IACR Communications in Cryptology
ISSN 3006-5496, Vol. 1, No. 3, 21 pages.

https://doi.org/10.62056/abkp2c3w9p

Information Theoretic Evaluation of Raccoon's Side-Channel Leakage

Dinal Kamel[1], François-Xavier Standaert[1] and Olivier Bronchain[2]

[1] UCLouvain, Crypto Group, Louvain-la-Neuve, Belgium
[2] NXP Semiconductors, Leuven, Belgium

Abstract. Raccoon is a lattice-based scheme submitted to the NIST 2022 call for additional post-quantum signatures. One of its main selling points is that its design is intrinsically easy to mask against side-channel attacks. So far, Raccoon's physical security guarantees were only stated in the abstract probing model. In this paper, we discuss how these probing security results translate into guarantees in more realistic leakage models. We also highlight that this translation differs from what is usually observed (e.g., in symmetric cryptography), due to the algebraic structure of Raccoon's operations. For this purpose, we perform an in-depth information theoretic evaluation of Raccoon's most innovative part, namely the`AddRepNoise` function which allows generating its arithmetic shares on-the-fly. Our results are twofold. First, we show that the resulting shares do not enforce a statistical security order (i.e., the need for the side-channel adversary to estimate higher-order moments of the leakage distribution), as usually expected when masking. Second, we observe that the first-order leakage on the (large) random coefficients manipulated by Raccoon cannot be efficiently turned into leakage on the (smaller) coefficients of its long-term secret. Concretely, our information theoretic evaluations for relevant leakage functions also suggest that Raccoon's masked implementations can ensure high security with less shares than suggested by a conservative analysis in the probing model.

Keywords: Post-Quantum Cryptography · Signature Schemes · Side-Channel Analysis · Masking Countermeasure · CRYSTALS-Dilithium · Raccoon

1 Introduction

In parallel to its standardization in 2022, the secure implementation of CRYSTALS-Dilithium has been shown to be non-trivial due to different attack vectors, for example including [RJH+18, LZS+21, UMTS22, BVC+23]. As a result it has been concluded that protecting Dilithium against leakage requires masking all its operations, which implies overheads that are roughly quadratic in the number of shares [MGTF19, ABC+23].

Originally introduced at IEEE S&P 2023 [dPPRS23], Raccoon is an alternative signature algorithm that has been submitted to the NIST call for additional post-quantum signature schemes.[1] One of the main selling points of Raccoon is that the fine-tuning of its parameters enables to make its design more amenable to the masking countermeasure, and therefore to reduce its implementation overheads from quadratic to quasi-linear.

From the leakage viewpoint, the critical operation that discriminates the design of Dilithium from the one of Raccoon, on which this paper is focused, is the generation of a signature $\mathbf{z}$ from a challenge $\mathbf{c}$, long-term secret $\mathbf{s}$ and randomness $\mathbf{r}$ as $\mathbf{z} = \mathbf{c} \circ \mathbf{s} + \mathbf{r}$.[2]

E-mail: dina.kamel@uclouvain.be (Dinal Kamel)
[1] See https://raccoonfamily.org/ for the latest updates and complete specification.
[2] The randomness is denoted as $\mathbf{y}$ in CRYSTALS-Dilithium.

Received: 2024-07-09 Accepted: 2024-09-02

In the case of Dilithium, and as detailed in [ABC+23], a Boolean-to-arithmetic conversion is necessary in order to generate the arithmetic shares of $\mathbf{r}$ based on binary random values (which are generated with SHA3 in Dilithium's standard version).

By contrast, in the case of Raccoon, a specific procedure denoted as `AddRepNoise` is used and allows generating the arithmetic shares of $\mathbf{r}$ on-the-fly, without relying on a Boolean-to-arithmetic conversion. The authors of Raccoon show that this procedure is secure in the probing model [ISW03]. Yet, contrary to the vast majority of existing proofs in this model, their simulation is not perfect. The authors rather use a statistical argument based on the Rényi divergence to show that the simulation error is small enough.[3]

While proofs in the probing model are usually an important step towards secure implementations, turning them into practical security guarantees requires analyses in more concrete leakage models, such as the noisy leakage model [PR13]. In the context of symmetric cryptography, where most simulations are perfect, it is possible to leverage security reductions from one leakage model to the other. As a result, the conceptual match between theory and practice is clear and well established [DDF14, DFS15].

Unfortunately, such a direct link between probing security and noisy leakage security is not (yet) proven in case of imperfect simulations. For example, the quite popular (though heuristic) notion of statistical security order [SM16], defined as the highest statistical moment of the leakage distribution that is independent of any secret manipulated by an implementation, cannot be guaranteed in this case. While this may not be a concrete issue, since security depends on the effective security order [DFS15, LBS19], defined as the order of the statistical moment that can be exploited with the smallest number of measurements, it implies that understanding the concrete security guarantees of masked Raccoon implementations still requires more in-depth investigations.

In this paper, we aim to contribute to this issue by analysing the leakage of masked Raccoon implementations using the information theoretic framework advocated in [SMY09]. Such information theoretic analyses can be seen as a specific counterpart to generic proofs in the noisy leakage model. That is, rather than proving the leakage security of any implementation given some (noise and independence) assumptions, it allows quantifying the leakage security against actual (practically-relevant or theoretically-insightful) leakage functions. Concretely, such analyses typically lead to plots where the (mutual) information leaked on sensitive variables is given in function of a noise parameter for different number of shares, where (if using a log/log scale) the slope of the curves captures the (effective) statistical security order of the implementations – see for example [SVO+10].

Since information theoretic analyses become computationally intensive as the size of the fields in which one operates grows, our investigations consider reduced versions of Raccoon and capture how the side-channel security of their masked implementation evolves with the size of their parameters. We also consider two illustrative leakage functions: first, the Hamming Weight (HW) function as a natural simplification of the (global) leakages that can be encountered in practice [MOP07]. Second, the Least Significant Bit (LSB) function as a possibly unrealistic example of very localized leakage. Our evaluations lead to observations that shed new light on the implementation features of Raccoon.

First, and for both the HW and LSB functions, there is always first-order side-channel leakage on the randomness $\mathbf{r}$. This leakage can be explained by the fact that `AddRepNoise` generates shares on-the-fly by means of additions (rather than modular additions), leading to non-uniform patterns that can be observed in the mean leakage values.

Second, this first-order leakage on $\mathbf{r}$ does not automatically translate into an exploitable first-order leakage on the ephemeral secret $\mathbf{x} = \mathbf{c} \cdot \mathbf{s}$ and leads to consequences for the side-channel security of Raccoon that depends on the leakage functions.

[3] The same argument conveniently allows Raccoon to be implemented without rejection sampling.

For the (more realistic) HW function, we show that side-channel security mostly depends on the ratio between the variance of the randomness $\mathbf{r}$ and the variance of the ephemeral secret $\mathbf{x}$. That is, security essentially comes from the fact that the information provided by a (large) HW on $\mathbf{r}$ is considerably reduced when compressed into information on the smaller $\mathbf{x}$. There is no "security order intuition" in this context, and therefore limited incentive to increase the number of shares for side-channel security, as initially foreseen by the Raccoon designers. But the situation is actually better than that, since we can show that a modest amount of noise (that decreases with the size of $\mathbf{r}$) is enough for the first-order side-channel leakage on $\mathbf{r}$ to become innocuous. That is, already for low number of shares, this physical leakage rapidly becomes less informative than the (provably small) mathematical leakage due to the knowledge of the public signature $\mathbf{z}$.

For the LSB function, the more localized information brings us back to the standard case where increasing the number of shares amplifies the physical leakage noise.

As complementary discussions, we put forward that applying the design ideas of Raccoon to Dilithium will not be possible without increasing the size of its parameters. In order to clarify the practical implications of our results, we then repeat our information theoretic evaluations for the (linear) leakage model of the ARM Cortex-M4 given in [CDSU23]. Despite such a device is a challenging target for secure implementation given its low level of noise, we show that in this case as well, the leakage of `AddRepNoise` is unlikely to be a concrete threat, even for Raccoon's smallest physical security parameters. In order to confirm that these positive results are not due to the linearity of the leakage model in [CDSU23], we also provide results using the ELMO leakage model [MOW17], which includes quadratic terms, and a simulated leakage model where the quadratic terms of adjacent wires contribute to the leakage as much as the linear terms corresponding to the wires. We conclude by observing that our analyzes are in line with the security arguments given by Raccoons' designers, and can be seen are the two faces of the same coin.

Related work. In the rest of the paper, we focus on the leakage on the randomness $\mathbf{r}$ and ephemeral secret $\mathbf{x}$, not on the long-term secret $\mathbf{s}$. Bronchain et al. show how leakage on the ephemeral $\mathbf{x}$ can be turned onto leakage on the long-term secret $\mathbf{s}$ [BAE+24].

2 Background

2.1 Notations

We denote by $\mathbb{Z}_q$ the ring of integers modulo the prime q, and by $\mathcal{R}_q = \mathbb{Z}_q[x]/(x^n+1)$ the polynomial ring in x modulo x^n+1, with n the degree of the polynomial. Polynomials in $\mathbb{Z}_q[x]/(x^n+1)$ are written in bold. For Raccoon, $q = 549824583172097$ and $n = 512$. The multiplication between two polynomials $\mathbf{a}$ and $\mathbf{b}$ is written as $\mathbf{c} = \mathbf{a} \circ \mathbf{b}$. Constants are denoted with Greek letters. Calligraphic letters like $\mathcal{S}$ denote sets. The i-th share of the j-th coefficient in a polynomial $\mathbf{a}$ is denoted as $\mathbf{a}_\mathbf{i}^\mathbf{j}$ or simply $\mathbf{a_i}$ when knowledge of the coefficient position is unnecessary. We further denote random variables with upper case letters, e.g., X, and their realizations with lower cases, e.g., $X = x$. We finally use the notation $\Pr[X = x] := \Pr[x]$ to denote the probability of a realization.

2.2 Masking countermeasure

Masking is one of the most investigated countermeasures against side-channel attacks [CJRR99, GP99]. It consists in splitting any sensitive value manipulated by an implementation into d shares. In lattice-based cryptography, two types of additive encodings are used, namely Boolean and Arithmetic. Let X be a sensitive random variable (i.e., an intermediate value that depends on a long-term secret). In a Boolean masking scheme, the

sharing of a variable $X \in \mathbb{F}_{2^n}$ is expressed as $X = \bigoplus_{i=0}^{d-1} X_i$, where X_i is the i-th share of X. Boolean masking is typically used for protecting symmetric primitives such as block ciphers or hash functions. In an arithmetic masking scheme, the sharing of a variable $X \in \mathbb{Z}_q$ is expressed as $X = \sum_{i=0}^{d-1} X_i \bmod q$, where X_i is the i-th share of X. Arithmetic masking is typically used to perform polynomial operations such as additions and subtractions, share-independent linear operations such as the Number Theoretic Transform (NTT) and multiplications with public variables or scalar constants. It has been proven that Boolean masking can amplify the noise of an implementation (and therefore its security) exponentially in the number of shares [ISW03, PR13, DDF14, DFS15, IUH22, BCG+23]. Furthermore, arithmetic masking (especially if operating in prime fields) can lead to security amplification even without noise [DFS16, MMMS23, FMM+24].

2.3 Raccoon parameters

We next detail the parameters that directly impact the security analysis in this work. Table 1 contains the parameters for Raccoon-128 (NIST post-quantum level I). We refer to the Raccoon specifications (https://raccoonfamily.org/) for the parameter sets of Raccoon-192 and Raccoon-256 (NIST post-quantum levels III and V, respectively).

Table 1: Raccoon-128 parameters' set.

Parameter	Raccoon-128	128-2	128-4	128-8	128-16	128-32
ω			19			
d	1	2	4	8	16	32
rep	8	4	2	4	2	4
T	8	8	8	32	32	128
u_t	6	6	6	5	5	4
u_w	41	41	41	40	40	39
size of $\mathbf{r}$ coef.	44	44	44	45	45	46
size of $\mathbf{x}$ coef.	14	14	14	15	15	16
variance of $\mathbf{r}$			3e24			
variance of $\mathbf{x}$			5e4			

The parameters in the table are as follows: ω is the number of non-zero coefficients in the challenge polynomial $\mathbf{c}$, d is the number of masking shares, rep is the number of iterations used to generate a Sum of Uniforms (SU) distribution (which, as explained in the following section, is close to Gaussian), T is the product of d and rep, u_t and u_w are distribution parameters for the long-term secret polynomial $\mathbf{s}$ and the randomness polynomial $\mathbf{r}$, respectively (i.e., the size of the small uniform added noise).

2.4 Randomness sampling

One of Raccoon's main design goals is its better amenability to masking, and one of the most challenging operation to mask in lattice-based signature schemes is the sampling of random errors and secrets. For example, in Dilithium, the (uniform) randomness is first sampled in a Boolean masked form and a Boolean-to-arithmetic converter is used to transform the shares afterwards. Such a mask conversion is not trivial to implement securely, and has a complexity of $O(d^2)$ – see, e.g., [BBE+18] – where d is the number of shares. Furthermore, the standard version of Dilithium requires to mask the Extendable-Output Function (XOF), based on SHA-3 standard, which generates the randomness and

has a complexity which is almost quadratic in the number of shares as well [ABC+23]. Raccoon, on the other hand, simplifies the randomness generation by directly outputting the shares in arithmetic form. Each share of each integer coefficient of the randomness is the sum of rep uniform random samples in a fixed interval. As a result, each randomness coefficient is the sum of $d \cdot rep$ uniform samples, and the small uniform random samples are generated using an XOF (based on SHA-3) which does not need to be masked as explained by the Raccoon authors on the webpage: https://raccoonfamily.org/.

The `AddRepNoise` gadget that is used to perform this operation, on which we focus next, is therefore especially critical from the side-channel analysis viewpoint.

2.5 AddRepNoise procedure

As explained in the Raccoon specifications, the `AddRepNoise` function can be used to generate the shares of the Sum of Uniforms distribution $\mathrm{SU}(u, d \cdot rep)$ on-the-fly, where the $\mathrm{SU}(u, T)$ distribution is given by:

$$\mathrm{SU}(u, T) = [T] \cdot \mathcal{U}(\{-2^{u-1}, \ldots, 2^{u-1} - 1\}),$$

with $[T] \cdot \mathcal{U}$ the distribution of the sum of $T = d \cdot rep$ independent random variables, each being sampled from the uniform distribution $\mathcal{U}$. For this purpose, `AddRepNoise` interleaves noise additions and refresh operations (which are necessary to ensure security in the probing model). For each masked coefficient of a polynomial, a small uniform noise is added to each share before being refreshed, and this operation is repeated rep times. It is important to note that the additions used in `AddRepNoise` are not modular, which is in contrast with arithmetic encodings (e.g., used in symmetric cryptography) and, as will be discussed later, is the source of low-order leakages. Before the deployment of the `AddRepNoise` function to generate the randomness, a call to a gadget named `ZeroEncoding` is made. The output samples of `AddRepNoise` belong to a small field set $(T \cdot (-2^{u-1}), \ldots, T \cdot (2^{u-1} - 1))$, whereas the shares of the randomness coefficients $\in \mathbb{Z}_q$. Using `ZeroEncoding` before `AddRepNoise`, which outputs random arithmetic shares $\in \mathbb{Z}_q$ of 0, guarantees that the shares of the randomness coefficients $\in \mathbb{Z}_q$.

2.6 Raccoon Signature generation

We now describe the operations required for generating the signature $\mathbf{z}$. Raccoon is based on the MLWE hard learning problem. Therefore, the signature $\mathbf{z}$, the randomness $\mathbf{r}$ and the secret key $\mathbf{s}$ are vectors of polynomials while the challenge $\mathbf{c}$ is a polynomial. In this work we consider a single polynomial operation. The signature polynomial $\mathbf{z}$ is:

$$\mathbf{z} = \mathbf{c} \circ \mathbf{s} + \mathbf{r}.$$

Randomness. The ephemeral randomness $\mathbf{r}$ is generated in masked form for every new signature. For this purpose, `AddRepNoise` is applied in order to gradually transform a d-share encoding of zero into a d-share encoding of $\mathbf{r}$. All coefficients in this polynomial are independently drawn from $\mathrm{SU}(u_w, T)$ which is approximately Gaussian. The mean and variance of the $\mathbf{r}$ distribution are computed as $-T/2$ and $T \times (2^{2u_w} - 1)/12$.

Secret key. The long-term secret key $\mathbf{s}$ is generated in masked form during key generation. It also leverages the `AddRepNoise` procedure and its coefficients are independently drawn from $\mathrm{SU}(u_t, T)$ which is approximately Gaussian. The mean and variance of the $\mathbf{s}$ distribution are computed as $-T/2$ and $T \times (2^{2u_t} - 1)/12$.

Challenge. The challenge $\mathbf{c}$ (denoted as c_{poly} in the Raccoon specifications) is computed deterministically during the signature process, as an expansion of a random bit string

thanks to `ChalPoly`. It is a "ternary" polynomial with a weight of ω elements, so that exactly ω coefficients of $\mathbf{c}$ are equal to either $+1$ or -1 and the rest are zeros.

Ephemeral secret. The signature generation involves computing the polynomial multiplication between $\mathbf{c}$ and $\mathbf{s}$, which we will denote as the ephemeral secret $\mathbf{x}$:

$$\mathbf{x^i} = \sum_{\mathbf{j=0}}^{\mathbf{i}} \mathbf{s^j} \cdot \mathbf{c^{i-j}} - \sum_{\mathbf{j=i+1}}^{\mathbf{n-1}} \mathbf{s^j} \cdot \mathbf{c^{n+i-j}}.$$

Since the challenge polynomial $\mathbf{c}$ contains exactly ω non-zero coefficients, each coefficient of $\mathbf{x}$ is the weighted sum of a subset of ω secret keys coefficients. Consequently, the distribution of the sensitive value $\mathbf{x}$ is also roughly Gaussian.

We next focus on the recovery of a single coefficient of the ephemeral secret $\mathbf{x}$. The same attack can be applied identically to extract information on all coefficients.

2.7 Evaluation metric

In order to characterize the worst-case side-channel information leakage of Raccoon's masked implementations, we carry out information theoretic analyzes using the Mutual Information (MI) metric [SMY09]. As consolidated in a sequence of works [ISW03, PR13, DDF14, DFS15, IUH22, BCG+23], this metric can serve as a good proxy for the complexity of Differential Power Analysis (DPA) attacks. It can be computed as:

$$\mathrm{MI}\left[X|L\right] = \mathrm{H}\left[X\right] - \sum_{x} \Pr\left[x\right] \cdot \sum_{l} \Pr\left[l|x\right] \cdot \log_2 \Pr\left[x|l\right],$$

where $\mathrm{H}\left[X\right]$ denotes the entropy of the sensitive value X and $\Pr\left[l|x\right]$ is the conditional Probability Density Function (PDF) of the leakage l given the secret x.

Assuming multivariate leakages with mean vectors μ and additive Gaussian noise with covariance matrix Σ, $\Pr\left[l|x\right]$ can be evaluated as a Gaussian mixture $\sum_{m \in M} \mathcal{N}(x|\mu, m, \Sigma)$, where $\mathcal{M}$ denotes the set of possible shares' combinations. The probabilities $\Pr\left[x|l\right]$ can then be directly derived from $\Pr\left[l|x\right]$ using Bayes' formula as $\frac{\Pr[l|x] \cdot \Pr[x]}{\sum_{x^*} \Pr[l|x^*] \cdot \Pr[x^*]}$.

3 Idealized implementation of AddRepNoise

In order to perform an information theoretic analysis of `AddRepNoise`, we additionally must define the actual leakage that can be collected by the adversary. At high-level, this function works by interleaving noise additions as the sum of uniform random samples and refresh operations for each masked coefficient of the polynomial $\mathbf{r}$. The idealized implementation that we will consider in our evaluations to generate the randomness $\mathbf{r}$ is illustrated in Figure 1 for $d = 2$ shares and a repetition count $rep = 4$ (so $T = 8$).

The shares of each polynomial coefficient of $\mathbf{r}$ are generated by gradually integrating small uniform noise samples from the set $\{-2^{u_w}, \ldots, 2^{u_w} - 1\}$ into a 2-share encoding of zero, denoted as $\mathbf{0_0}$ and $\mathbf{0_1}$. A refresh operation is applied after each such addition and this process is repeated $rep = 4$ times for each share. We next focus on the information leakage per coefficient that corresponds to the internal uniform random samples $[r_{i,rep_0}, \ldots, r_{i,rep_3}]$ during the generation of the randomness polynomial $\mathbf{r}$, which we denote as $l = [l_{i,rep_0}, \ldots, l_{i,rep_3}]$. For this purpose, we consider two leakage models motivated in the introduction: the HW and LSB models. Hence, the leakage of each coefficient is either $l_{i,rep_j} = \mathrm{HW}(r_{i,rep_j}) + \mathcal{N}(0, \sigma^2)$ or $l_{i,rep_j} = \mathrm{LSB}(r_{i,rep_j}) + \mathcal{N}(0, \sigma^2)$ for all i shares and j repetitions, where $\mathcal{N}(0, \sigma^2)$ denotes a simulated Gaussian noise with variance σ^2.[4]

[4] We do not include additional intermediate computations in our model because the refresh operation that precedes them makes them uniform which cancels the low-order leakages we aim to analyse.

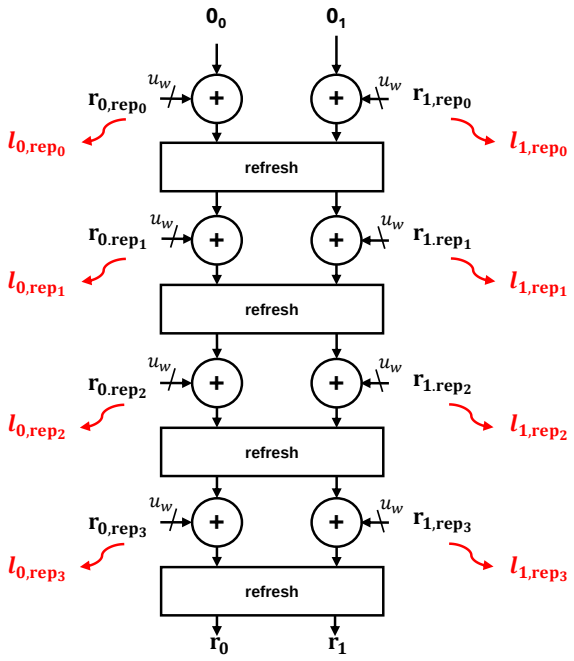

Figure 1: Idealized implementation of `AddRepNoise` with given $d = 2$ and $rep = 4$.

4 Reduced Raccoon

The parameters used in Raccoon, even for security level I, are large and make the direct estimation of the MI prohibitively expensive (in memory and computation). As a result, and in order to put forward how the security of Raccoon's masked implementations scale in function of these parameters, we introduce several simplified instances by decreasing the size of the uniform added noise u_w used to generate the randomness $\mathbf{r}$.

Table 2 summarizes the proposed reduced parameters for different instances. Since the impact of d and rep is the same on the final distribution of the randomness $\mathbf{r}$, we only use their product T as a parameter. For each instance in the table, we analysed what happens when we increase T and adapt u_w accordingly to keep the variance of the $\mathbf{r}$ distribution constant in the same way it is done in the Raccoon specifications.

We considered other simplifications such as omitting the `ZeroEncoding` function (on top of the `refresh` operations – see Footnote 4), since the analysed leakage is located at the uniform random generation step within the `AddRepNoise` function. Finally, in the Raccoon specifications, the distribution of the $\mathbf{x} = \mathbf{c} \circ \mathbf{s}$ coefficients are very close to Gaussian as explained in Section 2.6. In the reduced instances of Raccoon, we assumed for simplicity that these distributions are uniform of size $u_x = 3$ and variance $= 5.25$.

None of these variations affect our main conclusions. First, and in general, the shape of the secret distributions cannot impact the security order of its leakages. Second, and as will be clear next, the quantitative security levels that each instance leads to essentially depend on the size of the distribution (not their shape). Figure 2 shows the scheme of the simplified `AddRepNoise`. Similar to the idealized implementation, the information leakage per coefficient $l = [l_0, \ldots, l_{T-1}]$ corresponds to the uniform random samples of $[r_0, \ldots, r_{T-1}]$ during the generation of the ephemeral randomness polynomial $\mathbf{r}$.

Table 2: Proposed instances with reduced parameters compared to Raccoon.

Instance	T	u_w	size of r	variance of r
1	2/8/32/128	5/4/3/2	6/7/8/9	$1.7e2$
2	2/8/32/128	6/5/4/3	7/8/9/10	$6.8e2$
3	2/8/32/128	7/6/5/4	8/9/10/11	$2.7e3$
4	2/8/32/128	8/7/6/5	9/10/11/12	$1.1e4$
5	2/8/32/128	9/8/7/6	10/11/12/13	$4.4e4$
6	2/8/32/128	10/9/8/7	11/12/13/14	$1.7e5$
7	2/8/32/128	11/10/9/8	12/13/14/15	$7e5$

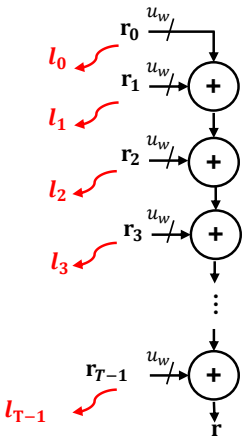

Figure 2: Reduced implementation of `AddRepNoise` with leakage.

We note that the difficulty to estimate the MI for Raccoon's (large) parameters is not a security argument per se, since more efficient (heuristic) attack vectors can be used as surrogates. Interestingly, and as will be clear in Section 7, the trends that information theoretic analyses on smaller instances put forward will allow us to give security arguments for Raccoon's main instances even without relying on such additional heuristics.

5 Global leakages (HW)

We now analyse the noisy leakage security of the `AddRepNoise` operation assuming a (global) HW leakage function. For this purpose, and as illustrated in Figure 2, we consider an adversary who collects leakage samples for the T components of the randomness $\mathbf{r}$.

5.1 First-order leakage on the randomness

First, we evaluate the information theoretic metric $\mathrm{MI}[R|L]$ for two simplified Raccoon instances at different noise levels in Figure 3. For each instance, we analyze the impact of increasing T while adapting the size of the small added noise components (u_w) in order to keep the variance constant (as in Raccoon). The key observation is that all instances lead to observable first-order leakage, reflected by the slope -1 of the information theoretic curves for any T value, which is in contrast with the standard expectation for masked implementations – see Figure 8 in [SVO+10]. This can be explained by the fact that `AddRepNoise` does not use modular additions. Hence, the addition of two uniform n-bit values gives a $n+1$-bit value that is not uniform, leading to first-order leakage.

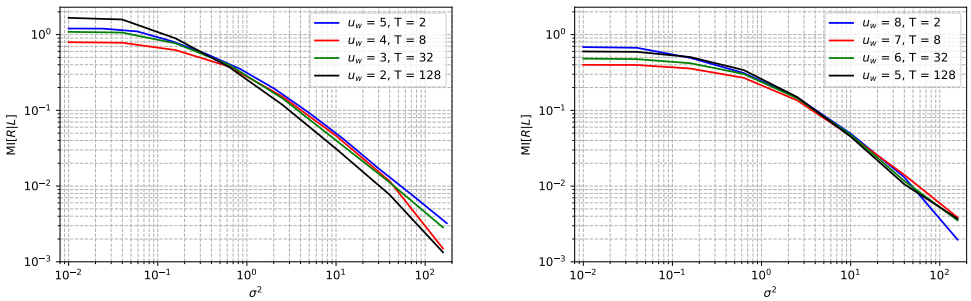

(a) Instance 1: var[r] = 1.7e2 (b) Instance 4: var[r] = 1.09e4

Figure 3: Information theoretic analysis of the randomness $\mathbf{r}$ under HW leakages, for two distinct simplified Raccoon instances and different $T - u_w$ combinations.

Note that the information theoretic evaluation of Figure 3 is for the serial implementation described in Section 4. This is natural choice given that post-quantum signatures schemes like Raccoon or Dilithium have been primarily implemented in software so far. Yet, our conclusions are not specific to this serial context. For example, in a parallel implementation where the leakage of different shares would be summed, the same first-order leakage would be observed. We illustrate this with Figure 4, which exhibits the (first-order) dependency of the leakage mean values in this parallel case (where the leakage distribution is univariate so easy to plot). As a result, the main (generic) difference between serial and parallel implementations is that the latter ones can have a better signal-to-noise ratio, which can result in a shift on the X-axes of our information theoretic plots.

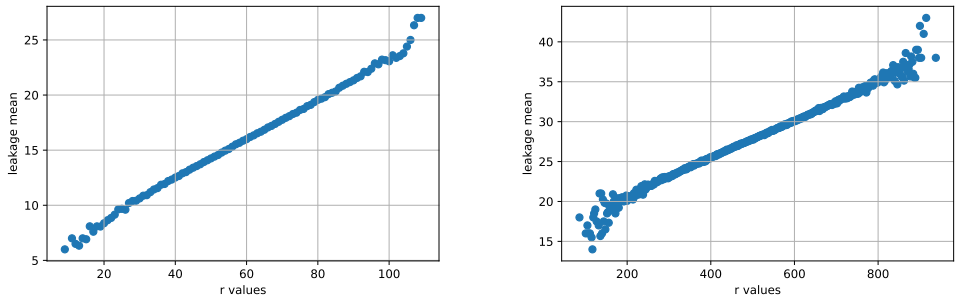

(a) Instance 1: var[r] = 1.7e2 (b) Instance 4: var[r] = 1.09e4

Figure 4: First-order dependency of the randomness under HW leakages ($T = 8$).

5.2 Leakage on the ephemeral secret

The first-order leakage on the randomness $\mathbf{r}$ raises the important question whether it can be turned into exploitable low-order information on the ephemeral secret $\mathbf{x}$ given the knowledge of the signature $\mathbf{z} = \mathbf{r} + \mathbf{x}$. Indeed, it is shown in [BAE+24] that such a leakage can itself be turned into exploitable leakage on the long-term secret $\mathbf{s}$.

In order to answer this question, we next evaluate the information theoretic metric $\mathrm{MI}[X|L, Z]$ for all our simplified Raccoon instances at different noise levels. The results of these investigations are in Figure 5 and lead to two main conclusions.

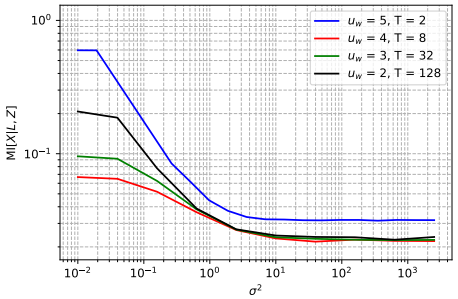

(a) Instance 1: var[r] = 1.7e2

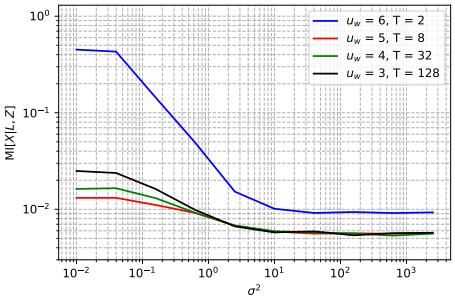

(b) Instance 2: var[r] = 6.82e2

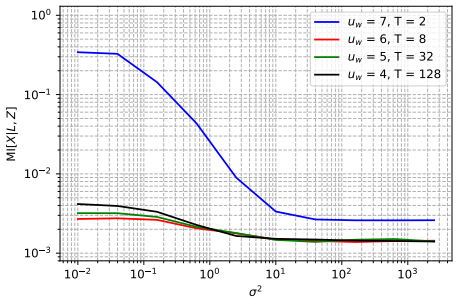

(c) Instance 3: var[r] = 2.73e3

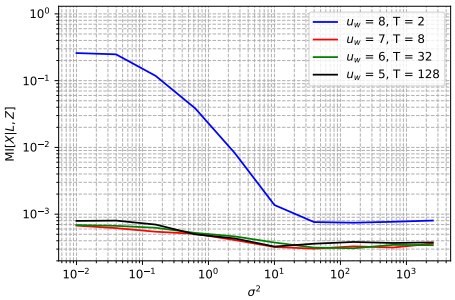

(d) Instance 4: var[r] = 1.09e4

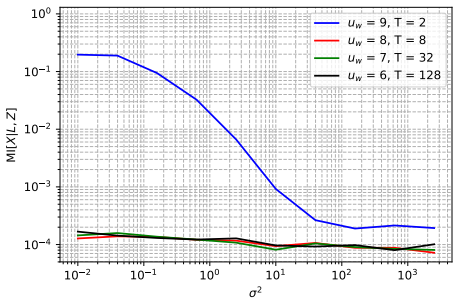

(e) Instance 5: var[r] = 4.37e4

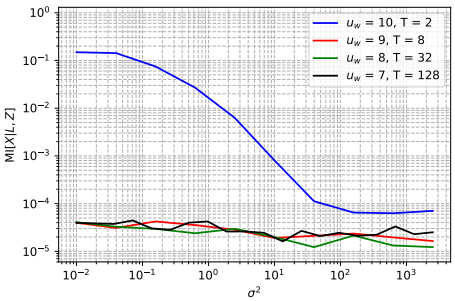

(f) Instance 6: var[r] = 1.75e5

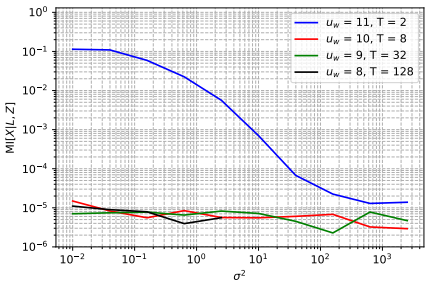

(g) Instance 7: var[r] = 6.99e54

Figure 5: Information theoretic analysis of the ephemeral secret $\mathbf{x}$ under HW leakages, for different simplified Raccoon instances and different $T - u_w$ combinations.

First, and focusing on the Y-axes of the plots, we can see that for sufficient noise levels, the amount of information leakage is limited to a *plateau*, of which the value is determined by the ratio between the variances of the $\mathbf{x}$ and $\mathbf{r}$ distributions. This holds independent of the number of shares considered (i.e., the T value) and has a simple and instructive intuition. Namely, as the noise variance increases, the physical information leaked on $\mathbf{r}$ via side-channels becomes less informative than the mathematical information leakage due to the knowledge of the signature $\mathbf{z}$. Interestingly, mathematical information leakage is precisely what the Raccoon designers have considered in their security analysis and they proved that it is small enough (we will further elaborate on that in conclusions).

Second, and focusing on the X-axes of the plots, we can see that the amount of noise that is needed to reach the plateau region of the information theoretic curves is actually quite modest. Furthermore, this level of noise conveniently decreases as we move towards instances with higher mathematical security (as Raccoon's main instance for example).

It turns out this second observation also has an intuitive explanation. Namely, it results from the need to compress the information on the (large) $\mathbf{r}$ into leakage on the (smaller) $\mathbf{x}$, which grows as the size of the $\mathbf{r}$ distribution grows compared to the one of the $\mathbf{x}$ distribution. In order to illustrate the reduction of information that this compression implies, we plot in Figure 6 the PDF of the leakage samples for all possible values of $\mathbf{r}$ (on the left) and $\mathbf{x}$ (on the right), for Instance 5 at $u_w = 8$ and $T = 8$, in a low-noise setting. While the PDF values of the left plot show clear dependencies on the randomness (confirming the results of Figure 3), the PDF values on the right plot show that the leakage samples conditioned on the ephemeral secret are close to uniform. Combined with the fact that the small information that remains after compression is due to leakage on the LSBs of the randomness, it also implies that the noise needed to "hide" this information is lower than the noise that would be needed to hide the HW of all the randomness bits.

Overall, these results are quite positive for Raccoon. They show that for a practically-relevant leakage function, increasing the noise level and number of shares in its masked implementations is actually not needed to the extent that is anticipated by an analysis in the probing model [dPPRS23]. The number of shares should just be large enough for the sum of uniforms distribution to be approximately Gaussian and the level of noise is significantly below the one needed for Boolean masking in the context of block ciphers. Concretely, the empirical evaluations in this section show that this starts for $T = 8$.

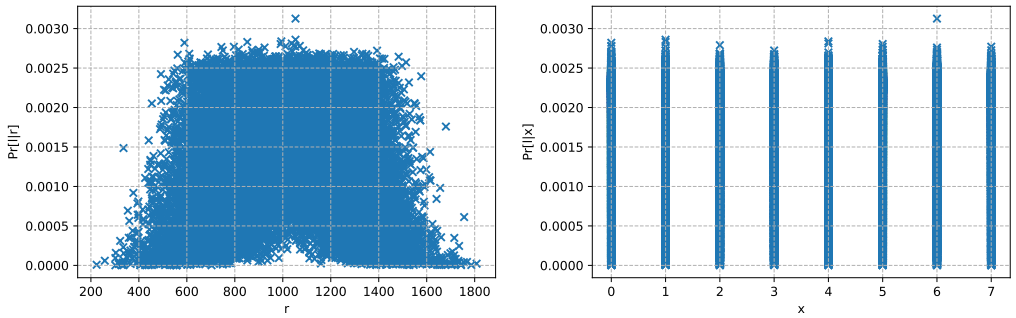

(a) $\Pr[l|r]$ for $u_w = 8$ and $T = 8$.

(b) $\Pr[l|x]$ for $u_w = 8$, $u_x = 3$ and $T = 8$.

Figure 6: Compressing information on $\mathbf{r}$ into information on $\mathbf{x}$.

6 Local leakages (LSB)

As a complement to the analysis of (global) HW leakages in the previous section, we next analyse (local) LSB leakages. The motivation in this section is more theoretical than the previous one, since it is widely acknowledged that extracting bit leakages from implementations is hard with power or electromagnetic measurements and requires different (more expensive) equipment [KGM+21]. Nevetheless, it is interesting to study in order to assess whether the security of Raccoon is maintained in cases where the compression argument of the previous section does not hold. Besides, LSB leakages are also of general interest since bit leakages are worst-case for additive secret sharing schemes [FMM+24].

We therefore repeated the information theoretic evaluation given Figure 5 for HW leakages, this time with LSB leakages. The results are reported in Figure 7.

On the one hand, we can notice that the same plateau as in the previous section pops up once the noise level and number of shares is large enough. On the other hand, the way to reach this plateau is fundamentally different. This time, we can observe a slope (and therefore a statistical security order) that evolves with the number of shares.

Once more, the reason of this phenomenon is simple and instructive. Namely, in the context of bit leakages, the relation $\mathrm{LSB}(r) \equiv \mathrm{LSB}(r_0) + \ldots + \mathrm{LSB}(r_{T-1}) \bmod 2$ holds. As a result, we are back to a situation where the leakage connects the shares according to a modular sum, which is similar to Boolean masking in symmetric cryptography. This explains why increasing the number of shares allows decreasing the information leakage (i.e., amplifying the noise) faster towards the plateau value. It also explains why in the low noise regions of the plots (i.e., in the left parts of the figure), security is actually quite low, since such Boolean masking is insecure without physical noise [MMMS23].

Overall, these results are again positive for Raccoon. They show that even in the case of (admittedly pessimistic) leakage functions such as the LSB, for which the information on the randomness $\mathbf{r}$ must not be compressed to be turned into information on the ephemeral secret $\mathbf{x}$, the standard security arguments of masked implementations take over. Furthermore, the noise needed for masking to be effective remains quite low as well (e.g., does not grow linearly with the size of $\mathbf{r}$), since it only needs to "hide" a few bits.

7 Discussion

As already mentioned, our empirical evaluations are limited by the sizes of the reduced instances for which we are able to evaluate the mutual information (for computational reasons). They also correspond to idealized leakage functions that may not perfectly capture the peculiarities of actual implementations. In this section, we first highlight that the trend we observe for the asymptotic security of masked implementations of reduced Raccoon can be extrapolated, leading to useful conclusions for larger instances. We then show that our positive conclusions for Raccoon are preserved even when considering powerful leakage models such as can be obtained for (quite low-noise) ARM Cortex-M4 devices, and that they do not critically depend on the shape of the leakage functions observed.

7.1 Extrapolation to large instances

For both the HW and LSB leakages, the previous sections put forward that the mutual information that can be collected on the ephemeral secret $\mathbf{x}$ thanks to leakage on the randomness $\mathbf{r}$ is bounded by the mathematical security level of Raccoon, quantified by the ratio between the variance of the randomness $\mathbf{r}$ and the variance of the ephemeral secret $\mathbf{x}$. This asymptotic (so-called plateau) value can be reached very fast, due to a "compression of information" argument for HW leakages and a "noise amplification" argument for LSB

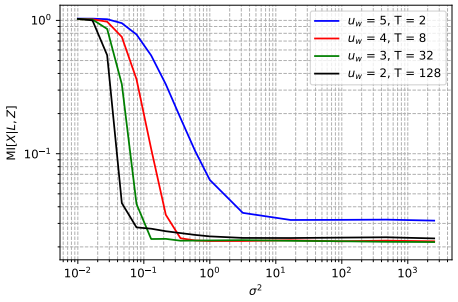

(a) Instance 1: var[r] = 1.7e2

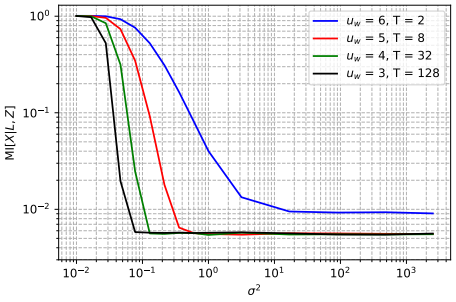

(b) Instance 2: var[r] = 6.82e2

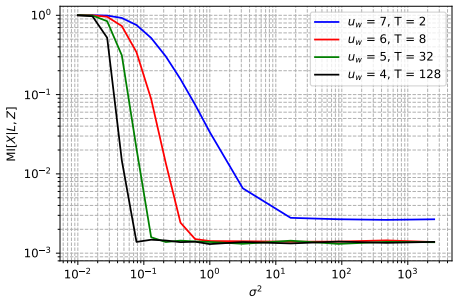

(c) Instance 3: var[r] = 2.73e3

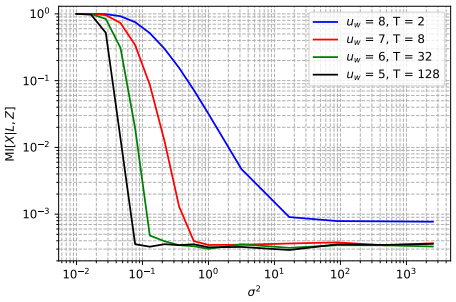

(d) Instance 4: var[r] = 1.09e4

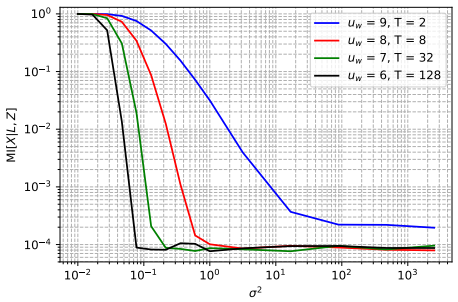

(e) Instance 5: var[r] = 4.37e4

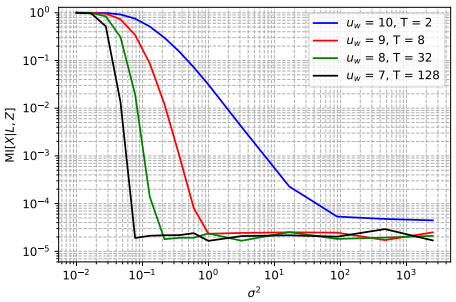

(f) Instance 6: var[r] = 1.75e5

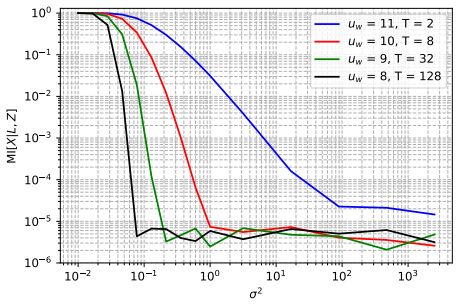

(g) Instance 7: var[r] = 6.98e5

Figure 7: Information theoretic analysis of the ephemeral secret $\mathbf{x}$ under LSB leakages, for different simplified Raccoon instances and different $T - u_w$ combinations.

leakages. In Figure 8, we plot this asymptotic/plateau value for the reduced instances we investigated, and notice that the dependency of the mutual information on the variance ratio can be easily extrapolated for the parameters of Raccoon's main instance.

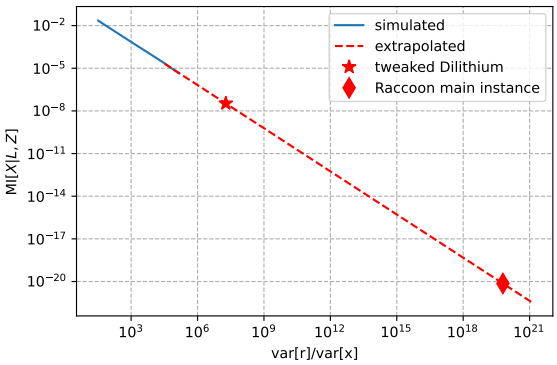

Figure 8: Asymptotic information theoretic analysis (i.e., "plateau value" for large enough number of shares and noise) as a function of the ratio between the variance of the randomness $\mathbf{r}$ and the variance of the ephemeral secret $\mathbf{x}$ and extrapolation.

Interestingly, such an extrapolation also suggests a kind of "no-free lunch" theorem. Namely, the good properties of Raccoon's masked implementations are indeed due to the larger parameters it uses. As a result, tweaking Dilithium in order to generate its shares with an `AddRepNoise`-like function, while maintaining the size of its randomness, ephemeral secret and signature, would lead to an insufficiently low information leakage (i.e., a plateau value that corresponds to concretely exploitable leakage). We note that this is is precisely the reason why Dilithium is only secure with rejection sampling.

So eventually, the interest of Raccoon over Dilithium will essentially depend on the tradeoff between their target physical security level and the cost of their implementations, with Raccoon gaining interest when the former increases. Accurately answering this question would require systematizing the variety of attack vectors introduced against Dilithium into a sound evaluation methodology, in order to efficiently estimate their worst-case security in practice, which is an interesting scope for further research. As a first step in this direction, we next show that the trends put forward by the previous information theoretic analyses for idealized leakage functions holds for more concrete ones.

7.2 Impact for concrete leakage functions

Since the results of Sections 5 and 6 show that Raccoon's implementation security guarantees evolve differently in the presence of global or local leakage, a natural question is whether the leakage models observed in practice are closer to one or the other case. In this section, we therefore repeat Raccoon's information theoretic evaluations using the profiled model applied in [CDSU23] to an ARM Cortex-M4 device. This model is interesting for two reasons. First, it is a quite powerful one, combining dimensionality reduction and linear regression to estimate accurate templates even for large target intermediate values. Second, it is applied to a challenging (low-noise) target for secure implementation. As a result, obtaining positive conclusions for such a model would be quite encouraging for the physical security of Raccoon. Besides, the authors of [CDSU23] performed atomic experiments and estimated models for buses of increasing bit widths, which is convenient for our purposes where the leakages of intermediate computations with different bus sizes

are considered for the different reduced instances that we analyzed. Their work also comes with an open source code which allows generating leakage models of the shape:

$$\mathsf{L}(r) = \sum_{i=1}^{|r|} \beta_i \cdot r(i) + \mathcal{N}(0, \sigma^2), \tag{1}$$

with $r(i)$ the ith bit of the coefficient r and $\mathcal{N}(0, \sigma^2)$ a Gaussian noise with variance σ^2. Such models generalize our previous experiments since the β_i coefficients are all equal to one in the HW case, and all but one are equal to zero in the LSB case.

Figure 9 shows the results of information theoretic evaluations for the ephemeral secret $\mathbf{x}$ integrated in Instances 4, 5 and 6, using regression-based models of the appropriate bitsizes and normalized so that the actual measurement noise is one in all cases.

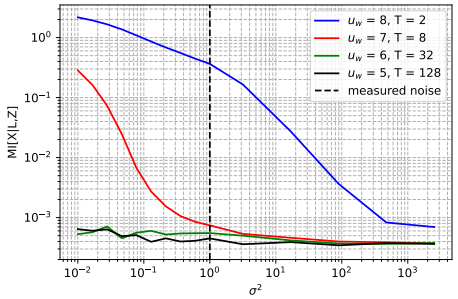

(a) Instance 4: var[r] = 1.09e4

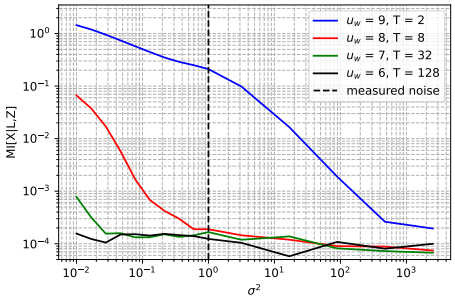

(b) Instance 5: var[r] = 4.37e4

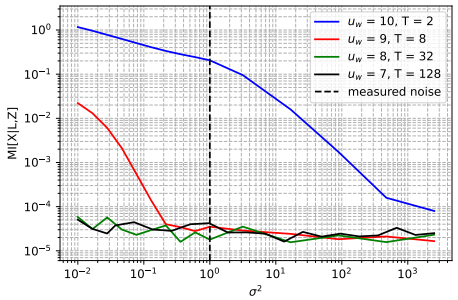

(c) Instance 6: var[r] = 1.75e5

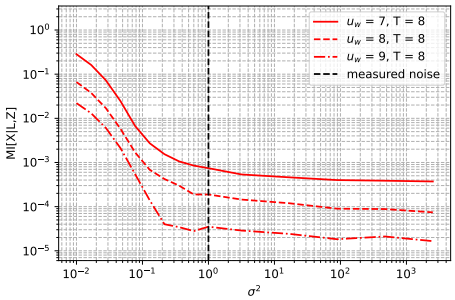

(d) Instance 4,5,6 at T = 8

Figure 9: Information theoretic analysis of the ephemeral secret $\mathbf{x}$ under ARM Cortex-4 leakages, for different simplified Raccoon instances and different $T - u_w$ combinations.

These results essentially show that for very low noise levels (i.e., below the ones observed for a low-noise target like an ARM Cortex-M4), the regression-based model is closer to the (local) LSB one. This is because a regression-based model without noise is bijective (or close enough to), a context where no security is possible at all. However, and more positively, the plots also show that for a modest amount of noise, the information leakage on the ephemeral secret $\mathbf{x}$ rapidly decreases in a way that is quite similar to what we observed for the (global) HW leakages. Interestingly, such a level of noise is already observed for the ARM Cortex-M4 device we consider (remember that all plots in Figure 9 are normalized so that the measured noise level is always one). This is because the fine-grain characterization of the β_i coefficient that makes the regression-based model differ from the HW one (i.e., the fact that every bit in the models consume in a slightly different manner) is more easily (i.e.,

require less noise to be) hidden than the coarser-grain differences due to the very activity of these bits or lack thereof. So these results suggest that even in the context of low-noise embedded microcontrollers, where the application of masking is usually challenging [BS21], implementing Raccoon with $T = 8$ (i.e., its smallest physical security parameters) is likely to be sufficient for the leakage of `AddRepNoise` not to be a critical security threat. This last conclusion is visually confirmed by the lower right (d) plot of Figure 9.

For completeness, we also analyzed two other leakage models. First, a model extracted from ELMO [MOW17], which simulates an accurate characterization of the ARM Cortex-M0 processor that features the weighted Hamming weight/distance on the data bus and second-degree bit interactions between bits/bit flips within the same operand (valid only in case of multiplications and shift instructions, so we extracted leakage models for the outputs of multiplication operations). Second, an idealized model where the quadratic terms of adjacent wires have the same amplitude as the linear ones, defined as:

$$\mathsf{L}(r) = (\sum_{i=1}^{|r|} r(i) + \sum_{i=1}^{|r|-1} r(i) \cdot r(i+1)) + \mathcal{N}(0, \sigma^2), \tag{2}$$

Concretely, we then reproduced Figure 9 (d), where the information theoretic metric is evaluated for the ephemeral secret $\mathbf{x}$ integrated in Instances 4, 5 and 6, using both the ELMO power model as shown in Figure 10 (a) and the idealized quadratic model as shown in Figure 10 (b). Both plots confirm that our conclusions are not significantly affected by the presence of quadratic terms in the power model that may appear in practice.

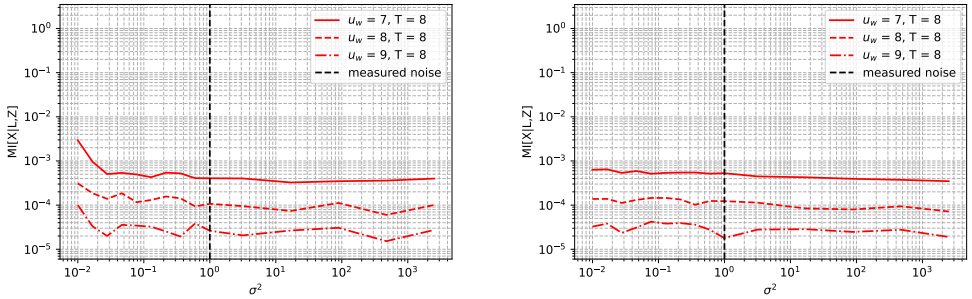

(a) ELMO leakage.

(b) Simple quadratic model leakage.

Figure 10: Information theoretic analysis of the ephemeral secret $\mathbf{x}$ under ELMO and simple quadratic model leakages, for the simplified Raccoon instances 4,5,6 at T = 8.

8 Conclusions

Implementing post-quantum cryptography efficiently is in general an important challenge, which is made more difficult if side-channel attacks are a concern. In view of the high overheads needed to protect future standard algorithms like CRYSTALS-Kyber or Dilithium with state-of-the-art countermeasures, long-term research has been initiated to identify properties that could facilitate the secure implementation of new algorithms [dPPRS23, HLM+23]. This research being in an early stage, it remains with gaps to fill in order to connect theoretical analyzes and practical security guarantees in a sound manner.

In this paper, we expose the (excellent) side-channel security guarantees that Raccoon, a post-quantum signature scheme designed with physical security in mind, offers. While

Raccoon's original security analysis in the abstract probing model already outlined such positive features, our results are in more concrete leakage models and therefore offer an interesting complementary view. Informally, rather than studying the black box and probing security of Raccoon in isolation, its designers performed this analysis jointly. This allowed them to show that the security of Raccoon with probes can be reduced to instances of Raccoon with smaller parameters. Our information theoretic evaluation of the `AddRepNoise` function, which is at the core of Raccoon's good leakage features, shows that this probing security analysis might be conservative. Precisely, it shows that by adding a mild amount of noise to idealized Hamming weight (or LSB) leakages, it is possible to make the information on Raccoon's long-term secret negligible, and this observation extends to more concrete leakage functions estimated from ARM Cortex devices.

Such results suggest a formalization in the (random probing or) noisy leakage model(s) as main open problem. That is, rather than providing an information theoretic evaluation for various leakage functions, it could be proven that Raccoon's implementation security is guaranteed for any leakage function satisfying some (noise and independence) requirement. Such an analysis could then formally confirm that increasing the number of shares as encouraged by a probing security analysis (and less by our information theoretic evaluations) is indeed not necessary beyond Raccoon's smallest physical security parameters.

Acknowledgments. François-Xavier Standaert is a research director of the Belgian Fund for Scientific Research (FNRS-F.R.S.). This work has been funded in part by the ERC Advanced Grant 101096871 (BRIDGE). Views and opinions expressed are those of the authors and do not necessarily reflect those of the European Union or the ERC. Neither the European Union nor the granting authority can be held responsible for them.

References

[ABC+23] Melissa Azouaoui, Olivier Bronchain, Gaëtan Cassiers, Clément Hoffmann, Yulia Kuzovkova, Joost Renes, Tobias Schneider, Markus Schönauer, François-Xavier Standaert, and Christine van Vredendaal. Protecting dilithium against leakage revisited sensitivity analysis and improved implementations. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2023(4):58–79, 2023. URL: https://doi.org/10.46586/tches.v2023.i4.58-79, doi:10.46586/TCHES.V2023.I4.58-79.

[BAE+24] Olivier Bronchain, Melissa Azouaoui, Mohamed ElGhamrawy, Joost Renes, and Tobias Schneider. Exploiting small-norm polynomial multiplication with physical attacks application to crystals-dilithium. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2024(2):359–383, 2024. URL: https://doi.org/10.46586/tches.v2024.i2.359-383, doi:10.46586/TCHES.V2024.I2.359-383.

[BBE+18] Gilles Barthe, Sonia Belaïd, Thomas Espitau, Pierre-Alain Fouque, Benjamin Grégoire, Mélissa Rossi, and Mehdi Tibouchi. Masking the GLP lattice-based signature scheme at any order. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part II*, volume 10821 of *Lecture Notes in Computer Science*, pages 354–384. Springer, 2018. doi:10.1007/978-3-319-78375-8_12.

[BCG+23] Julien Béguinot, Wei Cheng, Sylvain Guilley, Yi Liu, Loïc Masure, Olivier Rioul, and François-Xavier Standaert. Removing the field size loss from duc et al.'s conjectured bound for masked encodings. In Elif Bilge Kavun and

Michael Pehl, editors, *Constructive Side-Channel Analysis and Secure Design - 14th International Workshop, COSADE 2023, Munich, Germany, April 3-4, 2023, Proceedings*, volume 13979 of *Lecture Notes in Computer Science*, pages 86–104. Springer, 2023. doi:10.1007/978-3-031-29497-6_5.

[BS21] Olivier Bronchain and François-Xavier Standaert. Breaking masked implementations with many shares on 32-bit software platforms or when the security order does not matter. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2021(3):202–234, 2021. URL: https://doi.org/10.46586/tches.v2021.i3.202-234, doi:10.46586/TCHES.V2021.I3.202-234.

[BVC+23] Alexandre Berzati, Andersson Calle Viera, Maya Chartouny, Steven Madec, Damien Vergnaud, and David Vigilant. Exploiting intermediate value leakage in dilithium: A template-based approach. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2023(4):188–210, 2023. URL: https://doi.org/10.46586/tches.v2023.i4.188-210, doi:10.46586/TCHES.V2023.I4.188-210.

[CDSU23] Gaëtan Cassiers, Henri Devillez, François-Xavier Standaert, and Balazs Udvarhelyi. Efficient regression-based linear discriminant analysis for side-channel security evaluations towards analytical attacks against 32-bit implementations. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2023(3):270–293, 2023. URL: https://doi.org/10.46586/tches.v2023.i3.270-293, doi:10.46586/TCHES.V2023.I3.270-293.

[CJRR99] Suresh Chari, Charanjit S. Jutla, Josyula R. Rao, and Pankaj Rohatgi. Towards sound approaches to counteract power-analysis attacks. In Michael J. Wiener, editor, *Advances in Cryptology - CRYPTO '99, 19th Annual International Cryptology Conference, Santa Barbara, California, USA, August 15-19, 1999, Proceedings*, volume 1666 of *Lecture Notes in Computer Science*, pages 398–412. Springer, 1999. doi:10.1007/3-540-48405-1_26.

[DDF14] Alexandre Duc, Stefan Dziembowski, and Sebastian Faust. Unifying leakage models: From probing attacks to noisy leakage. In Phong Q. Nguyen and Elisabeth Oswald, editors, *Advances in Cryptology - EUROCRYPT 2014 - 33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Copenhagen, Denmark, May 11-15, 2014. Proceedings*, volume 8441 of *Lecture Notes in Computer Science*, pages 423–440. Springer, 2014. doi:10.1007/978-3-642-55220-5_24.

[DFS15] Alexandre Duc, Sebastian Faust, and François-Xavier Standaert. Making masking security proofs concrete - or how to evaluate the security of any leaking device. In Elisabeth Oswald and Marc Fischlin, editors, *Advances in Cryptology - EUROCRYPT 2015 - 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Sofia, Bulgaria, April 26-30, 2015, Proceedings, Part I*, volume 9056 of *Lecture Notes in Computer Science*, pages 401–429. Springer, 2015. doi:10.1007/978-3-662-46800-5_16.

[DFS16] Stefan Dziembowski, Sebastian Faust, and Maciej Skórski. Optimal amplification of noisy leakages. In Eyal Kushilevitz and Tal Malkin, editors, *Theory of Cryptography - 13th International Conference, TCC 2016-A, Tel Aviv, Israel, January 10-13, 2016, Proceedings, Part II*, volume 9563 of *Lecture Notes in Computer Science*, pages 291–318. Springer, 2016. doi:10.1007/978-3-662-49099-0_11.

[dPPRS23] Rafaël del Pino, Thomas Prest, Mélissa Rossi, and Markku-Juhani O. Saarinen. High-order masking of lattice signatures in quasilinear time. In *44th IEEE*

Symposium on Security and Privacy, SP 2023, San Francisco, CA, USA, May 21-25, 2023, pages 1168–1185. IEEE, 2023. doi:10.1109/SP46215.2023.10179342.

[FMM+24] Sebastian Faust, Loïc Masure, Elena Micheli, Maximilian Orlt, and François-Xavier Standaert. Connecting leakage-resilient secret sharing to practice: Scaling trends and physical dependencies of prime field masking. In Marc Joye and Gregor Leander, editors, *Advances in Cryptology - EUROCRYPT 2024 - 43rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zurich, Switzerland, May 26-30, 2024, Proceedings, Part IV*, volume 14654 of *Lecture Notes in Computer Science*, pages 316–344. Springer, 2024. doi:10.1007/978-3-031-58737-5_12.

[GP99] Louis Goubin and Jacques Patarin. DES and differential power analysis (the "duplication" method). In Çetin Kaya Koç and Christof Paar, editors, *Cryptographic Hardware and Embedded Systems, First International Workshop, CHES'99, Worcester, MA, USA, August 12-13, 1999, Proceedings*, volume 1717 of *Lecture Notes in Computer Science*, pages 158–172. Springer, 1999. doi:10.1007/3-540-48059-5_15.

[HLM+23] Clément Hoffmann, Benoît Libert, Charles Momin, Thomas Peters, and François-Xavier Standaert. POLKA: towards leakage-resistant post-quantum cca-secure public key encryption. In Alexandra Boldyreva and Vladimir Kolesnikov, editors, *Public-Key Cryptography - PKC 2023 - 26th IACR International Conference on Practice and Theory of Public-Key Cryptography, Atlanta, GA, USA, May 7-10, 2023, Proceedings, Part I*, volume 13940 of *Lecture Notes in Computer Science*, pages 114–144. Springer, 2023. doi:10.1007/978-3-031-31368-4_5.

[ISW03] Yuval Ishai, Amit Sahai, and David A. Wagner. Private circuits: Securing hardware against probing attacks. In Dan Boneh, editor, *Advances in Cryptology - CRYPTO 2003, 23rd Annual International Cryptology Conference, Santa Barbara, California, USA, August 17-21, 2003, Proceedings*, volume 2729 of *Lecture Notes in Computer Science*, pages 463–481. Springer, 2003. doi:10.1007/978-3-540-45146-4_27.

[IUH22] Akira Ito, Rei Ueno, and Naofumi Homma. On the success rate of side-channel attacks on masked implementations: Information-theoretical bounds and their practical usage. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security, CCS 2022, Los Angeles, CA, USA, November 7-11, 2022*, pages 1521–1535. ACM, 2022. doi:10.1145/3548606.3560579.

[KGM+21] Thilo Krachenfels, Fatemeh Ganji, Amir Moradi, Shahin Tajik, and Jean-Pierre Seifert. Real-world snapshots vs. theory: Questioning the t-probing security model. In *42nd IEEE Symposium on Security and Privacy, SP 2021, San Francisco, CA, USA, 24-27 May 2021*, pages 1955–1971. IEEE, 2021. doi:10.1109/SP40001.2021.00029.

[LBS19] Itamar Levi, Davide Bellizia, and François-Xavier Standaert. Reducing a masked implementation's effective security order with setup manipulations and an explanation based on externally-amplified couplings. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2019(2):293–317, 2019. URL: https://doi.org/10.13154/tches.v2019.i2.293-317, doi:10.13154/TCHES.V2019.I2.293-317.

[LZS+21] Yuejun Liu, Yongbin Zhou, Shuo Sun, Tianyu Wang, Rui Zhang, and Jingdian Ming. On the security of lattice-based fiat-shamir signatures in the presence of randomness leakage. *IEEE Trans. Inf. Forensics Secur.*, 16:1868–1879, 2021. doi:10.1109/TIFS.2020.3045904.

[MGTF19] Vincent Migliore, Benoît Gérard, Mehdi Tibouchi, and Pierre-Alain Fouque. Masking dilithium - efficient implementation and side-channel evaluation. In Robert H. Deng, Valérie Gauthier-Umaña, Martín Ochoa, and Moti Yung, editors, *Applied Cryptography and Network Security - 17th International Conference, ACNS 2019, Bogota, Colombia, June 5-7, 2019, Proceedings*, volume 11464 of *Lecture Notes in Computer Science*, pages 344–362. Springer, 2019. doi:10.1007/978-3-030-21568-2_17.

[MMMS23] Loïc Masure, Pierrick Méaux, Thorben Moos, and François-Xavier Standaert. Effective and efficient masking with low noise using small-mersenne-prime ciphers. In Carmit Hazay and Martijn Stam, editors, *Advances in Cryptology - EUROCRYPT 2023 - 42nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Lyon, France, April 23-27, 2023, Proceedings, Part IV*, volume 14007 of *Lecture Notes in Computer Science*, pages 596–627. Springer, 2023. doi:10.1007/978-3-031-30634-1_20.

[MOP07] Stefan Mangard, Elisabeth Oswald, and Thomas Popp. *Power analysis attacks - revealing the secrets of smart cards*. Springer, 2007.

[MOW17] David McCann, Elisabeth Oswald, and Carolyn Whitnall. Towards practical tools for side channel aware software engineering: 'grey box' modelling for instruction leakages. In Engin Kirda and Thomas Ristenpart, editors, *26th USENIX Security Symposium, USENIX Security 2017, Vancouver, BC, Canada, August 16-18, 2017*, pages 199–216. USENIX Association, 2017. URL: https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/mccann.

[PR13] Emmanuel Prouff and Matthieu Rivain. Masking against side-channel attacks: A formal security proof. In Thomas Johansson and Phong Q. Nguyen, editors, *Advances in Cryptology - EUROCRYPT 2013, 32nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Athens, Greece, May 26-30, 2013. Proceedings*, volume 7881 of *Lecture Notes in Computer Science*, pages 142–159. Springer, 2013. doi:10.1007/978-3-642-38348-9_9.

[RJH+18] Prasanna Ravi, Mahabir Prasad Jhanwar, James Howe, Anupam Chattopadhyay, and Shivam Bhasin. Side-channel assisted existential forgery attack on dilithium - A NIST PQC candidate. *IACR Cryptol. ePrint Arch.*, page 821, 2018. URL: https://eprint.iacr.org/2018/821.

[SM16] Tobias Schneider and Amir Moradi. Leakage assessment methodology - extended version. *J. Cryptogr. Eng.*, 6(2):85–99, 2016. URL: https://doi.org/10.1007/s13389-016-0120-y, doi:10.1007/S13389-016-0120-Y.

[SMY09] François-Xavier Standaert, Tal Malkin, and Moti Yung. A unified framework for the analysis of side-channel key recovery attacks. In Antoine Joux, editor, *Advances in Cryptology - EUROCRYPT 2009, 28th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cologne, Germany, April 26-30, 2009. Proceedings*, volume 5479 of *Lecture Notes in Computer Science*, pages 443–461. Springer, 2009. doi:10.1007/978-3-642-01001-9_26.

[SVO+10] François-Xavier Standaert, Nicolas Veyrat-Charvillon, Elisabeth Oswald, Benedikt Gierlichs, Marcel Medwed, Markus Kasper, and Stefan Mangard. The world is not enough: Another look on second-order DPA. In Masayuki Abe, editor, *Advances in Cryptology - ASIACRYPT 2010 - 16th International Conference on the Theory and Application of Cryptology and Information Security, Singapore, December 5-9, 2010. Proceedings*, volume 6477 of *Lecture Notes in Computer Science*, pages 112–129. Springer, 2010. doi:10.1007/978-3-642-17373-8_7.

[UMTS22] Vincent Quentin Ulitzsch, Soundes Marzougui, Mehdi Tibouchi, and Jean-Pierre Seifert. Profiling side-channel attacks on dilithium - A small bit-fiddling leak breaks it all. In Benjamin Smith and Huapeng Wu, editors, *Selected Areas in Cryptography - 29th International Conference, SAC 2022, Windsor, ON, Canada, August 24-26, 2022, Revised Selected Papers*, volume 13742 of *Lecture Notes in Computer Science*, pages 3–32. Springer, 2022. doi:10.1007/978-3-031-58411-4_1.