 \cdot 10^{-22}$ and $\ell_{\text{grid}} = 4.341452468542223 \cdot 10^{-36}$, outward bounds on the per-coordinate order-512 crown loss and pointwise grid-to-rounded-Gaussian loss, respectively; the latter includes finite-support normalization. An admissible outer-hop bound is

$$L_{o,P} = \frac{512C_{\text{cov}}Q_{\text{ID}}}{2r^2 \ln 2} ((\varepsilon_{1,P}^{\text{ctr}})^2 + (\varepsilon_{2,P}^{\text{ctr}})^2) + 2N_P Q_{\text{ID}}(\ell_{\text{crown}} + \ell_{\text{grid}}).$$

Lift the finite grid and its reference through the same conditional Gaussian within each rounding cell. The latent-variable moment is bounded by the crown loss and the pointwise grid comparison. Conditional on that latent variable, the center-displacement bound applies first to x_2 and then to x_1 for every preceding x_2 in the proof domain. Multiplication of these conditional moments, followed by adaptive composition over identities, yields the displayed bound. Conditioning on the proof domain and removing that conditioning are charged separately below.

To bound the reference tails, put $\ell_P = \sqrt{N_P}$ in a power-of-two ring and $\ell_P = \sqrt{3N_P/2}$ in a trinomial ring. Then $\|v\|_{\text{can}} \leq \ell_P \|v\|_{\infty}$. Uniformly over the center, let $\mathcal{T}_P(t)$ bound the probability that the discrete displacement has a coefficient of magnitude at least t . With $k = \lfloor t \rfloor - 1 \geq 300$, the scalar likelihood bound $L_s = 1 + 2^{-38}$ gives

$$\mathcal{T}_P(t) \leq \frac{2N_P L_s e^{-k^2/(2r_+^2)}}{(1 - e^{-(2k+1)/(2r_+^2)})e^{-1/(8r_-^2)}}$$

for a power-of-two profile, where one may take $r_- = r_+ = r$. For a trinomial profile, nearest-integer normalizer bounds give the conservative pair-mass envelope $2e^{-(u^2+uv+v^2)/(2r_{\text{tail}}^2)}$, with $r_{\text{tail}} = 1.3 > r$. Hence $u^2 + uv + v^2 \geq (u^2 + v^2)/2$ and a geometric tail sum give

$$\mathcal{T}_P(t) \leq \frac{4N_P(1 + 2r_{\text{tail}}\sqrt{\pi})e^{-k^2/(4r_{\text{tail}}^2)}}{1 - e^{-(2k+1)/(4r_{\text{tail}}^2)}}.$$

These bounds cover the exact scalar and pair intermediates; the finite-table laws have the bounded support established above. For bounded grid perturbations, the x_2 -center norm bound $D_{2,P}$ gives tail at most $\mathcal{T}_P((C_P - D_{2,P})/\ell_P)$.

The continuous endpoint requires a separate bound. Write $x_2 = \mu_2 - y_2 + e$, where $\widehat{\mu}_{2i} = \widehat{a}_i u_i$ and e is the exact discrete displacement. The nonnegative lift

and admitted cache imply

$$\|\mu_2\|_{\text{can}} \leq m_P := (1/2400 + 2^{-125})(q-1)\ell_P, \quad t_P := \frac{C_P - m_P}{2\ell_P} > 0.$$

The event $\|x_2\|_{\text{can}} \geq C_P$ implies that some coefficient of y_2 or e has magnitude at least t_P . The coefficient variances of y_2 are at most $\tau_{\max}^2$ in a power-of-two ring and $2\tau_{\max}^2$ in a trinomial ring, where

$$\tau_{\max} = r\sqrt{\alpha_{\text{ann}}^4 - 1} < 1.206.$$

Gaussian coefficient tails and $\mathcal{T}_P(t_P)$ therefore bound the continuous endpoint. Independence of e and y_2 is not required, since the displacement tail is uniform over its center.

Let δ_P be the sum of the bounded-grid and continuous per-extraction tail bounds. Directed evaluation gives $2^{48}\delta_P < 2^{-39000}$ in every profile. These bounds are uniform over the admitted master context and previous transcript. Conditioning each comparison on $G_{2,P}$ contributes at most $(1-\delta_P)^{-512/511}$ per extraction at order 512, and at most $(1-\delta_P)^{-1}$ for a pointwise comparison. Reserving four order-512 comparisons, four pointwise comparisons, and final deconditioning gives the common allowance

$$L_{\text{guard},P} = Q_{\text{ID}} \left(\frac{4 \cdot 512}{511} + 5 \right) \frac{\delta_P}{(1-\delta_P) \ln 2}.$$

This also covers the additional count comparison in Section B.6. The concrete law is unchanged by conditioning; deconditioning the reference uses $Q(E \mid G_{2,P}) \leq Q(E)/Q(G_{2,P})$.

Composition. Put $h_{512} = 511/512$. Applying Lemmas 4.1 and 4.2 successively to the three comparisons gives

$$\theta_P = h_{512}^3, \quad \beta_P = h_{512}L_{q,P} + h_{512}^2L_{e,P} + h_{512}^3L_{o,P} + L_{\text{guard},P},$$

and, for every event E ,

$$\Pr[\text{Law}_{\text{ext},P}^{\text{fx,raw}} \in E] \leq 2^{\beta_P} \Pr[\text{Law}_{\text{ext},P}^{\text{hyb}} \in E]^{\theta_P}. \quad (\text{B.16})$$

Extraction failure and complete bound. Restoring the witness gate and finite sampler caps yields the complete extraction bound.

Table B.2. Fixed-point extraction bounds for $Q_{\text{ID}} \leq 2^{48}$. Failure includes the witness gate and crown or discrete-sampler exhaustion; β_P is rounded upward.

profile	θ_P	β_P	per-extraction failure $p_{\text{ext},P}$	lifetime failure $\varepsilon_{\text{ext},P}$
IBS-1s	h_{512}^3	0.058676	$< 2^{-172.30}$	$< 2^{-124.71}$
IBS-1	h_{512}^3	0.373041	$< 2^{-224.79}$	$< 2^{-176.87}$
IBS-2	h_{512}^3	0.265520	$< 2^{-295.10}$	$< 2^{-247.29}$
IBS-3	h_{512}^3	1.310882	$< 2^{-331.74}$	$< 2^{-282.88}$

Lemma B.3 (Uniform extraction failure). Under the fresh private-tape model of this section, used in Eq. (B.16), the failure bounds in Table B.2 hold for every master context satisfying $\text{Valid}_P(B)$ and for adaptively selected syndromes, with at most $Q_{\text{ID}} \leq 2^{48}$ witness-generation states.

Proof. Fix B, h, a , take $S \leftarrow D_{\text{pub}}(\cdot \mid h, a)$, and let x be its normalized canonical image. Its law is a spherical discrete Gaussian on the corresponding affine lattice. Write $\|X^\ell S\|_2^2 = x^\top A_\ell x$. Appendix A.3 gives the spectrum of A_ℓ : 1 with multiplicity $2N_P$ in a power-of-two ring, and $2/3, 2$ each with multiplicity N_P in a trinomial ring. Put

$$\epsilon(r) = \left(\frac{e^{2\pi^2 r^2}}{4N_P} - 1 \right)^{-1}, \quad r_t = \frac{R_{\text{ext}}}{\sqrt{\max\{1, 1 - 2t\lambda_{\max}(A_\ell)\}}}.$$

Whenever $I - 2tA_\ell \succ 0$ and $0 < \epsilon(r_t) < 1$, Poisson summation and Lemma 2.1 give, uniformly in the affine translate,

$$\mathbb{E} \exp\left(\frac{t\|X^\ell S\|_2^2}{\sigma_1^2}\right) \leq \frac{1 + \epsilon(r_t)}{1 - \epsilon(R_{\text{ext}})} \det(I - 2tA_\ell)^{-1/2}.$$

The tilted Gaussian weight has covariance parameter $\sigma_1^2(I - 2tA_\ell)^{-1}$, whose least eigenvalue is at least $r_t^2 Q_{\text{hyb}}(B)^2$. Lemma 2.1 therefore bounds both Poisson normalization factors. Chernoff's inequality with positive and negative t yields the upper and lower gate tails, respectively.

The corresponding linear-moment bound, with coefficient variance proxy at most $2\sigma_1^2$, gives

$$\Pr[\mathcal{C}_q(S) \neq S] \leq 4N_P \frac{1 + \epsilon(R_{\text{ext}})}{1 - \epsilon(R_{\text{ext}})} e^{-q^2/(16\sigma_1^2)}.$$

There is one shift-norm event in a power-of-two profile and at most N_P in a trinomial profile. Combining their two tails with the centering bound gives gate-failure probability at most $g_P = 2^{-b_P}$ under $D_{\text{pub}}^{\text{ctr}}$, where, in profile order, $b_P = (173.998, 226.999, 297.999, 334.999)$. These conservative budgets use the quadratic-form upper-tail envelope underlying Lemma 4.12, with the affine-coset normalization above; optimizing the full moment bound gives stronger exponents. The bound is uniform in B, h, a , so the lifetime event has probability

at most $Q_{\text{ID}}g_P$ under adaptive syndrome selection. Lemma 4.3 at order 512 and Eq. (B.16) then give

$$\Pr[\text{any fixed-point extraction fails the gate}] \leq 2^{\beta_P} (r_p^{Q_{\text{ID}}} Q_{\text{ID}}g_P)^{h_{512}\theta_P}.$$

Here $r_p = r_p(512, \varepsilon_p)$ is the factor from Lemma 4.3. The conservative accounting in Table B.2 also includes the order-512 unit-label factor $r_u^\times$ from Lemma 4.4. Including this additional transfer and the crown- and discrete-cap union bounds gives

$$\begin{aligned} \Pr[\text{any fixed-point extraction fails}] &\leq 2^{\beta_P} \left(r_p^{Q_{\text{ID}}} ((r_u^\times)^{Q_{\text{ID}}})^{Q_{\text{ID}}g_P} \right)^{h_{512}\theta_P} \\ &\quad + \frac{Q_{\text{ID}}}{2^{48}} (2^{-938.5} + 2^{-1303.53}). \end{aligned}$$

For one witness, replace Q_{ID} by 1 and β_P by $\beta_P/2^{48}$. Table B.2 rounds the bounds outward to two decimals. $\square$

Theorem B.1 (Fixed-point extraction). Fix a master context satisfying $\text{Valid}_P(B)$, and let $\text{Law}_{\text{ext},P}^{\text{fx}}$ be the adaptive response law of Algorithm B.7, including $\perp$ on witness-gate failure or exhaustion of either sampler. For every event E ,

$$\Pr[\text{Law}_{\text{ext},P}^{\text{fx}} \in E] \leq 2^{\beta_P} \Pr[\text{Law}_{\text{ext},P}^{\text{hyb}} \in E]^{\theta_P} + \varepsilon_{\text{ext},P}. \quad (\text{B.17})$$

Proof. Couple $\text{Law}_{\text{ext},P}^{\text{fx}}$ and $\text{Law}_{\text{ext},P}^{\text{fx,raw}}$ through the same root and tapes for each extraction. Their transcripts coincide until the first gate failure or sampler exhaustion, whose lifetime probability is at most $\varepsilon_{\text{ext},P}$ by Lemma B.3. Applying Eq. (B.16) to the uncapped ungated law proves the claim. $\square$

By Rényi data processing, Lemma 4.3 also compares $\text{Law}_{\text{ext},P}^{\text{hyb}}$ with $D_{\text{pub}}^{\text{ctr}}$. Theorem 4.2 charges the subsequent witness-first transition to $\varepsilon_{\text{wrap}}^{\text{wf}}$. Witness-gate failure and exhaustion of either sampler contribute $\varepsilon_{\text{ext},P}$, separately from master-key validity.

B.5 Fixed-Point Signing

Fix a profile P and a public mode $\chi \in \{\text{mixed}, \text{Q128}\}$, selected independently of keys, messages, challenges, and secret state and fixed for the implementation's lifetime. Table B.3 defines the transport precision $p_{P,\chi}$.

In fixed-point signing, Gaussian mask sampling uses integer alias tables and Q128 exponential evaluations based on precomputed factors or range-reduced polynomial evaluation followed by repeated squaring. Transport probabilities are obtained from truncated alternating series evaluated in fixed-point arithmetic; directed rounding and truncation bounds provide enclosing probability intervals. Integer thresholds determine acceptance and signed updates, with unresolved transport comparisons causing rejection. Transport uses Q128 for every profile in Q128 mode, and Q96 for IBS-1s/1/2 with Q128 for IBS-3 in mixed mode.

Let $\text{Sign}_P^{\text{exact}}$ be Algorithm 3.3 with exact discrete-Gaussian masks, exact transports, and unbounded first-success semantics. The fixed-point signer uses

the profile-dependent attempt cap below. We compare their complete adaptive response laws under the ideal XOF $\mathcal{O}$: distinct input transcripts receive independent uniform tapes, whereas a repeated transcript reuses its tape. Section B.7 handles the joint call families and restores the computational transition.

Admitted signing contexts. Algorithm 3.2 writes the abstract identity key as (j, a, s_0, s_1) , but its canonical encoding omits the derivable label a . Signing-key preparation decodes (j, s_0, s_1) , computes $(j', a) = H_{\text{id}}(h, \text{id})$, requires $j' = j$, and verifies

$$s_0 + h s_1 = a \pmod{q}.$$

It then sets

$$A = (\hat{q}a^{-1}, \hat{q}ha^{-1}), \quad S = (s_0, s_1)^T.$$

For each $0 \leq \ell < N$, compute $v_\ell = X^\ell S$ and $n_\ell = \|v_\ell\|_2^2$, and check Eq. (B.14) before any message-dependent attempt. The lower endpoints are the chosen squares $150618^2, 209999^2, 217849^2, 259108^2$, which ensure that the finite-series schedules cover the admitted interval. The upper endpoints satisfy $n_{\max, P} < B_{\text{shift}, P}^2$ and $\lceil \sqrt{n_{\max, P}} \rceil = B_{\text{shift}, P}$, preserving the bounds used for M_{iter} and the abort guard.

Table B.3. Fixed-point signing parameters.

profile	n_{word}	$p_{P, \text{mixed}}$	$p_{P, \text{Q128}}$	T_P^{ref}	k_P	$[n_{\min, P}, n_{\max, P}]$
IBS-1s	2	96	128	32	2^{12}	[22685781924, 62677966353]
IBS-1	2	96	128	44	2^{15}	[44099580001, 131253033394]
IBS-2	2	96	128	48	2^{15}	[47458186801, 121200610493]
IBS-3	2	128	128	28	2^{14}	[67136955664, 195463297549]

Both precisions reserve $n_{\text{word}} = 2$ 64-bit words for the fractional part, with a separate bounded integral part. Q128 uses four little-endian base- 2^{32} limbs; Q96 uses three and zeros the fourth. The 32- and 64-bit kernels sum identical exact limb products before truncating at bit $p_{P, \chi}$.

Let $b_\perp = 20$. The fixed signer stores the exact rational $M_{\text{iter}, P} = M_P^{\text{num}}/M_P^{\text{den}}$ from Section 5.1 and an outward enclosure of its reciprocal. Directed real-interval arithmetic proves

$$(1 - 2^{-b_\perp}) \frac{M_P^{\text{num}}}{M_P^{\text{den}}} \geq M_{\sigma_y/B_{\text{shift}, P}}$$

for every profile.

Let $\rho_{\text{na}, P}$ be the certified lower bound on $\rho_{\text{pub}}/M_{\text{iter}, P}^{\kappa_P}$ in Table B.6. For $Q_S = 2^{64}$ signing queries, choose

$$T_{\text{sign}, P} = \left\lceil \frac{258 + \log_2 Q_S}{-\log_2(1 - \rho_{\text{na}, P})} \right\rceil.$$

This choice gives $Q_S(1 - \rho_{\text{na}, P})^{T_{\text{sign}, P}} \leq 2^{-258}$; the endpoint exhaustion analysis below accounts for finite sampling. The caps and additive errors use the certified

$\rho_{\text{na},P}$ bounds before tabular rounding. The certified expected-attempt upper bounds are 1.522898, 1.150179, 1.156671, and 2.380216 in profile order; Table B.6 lists the resulting caps.

Integer mask sampling. Put $\sigma_{\text{base},P} = \sigma_y/k_P$ and define

$$p_{P,b} = \frac{\exp(-b^2/(2\sigma_{\text{base},P}^2))}{\sum_{u \geq 0} \exp(-u^2/(2\sigma_{\text{base},P}^2))} \quad (b \geq 0).$$

The nominal cutoff is $J_P = \lfloor (Z_{\text{BOUND}} + \kappa_P \lceil \sqrt{n_{\text{max},P}} \rceil) / k_P \rfloor$, the least value with $(J_P + 1)k_P - 1 \geq Z_{\text{BOUND}} + \kappa_P \lceil \sqrt{n_{\text{max},P}} \rceil$. The dyadic base law has counts

$$m_{P,b} = \lfloor 2^{120} p_{P,b} \rfloor \quad (1 \leq b \leq J_P), \quad m_{P,0} = 2^{120} - \sum_{b=1}^{J_P} m_{P,b}.$$

Thus an integer alias table samples b with probability $m_{P,b}/2^{120}$. Given an independent uniform residue $r \in \{0, \dots, k_P - 1\}$, set $x = bk_P + r$ and put

$$e_{b,r} = \exp\left(-\frac{r(r + 2bk_P)}{2\sigma_y^2}\right). \quad (\text{B.18})$$

For $x \neq 0$, the mask exponential evaluator returns an outward interval $\mathbf{E}_{b,r}$ on the independent Q128 grid containing $e_{b,r}$. The IBS-1s/1/2 profiles multiply enclosed window factors; IBS-3 instead divides the exponent by 2^{22} , applies directed degree-22 Horner evaluation, and performs 22 directed squarings. The respective interval widths are at most $8192 \cdot 2^{-128}$ and 2^{-99} . A uniform 104-bit integer is accepted precisely when it is smaller than

$$t_{b,r} = \left\lfloor 2^{104} \mathbf{E}_{b,r}^- \right\rfloor.$$

The case $x = 0$ is accepted directly. A prefix test may settle a strict comparison early but invokes the full evaluator when necessary, leaving the accepted set unchanged. After acceptance, an independent sign is applied; a proposal representing negative zero is rejected. With the unquantized base law and exact probability $e_{b,r}$, this construction samples $D_{\mathbb{Z}, \sigma_y}$ exactly. The analysis below accounts for the dyadic base law, finite support, fixed-point enclosure, and 104-bit floor.

In the profile order IBS-1s/1/2/3, the nominal cutoffs, largest bases of nonzero dyadic mass, and resulting magnitude bounds are

$$\begin{aligned} (J_P) &= (1512, 360, 390, 1202), \\ (b_P^*) &= (857, 183, 183, 390), \\ (B_{y,P}) &= (3514367, 6029311, 6029311, 6406143). \end{aligned}$$

Every base $b > b_P^*$ has $m_{P,b} = 0$. Since $b_P^* < J_P$, the cutoff does not affect the sampled law. In the same profile order, set $(b_{\text{col},P})_P = (10, 8, 8, 9)$; then $2^{b_{\text{col},P}} \geq b_P^* + 1$, so the alias table covers the complete dyadic support.

Each lane processes (9, 14, 19, 20) blocks of 128 proposals and returns its first N accepted coefficients. A proposal stores the $b_{\text{col},P}$ -bit alias column, residue, sign, and 16-bit alias and Bernoulli prefixes. Prefix equality consumes the next $104 - b_{\text{col},P}$ -bit alias suffix or 88-bit Bernoulli suffix from independent block pools of (16, 18, 24, 29) slots each; all multibyte fields are little-endian. Strict prefix comparisons consume no suffix. This deferred construction reveals independent 120- and 104-bit uniforms in proposal order until pool overflow, which, like insufficient accepted proposals, rejects the attempt.

Let $R_{\text{pair},P}$ be the certified pointwise likelihood factor from one concrete mask-pair law to the corresponding ideal Gaussian mask-pair law. With exact transport in both games, conditioning on a successful signing call and composing at most Q_S adaptive calls gives the lifetime factor $R_{\text{mask},P}$. The fixed-batch, suffix-pool, and attempt-exhaustion events are charged additively rather than included in these factors.

Integer transport. For a selected shift v , write

$$n = \|v\|_2^2, \quad \delta = \langle z, v \rangle, \quad D = 2\sigma_y^2.$$

All transport probabilities are represented by outward dyadic intervals on the $2^{-P_{P,\chi}}$ grid. For $\delta \geq -n$, define

$$\text{Alt}_n(\delta) = \sum_{k \geq 0} (-1)^k \exp\left(-\frac{k^2 n + 2k\delta}{D}\right).$$

The reference recurrence is

$$r_0 = 1, \quad q_0 = e^{-(n+2\delta)/D}, \quad s = e^{-2n/D}, \quad r_{k+1} = r_k q_k, \quad q_{k+1} = q_k s. \quad (\text{B.19})$$

The recurrence sums $r_0, \dots, r_{T_P^{\text{ref}}-1}$ with alternating signs; replacing δ by $-\delta$ gives $\text{Alt}_n(-\delta)$. With $J_{\text{tr}} = T_P^{\text{ref}} - 1$ and $a = n/(2\sigma_y^2)$, Lemma 4.8 and Table B.3 give, for either realization,

$$\varepsilon_{\text{series}} \leq \frac{e^{-J_{\text{tr}}^2 a}}{e^{-a} - e^{-4a}} < 2^{-P_{P,\chi}}. \quad (\text{B.20})$$

Cached-series Horner evaluation. The hardened evaluator uses the same known-sign recurrence in every profile. Put $p = p_{P,\chi}$, let $m = T_{P,\chi}^{\text{ct}}$ be the series length in Table B.4, and define directed rounding to the fixed grid by

$$\text{rdn}_p(x) = 2^{-p} \lfloor 2^p x \rfloor, \quad \text{rup}_p(x) = 2^{-p} \lceil 2^p x \rceil.$$

For an admitted squared norm n , let $\underline{s}$ be the stored lower endpoint of $e^{-2n/D}$ and construct the nonnegative cache

$$P_0 = C_0 = 1, \quad P_{j+1} = \text{rdn}_p(P_j \underline{s}), \quad C_{j+1} = \text{rdn}_p(C_j P_j).$$

Thus $P_j \leq e^{-2jn/D}$ and $C_j \leq e^{-j(j-1)n/D}$. For $d \in \{\delta, -\delta\}$ in a branch where $d \geq -n$, let R denote the stored lower endpoint of $e^{-(n+2d)/D}$. The directed Horner recurrence is

$$H_{m-1} = (-1)^{m-1} C_{m-1}, \quad H_{k-1} = (-1)^{k-1} C_{k-1} + \text{rdn}_p(R H_k),$$

for $k = m - 1, \dots, 1$.

For $k \geq 2$, $d \geq -n$ implies $R \leq e^{n/D}$ and

$$P_{k-1}R \leq e^{-(2k-3)n/D} \leq e^{-n_{\min,P}/D} < 1.$$

The certified upper bound is 0.907. Backward induction gives $H_k \in (-1)^k[0, C_k]$ for $k \geq 1$: the initial value has this form, and when $C_{k-1} > 0$,

$$R|H_k| \leq RC_k \leq RC_{k-1}P_{k-1} < C_{k-1} \implies \text{rup}_p(R|H_k|) \leq C_{k-1}.$$

If $C_{k-1} = 0$, all subsequent C_j and H_j vanish. Using $\text{rdn}_p(-x) = -\text{rup}_p(x)$, each step therefore subtracts a magnitude at most C_{k-1} from C_{k-1} with sign $(-1)^{k-1}$. This evaluates the signed recurrence exactly. The $k = 1$ step uses general signed addition because H_0 has no prescribed sign. All product magnitudes are at most 1 for $k \geq 2$ and at most $R \leq e^{n_{\max,P}/D} < 1.50 < 2$ for $k = 1$, within the multiplier's specified $[0, 2)$ domain. At a fixed precision and term count, the signed recurrence and its magnitude-based implementation produce identical rounded values, preserving transport decisions, tape consumption, and retries.

The ordinary evaluator uses Table B.3's term counts, and the hardened evaluator uses Table B.4's parameters. Both enclose the same exact probabilities within the common error bound used below. Their transcripts can differ because different enclosures leave different comparisons unresolved.

For every profile, the computed partial sums are widened by the series radius in Table B.4. Let $\mathbf{Alt}_n(\delta)$ and $\mathbf{Alt}_n(-\delta)$ denote the resulting enclosures and $\mathbf{M}_P^{-1}$ an outward enclosure of $1/M_{\text{iter}}$. The minus-branch interval is

$$\mathbf{F} = \begin{cases} \mathbf{Alt}_n(\delta)\mathbf{M}_P^{-1}, & \delta \geq n, \\ (1 - \mathbf{Alt}_n(-\delta))\mathbf{M}_P^{-1}, & \delta < n. \end{cases}$$

The cumulative success interval is

$$\mathbf{C} = \begin{cases} \mathbf{M}_P^{-1}, & \delta \geq n \text{ or } \delta < -n, \\ (2 - \mathbf{Alt}_n(\delta) - \mathbf{Alt}_n(-\delta))\mathbf{M}_P^{-1}, & -n \leq \delta < n. \end{cases}$$

After intersection with $[0, 1]$, read a $p_{P,\chi}/8$ -byte string buf from the transport stream and set $U = 2^{-p_{P,\chi}} \text{LE}_{p_{P,\chi}}(buf)$, where LE_p denotes little-endian decoding. Return $z - v$ when $U < \mathbf{F}^-$, return $z + v$ when $U \geq \mathbf{F}^+$ and $U < \mathbf{C}^-$, and reject when $U \geq \mathbf{C}^+$. If U lies in either nondegenerate enclosure, reject as well. Thus every successful fixed-point decision agrees with the exact branch.

For **IBS-1s**/1/2 and $x \in [0, 160]$, the negative-direction exponential graph evaluates the degree- $d_{P,\chi}$ Taylor polynomial at $x2^{-s_{P,\chi}}$ with coefficients $\lfloor 2^{p_{P,\chi}}/k! \rfloor$, followed by $s_{P,\chi}$ squarings. Horner products and squarings truncate magnitudes. Table B.4 gives the schedule and point radius $r_{\text{pt},P,\chi}$, measured in ulps (units in the last place) of size $2^{-p_{P,\chi}}$. Directed propagation over 8192 outward cells covers the full domain, including truncation, rounding, and squaring amplification. For $x \geq 160$, return $[0, 2^{-p_{P,\chi}} B_{\text{exp},P,\chi}]$, where

$$B_{\text{exp},P,\chi} = \lceil 2^{p_{P,\chi}} e^{-160} \rceil + r_{\text{pt},P,\chi}.$$

For IBS-1s/1/2 these bounds are 4097, 262145, and 262145 ulps, respectively, in either realization.

For IBS-3, both exponentials use independent directed Q128 degree-22 Horner graphs after division by 2^{22} , followed by 22 directed squarings, with a paired multiplication schedule. Each series radius covers Eq. (B.20) at the specified term count and the recurrence error.

Table B.4. Transport approximation parameters. Radii and error bounds are in units of $2^{-p_{P,\chi}}$.

profile	realization	$p_{P,\chi}$	evaluator	$(s_{P,\chi}, d_{P,\chi})$	point radius / required	series terms / radius
IBS-1s	mixed	96	Taylor	(8, 11)	$2^{12}/2264$	$21/2^{21}$
IBS-1s	Q128	128	Taylor	(8, 15)	$2^{12}/2264$	$24/2^{23}$
IBS-1	mixed	96	Taylor	(8, 10)	$2^{18}/8955$	$25/2^{23}$
IBS-1	Q128	128	Taylor	(14, 8)	$2^{18}/117238$	$29/2^{24}$
IBS-2	mixed	96	Taylor	(8, 10)	$2^{18}/8955$	$24/2^{23}$
IBS-2	Q128	128	Taylor	(14, 8)	$2^{18}/112131$	$28/2^{24}$
IBS-3	mixed	128	direct Horner	(22, 22)	$2^{28}/45606177$	$28/2^{40}$
IBS-3	Q128	128	direct Horner	(22, 22)	$2^{28}/45606177$	$28/2^{40}$

For the ordinary table-based evaluator in Q96, the exponential enclosures have width at most 8192 ulps; the three cached-series errors, including tails, are below 517184, 7171306, and 6157982 ulps, respectively. The cached-series errors are below their chosen radii, and the point error bounds for the hardened evaluators in Table B.4 are below the corresponding chosen point radii. Together with the carrier and padding invariants, the directed recurrences prove enclosure of every exact value. The same argument applies at Q128.

Arithmetic ranges. The mask supports and admitted shift norms bound every carrier. Writing $z = y + \sum_{j < t} \epsilon_j X^{\ell_j} S$ with at most κ_P signed admitted shifts gives

$$|\delta| \leq B_{y,P} \sqrt{2N_P n_{\max,P}} + \kappa_P n_{\max,P} < 2^{49}.$$

Each transport step adds or subtracts a shift v_t with $\|v_t\|_\infty \leq \lceil \sqrt{n_{\max,P}} \rceil$, so intermediate response coefficients are at most $B_{y,P} + \kappa_P \lceil \sqrt{n_{\max,P}} \rceil$ in absolute value, namely 8521487, 15811114, 16821620, and 24090663 for IBS-1s/1/2/3. Across all profiles they are therefore below 2^{25} , so signed 32- and 64-bit carriers suffice for the response state and inner product; the wider products and fixed-limb accumulators remain within their stated carriers.

Complete fixed signer. Each signing call samples an independent uniform 2λ -bit seed and initializes

$$\begin{aligned} \mathcal{S}_T &= \text{SHAKE256}('S' \parallel \text{seed} \parallel \text{octet}(0)), \\ \mathcal{S}_{y_i} &= \text{SHAKE256}('S' \parallel \text{seed} \parallel \text{octet}(i+1)), \quad i \in \{0, 1\}. \end{aligned}$$

All attempts continue these streams without rewinding or deriving another seed.

Algorithm B.8 ($\text{Sign}_{P,\chi}^{\text{fx}}$)

Input: χ , an admitted context $(j, A, S, (n_\ell)_{\ell < N})$, a digest μ , and a fresh seed.
 Output: a signature (j, c, z_0, z_1) or $\perp$.

1. Initialize $(\mathcal{S}_T, \mathcal{S}_{y_0}, \mathcal{S}_{y_1})$ as above.
2. For at most $T_{\text{sign}, P}$ attempts:
 1. Draw (y_0, y_1) with the fixed-batch radix sampler from $\mathcal{S}_{y_0}$ and $\mathcal{S}_{y_1}$. If either lane rejects, continue to the next attempt.
 2. Compute $u = A_0 y_0 + A_1 y_1 \bmod q$, $w = [u]_{d_h}$, and $c = H_{\text{ch}}(\mu, w) = \sum_{t=1}^{\kappa} X^{\ell_t}$ as in Algorithm 3.3.
 3. Set $z = (z_0, z_1) = (y_0, y_1)$. For $t = 1, \dots, \kappa$, apply the fixed transport for mode χ with $v = X^{\ell_t} S$, the certified norm n_{ℓ_t} , and the next $p_{P, \chi}$ bits of $\mathcal{S}_T$. Abandon the attempt if any transport rejects.
 4. Continue to the next attempt unless $\|z_0\|_\infty, \|z_1\|_\infty \leq Z_{\text{BOUND}}$.
 5. Set $\tau = A_0 z_0 + A_1 z_1 + c\hat{q} \bmod q$ and continue unless $[\tau]_{d_h} = [\tau - c]_{d_h}$.
 6. Return (j, c, z_0, z_1) in its canonical fixed-length encoding.
3. Return $\perp$.

Lemma B.4 (Signing correctness). For either realization mode χ , every signature returned by $\text{Sign}_{P, \chi}^{\text{fx}}$ is accepted by Algorithm 3.4.

Proof. Let $\epsilon_t \in \{-1, 1\}$ be the branch selected for $X^{\ell_t} S$ and set

$$c' = \sum_{t=1}^{\kappa} \epsilon_t X^{\ell_t}, \quad d = \frac{c + c'}{2}.$$

Then $z = y + c' S$. Since $AS = \hat{q}$ and $2\hat{q} = 1 \bmod q$,

$$Az + c\hat{q} = u + (c' + c)\hat{q} = u + d \pmod{q}.$$

The filter in Step 5 is precisely $[u + d]_{d_h} = [u + d - c]_{d_h}$, so Lemma 4.11 gives $[u + d]_{d_h} = [u]_{d_h}$. Verification therefore reconstructs the original commitment and recomputes c . The label, response, and encoding checks then complete correctness. $\square$

Both modes share the public parameters, signature format, and verifier; χ is not encoded. Their seeded transcripts may differ: transport consumes $p_{P, \chi}/8$ bytes per comparison, and the evaluator schedule depends on the mode.

Distributional comparison. Let $\text{Law}_{\text{sign}, P, \chi}^{\text{fx}}$ and $\text{Law}_{\text{sign}, P}^{\text{ex}}$ be the adaptive oracle-response laws of Algorithm B.8 in mode χ and the unbounded $\text{Sign}_P^{\text{exact}}$, respectively.

Signing-seed reuse. Let $\text{Bad}_{\text{seed}, P}$ be the event that a call repeats a seed used by an earlier signing call, across all identities in the fixed profile. For at most Q_S calls,

$$\Pr[\text{Bad}_{\text{seed}, P}] \leq \varepsilon_{\text{seed}, P} := \binom{Q_S}{2} 2^{-2\lambda}.$$

For $Q_S \leq 2^{64}$, these bounds are strictly below 2^{-129} , 2^{-129} , 2^{-257} , and 2^{-385} in profile order. The comparison game stops before a repeated seed initializes a

sampler stream. Independent seed labels are added to the exact reference games solely to define this stopping rule; without the rule, their oracle-response laws are unchanged. Until the stop, calls use independent tapes, and attempts within a call consume unread suffixes without rewinding.

Oracle freshness. The projection-tail bound below must account for adaptive random-oracle histories. Put

$$Q_{\text{att},P} = Q_S T_{\text{sign},P}, \quad \varepsilon_{\text{raw},P} = ((2^{d_h+1} + 2)p_{q,\sigma_y})^{N_P}.$$

Work in the capped seed-stopped games used for the finite-precision comparison. Let $\text{Bad}_{\text{RO},P}$ occur when the random-oracle table already contains (μ, w) as an attempt reaches $H_{\text{ch}}(\mu, w)$. At each live prefix immediately before mask sampling, the unread sampler suffixes are independent of the preceding transcript. The table contains at most $Q_H + Q_{\text{att},P}$ relevant points. Lemma 4.15's raw-commitment bound, transferred to the fixed mask law by $R_{\text{pair},P}$, therefore gives

$$\Pr_{\text{seed-stop}} [\text{Bad}_{\text{RO},P}] \leq R_{\text{pair},P} Q_{\text{att},P} (Q_H + Q_{\text{att},P}) \varepsilon_{\text{raw},P} =: \varepsilon_{\text{RO},P}. \quad (\text{B.21})$$

For $Q_H \leq 2^{64}$, this oracle-collision probability in the seed-stopped game is below 2^{-5400} in all four profiles. We analyze the projection tail after additionally stopping on $\text{Bad}_{\text{RO},P}$. Coupling with the original capped game costs at most $\varepsilon_{\text{seed},P} + \varepsilon_{\text{RO},P}$. Each stopped law includes an absorbing proof-only flag. Events of the original response transcript exclude flagged executions, and the surviving subprobabilities are not renormalized. At an unflagged fresh query, the oracle answer is uniform conditional on the preceding history and the current mask.

Transport divergence and projection tails. In profile order, the transport target exponents are

$$(b_{\text{tr},P})_P = (91, 128, 192, 256).$$

The Rényi orders and lifetime budgets are

$$a_P = 2b_{\text{tr},P} - 1, \quad \vartheta_{\text{tr},P} = \frac{a_P - 1}{a_P},$$

$$Q_{\text{step},P} = Q_S \kappa_P T_{\text{sign},P}, \quad Q_{\text{win},P} = N_P Q_{\text{step},P}.$$

For a shift v and current accumulator z , write

$$\gamma = \frac{\|v\|_2^2}{2\sigma_y^2}, \quad u = \frac{\langle z, v \rangle}{\sigma_y^2}, \quad \gamma_{\min,P} = \frac{n_{\min,P}}{2\sigma_y^2}, \quad \gamma_{\max,P} = \frac{n_{\max,P}}{2\sigma_y^2}.$$

The projection window is

$$U_P = \sqrt{4\gamma_{\max,P} \ln 2 (\log_2 Q_{\text{win},P} + b_{\text{tr},P} + 8)}. \quad (\text{B.22})$$

For the exact three-output kernel $K_{v,z}$, the public guard gives $K_{v,z}(\perp) \geq 2^{-20}$. On $|u| \leq U_P$, the alternating-series bounds also give, for each retained output,

$$K_{v,z}(o) \geq p_{\min,P} := \frac{e^{-U_P - \gamma_{\max,P}} (1 - e^{-5\gamma_{\min,P}})}{M_{\text{iter},P}}, \quad o \in \{-v, +v\}. \quad (\text{B.23})$$

Here $U_P \geq 2 + 2\gamma_{\max, P}$, so the displayed outer-region bound also lower-bounds the two middle-region probabilities. For every profile, $p_{\min, P} < 2^{-20}$, so it lower-bounds all three output probabilities.

Let $\hat{K}_{v,z}^\chi$ be the specified kernel for realization χ , including rejection on an unresolved comparison. Directed interval propagation bounds the absolute error of every one of its three output probabilities by $e_{P,\chi}$, where the abort coordinate includes the errors of both retained thresholds. This bound includes the exponential, recurrence, reciprocal, interval-propagation, and $p_{P,\chi}$ -bit threshold-grid errors. Thus

$$\left| \frac{\hat{K}_{v,z}^\chi(o)}{K_{v,z}(o)} - 1 \right| \leq \eta_{P,\chi} := \frac{e_{P,\chi}}{p_{\min, P}} \quad (|u| \leq U_P).$$

For $\Delta_o = \hat{K}_{v,z}^\chi(o)/K_{v,z}(o) - 1$, normalization gives $\sum_o K_{v,z}(o)\Delta_o = 0$. The certified values satisfy $\eta_{P,\chi} < 1$. Writing $D_a = \ln R_a$, Taylor's theorem therefore gives

$$D_{a_P}(\hat{K}_{v,z}^\chi \| K_{v,z}) \leq \frac{a_P}{2} (1 + \eta_{P,\chi})^{a_P-2} \eta_{P,\chi}^2. \quad (\text{B.24})$$

At Rényi orders (181, 255, 383, 511) in profile order, the certified upper bounds on U_P are (14.320144, 13.383935, 14.613256, 19.272207), and the corresponding lower bounds on $\log_2 p_{\min, P}$ are

$$(-22.218878, -21.102506, -22.761294, -29.432011).$$

Let $D_{\text{life}, P, \chi}$ denote the right side of Eq. (B.24) multiplied by $Q_{\text{step}, P}$.

To account for the unbounded projection, define a hybrid that uses the specified fixed-point kernel inside the window and the exact kernel outside it. The hybrid and the implementation coincide until an evaluation leaves the window.

Fix the history immediately before an attempt with the ideal mask and exact transports. Let $\nu_t(z, c)$ be the joint subprobability mass of passing the freshness test and completing t transports with accumulator z and challenge c . Since a fresh oracle answer is uniform, $\nu_0(z, c) \leq p_{\sigma_y}(z)/|\mathcal{C}_\kappa|$. Each exact transport maps p_{σ_y} to $p_{\sigma_y}/M_{\text{iter}}$ by Eq. (6); its nonnegative transition kernel preserves pointwise domination. Induction gives

$$\nu_t(z, c) \leq \frac{M_{\text{iter}}^{-t}}{|\mathcal{C}_\kappa|} p_{\sigma_y}(z), \quad 0 \leq t \leq \kappa.$$

For any fixed admissible shift v , completing the square and maximizing Gaussian mass at the centered lattice give

$$\Pr_{Z \leftarrow D_{\sigma_y}^{2N_P}} \left[\frac{|\langle Z, v \rangle|}{\sigma_y^2} > U_P \right] \leq 2e^{-U_P^2/(4\gamma)}.$$

Summing over c bounds the unnormalized mass reaching a step with an out-of-window projection by the same tail. The pointwise factor $R_{\text{pair}, P}$ transfers this event from the ideal mask pair to the concrete radix mask pair. A union bound

over all N_P admissible monomial shifts at each of the $Q_{\text{step},P}$ potential transport steps gives

$$\tau_P = R_{\text{pair},P} Q_{\text{win},P} 2e^{-U_P^2/(4\gamma_{\max,P})}.$$

Sequential composition and Lemma 4.1 transfer the out-of-window event to the fixed-point hybrid with probability at most

$$2^{\vartheta_{\text{tr},P} D_{\text{life},P,\chi} / \ln 2} \tau_P^{\vartheta_{\text{tr},P}}.$$

Table B.5 summarizes the kernel divergence and the resulting projection-window term over the complete signing lifetime. All displayed decimal bounds in Tables B.5 and B.6 are rounded outward.

Table B.5. Certified transport-kernel and projection-window bounds for $Q_S \leq 2^{64}$.

profile	realization	$p_{P,\chi}$	upper	upper	transferred window event
			$\log_2 \eta_{P,\chi}$	$\log_2 D_{\text{life},P,\chi}$	
IBS-1s	mixed	96	-48.781119	-15.033105	$< 2^{-97.458520}$
IBS-1s	Q128	128	-78.781121	-75.033110	$< 2^{-97.458563}$
IBS-1	mixed	96	-47.897493	-13.264385	$< 2^{-134.470442}$
IBS-1	Q128	128	-78.897493	-75.264386	$< 2^{-134.470588}$
IBS-2	mixed	96	-46.238705	-9.134658	$< 2^{-198.477857}$
IBS-2	Q128	128	-77.238705	-71.134659	$< 2^{-198.480417}$
IBS-3	mixed	128	-54.567988	-23.137389	$< 2^{-262.485322}$
IBS-3	Q128	128	-54.567988	-23.137389	$< 2^{-262.485322}$

Attempt exhaustion. At either endpoint of each mask or transport distributional comparison, the game’s lifetime exhaustion probability is below 2^{-258} . The two endpoints of each mask and transport comparison contribute four terms totaling less than 2^{-256} .

Complete response bound. Let $\varepsilon_{\text{samp},P,\chi}$ collect the certified lifetime contributions from fixed-batch failure, suffix-pool overflow, the transferred projection-window event, and the endpoint attempt-exhaustion terms. The signing additive bound is chosen to satisfy

$$\varepsilon_{\text{seed},P} + \varepsilon_{\text{samp},P,\chi} \leq \varepsilon_{\text{sign},P,\chi}.$$

Table B.6 includes seed collisions in $\varepsilon_{\text{sign},P,\chi}$; $\varepsilon_{\text{RO},P}$ is charged separately.

For every profile, set

$$\vartheta_{\text{sign},P} = \vartheta_{\text{tr},P}.$$

Theorem B.2 (Fixed-point signing). Fix either public static realization mode χ . In the ideal-XOF hybrid, every adaptive event E involving $Q_S \leq 2^{64}$ signing calls on admitted statements satisfies

$$\Pr[\text{Law}_{\text{sign},P,\chi}^{\text{fx}} \in E] \leq 2^{\beta_{\text{sign},P,\chi}} \Pr[\text{Law}_{\text{sign},P}^{\text{ex}} \in E]^{\vartheta_{\text{sign},P}} + \varepsilon_{\text{sign},P,\chi} + \varepsilon_{\text{RO},P}. \quad (\text{B.25})$$

Table B.6. Signing-lifetime bounds for $Q_S \leq 2^{64}$.

profile	realization	lower $\rho_{\text{na},P}$	$T_{\text{sign},P}$	$\vartheta_{\text{sign},P}$	factor excess $2^{\beta_{\text{sign},P,\chi}} - 1$	additive error $\varepsilon_{\text{sign},P,\chi}$
IBS-1s	mixed	0.6566	209	180/181	$< 2^{-15.040703}$	$< 2^{-97.458520}$
IBS-1s	Q128	0.6566	209	180/181	$< 2^{-26.954314}$	$< 2^{-97.458563}$
IBS-1	mixed	0.8694	110	254/255	$< 2^{-13.269851}$	$< 2^{-128.967822}$
IBS-1	Q128	0.8694	110	254/255	$< 2^{-26.710371}$	$< 2^{-128.967825}$
IBS-2	mixed	0.8645	112	382/383	$< 2^{-9.137140}$	$< 2^{-198.477857}$
IBS-2	Q128	0.8645	112	382/383	$< 2^{-26.290490}$	$< 2^{-198.480417}$
IBS-3	mixed	0.4201	410	510/511	$< 2^{-20.403308}$	$< 2^{-256.321262}$
IBS-3	Q128	0.4201	410	510/511	$< 2^{-20.403308}$	$< 2^{-256.321262}$

Proof. Couple each proposal to independent 120-bit alias and 104-bit Bernoulli uniforms. Unless a suffix pool overflows, deferred comparison reveals exactly those uniforms in proposal order. Replacing the dyadic base law, finite support, and lower Bernoulli threshold by the ideal quantities in Eq. (B.18) gives the certified one-sided likelihood factor; pool overflow and fixed-batch exhaustion are additive events.

Every transport enclosure contains the exact branch probabilities. Under a common uniform threshold, a successful fixed-point branch is therefore the exact branch, while an unresolved cell contributes only rejection mass. These losses compose through the κ adaptive transports because the admitted norm interval is uniform over every message and challenge. Deterministic challenge generation, response testing, and public filtering cannot increase these losses.

For the fixed pair (P, χ) , work first in the game that stops on $\text{Bad}_{\text{seed},P}$ or $\text{Bad}_{\text{RO},P}$ and apply the exact-outside-window hybrid above. The local kernel bound in Eq. (B.24) holds uniformly over every live history, so Lemma 4.2 bounds the complete transport divergence. Lemma 4.1 contributes the exponent $\vartheta_{\text{sign},P}$ and the transport part of $\beta_{\text{sign},P,\chi}$. Applying the radix-mask one-sided likelihood bound inside that event transfer contributes $\vartheta_{\text{sign},P} \log_2 R_{\text{mask},P}$ to $\beta_{\text{sign},P,\chi}$. The transferred probability of entering the out-of-window branch is included in $\varepsilon_{\text{sign},P,\chi}$. Consequently,

$$\beta_{\text{sign},P,\chi} = \vartheta_{\text{sign},P} \left(\log_2 R_{\text{mask},P} + \frac{D_{\text{life},P,\chi}}{\ln 2} \right).$$

In the independent-challenge reference execution, an exact attempt succeeds with probability at least $\rho_{\text{na},P}$. Normalizing its successful-output submeasure amplifies the per-attempt statistical discrepancy by at most $1/\rho_{\text{na},P}$; this factor is included in $R_{\text{mask},P}$ and the certified endpoint success bounds. Stopping removes continuation paths, so the reference execution also bounds exhaustion before the stop. The endpoint exhaustion terms are included in $\varepsilon_{\text{samp},P,\chi}$.

Apply Lemmas 4.1 and 4.2 to the normalized laws with their absorbing flags, and count E only on unflagged executions. The completed, unflagged ideal transcripts form a subset of E in the unbounded exact game; the endpoint cap discrepancies are already included in $\varepsilon_{\text{samp},P,\chi}$. Thus the stopped comparison

contributes $2^{\beta_{\text{sign},P,\chi}} \Pr[\text{Law}_{\text{sign},P}^{\text{ex}} \in E]^{\vartheta_{\text{sign},P}}$ and $\varepsilon_{\text{samp},P,\chi}$. Restoring the original fixed-point game adds at most $\varepsilon_{\text{seed},P} + \varepsilon_{\text{RO},P}$ by the preceding couplings. The seed term is included in $\varepsilon_{\text{sign},P,\chi}$, yielding Eq. (B.25). $\square$

B.6 Hardened Realization

The hardened realization uses fixed-capacity arithmetic, public schedules, and masked selection in Setup, key loading, preparation, extraction, and signing. The schedule is fixed between designated outcome events: candidate rejection or failure in Setup, invalid-key rejection on loading, sampler acceptance and witness admission in extraction, and the attempt outcome in signing. Immutable norm and series caches are constructed during key preparation. These arithmetic refinements preserve the specified maps and the error bounds of Theorems B.1 and B.2. The count-transcript analysis below quantifies extraction leakage separately.

Setup. The Q128 evaluator uses fixed carry chains and fixed-step CORDIC and square-root computations. The Q96 annular predicate uses masked interval products and fixed-step division and square root. The NTRU solver follows the public schedules and capacities of Section B.2. Failed candidate stages terminate that candidate, so the complete Setup invocation is not claimed to have a constant instruction count. The leakage model exposes candidate-stage outcomes without assuming their independence from the returned key.

Loading and extraction. Master-key loading reuses the Q96 annular predicate. Preparation uses masked Q96 interval products and fixed-step division and square root. Division by the secret polynomial f uses fixed-schedule field inversion in both Setup and loading; inversion of a public identity label may use variable-time arithmetic.

The scalar and paired extraction kernels use fixed table scans, public Horner schedules, and complete proposal records. The continuous-perturbation loop has a key-independent acceptance pattern; the discrete rejection loop depends on the center.

Extraction count transcript. For an uncapped discrete-sampler call, conditional on a center, the accepted output and proposal count are independent: if w_z is the one-proposal accepted mass and $p = \sum_z w_z$, their joint mass is $(1-p)^{t-1}w_z$. The count is therefore geometric with parameter p .

The center-independent reference parameters are

$$p_{*,\text{sc}} = \frac{\kappa_{\text{acc}} r \sqrt{2\pi}}{2Z_0}, \quad p_{*,\text{pair}} = \frac{\pi r^2}{\sqrt{3}Z_0Z_1}.$$

Poisson summation bounds the relative variation of the ideal acceptance probabilities by $1.085 \cdot 10^{-14}$ and $5.718 \cdot 10^{-19}$, respectively. Adding the table, evaluator, and U56 errors gives certified intervals for the actual p . For $a = 512$, the exact geometric moment is

$$\log_2 R_a(G_p \| G_{p_*}) = \frac{1}{a-1} \log_2 \frac{p^a / p_*^{a-1}}{1 - (1-p)^a / (1-p_*)^{a-1}}.$$

The denominator is positive throughout each interval. The moment decreases up to $p = p_*$ and increases thereafter, so its endpoint values bound it uniformly in the center.

Let Tr_P record the perturbation acceptance patterns, discrete proposal counts, and resulting SHAKE consumption, and let Tr_P^0 replace each discrete count by an independent G_{p_*} count. Perturbation acceptance patterns have a key-independent law and are independent of their accepted values. They therefore contribute no divergence. The uniform conditional bounds compose adaptively for the discrete counts, and deterministic refills preserve the bound by data processing.

Table B.7. Lifetime count-transcript bounds for uncapped fixed-point extraction with $Q_{\text{ID}} \leq 2^{48}$.

profile	upper bound on $\log_2 R_{512}(\text{Tr}_P \parallel \text{Tr}_P^0)$
IBS-1s	$2.235024 \cdot 10^{-6}$
IBS-1	$3.247352 \cdot 10^{-13}$
IBS-2	$4.470047 \cdot 10^{-6}$
IBS-3	$4.871028 \cdot 10^{-13}$

If d_P is the bound in Table B.7 and Tr_P^{act} includes the actual caps and terminal admissibility decision, coupling and Lemma 4.1 give

$$\Pr[\text{Tr}_P^{\text{act}} \in E] \leq 2^{h_{512}d_P} \Pr[\text{Tr}_P^0 \in E]^{h_{512}} + \varepsilon_{\text{ext},P},$$

where the reference transcript always reports successful termination. For a joint view of outputs and counts, replace the counts first and then apply the three output comparisons of Section B.4. Its exponent is h_{512}^4 and its factor exponent is

$$h_{512}d_P + h_{512}^2L_{q,P} + h_{512}^3L_{e,P} + h_{512}^4L_{o,P} + L_{\text{guard},P}.$$

Table B.7 bounds the count and refill transcript, not the full instruction schedule of extraction.

Signing. Every hardened attempt completes all κ transport steps, the response-bound scans, the high-bit equality check, and candidate encoding before deciding whether to retry. A rejection flag records sampler-budget failures, transport aborts, unresolved comparisons, and failed checks. Exhaustion of a mask-sampling budget selects a zero mask pair before commitment generation; an aborted or unresolved transport comparison selects a zero update. In either case, the remaining arithmetic executes with rejection irrevocably recorded. Consuming the full transport-coin schedule after rejection preserves the ideal-XOF distribution because subsequent attempts draw from independent unread suffixes (Section B.5). Operational stream failures terminate the call as errors.

Alias selection in the mask sampler scans the table at fixed addresses. On AArch64, sixteen proposals are processed together: each group of 64 entries is

loaded, and selection takes place within vector registers. The scalar implementation scans all entries. An alias-prefix tie invokes a complete threshold scan, while deferred-suffix retrieval scans every reserved slot.

The hardened challenge sampler represents the support in Section 3.1 as a list of occupied positions. At iteration $t = 0, \dots, \kappa - 1$, put $i = N - \kappa + t$ and let $j \in \{0, \dots, i\}$ be the accepted draw. A scan of the t occupied positions replaces any entry equal to j by i through masked selection, then appends j . This implements the assignments $c_i \leftarrow c_j$, $c_j \leftarrow 1$ in Section 3.1: the occupied positions remain distinct and give the same support for every sequence of draws. A fixed sequence of masked adjacent compare-exchanges puts the support in canonical order. Construction uses $\kappa(\kappa-1)/2$ masked collision checks, and sorting uses the same number of compare-exchanges, with public loop bounds and list addresses. Thus collection and ordering preserve both the challenge encoding and the order of the transport steps.

Both ratio evaluations follow the known-sign Horner recurrence of Section B.5. Fixed-limb arithmetic, reciprocal division with masked correction, and fixed-schedule modular reduction give arithmetic loops whose lengths are independent of their operands. Mathematical rejections are accumulated in the attempt flag. An arithmetic-guard or invariant failure terminates the signing call and clears its output; it is not retried.

Monomial dot products and response updates scan all N coefficients, applying the ring's wrap and fold rules through masks. The direct AArch64 kernels execute $N/4$ iterations of four coefficients each. For IBS-3, a folded representation provides contiguous access to the shifted coefficients. For a witness polynomial s with coefficient halves $s^{(0)}, s^{(1)}$, preparation stores

$$\text{Fold}(s) = \text{concat}(-s^{(0)} - s^{(1)}, -s^{(1)}, s^{(0)}, s^{(0)}, s^{(0)} + s^{(1)}, s^{(1)}).$$

The relations $X^N = X^{N/2} - 1$ and $X^{3N/2} = -1$ imply that, for $0 \leq \ell < N$, the coefficient halves of $X^\ell s \bmod \phi$ are the length- $N/2$ windows beginning at offsets $N - \ell$ and $5N/2 - \ell$. Both windows lie within the $3N$ stored coefficients. A cached dot product or update traverses them in two loops, each with $N/8$ AArch64 vector iterations or $N/2$ scalar iterations. The stored sums have magnitude at most $2B_{\text{shift}} < 2^{31}$, and Section B.5 bounds the partial dot sums. Saturating-add checks and masked stores detect overflow and retain the previous value of an affected coefficient. Since the windows give exact monomial shifts, their use preserves every transport decision and, for a fixed arithmetic schedule and random tape, the complete signature and retry count.

The fixed-schedule guarantee covers the designated arithmetic kernels. Challenge-indexed addresses remain visible, and mask compaction, suffix refinement, challenge rejection draws, XOF refills, and outer retries remain variable. The analysis assumes fresh independent randomness and excludes RNG failure or rollback, faults (including forced-mask faults), value-dependent power leakage, erasure guarantees, external scheduling, and leakage through shared microarchitectural state.

B.7 Security of the Fixed-Point Scheme

This section combines the ideal IBS reduction with the finite-extraction and signing comparisons of Theorems B.1 and B.2 to bound forgery against the fixed-point implementation. Corollary B.2 applies to either realization mode under the stated query budgets and specified master-key law. It accounts for the Rényi factors and exponents, sampling failures, public-key screening, and the computational replacement of pseudorandom streams by ideal tapes. The numerical evaluation then quantifies the additional reduction loss due to finite precision under the conditional BMST work–success model of Section A.6.

Let $\text{ANSA-IBS}_{P,\chi}^{\text{fx}}$ denote the concrete scheme using Algorithms B.5 and B.7 together with Algorithm B.8 in mode χ .

Ideal security bound. Let $\mathcal{B}_P^{\text{ideal}}$ be the bound obtained from Corollary 4.1 and Theorem 4.2 for the unscreened specified AnnularTrapdoorGen, the centered ungated exact Hybrid-Sampler law $\text{Law}_{\text{ext},P}^{\text{hyb}}$, and exact signing with $\varepsilon_{\text{impl}} = 0$, omitting the extraction-tape term $\text{Adv}_{\text{ext-rand}}(\mathcal{F}_{G_{\text{ext}}})$. That advantage is charged separately below. The remaining bound includes the identity-layer, signing, digest-collision, hardness, and centering-change terms; finite-extraction failure and signing errors enter through Theorems B.1 and B.2.

Computational replacements. Let $\Delta_{\text{XOF},P,\chi}^{\text{fx}}$ be the joint distinguishing advantage for replacing the domain-separated SHAKE256 calls by the ideal XOF $\mathcal{O}$. Each distinct input receives an independent random tape, and repeated inputs reuse their tapes. This term excludes the keyed extraction-root hop.

Evaluate $\text{Adv}_{\text{ext-rand}}(\mathcal{F}_{G_{\text{ext}}})$ in the concrete instance of ExtRand^b from Section 4.1: Setup uses Algorithm B.5, extraction uses Algorithm B.7, and signing uses Algorithm B.8 in mode χ . For $b = 0$, extraction tapes are $G_{\text{ext}}(k_{\text{ext}}(B), \text{id}, j)$; for $b = 1$, they are independent cached 512-bit strings. The event $\text{Bad}_{\text{ext-coll}}$ includes extraction-root collisions.

The leading tags 'E' and 'S' make the extraction and signing sampler input families disjoint, including when a 512-bit extraction root equals a signing seed. Repeated roots and seeds within each family retain their respective collision bounds; cross-family overlap has probability zero.

Corollary B.2 (Concrete ANSA-IBS security). Fix either public static realization mode χ and the specified generator policy of Sections B.1–B.3. Set

$$\mathcal{B}_P^{\text{scr}} := \mathcal{B}_P^{\text{ideal}} + \Delta_{\text{stab},P} + \varepsilon_{\text{stab},P}.$$

For $Q_{\text{ID}} \leq 2^{48}$ and $Q_S \leq 2^{64}$,

$$\begin{aligned} & \text{Adv}_{\text{ANSA-IBS}_{P,\chi}^{\text{fx}}}^{\text{ID-UF-CMA,ROM}}(\mathcal{A}) \\ & \leq \Delta_{\text{XOF},P,\chi}^{\text{fx}} + \text{Adv}_{\text{ext-rand}}(\mathcal{F}_{G_{\text{ext}}}) \\ & \quad + 2^{\beta_{\text{sign},P,\chi}} \left(2^{\beta_P} (\mathcal{B}_P^{\text{scr}})^{\theta_P} + \varepsilon_{\text{ext},P} \right)^{\vartheta_{\text{sign},P}} \\ & \quad + \varepsilon_{\text{sign},P,\chi} + \varepsilon_{\text{RO},P}. \end{aligned} \tag{B.26}$$

Proof. Replace the internal SHAKE256 calls by $\mathcal{O}$ at cost $\Delta_{\text{XOF},P,\chi}^{\text{fx}}$. Next, as in game G_1 of Theorem 4.2, replace the keyed extraction roots and tapes by

independent cached uniform tapes at cost $\text{Adv}_{\text{ext-rand}}(\mathcal{F}_{G_{\text{ext}}})$. This hop precedes the application of Theorem B.1, whose bridge requires private fresh extraction tapes; its cost is therefore additive outside the Rényi powers. Assumption 4.1 applies to the specified finite candidate map, including its cutoff and NTRU-solver predicates. Proposition B.1 proves correctness on accepted candidates, and Corollary B.1 identifies the resulting returned-key law. Proposition B.2 makes preparation a deterministic construction on every valid key, so it introduces no further rejection or joint-law conditioning. Exhaustion of the Setup cap produces no public key and therefore cannot produce a forgery.

Apply Theorem B.2 to the complete signing transcript in mode χ . Theorem B.1 applies uniformly to every returned master key and transfers extraction to $\text{Law}_{\text{ext},P}^{\text{hyb}}$ with factor 2^{β_P} , exponent θ_P , and error $\varepsilon_{\text{ext},P}$. Removing the Galois screen in the centered, ungated exact game costs at most $\Delta_{\text{stab},P} + \varepsilon_{\text{stab},P}$. Lemma 4.3 and the witness transition in Theorem 4.2 bound the unscreened game by $\mathcal{B}_P^{\text{ideal}}$. Substitution gives Eq. (B.26). Table B.7 contributes no term because the ID-UF-CMA experiment provides no execution-transcript oracle. $\square$

Numerical calibration. Set the computational distinguishing advantages to zero and bound the extraction-root collision contribution by $\varepsilon_{\text{root}} = \binom{Q_{\text{ID}}}{2} 2^{-512} < 2^{-417}$. The function $\mathcal{E}_P$ of Section A.6, which includes this statistical term, also upper-bounds the ideal expression with that term removed. With $\Delta_{\text{stab},P} = 0$, Eq. (B.26) therefore gives the conservative transfer

$$\begin{aligned} \mathcal{F}_{P,\chi}(\epsilon) &:= \varepsilon_{\text{root}} + 2^{\beta_{\text{sign},P,\chi}} \left(2^{\beta_P} (\mathcal{E}_P(\epsilon) + \varepsilon_{\text{stab},P})^{\theta_P} + \varepsilon_{\text{ext},P} \right)^{\vartheta_{\text{sign},P}} \\ &\quad + \varepsilon_{\text{sign},P,\chi} + \varepsilon_{\text{RO},P}. \end{aligned}$$

The conditional probability exponent is

$$b_{\text{ref},P,\chi}^{\text{fx}} = -\log_2 \mathcal{F}_{P,\chi}(\varepsilon_{\text{ref},P}(2^{64}, 2^{64})).$$

At the displayed precision, this exponent is given by $\vartheta_{\text{sign},P}(\theta_P b_{\text{ref},P} - \beta_P) - \beta_{\text{sign},P,\chi}$. For Q128, the rounded probability exponents are 10.167, 44.284, 112.077, and 149.843 for IBS-1s/1/2/3. Under the same work–success calibration, adding 64 gives time–success exponents 74.167, 108.284, 176.077, and 213.843. The mixed-mode values differ by less than 0.003 bit. These are conditional evaluations of Eq. (A.2); Table A.2 separately gives the reduction losses at specified forgery-success probabilities.

C Falcon-Style Private-Key Generation with Unchanged User Algorithms

This appendix replaces annular trapdoor generation and Hybrid-Sampling at the private-key generator (PKG) by Gaussian NTRU generation and Klein–GPV preimage sampling, the sampling framework underlying Falcon [12,15,30]. The public identity map, witness relation, witness width, and admission checks

remain those of ANSA-IBS. Consequently, the replacement uses the existing signing and verification implementations, including their fixed-point parameters and encodings. Its master public key and identity keys are newly generated. The identity budgets are those of Section 5.1. The extraction bounds below are certified uniformly for $Q_{\text{ID}} \leq 2^{48}$ and therefore also apply to the smaller IBS-1s identity-label budget; $Q_S, Q_H, Q_\mu \leq 2^{64}$ throughout.

C.1 Reference Key Generation and Extraction

Fix a profile from Table 1 and put $d = 2N$ and $\alpha_F = 1.17$. For an NTRU basis B , let $C_B = (\mathbf{b}_1, \dots, \mathbf{b}_d)$ be the ordered real basis obtained by applying the normalized canonical embedding to

$$(X^j(g, -f))_{j=0}^{N-1}, \quad (X^j(G, -F))_{j=0}^{N-1}.$$

Write $\tilde{\mathbf{b}}_i$ for its Gram–Schmidt vectors. The replacement imposes the quality condition

$$g_F(B) := \max_{1 \leq i \leq d} \|\tilde{\mathbf{b}}_i\|_2 \leq \alpha_F \sqrt{q}. \quad (\text{C.1})$$

This is a Gram–Schmidt condition on the expanded real basis. Falcon’s usual coefficient metric agrees with this metric in the power-of-two rings; the trinomial profiles use the canonical Gram matrix H_N of Section 2.1.

Algorithm C.1 (Gaussian NTRU Setup). For at most 2^{16} independent candidates:

1. Sample $(f, g) \leftarrow D_{R^2, \tau}^{\text{amb}}$ with $\tau = \alpha_F \sqrt{q/(2N)}$.
2. Require $f \in R_q^\times$ and run the Gaussian NTRU-solver policy defined below. Apply the second-row norm and encoding limits of Section B.1. Reject unless $fG - gF = q$, and set $h = gf^{-1} \bmod q$.
3. Reject unless Eq. (C.1) holds and h passes the Galois screen of Algorithm B.2.
4. Return the encoded master keys. Return $\perp$ if every candidate is rejected.

In a trinomial ring, an exact canonical coefficient pair can be sampled by proposing independent $u, v \leftarrow D_{\mathbb{Z}, \sqrt{2}\tau}$ and accepting with probability $\exp(-(u+v)^2/(4\tau^2))$. Its accepted mass is proportional to $\exp(-(u^2+uv+v^2)/(2\tau^2))$, as required by G . Section C.2 supplies exact scalar sampling. The Gram matrix of C_B is rational, so rational Gram–Schmidt computation decides the squared version of Eq. (C.1) exactly. The Gaussian NTRU solver uses arbitrary-length integer norm descent and lifting along the ring tower. At the bottom, use the canonical Bézout representative specified in Section B.2 and reject if the gcd is not one. At each lifted level, compute the projection quotient $(Ff^* + Gg^*)/(ff^* + gg^*)$ exactly in the number field, round its power-basis coefficients to nearest integers with ties to even, and subtract the resulting multiple of (f, g) only if the coefficient squared norm decreases. Stop on a zero quotient or a nondecreasing update; reject if all 128 rounds (IBS-1s/1/2) or 256 rounds (IBS-3) improve

without stopping. Require the exact NTRU equation and signed-32-bit (F, G) coefficients. This policy is part of $\mathcal{H}_F$ and uses neither the annular candidate cut-off nor annularity-dependent finite capacities. The second-row norm threshold of Section B.1 is applied as a fixed public filter: its derivation there uses annularity, but here it is simply part of the specified key law $\mathcal{H}_F$ and of Assumption C.1, and it also guarantees the coefficient ranges required by the compact master-key encoding.

Let $\mathcal{H}_F$ be the public-key law of Algorithm C.1 conditioned on successful output when only its Galois screen is omitted. Independent candidate rejection gives the same conditional law for the capped loop as for one accepted candidate. Thus the NTRU-solver policy and quality test are part of the specified law.

If $p_{F,P}$ is the probability that a Gaussian candidate passes all tests in Algorithm C.1, including the Galois screen, Setup exhausts its cap with probability $(1 - p_{F,P})^{2^{16}}$. No lower bound on $p_{F,P}$ is established here; the substitution guarantees compatibility and security of returned keys, but does not establish reliable key-generation availability.

Assumption C.1 (Decisional Gaussian NTRU). For the fixed profile and generator above, define

$$\Delta_F(\mathcal{D}) := \left| \Pr_{h \leftarrow \mathcal{H}_F} [\mathcal{D}(h) = 1] - \Pr_{h \leftarrow U(R_q)} [\mathcal{D}(h) = 1] \right|.$$

We assume this advantage is negligible for efficient distinguishers. This assumption replaces Assumption 4.1 for the alternative generator; the unit-label and BMST-RSIS assumptions remain unchanged.

Algorithm C.2 (Gaussian preimage extraction). Retain the existing extraction interface and its input limits. Compute $(j, a) = H_{\text{id}}(h, \text{id})$ and the deterministic extraction root as in Algorithm 3.2. On identity-map failure, return $\perp$. Use the 'E' stream with index zero, consumed sequentially, for the following draws:

1. Let $\mathbf{t}$ be the normalized canonical image of $(a, 0)$, using the lift of a with coefficients in $\{0, \dots, q-1\}$ (the output law is lift-independent, as in Section B.4), and set $\mathbf{v} = 0$.
2. For $i = d, d-1, \dots, 1$, compute

$$c_i = \frac{\langle \mathbf{t} - \mathbf{v}, \tilde{\mathbf{b}}_i \rangle}{\|\tilde{\mathbf{b}}_i\|_2^2}, \quad \varsigma_i^2 = \frac{\sigma_1^2}{\|\tilde{\mathbf{b}}_i\|_2^2}.$$

Draw $x_i \leftarrow D_{\mathbb{Z}, \varsigma_i, c_i}$ and update $\mathbf{v} \leftarrow \mathbf{v} + x_i \mathbf{b}_i$.

3. Convert $\mathbf{t} - \mathbf{v}$ to its integral polynomial pair and apply $\mathcal{C}_q$. Unless every monomial shift satisfies Eq. (B.14), return $\perp$ and cache $\perp$ for this identity; no redraw is made.
4. Return (j, a, s_0, s_1) in the existing user-key format and cache it for this identity.

Every returned key satisfies $s_0 + hs_1 = a \bmod q$ because $\mathbf{v}$ represents an element of $\Lambda_{h,q}$. Equation (C.1) gives $\varsigma_i \geq \sigma_1 / (\alpha_F \sqrt{q}) = R_{\text{ext}} = 1.29$. All centers

c_i and variances ς_i^2 are rational: their computation uses the rational coefficient Gram matrix and the rational value of σ_1^2 . The same width therefore defines the intended canonical witness law in both ring families.

C.2 Exact Sampling and Extraction Bounds

The scalar draws admit an exact implementation with finite computations almost surely. For center $c \in \mathbb{Q}$ and variance $v \in \mathbb{Q}_{>0}$, set $m = \lfloor c \rfloor$ and $b = \lceil \sqrt{v} \rceil$. Draw two independent geometric variables

$$\Pr[G_i = k] = (1 - e^{-1/b})e^{-k/b}, \quad k \geq 0,$$

and propose $z = m + G_1 - G_2$. Accept with probability e^{-T} , where

$$T = \frac{(z - c)^2}{2v} - \frac{|z - m|}{b} + \frac{v}{2b^2} + \frac{1}{b} \geq 0. \quad (\text{C.2})$$

Otherwise repeat. One obtains $G_i = \lfloor -b \ln U_i \rfloor$ from independent lazy uniform bit tapes $U_i \in (0, 1)$. Directed interval refinement determines this integer and compares a fresh uniform tape with e^{-T} . Unresolved comparisons increase precision; they do not reject a proposal. Boundary equalities have probability zero, so refinement terminates almost surely and introduces no numerical sampling error.

Lemma C.1 (Exact scalar sampling). This procedure outputs $D_{\mathbb{Z}, \sqrt{v}, c}$. For $\sqrt{v} \geq 1.29$, its acceptance probability exceeds $1/4$.

Proof. The proposal mass is $\tanh(1/(2b))e^{-|z-m|/b}$. Multiplying by e^{-T} cancels the linear exponent and leaves a constant times $e^{-(z-c)^2/(2v)}$. Nonnegativity of T follows from $|z - m| \leq |z - c| + 1$ and completing the square. Put $r = 1.29$ and

$$u_r = \frac{2e^{-2\pi^2 r^2}}{1 - e^{-6\pi^2 r^2}}.$$

Poisson summation lower-bounds the Gaussian normalizer by $\sqrt{2\pi v}(1 - u_r)$. Since $b \geq 2$, $b \leq \sqrt{v} + 1$, and $\tanh x \geq x - x^3/3$ for $x \geq 0$, the acceptance probability is at least

$$\frac{47\sqrt{2\pi} r}{96(r+1)e}(1 - u_r) > 0.2543 > \frac{1}{4}.$$

Thus the proposal loop has fewer than four iterations in expectation. $\square$

Let $K_F(\cdot \mid B, h, a)$ denote Algorithm C.2's uncentered output law before its admission test, and retain $D_{\text{pub}}(\cdot \mid h, a)$ from Section 2.2. Define

$$\begin{aligned} \varepsilon_{F,P} &= \left(\frac{e^{2\pi^2 R_{\text{ext}}^2}}{4N_P} - 1 \right)^{-1}, \quad \Gamma_{F,P} = \left(\frac{1 + \varepsilon_{F,P}/(2N_P)}{1 - \varepsilon_{F,P}/(2N_P)} \right)^{2N_P}, \\ \delta_{F,P} &= \Gamma_{F,P} - 1. \end{aligned} \quad (\text{C.3})$$

Lemma C.2 (Extraction comparison). For every admitted basis and syndrome, the likelihood ratio K_F/D_{pub} lies in $[\Gamma_{F,P}^{-1}, \Gamma_{F,P}]$. For $a_p > 1$,

$$R_{a_p}(K_F \| D_{\text{pub}}) \leq r_{F,P}(a_p) := \left(1 + \frac{a_p(a_p - 1)\delta_{F,P}^2}{2(1 - \delta_{F,P})^{a_p+1}} \right)^{1/(a_p-1)}. \quad (\text{C.4})$$

The same bounds hold after coefficient centering. Over Q_{ID} adaptive extractions, the Rényi factor is at most $r_{F,P}(a_p)^{Q_{\text{ID}}}$.

Proof. The Gram–Schmidt smoothing bound gives $R_{\text{ext}} \geq \eta_{\varepsilon_{F,P}}(\mathbb{Z}^{2N_P})$. Prest’s comparison for Klein’s sampler [29, Lemma 6] therefore applies to C_B . Its proof bounds every scalar Gaussian normalizer and gives the displayed likelihood interval uniformly over the affine center. The map from the sampled lattice point to the preimage preserves likelihood ratios. The relative-error lemma [29, Lemma 3] gives Eq. (C.4); coefficient centering uses data processing, and Lemma 4.2 gives adaptive composition. The certified values satisfy $\varepsilon_{F,P} < 1/4$ and $\delta_{F,P} < 1$. $\square$

The quality condition also gives $\eta_{\varepsilon_u}(\Lambda_{h,q}) \leq \rho_u(N, \varepsilon_u)\alpha_F\sqrt{q}$. Thus the quotient-smoothing and witness-programming bounds of Lemma 4.4 hold with their existing parameters. For the witness gate, the proof of Lemma B.3 uses the master basis only through this Gram–Schmidt smoothing bound. Under an invertible linear map T , the distance of $T\mathbf{b}_i$ to its preceding span is at most $\|T\|_{\text{op}}\|\tilde{\mathbf{b}}_i\|_2$, so the same bound controls the tilted Gaussian covariances in that proof. Replacing $Q_{\text{hyb}}(B)$ there by $g_F(B)$ gives the same ideal gate-failure bounds $g_P = 2^{-b_P}$, with $b_P = (173.998, 226.999, 297.999, 334.999)$. The likelihood bound in Lemma C.2 then gives

$$\Pr[\text{an extraction fails the witness gate}] \leq \varepsilon_{\text{gate},F,P} := Q_{\text{ID}}\Gamma_{F,P}g_P. \quad (\text{C.5})$$

This bound includes coefficient-centering changes and both endpoints of the shift-norm interval. It is uniform over the chosen syndromes and prior transcript.

For comparison, Falcon’s nominal width prescription [30, Section 2.6] uses $\epsilon = (Q_{\text{ID}}\lambda)^{-1/2}$ and

$$\sigma_F = \frac{\alpha_F\sqrt{q}}{\pi} \sqrt{\frac{\ln(4N(1+1/\epsilon))}{2}}.$$

Table C.1 compares this prescription with the retained width and gives directed-interval bounds for the replacement at σ_1 . The smaller nominal widths demonstrate the available margin; Algorithm C.2 retains σ_1 so that the existing lower and upper witness-admission checks continue to apply.

C.3 Security with the Existing Signing Implementation

Let $\mathcal{B}_{F,P}$ be the bound obtained from Corollary 4.1 combined with Theorem 4.2, with r_p replaced by $r_{F,P}(a_p)$, Δ_{ann} replaced by Δ_F , and the extraction-tape term removed. Eq. (C.6) charges that term, $\text{Adv}_{\text{ext-rand}}^F$, separately. This

Table C.1. Gaussian PKG extraction at $q = 8380417$ and $Q_{\text{ID}} \leq 2^{48}$. Widths are rounded to three decimals; Rényi-factor bounds are rounded upward and failure exponents downward.

quantity	IBS-1s	IBS-1	IBS-2	IBS-3
nominal Falcon width σ_F	3989.010	4018.438	4053.747	4072.458
retained witness width σ_1	4369.265	4369.265	4369.265	4369.265
$Q_{\text{ID}} \log_2 r_{F,P}(512)$	< 0.000206	< 0.000462	< 0.000821	< 0.001039
$-\log_2 \varepsilon_{\text{gate},F,P}$	> 125.99	> 178.99	> 249.99	> 286.99

advantage is evaluated in ExtRand^b using Algorithm C.1 for Setup, Algorithm C.2 for extraction, and Algorithm B.8 for signing in mode χ , with the same tape rule for $b = 0, 1$. The witness-first distribution uses the same σ_1 , the uniform-key transition uses Assumption 4.2, and the signing relation uses the same BMST-RSIS bounds.

Let $\Delta_{\text{stab},F,P}$ be the Assumption C.1 advantage of the public stabilizer test from Section 5.3. The same argument bounds the distance between the screened and unscreened new key laws by $\Delta_{\text{stab},F,P} + \varepsilon_{\text{stab},P}$. Put

$$\mathcal{U}_{F,P} = \mathcal{B}_{F,P} + \Delta_{\text{stab},F,P} + \varepsilon_{\text{stab},P} + \varepsilon_{\text{gate},F,P}.$$

Theorem C.1 (PKG substitution with unchanged signing and verification). Algorithms C.1 and C.2 use the existing public-key and user-key formats and the signing and verification algorithms of Sections 3.4 and 3.5. Every returned key is admitted by the existing fixed-point signer. Under the stated query budgets and for either static signing mode χ , the resulting scheme satisfies

$$\begin{aligned} \text{Adv}_{\text{ANSA-IBS}_{F,\chi}^{\text{ID-UF-CMA,ROM}}}(\mathcal{A}) &\leq \Delta_{\text{XOF},P,\chi}^F + \text{Adv}_{\text{ext-rand}}^F + 2^{\beta_{\text{sign},P,\chi}} \mathcal{U}_{F,P}^{\vartheta_{\text{sign},P}} \\ &\quad + \varepsilon_{\text{sign},P,\chi} + \varepsilon_{\text{RO},P}. \end{aligned} \quad (\text{C.6})$$

Here $\Delta_{\text{XOF},P,\chi}^F$ is the joint XOF replacement advantage for this realization, excluding the separately charged keyed extraction-root transition. The signing factors and additive errors are exactly those in Theorem B.2 and Table B.6.

Proof. Algorithm C.1 returns a valid NTRU public key in the same ring and modulus. Algorithm C.2 produces its identity preimage, centers it, and applies precisely the admission predicate required by Sections B.5 and B.6. Normalization still gives $A_0 s_0 + A_1 s_1 = \hat{q} \bmod q$. The existing user-key and public-key loaders, monomial transport kernels, response encoding, and verifier therefore apply without modification, and correctness follows from Lemma 4.11. Gaussian master bases need not pass the annular master-key loader, which applies the Q96 annular predicate and Valid_P ; Algorithm C.1 shares only the byte format of EncodeMSK ; its master-key loader must omit the annular check.

In the ideal-XOF game, first replace the keyed extraction roots and tapes by independent cached uniform tapes at cost $\text{Adv}_{\text{ext-rand}}^F$, before any extraction comparison. Then couple gated and ungated extraction until the first gate failure;

Eq. (C.5) bounds its probability. The Gaussian-to-ideal comparison is Lemma C.2. Quotient smoothing follows from Eq. (C.1), so all subsequent identity-layer hybrids in Theorem 4.2 apply with the substitutions defining $\mathcal{B}_{F,P}$. Assumption C.1 supplies the public-key switch. Removing the public Galois screen costs the stated single-key term, giving $\mathcal{U}_{F,P}$ for exact signing.

Theorem B.2 holds uniformly for every admitted signing statement and hence for the keys returned by Algorithm C.2. Applying it to the entire adaptive signing transcript yields Eq. (C.6); the joint XOF hop restores the concrete stream realization. Lemma C.1 and exact rational preparation add no numerical extraction error. A Setup invocation that exhausts its candidate cap produces no public key and cannot yield a successful forgery. $\square$

The replacement changes the assumed master-key law and the extraction comparison, while preserving the complete user-side implementation and signature sizes. Its reference extractor uses exact scalar decisions; the fixed-point Hybrid-Sampler bounds and PKG timings in Appendix B and Table 3 describe the annular implementation.

References

1. Abapour, N.: Sharper and Closed-Form Attacks on SIS When Modulus Is Small. Cryptology ePrint Archive, Paper 2026/1349 (2026)
2. Barbosa, M., Barthe, G., Doczkal, C., et al.: Fixing and Mechanizing the Security Proof of Fiat-Shamir with Aborts and Dilithium. In: CRYPTO 2023. LNCS, vol. 14085, pp. 358–389. Springer (2023)
3. Bellare, M., Namprempre, C., Neven, G.: Security Proofs for Identity-Based Identification and Signature Schemes. *Journal of Cryptology* **22**, 1–61 (2009)
4. Bellare, M., Neven, G.: Multi-Signatures in the Plain Public-Key Model and a General Forking Lemma. In: Proceedings of the 13th ACM Conference on Computer and Communications Security. pp. 390–399. ACM (2006)
5. Cheng, Z., Xiong, K.: ANSA-IBKEM: Practical Quantum-Safe Identity-Based Key Encapsulation via Annular NTRU Trapdoors and Standardized PQC Arithmetic Reuse. Cryptology ePrint Archive, Paper 2026/1449 (2026)
6. Cheon, J.H., Choe, H., Devevey, J., et al.: HAETAE: Shorter Lattice-Based Fiat-Shamir Signatures. Cryptology ePrint Archive, Paper 2023/624 (2023)
7. Cramer, R., van Gent, D., Lesavourey, A., Pellet-Mary, A.: On Module Lattices with Galois-Symmetries: What You See Is Not What You Get. Cryptology ePrint Archive, Paper 2026/1805 (2026)
8. Devevey, J., Fallahpour, P., Passelègue, A., Stehlé, D.: A Detailed Analysis of Fiat-Shamir with Aborts. In: CRYPTO 2023. LNCS, vol. 14085, pp. 327–357. Springer (2023)
- 9.ucas, L., Durmus, A., Lepoint, T., Lyubashevsky, V.: Lattice Signatures and Bimodal Gaussians. In: CRYPTO 2013. LNCS, vol. 8042, pp. 40–56. Springer (2013)
- 10.ucas, L., Espitau, T., Postlethwaite, E.W.: Finding Short Integer Solutions When the Modulus Is Small. In: CRYPTO 2023. pp. 150–176. Springer (2023)
- 11.ucas, L., Kiltz, E., Lepoint, T., et al.: CRYSTALS-Dilithium: A Lattice-Based Digital Signature Scheme. *IACR Transactions on Cryptographic Hardware and Embedded Systems* **2018**(1), 238–268 (2018)

12. Ducas, L., Lyubashevsky, V., Prest, T.: Efficient Identity-Based Encryption over NTRU Lattices. In: ASIACRYPT 2014. LNCS, vol. 8874, pp. 22–41. Springer (2014)
13. Ducas, L., van Woerden, W.P.J.: NTRU Fatigue: How Stretched is Overstretched? In: ASIACRYPT 2021. LNCS, vol. 13093, pp. 3–32. Springer (2021)
14. Espitau, T., Nguyen, T.T.Q., Sun, C., et al.: Antrag: Annular NTRU Trapdoor Generation. In: ASIACRYPT 2023. LNCS, vol. 14444, pp. 3–36. Springer (2023)
15. Gentry, C., Peikert, C., Vaikuntanathan, V.: Trapdoors for hard lattices and new cryptographic constructions. In: Ladner, R.E., Dwork, C. (eds.) 40th ACM STOC. pp. 197–206. ACM Press (May 2008)
16. Gärtner, J.: Compact Lattice Signatures via Iterative Rejection Sampling. Cryptology ePrint Archive, Paper 2024/2052 (2024), revised 2025
17. Günther, F., Lyubashevsky, V., Schmidt, R.: FALCON with Message Recovery, a Specification. Cryptology ePrint Archive, Paper 2026/420 (2026)
18. Karenin, A., Kirshanova, E.: Finding Dense Submodules with Algebraic Lattice Reduction. In: AFRICACRYPT 2024. LNCS, vol. 14861, pp. 403–427. Springer (2024)
19. Kim, J., Park, J.H.: NTRU+: Compact Construction of NTRU Using Simple Encoding Method. IEEE TIFS **18**, 4760–4774 (2023)
20. Kirchner, P., Fouque, P.A.: Revisiting Lattice Attacks on Overstretched NTRU Parameters. In: EUROCRYPT 2017. LNCS, vol. 10210, pp. 3–26. Springer (2017)
21. Liu, Y., Zhang, Y., Lu, X., Cheng, Y., Yin, Y.: DAWN: Smaller and Faster NTRU Encryption via Double Encoding. In: ASIACRYPT 2025. LNCS, vol. 16247, pp. 396–427. Springer (2026)
22. Luo, P., Chen, X., Susilo, W., Huang, Q.: Identity-Based Signature from Lattices Without Trapdoors. In: ICICS 2024. LNCS, vol. 15057, pp. 211–230. Springer (2024)
23. Lyubashevsky, V.: Fiat-Shamir with Aborts: Applications to Lattice and Factoring-Based Signatures. In: ASIACRYPT 2009. LNCS, vol. 5912, pp. 598–616. Springer (2009)
24. Lyubashevsky, V.: Lattice Signatures Without Trapdoors. In: EUROCRYPT 2012. LNCS, vol. 7237, pp. 738–755. Springer (2012)
25. Micciancio, D., Regev, O.: Worst-case to average-case reductions based on Gaussian measures. SIAM Journal on Computing **37**(1), 267–302 (2007)
26. National Institute of Standards and Technology: Module-Lattice-Based Digital Signature Standard. FIPS 204, National Institute of Standards and Technology (2024)
27. Pornin, T.: NTRUGEN: NTRU key pair generation. <https://github.com/pornin/ntruigen>, software, revision f92722f; accessed 4 October 2026
28. Prest, T.: Gaussian Sampling in Lattice-Based Cryptography. Ph.D. thesis, École Normale Supérieure, Paris, France (2015)
29. Prest, T.: Sharper Bounds in Lattice-Based Cryptography Using the Rényi Divergence. In: ASIACRYPT 2017. LNCS, vol. 10624, pp. 347–374. Springer (2017)
30. Prest, T., Fouque, P.A., Hoffstein, J., et al.: FALCON (2017), <https://falcon-sign.info/falcon.pdf>, technical report
31. Woo, J., Kim, J., Hong, G.H., et al.: NTRU+Sign: Compact NTRU-Based Signatures Using Bimodal Distributions. Des. Codes Cryptogr. **94**(1), Article 2 (2026)
32. Xie, J., Hu, Y.P., Gao, J.T., Gao, W.: Efficient Identity-Based Signature over NTRU Lattice. Frontiers of Information Technology & Electronic Engineering **17**(2), 135–142 (2016)