




Raccoon: A Masking-Friendly Signature Proven in the Probing Model

Rafaël del Pino¹, Shuichi Katsumata^{1,2} 
, Thomas Prest¹ , and Mélissa Rossi³ 

¹ PQShield (`firstname.lastname@pqshield.com`)

² AIST

³ ANSSI (`firstname.lastname@ssi.gouv.fr`)

Abstract. This paper presents Raccoon, a lattice-based signature scheme submitted to the NIST 2022 call for additional post-quantum signatures. Raccoon has the specificity of always being masked. Concretely, all sensitive intermediate values are shared into d parts. The main design rationale of Raccoon is to be easy to mask at high orders, and this dictated most of its design choices, such as the introduction of new algorithmic techniques for sampling small errors. As a result, Raccoon achieves a masking overhead $O(d \log d)$ that compares favourably with the overheads $O(d^2 \log q)$ observed when masking standard lattice signatures.

In addition, we formally prove the security of Raccoon in the t -probing model: an attacker is able to probe $t \leq d-1$ shares during each execution of the main algorithms (key generation, signing, verification). While for most cryptographic schemes, the black-box t -probing security can be studied in isolation, in Raccoon this analysis is performed jointly.

To that end, a bridge must be made between the black-box game-based EUF-CMA proof and the usual simulation proofs of the ISW model (CRYPTO 2003). We formalize an end-to-end masking proof by deploying the probing EUF-CMA introduced by Barthe et al. (Eurocrypt 2018) and exhibiting the simulators of the non-interference properties (Barthe et al. CCS 2016). The proof is divided into three novel parts:

- a simulation proof in the ISW model that allows to propagate the dependency to a restricted number of inputs and random coins,
- a game-based proof showing that the security of Raccoon with probes can be reduced to an instance of Raccoon with smaller parameters,
- a parameter study to ensure that the smaller instance is secure, using a robust generalization of the Rényi divergence.

While we apply our techniques to Raccoon, we expect that the algorithmic and proof techniques we introduce will be helpful for the design and analysis of future masking-friendly schemes.

Keywords: Raccoon signature; t -probing model; side-channel attacks.

1 Introduction

In the past decade, post-quantum cryptography has reached quickly grown from a mostly theoretical field to one with sufficient maturity to be deployed on a wide

scale. This is epitomized by NIST’s standardization in 2020 of the hash-based signatures XMSS and LMS, as well as its announcement in 2022 of the future standardization of the lattice-based KEM Kyber, the lattice-based signatures Dilithium and Falcon, and the hash-based signature SPHINCS+. Whilst the efficiency profiles and black-box security of these schemes are well-understood, resistance against side-channel attacks remains a weak spot.

Side-channel attacks. In a side-channel attack (SCA), an attacker can learn information about the physical execution of an algorithm, such as its running time or its effect on the power consumption, electromagnetic or acoustic emission of the device running it. This information can then be leveraged to recover sensitive information, for example, cryptographic keys.

SCAs can be devastating against cryptographic implementations, and post-quantum schemes are no exception. See Section 1.3 for references of concrete SCAs against Dilithium.

Masking. The main countermeasure against side-channel attacks is masking [27]. It consists of splitting sensitive information in d shares (concretely: $x = x_0 + \dots + x_{d-1}$), and performing secure computation using MPC-based techniques. Masking provides a *trade-off* between efficiency and SCA resistance: the computational efficiency of the implementation is reduced by a polynomial factor in d , but the cost of a side-channel attack is expected to grow exponentially [19,28].

Unfortunately, lattice-based signatures contain subroutines that are extremely expensive to mask, such as (a) sampling from a small set, (b) bit-decomposition, and (c) rejection sampling. Currently, the best known ways to perform these operations is to rely on mask conversions [26,13], which convert between arithmetic and boolean masking. This typically incurs an overhead $O(d^2 \log q)$ [14] or $O(2^{d/2})$ [11], and quickly becomes the efficiency bottleneck. As an illustration, the only publicly available *masked* implementation of Dilithium [12] is 53 (resp. 200) times slower than unmasked Dilithium for $d = 2$ (resp. $d = 4$).

Masking-friendly schemes. In order to overcome these limitations, a natural research direction is to design lattice-based signatures that are naturally amenable to masking. However, this is easier said than done. The few designs that exist have either been shown insecure or lack a formal security proof, see Section 1.3 for a more detailed discussion. Thus having a masking-friendly signature with a formal proof has been an elusive goal.

1.1 Our Contributions

We propose Raccoon, a masking-friendly signature, and provide a formal security proof in the t -probing model [27]. While Raccoon is inspired from the similarly named scheme from [17], we have heavily modified its design in order to make it more efficient and provably secure under standard assumptions. The design presented in this paper is exactly the same as the one submitted to the NIST on-ramp standardization campaign [16].

Blueprint. Raccoon is based on the “Lyubashevsky signature without aborts” blueprint, also found in works on threshold signatures [1], and which we recall below. Assume the public key vk is a Learning With Errors (LWE) sample $(\mathbf{A}, \mathbf{t} = [\mathbf{A} \ \mathbf{I}] \cdot \mathbf{s})$, where \mathbf{s} is a small vector, \mathbf{I} is the identity matrix and \mathbf{A} is a uniform matrix (precise definitions will be provided later in the paper). Signing proceeds as follows:

- (S1) Sample \mathbf{r} , compute a commitment $\mathbf{w} = [\mathbf{A} \ \mathbf{I}] \cdot \mathbf{r}$;
- (S2) Compute a challenge $c = H(\mathbf{w}, \text{vk}, \text{msg})$;
- (S3) Compute a response $\mathbf{z} = \mathbf{s} \cdot c + \mathbf{r}$.

The verification procedure checks that $H(\mathbf{A} \cdot \mathbf{z} - \mathbf{t} \cdot c, \text{vk}, \text{msg}) = c$ and that \mathbf{z} is short. Using a Rényi divergence argument, we can argue security if the modulus q grows as the square root of the number of queries Q_s , that is $q = \Omega(\sqrt{Q_s})$. By eliminating the need for rejection sampling, this sidesteps the issue of masking it. In addition, unlike in Dilithium, the security argument does not rely on bit-decomposition. This eliminates the need to *mask* bit-dropping, which we now employ purely for efficiency reasons. We note that our final modulus has 49 bits, which is larger than the standard precision (32-bit or less) on many embedded platforms. We mitigate this by taking $q = q_1 \cdot q_2$, where q_1 and q_2 are 24-bit and 25-bit NTT-friendly prime moduli.

We note that rejection sampling in Dilithium requires a smaller modulus $q = \Omega(\dim(\mathbf{s}))$, in practice $\log q \approx 23$ in Dilithium. Our design choice entails a trade-off between compactness (Dilithium) and ease of masking (Raccoon).

The problem with Gaussians. Standard Rényi divergence arguments as in [1] require \mathbf{r} to be sampled from a discrete Gaussian distribution. However, Gaussians are notoriously difficult to generate in a way that is robust to SCA. The most common method for sampling Gaussians in a constant-time manner is via probability distribution tables (PDT), see for example FrodoKEM [34] or Falcon [36]. For signatures, the PDT would require a precision $p \approx \log(Q_s)$, for example Falcon takes $p = 72$. Masking this step would incur a prohibitive overhead $O(d^2 \log q)$. Similarly, all other existing sampling methods (see e.g. “Related works” in [25]) comprise at least one step that is expensive to mask. We *could* use Gaussians, and from a purely theoretic perspective the security proof would go through, but from a practical point of view this would show little relevance to the real-world issues that masking is trying to solve in the first place.

Sums of uniforms. Our solution is to pick a distribution that has Gaussian-style properties, but is easier to sample securely on embedded devices. As it turns out, sampling \mathbf{r} as a sum of uniform variates (over a small set) produces remarkably Gaussian-like distributions, which is unsurprising and a straightforward consequence of the central limit theorem. Unfortunately, standard Rényi divergence arguments fail for these distributions since they have finite support.

We resolve this analytical issue by introducing the *smooth Rényi divergence*, a more robust generalization of the Rényi divergence that is able to provide cryptographically useful statements about sums of uniform distributions. We

define it as a simple combination of the statistical distance and the Rényi divergence. This generalisation achieves the best of both worlds: the robustness of the statistical distance and the power of the Rényi divergence.

Probing-resilient sampling via AddRepNoise. Now that we have identified a suitable distribution (that is, sum of uniforms) for \mathbf{r} , the final step is to sample it in a way that is resilient to t -probing adversaries. A naive approach would be to sample in parallel each share \mathbf{r}_i of $\llbracket \mathbf{r} \rrbracket$ as the sum of rep small uniform variates, so that \mathbf{r} is the sum of $d \cdot \text{rep}$ small uniform variates. However, a probing adversary is allowed to probe $t \leq d - 1$ individual shares \mathbf{r}_i . This would reduce the standard deviation of the conditional distribution of \mathbf{r} by a factor \sqrt{d} , and lead to worse parameters.

We resolve this by proposing a new algorithm, called **AddRepNoise**, which interleaves (a) parallel generation of individual noises and (b) refreshing the masked vector, and repeats this rep times. We can formally prove that a t -probing adversary only learns t individual uniform variates, so that the standard deviation of \mathbf{r} conditioned to these variates is the sum of $d \cdot \text{rep} - d + 1$ uniform variates, which allows to prove security with a minimal loss in tightness.

1.2 Overview of the Security Proof

We recall that a high-level description of Raccoon is given in Section 1.1. Now, in a masked form, the secret is shared as $\mathbf{s} = \sum_{i \in [d]} \mathbf{s}_i$ where the coefficients of the vectors \mathbf{s}_i are sampled in a short interval. This is a deliberate choice of Raccoon that allows good sampling performance.

At first sight, if the \mathbf{s}_i are safely manipulated in the signature algorithm and never recombined, the masking security seems guaranteed as the exact value of \mathbf{s} cannot be recombined. However, if an adversary probes $d - 1$ shares of \mathbf{s}_i , say $\{\mathbf{s}_0, \dots, \mathbf{s}_{d-2}\}$, he can compute $\mathbf{vk}' = \mathbf{vk} - [\mathbf{A} \ \mathbf{I}] \sum_{i=0}^{d-2} \cdot \mathbf{s}_i = [\mathbf{A} \ \mathbf{I}] \mathbf{s}_{d-1}$. Key recovery is significantly easier as the updated secret is now from a narrower distribution. Hence, while the exact value of \mathbf{s} is inaccessible, the knowledge of the probes combined with the knowledge of the public key can lead to a simpler key recovery. This aspect makes a link between two families of proofs that are typically separated in other works: the black-box game-based EUF-CMA proofs and the simulation proofs of masking. The former quantifies the advantage of a black-box attacker and provides a security statement conditioned to the hardness of well-defined mathematical problems (like LWE). The latter provides a statistical statement showing that any probing attacker limited to $d - 1$ probes have no statistical advantage to recover the sensitive information.

To prove the security of Raccoon, it is important to link these two notions. For that, we detail and formalize the probing security from a game-based perspective, i.e. with well-defined simulators and reuse the notion of probing EUF-CMA provided in [5]. Such a notion has been defined but it was not formally used in a game-based proof before. The main contribution of this paper is the proof of the probing EUF-CMA security of Raccoon. It will consist in several steps.

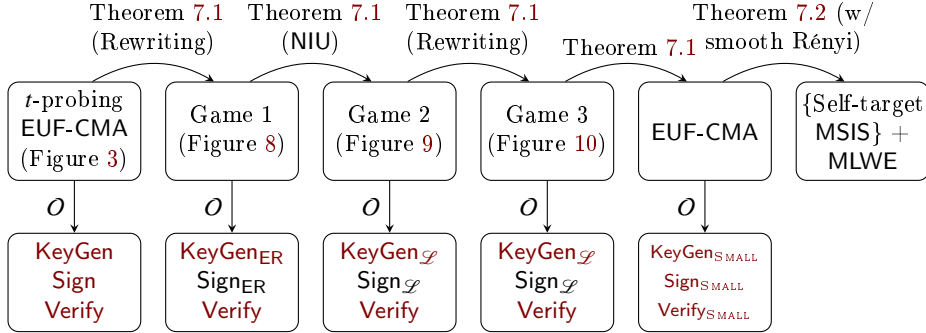


Fig. 1: Proof overview. Jump 1 consists in moving randomness to inputs as per Definition 5.2. Jump 2 uses Lemma 5.2 to move all probes to inputs. Jump 3 is a simple rewriting step. Jump 4 is a black-box reduction to a simpler unmasked signature Small Raccoon. Jump 5 is the security proof of Small Raccoon. \mathcal{O} denote access to an oracle to the corresponding algorithm.

1. **Non-uniform masks and sNIU:** First, one needs to handle the sensitive small uniforms that are deviating from the classical ISW model [27] and other masking proof techniques [4]. For that, all the small uniforms will not be considered as a sharing of a secret value but as several random coins provided in input. The notion of sNIU introduced in [20] (detailed later on in the paper) will come handy. That way, we will be able to prove the masking security of the key generation and signature algorithm when the small uniforms are provided as inputs in Section 6.
2. **Reduction from t-probing EUF-CMA to standard EUF-CMA:** Next, we will use this probe simulation property offered by the NIU model (cf. Lemma 5.2) as part of a game based proof in the probing-EUF-CMA security model. Through a sequence of games, we prove that the probing-EUF-CMA security of Raccoon reduces on the black-box-EUF-CMA of a different version of Raccoon with smaller noise distributions, called small Raccoon. This reduction lets us include the probing adversary in the attack and reduce to a standard (non probing) EUF-CMA adversary. This proof is presented in Section 7.
3. **Unforgeability and smooth Rényi divergence:** Finally, the proof concludes with the black-box security of small Raccoon. Such a proof is close to existing EUF-CMA proofs of signatures following the Fiat-Shamir with aborts framework with a significative difference. To allow a complete end-to-end proof, we avoid any heuristic assessments and introduce the notion of smooth Rényi divergence for obtaining provable and tighter parameters. This proof is presented in Section 7.3.

In Section 8, we instantiate the parameters to valide our proof and confirm that the current NIST submission is secure.

1.3 Related works

SCA against Dilithium. Several side-channel attacks against post-quantum schemes have been published. For concision, we only mention those related to Dilithium, which shares similarities with Raccoon. Since its initial publication, a string of increasingly practical side-channel attacks have been proposed against unprotected implementations of Dilithium: see for example [22], [29], [31], [8], [37], [9], [38].

Masking lattice schemes. The formal study of masking lattice-based signatures has been initiated by Barthe et al. [5], which studied the GLP signature. Since then, BLISS [6] and qTESLA [23] have also been studied from a masking perspective. Masked implementations of Dilithium have been proposed in [33], [3], [12].

Masking-friendly signatures. A few masking-friendly signatures have been proposed in the past two years.

- *Mitaka.* Espitau et al. [21] proposed the Mitaka scheme, a masking-friendly variant of Falcon. A flaw in the security proof of Mitaka, as well as a practical attack in the t -probing model, was later demonstrated by Prest [35].
- *IEEE SP Raccoon.* At IEEE S&P 2023, del Pino et al. [17] presented a lattice-based masking-friendly signature, also called Raccoon. Our scheme is a conceptual descendent of the scheme from [17], with significant improvements. While both versions of Raccoon are Fiat-Shamir lattice-based signatures, the security proof of [17] relies on several heuristic arguments, and the scheme itself is less compact than ours due to the use of a variant of *uniform secret LWR*. In comparison, our design is more streamlined, more compact, relies on standard assumptions and has a formal security proof.
- *Plover.* Since the original publication of Raccoon as a NIST candidate [16], Esgin et al. [20] have proposed Plover, a signature scheme heavily inspired from our scheme, including the use of **AddRepNoise**. The key insight of Plover is to realize that our techniques are not limited to Fiat-Shamir signatures, and can also be applied in a hash-then-sign setting. Conversely, [20] introduced the NIU notion, a useful abstraction that we re-use in our analysis.

2 Preliminaries

We provide the minimal set of preparation. We refer the readers to the full version for more details. First, let us prepare some notations. We note \mathbb{N} the set of non-negative integers, including zero. Given $n \in \mathbb{N}$, we denote by $[n]$ the set $\{0, 1, \dots, n-1\}$. Let $f : X \rightarrow Y$ be a function, and $x \in X$. When f is deterministic, we use the notation $y := f(x)$ to indicate that we assign the output of $f(x)$ to y . When f is randomized, we instead use the notation $y \leftarrow f(x)$. From a programming viewpoint, both of these notations indicate an assignment of the result to the variable on the left. Given a probability distribution \mathcal{D} over Y , we note $y \leftarrow \mathcal{D}$ to express that $y \in Y$ is sampled from \mathcal{D} .

2.1 Hardness Assumptions

The security of Raccoon is based on the Module Learning with Errors (MLWE) and Module Short Integer Solutions (MSIS) assumptions. More precisely, we rely on the Self Target MSIS (SelfTargetMSIS) problem, a variant of the MSIS problem, where the problem is defined relative to some hash function modeled as a random oracle. This assumption also underlies the security of Dilithium.

2.2 Masking Preliminaries

We consider all operations and variables used in algorithms to be over the scalar ring \mathcal{R}_q (i.e. we consider that basic operations are done directly on polynomials in \mathcal{R}_q), this entails that we consider that probes leak full polynomials in \mathcal{R}_q and not bits or even coefficients (leading to a stronger attacker model). An algorithm is defined as a sequence of gadget calls, each gadget being a sequence of (probabilistic or deterministic) assignments of expressions to local variables.

Well-formed gadgets. We say a gadget is well-formed if it is written in SSA (single static assignment) form, i.e. if its scalar variables appear at most once on the left-hand side of an assignment, and if all assignments are three-address code instructions, i.e. of the form $a = b * c$ with $*$ an operator. These restrictions ensure that all intermediate values are captured by local variables at some point in the code. An algorithm is well formed if in all gadget calls $\mathbf{b} = G(\mathbf{x}_1, \dots, \mathbf{x}_k)$ the variables $\mathbf{b}, \mathbf{x}_1, \dots, \mathbf{x}_k$ are pairwise disjoint. While some algorithms we provide are not well formed (e.g., Algorithms 1 and 2), it is clear that this can be easily remedied by indexing variables and adding new local variables.

We use the notation $\llbracket \mathbf{x} \rrbracket = (\mathbf{x}_i)_{i \in [d]}$ to denote a tuple of d values in \mathcal{R}_q , which implicitly defines the value $\mathbf{x} = \sum_0^{d-1} \mathbf{x}_i \in \mathcal{R}_q$. This notation is used to express that the secret value \mathbf{x} is shared as d additive shares as the encoding $\llbracket \mathbf{x} \rrbracket$.

Variables' values and names. We will distinguish variables (designated by a binary string representing their name) from the values they take (in the scalar ring \mathcal{R}_q), all objects pertaining to variables (singular variables, vectors, sets, etc...) will have a name with a bar (e.g. $\bar{x} \in \{0, 1\}^*$, $\bar{\mathcal{V}} \subset \{0, 1\}^*$), while the corresponding value will not (e.g. $x \in \mathcal{R}_q$).

For a gadget G we define the local variables of G as $\bar{\mathcal{V}}_G \subset \{0, 1\}^*$ (noted $\bar{\mathcal{V}}$ when the gadget is clear from the context), since all variables are assigned only once we can match the position of a variable with its name. For a program P with input scalar variables $(\bar{a}_1, \dots, \bar{a}_N)$ that calls the gadgets G_1, \dots, G_k , (with $N, k \in \mathbb{N}$), we will consider the set of variables $\bar{\mathcal{V}}_P = \{\bar{a}_1, \dots, \bar{a}_N\} \uplus \bar{\mathcal{V}}_{G_1} \uplus \dots \uplus \bar{\mathcal{V}}_{G_k}$ (where the local variables of G_i are additionally labelled with i to differentiate between gadgets and \uplus is the disjoint union). Note that since all gadgets are written in three-address code SSA form, all intermediate computations and output variables are at some point stored locally in a uniquely defined local variable $\bar{v} \in \bar{\mathcal{V}}_P$. We thus define the set of all possible probes as the set $\bar{\mathcal{V}}_P$ of all local variables as well as the input variables.

Remark 2.1. We will consider that a program P always outputs all unmasked values it computes even if they are not explicitly returned by P .

Definition 2.1 (Probes). For a well-formed program P with variables $\bar{\mathcal{V}}_P$ and input variables $\bar{a}_1, \dots, \bar{a}_N$, a set of probes is a set $\bar{\mathcal{F}} \subset \bar{\mathcal{V}}_P$. For any set $\bar{\mathcal{F}} \subset \bar{\mathcal{V}}_P$ and any scalars $\mathcal{X} = (a_1, \dots, a_N)$ we will denote as $\text{ExecObs}(P, \bar{\mathcal{F}}, \mathcal{X})$ the joint distribution of the (masked and unmasked) outputs of $P(a_1, \dots, a_N)$ and of all the values taken by the variables in $\bar{\mathcal{F}}$. In particular for

$$(\text{out}_{\text{masked}}, \text{out}_{\text{unmasked}}, \mathcal{L}) \leftarrow \text{ExecObs}(P, \bar{\mathcal{F}}, \mathcal{X}),$$

$\text{out}_{\text{masked}}$ (resp. $\text{out}_{\text{unmasked}}$) is the masked (resp. unmasked) output of $P(a_1, \dots, a_N)$ for some internal random coins and \mathcal{L} is the value taken by the variables in $\bar{\mathcal{F}}$ for these random coins.

Probing model. We recall standard non-interference results from [4].

Definition 2.2 (Perfect simulatability, reformulation of [4]). Let $\bar{\mathcal{F}}$ be a set of probes of a gadget \mathbf{G} with input shares \bar{X} . We say that the PPT simulator $(\text{SimIn}, \text{SimOut})$ perfectly simulates the probes $\bar{\mathcal{F}}$ if and only if for any input values X , $\text{SimIn}(\mathbf{G}, \bar{\mathcal{F}})$ outputs a subset $\bar{X}' \subset X$ of the input variables of \mathbf{G} , and $\text{SimOut}(\mathbf{G}, X')$ (where X' is the values taken by X at indices \bar{X}') outputs a tuple of values such that the marginal distribution of \mathcal{L} , for $(\text{out}_{\text{masked}}, \text{out}_{\text{unmasked}}, \mathcal{L}) \leftarrow \text{ExecObs}(P, \bar{\mathcal{F}}, X)$, and $\text{SimOut}(\mathbf{G}, X')$ are identical.

Definition 2.3 (Non Interference [4]). A gadget is said $(d-1)$ -non-interfering (written $(d-1)$ -NI for short) iff any set of probes $\bar{\mathcal{F}}$ such that $|\bar{\mathcal{F}}| \leq d-1$ can be perfectly simulated (See Definition 2.2) by a simulator $(\text{SimIn}, \text{SimOut})$ such that $\text{SimIn}(\mathbf{G}, \bar{\mathcal{F}})$ outputs a set \bar{X}' of at most $d-1$ shares of each input.

Definition 2.4 (Strong Non Interference [4]). A gadget is said $(d-1)$ -strongly-non-interfering (written $(d-1)$ -sNI for short) iff any set $\bar{\mathcal{F}}$ of at most $d-1 = d_{\text{int}} + d_{\text{out}}$ probes, where d_{int} are made on internal data and d_{out} are made on the outputs, can be perfectly simulated by a simulator $(\text{SimIn}, \text{SimOut})$ such that $\text{SimIn}(\mathbf{G}, \bar{\mathcal{F}})$ outputs a set \bar{X}' of at most d_{int} shares of each input.

Lemma 2.1 (Composability of NI and sNI gadgets [5]). A well-formed algorithm is NI if all of its gadgets are NI or sNI and each sharing is used at most once as input of a non-sNI gadget. Moreover, a well-formed algorithm is sNI if it is NI and its output sharings are issued from a sNI gadget.

Lastly, in this paper, the masking order is fixed at $d-1$ where d is the number of shares. For simplicity, we omit the $d-1$ when referring to NI/sNI properties.

2.3 Sum of Uniforms

Given a distribution \mathcal{D} of support included in an additive group, we note $[T] \cdot \mathcal{D}$ the convolution of T identical copies of \mathcal{D} ; in other words, $[T] \cdot \mathcal{D}$ is the

distribution of the sum of T independent random variables, each being sampled from \mathcal{D} . Given integers $u, T > 0$, and if we note $\mathcal{U}(S)$ the uniform distribution over a finite S , we note:

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

The acronym SU stands for “sum of uniforms”. This class of distributions is




Fig. 2: The distribution $\text{SU}(4, T)$, for $T \in \{1, 2, 4, 8\}$

illustrated in Figure 2. This distribution is highly desirable for our purposes, since for $T \geq 4$ it verifies statistical properties in the same way as Gaussians do. However, unlike Gaussians, they are straightforward to sample in constant-time and without requiring tables or elaborate mathematical machinery. This makes them adequate for Raccoon. Finally, we note $\text{RSU}(u, 1)$ the distribution over \mathcal{R} obtained by sampling each integer coefficient of $a \in \mathcal{R}$ according to $\text{SU}(u, 1)$, and outputting a . More details about sums of uniforms can be found the full version of this paper.

3 The Raccoon Signature Scheme

In this section, we present our masking-friendly signature scheme called Raccoon. We describe the key generation (Algorithm 1), signing (Algorithm 2) and verification (Algorithm 3). Key generation and signing are always performed in a masked manner; when $d = 1$, the algorithmic descriptions remain valid but the algorithms are, in effect, unmasked.

We reference relevant variables and parameters in Table 1.

3.1 Key Generation

Masked key generation process is described by Algorithm 1. At a high-level, **KeyGen** generates d -sharings ($\llbracket \mathbf{s} \rrbracket, \llbracket \mathbf{e} \rrbracket$) of small errors (\mathbf{s}, \mathbf{e}), computes the verification key as an LWE sample ($\mathbf{A}, \mathbf{t} = \mathbf{A} \cdot \mathbf{s} + \mathbf{e}$), and rounds \mathbf{t} for efficiency. A key technique is that $\llbracket \mathbf{s} \rrbracket, \llbracket \mathbf{e} \rrbracket$ are generated in Lines 4 and 6 using our novel algorithm **AddRepNoise** (Algorithm 5).

Parameter	Explanation
(\mathcal{R}_q, n)	Polynomial ring $\mathcal{R}_q = \mathbb{Z}[X]/(q, X^n + 1)$
(k, ℓ)	Dimension of public matrix $\mathbf{A} \in \mathcal{R}_q^{k \times \ell}$
d	Number of shares used, corresponding to a masking order $d - 1$
$\text{RSU}(a, b)$	Sum of a polynomials with coefficients uniform in $\{-2^{u-1}, \dots, 2^{u-1} - 1\}$
u_t, u_w	Parameter and repetition rate used for the sum of uniform in the secret/signature $\mathbf{s} \leftarrow \text{RSU}^\ell(u_t, \text{rep}), \mathbf{r} \leftarrow \text{RSU}^\ell(u_w, \text{rep})$
rep	
ν_t	Amount of bit dropping performed on verification key
ν_w	Amount of bit dropping performed on (aggregated) commitment
(q_t, q_w)	Rounded moduli satisfying $(q_t, q_w) := (\lfloor q/2^{\nu_t} \rfloor, \lfloor q/2^{\nu_w} \rfloor)$
(C, ω)	Challenge set $\{c \in \mathcal{R}_q \mid \ c\ _\infty = 1 \wedge \ c\ _1 = \omega\}$ s.t. $ C \geq 2^{2\kappa}$
(B_2, B_∞)	Two-norm and infinity-norm bounds on the signature

Table 1: Overview of parameters used in the Raccoon signature.

Algorithm 1 $\text{KeyGen}(\phi) \rightarrow (\text{vk}, \text{sk})$ **Output:** Keypair vk, sk

- 1: $\text{seed} \leftarrow \{0, 1\}^\kappa$ $\triangleright \kappa$ -bit random seed for \mathbf{A} .
- 2: $\mathbf{A} := \text{ExpandA}(\text{seed})$ \triangleright Similar to ExpandA in Dilithium. $\mathbf{A} \in \mathcal{R}_q^{k \times \ell}$.
- 3: $\llbracket \mathbf{s} \rrbracket \leftarrow \ell \times \text{ZeroEncoding}(d)$ \triangleright Masked zero vector $\llbracket \mathbf{s} \rrbracket \in (\mathcal{R}_q^\ell)^d$. Algorithm 8.
- 4: $\llbracket \mathbf{s} \rrbracket \leftarrow \text{AddRepNoise}(\llbracket \mathbf{s} \rrbracket, u_t, \text{rep})$ \triangleright Generate the secret distribution.
Algorithm 5.
- 5: $\llbracket \mathbf{t} \rrbracket := \mathbf{A} \cdot \llbracket \mathbf{s} \rrbracket$ \triangleright Compute masked product $\llbracket \mathbf{t} \rrbracket \in (\mathcal{R}_q^k)^d$.
- 6: $\llbracket \mathbf{t} \rrbracket \leftarrow \text{AddRepNoise}(\llbracket \mathbf{t} \rrbracket, u_t, \text{rep})$ \triangleright Add masked noise to $\llbracket \mathbf{t} \rrbracket$. Algorithm 5.
- 7: $\mathbf{t} := \text{Decode}(\llbracket \mathbf{t} \rrbracket)$ \triangleright Collapse $\mathbf{t} \in \mathcal{R}_q^k$. Algorithm 6.
- 8: $\mathbf{t} := \lfloor \mathbf{t} \rfloor_{\nu_t}$ \triangleright Rounding and right-shift to modulus $q_t = \lfloor q/2^{\nu_t} \rfloor$.
- 9: **return** $(\text{vk} := (\text{seed}, \mathbf{t}), \text{sk} := (\text{vk}, \llbracket \mathbf{s} \rrbracket))$ \triangleright Return serialized key pair.

3.2 Signing Procedure

The masked signing process is described by Algorithm 2. This signing procedure is similar to the “Lyubashevsky’s Signature Without Aborts” in [1, Fig. 2]. Again, the use of `AddRepNoise` is crucial in this procedure. The challenge computation is divided in two parts, first a 2κ bitstring is computed using the hash function `ChalHash`, then this bitstring is mapped to a ternary polynomial with fixed hamming weight using `ChalPoly`. As in previous works this distinction is made for ease of implementation and storage.

3.3 Verification Procedure

Algorithm 3 describes the signature verification process. Signature verification is not masked, and its parameters are independent of the number of shares d used when creating the signature. As is usual in lattice signatures, verification performs a bound check and an equality check.

It is easy to check that the equation of line 7 verifies by construction when the signature algorithm is run honestly, we will fix the bounds B_∞ and B_2 such that honest signatures verify with overwhelming probability (this is necessary for the reduction of Section 7.2 to go through).

Algorithm 2 $\text{Sign}(\llbracket \text{sk} \rrbracket, \text{msg}) \rightarrow \text{sig}$

Input: Secret signing key $\text{sk} = (\text{vk}, \llbracket \text{s} \rrbracket)$, message to be signed $\text{msg} \in \{0, 1\}^*$.
Output: Signature $\text{sig} = (c_{\text{hash}}, \mathbf{h}, \mathbf{z})$ of msg under sk .

- 1: $\mu := \text{H}(\text{H}(\text{vk}) \parallel \text{msg})$ ▷ Bind vk with msg to form $\mu \in \{0, 1\}^{2\kappa}$.
- 2: $\mathbf{A} := \text{ExpandA}(\text{seed})$ ▷ Similar to ExpandA in Dilithium. $\mathbf{A} \in \mathcal{R}_q^{k \times \ell}$.
- 3: $\llbracket \mathbf{r} \rrbracket \leftarrow \ell \times \text{ZeroEncoding}(d)$ ▷ Masked zero vector $\llbracket \mathbf{r} \rrbracket \in (\mathcal{R}_q^\ell)^d$. Algorithm 8.
- 4: $\llbracket \mathbf{r} \rrbracket \leftarrow \text{AddRepNoise}(\llbracket \mathbf{r} \rrbracket, u_{\mathbf{w}}, \text{rep})$ ▷ Add masked noise to $\llbracket \mathbf{r} \rrbracket$. Algorithm 5.
- 5: $\llbracket \mathbf{w} \rrbracket := \mathbf{A} \cdot \llbracket \mathbf{r} \rrbracket$ ▷ Compute masked product $\llbracket \mathbf{w} \rrbracket \in (\mathcal{R}_q^k)^d$.
- 6: $\llbracket \mathbf{w} \rrbracket \leftarrow \text{AddRepNoise}(\llbracket \mathbf{w} \rrbracket, u_{\mathbf{w}}, \text{rep})$ ▷ Add masked noise to $\llbracket \mathbf{w} \rrbracket$. Algorithm 5.
- 7: $\mathbf{w} := \text{Decode}(\llbracket \mathbf{w} \rrbracket)$ ▷ Collapse LWE commitment \mathbf{w} . Algorithm 6.
- 8: $\mathbf{w} := \lfloor \mathbf{w} \rfloor_{v_{\mathbf{w}}}$ ▷ Rounding and right-shift to modulus $q_{\mathbf{w}} = \lfloor q/2^{v_{\mathbf{w}}} \rfloor$.
- 9: $c_{\text{hash}} := \text{ChalHash}(\mathbf{w}, \mu)$ ▷ Map \mathbf{w} and μ to $c_{\text{hash}} \in \{0, 1\}^{2\kappa}$.
- 10: $c_{\text{poly}} := \text{ChalPoly}(c_{\text{hash}})$ ▷ Map c_{hash} to $c_{\text{poly}} \in \mathcal{C}$.
- 11: $\llbracket \mathbf{s} \rrbracket \leftarrow \text{Refresh}(\llbracket \mathbf{s} \rrbracket)$ ▷ Refresh $\llbracket \mathbf{s} \rrbracket$ before re-use. Algorithm 7.
- 12: $\llbracket \mathbf{r} \rrbracket \leftarrow \text{Refresh}(\llbracket \mathbf{r} \rrbracket)$ ▷ Refresh $\llbracket \mathbf{r} \rrbracket$ before re-use. Algorithm 7.
- 13: $\llbracket \mathbf{z} \rrbracket := c_{\text{poly}} \cdot \llbracket \mathbf{s} \rrbracket + \llbracket \mathbf{r} \rrbracket$ ▷ Masked response $\llbracket \mathbf{z} \rrbracket \in (\mathcal{R}_q^\ell)^d$.
- 14: $\llbracket \mathbf{z} \rrbracket \leftarrow \text{Refresh}(\llbracket \mathbf{z} \rrbracket)$ ▷ Refresh $\llbracket \mathbf{z} \rrbracket$ before collapsing it. Algorithm 7.
- 15: $\mathbf{z} := \text{Decode}(\llbracket \mathbf{z} \rrbracket)$ ▷ Collapse into response $\mathbf{z} \in \mathcal{R}_q^\ell$. Algorithm 6.
- 16: $\mathbf{y} := \mathbf{A} \cdot \mathbf{z} - 2^{v_{\mathbf{t}}} \cdot c_{\text{poly}} \cdot \mathbf{t}$ ▷ “Noisy” LWE commitment.
- 17: $\mathbf{h} := \mathbf{w} - \lfloor \mathbf{y} \rfloor_{v_{\mathbf{w}}}$ ▷ Compute hint $\mathbf{h} \in \mathcal{R}_{q_{\mathbf{w}}}^k$. Subtraction mod $q_{\mathbf{w}}$.
- 18: $\text{sig} := (c_{\text{hash}}, \mathbf{h}, \mathbf{z})$
- 19: **if** $\{\text{CheckBounds}(\text{sig}) = \text{FAIL}\}$ **goto** Line 3 ▷ Sanity check on the signature. Algorithm 4.
- 20: **return** sig ▷ Return encoded signature triplet.

Algorithm 3 $\text{Verify}(\text{sig}, \text{msg}, \text{vk}) \rightarrow \{\text{OK or FAIL}\}$

Input: Signature $\text{sig} = (c_{\text{hash}}, \mathbf{h}, \mathbf{z})$, message $\text{msg} \in \{0, 1\}^*$, public key $\text{vk} = (\text{seed}, \mathbf{t})$.
Output: Signature validity: OK (accept) or FAIL (reject).

- 1: **if** $\text{CheckBounds}(\text{sig}) = \text{FAIL}$ **return** FAIL ▷ Norms check. Algorithm 4.
- 2: $\mu := \text{H}(\text{H}(\text{vk}) \parallel \text{msg})$; $\mathbf{A} := \text{ExpandA}(\text{seed})$
- 3: $c_{\text{poly}} := \text{ChalPoly}(c_{\text{hash}})$ ▷ Map c_{hash} to $c_{\text{poly}} \in \mathcal{C}$.
- 4: $\mathbf{y} := \mathbf{A} \cdot \mathbf{z} - 2^{v_{\mathbf{t}}} \cdot c_{\text{poly}} \cdot \mathbf{t}$ ▷ Scale \mathbf{t} from $\mathbb{Z}_{q_{\mathbf{t}}}$ to \mathbb{Z}_q and recompute the commitment.
- 5: $\mathbf{w}' := \lfloor \mathbf{y} \rfloor_{v_{\mathbf{w}}} + \mathbf{h}$ ▷ Adjust the LWE commitment with hint (mod $q_{\mathbf{w}}$).
- 6: $c'_{\text{hash}} := \text{ChalHash}(\mathbf{w}', \mu)$ ▷ Recompute $c'_{\text{hash}} \in \{0, 1\}^{2\kappa}$.
- 7: **if** $c_{\text{hash}} \neq c'_{\text{hash}}$ **return** FAIL ▷ Check commitment.
- 8: **return** OK ▷ Signature is accepted.

3.4 Helper Algorithms

The following are algorithms used within our key generation (Algorithm 1), signing (Algorithm 2) and verification (Algorithm 3). The algorithm **AddRepNoise** (Algorithm 5) is the most interesting one, which we come back later when discussing probing security.

Checking Bounds. The function **CheckBounds** (Algorithm 4) is used to check the norm bounds and encoding soundness of signatures by both the verification

function (Algorithm 3), but also by the signing function (Algorithm 2). Note that unlike rejection, **CheckBounds** is used to enforce the zero-knowledge property, and therefore it does need to be masked. Rather, it detects signatures that are a bit too large. Note that **CheckBounds** could be removed entirely at the cost of a slight increase in signature size (and therefore a slight decrease in security).

Algorithm 4 **CheckBounds**(sig) \rightarrow {OK or FAIL}

Input: Signature sig = ($c_{\text{hash}}, \mathbf{h}, \mathbf{z}$).

Output: Format validity check OK or FAIL.

1: **if** ($\|(\mathbf{z}, 2^{\nu_w} \cdot \mathbf{h})\|_{\infty} > B_{\infty}$) **or** ($\|(\mathbf{z}, 2^{\nu_w} \cdot \mathbf{h})\|_2 > B_2$) **return** FAIL **else return** OK

Error Distributions. **AddRepNoise** (Algorithm 5) implements the Sum of Uniforms (SU) distribution $SU(u, d \cdot \text{rep})$ (Section 2.3) in a masked implementation. **AddRepNoise** interleaves noise additions and refresh operations; more precisely, for each (masked) coefficient $\llbracket a \rrbracket$ of $\llbracket \mathbf{v} \rrbracket$, small uniform noise is added to each share of $\llbracket a \rrbracket$, then $\llbracket a \rrbracket$ is refreshed, and this operation is repeated rep times. The security properties of **AddRepNoise** is analyzed in Section 6.2.

Algorithm 5 **AddRepNoise**($\llbracket \mathbf{v} \rrbracket, u, \text{rep}$) \rightarrow $\llbracket \mathbf{v} \rrbracket$

Input: Masked vector $\llbracket \mathbf{v} \rrbracket = (\mathbf{v}_j)_{j \in [d]} = (v_{i,j})_{i \in [\text{len}(\mathbf{v})], j \in [d]}$.

Input: Bit size (distribution parameter) u .

Input: Global repetition count parameter rep .

Output: Updated $\llbracket \mathbf{v} \rrbracket$ with $SU(u, d \cdot \text{rep})$ distribution added to each coefficient of \mathbf{v} .

```

1: for  $i \in [\text{len}(\mathbf{v})]$  do                                ▷ Vector index.
2:   for  $i_{\text{rep}} \in [\text{rep}]$  do                            ▷ Repetition index.
3:     for  $j \in [d]$  do                                    ▷ Share index.
4:        $\rho \leftarrow \text{RSU}(u, 1)$                           ▷ uniform sample of  $u$  bits
5:        $v_{i,j} \leftarrow v_{i,j} + \rho$                     ▷ Add small uniform to the polynomial.
6:        $\llbracket \mathbf{v}_i \rrbracket \leftarrow \text{Refresh}(\llbracket \mathbf{v}_i \rrbracket)$     ▷ Refresh polynomial on each repeat.
7: return  $\llbracket \mathbf{v} \rrbracket$ 

```

Challenge Computation. As in Dilithium, the challenge computation is split in two subroutines: **ChalHash** computes a hash digest, and **ChalPoly** expands it into a challenge polynomial c_{poly} that is (pseudo-randomly) uniform in the set $C = \{c \in \mathcal{R}, \|c\|_1 = \omega\}$. These functions do not need to be masked.

Refresh and Decoding Gadgets. Lastly, we recall some useful gadgets. **Refresh** (Algorithm 7) generates a fresh d -sharing of a value in \mathcal{R}_q , or “refresh” the d -sharing. This operation is important for security against t -probing adversaries.

Refresh uses **ZeroEncoding** (Algorithm 8) as a subroutine. Both algorithms perform $O(d \log d)$ basic operations over \mathcal{R}_q and require $O(d \log(d) \log(q))$ bits of entropy. While we present **ZeroEncoding** as a recursive algorithm, one can see that it can be computed in-place and its memory requirement is $O(d)$. Remark that our **ZeroEncoding** algorithm entails that the number of shares d is a power of 2, as the rest of our algorithms are agnostic to this property we could use a **ZeroEncoding** that produces a more fine-grained number of shares to obtain different parameters (e.g. by using Algorithm 8 and collapsing some of the shares).

We describe in Algorithm 6 a **Decode** gadget that takes $\llbracket \mathbf{x} \rrbracket = (\mathbf{x}_i)_{i \in [t+1]}$ as input, refreshes it with Algorithm 7, then computes the sum $\mathbf{x}_0 + \dots + \mathbf{x}_{d-1} \bmod q$. When the decoding gadget is already preceded by a refresh gadget, one of them may be omitted. **Decode** is similar to the algorithm “FullAdd” from [5, Alg. 16].

<hr/> <p>Algorithm 6 Decode($\llbracket x \rrbracket$) $\rightarrow x$</p> <hr/> <p>Input: d-sharing $\llbracket x \rrbracket = (x_i)_i$ of $x \in \mathcal{R}_q$ Output: The clear value $x \in \mathcal{R}_q$ 1: $\llbracket x \rrbracket \leftarrow \text{Refresh}(\llbracket x \rrbracket)$ 2: return $x := \sum_{i \in [d]} x_i$</p> <hr/>	<hr/> <p>Algorithm 7 Refresh($\llbracket x \rrbracket$) $\rightarrow \llbracket x \rrbracket'$</p> <hr/> <p>Input: A d-sharing $\llbracket x \rrbracket$ of $x \in \mathcal{R}_q$ Output: A fresh d-sharing $\llbracket x \rrbracket'$ of x 1: $\llbracket z \rrbracket \leftarrow \text{ZeroEncoding}(d)$ 2: return $\llbracket x \rrbracket' := \llbracket x \rrbracket + \llbracket z \rrbracket$</p> <hr/>
<hr/> <p>Algorithm 8 ZeroEncoding(d) $\rightarrow \llbracket z \rrbracket_d$</p> <hr/> <p>Input: A power-of-two integer d, a ring \mathcal{R}_q Output: A uniform d-sharing $\llbracket z \rrbracket \in \mathcal{R}_q^d$ of $0 \in \mathcal{R}_q$ 1: if $d = 1$ then 2: return $\llbracket z \rrbracket_1 := (0)$ \triangleright There is only one way to encode zero into 1 share. 3: $\llbracket z_1 \rrbracket_{d/2} \leftarrow \text{ZeroEncoding}(d/2)$ \triangleright Recursively obtain left side. 4: $\llbracket z_2 \rrbracket_{d/2} \leftarrow \text{ZeroEncoding}(d/2)$ \triangleright Recursively obtain right side. 5: $\llbracket r \rrbracket_{d/2} \xleftarrow{M} \mathcal{R}_q^{d/2}$ \triangleright Sampled using a Mask Random Generator (MRG). 6: $\llbracket z_1 \rrbracket_{d/2} := \llbracket z_1 \rrbracket_{d/2} + \llbracket r \rrbracket_{d/2}$ \triangleright Add to the left side. 7: $\llbracket z_2 \rrbracket_{d/2} := \llbracket z_2 \rrbracket_{d/2} - \llbracket r \rrbracket_{d/2}$ \triangleright Subtract from the right side. 8: return $\llbracket z \rrbracket_d := (\llbracket z_1 \rrbracket_{d/2} \parallel \llbracket z_2 \rrbracket_{d/2})$ \triangleright Concatenate the two.</p> <hr/>	

4 Smooth Rényi Divergence and Useful Bounds

Raccoon’s core design choice is using the sum of uniform distributions as opposed to the discrete Gaussian distributions. From a practical point of view, the sum of uniform distribution is a much simpler distribution to mask and implement. On the other hand, from a theoretical point of view, it poses more challenges, as there are far fewer established statistical guarantees usable in cryptography. Notably, since the sum of uniform distribution only has finite support, a standard proof technique used in lattice-based cryptography relying on the Rényi divergence breaks down. To this end, we generalize the Rényi divergence and prepare useful statistical bounds on the sum of uniform distribution.

4.1 Smooth Rényi Divergence

The usual Rényi divergence is undefined for distributions P, Q of supports not included in one another. For example, this happens when $P = \text{SU}(u, T)$ and $Q = P+a$, for any $a \neq 0$. The *smooth* Rényi divergence (Definition 4.1) addresses these limitations by combining the statistical distance and the Rényi divergence. The statistical distance component captures problematic sets (typically, distribution tails), while the Rényi divergence component benefits from the same efficiency as the usual Rényi divergence over unproblematic parts of the supports.

Definition 4.1 (Smooth Rényi divergence). *Let $\epsilon \geq 0$ and $1 < \alpha < \infty$. Let P, Q be two distributions of countable supports $\text{Supp}(P) \subseteq \text{Supp}(Q) = X$. The smooth Rényi divergence of parameters (α, ϵ) between P and Q is defined as:*

$$R_\alpha^\epsilon(P; Q) = \min_{\substack{\Delta_{\text{SD}}(P'; P) \leq \epsilon \\ \Delta_{\text{SD}}(Q'; Q) \leq \epsilon}} R_\alpha(P'; Q'), \quad (1)$$

where Δ_{SD} and R_α denote the statistical distance and the Rényi divergence, respectively:

$$\Delta_{\text{SD}}(P; Q) = \frac{1}{2} \sum_{x \in X} |P(x) - Q(x)|, \quad R_\alpha(P; Q) = \left(\sum_{x \in X} \frac{P(x)^\alpha}{Q(x)^{\alpha-1}} \right)^{\frac{1}{\alpha-1}}.$$

While [18] has also provided a definition of smooth Rényi divergence, we argue that our definition is more natural. Indeed, it satisfies variations of properties that are expected from classical Rényi divergences. These are listed in Lemma 4.1.

Tools for smooth Rényi divergence. We review some basic properties of the smooth Rényi divergence.

Lemma 4.1. *The smooth Rényi divergence satisfies the following properties.*

1. **Data processing inequality.** *Let P, Q be two distributions, let $\epsilon \geq 0$, and g be a randomized function over (a superset of) $\text{Supp}(P) \cup \text{Supp}(Q)$.*

$$R_\alpha^\epsilon(g(P); g(Q)) \leq R_\alpha^\epsilon(P; Q). \quad (2)$$

2. **Probability preservation.** *For any event $E \subseteq \text{Supp}(Q)$:*

$$P(E) \leq (Q(E) + \epsilon)^{(\alpha-1)/\alpha} \cdot R_\alpha^\epsilon(P; Q) + \epsilon. \quad (3)$$

3. **Tensorization.** *Let $(P_i)_{i \in I}, (Q_i)_{i \in I}$ be two finite families of distributions, let $\epsilon_i \geq 0$ for $i \in I$, and let $\epsilon = \sum_{i \in I} \epsilon_i$.*

$$R_\alpha^\epsilon \left(\prod_{i \in I} P_i; \prod_{i \in I} Q_i \right) \leq \prod_{i \in I} R_\alpha^{\epsilon_i}(P_i; Q_i). \quad (4)$$

Proof. We recall that Δ_{SD} and $(R_\alpha^\alpha - 1)$ can be cast as f -divergences, following Csiszár's terminology [15]. Item 1 follows from the data processing inequality for f -divergences. Item 2 is a special case of Item 1. Finally, Item 3 follows from tensorization properties of the statistical distance and the Rényi divergence. \square

4.2 Useful Bounds on Sum of Uniforms

We bound the smooth Rényi divergence between two sums of uniforms, centered at either 0 or a small offset. This will be a key lemma establishing the hardness of standard EUF-CMA security of the small Raccoon (cf. Section 7.3). Due to page limitation, the proof is provided in the full version of this paper.

Lemma 4.2. *Let $T, u, N \in \mathbb{N}$ and $c \in \mathbb{Z}$ such that $T \geq 4$ and $N = 2^u$. Let $P = \text{SU}(u, T)$ and Q the distributions corresponding to shifting the support of P by c . Let $\alpha \geq 2$ and $\tau > 0, \epsilon > 0$ be such that:*

1. $\alpha |c| \leq \tau = o(N/(T-1))$;
2. $\epsilon = \frac{(\tau+T)^T}{N^T T!}$.

Then:

$$R_\alpha^\epsilon(P; Q) \leq \left(1 + \frac{\alpha(\alpha-1)}{2} \left(\frac{Tc}{N} \right)^2 + \frac{2}{T!} \left(\frac{T\alpha c}{N} \right)^2 + \epsilon + O \left(\left(\frac{T\alpha c}{N} \right)^3 \right) \right)^{1/(\alpha-1)} \quad (5)$$

Gap with practice. In practice, Lemma 4.2 is a bit sub-optimal. Let us note $\sigma^2 = \frac{T(N^2-1)}{12}$ the variance of P and $Tc = o(N)$, which follows from Item 1 above. We also use the notation $a \lesssim b$ for $a \leq b + o(b)$. Then, Lemma 4.2 essentially tells us that $\log R_\alpha^\epsilon(P; Q) \lesssim \frac{\alpha}{2} \left(\frac{Tc}{N} \right)^2 \sim \frac{\alpha c^2 T^3}{24 \sigma^2}$. In comparison, [1, Lemma 2.28] tells that if P is instead a Gaussian of parameter σ , then $\log R_\alpha(P; Q) \leq \frac{\alpha c^2}{2 \sigma^2}$. Thus there is a gap $O(T^3)$ between Lemma 4.2 and [1, Lemma 2.28].

One could assume that this gap is caused by a fundamental difference between Gaussians and sums of uniforms. However we performed extensive experiments and found that this gap does not exist in practice, i.e., it seems to be an artifact of our proof. For this reason, we put forward the following conjecture which ignores this gap and which we use when setting our concrete parameters. Due to page limitation, we expand upon Conjecture 4.1 in the full version.

Conjecture 4.1. Under the conditions of Lemma 4.2, we have

$$R_\alpha^\epsilon(P; Q) \lesssim \exp \left(\frac{C_{\text{RÉNYI}} \cdot \alpha \cdot c^2 \left(1 + \frac{2}{\alpha-1} \right)}{T \cdot N^2} \right) \quad (6)$$

for a constant $C_{\text{RÉNYI}} \approx 6$. Therefore, for any M -dimensional vector \mathbf{c} , $\mathcal{P} = P^M$ and $\mathcal{Q} = \mathbf{c} + Q^M$, and further assuming $\alpha = \omega_{\text{asympt}}(1)$ and $T = o(\alpha |c_i|)$ for all the i -th ($i \in [M]$) entry of \mathbf{c} , we have:

$$R_\alpha^\epsilon(\mathcal{P}; \mathcal{Q}) \lesssim \exp \left(\frac{C_{\text{RÉNYI}} \cdot \alpha \cdot \|\mathbf{c}\|_2^2}{T \cdot N^2} \right), \quad (7)$$

$$\text{where } \epsilon \approx \frac{\alpha^T \|\mathbf{c}\|_T^T}{N^T T!} \lesssim \frac{1}{\sqrt{2\pi T}} \left(\frac{\alpha e \|\mathbf{c}\|_2}{NT} \right)^T \quad (8)$$

and where $\|\mathbf{c}\|_T \leq \|\mathbf{c}\|_2$ is the L_T norm.

5 Enhancing NI/sNI for Probing EUF-CMA Security

We first formally define NI security against a probing adversary, the security model in which Raccoon will later be proved in. We then argue that existing probing tools/models discussed in Section 2.2 are insufficient to prove EUF-CMA security and prepare useful tools that may be of independent interest. Our tools build on the recent techniques developed by [20] (cf. Section 1.3).

5.1 EUF-CMA Security in the Probing Model

We use the definition of [5] that captures the fact that no PPT adversary with access to less than $d - 1$ probes on **KeyGen** and **Sign** should be able to break EUF-CMA security (i.e., unforgeability). Below, our definition slightly deviates from theirs as we rely on more generalized (and formal) notion of probes captured by the function **ExecObs** (cf. Definition 2.1).

Definition 5.1. *Let $d \geq 1$ an integer, Q_s be a fixed maximum amount of signature queries. A signature scheme (**KeyGen**, **Sign**, **Verify**) with an efficient signing key update algorithm **KeyUpdate** is EUF-CMA-secure in the $(d-1)$ -probing model if any probabilistic polynomial time adversary has a negligible probability of winning the game presented in Figure 3.*

As in [5], we assume a **KeyUpdate** algorithm that refreshes the secret key between signature queries and cannot be probed by the attacker. This is performed to avoid attackers probing more than $d - 1$ shares of the secret across different signature queries. See [5, Remark 3] for more details.

Remark 5.1 (Standard EUF-CMA security). We note that Definition 5.1 incorporates the standard notion of standard EUF-CMA (i.e., 0-probing). For this, we define **KeyUpdate** to be the identity function; the restriction that the adversary can only query an empty set for the set of probes is enforced by the winning condition.

5.2 Insufficiency of the NI/sNI Models

At first glance, all subroutines of Raccoon can be proven composable in the NI model. However, careful consideration shows that the NI model does not capture security when the intermediate values are not uniformly distributed and biased with the knowledge of the public output. Indeed, for example in the **KeyGen**, the combined knowledge of some shares of $\llbracket \mathbf{s} \rrbracket$ and of the public key \mathbf{vk} allows one to decrease the key-recovery security (decreasing the standard deviation of the short vector in a lattice) as presented in the technical overview in Section 1.

The gist of the problem when taking the output of an algorithm into account comes from the fact that the NI model proves that there exists a simulator that can simulate any set of probes from a subset of the input shared secrets of the

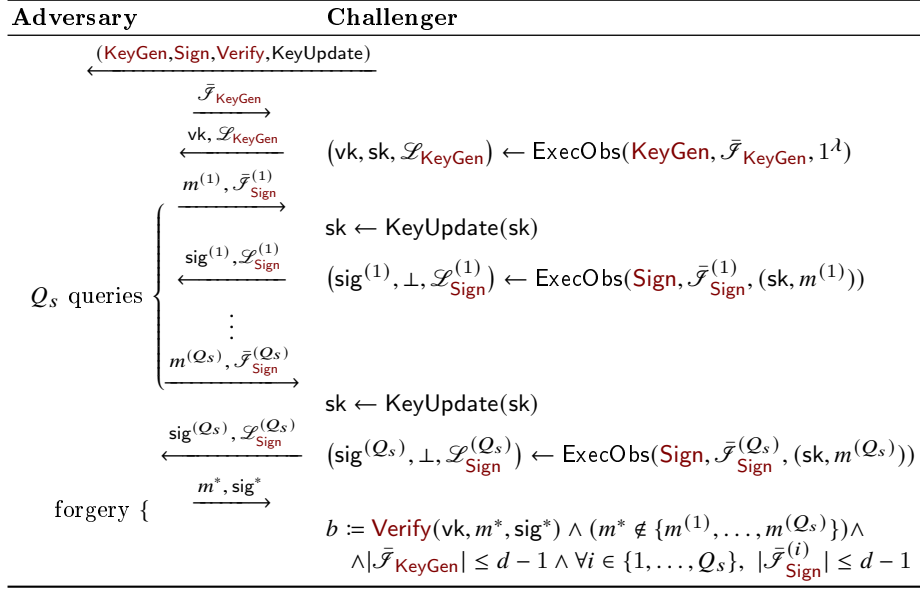


Fig. 3: EUF-CMA security game in the $d - 1$ -probing model. See Definition 2.1 for the definition of ExecObs.

algorithm. However, the aforementioned property does not entail that the distribution of the probes can be simulated when taking into account the output. This is clearly apparent in Definition 2.2 where the definition requires $\text{SimOut}(\mathbf{G}, \mathcal{X}')$ and \mathcal{L} to be identically distributed, but not $(\text{out}_{\text{unmasked}}, \text{SimOut}(\mathbf{G}, \mathcal{X}'))$ and $(\text{out}_{\text{unmasked}}, \mathcal{L})$.




Fig. 4: Example of an algorithm without unshared inputs (left), and its equivalent where randomnesses are explicitly passed as unshared inputs (right).

To see that the marginal distributions being identical is insufficient we give a simple example in Figure 4: both algorithms are trivially NI since any probe $\bar{\rho}_j$ or \bar{v}'_j can be simulated by sampling a small uniform and outputting it or adding it to the corresponding input v_j . However, if we consider the unmasked value

w as a public output, a simulator taking as input shares of $\llbracket \mathbf{v} \rrbracket$ cannot output probes that are correlated to w . For example, in gadget `non-NIU`, consider the set of probes $\bar{\mathcal{F}} = \{\bar{v}'_1\}$ which corresponds to the sum of v_1 and ρ_1 . A simulator $(\text{SimIn}, \text{SimOut})$ can perfectly simulate $\bar{\mathcal{F}}$ by setting $\text{SimIn}(\text{non-NIU}, \bar{\mathcal{F}}) := \bar{v}_1$, and $\text{SimOut}(\text{non-NIU}, v_1) := v_1 + \text{RSU}(u, 1)$. Then the variable $\mathcal{L} = v'_1$ being probed has the same distribution as $\text{SimOut}(\text{non-NIU}, v_1)$. However the distribution of $(\mathcal{L}, \text{out}_{\text{unmasked}}) = (v_1 + \rho_1, v + \rho_1 + \dots + \rho_d)$ is clearly not the same as that of $(\text{SimOut}(\text{non-NIU}, v_1), \text{out}_{\text{unmasked}}) = (v_1 + \text{RSU}(u, 1), v + \rho_1 + \dots + \rho_d)$.

5.3 NI/sNI with Unshared Inputs

To be able to handle cases where the values being probed are correlated with the public output we will modify the relevant gadgets and consider that any correlated random variables will be considered as inputs. We will formalize this idea with a model named Non-Interference with Unshared Inputs (NIU) (see Definitions 5.2 and 5.3 below), in which we will consider a variant of the algorithm where all random values that can affect the distribution of the output will be considered as inputs of the algorithm. While this model is stronger than the NI model, as it can be used to prove security even in the presence of leakage (see Lemma 5.2), we note that once an algorithm P has been modified to have its relevant randomness moved to inputs, the difference with the NI model becomes mostly syntactical since the new inputs of the algorithm and gadgets can be considered as just an additional shared secret input.

As an example, see the algorithm `NIU` in Figure 4 where we parse the random samples ρ_i as inputs rather than local variables. `NIU` thus takes two tuples of d values as input, and can as before be proven NI (where we artificially consider the tuple $(\rho_i)_{i \in [d]}$ as a shared input). However this time the NI proof does entail that the joint distribution of the probes and the output is identical to that of the simulator and output, because the output is a deterministic function of the input. Using the same set of probes $\bar{\mathcal{F}} = \bar{v}'_1$ as before, this time the simulator needs to use two input values to simulate the probe: $\text{SimIn}(\text{NIU}, \bar{\mathcal{F}}) := \{\bar{v}_1, \bar{\rho}_1\}$, however since each input variable is in a different shared input this simulator fits the definition of 2-NI in Definition 2.3, and we can set $\text{SimOut}(\text{NIU}, \{v_1, \rho_1\}) := v_1 + \rho_1$. It is obvious that in this case $(\mathcal{L}, \text{out}_{\text{unmasked}}) = (v_1 + \rho_1, v + \rho_1 + \dots + \rho_d) = (\text{SimOut}(\text{NIU}, \{v_1, \rho_1\}), \text{out}_{\text{unmasked}})$.

We will now first formalize the $(d - 1)$ -NIU notion, introduced in [20], in Definitions 5.2 and 5.3. Using the formalism of Section 2.2 we can then state and prove composition properties in Lemma 5.1, which are straightforward though never made explicit in [20]. Finally we can prove the core simulatability property of Lemma 5.2 which shows that when passing appropriate random variables as input NIU is sufficient to simulate the joint distribution of the probes and outputs of an algorithm. While this property was implicitly used in [20], it was actually never proven.

Definition 5.2 (Non Interference with Unshared input [20]). *Let G be a gadget taking two types of inputs:*

1. shared inputs \mathcal{X} , where all elements in \mathcal{X} are d -tuples of elements in \mathcal{R}_q
2. unshared input \mathcal{Y} , where all elements in \mathcal{Y} are tuples (not of fixed size) of elements in \mathcal{R}_q

A gadget \mathbf{G} with shared and unshared inputs is said $(d-1)$ -non-interfering with unshared inputs (written $(d-1)$ -NIU for short) iff any set of probes $\bar{\mathcal{F}}$ such that $|\bar{\mathcal{F}}| \leq d-1$ can be perfectly simulated (See Definition 2.2) by a simulator $(\text{SimIn}, \text{SimOut})$ such that $\text{SimIn}(\mathbf{G}, \bar{\mathcal{F}})$ outputs a set $\bar{\mathcal{X}}' \cup \bar{\mathcal{U}}$ of at most $d-1$ shares of each shared input ($\bar{\mathcal{X}}'$) and each unshared input ($\bar{\mathcal{U}}$).

Definition 5.3 (Strong Non Interference with Unshared input [20]).

A gadget is said $(d-1)$ -strongly-non-interfering with unshared inputs (written $(d-1)$ -sNIU for short) iff any set $\bar{\mathcal{F}}$ of at most $d-1 = d_{\text{int}} + d_{\text{out}}$ probes where d_{int} are made on internal data and d_{out} are made on the outputs can be simulated as in Definition 5.2 with d_{int} instead of $d-1$.

Since unshared inputs only differ from shared inputs by semantics (the distinction comes mostly from the fact that they do not represent a secret being used by the algorithm but internal randomnesses), one can note that if we ignore this distinction, the definitions of NIU and NI are identical. The interesting property of NIU comes from the fact that first transforming the relevant gadgets (namely **AddRepNoise**) to include the randomness as unshared inputs allows NIU to prove a meaningful statement on the joint distribution of the probes and the output. A key property we use to prove EUF-CMA in the probing model.

As argued earlier once the randomness is moved to inputs the definition of NIU becomes identical to the one of NI with the difference that inputs are separated in two sets by whether they are shared or unshared. Since this difference is purely syntactical the composition lemma of NI naturally extends to NIU.

Lemma 5.1 (Composability of NIU and sNIU gadgets). *A well-formed algorithm is NIU if all of its gadgets are NIU or sNIU and each sharing and each unshared variable is used at most once as input of a non-sNIU gadget. Moreover, a well-formed algorithm is sNIU if it is NIU and its output sharings are issued from an sNIU gadget.*

We now give a core lemma to use NIU. In essence the following lemma states that by passing the relevant randomnesses of a program to inputs, proving NIU becomes sufficient to prove that probes can be simulated even in the presence of outputs.

Lemma 5.2. *Let P be an algorithm with shared inputs \mathcal{X} and unshared inputs \mathcal{U} . If P is $(d-1)$ -NIU, and the public output of P is a deterministic function of $(\mathcal{X}, \mathcal{U})$. Then for any input \mathcal{X} and any probes $\bar{\mathcal{F}}$ (with $|\bar{\mathcal{F}}| \leq d-1$), the distribution of $(\text{out}_{\text{unmasked}}, \text{SimOut}(P, (\mathcal{X}', \mathcal{U}')))$ and $(\text{out}_{\text{unmasked}}, \mathcal{L})$ over the randomness \mathcal{U} and the random coins of P and SimOut are identical, where $(\text{out}_{\text{masked}}, \text{out}_{\text{unmasked}}, \mathcal{L}) \leftarrow \text{ExecObs}(P, \bar{\mathcal{F}}, \mathcal{X})$ and $(\bar{\mathcal{X}}', \bar{\mathcal{U}}') \leftarrow \text{SimIn}(P, \bar{\mathcal{F}})$.*

Proof. We will fix the input \mathcal{X} and not \mathcal{D} the distribution from which \mathcal{U} is sampled. \mathcal{L} and $\text{out}_{\text{unmasked}}$ are random variables over the choice of \mathcal{U} and

the random coins of P which we will note rc_P , and $\text{SimOut}(P, (\mathcal{X}', \mathcal{U}'))$ is a random variable over the choice of \mathcal{U} and the random coins of SimOut which we will note rc_S (SimOut only uses the randomness in $\mathcal{U}' \subset \mathcal{U}$ but we can consider it as a variable of \mathcal{U} since \mathcal{U}' is a marginal of \mathcal{U}). First we observe that since the definition of NI and NIU are identical if we simply consider the extra randomness as another input we have that the marginal distributions of \mathcal{L} and $\text{SimOut}(P, (\mathcal{X}', \mathcal{U}'))$ are identical, i.e. for any possible leakage Λ we have:

$$\Pr_{\mathcal{U} \leftarrow \mathcal{D}, rc_P \leftarrow \{0,1\}^*} [\mathcal{L}(\mathcal{X}, \mathcal{U}, rc_P) = \Lambda] = \Pr_{\mathcal{U} \leftarrow \mathcal{D}, rc_S \leftarrow \{0,1\}^*} [\text{SimOut}(\mathcal{X}, \mathcal{U}, rc_S) = \Lambda]$$

Since the algorithm P is deterministic when given $(\mathcal{X}, \mathcal{U})$, we have that for any possible leakage value Λ and output value θ :

$$\begin{aligned} & \Pr_{\mathcal{U} \leftarrow \mathcal{D}, rc_P \leftarrow \{0,1\}^*} [\mathcal{L}(\mathcal{X}, \mathcal{U}, rc_P) = \Lambda, out_{\text{unmasked}}(\mathcal{X}, \mathcal{U}) = \theta] \\ &= \sum_{\mathcal{U} \text{ s.t. } out_{\text{unmasked}}(\mathcal{X}, \mathcal{U}) = \theta} \Pr_{rc_P \leftarrow \{0,1\}^*} [\mathcal{L}(\mathcal{X}, \mathcal{U}, rc_P) = \Lambda] \\ &= \sum_{\mathcal{U} \text{ s.t. } out_{\text{unmasked}}(\mathcal{X}, \mathcal{U}) = \theta} \Pr_{rc_S \leftarrow \{0,1\}^*} [\text{SimOut}(\mathcal{X}, \mathcal{U}, rc_S) = \Lambda] \\ &= \Pr_{\mathcal{U} \leftarrow \mathcal{D}, rc_S \leftarrow \{0,1\}^*} [\text{SimOut}(\mathcal{X}, \mathcal{U}, rc_S) = \Lambda, out_{\text{unmasked}}(\mathcal{X}, \mathcal{U}) = \theta] \end{aligned}$$

which is the desired result. \square

6 NIU Property of Raccoon's KeyGen and Sign

Before establishing EUF-CMA security of Raccoon in the probing model, we prove that the **KeyGen** and **Sign** algorithms are NIU. Looking ahead, this allows a reduction to simulate the probes $\mathcal{L}_{\text{KeyGen}}$ and $\mathcal{L}_{\text{Sign}}^{(i)}$ in the EUF-CMA security game in the probing model (cf. Figure 3).

6.1 Existing Security Properties

Thanks to the composability of the sNI/NIU models, we can focus on the smaller gadgets comprising the **KeyGen** and **Sign** algorithms. Table 2 summarizes the security properties of the gadgets used in Raccoon, where we can rely on prior works to establish the security of every gadget, except for **AddRepNoise**. We refer to the cited papers for more information about the proofs.

6.2 Security Property of the AddRepNoise Gadget

Let us start with an intuition on the role of the **Refresh** operations in **AddRepNoise**. When considering unmasked coefficients, **AddRepNoise** is functionally equivalent to performing $a \leftarrow a + \text{SU}(u, T)$ for each coefficient a , for $T = d \cdot \text{rep}$. The internal

Table 2: Security properties of the known and new gadgets. No security property is necessary for the other unmasked operations (`ExpandA`, `ChalHash`, `ChalPoly`, `CheckBounds`, Computing the hint `h`).

Name	Property	Proof reference
<code>×A</code> and Line 13 of Algorithm 2	NI	\mathbb{Z}_q -linear
<code>Refresh</code> (Algorithm 7)	sNI	[7, 32, 24]
<code>ZeroEncoding</code> (Algorithm 8)	sNI	[32]
<code>Decode</code> (Algorithm 6)	NI	[5, Alg. 16]
<code>AddRepNoise</code> (Algorithm 5)	sNIU	Proved in Section 6.2, Lemma 6.1

use of `Refresh` operations does not affect this behavior but is meant to offer some resilience to probing adversaries.

Without `Refresh`, a viable strategy would be to probe individual shares of $\llbracket a \rrbracket$ at the start and at the end of `AddRepNoise`, allowing to learn the sum b of $\text{rep} \cdot (d-1)/2$ small uniform errors. The conditional distribution of the additive noise (conditioned on the $d-1$ probed values) is now $b + \text{SU}(u, T - (d-1) \cdot \text{rep}/2)$. With `Refresh`, this strategy is not possible anymore but a probing adversary can still probe individual errors, which in the end gives out no more than the sum b of $d-1$ small uniform errors. The conditional distribution of the additive noise (conditioned on the $d-1$ probed values) is now $b + \text{SU}(u, T - (d-1))$, where the adversary learns b but knows nothing about the realization of $\text{SU}(u, T - (d-1))$.

While `AddRepNoise` performs operations share by share, the underlying distributions are not uniform. The addition of short noise values are added biases the *a posteriori* distribution of the final noise. Hence, one cannot prove that this gadget is probing secure. We resolve this issue by moving the short noise values as random coin inputs of the algorithm, introducing `AddRepNoiseER` in Algorithm 9, an instance of `AddRepNoise` with explicit randomness (ER) for the small uniforms. Note that the complete set of small uniforms is considered as a single unshared input. We can now formally show in Lemma 6.1 that `AddRepNoiseER` is sNIU. A similar result was proven in [20] but our proof strategy is different and perhaps a bit more formal. Later, these inputs will be handled in the general composition proof.

Lemma 6.1. *The `AddRepNoiseER` gadget is $(d-1)$ -sNIU.*

Proof. We represent `AddRepNoiseER` as a sequential succession of `MiniAddRepNoise` and `Refresh` as presented in Figure 5. To prove the sNIU property, we exhibit the randomness $\rho_{i, i_{\text{rep}}, j}$ in the input. Let us remark that the randomness involved in `Refresh` (and thus in `ZeroEncoding`) are not explicated as the algorithm is already proved sNI. Hence, `AddRepNoiseER` is partially derandomized. Our proof proceeds in two steps; we first study the `MiniAddRepNoise` sub-gadget, then `AddRepNoiseER`.

Step 1: MiniAddRepNoise. We first show that any probe inside `MiniAddRepNoise` can be perfectly simulated (see Definition 2.2) with $\rho_{i, i_{\text{rep}}, j}$ and the input \mathbf{v}_j , where (i, i_{rep}, j) corresponds to the targeted loop. Indeed, let p be a probe inside `MiniAddRepNoise`. The description of this probe necessarily includes (i, i_{rep}, j) to specify the involved loop. The intermediate value targeted by p can be




Fig. 5: Structure of $\text{AddRepNoise}_{\text{ER}}$ (using Algorithm 10). A gadget proven sNI is noted $\boxed{\text{gadget}}$. The gadgets with no proven property are noted $\boxed{\text{gadget}}$. Single arrows (\rightarrow) and double arrows (\Rightarrow) represent plain and masked values, respectively.

Algorithm 9 $\text{AddRepNoise}_{\text{ER}}(\llbracket \mathbf{v} \rrbracket, (\rho_{i,i_{\text{rep}},j})) \rightarrow \llbracket \mathbf{v}' \rrbracket$, w/ partial explicit randomness

Input: Masked vector $\llbracket \mathbf{v} \rrbracket = (\mathbf{v}_j)_{j \in [d]} = (v_{i,j})_{i \in [\text{len}(\mathbf{v})], j \in [d]}$.

Input: Randomness $(\rho_{i,i_{\text{rep}},j})_{i \in [\text{len}(\mathbf{v})], i_{\text{rep}} \in [\text{rep}], j \in [d]}$

Output: Updated $\llbracket \mathbf{v} \rrbracket$ with $\text{SU}(u, d \cdot \text{rep})$ distribution added to each coefficient of \mathbf{v} .

- 1: **for** $(i, i_{\text{rep}}) \in [\text{len}(\mathbf{v})] \times [\text{rep}]$ **do** \triangleright Vector index.
 - 2: **for** $i_{\text{rep}} \in [\text{rep}]$ **do**
 - 3: $\llbracket \mathbf{v}_i \rrbracket \leftarrow \text{MiniAddRepNoise}(\llbracket \mathbf{v}_i \rrbracket, (\rho_{i,i_{\text{rep}},j})_{i \in [\text{len}(\mathbf{v})], j \in [d]})$
 - 4: $\llbracket \mathbf{v}_i \rrbracket \leftarrow \text{Refresh}(\llbracket \mathbf{v}_i \rrbracket)$ \triangleright Refresh polynomial on each repeat.
 - 5: **return** $\llbracket \mathbf{v} \rrbracket$
-

Algorithm 10 $\text{MiniAddRepNoise}(\llbracket \mathbf{v} \rrbracket, i_{\text{rep}}, (\rho_{i,i_{\text{rep}},j})) \rightarrow \llbracket \mathbf{v}' \rrbracket$

Input: Masked vector $\llbracket \mathbf{v}' \rrbracket$, index $i_{\text{rep}} \in [\text{rep}]$, randomness $(\rho_{i,i_{\text{rep}},j})_{i \in [\text{len}(\mathbf{v})], j \in [d]}$

Output: Updated $\llbracket \mathbf{v} \rrbracket$.

- 1: **for** $j \in [d]$ **do**
 - 2: $v'_j \leftarrow v_j + \rho_{i,i_{\text{rep}},j}$
 - 3: **return** $\llbracket \mathbf{v}' \rrbracket$
-

1. the randomness $\rho_{i,i_{\text{rep}},j}$,
2. the value v_j or v'_j .

It is easy to conclude that any of these values can be perfectly simulated from $\rho_{i,i_{\text{rep}},j}$ and the input \mathbf{v}_j . The only intermediate value that needs both is v'_j as it needs $\rho_{i,i_{\text{rep}},j}$.

Step 2: AddRepNoise_{ER}. Let us now look at the bigger picture. In this proof, we will perform a composition proof by propagating the dependency of the intermediate variables to shares of $\rho_{i,i_{\text{rep}},j}$ and \mathbf{v}_j . Let $\bar{\mathcal{F}}$ be the given set of at most $d - 1$ probes in AddRepNoise . We decompose $\bar{\mathcal{F}}$ as follows.

- Let $\delta_{\text{MiniAddRepNoise}}^{i,i_{\text{rep}}}$ be the number intermediate variables that are probed inside the MiniAddRepNoise gadget of the loop with indexes i, i_{rep} .

- Let $\delta_{\text{Refresh}}^{i, i_{\text{rep}}}$ be the number intermediate variables that are probed inside the **Refresh** gadget of the loop with indexes i, i_{rep} .

By definition,

$$\sum_{i=0}^{\text{len}(\mathbf{v})} \sum_{i_{\text{rep}}=0}^{\text{rep}} \left(\delta_{\text{MiniAddRepNoise}}^{i, i_{\text{rep}}} + \delta_{\text{Refresh}}^{i, i_{\text{rep}}} \right) \leq d - 1. \quad (9)$$

Going from right to left in Figure 5, we first consider the last **Refresh** of the last loop (where $i = \text{len}(\mathbf{v})$ and $i_{\text{rep}} = \text{rep}$). Thanks to the sNI property of the last **Refresh** algorithm, all the $\delta_{\text{Refresh}}^{\text{len}(\mathbf{v}), \text{rep}}$ probes can be perfectly simulated from $\delta_{\text{Refresh}}^{\text{len}(\mathbf{v}), \text{rep}}$ shares of \mathbf{v}' , which is also the output of the last **MiniAddRepNoise**. So, thanks to the above paragraph about **MiniAddRepNoise**, all the probes from the last **MiniAddRepNoise**, can be perfectly simulated from two sets of probes:

- $\bar{\mathcal{F}}_{\text{len}(\mathbf{v}), \text{rep}}$ defined as the description of at most $\delta_{\text{MiniAddRepNoise}}^{\text{len}(\mathbf{v}), \text{rep}} + \delta_{\text{Refresh}}^{\text{len}(\mathbf{v}), \text{rep}}$ values of $\rho_{\text{len}(\mathbf{v}), \text{rep}, j}$ (with several different j 's),
- $\bar{\mathcal{F}}'_{\text{len}(\mathbf{v}), \text{rep}}$ defined as the set of to at most $\delta_{\text{MiniAddRepNoise}}^{\text{len}(\mathbf{v}), \text{rep}} + \delta_{\text{Refresh}}^{\text{len}(\mathbf{v}), \text{rep}}$ shares of \mathbf{v} , the input of the last **MiniAddRepNoise**.

The set of $\bar{\mathcal{F}}'_{\text{len}(\mathbf{v}), \text{rep}}$ can also be seen as probes of the output of the penultimate **Refresh**. But, thanks to the sNI property of the penultimate **Refresh** algorithm, they can be simulated independently from the $\delta_{\text{Refresh}}^{\text{len}(\mathbf{v})-1, \text{rep}-1}$ intermediate variables probed inside the penultimate **Refresh** algorithm. In conclusion, the $\bar{\mathcal{F}}'_{\text{len}(\mathbf{v}), \text{rep}}$ probes can be simulated from uniform random.

Applying the same reasoning for all the subsequent loops, the set of $\bar{\mathcal{F}}$ probes can be perfectly simulated from

- $\bar{\mathcal{F}}_{i, i_{\text{rep}}}$ defined as the description of at most $\delta_{\text{MiniAddRepNoise}}^{i, i_{\text{rep}}} + \delta_{\text{Refresh}}^{i, i_{\text{rep}}}$ values of $\rho_{i, i_{\text{rep}}, j}$ (with several different j 's),
- $\bar{\mathcal{F}}'_{0,0}$ defined as the set of to at most $\delta_{\text{MiniAddRepNoise}}^{0,0} + \delta_{\text{Refresh}}^{0,0}$ shares of \mathbf{v} , the input of the **AddRepNoise_{ER}**.

We define $\bar{\mathcal{U}} = \bar{\mathcal{F}}_{0,0} \cup \dots \cup \bar{\mathcal{F}}_{\text{len}(\mathbf{v}), \text{rep}}$ and $\bar{\mathcal{X}}' = \bar{\mathcal{F}}'_{0,0}$. Thanks to Eq. (9) and Lemma 2.1, we have shown that **AddRepNoise_{ER}** is (d-1)-sNIU. \square

6.3 Security Property of **KeyGen** and **Sign**

Now that **AddRepNoise_{ER}** is proved, one needs to derive the security of the key generation and signature algorithms with a composition proof. Let us first introduce **KeyGen_{ER}** and **Sign_{ER}**, simple modifications of **KeyGen** and **Sign** algorithms where the small uniform randomness is provided as input. **KeyGen_{ER}** is formally described in Algorithm 11. Due to space constraints, the formal description of **Sign_{ER}** is deferred to the full version. We provide a proof of Lemma 6.2 for **KeyGen_{ER}**. The proof for **Sign_{ER}** proceeds in a similar fashion and is included in the full version of this paper.

Lemma 6.2. *The algorithms **KeyGen_{ER}** and **Sign_{ER}** are (d - 1)-NIU.*

Algorithm 11 $\text{KeyGen}_{\text{ER}}((\rho_{i,i_{\text{rep}},j}^{(0)}), (\rho_{i,i_{\text{rep}},j}^{(1)})) \rightarrow (\text{vk}, \text{sk})$

▷ KeyGen with explicit randomness for AddRepNoise

Input: Randomness $(\rho_{i,i_{\text{rep}},j}^{(0)})_{i \in [\text{len}(\mathbf{v})], i_{\text{rep}} \in [\text{rep}], j \in [d]}$, $(\rho_{i,i_{\text{rep}},j}^{(1)})_{i \in [\text{len}(\mathbf{v})], i_{\text{rep}} \in [\text{rep}], j \in [d]}$

Output: Keypair vk, sk

- 1: $\text{seed} \leftarrow \{0, 1\}^{\kappa}$; $\mathbf{A} := \text{ExpandA}(\text{seed})$
 - 2: $\llbracket \mathbf{s} \rrbracket \leftarrow \ell \times \text{ZeroEncoding}(d)$
 - 3: $\llbracket \mathbf{s} \rrbracket \leftarrow \text{AddRepNoise}_{\text{ER}}(\llbracket \mathbf{s} \rrbracket, u_{\text{t}}, \text{rep}, (\rho_{i,i_{\text{rep}},j}^{(0)}))$ ▷ Partially derandomized
 AddRepNoise.
 - 4: $\llbracket \mathbf{t} \rrbracket := \mathbf{A} \cdot \llbracket \mathbf{s} \rrbracket$
 - 5: $\llbracket \mathbf{t} \rrbracket \leftarrow \text{AddRepNoise}_{\text{ER}}(\llbracket \mathbf{t} \rrbracket, u_{\text{t}}, \text{rep}, (\rho_{i,i_{\text{rep}},j}^{(1)}))$ ▷ Partially derandomized
 AddRepNoise.
 - 6: $\mathbf{t} := \text{Decode}(\llbracket \mathbf{t} \rrbracket)$
 - 7: $\mathbf{t} := \llbracket \mathbf{t} \rrbracket_{v_{\text{t}}}$
 - 8: **return** $(\text{vk} := (\text{seed}, \mathbf{t}), \text{sk} := (\text{vk}, \llbracket \mathbf{s} \rrbracket))$
-




Fig. 6: Structure of KeyGen (Algorithm 11). Gadgets proven NI (resp. sNIU) is noted gadget (resp. gadget). Triangular gadgets either start from a masked input and output an unmasked value, or the other way around.

Proof (Lemma 6.2). Let us decompose the key generation as a succession of gadgets. The gadgets may be represented as in Figure 6. We assume the respective NI/sNI/sNIU properties of each gadget as presented in Table 2.

Recall that given a set $\bar{\mathcal{F}}$ of at most $d - 1$ probes inside $\text{KeyGen}_{\text{ER}}$, we aim at proving that they can be perfectly simulated with at most $d - 1$ shares of $(\rho_{i,i_{\text{rep}},j}^{(0)})$ and $d - 1$ shares of $(\rho_{i,i_{\text{rep}},j}^{(1)})$. In other words we will exhibit two sets $\bar{\mathcal{F}}_0$ of at most $d - 1$ values of $(\rho_{i,i_{\text{rep}},j}^{(0)})$, and $\bar{\mathcal{F}}_1$ of at most $d - 1$ values of $(\rho_{i,i_{\text{rep}},j}^{(1)})$ which will be enough to perfectly simulate $\bar{\mathcal{F}}$.

Let us decompose the set $\bar{\mathcal{F}}$ of at most $d - 1$ probes in $\text{KeyGen}_{\text{ER}}$ among the different gadgets. By convention, to avoid counting certain probes twice (once as output of a gadget and once as input of the subsequent gadget), we do not count the probes on the outputs. For example, if a probe is made on the output of a gadget \mathbf{G} , we will consider that it is actually made on the input of the subsequent gadget. We note:

- δ_0 the number of intermediate variables probed in Line 6 (final **Decode** gadget);
 - δ_1 the number of intermediate variables probed in Line 5 (second **AddRepNoise_{ER}**);
 - δ_2 the number of intermediate variables probed in Line 4 (multiplication with **A**);
 - δ_3 the number of intermediate variables probed in Line 3 (first **AddRepNoise_{ER}**);
 - δ_4 the number of intermediate variables probed in Line 2 (**ZeroEncoding**);
- We recall that by definition of $\bar{\mathcal{F}}$, $\sum_{i=0}^4 \delta_i \leq d - 1$.

The proof is similar to a standard composition proof. Thanks to the NI property of the **Decode** gadget, all the δ_0 intermediate variables can be perfectly simulated (see Definition 2.2) with at most δ_0 shares of $\llbracket t \rrbracket$. Since the second **AddRepNoise_{ER}** is $d - 1$ -sNIU, the $\delta_1 + \delta_0$ intermediate variables observed during **Decode** and the last **AddRepNoise_{ER}** may be perfectly simulated with δ_1 shares of $\llbracket t \rrbracket$ (the output of the $\times \mathbf{A}$ operation) and δ_1 shares of $(\rho_{i,i_{\text{rep}}}^{(1)})$. We note $\bar{\mathcal{F}}_1$ this set. Note that δ_0 is discarded as it concerns the output of a sNIU gadget.

With the same reasoning, all the $\delta_0 + \delta_1 + \delta_2 + \delta_3$ intermediate variables observed after the first **AddRepNoise_{ER}** can be perfectly simulated with at most δ_3 shares of $\llbracket s \rrbracket$ (which are also the output of **ZeroEncoding**) and at most δ_3 shares of $(\rho_{i,i_{\text{rep}}}^{(0)})$. We note $\bar{\mathcal{F}}_0$ this sets. In addition, the δ_4 intermediate variables in the **ZeroEncoding** gadget may be perfectly simulated from the public parameters as **ZeroEncoding** is NI and does not take any input.

Putting everything together, we have proved that the distribution of the intermediate variables in $\bar{\mathcal{F}}$ may be perfectly simulated from :

- the set $\bar{\mathcal{F}}_0$ containing at most δ_3 shares of $(\rho_{i,i_{\text{rep}}}^{(0)})$
- the sets $\bar{\mathcal{F}}_1$ containing at most δ_1 shares of $(\rho_{i,i_{\text{rep}}}^{(1)})$

Since $\delta_3 + \delta_1 \leq \sum_{i=0}^4 \delta_i \leq d - 1$, we have exhibited a set $\bar{\mathcal{U}}$ of at most $d - 1$ of the unshared input which concludes the proof. \square

7 EUF-CMA Security of Raccoon in the Probing Model

We are finally ready to prove EUF-CMA security of Raccoon in the probing model. This is done in two steps. We first reduce EUF-CMA security of Raccoon in the probing model to the *standard* EUF-CMA security of *small* Raccoon, formally defined in Figure 7. We then establish that this small Raccoon is EUF-CMA secure. Technically, the first part relies on the NIU property of **KeyGen** and **Sign** (cf. Section 6), a purely statistical step claiming that given a small Raccoon key and signature, we can simulate the leakage of Raccoon. The second part relies on the smooth Rényi divergence for the sum of uniform distributions (cf. Section 4), and reduces to computational problems.

7.1 Description of a Non-Masked Small Raccoon

We first formally define a *non-masked* and simplified variant of Raccoon, called *small* Raccoon, depicted in Figure 7. Notice that there are no more masking or

bit-droppings applied. More importantly, it is “small” since the sum of uniform distribution is smaller. We effectively modify the bounds on the signature size to be smaller, using \bar{B}_∞ and \bar{B}_2 , whose formal definition appears in Theorem 7.1.

<p>Algorithm 12 $\text{KeyGen}_{\text{SMALL}}(\emptyset) \rightarrow (\text{vk}, \text{sk})$</p> <hr/> <p>Output: Keypair vk, sk</p> <p>1: $\text{seed} \leftarrow \{0, 1\}^k$</p> <p>2: $\mathbf{A} := \text{ExpandA}(\text{seed})$</p> <p>3: $(s, e) \leftarrow \text{RSU}(u_t, d(\text{rep} - 1) + 1)^\ell \times \text{RSU}(u_t, d(\text{rep} - 1) + 1)^k$</p> <p>4: $\mathbf{t} := \mathbf{A} \cdot s + e$ \triangleright No rounding of $\mathbf{t} \in \mathcal{R}_q^k$</p> <p>5: return $(\text{vk} := (\text{seed}, \mathbf{t}), \text{sk} = (\text{vk}, s))$</p> <hr/> <p>Algorithm 13 $\text{Sign}_{\text{SMALL}}(\text{sk}, \text{msg}) \rightarrow \text{sig}$</p> <hr/> <p>Input: Secret signing key $\text{sk} = (\text{vk}, s)$, message $\text{msg} \in \{0, 1\}^*$.</p> <p>Output: Signature $\text{sig} = (c_{\text{hash}}, \mathbf{h}, \mathbf{z})$ of msg under sk.</p> <p>1: $\mu := \text{H}(\text{H}(\text{vk}) \parallel \text{msg})$</p> <p>2: $(\mathbf{r}, \mathbf{e}') \leftarrow \text{RSU}(u_w, d(\text{rep} - 1) + 1)^\ell \times \text{RSU}(u_w, d(\text{rep} - 1) + 1)^k$</p> <p>3: $\mathbf{w} = \mathbf{A} \cdot \mathbf{r} + \mathbf{e}'$ \triangleright No rounding of $\mathbf{w} \in \mathcal{R}_q^k$</p> <p>4: $c_{\text{hash}} := \text{ChalHash}(\mathbf{w}, \mu)$ \triangleright ChalHash redefined to take $\mathbf{w} \in \mathcal{R}_q^k$</p> <p>5: $c_{\text{poly}} := \text{ChalPoly}(c_{\text{hash}})$</p> <p>6: $\mathbf{z} := c_{\text{poly}} \cdot s + \mathbf{r}$</p> <p>7: $\mathbf{y} := \mathbf{A} \cdot \mathbf{z} - c_{\text{poly}} \cdot \mathbf{t}$ \triangleright No need to lift \mathbf{t} anymore</p> <p>8: $\mathbf{h} := \mathbf{w} - \mathbf{y}$ \triangleright Hint \mathbf{h} now defined over \mathcal{R}_q^k</p> <p>9: $\text{sig} := (c_{\text{hash}}, \mathbf{z}, \mathbf{h})$</p> <p>10: if $(\ \mathbf{z}, \mathbf{h}\ _\infty > \bar{B}_\infty)$ or $(\ \mathbf{z}, \mathbf{h}\ _2 > \bar{B}_2)$ goto Line 2 \triangleright Check smaller bound</p> <p>11: return sig</p> <hr/> <p>Algorithm 14 $\text{Verify}_{\text{SMALL}}(\text{sig}, \text{msg}, \text{vk}) \rightarrow \{\text{OK or FAIL}\}$</p> <hr/> <p>Input: Signature $\text{sig} = (c_{\text{hash}}, \mathbf{h}, \mathbf{z}) := \text{sig}$.</p> <p>Output: Signature validity: OK (accept) or FAIL (reject).</p> <p>1: if $(\ \mathbf{z}, \mathbf{h}\ _\infty > \bar{B}_\infty)$ or $(\ \mathbf{z}, \mathbf{h}\ _2 > \bar{B}_2)$ return FAIL else return OK</p> <p>2: $\mu := \text{H}(\text{H}(\text{vk}) \parallel \text{msg}); \mathbf{A} := \text{ExpandA}(\text{seed})$</p> <p>3: $c_{\text{poly}} := \text{ChalPoly}(c_{\text{hash}})$</p> <p>4: $\mathbf{y} := \mathbf{A} \cdot \mathbf{z} - c_{\text{poly}} \cdot \mathbf{t}$</p> <p>5: $\mathbf{w} := \mathbf{y} + \mathbf{h}$</p> <p>6: $c'_{\text{hash}} := \text{ChalHash}(\mathbf{w}, \mu)$</p> <p>7: if $c_{\text{hash}} \neq c'_{\text{hash}}$ return FAIL</p> <p>8: return OK</p>

Fig. 7: A non-masked and simplified Raccoon, named *small* Raccoon. While we used the notation from the masked Raccoon for consistency, notice above that \mathbf{h} simply becomes $c_{\text{poly}} \cdot \mathbf{e} + \mathbf{e}'$ without rounding errors.

7.2 EUF-CMA Security of Small Raccoon \Rightarrow Probing EUF-CMA Security of Raccoon

This consists of the first step. Once the following theorem is established, we only need to prove standard EUF-CMA security of small Raccoon.

Theorem 7.1. *Let B_∞ and B_2 satisfying:*

$$\begin{aligned} - \bar{B}_\infty &\geq B_\infty + \omega \cdot (d-1) \cdot \left(\frac{1}{2} + \frac{2^{3u}}{3}\right) \cdot (\kappa + \log(n(k+\ell))) + 2^{\gamma_w} + \omega \cdot 2^{\gamma_t} \\ - \bar{B}_2 &\geq B_2 + \omega \cdot \sqrt{n(k+\ell)} \cdot (d-1) \cdot \left(\frac{1}{2} + \frac{2^{3u}}{3}\right) \cdot (\kappa + \log(n(k+\ell))) + 2^{\gamma_w} \cdot \sqrt{nk} + \omega \cdot 2^{\gamma_t} \cdot \sqrt{nk} \end{aligned}$$

Let Q_H and Q_S denote the number of random oracle queries and signing queries performed by \mathcal{A} . For any PPT adversary \mathcal{A} against the EUF-CMA security on Raccoon in the $(d-1)$ -probing model with time T and advantage ε , there exists a PPT adversary \mathcal{B} against the EUF-CMA security on small Raccoon (cf. Figure 7) with time $O(T)$ and advantage:

$$\text{Adv}_{\mathcal{B}} \geq \text{Adv}_{\mathcal{A}} - 4Q_H Q_S \cdot 2^{-2\kappa} - 2^{-\kappa+1} - \frac{1}{|C|}.$$

We will use a series of hybrids defined below to prove the theorem.

Hybrid₀: This hybrid corresponds to real the EUF-CMA security game in the $(d-1)$ -probing model (cf. Figure 3).

Hybrid₁: In this hybrid we replace **KeyGen** with **KeyGen_{ER}** and **Sign** with **Sign_{ER}**, in which all randomnesses are sampled prior to running the algorithm. Since the algorithms are functionnaly identical the advantage is unchanged.

Hybrid₂: This hybrid corresponds to Figure 8, in which all the probes queried by the adversary during either key generation or signature are mapped to probes that target only the randomness used in the **AddRepNoise** gadgets. We prove that the values output by these probes can be used to perfectly simulate the output queried by the adversary in Lemma 6.2.

More precisely there is a first PPT simulator ($\text{SimIn}_{\text{KeyGen}}, \text{SimOut}_{\text{KeyGen}}$) such that for any probe set $|\bar{\mathcal{F}}_{\text{KeyGen}}| \leq t$ in **KeyGen**(1^κ), all probes in $\bar{\mathcal{F}}' := (\bar{\mathcal{F}}'_s, \bar{\mathcal{F}}'_e) := \text{SimIn}_{\text{KeyGen}}(\bar{\mathcal{F}}_{\text{KeyGen}})$ are of the form $\bar{\rho}_{s,i,i_{\text{rep}},j} \in \bar{\mathcal{F}}'_s$ for some $(i, i_{\text{rep}}, j) \in [\ell, \text{rep}, d]$, and $\bar{\rho}_{e,i,i_{\text{rep}},j} \in \bar{\mathcal{F}}'_e$ for some $(i, i_{\text{rep}}, j) \in [k, \text{rep}, d]$ (note that the variable names $\bar{\rho}$ are also indexed by the **AddRepNoise** gadget to which they belong to ensure unique namings), and $\max(|\bar{\mathcal{F}}'_s|, |\bar{\mathcal{F}}'_e|) \leq d-1$. Using Lemma 5.2 we have that $(\text{vk}, \text{SimOut}(\text{KeyGen}_{\text{ER}}, \bar{\mathcal{F}}'))$ follows the same distribution as (vk, \mathcal{L}) , where $(\text{sk}, \text{vk}, \mathcal{L}) \leftarrow \text{ExecObs}(\bar{\mathcal{F}}_{\text{KeyGen}}, \text{KeyGen}_{\text{ER}}, 1^\lambda)$. Similarly there is a second PPT simulator ($\text{SimIn}_{\text{Sign}}, \text{SimOut}_{\text{Sign}}$) such that for any message msg , masked secret key $\llbracket \text{sk} \rrbracket$, and probe set $|\bar{\mathcal{F}}_{\text{Sign}}| \leq t$ in **Sign**($\llbracket \text{sk} \rrbracket, \text{msg}$), all probes in $\bar{\mathcal{F}}' := (\bar{\mathcal{F}}'_r, \bar{\mathcal{F}}'_e, \bar{\mathcal{F}}'_{\text{sk}}) := \text{SimIn}_{\text{Sign}}(\bar{\mathcal{F}}_{\text{Sign}})$ are of the form $\bar{\rho}_{r,i,i_{\text{rep}},j} \in \bar{\mathcal{F}}'_r$ for some $(i, i_{\text{rep}}, j) \in [\ell, \text{rep}, d]$, $\bar{\rho}_{e',i,i_{\text{rep}},j} \in \bar{\mathcal{F}}'_e$ for some $(i, i_{\text{rep}}, j) \in [k, \text{rep}, d]$, and $\bar{s}_i \in \bar{\mathcal{F}}'_{\text{sk}}$ for some $i \in [d]$, and $\max(|\bar{\mathcal{F}}'_r|, |\bar{\mathcal{F}}'_e|, |\bar{\mathcal{F}}'_{\text{sk}}|) \leq t$.

It also holds that $(\text{sig}, \text{SimOut}(\text{ExecObs}(\bar{\mathcal{F}}', \text{Sign}, 1^\lambda)))$ follows the same distribution as $\text{ExecObs}(\bar{\mathcal{F}}_{\text{Sign}}, \text{Sign}, 1^\lambda)$. From Lemma 5.2, $\text{SimOut}(\text{Sign}_{\text{ER}}, \bar{\mathcal{F}}')$ follows the same distribution as $(\text{sig}, \mathcal{L})$, where $(\text{sig}, \mathcal{L}) \leftarrow \text{ExecObs}(\bar{\mathcal{F}}_{\text{Sign}}, \text{Sign}_{\text{ER}}, \text{msg})$. Thus the two hybrids are identical.

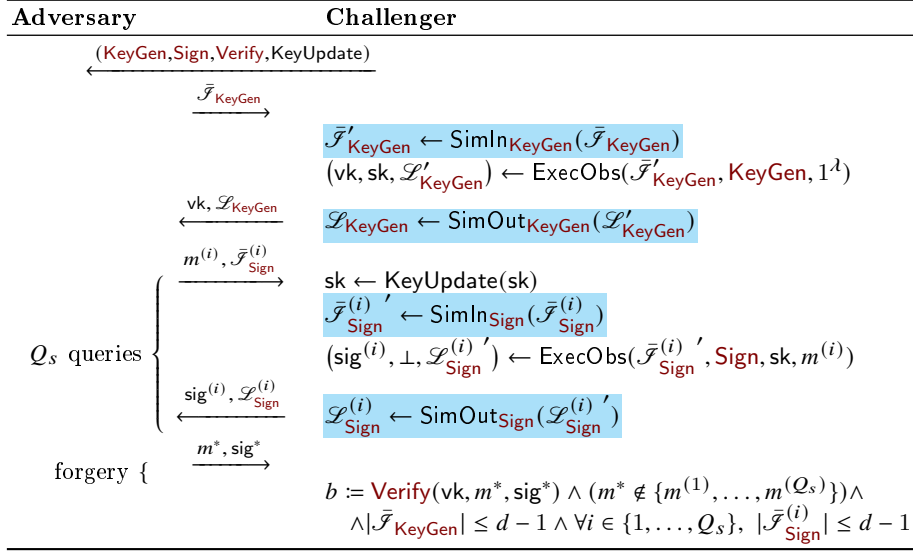


Fig. 8: Hybrid₂: The NIU properties proven in Lemma 6.2 ensure the existence of two PPT simulators ($\text{SimIn}_{\text{KeyGen}}, \text{SimOut}_{\text{KeyGen}}$) and ($\text{SimIn}_{\text{Sign}}, \text{SimOut}_{\text{Sign}}$). This ensures all probes can be moved to the randomness in the **AddRepNoise** gadgets in **KeyGen** and **Sign**. Differences with the EUF-CMA security game in the $(d-1)$ -probing model (Figure 3) are highlighted.

Algorithm 15 $\text{KeyGen}_{\mathcal{L}}(1^\kappa, \bar{\mathcal{F}}) \rightarrow (\text{vk}, \text{sk}, \mathcal{L})$

Input: Probe set $\mathcal{F} = (\mathcal{F}_s, \mathcal{F}_e)$, $\bar{\mathcal{F}}_s \subset \{\bar{\rho}_{s,i,i_{\text{rep}},j}; (i,i_{\text{rep}},j) \in [\ell] \times [\text{rep}] \times [d]\}$,

$\bar{\mathcal{F}}_e \subset \{\bar{\rho}_{e,i,i_{\text{rep}},j}; (i,i_{\text{rep}},j) \in [k] \times [\text{rep}] \times [d]\}$

Output: Keypair vk, sk and Leakage \mathcal{L}

- 1: $\text{seed} \leftarrow \{0,1\}^\kappa$; $\mathbf{A} := \text{ExpandA}(\text{seed})$
 - 2: $\llbracket \mathbf{s} \rrbracket = (s_1, \dots, s_d) := (0, \dots, 0) \in (\mathcal{R}_q^\ell)^d$
 - 3: **for** $(i, i_{\text{rep}}, j) \in [\ell] \times [\text{rep}] \times [d]$ **do**
 - 4: $\rho_{s,i,i_{\text{rep}},j} \leftarrow \text{RSU}(u, 1)$
 - 5: $s_{j,i} \leftarrow s_{j,i} + \rho_{s,i,i_{\text{rep}},j}$
 - 6: $\llbracket \mathbf{t} \rrbracket := \mathbf{A} \cdot \llbracket \mathbf{sk} \rrbracket \in (\mathcal{R}_q^k)^d$
 - 7: **for** $(i, i_{\text{rep}}, j) \in [k] \times [\text{rep}] \times [d]$ **do**
 - 8: $\rho_{e,i,i_{\text{rep}},j} \leftarrow \text{RSU}(u, 1)$
 - 9: $t_{j,i} \leftarrow t_{j,i} + \rho_{e,i,i_{\text{rep}},j}$
 - 10: $\mathbf{t} := \text{Decode}(\llbracket \mathbf{t} \rrbracket)$
 - 11: $\mathbf{t} := \lfloor \mathbf{t} \rfloor_{v_t}$
 - 12: $\mathcal{L} := \left\{ (\rho_{s,i,i_{\text{rep}},j}, \rho_{e,i',i'_{\text{rep}},j'}) \right\}_{(\bar{\rho}_{s,i,i_{\text{rep}},j}, \bar{\rho}_{e,i',i'_{\text{rep}},j'}) \in \bar{\mathcal{F}}}$
 - 13: **return** $(\text{vk} := (\text{seed}, \mathbf{t}), \text{sk} := (\text{vk}, \llbracket \mathbf{s} \rrbracket), \mathcal{L})$
-

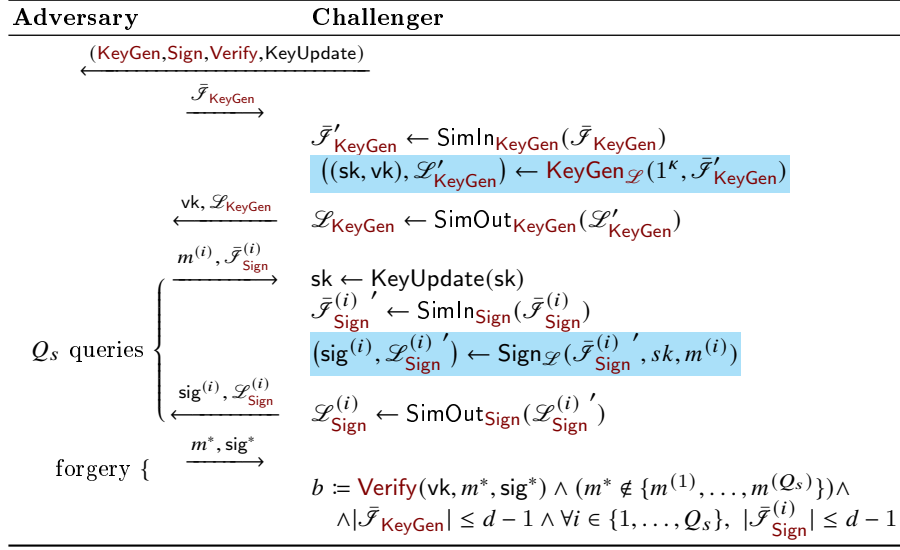


Fig. 9: Hybrid₃: We replace the ExecObs calls with the functionally identical algorithms $\text{KeyGen}_{\mathcal{G}}$ (cf. Algorithm 15) and $\text{Sign}_{\mathcal{G}}$ (cf. full version).

Hybrid₃: This hybrid corresponds to Figure 9, in which the algorithms $\text{ExecObs}(\bar{\mathcal{F}}, \text{KeyGen}, 1^\kappa)$ and $\text{ExecObs}(\bar{\mathcal{F}}, \text{Sign}, \text{sk}, \text{msg})$ are replaced by $\text{KeyGen}_{\mathcal{G}}(1^\kappa, \bar{\mathcal{F}})$ and $\text{Sign}_{\mathcal{G}}(\text{sk}, \text{msg}, \bar{\mathcal{F}})$, respectively. The former is presented in Algorithm 15. The latter is defined analogously and deferred to the full version due to page limitations. Observe that since $\text{ExecObs}(\bar{\mathcal{F}}, \text{KeyGen}, 1^\kappa)$ outputs the same output as $\text{KeyGen}(1^\kappa)$ as well as the value of the variables at indices $\bar{\mathcal{F}}$, any algorithm that outputs the same distribution is semantically identical. Since the variables in $\bar{\mathcal{F}}$ are now restricted to the randomness used in AddRepNoise it is clear that the algorithm $\text{KeyGen}_{\mathcal{G}}$ outputs the same distribution. The same argument goes for $\text{ExecObs}(\bar{\mathcal{F}}, \text{Sign}, \text{sk}, \text{msg})$. Hence, the two hybrids are identical.

Hybrid₄: This hybrid corresponds to Figure 10, in which the challenger artificially extends the set of probes queried to the key generation and signing algorithm. More specifically, we define Extend so that for any $\rho_{s,i,i_{\text{rep}},j} \in \bar{\mathcal{F}}_s$, all variables $\rho_{s,i',i_{\text{rep}},j}$ for $i' \in [l]$ are in $\text{Extend}(\bar{\mathcal{F}}_s)$ (same for $\text{Extend}(\bar{\mathcal{F}}_e)$, $\text{Extend}(\bar{\mathcal{F}}_r)$, $\text{Extend}(\bar{\mathcal{F}}_e)$). Conversely $\text{Collapse}(\mathcal{L}'_s)$ discards the values of any variables that are in $\bar{\mathcal{F}}_r$ but not $\bar{\mathcal{F}}'_r$. Clearly, this does not modify the view of the adversary. This conceptual change will be necessary to reduce to a simpler signing algorithm in the following section.

Lastly, we prove that for any PPT adversary \mathcal{A} against the game described in Hybrid₄ (cf. Figure 10), we can construct an adversary \mathcal{B} against the standard EUF-CMA security of small Raccoon in Figure 7. At a high level a challenger can simulate queries from $\text{KeyGen}_{\mathcal{G}}$ by querying the public key $\bar{\mathbf{t}}$ from the oracle

Adversary	Challenger
$(\text{KeyGen}, \text{Sign}, \text{Verify}, \text{KeyUpdate})$	
$\xleftarrow{\bar{\mathcal{F}}_{\text{KeyGen}}}$	$(\bar{\mathcal{F}}'_s, \bar{\mathcal{F}}'_e) := \bar{\mathcal{F}}'_{\text{KeyGen}} \leftarrow \text{SimIn}_{\text{KeyGen}}(\bar{\mathcal{F}}_{\text{KeyGen}})$ $\bar{\mathcal{F}}'_s = \text{Extend}(\bar{\mathcal{F}}'_s)$ $\bar{\mathcal{F}}'_e = \text{Extend}(\bar{\mathcal{F}}'_e)$
$\xleftarrow{\text{vk}, \mathcal{L}_{\text{KeyGen}}}$	$((\text{sk}, \text{vk}), (\mathcal{L}'_s, \mathcal{L}'_e)) \leftarrow \text{KeyGen}_{\mathcal{G}}((\bar{\mathcal{F}}'_s, \bar{\mathcal{F}}'_e), 1^\lambda)$ $\mathcal{L}_{\text{KeyGen}} \leftarrow \text{SimOut}_{\text{KeyGen}}(\text{Collapse}(\mathcal{L}'_s), \text{Collapse}(\mathcal{L}'_e))$
$\xleftarrow{m^{(i)}, \bar{\mathcal{F}}_{\text{Sign}}^{(i)}}$	$\text{sk} \leftarrow \text{KeyUpdate}(\text{sk})$ $(\bar{\mathcal{F}}'_r, \bar{\mathcal{F}}'_{e'}, \bar{\mathcal{F}}'_{\text{sk}}) := \bar{\mathcal{F}}'^{(i)} \leftarrow \text{SimIn}_{\text{Sign}}(\bar{\mathcal{F}}_{\text{Sign}}^{(i)})$ $\bar{\mathcal{F}}'_r = \text{Extend}(\bar{\mathcal{F}}'_r)$ $\bar{\mathcal{F}}'_{e'} = \text{Extend}(\bar{\mathcal{F}}'_{e'})$
$\xleftarrow{\text{sig}^{(i)}, \mathcal{L}_{\text{Sign}}^{(i)}}$	$(\text{sig}^{(i)}, (\mathcal{L}'_r, \mathcal{L}'_{e'}, \mathcal{L}'_{\text{sk}})) \leftarrow \text{Sign}_{\mathcal{G}}((\bar{\mathcal{F}}'_r, \bar{\mathcal{F}}'_{e'}, \bar{\mathcal{F}}'_{\text{sk}}), \text{sk}, m^{(i)})$ $\mathcal{L}_{\text{Sign}}^{(i)} \leftarrow \text{SimOut}_{\text{Sign}}(\text{Collapse}(\mathcal{L}'_r), \text{Collapse}(\mathcal{L}'_{e'}), \mathcal{L}'_{\text{sk}})$
$\xleftarrow{m^*, \text{sig}^*}$	$b := \text{Verify}(\text{vk}, m^*, \text{sig}^*) \wedge (m^* \notin \{m^{(1)}, \dots, m^{(Q_s)}\}) \wedge$ $\wedge \bar{\mathcal{F}}_{\text{KeyGen}} \leq d - 1 \wedge \forall i \in \{1, \dots, Q_s\}, \bar{\mathcal{F}}_{\text{Sign}}^{(i)} \leq d - 1$

Fig. 10: Hybrid_4 : In this game, for any variable name $\bar{\rho}_{s,i,\text{rep},j}$ the challenger artificially leaks all variables $\rho_{s,i,\text{rep},j'}$ for $j' \in [\ell]$ (and similarly when s is replaced by e, r, e'). He then discards the extra leakage before sending it to the adversary. The view of the adversary is unchanged.

for $\text{KeyGen}_{\text{Small}}$ and artificially sampling additional noises (\tilde{s}, \tilde{e}) as the sum of $d - 1$ small uniforms and outputting the public key $\mathbf{t} := \lfloor \tilde{\mathbf{t}} + \mathbf{A}\tilde{s} + \tilde{\mathbf{e}} \rfloor_{v_t}$ which will be distributed exactly as a public key for $\text{KeyGen}_{\mathcal{G}}$. Similarly, a signature from $\text{Sign}_{\text{Small}}$ can be mapped to a signature for $\text{Sign}_{\mathcal{G}}$ by sampling the appropriate sums of uniform $(\tilde{\mathbf{r}}, \tilde{\mathbf{e}}')$ and setting $\mathbf{w} = \lfloor \tilde{\mathbf{w}} + \mathbf{A}\tilde{\mathbf{r}} + \tilde{\mathbf{e}}' \rfloor_{v_w}$. Finally we show a forgery for $\text{Sign}_{\mathcal{G}}$ can be mapped to a forgery for $\text{Sign}_{\text{Small}}$. The formal proof is given in the full version. This completes the proof.

7.3 MLWE + SelfTargetMSIS \Rightarrow EUF-CMA Security of Small Raccoon

Notations for smooth Rényi divergence. We further define some useful notations to aid the readability. For any $c \in \mathcal{C}$, $\mathbf{s} \in \mathcal{R}_q^\ell$, and $\mathbf{e} \in \mathcal{R}_q^k$, we note $\text{center} := c \cdot \begin{bmatrix} \mathbf{s} \\ \mathbf{e} \end{bmatrix} \in \mathcal{R}_q^{\ell+k}$ and recall $T = d \cdot (\text{rep} - 1) + 1$. We define two distributions: $\mathcal{P} := \text{SU}(u_w, T)^{n(\ell+k)}$ and $\mathcal{Q}(\text{center}) := \text{center} + \mathcal{P}$.

We bound the smooth Rényi divergence of \mathcal{P} and \mathcal{Q} . For any $\alpha = \omega_{\text{asymp}}(1)$ and $\epsilon_{\text{TAIL}}(\text{center}) = \frac{1}{\sqrt{2\pi T}} \left(\frac{\alpha e \|\text{center}\|_2}{2^{u_w \cdot T}} \right)^T$ (see Conjecture 4.1 or Lemma 4.2), we

define ϵ_{TAIL} and $R_\alpha^{\epsilon_{\text{TAIL}}}(\mathcal{P}; \mathcal{Q})$ to be any two values satisfying

$$\Pr \left[\epsilon_{\text{TAIL}} \geq \max_{c \in \mathcal{C}} \epsilon_{\text{TAIL}}(\text{center}) \right] \geq 1 - \text{negl}(\kappa). \quad (10)$$

$$\Pr \left[R_\alpha^{\epsilon_{\text{TAIL}}}(\mathcal{P}; \mathcal{Q}) \geq \max_{c \in \mathcal{C}} R_\alpha^{\epsilon_{\text{TAIL}}(\text{center})}(\mathcal{P}; \mathcal{Q}(\text{center})) \right] \geq 1 - \text{negl}(\kappa). \quad (11)$$

where both probabilities are taken under the randomness of $(\mathbf{s}, \mathbf{e}) \leftarrow \text{RSU}(u_t, T)^\ell \times \text{RSU}(u_t, T)^k$. For efficiency and better parameters, we set ϵ_{TAIL} and $R_\alpha^{\epsilon_{\text{TAIL}}}(\mathcal{P}; \mathcal{Q})$ to be the smallest values satisfying the above inequality. The above parameters we provide is one set of candidate asymptotic parameters.

It remains to prove that small Raccoon in Figure 7 is (standard) EUF-CMA secure. This is established in the following theorem.

Theorem 7.2. *The small Raccoon in Figure 7 is EUF-CMA secure under the $\text{MLWE}_{q,\ell,k,\text{SU}(u_t,T)}$ and $\text{SelfTargetMSIS}_{q,\ell+1,k,C,\beta}$ assumptions.*

Formally, for any adversary \mathcal{A} against the EUF-CMA security game making at most Q_h random oracle queries and Q_s signing queries, and ϵ_{TAIL} and $R_\alpha^{\epsilon_{\text{TAIL}}}(\mathcal{P}; \mathcal{Q})$ satisfying Eqs. (10) and (11), there exists adversaries \mathcal{B} and \mathcal{B}' against the $\text{MLWE}_{q,\ell,k,\text{SU}(u_t,T)}$ and $\text{SelfTargetMSIS}_{q,\ell+1,k,C,\beta}$ problems such that

$$\begin{aligned} \text{Adv}_{\mathcal{A}}^{\text{EUF-CMA}} &\leq 2^{-\kappa} \cdot Q_h \cdot (1 + 2^{-\kappa+1} \cdot Q_s) + Q_s \cdot \epsilon_{\text{TAIL}} \\ &\quad + \left(\text{Adv}_{\mathcal{B}}^{\text{MLWE}} + \text{Adv}_{\mathcal{B}'}^{\text{SelfTargetMSIS}} + Q_s \cdot \epsilon_{\text{TAIL}} \right)^{\frac{\alpha-1}{\alpha}} \cdot \left(R_\alpha^{\epsilon_{\text{TAIL}}}(\mathcal{P}; \mathcal{Q}) \right)^{Q_s}, \end{aligned} \quad (12)$$

where $\text{Time}(\mathcal{A}) \approx \text{Time}(\mathcal{B}) \approx \text{Time}(\mathcal{B}')$.

We now present an overview of the proof which, due to page constraints, is left to the full version. As a first step we replace the hash function by a random oracle which we will program by first sampling $c_{\text{poly}} \leftarrow \mathcal{C}$ and setting the hash function accordingly. Once this is done \mathbf{w} can be defined as a function of c_{poly} rather than $(\mathbf{r}, \mathbf{e}')$, using the equation $\mathbf{w} := \mathbf{A} \cdot \mathbf{z} - c_{\text{poly}} \cdot \mathbf{t} + \mathbf{z}'$ where $\mathbf{z}' := c_{\text{poly}} \cdot \mathbf{e} + \mathbf{e}'$. We now observe that all variables can be computed as deterministic functions of $(\mathbf{z}, \mathbf{z}')$, we thus want to prove that $(\mathbf{z}, \mathbf{z}')$ are independent of (\mathbf{s}, \mathbf{e}) . Using the Smooth-Rényi divergence property of Lemma 4.2 we can bound the divergence between $(\mathbf{z}, \mathbf{z}') = (c_{\text{poly}} \cdot \mathbf{s} + \mathbf{r}, c_{\text{poly}} \cdot \mathbf{e} + \mathbf{e}')$ and $(\mathbf{r}, \mathbf{e}')$ which are sums of uniforms independent of the secret. Finally we can replace the public key with a uniform vector using MLWE, and use the forgery output by the adversary to break MSIS.

8 Concrete Instantiation

Looking at Theorem 7.2, it is clear that the security bottlenecks in Theorem 7.2 are the hardness of MLWE, of SelfTargetMSIS, and the smooth Rényi divergence (ϵ_{TAIL} and $R_\alpha^{\epsilon_{\text{TAIL}}}$). Instantiating Raccoon boils down to an optimization problem where we need to balance the hardness assumptions (MLWE, SelfTargetMSIS), the smooth Rényi divergence and the performance metrics (size of vk and sig).

- Our analysis of MLWE and SelfTargetMSIS is fairly standard. We rely on the lattice estimator [2] for the concrete analysis of MLWE. Following the Dilithium methodology [30, §C.3], we assume that breaking SelfTargetMSIS requires to either (a) break the second-preimage resistance of the hash function, or (b) break an inhomogeneous MSIS instance, for which the best known attack is in [10, §4.2].
- For the smooth Rényi divergence, one could use Lemma 4.2 for a provable bound. However, it is not tight so we opt instead to use Conjecture 4.1.

We refer the reader to the full version of this paper where we provide the relationship between parameters the security/efficiency metrics is in. In addition, we provide example parameters for the NIST security level I.

Table 3: Parameters for Raccoon-128, NIST Post-Quantum security strength category 1. For all Raccoon-128 masking orders, we fix: $\kappa = 128$, $Q_s = 2^{53}$, $q = (2^{24} - 2^{18} + 1) \cdot (2^{25} - 2^{18} + 1)$, $n = 512$, $k = 5$, $\ell = 4$, $\nu_t = 42$, $\nu_w = 44$, $\omega = 19$, $2^{-64}B_2^2 = 14656575897$, $B_\infty = 41954689765971$.

Parameter	Raccoon-128	128-2	128-4	128-8	128-16	128-32
$ \text{sig} $ (bytes)	11524	=	=	=	=	=
$ \text{vk} $ (bytes)	2256	=	=	=	=	=
d	1	2	4	8	16	32
rep	8	4	2	4	2	4
u_t	6	6	6	5	5	4
u_w	41	41	41	40	40	39
$ \text{sk} $ (bytes)	14800	14816	14848	14912	15040	15296

9 Conclusion and Next Steps

We have presented Raccoon, a masking-friendly signature scheme with a formal security proof in the t -probing model based on standard lattice assumptions. We present a few natural extensions of our work:

- **Tighter proof.** The recent Hint-MLWE assumption by Kim et al. [?] seems perfectly suited to study Raccoon, as illustrated by a thresholdized variant of Raccoon [?]. For Raccoon itself, an obstacle to a direct application is that [?] provided security reductions for Gaussian distributions, whereas Raccoon uses sums of uniform distributions.
- **More realistic models.** While the t -probing model is a simple and convenient abstraction of real-world leakage, there exist more realistic models such as the random probing and noisy leakage models. We expect a security analysis in these models to be informative and to raise its own challenges.
- **Real-world assessment.** Since side-channel analysis are grounded in real-world deployment, this work needs to be completed with a study of the concrete leakage of Raccoon when implemented on real-world devices.

References

1. Agrawal, S., Stehlé, D., Yadav, A.: Round-optimal lattice-based threshold signatures, revisited. In: Bojanczyk, M., Merelli, E., Woodruff, D.P. (eds.) ICALP 2022. LIPIcs, vol. 229, pp. 8:1–8:20. Schloss Dagstuhl (Jul 2022). <https://doi.org/10.4230/LIPIcs.ICALP.2022.8>

2. Albrecht, M.R., Player, R., Scott, S.: On the concrete hardness of Learning with Errors. *Journal of Mathematical Cryptology* **9**(3), 169–203 (2015). <https://doi.org/doi:10.1515/jmc-2015-0016>, <https://doi.org/10.1515/jmc-2015-0016>
3. Azouaoui, M., Bronchain, O., Cassiers, G., Hoffmann, C., Kuzovkova, Y., Renes, J., Schneider, T., Schönauer, M., Standaert, F.X., van Vredendaal, C.: Protecting dilithium against leakage revisited sensitivity analysis and improved implementations. *IACR TCHES* **2023**(4), 58–79 (2023). <https://doi.org/10.46586/tches.v2023.i4.58-79>
4. Barthe, G., Belaïd, S., Dupressoir, F., Fouque, P.A., Grégoire, B., Strub, P.Y., Zucchini, R.: Strong non-interference and type-directed higher-order masking. In: Weippl, E.R., Katzenbeisser, S., Kruegel, C., Myers, A.C., Halevi, S. (eds.) *ACM CCS 2016*. pp. 116–129. ACM Press (Oct 2016). <https://doi.org/10.1145/2976749.2978427>
5. Barthe, G., Belaïd, S., Espitau, T., Fouque, P.A., Grégoire, B., Rossi, M., Tibouchi, M.: Masking the GLP lattice-based signature scheme at any order. In: Nielsen, J.B., Rijmen, V. (eds.) *EUROCRYPT 2018, Part II. LNCS*, vol. 10821, pp. 354–384. Springer, Heidelberg (Apr / May 2018). https://doi.org/10.1007/978-3-319-78375-8_12
6. Barthe, G., Belaïd, S., Espitau, T., Fouque, P.A., Rossi, M., Tibouchi, M.: GALACTICS: Gaussian sampling for lattice-based constant-time implementation of cryptographic signatures, revisited. In: Cavallaro, L., Kinder, J., Wang, X., Katz, J. (eds.) *ACM CCS 2019*. pp. 2147–2164. ACM Press (Nov 2019). <https://doi.org/10.1145/3319535.3363223>
7. Battistello, A., Coron, J.S., Prouff, E., Zeitoun, R.: Horizontal side-channel attacks and countermeasures on the ISW masking scheme. In: Gierlichs, B., Poschmann, A.Y. (eds.) *CHES 2016. LNCS*, vol. 9813, pp. 23–39. Springer, Heidelberg (Aug 2016). https://doi.org/10.1007/978-3-662-53140-2_2
8. Berzati, A., Viera, A.C., Chartouny, M., Madec, S., Vergnaud, D., Vigilant, D.: Exploiting intermediate value leakage in dilithium: A template-based approach. *IACR TCHES* **2023**(4), 188–210 (2023). <https://doi.org/10.46586/tches.v2023.i4.188-210>
9. Bronchain, O., Azouaoui, M., ElGhamrawy, M., Renes, J., Schneider, T.: Exploiting small-norm polynomial multiplication with physical attacks: Application to crystals-dilithium. *Cryptology ePrint Archive*, Paper 2023/1545 (2023), <https://eprint.iacr.org/2023/1545>, <https://eprint.iacr.org/2023/1545>
10. Chuengsatiansup, C., Prest, T., Stehlé, D., Wallet, A., Xagawa, K.: ModFalcon: Compact signatures based on module-NTRU lattices. In: Sun, H.M., Shieh, S.P., Gu, G., Ateniese, G. (eds.) *ASIACCS 20*. pp. 853–866. ACM Press (Oct 2020). <https://doi.org/10.1145/3320269.3384758>
11. Coron, J.S.: High-order conversion from Boolean to arithmetic masking. In: Fischer, W., Homma, N. (eds.) *CHES 2017. LNCS*, vol. 10529, pp. 93–114. Springer, Heidelberg (Sep 2017). https://doi.org/10.1007/978-3-319-66787-4_5
12. Coron, J.S., Gérard, F., Trannoy, M., Zeitoun, R.: Improved gadgets for the high-order masking of dilithium. *IACR TCHES* **2023**(4), 110–145 (2023). <https://doi.org/10.46586/tches.v2023.i4.110-145>
13. Coron, J.S., Großschädl, J., Tibouchi, M., Vadnala, P.K.: Conversion from arithmetic to Boolean masking with logarithmic complexity. In: Leander, G. (ed.) *FSE 2015. LNCS*, vol. 9054, pp. 130–149. Springer, Heidelberg (Mar 2015). https://doi.org/10.1007/978-3-662-48116-5_7

14. Coron, J.S., Großschädl, J., Vadnala, P.K.: Secure conversion between Boolean and arithmetic masking of any order. In: Batina, L., Robshaw, M. (eds.) CHES 2014. LNCS, vol. 8731, pp. 188–205. Springer, Heidelberg (Sep 2014). https://doi.org/10.1007/978-3-662-44709-3_11
15. Csiszár, I.: Eine informationstheoretische Ungleichung und ihre Anwendung auf den Beweis der Ergodizität von Markoffschen Ketten. Magyar. Tud. Akad. Mat. Kutató Int. Közl **8**, 85–108 (1963)
16. del Pino, R., Espitau, T., Katsumata, S., Maller, M., Mouhartem, F., Prest, T., Rossi, M., Saarinen, M.: Raccoon. Tech. rep., National Institute of Standards and Technology (2023), available at <https://csrc.nist.gov/Projects/pqc-dig-sig/round-1-additional-signatures>
17. del Pino, R., Prest, T., Rossi, M., Saarinen, M.J.O.: High-order masking of lattice signatures in quasilinear time. In: 2023 IEEE Symposium on Security and Privacy. pp. 1168–1185. IEEE Computer Society Press (May 2023). <https://doi.org/10.1109/SP46215.2023.10179342>
18. Devevey, J., Fawzi, O., Passelègue, A., Stehlé, D.: On rejection sampling in lyubashevsky’s signature scheme. In: Agrawal, S., Lin, D. (eds.) ASIACRYPT 2022, Part IV. LNCS, vol. 13794, pp. 34–64. Springer, Heidelberg (Dec 2022). https://doi.org/10.1007/978-3-031-22972-5_2
19. Duc, A., Faust, S., Standaert, F.X.: Making masking security proofs concrete (or how to evaluate the security of any leaking device), extended version. Journal of Cryptology **32**(4), 1263–1297 (Oct 2019). <https://doi.org/10.1007/s00145-018-9277-0>
20. Esgin, M., Espitau, T., Niot, G., Prest, T., Sakzad, A., Steinfeld, R.: Plover: Masking-friendly hash-and-sign lattice signatures. In: EUROCRYPT (2024), <https://tprest.github.io/pdf/pub/plover.pdf>
21. Espitau, T., Fouque, P.A., Gérard, F., Rossi, M., Takahashi, A., Tibouchi, M., Wallet, A., Yu, Y.: Mitaka: A simpler, parallelizable, maskable variant of falcon. In: Dunkelman, O., Dziembowski, S. (eds.) EUROCRYPT 2022, Part III. LNCS, vol. 13277, pp. 222–253. Springer, Heidelberg (May / Jun 2022). https://doi.org/10.1007/978-3-031-07082-2_9
22. Fournaris, A.P., Dimopoulos, C., Koufopavlou, O.G.: Profiling Dilithium Digital Signature Traces for Correlation Differential Side Channel Attacks. In: Orailoglu, A., Jung, M., Reichenbach, M. (eds.) Embedded Computer Systems: Architectures, Modeling, and Simulation - 20th International Conference, SAMOS 2020, Samos, Greece, July 5-9, 2020, Proceedings. Lecture Notes in Computer Science, vol. 12471, pp. 281–294. Springer (2020). https://doi.org/10.1007/978-3-030-60939-9_19, https://doi.org/10.1007/978-3-030-60939-9_19
23. Gérard, F., Rossi, M.: An efficient and provable masked implementation of qtesla. In: Belaïd, S., Güneysu, T. (eds.) Smart Card Research and Advanced Applications - 18th International Conference, CARDIS 2019, Prague, Czech Republic, November 11-13, 2019, Revised Selected Papers. Lecture Notes in Computer Science, vol. 11833, pp. 74–91. Springer (2019). https://doi.org/10.1007/978-3-030-42068-0_5, https://doi.org/10.1007/978-3-030-42068-0_5
24. Goudarzi, D., Prest, T., Rivain, M., Vergnaud, D.: Probing security through input-output separation and revisited quasilinear masking. IACR TCHES **2021**(3), 599–640 (2021). <https://doi.org/10.46586/tches.v2021.i3.599-640>, <https://tches.iacr.org/index.php/TCHES/article/view/8987>
25. Howe, J., Prest, T., Ricosset, T., Rossi, M.: Isochronous gaussian sampling: From inception to implementation. In: Ding, J., Tillich, J.P. (eds.) Post-Quantum Cryptography (2023), pp. 1–24. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-37080-9_1

- tography - 11th International Conference, PQCrypto 2020. pp. 53–71. Springer, Heidelberg (2020). https://doi.org/10.1007/978-3-030-44223-1_5
26. Hutter, M., Tunstall, M.: Constant-time higher-order Boolean-to-arithmetic masking. *Journal of Cryptographic Engineering* **9**(2), 173–184 (Jun 2019). <https://doi.org/10.1007/s13389-018-0191-z>
 27. Ishai, Y., Sahai, A., Wagner, D.: Private circuits: Securing hardware against probing attacks. In: Boneh, D. (ed.) *CRYPTO 2003*. LNCS, vol. 2729, pp. 463–481. Springer, Heidelberg (Aug 2003). https://doi.org/10.1007/978-3-540-45146-4_27
 28. Ito, A., Ueno, R., Homma, N.: On the success rate of side-channel attacks on masked implementations: Information-theoretical bounds and their practical usage. In: Yin, H., Stavrou, A., Cremers, C., Shi, E. (eds.) *ACM CCS 2022*. pp. 1521–1535. ACM Press (Nov 2022). <https://doi.org/10.1145/3548606.3560579>
 29. Karabulut, E., Alkim, E., Aysu, A.: Single-Trace Side-Channel Attacks on ω -Small Polynomial Sampling: With Applications to NTRU, NTRU Prime, and CRYSTALS-DILITHIUM. In: *IEEE International Symposium on Hardware Oriented Security and Trust, HOST 2021, Tysons Corner, VA, USA, December 12–15, 2021*. pp. 35–45. IEEE (2021). <https://doi.org/10.1109/HOST49136.2021.9702284>, <https://doi.org/10.1109/HOST49136.2021.9702284>
 30. Lyubashevsky, V., Ducas, L., Kiltz, E., Lepoint, T., Schwabe, P., Seiler, G., Stehlé, D., Bai, S.: CRYSTALS-DILITHIUM. Tech. rep., National Institute of Standards and Technology (2022), available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022>
 31. Marzougui, S., Ulitzsch, V., Tibouchi, M., Seifert, J.P.: Profiling side-channel attacks on Dilithium: A small bit-fiddling leak breaks it all. *Cryptology ePrint Archive, Report 2022/106* (2022), <https://eprint.iacr.org/2022/106>
 32. Mathieu-Mahias, A.: *Securisation of implementations of cryptographic algorithms in the context of embedded systems. (Sécurisation des implémentations d’algorithmes cryptographiques pour les systèmes embarqués)*. Ph.D. thesis, University of Paris-Saclay, France (2021), <https://tel.archives-ouvertes.fr/tel-03537322>
 33. Migliore, V., Gérard, B., Tibouchi, M., Fouque, P.A.: Masking Dilithium - efficient implementation and side-channel evaluation. In: Deng, R.H., Gauthier-Umaña, V., Ochoa, M., Yung, M. (eds.) *ACNS 19*. LNCS, vol. 11464, pp. 344–362. Springer, Heidelberg (Jun 2019). https://doi.org/10.1007/978-3-030-21568-2_17
 34. Naehrig, M., Alkim, E., Bos, J., Ducas, L., Easterbrook, K., LaMacchia, B., Longa, P., Mironov, I., Nikolaenko, V., Peikert, C., Raghunathan, A., Stebila, D.: FrodoKEM. Tech. rep., National Institute of Standards and Technology (2017), available at <https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization/round-1-submissions>
 35. Prest, T.: A key-recovery attack against mitaka in the t -probing model. In: Boldyreva, A., Kolesnikov, V. (eds.) *PKC 2023, Part I*. LNCS, vol. 13940, pp. 205–220. Springer, Heidelberg (May 2023). https://doi.org/10.1007/978-3-031-31368-4_8
 36. Prest, T., Fouque, P.A., Hoffstein, J., Kirchner, P., Lyubashevsky, V., Pornin, T., Ricosset, T., Seiler, G., Whyte, W., Zhang, Z.: FALCON. Tech. rep., National Institute of Standards and Technology (2022), available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022>
 37. Steffen, H.M., Land, G., Kogelheide, L.J., Güneysu, T.: Breaking and protecting the crystal: Side-channel analysis of dilithium in hardware. In: *PQCrypto. Lecture Notes in Computer Science*, vol. 14154, pp. 688–711. Springer (2023)

38. Wang, R., Ngo, K., Gärtner, J., Dubrova, E.: Single-trace side-channel attacks on crystals-dilithium: Myth or reality? Cryptology ePrint Archive, Paper 2023/1931 (2023), <https://eprint.iacr.org/2023/1931>, <https://eprint.iacr.org/2023/1931>