Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Centered Arithmetic Is Not Leakage-Neutral: Finite-Noise Optimal Attacks and Minimax Laws for Modular Representations

José Luis Delgado

jdelgado13@uoc.edu

Universitat Oberta de Catalunya

Research Article

Keywords: side-channel analysis, arithmetic masking, centered reduction, Hamming weight, Fourier analysis, lattice cryptography

Posted Date: August 31st, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10483885/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Centered Arithmetic Is Not Leakage-Neutral

Finite-Noise Optimal Attacks and Minimax Laws for Modular Representations

José Luis Delgado[1*]

[1*]Universitat Oberta de Catalunya, Barcelona, Spain; ORCID: 0009-0009-2771-7507.

Corresponding author(s). E-mail(s): jdelgado13@uoc.edu;

Abstract

Arithmetic masking represents a sensitive residue as the modular sum of random shares. In software, each share may be stored either as an unsigned residue or as a centered signed integer in two's-complement, even though both encodings implement the same modular arithmetic. We study how this choice changes Hamming-weight leakage. Under independent additive noise, we derive the exact likelihood of the masked sum and evaluate every candidate secret by a circular convolution. This gives an optimal finite-noise attack in quasi-linear time and an exact condition under which the commonly used product statistic retains all likelihood information.

We then compare the modular representations themselves. Among contiguous representatives, unsigned storage uniquely minimizes first-order leakage unless the modulus occupies the full machine word, whereas at higher orders it ceases to be optimal. Among arbitrary injective binary encodings, leakage vanishes exactly for constant-Hamming-weight codebooks, whose existence is determined by a sharp capacity bound. Our main result, the Centered-Reduction Leakage Law, gives an exact spectral decomposition of centered two's-complement leakage into sign-extension and carry terms and shows how masking propagates this spectrum across any number of shares. It yields a uniform centered-versus-unsigned separation once three excess bits are available, exact robustness radii under weighted Hamming-weight leakage, and finite-noise sample-complexity bounds. Exact computations, full-size ML-KEM and ML-DSA implementations, synthetic experiments, and public physical traces confirm the analysis.

Keywords: side-channel analysis, arithmetic masking, centered reduction, Hamming weight, Fourier analysis, lattice cryptography

1 Introduction

By encoding a sensitive value as random shares, masking ensures that no proper subset determines it. Security analyses typically specify the number of shares, gadget composition, independence, and noise [1–3]. Encoding each share as a machine word additionally requires a choice of section of $\mathbb{Z}_q$.

When lattice implementations store a small negative integer in two's-complement form, the word contains many leading ones, whereas the corresponding residue in $[0, q)$ does not. Tosun, Moradi, and Savas exploited this effect in central reduction and derived an absolute-value predictor [4]. Higher-order work subsequently developed recursive predictors and near-optimal sine/cosine

distinguishers [5]. Nagpal et al. studied the information and implementation consequences of signed, unsigned, and redundant representations [6]. The common mathematical problem is to determine the leakage imposed by a representation, its amplification under masking, and the representation that minimizes the resulting experiment.

Accordingly, our analysis connects the complete finite-noise channel, the global choice of representation, and the centered-reduction case. At the channel level, it derives an optimal all-hypothesis attack and an exact sufficiency boundary for product compression; at the representation level, it optimizes first over all contiguous sections and then over all injective binary codebooks. The centered-reduction case is governed by an exact law in (q, β, d), including the carry regimes that determine its degeneracies.

1.1 Contributions

1. **Optimal finite-noise inference.** The complete masked-sum likelihood is a cyclic convolution that evaluates all hypotheses in $O(dq \log q)$. A finite necessary-and-sufficient criterion characterizes Gaussian product sufficiency, and nonnegative leakage tables with a unique zero satisfy a global strictness result. Given the q-entry trace likelihoods, a finite-field Radon algorithm scores all q^2 incomplete number-theoretic-transform (NTT) coefficient pairs in $O(Tq+q^2 \log q)$, compared with $O(Tq^2)$ for direct scoring.
2. **Global representation minimax.** Unsigned storage is the unique first-order optimum over all contiguous sections, apart from the complete power-of-two tie. An exact permanent phase change governs higher orders, and a lexicographic law applies to every finite representation family. For arbitrary binary codebooks, constant weight is equivalent to zero finite-noise leakage, with an explicit minimum width and a sharp obstruction beyond its capacity.
3. **Centered-Reduction Leakage Law.** The exact d-share spectrum decomposes into sign-extension and carry terms. The uniform phase-independent dominance threshold is three excess bits, and a finite near-dyadic carry identity yields sharp near-Mersenne, fixed-defect, and thin-defect asymptotics. Representation neutrality occurs only for an exact-width power-of-two modulus.
4. **Robustness and operational consequences.** The weighted Hamming-weight model has an exact robustness minimax, and transition leakage inherits the corresponding representation law. An exponential-confidence harmonic test has a high-noise sample ratio equal to the $2d$-th power of the coefficient ratio. Exact oracles, interval proofs, standardized-parameter executions, and paired finite-noise experiments provide reproducible certificates.

Beyond the main exposition, Online Resource 1 contains the full proofs, whereas Online Resource 2 contains the source code, certificates, and result tables; the structural arguments and all principal theorem statements appear in the main paper.

2 Model and Objectives

Throughout the paper, let $q \geq 3$, $G = \mathbb{Z}_q$, $m = \lceil \log_2 q \rceil$, $r = 2^m - q$, $h = \lfloor q/2 \rfloor$, and $\beta \geq m$, and write $s = \beta - m$ for the excess word width. The centered representative satisfies

$$r_c(x) \in [-\lfloor q/2 \rfloor, \lceil q/2 \rceil - 1],$$

where $q/2$ is represented by $-q/2$ for even q. Let HW denote Hamming weight and $\mathrm{tc}_\beta(z)$ the β-bit two's-complement word, and define

$$w_c(x) = \mathrm{HW}(\mathrm{tc}_\beta(r_c(x))), \quad w_u(x) = \mathrm{HW}(x),$$
$$\Delta = w_c - w_u.$$

For uniform arithmetic shares $X_0, \ldots, X_{d-1}$ conditioned on $S = \sum_i X_i \pmod q$, and with $\bar{w} = w - \mathbb{E}w$, define their conditional product kernel by

$$K_d^w(a) = \mathbb{E}\left[\prod_i \bar{w}(X_i) \mid S = a\right].$$

The unnormalized discrete Fourier transform is $\widehat{f}(k) = \sum_x f(x) e^{-2\pi i k x/q}$, with inverse IDFT. Define

$$D(q, \beta) = q^{-1}|\widehat{\Delta}(1)|, \qquad U(q) = q^{-1}|\widehat{w}_u(1)|.$$

For representation design, write every contiguous section as $R_t = [-t, q-1-t]$, $0 \le t < q$, and let w_t denote its stored Hamming weight. The complete nonzero-frequency order-d energy is

$$E_d(t) = \sum_{k \neq 0} |q^{-1} \widehat{w}_t(k)|^{2d}. \tag{1}$$

At finite noise, the shares have conditionally independent local channels $\ell_i(y_i \mid x_i)$.

3 The Complete Finite-Noise Experiment

Theorem 1 (Exact modular likelihood) *For $a_i(x) = \ell_i(y_i \mid x)$, all conditional sum likelihoods are*

$$p_a(y_0, \dots, y_{d-1}) = q^{1-d}(a_0 * \cdots * a_{d-1})(a).$$

One observation is therefore scored for every $a \in \mathbb{Z}_q$ by $q^{1-d}\,\mathrm{IDFT}(\prod_i \widehat{a}_i)$ in $O(dq \log q)$ operations and $O(q)$ memory. For independent traces and public maps $a_j(h)$, $\sum_j \log p_{a_j(h)}(y_j)$ is the maximum a posteriori (MAP) score and every pairwise difference is the Neyman–Pearson statistic.

Proof Because each prescribed sum has exactly q^{d-1} share tuples, summing the product channel over the corresponding affine fiber produces the circular convolution, whose Fourier transform yields the stated algorithm. MAP and Neyman–Pearson optimality then follow from the likelihood principle. □

For an incomplete-NTT pair $h \in \mathbb{F}_q^2$, each trace contributes a ridge function $g_j(c_j \cdot h + b_j)$. After every nonzero $c_j = \lambda_j u_j$ is normalized by its projective direction, accumulation of $A_u(z) = \sum_{j : u_j = u} g_j(\lambda_j z + b_j)$ produces

$$S(h) = C + \sum_{u \in \mathbb{P}^1(\mathbb{F}_q)} A_u(u \cdot h).$$

The Fourier-slice identity $\widehat{S}(ku) = q\widehat{A}_u(k)$ fills the complete two-dimensional spectrum with $q+1$ length-q transforms, so all q^2 pairs are evaluated in $O(Tq + q^2 \log q)$ time and $O(q^2)$ memory [7].

Applying Fourier inversion to each ridge yields

$$A_u(u \cdot h) = q^{-1} \sum_{k \in \mathbb{F}_q} \widehat{A}_u(k) e^{2\pi \mathrm{i} k u \cdot h / q}.$$

Character orthogonality eliminates every term except the unique projective decomposition of a nonzero frequency $\xi = ku$, for which the two character sums contribute $q\widehat{A}_u(k)$. At frequency zero, each direction contributes its zero-frequency coefficient and the constant contributes $q^2 C$. The ridge tables cost $O(Tq)$, after which one two-dimensional inverse transform completes the evaluation. Including construction of the exact masked likelihood rows, the total cost is $O(Tdq \log q + q^2 \log q)$, with the same accumulator serving affine offsets, zero directions, and every trace prefix.

For the comparison between the complete experiment and the centered product, consider Gaussian channels $Y_i = w_i(X_i) + N_i$, $N_i \sim \mathcal{N}(0, \sigma_i^2)$. Fix centers c_i, put $z_i = Y_i - c_i$, and define

$$\begin{aligned} F_a(z) &= \sum_{x : \sum_i x_i = a} A_x e^{\langle \lambda_x, z \rangle}, \\ \lambda_{x,i} &= w_i(x_i)/\sigma_i^2, \\ A_x &= \prod_i \exp\!\left(-\frac{(c_i - w_i(x_i))^2}{2\sigma_i^2} \right), \\ C_i(a, b, \Lambda) &= \sum_{\substack{x : \sum x_i = a,\ y : \sum y_i = b \\ \lambda_x + \lambda_y = \Lambda}} A_x A_y (\lambda_{x,i} - \lambda_{y,i}). \end{aligned}$$

Theorem 2 (Product sufficiency and strictness) *For $d \ge 2$, $P(z) = \prod_i z_i$ is sufficient for the family $\{p_a : a \in \mathbb{Z}_q\}$ if and only if $C_i(a, b, \Lambda) = 0$ for every a, b, Λ, i. If every local table is nonnegative, has its unique zero at residue zero, and satisfies $w_i(1) > 0$, this condition fails for every $q \ge 2$, $d \ge 2$, and finite positive variance. For $S = 0$ versus $S = 1$, the complete observation then has strictly larger Kullback–Leibler (KL) divergence, Rényi-α divergence for every $0 < \alpha < 1$, Bhattacharyya distance, and Chernoff information than the product statistic.*

Proof Under Gaussian expansion, $p_a(z) = K(z) F_a(z)$, where $K > 0$ is independent of a, so every likelihood ratio is F_a / F_b. The product-preserving vector fields $D_{ij} = z_i \partial_i - z_j \partial_j$ characterize functions that are constant on the fibers of $P(z) = \prod_i z_i$. Expansion of the numerator produces

$$(D_{ij} F_a) F_b - F_a (D_{ij} F_b) = \sum_{\Lambda} (z_i C_i - z_j C_j) e^{\langle \Lambda, z \rangle}.$$

Linear independence of polynomial-exponential functions establishes the necessity of the coefficient test.

Conversely, if every coefficient vanishes, all D_{ij} annihilate every likelihood ratio. Analytic expansion at the origin joins the disconnected product fibers, making each ratio a function of P and thereby establishing sufficiency.

For strictness, choose the combined slope whose only nonzero coordinate is $\Lambda_i = w_i(1)/\sigma_i^2$. Under $(a, b) = (0, 1)$, nonnegativity and the unique zero force the only contributing pair to be the all-zero tuple and the tuple with residue one at coordinate i. Therefore

$$C_i(0, 1, \Lambda) = -A_x A_y w_i(1)/\sigma_i^2 < 0.$$

Because the nonzero coefficient excludes product sufficiency, strict conditional Jensen establishes the KL and Rényi separations, while minimization of the strict affinity over its interior parameter establishes the Bhattacharyya and Chernoff separations. □

Beyond Gaussian noise, the product-fiber calculus yields an exact Taylor-coefficient criterion for every positive real-analytic location channel. When $q = 2$, this criterion excludes sufficiency whenever all local leakage tables are nonconstant. Under Laplace noise, sufficiency is equivalent to equality of the finite leakage-vector histograms conditioned on the two sums; the supplementary material contains the complete non-Gaussian proofs.

Equation (1) also determines the leading finite-noise objective: for $Y_i = w(X_i) + \tau N_i$ with standard Gaussian N_i, uniformly over a, b,

$$C_{\mathrm{Ch}}(p_a, p_b) = \frac{(K_d^w(a) - K_d^w(b))^2}{8\tau^{2d}} + O(\tau^{-2d-1}), \tag{2}$$

and the mean pairwise Chernoff information is $E_d/(4\tau^{2d}) + O(\tau^{-2d-1})$. The representation minimax is therefore the exact leading finite-noise design problem.

4 Representation Minimax

At order one, Equation (1) coincides by Parseval with the variance of the stored word. Each section removes a cyclic interval from the complete β-bit table; among intervals of a fixed length, the initial and terminal intervals maximize both the absolute centered digit sum and its squared moment.

Theorem 3 (Global contiguous-section minimax) *For every $q \geq 2$, $\beta \geq \lceil \log_2 q \rceil$, and $0 \leq t < q$,*

$$E_1(t) \geq E_1(0).$$

If $q < 2^\beta$, unsigned storage $t = 0$ is the unique minimizer. If $q = 2^\beta$, every section induces a cyclic permutation of the complete word table and all sections tie. These are the only equality cases.

Proof For exact width, put $M = 2^m$, $r = M - q$, and write $y(x) = \mathrm{HW}(x) - m/2$. If $r = 0$, every section is a cyclic permutation of the complete m-bit word table. For $r > 0$, a section deletes an interval I of length r from that table. With $a_I = \sum_I y$ and $b_I = \sum_I y^2$, its unnormalized variance is

$$V(I) = q(Mm/4 - b_I) - a_I^2.$$

The binary interval inequalities

$$|a_I| \leq |a_{[0,r)}|, \qquad b_I \leq b_{[0,r)}$$

follow simultaneously by splitting $[0, M)$ into its two dyadic halves, complementing the upper half, and applying induction on m. In the crossing case, reduction to the two shorter endpoint intervals implies $V(I) \geq V([0, r))$. The terminal interval deleted by unsigned storage is the bit complement of this prefix and has the same two moments.

To establish strictness, take an interior interval $I = [a, a + r)$ and apply Graham's binary digital-sum inequality [8], which states that

$$H(a + r) - H(a) - H(r) \geq \min\{a, r\} > 0,$$
$$H(n) = \sum_{x<n} \mathrm{HW}(x).$$

By complementation, the symmetric strict inequality holds at the other endpoint, so $|a_I| < |a_{[0,r)}|$ and $V(I) > V([0, r))$. The prefix corresponds to the duplicate section $t = q$, leaving unsigned as the only minimizer in $0 \leq t < q$.

For additional leading bits, the first two section moments are

$$S_t = H(q - t) + \beta t - H(t),$$
$$T_t = Q(q - t) + \beta^2 t - 2\beta H(t) + Q(t),$$

where $Q(n) = \sum_{x<n} \mathrm{HW}(x)^2$, and $q^2 E_1(t) = qT_t - S_t^2$. Direct differentiation yields

$$\frac{\partial(qT_t - S_t^2)}{\partial\beta} = 2t(q - t)\left(\beta - \frac{H(t)}{t} - \frac{H(q - t)}{q - t}\right) > 0$$

for $0 < t < q$, since $H(n)/n \leq \log_2(n)/2$. Leading zeros do not change unsigned storage, which completes the inequality and the classification of equality. □

Although unsigned storage minimizes the order-one energy, increasing the masking order can permanently change the representation that minimizes E_d.

Theorem 4 (Exact all-order phase change) *At $(q,\beta) = (14,4)$, unsigned is the unique minimizer of E_d for $d = 1,2,3$, whereas centered is the unique minimizer for every $d \geq 4$.*

Proof For an integer table, Parseval and circular convolution establish the identity

$$14^{2d} E_d(t) = 14 \sum_s (w_t^{*d}(s))^2 - \left(\sum_x w_t(x) \right)^{2d}.$$

Through $d = 8$, exact integer convolution identifies the stated unique unsigned and centered minimizers. At the transition, their respective numerators are

$$(131473, 134848) \quad (d = 3),$$
$$(5950145, 4275348) \quad (d = 4).$$

For centered storage, let Y denote an unnormalized squared Fourier magnitude, with $Y = 36$ at the Nyquist mode. At the remaining even and odd modes, reduction modulo Φ_7 and Φ_{14}, respectively, produces

$$Y^3 - 35Y^2 + 378Y - 1183 = 0,$$
$$Y^3 - 45Y^2 + 402Y - 169 = 0.$$

Both polynomials are positive at 36 and strictly increasing above it, which sets the centered squared spectral radius to $9/49$. Every other section has a Nyquist coefficient of squared magnitude at least $1/4$, and therefore

$$E_d(7) \leq 13(9/49)^d, \qquad E_d(t) \geq (1/4)^d \quad (t \neq 7).$$

Since $13(36/49)^9 < 1$, centered storage is optimal for all $d \geq 9$, while exact integer convolution covers $4 \leq d \leq 8$. □

For a finite family $\mathcal{R}$ of real leakage tables on $\mathbb{Z}_q$, sort each squared nonzero-frequency spectrum as $\Lambda_w = (\lambda_{w,1} \geq \cdots \geq \lambda_{w,q-1})$. The resulting lexicographic order determines the complete high-order law.

Theorem 5 (Lexicographic high-order law) *The eventual minimizers of E_d over $\mathcal{R}$ are exactly the tables with lexicographically smallest Λ_w. If a minimizer and a competitor first differ at index j, with $a < b$, the minimizer wins for every*

$$d > \frac{\log(q-j)}{\log(b/a)}.$$

If $a = 0 < b$, the minimizer wins at every order. Identical sorted spectra tie permanently. Every spectral-radius minimizer is a $(q-1)$-approximation to $\min_{w \in \mathcal{R}} E_d(w)$ at every order.

Proof In a pairwise comparison, the first $j-1$ powers cancel, while sortedness implies

$$\sum_{r=j}^{q-1} \lambda_{\min,r}^d \leq (q-j)a^d < b^d \leq \sum_{r=j}^{q-1} \lambda_{w,r}^d.$$

Taking the maximum over the finitely many pairwise thresholds establishes eventual lexicographic optimality and permanent ties for identical spectra. If $M = \min_{w \in \mathcal{R}} \max_k |q^{-1}\widehat{w}(k)|^2$, a radius minimizer has energy at most $(q-1)M^d$, whereas every table has energy at least M^d. □

For arbitrary binary encodings, let $\mathcal{C}_\beta(q)$ contain all injective maps $c : \mathbb{Z}_q \to \{0,1\}^\beta$, with local leakage $w_c(x) = \mathrm{HW}(c(x))$, and put $M_\beta = \binom{\beta}{\lfloor \beta/2 \rfloor}$.

Theorem 6 (Constant-weight finite-noise minimax) *For every $q \geq 2$, $d \geq 2$, and finite positive Gaussian noise, a fixed codebook $c \in \mathcal{C}_\beta(q)$ has identical conditional masked-sum laws if and only if its mean pairwise Chernoff information is zero, if and only if c has constant Hamming weight. Such a codebook exists exactly when $q \leq M_\beta$. Hence the exact minimum width is*

$$\beta_{\mathrm{cw}}(q) = \min\{\beta : M_\beta \geq q\} = F(q) + O(1),$$
$$F(q) = \log_2 q + \tfrac{1}{2}\log_2 \log_2 q + \tfrac{1}{2}\log_2(\pi/2).$$

In particular, $\beta_{\mathrm{cw}}(3329) = 14$ and $\beta_{\mathrm{cw}}(8380417) = 26$. If $q > M_\beta$, every codebook satisfies the sharp bound

$$\sum_{k \neq 0} |q^{-1}\widehat{w}_c(k)|^{2d} \geq \frac{q-1}{q^{2d}},$$

with equality exactly at $q = M_\beta + 1$ for one adjacent-layer word added to a largest Hamming layer. Equivalently, up to adding a constant and changing sign, the leakage table is the indicator of one residue.

Proof A constant-weight codebook has constant local noiseless leakage, so the complete noisy law is independent of the share tuple and its sum. If all conditional laws are identical, centering each observation and taking the conditional expectation of their product leaves $K_d^{w_c}(s)$ by independence and zero-mean noise. This kernel is constant in s and has average zero; consequently, Theorem 7 implies

$$0 = \widehat{K}_d^{w_c}(k) = q^{1-d}\widehat{w}_c(k)^d \quad (k \neq 0).$$

Hence w_c is constant, and because Chernoff information is nonnegative and vanishes exactly for identical laws, the finite-noise equivalences characterize every minimizer.

Because a weight-t Hamming layer contains $\binom{\beta}{t}$ words, adjacent-layer ratios identify the central layer

as the largest and establish the exact capacity criterion. Stirling's formula with uniform remainder yields

$$\binom{\beta}{\lfloor \beta/2 \rfloor} = 2^{\beta}\sqrt{2/(\pi\beta)}(1 + O(\beta^{-1})),$$

and logarithmic inversion yields the stated width. For the two standard moduli, $\binom{13}{6} < 3329 \le \binom{14}{7}$ and $\binom{25}{12} < 8380417 \le \binom{26}{13}$.

For the obstruction beyond constant-weight capacity, set $y_x = w_c(x)$ and use the integer identity

$$q^2 \operatorname{Var}(y) = \sum_{x<t} (y_x - y_t)^2$$

implies $\operatorname{Var}(y) \ge (q-1)/q^2$ whenever y is nonconstant, with equality exactly for one exceptional entry differing by one. Parseval and the power-mean inequality therefore imply

$$\sum_{k\neq 0} |q^{-1}\widehat{w}_c(k)|^{2d} \ge (q-1)q^{-2d}\left(\frac{\sum_{k\neq 0}|\widehat{w}_c(k)|^2}{q-1}\right)^d \ge \frac{q-1}{q^{2d}}.$$

Equality requires the one-exception table, whose nonzero Fourier coefficients all have modulus one. Its realization under $q > M_\beta$ uses exactly $q - 1 = M_\beta$ words in a largest layer and one word in an adjacent layer. □

With an unbiased selector, pairing each base word with its bitwise complement produces a width-optimal randomized construction. For arbitrary fixed bit sensitivities, every residue then has the same conditional leakage mean and zero conditional-mean product-kernel energy. Uniquely decodable randomized binary encodings with common coordinate marginals require $n \ge \lceil \log_2 q \rceil + 1$ stored bits and at least one bit of seed entropy. Antipodal pairs attain both bounds simultaneously, within the standard Module-Lattice-Based Key-Encapsulation Mechanism (ML-KEM) and Module-Lattice-Based Digital Signature Algorithm (ML-DSA) word widths.

5 The Centered-Reduction Leakage Law

At the interface between masking and machine representation, conditional product kernels and stored-word leakage obey compatible exact spectral laws.

Theorem 7 (Exact masked spectrum) *For arbitrary share leakages $w_0, \dots, w_{d-1}$ and every $k \neq 0$,*

$$K^{w_0,\dots,w_{d-1}} = q^{1-d}\bar{w}_0 * \cdots * \bar{w}_{d-1},$$
$$\widehat{K}^{w_0,\dots,w_{d-1}}(k) = q^{1-d}\prod_i \widehat{w}_i(k).$$

For identical shares, the coefficient is $q^{1-d}\widehat{w}(k)^d$.

Proof Since each prescribed sum has q^{d-1} tuples, the conditional product equals the displayed circular convolution; Fourier transformation converts convolution into the stated product, and mean subtraction affects only mode zero. □

On the negative centered half, for $0 \le n < h$, the representative of $q-1-n$ is $-(n+1)$, whose word is the bitwise complement of n, so $w_c(q-1-n) = \beta - \mathrm{HW}(n)$. Since $q-1-n = (2^m-1)-(r+n)$, one also has $w_u(q-1-n) = m - \mathrm{HW}(r+n)$; combining these two relations produces the core identity.

Theorem 8 (Exact representation law) *On the negative centered half,*

$$\Delta(q-1-n) = s + \mathrm{HW}(r+n) - \mathrm{HW}(n), \quad (3)$$

and $\Delta = 0$ on the nonnegative half. With

$$S_q(k) = \sum_{n<h} e^{2\pi \mathrm{i} kn/q},$$
$$C_{q,r}(k) = \sum_{n<h} \big(\mathrm{HW}(r+n) - \mathrm{HW}(n)\big)e^{2\pi \mathrm{i} kn/q},$$

the complete spectrum is, up to a unit phase,

$$|\widehat{\Delta}(k)| = |sS_q(k) + C_{q,r}(k)|.$$

Each carry term is a sum of $O(\log q)$ finite geometric intervals.

The resulting spectrum separates into a macroscopic signed half interval and a binary carry defect, while Equation (3) characterizes $\Delta \equiv 0$ exactly by $q = 2^m$ and $\beta = m$. Near a dyadic boundary, the carry further admits a sharp finite law. Let $q = 2^m - r$, $1 \le r < 2^{m-1}$, $L = m-1$, $M = 2^L$, $c = \lceil r/2 \rceil$, $f = \lfloor r/2 \rfloor$, and $H(t) = \sum_{0\le n<t} \mathrm{HW}(n)$, and define

$$\kappa = \sum_{a\ge 1} \tan\Big(\frac{\pi}{2^{a+1}}\Big) = 1.8098372016777826\ldots.$$

Theorem 9 (Near-dyadic carry and thin-defect law) *Put $h = M - c$, $z_k = e^{2\pi \mathrm{i} k/q}$, and*

$$\begin{aligned} F_L(z) &= \sum_{t<M} \mathrm{HW}(t) z^t, \\ A_c(z) &= \sum_{u<c} z^{M-c+u}, \\ E_c(z) &= \sum_{u<c} \mathrm{HW}(c-1-u) \\ &\quad \cdot z^{M-c+u}, \\ D_f(z) &= \sum_{v<f} (1+\mathrm{HW}(v)) z^v, \\ B_r(z) &= \sum_{v<r} \mathrm{HW}(v) z^v. \end{aligned}$$

For every Fourier mode,

$$\begin{aligned} C_{q,r}(k) &= (z_k^{-r} - 1) F_L(z_k) + L A_c(z_k) - E_c(z_k) \\ &\quad + z_k^{-r} \big(z_k^M D_f(z_k) - B_r(z_k) \big). \end{aligned} \tag{4}$$

If $2^m \geq 4r$, then at mode one it satisfies

$$\begin{aligned} &\big| C_{q,r}(1) - \big(fL - 2f - 2H(f) + \mathrm{i} r\kappa \big) \big| \\ &\quad \leq \frac{2\pi r^2}{q} \left(\frac{7}{4} L + 6 + \frac{7}{4} \lceil \log_2(r+1) \rceil \right). \end{aligned}$$

For $q = 2^m - 1$, $qD(q,m) \to \kappa$. For fixed $r \geq 2$,

$$\frac{qD(q,m)}{m-1} \longrightarrow \lfloor r/2 \rfloor.$$

If $mr/2^m \to 0$, exact width has $D(q,m) \to 0$, while every fixed positive excess width has

$$\begin{aligned} &D(q, m+s) \to s/\pi, \qquad U(q) \to 1/\pi, \\ &|\widehat{w}_c(1)| / |\widehat{w}_u(1)| \to s+1. \end{aligned}$$

Proof With $z = z_k$, shifting the first digit-sum interval in the carry yields

$$\begin{aligned} C_{q,r}(k) &= (z^{-r} - 1) \sum_{t<h} \mathrm{HW}(t) z^t \\ &\quad + z^{-r} \sum_{h \leq t < h+r} \mathrm{HW}(t) z^t \\ &\quad - z^{-r} \sum_{t<r} \mathrm{HW}(t) z^t. \end{aligned}$$

The identities $h = M - c$, $h + r = M + f$, complementation below M, and $\mathrm{HW}(M+v) = 1 + \mathrm{HW}(v)$ imply

$$\begin{aligned} \sum_{t<h} \mathrm{HW}(t) z^t &= F_L - L A_c + E_c, \\ \sum_{h \leq t < h+r} \mathrm{HW}(t) z^t &= L A_c - E_c + z^M D_f. \end{aligned}$$

Substitution establishes Equation (4) at every mode.

At $k = 1$, marking binary digits in $P_L(u,z) = \prod_{j<L} (1 + u z^{2^j})$, differentiating at $u = 1$, and using $\prod_{j<L} (1 + z^{2^j}) = (1 - z^M)/(1 - z)$ yield

$$\begin{aligned} F_L(z) &= \frac{1 - z^M}{1 - z} \left(\frac{L}{2} + \frac{\mathrm{i}}{2} T_{m,r} \right), \\ T_{m,r} &= \sum_{j<L} \tan\left(\frac{\pi 2^j}{q} \right). \end{aligned} \tag{5}$$

Using $e^{\mathrm{i}x}/(1 + e^{\mathrm{i}x}) = \frac{1}{2} + \frac{\mathrm{i}}{2} \tan(x/2)$, put

$$\begin{aligned} A &= z^{-r} \left(\sum_{t<r} z^t \right) (1 - z^M), \\ Q &= A/2 + A_c(z), \\ R &= -E_c + z^{-r} (z^M D_f - B_r). \end{aligned}$$

Equations (4)–(5) become

$$C_{q,r} = LQ + \frac{\mathrm{i}}{2} A T_{m,r} + R. \tag{6}$$

For $\theta = 2\pi/q$ and $\ell_r = \lceil \log_2(r+1) \rceil$, the identities $z^M = -e^{\mathrm{i}\theta r/2}$ and $|e^{\mathrm{i}x} - 1| \leq |x|$, combined with direct geometric comparison, imply

$$|A - 2r| \leq \tfrac{5}{2} \theta r^2, \qquad |Q - f| \leq \tfrac{7}{4} \theta r^2. \tag{7}$$

Reindexing $T_{m,r}$ from the most significant bit produces

$$T_{m,r} = \sum_{a=1}^{L} \tan\left(\frac{\pi}{2^{a+1}(1 - r/2^m)} \right).$$

For $2^m \geq 4r$, monotonicity of $\tan x/x$, the mean-value theorem, and the geometric tail imply

$$0 < T_{m,r} < 3, \qquad r |T_{m,r} - \kappa| \leq 2\theta r^2. \tag{8}$$

Since every digit sum in the boundary polynomials is at most ℓ_r, comparison of their phases with $-1, -1, +1$, respectively, yields

$$\big| R - \big(H(c) - f - H(f) - H(r) \big) \big| \leq \theta r^2 \left(\tfrac{1}{4} + \tfrac{7}{4} \ell_r \right).$$

The recurrences $H(2t) = 2H(t) + t$ and $H(2t+1) = 2H(t) + t + \mathrm{HW}(t)$ reduce the center to $-2f - 2H(f)$ in both parity cases. Subtracting $fL - 2f - 2H(f) + \mathrm{i} r\kappa$ from Equation (6) and applying Equations (7)–(8) establishes the displayed disk bound.

For $r = 1$, the disk has center $\mathrm{i}\kappa$ and radius $O(m/q)$, from which $qD(q,m) \to \kappa$. For fixed $r \geq 2$, its real center is $f(m-3) - 2H(f)$, from which the second limit follows. Under $mr/2^m \to 0$, the disk is $o(q)$ and the half-interval Riemann sum satisfies $S_q/q \to \mathrm{i}/\pi$. Only the most significant bit of the complete dyadic unsigned table contributes at mode one, and deletion of r endpoints perturbs the normalized coefficient by $O(mr/2^m)$, so $\widehat{w}_u(1)/q \to \mathrm{i}/\pi$. The identity $\widehat{\Delta}(1) = z(sS_q + C_{q,r})$ then implies the remaining limits and the ratio $s+1$. □

Across all moduli, the first-mode separation is governed by a sharp uniform threshold in the excess word width.

Theorem 10 (Sharp dominance and global leakage law) *For* $\beta_s(q) = \lceil \log_2 q \rceil + s$,

$$D(q, \beta_s(q)) > 2U(q) \quad \textit{for every } q \geq 3 \iff s \geq 3.$$

Whenever the inequality holds and $U(q) > 0$,

$$\frac{|\widehat{K}_d^{w_c}(1)|}{|\widehat{K}_d^{w_u}(1)|} \geq \left(\frac{D(q,\beta) - U(q)}{U(q)} \right)^d > 1.$$

For every $q \geq 3$, $\beta \geq m$, *and* $d \geq 2$, *centered HW has a nonzero mode with* $|\widehat{K}_d^{w_c}(k)| \geq q^{1-d}$. *Outside the unique neutral case* $(q, \beta) = (2^m, m)$, *the isolated representation component* Δ *has such a mode as well.*

Proof By the reverse triangle inequality, $|\widehat{w}_c(1)| \geq q(D - U)$, which establishes the ratio whenever $D > 2U$. For $s = 3$, set $\lambda = q/2^m$. Decomposition of the bit planes into periodic square waves produces the continuous limiting margin

$$G_3(\lambda) = |3\mathrm{i}/\pi + C_\infty(\lambda)| - 2|U_\infty(\lambda)|.$$

Dyadic truncation after ten bit planes, interval evaluation on 20,000 rational subintervals of $[1/2, 1]$, and analytic geometric tails certify $\inf_\lambda G_3(\lambda) \geq 0.2400348288848567$. Exact geometric-to-integral error bounds transfer this margin to every $m \geq 17$, with finite margin at least 0.2336562411534776. A 160-bit direct interval sweep of the exact formulas covers every $3 \leq q \leq 2^{16}$. For $s \geq 4$, the phase-independent bound

$$D(q, m + s) \geq s|S_q(1)|/q - |C_{q,r}(1)|/q$$

exceeds $2U(q)$ for every modulus and increases with s, covering all $s \geq 3$.

For $q = 2^m$, every unsigned bit below the most significant one completes full periods at mode one, and the most significant bit is the half-interval indicator. Hence $D = sU$, so strict dominance fails for $s \leq 2$. For a nonconstant integer table y, $q^2 \operatorname{Var}(y) = \sum_{x<t}(y_x - y_t)^2 \geq q - 1$. Parseval and power mean imply that some nonzero-frequency coefficient has modulus at least one. Centered HW is nonconstant for every $q \geq 3$, and Equation (3) makes Δ constant only in the stated neutral case. Theorem 7 raises these coefficients to order d. □

For the two National Institute of Standards and Technology (NIST) moduli, Table 1 reports the first-mode separation, while Figure 1 plots the certified excess-width-three margin over the full normalized modulus range.

Table 1 Exact first-mode geometry at the NIST moduli, with $B_{\min}$ equal to the certified lower bound on $|\widehat{w}_c(1)|/|\widehat{w}_u(1)|$

q	β	excess	D	U	$B_{\min}$
3329	16	4	1.479	0.2902	4.098
8380417	32	9	2.867	0.3175	8.030

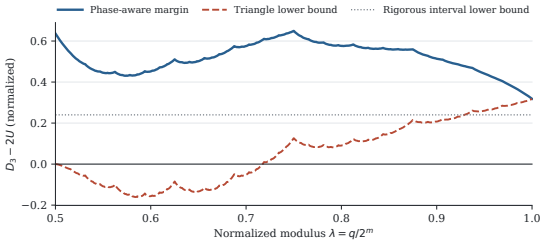


Fig. 1 Exact excess-width-three margin over the normalized modulus. The phase-aware margin is positive throughout; the dotted line is the rigorous interval lower bound used in Theorem 10

6 Robustness and Sample Complexity

For bit sensitivities $a = (a_j)$, let $w_{u,a} = \sum_j a_j \operatorname{bit}_j$ and define $w_{c,a}$ analogously. Their exact weighted representation law is

$$\Delta_a(q - 1 - n) = \sum_{j=m}^{\beta-1} a_j + \sum_{j<m} a_j \big(\operatorname{bit}_j(r + n) - \operatorname{bit}_j(n)\big).$$

Set $u_j = q^{-1}\widehat{b}_j^u(1)$ and $\delta_j = q^{-1}\widehat{(b_j^c - b_j^u)}(1)$, so that $U_a = |\sum_j a_j u_j|$ and $D_a = |\sum_j a_j \delta_j|$. Define

$$\rho_\infty(q, \beta) = \sup\{\rho \geq 0 : D_a > 2U_a \text{ whenever } \|a - \mathbf{1}\|_\infty < \rho\},$$

with $\rho_\infty = 0$ when HW itself does not satisfy the strict inequality.

Theorem 11 (Exact weighted robustness) *If* $D_{\mathbf{1}} > 2U_{\mathbf{1}}$, *the exact radius is*

$$\mathcal{Y}_\phi = \{y : \sum_j |\Re(\bar{y} v_j(\phi))| > 0\},$$

$$\rho_\infty(q,\beta) = \min_{\phi\in[0,2\pi)} \sup_{y\in\mathcal{Y}_\phi} \frac{|\Re(\bar{y}\sum_j v_j(\phi))|}{\sum_j |\Re(\bar{y}v_j(\phi))|},$$
$$v_j(\phi) = \delta_j - 2e^{\mathrm{i}\phi}u_j.$$

The supremum is attained on finitely many determinant directions. For $q = 2^m$, $\beta = m + s$,

$$\rho_\infty = 0 \ (s \le 2), \qquad \rho_\infty = (s-2)/(s+2) \ (s \ge 3).$$

Uniformly over all $q \ge 3$, *excess width three retains strict separation for every* $\alpha \ne 0$ *and* $|a_j/\alpha - 1| \le 0.077$. *For the ML-KEM and ML-DSA moduli, respectively,*

$$0.47444937499013233220 < \rho_\infty(3329, 16) < 0.47444937499013233221,$$

$$0.63672613882631207518 < \rho_\infty(8380417, 32) < 0.63672613882631207519.$$

Proof The failure set is the union, over phases ϕ, of the affine complex equations

$$\sum_j a_j v_j(\phi) = 0.$$

For $a = \mathbf{1} + e$, its ℓ_∞ distance from the origin is the least ρ for which the zonotope

$$\{\textstyle\sum_j e_j v_j : |e_j| \le \rho\}$$

contains $-\sum_j v_j$. Applying support-function separation to this planar zonotope establishes the minimax formula. In rank two, the extremal directions are orthogonal to one generator and produce the finite determinant formula; in rank one, the expression reduces to $|\sum c_j|/\sum |c_j|$.

At $q = 2^m$, only the most significant unsigned bit and the s extension bits contribute. The rank-one formula becomes $\max\{0, (s-2)/(s+2)\}$. For general q, every determinant is $A + B\cos\phi + C\sin\phi$. Under $t = \tan(\phi/2)$, the active branches are ratios of absolute quadratics. Their stationary points, intersections, rank drops, and cell endpoints form a finite algebraic candidate set. The outward-rounded evaluator certifies the uniform 0.077 box and isolates the two NIST minima in intervals of width below 10^{-43}. □

To compare the two robustness notions directly, Figure 2 plots both the uniform and phase-aware radii, with the all-modulus guarantee separated from the exact parameter-specific minima.

For transition leakage from an independent previous word P_i, let HD denote Hamming distance. The expectation $\mathbb{E}[\mathrm{HD}(y, P_i)]$ differs by a constant from weighted HW with $a_{i,j} = 1 - 2\Pr[\mathrm{bit}_j(P_i) = 1]$. The heterogeneous spectrum theorem therefore applies share by share: if each a_i lies in the certified box, the centered/unsigned kernel ratio equals the product of the local ratios when the unsigned coefficients are nonzero. A zero unsigned factor produces exact separation, whereas unbiased previous bits make every conditional-mean numerator vanish.

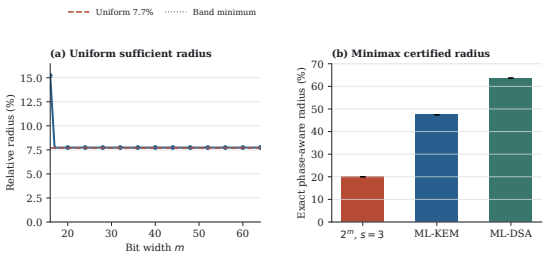


Fig. 2 Weighted Hamming-weight robustness, with the uniform all-modulus radius in the left panel and the exact phase-aware minimax radii for the dyadic obstruction and the two NIST parameter sets in the right panel

The harmonic numerator is invariant under noise when S is sampled uniformly and the uniform shares are subsequently conditioned on their sum. For $L_i = \alpha\bar{w}(X_i) + \tau N_i$, where the independent zero-mean noises have variance σ^2, and $Z_w = e^{2\pi\mathrm{i}kS/q}\prod_i L_i$,

$$|\mathbb{E}Z_w| = |\alpha|^d(|\widehat{w}(k)|/q)^d, \tag{9}$$
$$\mathrm{Var}(Z_w) = (\alpha^2\,\mathrm{Var}(w) + \tau^2\sigma^2)^d - |\mathbb{E}Z_w|^2. \tag{10}$$

Theorem 12 (Exponential-confidence harmonic test) *Let* $q \ge 3$, $d \ge 2$, *and* $2k \not\equiv 0 \pmod q$. *Let* $\mu_w = |\mathbb{E}Z_w| > 0$, *and let* V_w^* *be the larger phase-aligned real-score variance under the matched-label and independent-uniform-label laws. Divide*

$$n_w = b_w J_\delta,$$
$$b_w = \max\{1, \lceil 32V_w^*/\mu_w^2\rceil\},$$
$$J_\delta = 2\lceil(16/9)\log(2/\delta)\rceil + 1$$

independent traces into blocks and test the median block mean against $\mu_w/2$. *Both errors are at most* $\delta/2$. *If centered and unsigned mode-k coefficients are nonzero, then for fixed* δ,

$$\lim_{\tau\to\infty}\frac{n_u(\tau,\delta)}{n_c(\tau,\delta)} = \left(\frac{|\widehat{w}_c(k)|}{|\widehat{w}_u(k)|}\right)^{2d}.$$

At mode one, Theorem 10 lower-bounds this limit by $((D-U)/U)^{2d}$.

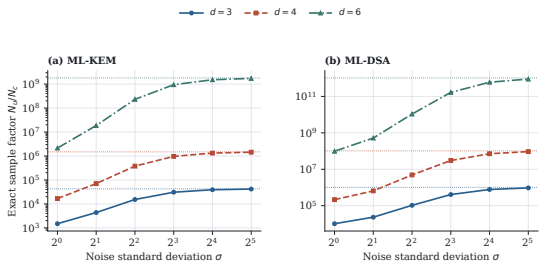


Fig. 3 Exact finite-noise centered/unsigned sample factors for orders $3, 4, 6$, with dotted lines marking their proved high-noise limits

Table 2 Paired ML-KEM synthetic rank-1 rates, where d is the share count and T the trace count. "Abs" counts membership in the two-element $\{h, -h\}$ top class; "Bayes" splits the tie

d	σ	T	full	product	OPF	abs/Bayes
2	1	100	100%	65%	43%	66%/33%
2	2	200	98%	58%	49%	81%/40.5%
3	1	200	82%	63%	28%	0%/0%

Proof After phase alignment, Z_w has a real score with mean μ_w, whereas an independent uniform label has mean zero. For one block average, Chebyshev's inequality bounds either tail at threshold $\mu_w/2$ by $4V_w^*/(b_w\mu_w^2) \le 1/8$. The block failures are independent, and Hoeffding bounds the probability that at least half fail by

$$\exp[-2J_\delta(1/2 - 1/8)^2] = \exp(-9J_\delta/32) \le \delta/2.$$

The condition $2k \not\equiv 0 \pmod q$ removes the constant-noise term from the second complex moment. At high noise, both score variances equal $\frac{1}{2}\tau^{2d}\sigma^{2d}(1 + o(1))$; the common median factor cancels, the ceilings become negligible, and Equation (9) yields the coefficient ratio to power $2d$. The stated lower bound then follows from the reverse triangle inequality. □

For a direct finite-noise comparison, Figure 3 plots the exact centered-to-unsigned sample factors together with their high-noise limits, so the finite-noise approach to the proved asymptotic ratios is visible at each reported order.

7 Algorithmic Validation

For fully paired comparisons, all methods use identical generated observations and a common complete hypothesis space. Each ML-KEM trial samples one secret coefficient, public nonzero multipliers, fresh uniform shares, and Gaussian Hamming-weight leakage. The compared methods are full likelihood, exact product likelihood, the recursive optimal prediction function (OPF), sine/cosine scores, and the absolute-value class predictor; Table 2 reports the three preselected 100-trial cells.

Exact product likelihood uses real log-magnitude convolution and a cyclic residue/sign fast Fourier transform (FFT). Independent direct convolution agrees within $1.67 \cdot 10^{-16}$; grids of 1,025 and 2,049 points produce identical rank-1 success/failure decisions on all 1,500 stored rows. Paired full-versus-product McNemar p-values for the displayed cells are $5.82 \cdot 10^{-11}$, $1.82 \cdot 10^{-12}$, and $1.57 \cdot 10^{-4}$.

Against direct $O(Tq^2)$ scoring, the finite-field ridge implementation agrees in 144 exhaustive cases. For $q = 3329$, 100 observations score all 11,082,241 ordered incomplete-NTT pairs and place the true pair at unique rank one for two and three shares in 19.34 and 25.41 seconds. For $q = 8{,}380{,}417$, linear convolution followed by modular folding ranks all hypotheses below 1 GiB. The complete reference run takes 34.86 seconds, including 16.60 seconds for the all-q likelihood pass, and small-modulus direct oracles cover two and three shares under both representations.

As a separate validation, two million tuples are used per configuration for both NIST moduli, every $d = 2, \ldots, 6$, three noise levels, and both representations. All 60 exact complex means lie in their 99% intervals. Runs under Gaussian, Laplace, Rademacher, and scaled Student-t_5 noise contain 160 million statistics and match the exact variance law. Native C sweeps cover all 3,329 ML-KEM and 8,380,417 ML-DSA residues without a representation-law mismatch.

Using public measurements, the physical validation replays the masked ML-DSA traces from the concealed inhomogeneous learning-with-errors (ILWE) study [9, 10]. On a disjoint held-out half, the sign predicate has Welch $|t| = 155.73$, standardized effect 0.923, and 90.80% balanced accuracy at a point selected on the other half, whereas a fixed seeded label permutation has $|t| = 0.207$ and 50.07%. The permutation control independently checks the sign-sensitive channel underlying the representation law.

8 Related Work

Beginning with central-reduction attacks, two's-complement sign extension was connected to masked lattice implementations [4]. Recursive optimal-correlation predictors and sine/cosine higher-order correlation power analysis (HOCPA) subsequently sharpened scalar attacks [5]. The likelihood theorem here operates on the complete finite-noise factor, and the sufficiency theorem classifies when a product-only method preserves that experiment. Related general results address optimal multivariate leakage and statistical higher-order differential power analysis (DPA) [11–13].

Beyond centered storage, representation effects have been studied for arithmetic encodings, prime-field masking, and redundant-number randomization [6, 14–16]. The section theorem optimizes all contiguous representatives at first order, the lexicographic theorem determines the high-order limit of every finite family, and the codebook theorem identifies the unrestricted deterministic finite-noise optimum.

More recently, attacks on masked ML-KEM and ML-DSA have included public-trace key recovery, sparse chosen-input methods, and profiling-device-free inference [17–21]. Contemporary countermeasures span prime-field hardware and specialized reductions [22–24]. Across these settings, the centered-reduction law determines representation neutrality, dominance, and robustness from the modulus, word width, and leakage weights.

9 Conclusion

Taken together, the results establish machine representation as part of the masked leakage channel, whose complete finite-noise experiment is efficiently computable and strictly richer than product compression throughout this nonnegative leakage class. Unsigned storage solves the global first-order contiguous-section problem, and constant-weight codebooks solve the unrestricted deterministic finite-noise problem. For centered storage, sign extension and carry produce an exact spectrum with a sharp uniform dominance threshold of three excess bits, together with classifications of the near-dyadic, weighted, and transition regimes; centered arithmetic is therefore not leakage-neutral.

Supplementary information

Online Resource 1. Complete proofs and detailed derivations for all theorems.

Online Resource 2. Source code, formal and interval certificates, exact oracles, experiment scripts, and result tables.

Statements and Declarations

Funding. No funding was received for conducting this study.

Competing Interests. The author has no relevant financial or non-financial interests to disclose.

Author contributions. José Luis Delgado conceived the study, developed the theory and software, performed the analyses, and wrote and approved the manuscript.

Data availability. All data generated by the exact, synthetic, and interval-certified analyses are included in Online Resource 2. The public physical traces analyzed in the paper are available from Zenodo under DOI 10.5281/zenodo.17291471 [10].

Code availability. Online Resource 2 contains the source code required to reproduce the reported calculations, certificates, figures, and tables, including the Matplotlib scripts that generate the figures.

Use of artificial intelligence tools. OpenAI GPT 5.5 assisted with software implementation, and language editing. The author independently checked the code and executions and takes full responsibility for the manuscript.

References

[1] Ishai, Y., Sahai, A., Wagner, D.: Private circuits: Securing hardware against probing attacks. In: Advances in Cryptology – CRYPTO 2003, pp. 463–481 (2003). https://doi.org/10.1007/978-3-540-45146-4_27

[2] Prouff, E., Rivain, M.: Masking against side-channel attacks: A formal security proof.

In: Advances in Cryptology – EUROCRYPT 2013, pp. 142–159 (2013). https://doi.org/10.1007/978-3-642-38348-9_9

[3] Duc, A., Dziembowski, S., Faust, S.: Unifying leakage models: From probing attacks to noisy leakage. In: Advances in Cryptology – EUROCRYPT 2014, pp. 423–440 (2014). https://doi.org/10.1007/978-3-642-55220-5_24

[4] Tosun, T., Moradi, A., Savas, E.: Exploiting the central reduction in lattice-based cryptography. IEEE Access **12**, 166814–166833 (2024) https://doi.org/10.1109/ACCESS.2024.3494593

[5] Tosun, T., Oswald, E., Savaş, E.: Non-profiled higher-order side-channel attacks against lattice-based post-quantum cryptography. IACR Communications in Cryptology **2**(3) (2025) https://doi.org/10.62056/a0txl8n4e

[6] Nagpal, R., Hadžić, V., Primas, R., Mangard, S.: Efficient SPA countermeasures using redundant number representation with application to ML-KEM. In: Selected Areas in Cryptography – SAC 2025, pp. 753–780 (2026). https://doi.org/10.1007/978-3-032-10536-3_28

[7] Kingston, A., Svalbe, I.: Generalised finite radon transform for $N \times N$ images. Image and Vision Computing **25**(10), 1620–1630 (2007) https://doi.org/10.1016/j.imavis.2006.03.002

[8] Graham, R.L.: On primitive graphs and optimal vertex assignments. Annals of the New York Academy of Sciences **175**(1), 170–186 (1970) https://doi.org/10.1111/j.1749-6632.1970.tb56468.x

[9] Damm, S., Fischer, A., May, A., Marzougui, S., Schwarz, L., Seidler, H., Seifert, J.-P., Thietke, J., Ulitzsch, V.Q.: Solving concealed ILWE and its application for breaking masked Dilithium. In: Advances in Cryptology – ASIACRYPT 2025, pp. 608–640 (2026). https://doi.org/10.1007/978-981-95-5096-8_19

[10] Ulitzsch, V.Q., Schwarz, L.: Attack Data for the Artifact Evaluation of Paper: Solving Concealed ILWE and its Application for Breaking Masked Dilithium. Zenodo (2025). https://doi.org/10.5281/zenodo.17291471

[11] Doget, J., Prouff, E., Rivain, M., Standaert, F.-X.: Univariate side channel attacks and leakage modeling. Journal of Cryptographic Engineering **1**(2), 123–144 (2011) https://doi.org/10.1007/s13389-011-0010-2

[12] Bruneau, N., Guilley, S., Heuser, A., Marion, D., Rioul, O.: Optimal side-channel attacks for multivariate leakages and multiple models. Journal of Cryptographic Engineering **7**(4), 331–341 (2017) https://doi.org/10.1007/s13389-017-0170-9

[13] Ding, A.A., Zhang, L., Fei, Y., Luo, P.: A statistical model for higher order DPA on masked devices. In: Cryptographic Hardware and Embedded Systems – CHES 2014, pp. 147–169 (2014). https://doi.org/10.1007/978-3-662-44709-3_9

[14] Pay, D., Standaert, F.-X.: Side-channel analysis of arithmetic encodings for post-quantum cryptography: Cautionary notes with application to kyber. In: Progress in Cryptology – AFRICACRYPT 2024, pp. 260–281 (2024). https://doi.org/10.1007/978-3-031-64381-1_12

[15] Cassiers, G., Masure, L., Momin, C., Moos, T., Standaert, F.-X.: Prime-field masking in hardware and its soundness against low-noise SCA attacks. IACR Transactions on Cryptographic Hardware and Embedded Systems **2023**(2), 482–518 (2023) https://doi.org/10.46586/tches.v2023.i2.482-518

[16] Duparc, M., Taha, M.: Improved NTT and CRT-based RNR Blinding for Side-Channel and Fault Resistant Kyber. Cryptology ePrint Archive, Paper 2025/181 (2025). https://eprint.iacr.org/2025/181

[17] Hermelink, J., Ning, K.-C., Petri, R.: Finding and protecting the weakest link. In: Advances in Cryptology – CRYPTO 2025, pp. 3–37 (2025). https://doi.org/10.1007/

978-3-032-01901-1_1

[18] Kuo, Y.-T., Takayasu, A.: Improving Correlation Power Analysis on Masked CRYSTALS-Kyber with Lattice Attack. Cryptology ePrint Archive, Paper 2026/820 (2026). https://eprint.iacr.org/2026/820

[19] Soulami, S.B.E.H., Connan, Y., Duquesne, S.: Full Secret Key Recovery of First-order Masked Crystals-Kyber implementation using multiple distinct chosen-ciphertexts. Cryptology ePrint Archive, Paper 2026/528 (2026). https://eprint.iacr.org/2026/528

[20] Wang, Y.: Profiling-Device-Free SASCA Framework for ML-KEM. Cryptology ePrint Archive, Paper 2026/981 (2026). https://eprint.iacr.org/2026/981

[21] Zhao, Y., Cheng, W., Qiao, Z., Liu, Y., Zhou, Y.: Rejection matters: Efficient non-profiling side-channel attack on ML-DSA via exploiting public templates. In: 2026 Design, Automation & Test in Europe Conference (DATE), pp. 1–7 (2026). https://doi.org/10.23919/DATE69613.2026.11539074

[22] Buschkowski, F., Höher, N., Sasdrich, P., Güneysu, T.: A Billion Hard CRYSTALS: Exploring Practical Aspects of Arithmetic Masking for PQC in Hardware. Cryptology ePrint Archive, Paper 2026/1265 (2026). https://eprint.iacr.org/2026/1265

[23] Karabulut, M., Azarderakhsh, R.: Drop-In Masked Modular Reduction for ML-DSA: Cutting Side-Channel Cost in the Root-of-Trust. Cryptology ePrint Archive, Paper 2026/711 (2026). https://eprint.iacr.org/2026/711

[24] Cassiers, G., Moos, T., Moradi, A., Müller, N., Standaert, F.-X.: Formal Verification, Integration and Physical Evaluation of Prime-Field Masking on Silicon. Cryptology ePrint Archive, Paper 2026/809 (2026). https://eprint.iacr.org/2026/809

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- ESM1.pdf
- ESM2.zip