

High-Order Masking of Lattice Signatures in Quasilinear Time


Rafaël del Pino ¹ Thomas Prest ¹
Mélissa Rossi ² Markku-Juhani O. Saarinen ¹

¹ PQShield ²ANSSI

44th IEEE Symposium on Security and Privacy
22 May 2023, San Francisco, CA

- **Side-Channel Attacks (SCA)** use external measurements such as latency (TA), power consumption (SPA/DPA), or electromagnetic emissions ([S/D]EMA) to extract secrets.
- *SCA resistance is important for PC, IoT, and mobile device “platform security” (secure boot, firmware updates, attestation), authentication tokens, smart cards, HSMs / secure elements..*
- Common compliance & market requirement for hardware (Common Criteria / AVA_VAN, FIPS 140-3 / ISO 17825).

- **Post-Quantum Cryptography (PQC)** implementations – e.g. lattice-based signature schemes **Dilithium** and **Falcon** inherit all of the security and compliance requirements of Elliptic Curve or RSA based solutions in applications.



- 1 **Masked Raccoon** is a member of the new **Raccoon** family of lattice-based PQC signature schemes.
- 2 **Side-Channel Security** is proved in the Strong Non-Interference (SNI) framework.
- 3 **Cryptanalytic Security** is proved in relation to well-studied MLWE and SelfTargetMSIS problems.
- 4 **Performance** is evaluated with both PC and a constrained FPGA hardware target.



- **Masking:** Secret data $[[\mathbf{s}]]$ is processed in $d = \text{order} + 1$ randomized shares \mathbf{s}_i .

$$\text{Boolean Masking: } [[\mathbf{s}]] = \mathbf{s}_1 \oplus \mathbf{s}_2 \oplus \cdots \oplus \mathbf{s}_d$$

$$\text{Arithmetic Masking: } [[\mathbf{s}]] = \mathbf{s}_1 + \mathbf{s}_2 + \cdots + \mathbf{s}_d \pmod{q}.$$

- Like secret sharing: Knowledge of $d - 1$ shares \mathbf{s}_i does not reveal anything about $[[\mathbf{s}]]$.
- If you only have partial or “noisy” measurements (traces), it has been shown that the number of such observations required to learn $[[\mathbf{s}]]$ grows exponentially with d .
- **Masking proofs** give formal, algorithm-level assurance against side-channel leakage.
- The proofs can be made in several models; the Ishai-Sahai-Wagner (ISW) t -probing security requires that any t internal intermediate values don't reveal secrets.
- The noisy leakage model is an alternative; links have been proven between t -probing security, noisy leakage model, and information-theoretic attack complexity bounds.

- Linear operations only need **linear** $O(d)$ effort to mask:
Addition / subtraction / XOR of masked variables ($\llbracket \mathbf{s} \rrbracket + \llbracket \mathbf{r} \rrbracket$), multiplication (or Boolean AND, OR) with a scalar constant or a public variable ($\mathbf{c} \cdot \llbracket \mathbf{s} \rrbracket$), or share-independent linear operations such as NTT (Number Theoretic Transform.)
- Non-linear operations generally require **quadratic** $O(d^2)$ effort:
Multiplication (Boolean AND, OR) between secret variables ($\llbracket \mathbf{s} \rrbracket \cdot \llbracket \mathbf{r} \rrbracket$), conversion between Arithmetic and Boolean masking (A2B and B2A), or symmetric cryptography like SHA3.
- But some non-linear operations can be done with **quasilinear** $O(d \log d)$ effort:
Practical quasilinear techniques are known only for a limited number of computational tasks.

- **Dilithium** requires a masked SHAKE; mixes bit manipulations with (mod q) arithmetic, requiring A2B and B2A; has masked comparisons / rejection sampler.

(For these non-linear operations only quadratic $O(d^2)$ gadgets are known.)

- **Raccoon** avoids quadratic operations. The cost of additional shares is nearly constant. *(Cycles/share even decreases initially due to a small constant overhead.)*

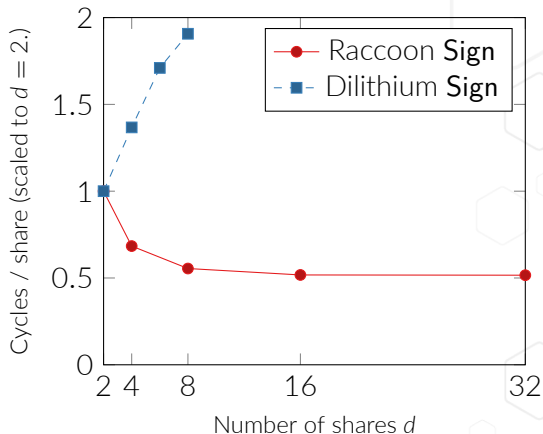


Figure 1: Cost of masking: Signing cycle count divided by d , normalized to a common start at 1 for $d = 2$. Dilithium data from [24, Table 3].

- Blueprint from Lyubashevsky [15,16], refined by Bai and Galbraith [17], and used in Dilithium and this work.
- Public key $\mathbf{vk} = (\mathbf{A}, \mathbf{t} = \mathbf{A} \cdot \mathbf{s} + \mathbf{e})$ is a *Module Learning With Errors*, or **MLWE** sample. Additionally, the security proof uses **SelfTargetMSIS** (as in Dilithium).
- There actually aren't “secret-secret” multiplications in the blueprint! *Could we build it entirely with quasilinear gadgets?*

Algorithm 1 PrototypeSign($\mathbf{sk}, \mathbf{vk}, \text{msg}$)

Input: A signing key $\mathbf{sk} = \mathbf{s}$, a verification key $\mathbf{vk} = (\mathbf{A}, \mathbf{t})$, a message msg .

Output: A signature sig of msg under \mathbf{sk} .

- 1: Sample \mathbf{r} uniformly in a small set S
 - 2: $\mathbf{u} := \mathbf{A} \cdot \mathbf{r}$
 - 3: $\mathbf{w} := \text{Truncate}(\mathbf{u})$ ▷ Commitment
 - 4: $c := H(\mathbf{w}, \text{msg})$ ▷ Challenge
 - 5: $\mathbf{z} := \mathbf{r} + c \cdot \mathbf{s}$ ▷ Response
 - 6: $\mathbf{y} := \mathbf{A} \cdot \mathbf{z} - c \cdot \mathbf{t}$
 - 7: **if** CheckCondition(\mathbf{z}, \mathbf{y}) = **False** **then**
 - 8: **goto** Line 1 ▷ Rejection sampling
 - 9: **return** $\text{sig} := (c, \mathbf{z})$
-

- Cryptanalytic sensitivity analysis: Which variables need to be protected?
- Raccoon signature and key generation functions are composed of **Masking Gadgets** that are individually t -non-interfering ($t - \text{NI}$) or t -strong non-interfering ($t - \text{SNI}$).
- The scheme is designed to be “masking friendly,” so the proofs are quite standard.




Figure 2: Raccoon signature function. Colors: gadget proven $t - \text{NI}$, gadget proven $t - \text{SNI}$, gadget unmasked. Single arrows (\rightarrow) and double arrows (\Rightarrow) represent plain and masked values.

Example: $O(d \log d)$ masking refresh (re-randomization) gadget, proven t – SNI [35,36,37].

Algorithm 2 Refresh()

Input: A d -shared $\llbracket x \rrbracket$ of $x \in \mathbb{Z}_q$

Output: A fresh d -shared $\llbracket x \rrbracket$ of x

- 1: $\llbracket z \rrbracket \leftarrow \text{ZeroEncoding}()$
 - 2: **return** $\llbracket x \rrbracket = \llbracket x \rrbracket + \llbracket z \rrbracket$
-

- Proofs examine correlations between intermediate variables, input/output.
- (Hardware implementation has circuits to generate masking randomness efficiently and perform all the ring arithmetic ops.)

Algorithm 3 ZeroEncoding()

Input: A power-of-two integer d , a ring \mathbb{Z}_q

Output: A d -shared $\llbracket z \rrbracket \in \mathbb{Z}_q^d$ of $0 \in \mathbb{Z}_q$

- 1: **if** $d = 1$ **then**
 - 2: **return** $\llbracket z_1 \rrbracket = (0)$ ▷ Order zero.
 - 3: $\llbracket z_1 \rrbracket_{d/2} \leftarrow \text{ZeroEncoding}(d/2)$
 - 4: $\llbracket z_2 \rrbracket_{d/2} \leftarrow \text{ZeroEncoding}(d/2)$
 - 5: $\llbracket r \rrbracket_{d/2} \leftarrow \mathbb{Z}_q^{d/2}$ ▷ Uniform random vector.
 - 6: $\llbracket z_1 \rrbracket_{d/2} = \llbracket z_1 \rrbracket_{d/2} + \llbracket r \rrbracket_{d/2}$
 - 7: $\llbracket z_2 \rrbracket_{d/2} = \llbracket z_2 \rrbracket_{d/2} - \llbracket r \rrbracket_{d/2}$
 - 8: **return** $\llbracket z \rrbracket_d = (\llbracket z_1 \rrbracket_{d/2} \parallel \llbracket z_2 \rrbracket_{d/2})$
-

- Hybrid Lemma 3 bounds a forger $\frac{\text{Adv}_A^{\text{Sign}}}{Q_s}$ to distinguishing public key from uniform Adv_A^{PK} .
- Thm. 1 provides a reduction to MSIS. Further consideration of SelfTargetMSIS is used in parameter selection (BKZ attack Core-SVP.)
- Thm. 2 reduces PK distinguishability to MLWE.



Parameter Selection

- As usual for lattice schemes, parameter selection for each security target λ_{target} is a complex multi-objective optimization problem.
- However, the core problems and high-level structure are well-studied, so we can rely on a large body of existing research.

Name	Raccoon- λ_{target}		
λ_{target}	128	192	256
Q_s	2^{48}	2^{48}	2^{49}
d	32	-	-
$\log q$	49 [†]	-	-
$\log p_t$	10	6	7
$\log p_w$	43	40	42
n	512	-	-
k	8	11	14
ℓ	3	5	6
ω	19	31	44
B_2^2	2^{14}	2^{14}	2^{15}
B_∞	8	-	-
$ \mathbf{vk} $	19 968	30 272	37 632
$ \mathbf{sig} $	12 000	19 232	23 328

[†]Across all parameter sets, we set $q = (2^{25} - 2^{18} + 1) \cdot (2^{24} - 2^{18} + 1)$.

- Portable C Implementation was developed to assess the relative speed to other algorithms. Unmasked Dilithium runs at about 1/2 time of than Raccoon with $d = 2$ masking. Unfortunately not many comparison points (no open masked SW Dilithium.)
- Artix7 FPGA target implements Raccoon up $d = 32$ and also has $d = 2$ proprietary Dilithium HW. Raccoon is already faster at first order, tens of times faster with higher d .
- No secret key leakage was detected in a 200,000-trace ISO 17825 / “TVLA” style leakage assessment of Raccoon-128 ($d = 2$) signature function on the FPGA target.



Contributions:

- 1 In this work, we have shown that lattice-based signature schemes can be masked with quasilinear complexity – giving the “defenders” a significant asymptotic advantage.
- 2 Proposed new algorithmic techniques, as well as new proof techniques.
- 3 Software and hardware experiments show that the performance and concrete leakage profile of Raccoon are consistent with our theoretical analyses (+new masking records!)

Note: *We have further developed the Raccoon framework since this work was submitted and have found new techniques and applications. Also, the parameter selection has changed.*

We are currently working (with an expanded team) to release a new version of Raccoon.