Online Resource 1: Freshness, Restart, and Public-Rank Details for Lee-Shell Rejection Sampling

Shunqi Lu[1*]

[1*]School of Artificial Intelligence, Capital University of Economics and Business, 121 Zhangjialukou, Huaxiang, Fengtai District, Beijing, 100070, China.

Corresponding author(s). E-mail(s): 32023210246@cueb.edu.cn;

Abstract

Online Resource 1 to *Exact Coordinate-Balancing Calculus and Sharp Stability for Symmetric Lattice Cross-Covariograms*. This online resource contains the probabilistic and algebraic interfaces used by the brief rejection-sampling application in the main manuscript. It states and proves classical fresh-query and adaptive-restart bounds, records the corresponding scoped quantum-random-oracle statements, and derives a public rank–bucket certificate for the ML-DSA commitment map. None of these results is used in the proofs of the Lee-ball cross-covariogram, majorization, stability, or radial-correlation theorems.

Keywords: rejection sampling, random oracle, adaptive restart, rank certificate

1 Classical and quantum freshness and restart details

This resource first states and proves the two classical interfaces cited in the main manuscript. It then supplies the adaptive-reprogramming specialization of Grilo et al. [1], the complete-instrument restart hybrid, and the exact scope conventions used by the QROM statements.

Theorem 1.1 (Fresh-query interface) *Let $\mathcal{X}$ and $\mathcal{C}$ be finite, let $S \subset \mathbb{Z}^d$ be finite, let $g_\theta : S \to \mathcal{X}$, and let $\phi_\theta : \mathcal{C} \to \mathbb{Z}^d$. Suppose*

$$\varnothing \neq T \subseteq \bigcap_{\theta,c} (S + \phi_\theta(c)), \qquad a = |T|/|S|.$$

Let $Y \sim \mathsf{U}(S)$ and, after an arbitrary classical random-oracle table $H : \mathcal{X} \to \mathcal{C}$ with queried domain Q_H, put

$$C = H(g_\theta(Y)), \qquad Z = Y + \phi_\theta(C).$$

Put $\eta_\theta = \Pr[g_\theta(Y) \in Q_H]$ and assume $\eta_\theta < a$. If the trial accepts exactly when $Z \in T$, then

$$|\Pr[Z \in T] - a| \le \eta_\theta, \tag{1}$$

$$\mathsf{TV}\big(\mathcal{L}_\theta(C, Z \mid Z \in T), \mathsf{U}(\mathcal{C}) \otimes \mathsf{U}(T)\big) \le \eta_\theta / a. \tag{2}$$

If $q_H = |Q_H|$ and every commitment fiber has size at most M_θ, then

$$\eta_\theta \le \min\{1, q_H M_\theta / |S|\} = \min\{1, q_H 2^{-H_\infty(g_\theta(Y))}\}. \tag{3}$$

The same conclusion holds after a random adaptive classical transcript by conditioning and averaging its hit probability.

Proof of Theorem 1.1 Couple the real challenge to an independent uniform challenge. They agree unless $g_\theta(Y)$ is already in the table, an event of probability η_θ. In the fresh game, every translated source contains T, so acceptance is exactly a and the accepted law is $\mathsf{U}(\mathcal{C}) \otimes \mathsf{U}(T)$. Restricting two laws at distance η_θ to an event of ideal mass a gives the factor $1/a$. The fiber union bound proves (3). □

Theorem 1.2 (Adaptive restart composition) *Suppose the ideal accepted sublaw at every retry history is the same aQ, while the real and ideal full-trial kernels at round i couple with error at most e_i. Then the joint law of the stopping time and accepted output is within*

$$\min\left\{1, \sum_{i \ge 1} (1-a)^{i-1} e_i\right\} \tag{4}$$

of $\mathsf{Geom}(a) \otimes Q$. In particular, if each rejection adds one table entry and the conditional commitment point mass is at most 2^{-h}, the classical error is at most

$$2^{-h}\left(\frac{q_0}{a} + \frac{1-a}{a^2}\right). \tag{5}$$

Proof of Theorem 1.2 Charge the first coupling disagreement at round i. The ideal process reaches that round with probability $(1-a)^{i-1}$, which proves (4). Substitute $e_i = (q_0 + i - 1)2^{-h}$ and evaluate the geometric and arithmetico-geometric series to obtain (5). □

Corollary 1.3 (Common-target interface from adaptive reprogramming) *Use the notation of Theorem 1.1, let $\mathcal{X}$ and $\mathcal{C}$ be finite, and let a quantum algorithm make q_H queries to a random oracle $H : \mathcal{X} \to \mathcal{C}$ before one signing trial. Here q_H counts oracle calls strictly before the reprogramming instruction; parallel calls are counted separately. After those calls have*

completed, sample $Y \sim \mathsf{U}(S)$ from a source independent of the joint classical–quantum pretrial state and of H, and put $X = g_\theta(Y)$, $C = H(X)$, and $Z = Y + \phi_\theta(C)$. Suppose

$$\max_{x \in \mathcal{X}} \Pr[X = x] \le 2^{-h_\theta}. \tag{6}$$

Let $\varnothing \neq T \subseteq \bigcap_{\theta,c}(S + \phi_\theta(c))$ and $a = |T|/|S|$. Then the real pretrial-state/challenge/response state is within

$$\varepsilon_\theta^{\mathrm{Q}} := \min\left\{1, \sqrt{q_H 2^{-h_\theta}} + \frac{1}{2} q_H 2^{-h_\theta}\right\} \tag{7}$$

in trace distance of the adaptive-reprogramming game in which C is freshly uniform and the oracle is consistently set at X to C before $C = H(X)$ is read. The cited theorem permits arbitrary finitely many oracle queries after this instruction; they do not enter q_H. Consequently, for the measurement that accepts exactly when $Z \in T$,

$$|p_\theta - a| \le \varepsilon_\theta^{\mathrm{Q}}, \tag{8}$$

$$\mathsf{D}\Big(\rho^{\theta}_{E,C,Z\,|\,Z\in T}, \rho^{0}_{E} \otimes \mathsf{U}(\mathcal{C}) \otimes \mathsf{U}(T)\Big) \le \frac{\varepsilon_\theta^{\mathrm{Q}}}{a}. \tag{9}$$

Here E contains the adversary's retained quantum side information. If Y is uniform and g_θ has maximum fiber M_θ, one may take

$$h_\theta = \log_2 \frac{|S|}{M_\theta}. \tag{10}$$

Theorem 1.4 (Adaptive quantum restart composition) *At retry round i, let the real and ideal trials be binary quantum instruments on an arbitrary classical–quantum history register. Include the accept/reject flag, the accepted public output X, and the updated history in their outputs. Suppose that on every normalized history state reachable in a hybrid, the two complete instrument outputs have trace distance at most e_i. Assume further that the ideal accepted branch has the product form*

$$a\, Q_X \otimes \mathcal{E}_i(\rho), \tag{11}$$

where $a \in (0, 1]$ and the public law Q_X are independent of the input history, while $\mathcal{E}_i$ may update arbitrary private side information; the ideal rejection branch may be stateful but has trace $1 - a$.

Then the real final state of the retry count, accepted public output, and all retained side information is within

$$\min\left\{1, \sum_{i \ge 1} (1-a)^{i-1} e_i\right\} \tag{12}$$

in trace distance of the ideal restart state. In that ideal state, N is geometric with parameter a and, conditional on every $N = i$ and every retained side-information state, X has law Q_X. The same bound holds after tracing out N or any private registers.

Corollary 1.5 (Growing-query QROM retry certificate) *Repeat the trial of Corollary 1.3 with a fresh uniform mask that is independent of the current quantum history at every attempt. Suppose (6) holds with the same h after every rejection, the ideal common-target acceptance is a, and one signing-oracle invocation is atomic: after it begins, the adversary cannot interleave additional oracle queries between its internal retry attempts. If at most q_0 oracle queries*

precede the invocation and each failed attempt adds exactly the signer's one classical challenge query, then round i has $q_i = q_0 + i - 1$ prior queries. The complete restart state obeys

$$\begin{aligned}\mathsf{D}(\rho^{\mathrm{real}}_{N,E,C,Z}, \rho^{0}_{N,E} \otimes \mathsf{U}(\mathcal{C}) \otimes \mathsf{U}(T)) &\le \min\left\{1, \sum_{i\ge 1}(1-a)^{i-1}\left(\sqrt{q_i 2^{-h}} + \frac{1}{2} q_i 2^{-h}\right)\right\} \\ &\le \min\left\{1, \frac{2^{-h/2}}{a}\sqrt{q_0 + \frac{1-a}{a}} + \frac{2^{-h}}{2}\left(\frac{q_0}{a} + \frac{1-a}{a^2}\right)\right\}.\end{aligned} \tag{13}$$

Here the notation for the ideal state permits N and E to be correlated; the displayed challenge–response factor is independent of both. If a model does permit adversarial interleaving, replace $q_0 + i - 1$ by a certified bound q_i on all calls preceding attempt i and use the first, weighted sum.

Proof of Corollary 1.3 Use the notation of the cited adaptive-reprogramming game and let its single distribution p draw $(x, x') = (g_\theta(Y), Y)$. Thus $R = 1$, its marginal point probability satisfies $p_{\max} = \max_x \Pr[X = x] \le 2^{-h_\theta}$, and its additional output $x' = Y$ retains all source information used to form Z. Exactly q_H oracle calls precede the reprogramming instruction, so the cited equation (2) has $\widehat{q}_1 = q_H$ and gives

$$\sqrt{\widehat{q}_1 p_{\max}} + \frac{1}{2}\widehat{q}_1 p_{\max} \le \sqrt{q_H 2^{-h_\theta}} + \frac{1}{2} q_H 2^{-h_\theta}.$$

In its game $b = 0$ the oracle is unchanged and the subsequent read returns the real $H(X)$; in game $b = 1$ the value at X is replaced by a fresh uniform C before that read. Theorem 1 is information-theoretic and explicitly allows the correlated additional output x'. Appending an optimal final binary measurement therefore shows, by the operational characterization of trace distance, that the complete retained state E, C, Z obeys the same bound. This also explains why no unmentioned conversion from event probability or Bures distance is present in (7).

In that game, the same counting argument as in Theorem 1.1 gives acceptance probability a and accepted classical law $\mathsf{U}(\mathcal{C}) \otimes \mathsf{U}(T)$. Apply the accept/reject measurement while retaining its classical flag. Trace distance is contractive, controls the difference of the two acceptance probabilities, and decomposes over the two flag blocks. Hence the trace norm of the accepted blocks plus the acceptance-mass difference is at most twice $\varepsilon^{\mathrm{Q}}_\theta$. Normalizing through the ideal mass a gives (9). Finally, uniformity of Y makes the largest point mass of $g_\theta(Y)$ equal to $M_\theta/|S|$, proving (10). □

Proof of Theorem 1.4 Use hybrids whose first $i - 1$ trial instruments are ideal and whose suffix is real, and replace the first real instrument in that suffix. Before round i the two adjacent hybrids therefore have the same live subnormalized history state, of trace $(1 - a)^{i-1}$. Homogeneity of the trace norm weights the conditional instrument error e_i by this trace. Every later real instrument is completely positive and trace-nonincreasing, so trace distance cannot increase after that replacement. The triangle inequality over rounds gives (12). Equation (11) makes ideal survival geometric and factors Q_X from the retry count and retained state. Partial trace proves the marginal statements. □

Proof of Corollary 1.5 Corollary 1.3 makes the complete round-i instrument error at most $e_i = \sqrt{q_i 2^{-h}} + q_i 2^{-h}/2$. Apply Theorem 1.4. For the closed upper bound, normalize the

weights $a(1-a)^k$ to the geometric law of a failure count K, use concavity of the square root, and then evaluate the linear series:

$$\sum_{k\geq 0}(1-a)^k\sqrt{q_0+k}=\frac{1}{a}\mathbb{E}\sqrt{q_0+K}\leq\frac{1}{a}\sqrt{q_0+\frac{1-a}{a}},\qquad \mathbb{E}K=\frac{1-a}{a}.$$

□

Remark 1.6 (Exact scope of the QROM statement) Corollary 1.3 and Theorem 1.4 are accepted-state comparison theorems, not Fiat–Shamir extraction theorems. Their independence hypothesis is satisfied by a private classical mask source or by a logically independent typed mask oracle whose output is not entangled with the adversary's pretrial state. The finite-domain hypothesis in Corollary 1.3 is literal. A variable-length XOF model must use an injective type-and-length encoding into one finite security-parameter domain and count every prior call to that global oracle in q_H; distinct types are independent only when the model declares separate oracles.

Although the literal FIPS 204 input lengths are disjoint, a single global quantum-accessible SHAKE/XOF model also has to control adversarial side information about the private 128- and 66-byte source restrictions. The classical source-hit certificate used in the main manuscript does not establish that quantum conditional-min-entropy statement. We therefore instantiate the QROM theorems only in the independent-mask-source experiment and do not relabel them as a proof of the complete ML-DSA signing algorithm. The register E denotes retained adversary-accessible side information, not the hidden internal oracle database. The cited adaptive-reprogramming game already permits a finite query-bounded continuation after the programming instruction; the displayed product form describes the accepted boundary before such a continuation.

Figure 1 summarizes the two restart interfaces. Its rank-to-entropy input is derived in Section 2; the figure is an overview, not an additional reduction.

2 Public-rank and generating-function details

This resource derives the general rank–bucket bound, specializes it to the public ML-DSA commitment map [2], and records the finite-field rank generating function used for the ideal-XOF tail. The main text uses only the resulting point-mass certificate.

Theorem 2.1 (Rank–bucket fiber certificate) *Let $L:\mathbb{F}_q^D\to\mathbb{F}_q^K$ have rank r, let $\Omega\subseteq\mathbb{F}_q^D$, and let each coordinate of a quantizer Q have fibers of size at most B. For every injective encoding E,*

$$\max_x\,|(E\circ Q\circ L)^{-1}(x)\cap\Omega|\leq\min\{|\Omega|,B^rq^{D-r}\}. \tag{14}$$

Consequently a uniform $Y\in\Omega$ satisfies

$$H_\infty(E(Q(LY)))\geq\log_2|\Omega|-\log_2\min\{|\Omega|,B^rq^{D-r}\}. \tag{15}$$

Proof Choose r independent output coordinates. Fixing their quantized values permits at most B^r exact outputs, and each exact output has a kernel coset of size q^{D-r}. Intersect with Ω and take logarithms. □

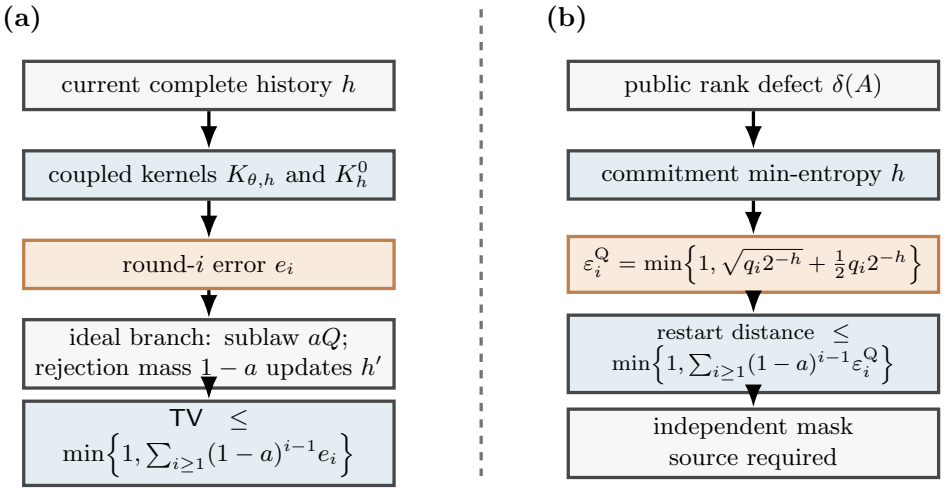

Fig. 1 Classical and QROM restart interfaces. (a) At history h, the real and ideal full-trial kernels are compared within e_i; only the ideal accepted branch has sublaw aQ. The restart hybrid weights the round-i discrepancy by the ideal survival probability $(1-a)^{i-1}$. (b) The public rank defect supplies commitment min-entropy h; q_i bounds the prior quantum queries in round i, adaptive reprogramming gives $\varepsilon_i^{\mathrm{Q}}$, and the same survival weights control the restarted state. The QROM branch requires an independent mask source and is not a reduction for the complete ML-DSA signing algorithm.

For ML-DSA, Y is uniform on $I_{\gamma_1}^D$ and the encoded commitment is

$$g_{A,\mu}(y) = \mu \,\|\, \mathsf{w1Encode}(\mathsf{HighBits}(Ay)). \tag{16}$$

For this map, write $q = 8380417$, $n = 256$, $D = n\ell$, and let $\widehat{A}[j] \in \mathbb{F}_q^{k\times\ell}$ be the public NTT lanes. Define

$$\delta(A) = \sum_{j=0}^{n-1} \big(\ell - \operatorname{rank}_{\mathbb{F}_q} \widehat{A}[j]\big). \tag{17}$$

Proposition 2.2 (ML-DSA commitment-fiber certificate) *Let $B_{\gamma_2} = 2\gamma_2+1$. For every public matrix A and message representative μ, the largest fiber $M_{A,\mu}$ of (16) and its min-entropy obey*

$$M_{A,\mu} \leq \min\{(2\gamma_1)^D, B_{\gamma_2}^{D-\delta(A)} q^{\delta(A)}\}, \tag{18}$$

$$h_{A,\mu} \geq \max\left\{0, D\log_2\frac{2\gamma_1}{B_{\gamma_2}} - \delta(A)\log_2\frac{q}{B_{\gamma_2}}\right\}. \tag{19}$$

The certificate is deterministic and checkable from the public key.

Proof of Proposition 2.2 FIPS 204 has $2\gamma_2 \mid q-1$. Directly from its DECOMPOSE exceptional branch, every HIGHBITS value has either $2\gamma_2$ or $2\gamma_2+1$ residue preimages; hence the maximum bucket size is B_{γ_2}. Let $\mathcal{N}_\ell : R_q^\ell \to \mathbb{F}_q^{n\ell}$ and $\mathcal{N}_k : R_q^k \to \mathbb{F}_q^{nk}$ be the vector NTT isomorphisms. If

Table 1 Checklist for the ML-DSA commitment-fiber instantiation.

Interface	Object used in the proof	Check
mask embedding	$I_{\gamma_1}^D \hookrightarrow \mathbb{F}_q^D$	injective because $2\gamma_1 < q$
coefficient linear map	$L_A : R_q^\ell \to R_q^k,\ y \mapsto Ay$	an $\mathbb{F}_q$-linear map before HighBits
NTT rank transfer	$\mathcal{N}_k L_A \mathcal{N}_\ell^{-1} = \bigoplus_j \widehat{A}[j]$	$\mathcal{N}_\ell, \mathcal{N}_k$ are bijections, so ranks are equal
scalar quantizer	coefficientwise HighBits	every bucket has at most $B_{\gamma_2} = 2\gamma_2+1$ residues
protocol encoding	$\mu \Vert \mathsf{w1Encode}(\mathsf{HighBits}(Ay))$	fixed prefix and injective fixed-width encoding preserve fibers

$L_A(y) = Ay$ is written in the coefficient basis, then the defining lane-wise NTT multiplication gives the exact conjugacy

$$\mathcal{N}_k \circ L_A \circ \mathcal{N}_\ell^{-1} = \bigoplus_{j=0}^{n-1} \widehat{A}[j]. \tag{20}$$

The two outer maps are bijections. Therefore rank invariance under composition with isomorphisms and rank additivity for a direct sum give

$$\begin{aligned} \operatorname{rank}_{\mathbb{F}_q} L_A &= \sum_{j=0}^{n-1} \operatorname{rank}_{\mathbb{F}_q} \widehat{A}[j] \\ &= \sum_{j=0}^{n-1} \big(\ell - (\ell - \operatorname{rank} \widehat{A}[j])\big) = n\ell - \delta(A) = D - \delta(A). \end{aligned}$$

This equality concerns only the linear rank. We now return to the coefficient basis, where FIPS applies HighBits coefficientwise, and apply Theorem 2.1 with the displayed rank and bucket bound. The mask interval embeds injectively in R_q^ℓ because $2\gamma_1 < q$. Finally, prepending the fixed μ and applying the injective fixed-width w1Encode neither merge nor enlarge a fiber. These five steps prove (18); taking logarithms and clipping by the source size proves (19). □

Ideal-XOF defect distribution

If

$$\mathcal{N}_{k,\ell}(r) = \prod_{i=0}^{r-1} \frac{(q^k - q^i)(q^\ell - q^i)}{q^r - q^i} \tag{21}$$

denotes the number of rank-r matrices in $\mathbb{F}_q^{k\times\ell}$, then in the ideal-XOF model the total NTT-lane defect has the exact probability generating function

$$\mathbb{E}[X^\delta] = \left(\sum_{e=0}^{\ell} \frac{\mathcal{N}_{k,\ell}(\ell - e)}{q^{k\ell}} X^e \right)^n. \tag{22}$$

Indeed, the n NTT lanes are independent uniform rectangular matrices. Applying the standard finite-field rank count to each lane and multiplying the one-lane generating functions proves (22).

3 Scoped ML-DSA strict-z illustration

As a deliberately scoped illustration of the interfaces above, consider only the strict z predicate of ML-DSA [2]. Put $q = 8380417$, $n = 256$, $D = n\ell$, $\beta = \tau\eta$, and $B_{\gamma_2} = 2\gamma_2 + 1$. With the public NTT-rank defect $\delta(A)$ of (17), Proposition 2.2 bounds the commitment point mass by

$$\epsilon_A := \min\left\{1, \left(\frac{B_{\gamma_2}}{2\gamma_1}\right)^{D-\delta(A)} \left(\frac{q}{2\gamma_1}\right)^{\delta(A)}\right\}. \tag{23}$$

The literal strict-z common interval gives

$$G_z = \left(\frac{2(\gamma_1 - \beta) - 1}{2\gamma_1}\right)^D. \tag{24}$$

In the separated classical-oracle model, Theorems 1.1 and 1.2 yield, for the counterfactual process that retries until this predicate alone accepts,

$$\mathsf{TV}_{\mathrm{restart}} \le \epsilon_A \left(\frac{q_0}{G_z} + \frac{1 - G_z}{G_z^2}\right). \tag{25}$$

Table 2 Scoped ML-DSA strict-z illustration in the separated classical model with $q_0 \le 2^{64}$. The budget is the largest rank defect for which (25) is below 2^{-128}; ACVP entries are audited keys, not an all-key claim.

set	D	β	γ_1	G_z	defect budget	ACVP $\delta = 0$
ML-DSA-44	1024	78	2^{17}	.5414714311	51	25/25
ML-DSA-65	1280	196	2^{19}	.6188909137	272	25/25
ML-DSA-87	1792	120	2^{19}	.6623821841	400	25/25

All 75 audited keys have zero defect; the corresponding restart exponents are at least 128.578, 129.109, and 129.928 bits. This calculation concerns only the strict z predicate: it neither analyzes the conjunction of Algorithm 7 filters nor constitutes a security reduction.

References

[1] Grilo, A.B., Hövelmanns, K., Hülsing, A., Majenz, C.: Tight adaptive reprogramming in the QROM. In: Advances in Cryptology—ASIACRYPT 2021, Part I.

Lecture Notes in Computer Science, vol. 13090, pp. 637–667. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-92062-3_22

[2] National Institute of Standards and Technology: Module-lattice-based digital signature standard. Federal Information Processing Standards Publication FIPS 204, U.S. Department of Commerce (2024). https://doi.org/10.6028/NIST.FIPS.204