

Masking-Friendly Lattice Schemes and Lattice-Friendly Masking Schemes

Thomas Prest (joint work with/by Rafael del Pino, Mélissa Rossi, Markku Saarinen & Shuichi Katsumata)



April 3, 2025

Why this Talk?

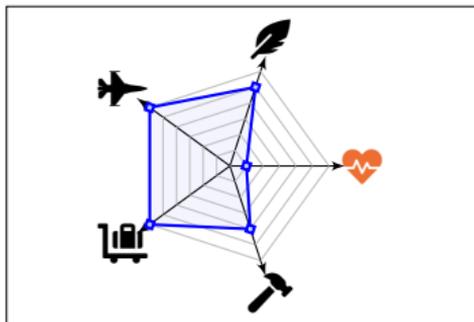
Observation: masking and post-quantum standards have poor compatibility.

- 1 Can we design **lattice-based cryptosystems** more suitable for **masking**?
- 2 Can we design **masking schemes** more suitable for **lattice cryptosystems**?

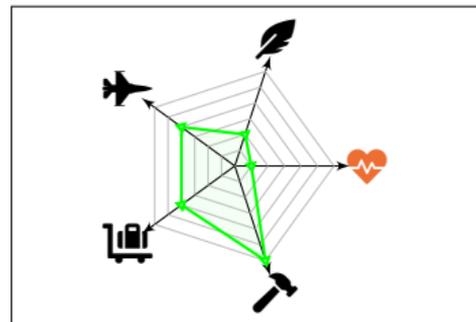
Lattice Schemes & Masking



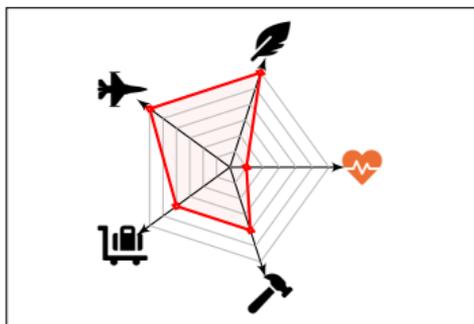
ML-DSA



SLH-DSA



FN-DSA

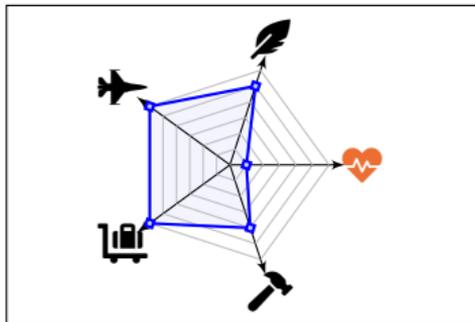


NIST PQC standards, selected in 2022, strike a balance between several criteria.

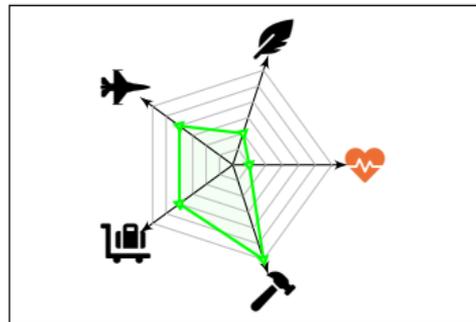
But what about :

Side-channel protection?

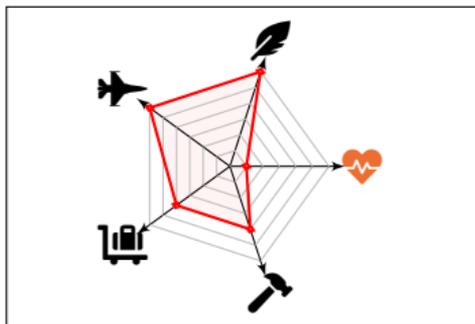
ML-DSA



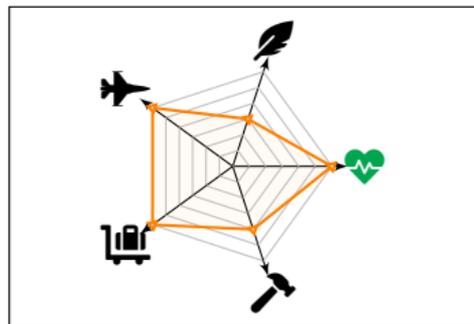
SLH-DSA



FN-DSA



Raccoon (2023)



Size



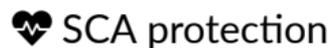
Speed



Portability



Assumptions

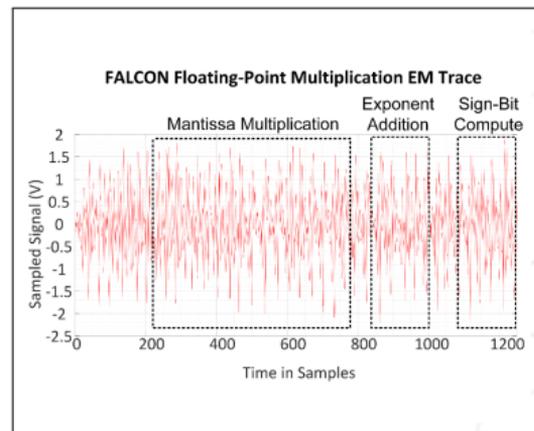
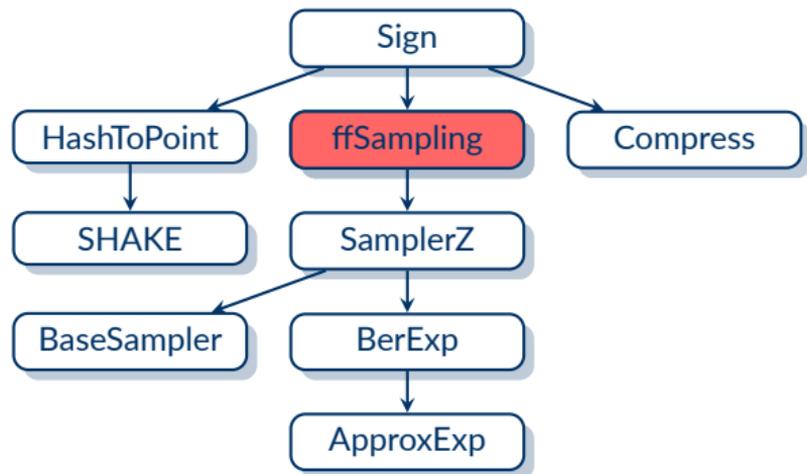


SCA protection

In Falcon, a signature sig is distributed as a Gaussian.

The signing key sk should remain private.

The power consumption leaks information about the dot product $\langle \text{sig}, \text{sk} \rangle$, or sk itself.



Learning sk directly

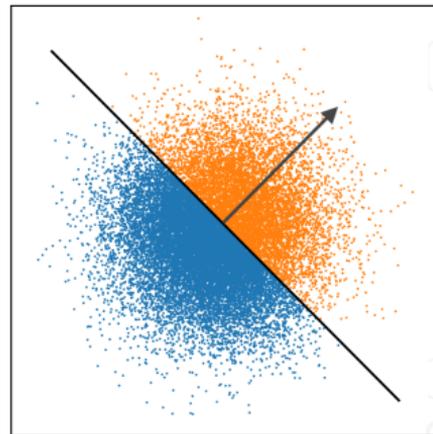
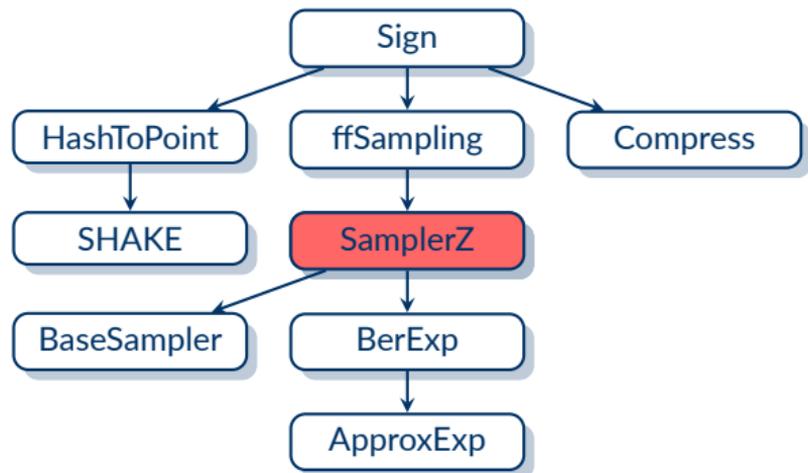
Figure 1: Flowchart of the signature

¹FALCON Down: Breaking FALCON Post-Quantum Signature Scheme through Side-Channel Attacks [KA21]

In Falcon, a signature sig is distributed as a Gaussian.

The signing key sk should remain private.

The power consumption leaks information about the dot product $\langle \text{sig}, \text{sk} \rangle$, or sk itself.



Filtering $\langle \text{sig}, \text{sk} \rangle > 0$

Figure 1: Flowchart of the signature

Dilithium-Sign

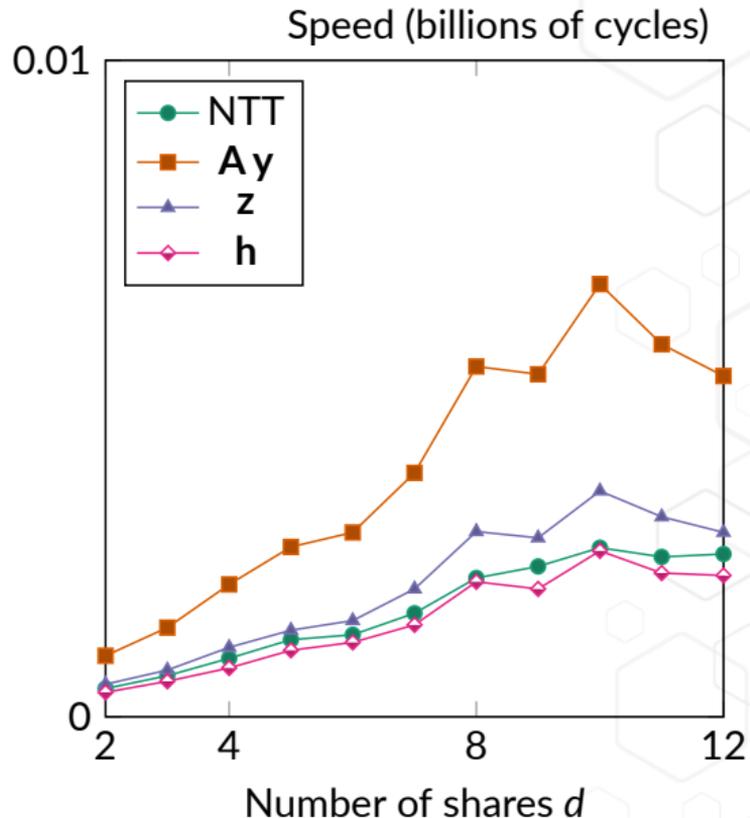
- 1 Sample $\mathbf{y} \leftarrow \text{Uniform}(S)$
- 2 $\mathbf{w} := \mathbf{A} \cdot \mathbf{y}$
- 3 $\mathbf{w}_0, \mathbf{w}_1 := \text{Decompose}(\mathbf{w})$
- 4 $\mathbf{c} := H(\mathbf{w}_1, \text{msg})$
- 5 $\mathbf{z} := \mathbf{y} + \mathbf{s}_1 \cdot \mathbf{c}$
- 6 $\tilde{\mathbf{r}} := \mathbf{w}_0 - \mathbf{s}_0 \cdot \mathbf{c}$
- 7 If $\|\mathbf{z}\|_\infty$ or $\|\tilde{\mathbf{r}}\|_\infty$ are too large,
goto 1
- 8 $\mathbf{h} := \mathbf{w}_1 - \lfloor \mathbf{A} \cdot \mathbf{z} - \mathbf{t} \cdot \mathbf{c} \rfloor_k$
- 9 Output sig = (c, z, h)

Observations:

- Some operations don't need to be masked (or conjectured to)
- Some operations are linear and are therefore easy to mask
- Three operations require mask conversions (overhead: $O(d^2 \log q)$):
 - 1 Sampling
 - 3 Decomposition
 - 6 Rejection sampling

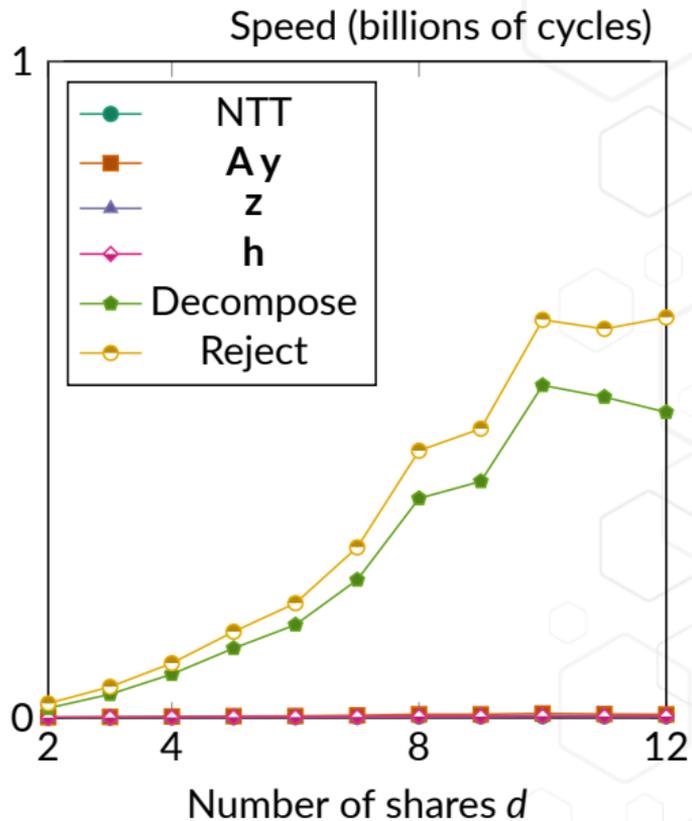
Dilithium-Sign

- 1 Sample $\mathbf{y} \leftarrow S$
- 2 $\mathbf{w} := \mathbf{A} \cdot \mathbf{y}$ ▷ $\tilde{O}(d)$
- 3 $\mathbf{w}_0, \mathbf{w}_1 := \text{Decompose}(\mathbf{w})$
- 4 $c := H(\mathbf{w}_1, \text{msg})$ ▷ No mask
- 5 $\mathbf{z} := \mathbf{y} + \mathbf{s}_1 c$ ▷ $\tilde{O}(d)$
- 6 $\tilde{\mathbf{r}} := \mathbf{w}_0 - \mathbf{s}_0 \cdot c$ ▷ $\tilde{O}(d)$
- 7 If $\|\mathbf{z}\|_\infty$ or $\|\tilde{\mathbf{r}}\|_\infty$ are too large, goto 1
- 8 $\mathbf{h} := \mathbf{w}_1 - \lfloor \mathbf{A} \cdot \mathbf{z} - \mathbf{t} \cdot c \rfloor_k$ ▷ No mask
- 9 Output sig = $(c, \mathbf{z}, \mathbf{h})$



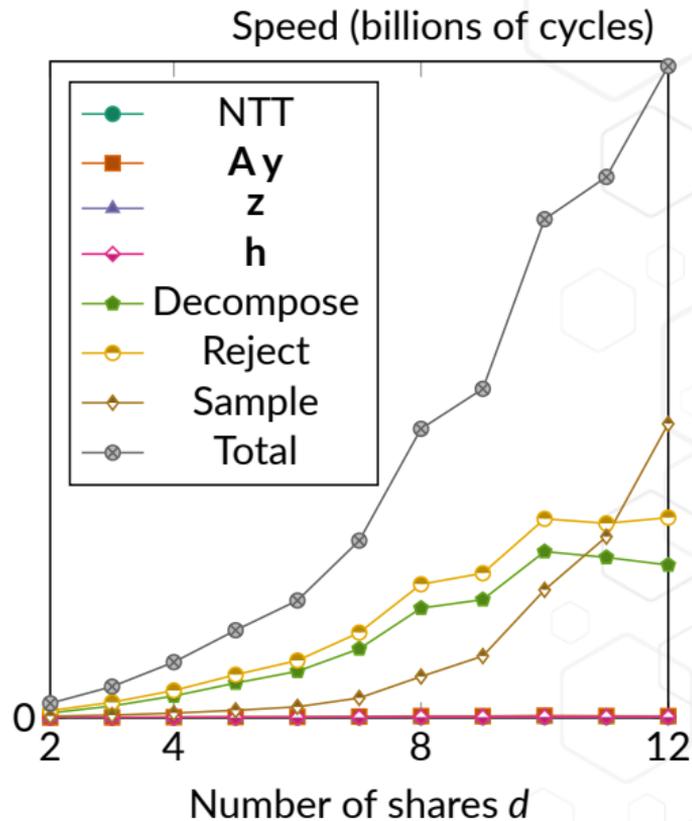
Dilithium-Sign

- 1 Sample $\mathbf{y} \leftarrow S$
- 2 $\mathbf{w} := \mathbf{A} \cdot \mathbf{y}$ $\triangleright \tilde{O}(d)$
- 3 $\mathbf{w}_0, \mathbf{w}_1 := \text{Decompose}(\mathbf{w})$ $\triangleright O(d^2 \log q)$
- 4 $c := H(\mathbf{w}_1, \text{msg})$ \triangleright No mask
- 5 $\mathbf{z} := \mathbf{y} + \mathbf{s}_1 c$ $\triangleright \tilde{O}(d)$
- 6 $\tilde{\mathbf{r}} := \mathbf{w}_0 - \mathbf{s}_0 \cdot c$ $\triangleright \tilde{O}(d)$
- 7 If $\|\mathbf{z}\|_\infty$ or $\|\tilde{\mathbf{r}}\|_\infty$ are too large, goto 1 $\triangleright O(d^2 \log q)$
- 8 $\mathbf{h} := \mathbf{w}_1 - \lfloor \mathbf{A} \cdot \mathbf{z} - \mathbf{t} \cdot c \rfloor_k$ \triangleright No mask
- 9 Output sig = $(c, \mathbf{z}, \mathbf{h})$



Dilithium-Sign

- 1 Sample $\mathbf{y} \leftarrow S$ $\triangleright O(d^2 \log q)$
- 2 $\mathbf{w} := \mathbf{A} \cdot \mathbf{y}$ $\triangleright \tilde{O}(d)$
- 3 $\mathbf{w}_0, \mathbf{w}_1 := \text{Decompose}(\mathbf{w})$ $\triangleright O(d^2 \log q)$
- 4 $c := H(\mathbf{w}_1, \text{msg})$ \triangleright No mask
- 5 $\mathbf{z} := \mathbf{y} + \mathbf{s}_1 c$ $\triangleright \tilde{O}(d)$
- 6 $\tilde{\mathbf{r}} := \mathbf{w}_0 - \mathbf{s}_0 \cdot c$ $\triangleright \tilde{O}(d)$
- 7 If $\|\mathbf{z}\|_\infty$ or $\|\tilde{\mathbf{r}}\|_\infty$ are too large, goto 1 $\triangleright O(d^2 \log q)$
- 8 $\mathbf{h} := \mathbf{w}_1 - \lfloor \mathbf{A} \cdot \mathbf{z} - \mathbf{t} \cdot c \rfloor_k$ \triangleright No mask
- 9 Output sig = $(c, \mathbf{z}, \mathbf{h})$



Raccoon

Raccoon.Keygen() \rightarrow sk, vk

- 1 $vk = [\mathbf{A} \ 1] \cdot sk$, for sk short.

Schnorr.Keygen() \rightarrow sk, vk

- 1 $vk = g^{sk}$, for sk uniform.

Raccoon.Sign(sk, msg) \rightarrow sig

- 1 Sample a short \mathbf{r}
- 2 $\mathbf{w} = [\mathbf{A} \ 1] \cdot \mathbf{r}$
- 3 $c = H(\mathbf{w}, \text{msg})$
- 4 $\mathbf{z} = \mathbf{r} + c \cdot sk$
- 5 Output sig = (c, z)

Schnorr.Sign(sk, msg) \rightarrow sig

- 1 Sample r
- 2 $w = g^r$
- 3 $c = H(w, \text{msg})$
- 4 $z = r + c \cdot sk$
- 5 Output sig = (c, z)

Raccoon.Verify(vk, msg, sig)

- 1 $\mathbf{w}' = [\mathbf{A} \ 1] \cdot \mathbf{z} - c \cdot vk$
- 2 Assert $H(\mathbf{w}', \text{msg}) = c$
- 3 Assert z is short

Schnorr.Verify(vk, msg, sig)

- 1 $w' = g^z \cdot vk^{-c}$
- 2 Assert $H(w', \text{msg}) = c$

Raccoon.Keygen() \rightarrow sk, vk

1 vk = $[\mathbf{A} \ 1] \cdot \text{sk}$, for sk short.

Raccoon.Sign(sk, msg) \rightarrow sig

1 Sample a short \mathbf{r}

2 $\mathbf{w} = [\mathbf{A} \ 1] \cdot \mathbf{r}$

3 $c = H(\mathbf{w}, \text{msg})$

4 $\mathbf{z} = \mathbf{r} + c \cdot \text{sk}$

5 Output sig = (c, z)

Raccoon.Verify(vk, msg, sig)

1 $\mathbf{w}' = [\mathbf{A} \ 1] \cdot \mathbf{z} - c \cdot \text{vk}$

2 Assert $H(\mathbf{w}', \text{msg}) = c$

3 Assert z is short

Security: Raccoon is EUF-CMA assuming:

- 1 **Hint-MLWE** [KLSS23] (next slide)
 - > Implied by lack of rejection sampling
 - > Ensures uniformity of the public key
- 2 **Self-target MSIS** [KLS18]
 - > Unforgeability

Rounding: we can round vk and w:

- ✓ Reduces the size of vk and sig
- ✓ Zero impact on Hint-MLWE
- ✓ Minor impact on unforgeability
- ✓ Not a sensitive information
 - > Will not need to be masked

(Hint-)MLWE [KLSS23]

It is difficult to distinguish both distributions:

$$\left\{ (\mathbf{A}, \mathbf{b}) \mid \mathbf{A} \leftarrow \mathcal{R}_q^{k \times \ell}, \text{sk} \leftarrow \chi_{\text{sk}}, \mathbf{b} := [\mathbf{A} \quad \mathbf{I}] \cdot \text{sk} \right\}$$

$$\left\{ (\mathbf{A}, \mathbf{b}) \mid \mathbf{A} \leftarrow \mathcal{R}_q^{k \times \ell}, \text{sk} \leftarrow \chi_{\text{sk}}, \mathbf{b} \leftarrow \mathcal{R}_q^k \right\}$$

In Hint-MLWE, the adversary is additionally given Q “hints” of the shape:

$$(c_i, \mathbf{z}_i \leftarrow c_i \cdot \text{sk} + \mathbf{r}_i), \quad \text{where } c_i \leftarrow \mathcal{C}, \mathbf{r}_i \leftarrow \chi_r$$

(Hint-)MLWE [KLSS23]

It is difficult to distinguish both distributions:

$$\left\{ (\mathbf{A}, \mathbf{b}) \mid \mathbf{A} \leftarrow \mathcal{R}_q^{k \times \ell}, \text{sk} \leftarrow \chi_{\text{sk}}, \mathbf{b} := [\mathbf{A} \quad \mathbf{I}] \cdot \text{sk} \right\}$$

$$\left\{ (\mathbf{A}, \mathbf{b}) \mid \mathbf{A} \leftarrow \mathcal{R}_q^{k \times \ell}, \text{sk} \leftarrow \chi_{\text{sk}}, \mathbf{b} \leftarrow \mathcal{R}_q^k \right\}$$

In Hint-MLWE, the adversary is additionally given Q “hints” of the shape:

$$(c_i, \mathbf{z}_i \leftarrow c_i \cdot \text{sk} + \mathbf{r}_i), \quad \text{where } c_i \leftarrow \mathcal{C}, \mathbf{r}_i \leftarrow \chi_r$$

Attack on Hint-MLWE

Assume $\forall i \in [Q], \|c_i\|^2 = \omega$. If we note $c^*(x) = c(x^{-1})$, we can recover sk by constructing this accumulator:

$$\begin{aligned} \text{acc} &= \sum_i c_i^* \cdot \mathbf{z}_i \\ &= \sum_i c_i^* c_i \cdot \text{sk} + \sum_i c_i^* \cdot \mathbf{r}_i \\ &\approx Q \cdot \omega \cdot \text{sk} + O(\sqrt{Q \cdot \omega} \cdot \|\mathbf{r}\|) \end{aligned}$$

If $\|\mathbf{r}\| = o(\sqrt{Q \cdot \omega})$, rounding acc to the closest multiple of $Q \cdot \omega$ gives sk.

(Hint-)MLWE [KLSS23]

It is difficult to distinguish both distributions:

$$\left\{ (\mathbf{A}, \mathbf{b}) \mid \mathbf{A} \leftarrow \mathcal{R}_q^{k \times \ell}, \text{sk} \leftarrow \chi_{\text{sk}}, \mathbf{b} := [\mathbf{A} \quad \mathbf{I}] \cdot \text{sk} \right\}$$

$$\left\{ (\mathbf{A}, \mathbf{b}) \mid \mathbf{A} \leftarrow \mathcal{R}_q^{k \times \ell}, \text{sk} \leftarrow \chi_{\text{sk}}, \mathbf{b} \leftarrow \mathcal{R}_q^k \right\}$$

In Hint-MLWE, the adversary is additionally given Q “hints” of the shape:

$$(c_i, \mathbf{z}_i \leftarrow c_i \cdot \text{sk} + \mathbf{r}_i), \quad \text{where } c_i \leftarrow \mathcal{C}, \mathbf{r}_i \leftarrow \chi_r$$

Attack on Hint-MLWE

Assume $\forall i \in [Q], \|c_i\|^2 = \omega$. If we note $c^*(x) = c(x^{-1})$, we can recover sk by constructing this accumulator:

$$\begin{aligned} \text{acc} &= \sum_i c_i^* \cdot \mathbf{z}_i \\ &= \sum_i c_i^* c_i \cdot \text{sk} + \sum_i c_i^* \cdot \mathbf{r}_i \\ &\approx Q \cdot \omega \cdot \text{sk} + O(\sqrt{Q \cdot \omega} \cdot \|\mathbf{r}\|) \end{aligned}$$

If $\|\mathbf{r}\| = o(\sqrt{Q \cdot \omega})$, rounding acc to the closest multiple of $Q \cdot \omega$ gives sk .

Security reduction, simplified [KLSS23, DKM⁺24]

If \mathbf{s} and \mathbf{r}_i are sampled from gaussians of standard deviation σ_{sk} and σ_r , then:

$$\text{Hint-MLWE}_{\mathcal{R}_q, k, \ell, \sigma_{\text{sk}}, \sigma_r, Q} \geq \text{MLWE}_{\mathcal{R}_q, k, \ell, \sigma_0}, \quad \text{where } \frac{1}{\sigma_0^2} \approx 2 \left(\frac{1}{\sigma_{\text{sk}}^2} + \frac{Q \cdot \omega}{\sigma_r^2} \right) \quad (1)$$

Raccoon.Sign(sk, msg) \rightarrow sig

- 1 Sample a short \mathbf{r}
- 2 $\mathbf{w} = [\mathbf{A} \ 1] \cdot \mathbf{r}$
- 3 $c = H(\mathbf{w}, \text{msg})$
- 4 $\mathbf{z} = \mathbf{r} + c \cdot \text{sk}$
- 5 Output sig = (c, \mathbf{z})

Starting point is “Schnorr over lattices”:

- ✓ No Rejection sampling
- ✓ Steps 2 and 4 are easy to mask
- ✓ Steps 3 does not need to be masked (no conjecture!)
- ? What about Sampling (step 1)?

Raccoon.Sign(sk, msg) \rightarrow sig

- 1 Sample a short \mathbf{r}
- 2 $\mathbf{w} = [\mathbf{A} \ 1] \cdot \mathbf{r}$
- 3 $c = H(\mathbf{w}, \text{msg})$
- 4 $\mathbf{z} = \mathbf{r} + c \cdot \text{sk}$
- 5 Output sig = (c, z)

MaskSign([[sk]], vk, msg) \rightarrow sig

- 1 $[\mathbf{r}] = [\mathbf{0}]$
- 2 For $i \in [\text{rep}]$:
 - 1 $[\mathbf{r}_i] = (r_{i,1}, \dots, r_{i,d}) \leftarrow \chi_r^d$
 - 2 $[\mathbf{r}] = [\mathbf{r}] + [\mathbf{r}_i]$
 - 3 Refresh($[\mathbf{r}]$)
- 3 $[\mathbf{w}] = [\mathbf{A} \ \mathbf{I}] \cdot [\mathbf{r}]$
- 4 Refresh($[\mathbf{w}]$)
- 5 $\mathbf{w} = \text{Decode}([\mathbf{w}])$
- 6 $c = H(\mathbf{w}, \text{msg}, \text{vk})$
- 7 $[\mathbf{z}] = [\text{sk}] \cdot c + [\mathbf{r}]$
- 8 Refresh($[\mathbf{z}], [\text{sk}]$)
- 9 $\mathbf{z} = \text{Decode}([\mathbf{z}])$
- 10 Output sig = (c, z)

Raccoon.Sign(sk, msg) \rightarrow sig

- 1 Sample a short \mathbf{r}
- 2 $\mathbf{w} = [\mathbf{A} \ \mathbf{1}] \cdot \mathbf{r}$
- 3 $\mathbf{c} = H(\mathbf{w}, \text{msg})$
- 4 $\mathbf{z} = \mathbf{r} + \mathbf{c} \cdot \text{sk}$
- 5 Output sig = (c, z)

We note $[[x]]$ a d -sharing of x .

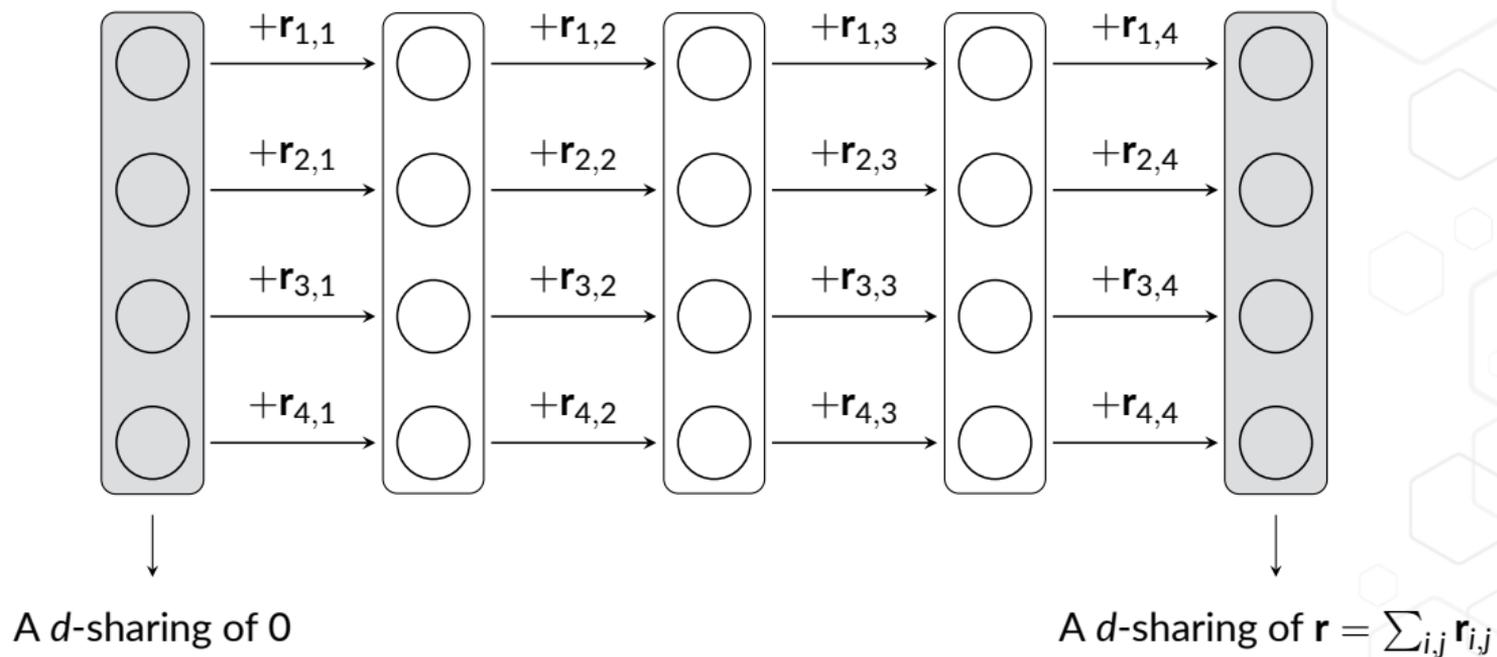
- \rightarrow **AddRepNoise** in lime green
 - \triangleright A t -probing adversary learns at most t of the $(d \cdot \text{rep})$ values $\mathbf{r}_{i,j}$
 - \triangleright Formal analysis in [EEN+24]
- \rightarrow **Refresh** is useful for:
 - \triangleright Concrete security
 - \triangleright Composing gadgets (SNI)
 - \triangleright Moving probes around (SNI)

All operations take time $O(d \log d)$.

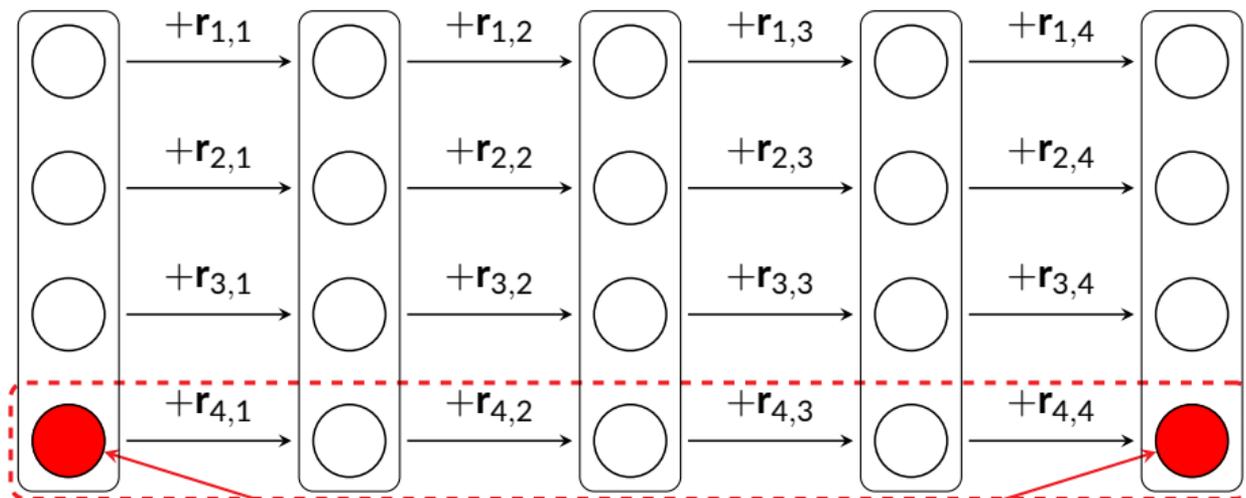
MaskSign($[[\text{sk}]]$, vk, msg) \rightarrow sig

- 1 $[[\mathbf{r}]] = [[\mathbf{0}]]$
- 2 For $i \in [\text{rep}]$:
 - 1 $[[\mathbf{r}_i]] = (\mathbf{r}_{i,1}, \dots, \mathbf{r}_{i,d}) \leftarrow \chi_{\mathbf{r}}^d$
 - 2 $[[\mathbf{r}]] = [[\mathbf{r}]] + [[\mathbf{r}_i]]$
 - 3 Refresh($[[\mathbf{r}]]$)
- 3 $[[\mathbf{w}]] = [\mathbf{A} \ \mathbf{1}] \cdot [[\mathbf{r}]]$
- 4 Refresh($[[\mathbf{w}]]$)
- 5 $\mathbf{w} = \text{Decode}([[\mathbf{w}]])$
- 6 $\mathbf{c} = H(\mathbf{w}, \text{msg}, \text{vk})$
- 7 $[[\mathbf{z}]] = [[\text{sk}]] \cdot \mathbf{c} + [[\mathbf{r}]]$
- 8 Refresh($[[\mathbf{z}]]$, $[[\text{sk}]]$)
- 9 $\mathbf{z} = \text{Decode}([[\mathbf{z}]])$
- 10 Output sig = (c, z)

What happens inside AddRepNoise?

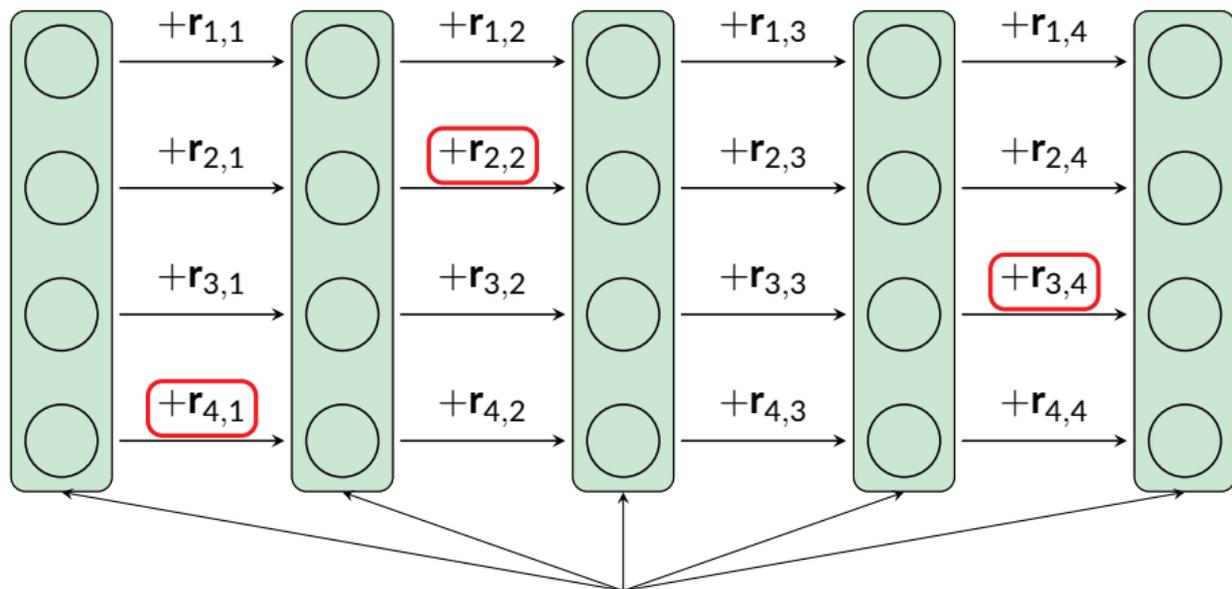


What happens inside AddRepNoise?

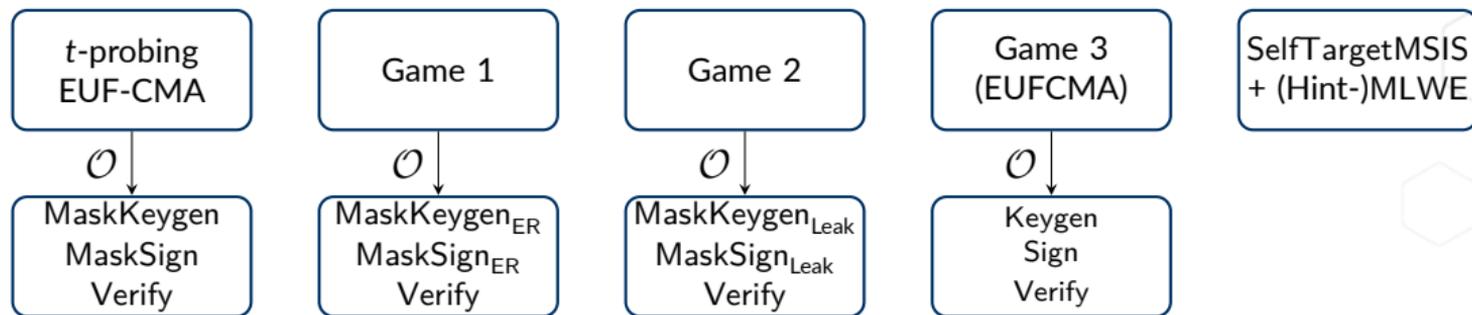


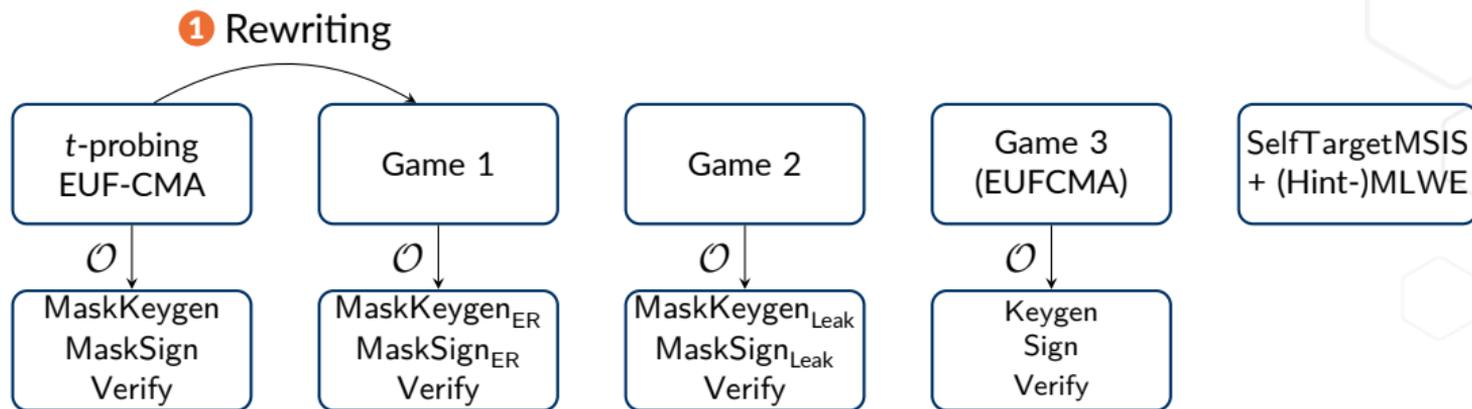
Without Refresh, a probing adversary could learn the sum of T random in 2 probes.

What happens inside AddRepNoise?

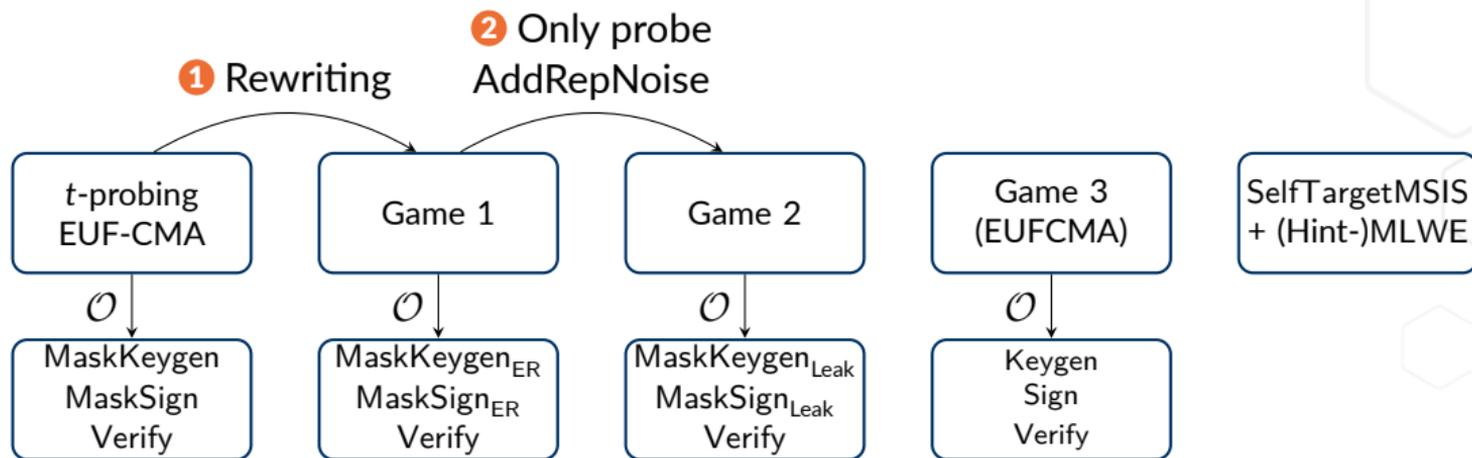


Solution: add refresh gadgets to separate the algorithm in independent layers
Now a probing adversary learns at most (the sum of) t short noises.

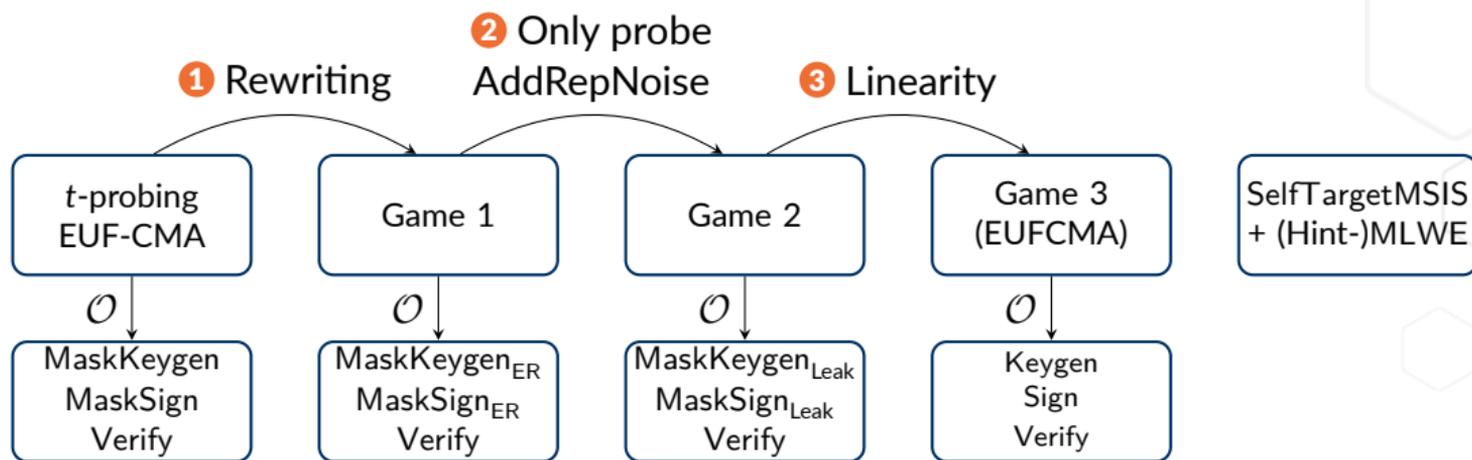




① **Rewriting:** make randomness explicit as input



- ① **Rewriting:** make randomness explicit as input
- ② **SNI(u) property:** move all probes to AddRepNoise randomness



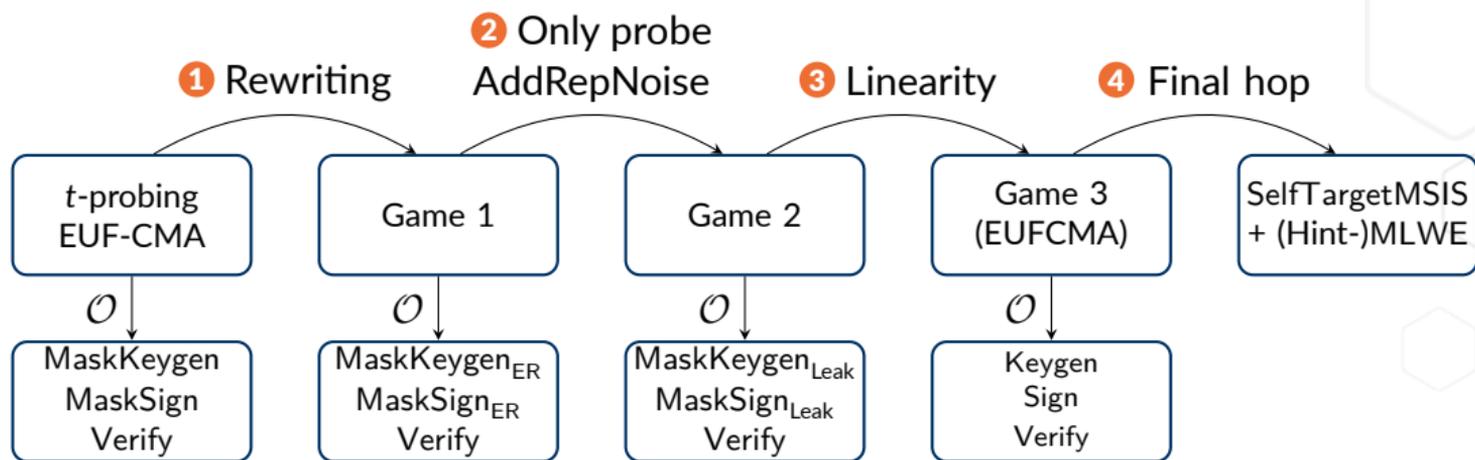
1 **Rewriting:** make randomness explicit as input

2 **SNI(u) property:** move all probes to AddRepNoise randomness

3 **Linearity:** we argue that we can simulate Game 2 from Game 3

Game 2: $\mathbf{w} = \begin{bmatrix} \mathbf{A} & \mathbf{I} \end{bmatrix} \cdot \mathbf{r}$ where $\mathbf{r} = \sum_{i \in [d \cdot \text{rep}]} \mathbf{r}_i$ and we leak $(\mathbf{r}_i)_{i \in S}$ for $|S| = t$

Game 3: $\mathbf{w} = \begin{bmatrix} \mathbf{A} & \mathbf{I} \end{bmatrix} \cdot \mathbf{r}'$ where $\mathbf{r}' = \sum_{i \in [d \cdot \text{rep} - t]} \mathbf{r}_i$



- 1 **Rewriting:** make randomness explicit as input
- 2 **SNI(u) property:** move all probes to AddRepNoise randomness
- 3 **Linearity:** we argue that we can simulate Game 2 from Game 3
 - Game 2: $\mathbf{w} = \begin{bmatrix} \mathbf{A} & \mathbf{I} \end{bmatrix} \cdot \mathbf{r}$ where $\mathbf{r} = \sum_{i \in [d \cdot \text{rep}]} \mathbf{r}_i$ and we leak $(\mathbf{r}_i)_{i \in S}$ for $|S| = t$
 - Game 3: $\mathbf{w} = \begin{bmatrix} \mathbf{A} & \mathbf{I} \end{bmatrix} \cdot \mathbf{r}'$ where $\mathbf{r}' = \sum_{i \in [d \cdot \text{rep} - t]} \mathbf{r}_i$
- 4 **Final hop:** {EUFCMA of Raccoon} \geq {SelfTargetMSIS + (Hint-)MLWE }

Signature sizes are quadratic in $\log q$ (trust me), so we want to minimize q .

Method	Modulus q (logarithmic view)
ML-DSA	<p> $\sigma(\text{sig})$ MSIS (forgery) $\sigma(\text{sk})$ $\ c\ _1 \cdot \dim(\text{sk})$ $\Omega(1)$ MLWE Rejection sampling </p>
Raccoon, Smooth Rényi [Proven]	<p> $\sigma(\text{sig})$ MSIS (forgery) $\sigma(\text{sk})$ $\ c\ \sqrt{Q_s} \cdot \dim(\text{sk}) \cdot \lambda \cdot d^3$ $\sqrt{2}$ $\Omega(1)$ MLWE (key rec.) Smooth Rényi divergence Probing </p>
Raccoon, Hint-MLWE [Heuristic]	<p> $\sigma(\text{sig})$ MSIS $\sigma(\text{sk})$ $\ c\ \sqrt{Q_s}$ $\sqrt{2}$ $\Omega(1)$ MLWE Hint-MLWE reduction (heur.) Probing </p>

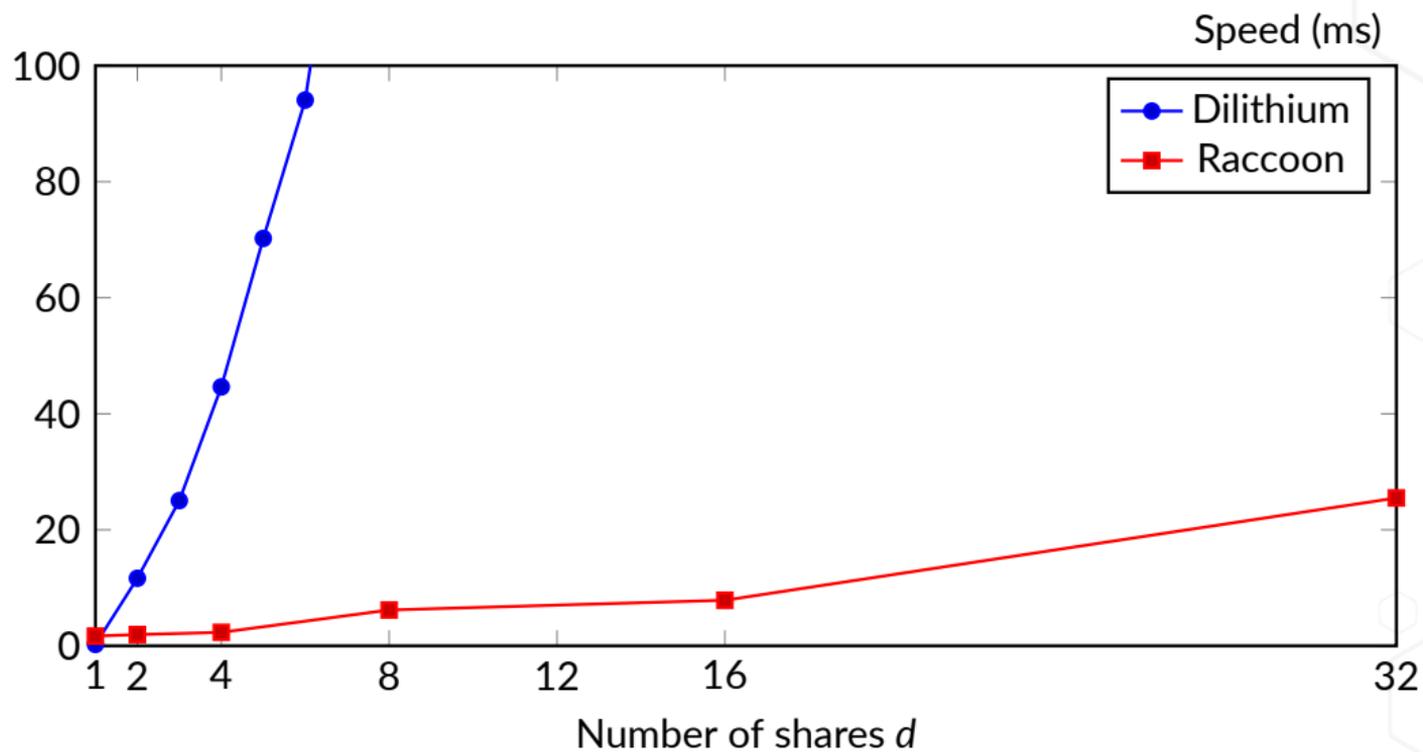
Parameter selection and the modulus q .

Signature sizes are quadratic in $\log q$ (trust me), so we want to minimize q .

Method	Modulus q (logarithmic view)
ML-DSA	$\underbrace{\underbrace{\sigma(\text{sk})}_{\text{MLWE}} \parallel \underbrace{\ c\ _1 \cdot \text{dim}(\text{sk})}_{\text{Rejection sampling}}}_{\sigma(\text{sig})} \parallel \underbrace{\Omega(1)}_{\text{MSIS (forgery)}}$
Raccoon, Smooth Rényi [Proven]	$\underbrace{\underbrace{\sigma(\text{sk})}_{\text{MLWE (key rec.)}} \parallel \underbrace{\ c\ \sqrt{Q_s} \cdot \text{dim}(\text{sk}) \cdot \lambda \cdot d^3}_{\text{Smooth Rényi divergence}}}_{\sigma(\text{sig})} \parallel \underbrace{\sqrt{2}}_{\text{Probing}} \parallel \underbrace{\Omega(1)}_{\text{MSIS (forgery)}}$
Raccoon, Hint-MLWE [Heuristic]	$\underbrace{\underbrace{\sigma(\text{sk})}_{\text{MLWE}} \parallel \underbrace{\ c\ \sqrt{Q_s}}_{\text{Hint-MLWE reduction (heur.)}}}_{\sigma(\text{sig})} \parallel \underbrace{\sqrt{2}}_{\text{Probing}} \parallel \underbrace{\Omega(1)}_{\text{MSIS}}$

→ **ML-DSA:** $q = 23$ bits, $|\text{sig}| = 2420$ bytes

→ **Raccoon:** $q = 49$ bits, $|\text{sig}| = 11524$ bytes



Mask Compression



FIPS 204

Federal Information Processing Standards Publication

Module-Lattice-Based Digital Signature Standard

Category: Computer Security Subcategory: Cryptography

Information Technology Laboratory
National Institute of Standards and Technology
Gaithersburg, MD 20899-8900

This publication is available free of charge from:
<https://doi.org/10.6028/NIST.FIPS.204>

Published August 13, 2024



U.S. Department of Commerce
Gina M. Raimondo, Secretary
National Institute of Standards and Technology
Laurie E. Locantore, NIST Director and Under Secretary of Commerce for Standards and Technology



Right now, we need to implement ML-DSA

Disclaimer: I am not involved in any of the works/techniques in this section. I just think they're neat.

- ✓ Implementing Dilithium/ML-DSA in **RAM-constrained devices**
 - Reference implementation has a footprint ≥ 50 KiB
 - Using several tricks, [BRS22] compress it down to ≤ 7 KiB

- ✓ Implementing **high-order masked** Dilithium/ML-DSA
 - Feasible with $O(d^2 \log q)$ overhead [CGTZ23, CGL⁺24]

- ? Implementing **high-order masked** Dilithium/ML-DSA in **RAM-constrained devices?**

Consider ML-DSA-87:

One ring element	768 B
Secret key (s_1, s_2)	3,840 B
Randomness y	4,480 B
Commitment w	5,888 B
Fully expanded matrix A	41,216 B
...	...

Solutions:

- (Re-)generate everything from seeds (A, s_1, s_2, y, \dots)
- Memory laziness (throw away values after usage)

Masked sensitive values are **expensive** to store. At order 4:

- $[[y]]$: between 17,920 B and 20,608 B
- $[[s_1]], [[s_2]]$: between 8,192 B and 44,160 B

Makes it impractical to implement ML-DSA-87 on devices with ≤ 32 KiB of RAM.

Is there a chance we can use **seeds** to reduce storage?

-  Markku-Juhani O. Saarinen, Mélissa Rossi: *Mask Compression: High-Order Masking on Memory-Constrained Devices*. SAC 2023. [SR24]

Mask Compression

-  **Key idea:** make all shares (except one) pseudorandom:
 - > $[[\mathbf{x}]]_d = (\mathbf{x}_0, \text{seed}_1, \dots, \text{seed}_j)$
 - > $\mathbf{x}_0 = \mathbf{x} - \sum_{i>0} \text{PRF}(\text{seed}_i)$
-  **Security (in isolation):** \mathbf{x} remains secret even if $t < d$ values are probed

-  Markku-Juhani O. Saarinen, Mélissa Rossi: *Mask Compression: High-Order Masking on Memory-Constrained Devices*. SAC 2023. [SR24]

Mask Compression

-  **Key idea:** make all shares (except one) pseudorandom:
 - > $[[\mathbf{x}]]_d = (\mathbf{x}_0, \text{seed}_1, \dots, \text{seed}_j)$
 - > $\mathbf{x}_0 = \mathbf{x} - \sum_{i>0} \text{PRF}(\text{seed}_i)$
-  **Security (in isolation):** \mathbf{x} remains secret even if $t < d$ values are probed
-  **Efficiency:** Decrease the bitsize from $d \cdot |\mathbf{x}|$ down to $|\mathbf{x}| + (d - 1) \cdot \lambda$.
If \mathbf{x} has k coefficients, we may either:
 - > Use one seed per coef \Rightarrow bitsize becomes $|\mathbf{x}| + (d - 1) \cdot k \cdot \lambda$
 - > Use different PRFs \Rightarrow the j -th coef of the i -th share would be $\text{PRF}_j(\text{seed}_i)$

-  Markku-Juhani O. Saarinen, Mélissa Rossi: *Mask Compression: High-Order Masking on Memory-Constrained Devices*. SAC 2023. [SR24]

Mask Compression

-  **Key idea:** make all shares (except one) pseudorandom:
 - > $[[\mathbf{x}]]_d = (\mathbf{x}_0, \text{seed}_1, \dots, \text{seed}_j)$
 - > $\mathbf{x}_0 = \mathbf{x} - \sum_{i>0} \text{PRF}(\text{seed}_i)$
-  **Security (in isolation):** \mathbf{x} remains secret even if $t < d$ values are probed
-  **Efficiency:** Decrease the bitsize from $d \cdot |\mathbf{x}|$ down to $|\mathbf{x}| + (d - 1) \cdot \lambda$.
If \mathbf{x} has k coefficients, we may either:
 - > Use one seed per coef \Rightarrow bitsize becomes $|\mathbf{x}| + (d - 1) \cdot k \cdot \lambda$
 - > Use different PRFs \Rightarrow the j -th coef of the i -th share would be $\text{PRF}_j(\text{seed}_i)$
-  **Computations?**
 - > [SR24] show how to perform a SNI refresh (compatible w/ this structure)
 - > Other ML-DSA operations (decompose, addition, rejection) may require all shares at once \rightarrow see efficiency trade-offs

-  Markku-Juhani O. Saarinen, Mélissa Rossi: *Mask Compression: High-Order Masking on Memory-Constrained Devices*. SAC 2023. [SR24]

Mask Compression

-  **Key idea:** make all shares (except one) pseudorandom:
 - $\gg \llbracket \mathbf{x} \rrbracket_d = (\mathbf{x}_0, \text{seed}_1, \dots, \text{seed}_d)$
 - $\gg \mathbf{x}_0 = \mathbf{x} - \sum_{i>0} \text{PRF}(\text{seed}_i)$
-  **Security (in isolation):** \mathbf{x} remains secret even if $t < d$ values are probed
-  **Efficiency:** Decrease the bitsize from $d \cdot |\mathbf{x}|$ down to $|\mathbf{x}| + (d - 1) \cdot \lambda$.
If \mathbf{x} has k coefficients, we may either:
 - \gg Use one seed per coef \Rightarrow bitsize becomes $|\mathbf{x}| + (d - 1) \cdot k \cdot \lambda$
 - \gg Use different PRFs \Rightarrow the j -th coef of the i -th share would be $\text{PRF}_j(\text{seed}_i)$
-  **Computations?**
 - \gg [SR24] show how to perform a SNI refresh (compatible w/ this structure)
 - \gg Other ML-DSA operations (decompose, addition, rejection) may require all shares at once \rightarrow see efficiency trade-offs

This allows to implement ML-DSA-87 with 4 shares and 16 KiB of RAM.

Conclusion



Masking-friendly lattice schemes

- Requires flexibility in exploration of design space and security notions
- Can be extremely efficient
- ? Concrete SCA resilience?
- ? Better proofs/constructions in alternative leakage models?
- ? Masking-friendly KEMs?

∞ Lattice-friendly masking schemes

- We have barely scratched the surface
- ? What would be **really** nice is a masking scheme that:
 - Can be converted efficiently from/to arithmetic masking
 - Allows to perform efficiently decompose/sample/reject

Questions?

<https://raccoonfamily.org>

<https://ia.cr/2024/1291>

<https://ia.cr/2023/1117>

<https://tprest.github.io>



- 
-  Joppe W. Bos, Joost Renes, and Amber Sprenkels.
Dilithium for memory constrained devices.
In Lejla Batina and Joan Daemen, editors, *AFRICACRYPT 22*, volume 2022 of *LNCS*, pages 217–235. Springer, Cham, July 2022.
 -  Jean-Sébastien Coron, François Gérard, Tancrède Lepoint, Matthias Trannoy, and Rina Zeitoun.
Improved high-order masked generation of masking vector and rejection sampling in dilithium.
IACR Transactions on Cryptographic Hardware and Embedded Systems, 2024(4):335–354, Sep. 2024.
 -  Jean-Sébastien Coron, François Gérard, Matthias Trannoy, and Rina Zeitoun.
Improved gadgets for the high-order masking of Dilithium.
IACR TCHES, 2023(4):110–145, 2023.
 -  Rafaël Del Pino, Shuichi Katsumata, Mary Maller, Fabrice Mouhartem, Thomas Prest, and Markku-Juhani O. Saarinen.
Threshold raccoon: Practical threshold signatures from standard lattice assumptions.

In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part II*, volume 14652 of *LNCS*, pages 219–248. Springer, Cham, May 2024.

 Rafaël del Pino, Shuichi Katsumata, Thomas Prest, and Mélissa Rossi.
Raccoon: A masking-friendly signature proven in the probing model.
In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part I*, volume 14920 of *LNCS*, pages 409–444. Springer, Cham, August 2024.

 Muhammed F. Esgin, Thomas Espitau, Guilhem Niot, Thomas Prest, Amin Sakzad, and Ron Steinfeld.
Plover: Masking-friendly hash-and-sign lattice signatures.
In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part VII*, volume 14657 of *LNCS*, pages 316–345. Springer, Cham, May 2024.

 Emre Karabulut and Aydin Aysu.
FALCON down: Breaking FALCON post-quantum signature scheme through side-channel attacks.
In *58th ACM/IEEE Design Automation Conference, DAC 2021, San Francisco, CA, USA, December 5-9, 2021*, pages 691–696. IEEE, 2021.

 Eike Kiltz, Vadim Lyubashevsky, and Christian Schaffner.

A concrete treatment of Fiat-Shamir signatures in the quantum random-oracle model.

In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part III*, volume 10822 of *LNCS*, pages 552–586. Springer, Cham, April / May 2018.

 Duhyeong Kim, Dongwon Lee, Jinyeong Seo, and Yongsoo Song.
Toward practical lattice-based proof of knowledge from hint-MLWE.
In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part V*, volume 14085 of *LNCS*, pages 549–580. Springer, Cham, August 2023.

 Markku-Juhani O. Saarinen and Mélissa Rossi.
Mask compression: High-order masking on memory-constrained devices.
In Claude Carlet, Kalikinkar Mandal, and Vincent Rijmen, editors, *SAC 2023*, volume 14201 of *LNCS*, pages 65–81. Springer, Cham, August 2024.

 Shiduo Zhang, Xiuhan Lin, Yang Yu, and Weijia Wang.
Improved power analysis attacks on falcon.
In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part IV*, volume 14007 of *LNCS*, pages 565–595. Springer, Cham, April 2023.