Improved Side-Channel Analysis of Finite-Field Multiplication

Sonia Belaïd[1(✉)], Jean-Sébastien Coron[2], Pierre-Alain Fouque[3], Benoît Gérard[4], Jean-Gabriel Kammerer[5], and Emmanuel Prouff[6]

[1] École Normale Supérieure and Thales Communications and Security, Gennevilliers, France
sonia.belaid@live.fr
[2] University of Luxembourg, Walferdange, Luxembourg
[3] Université de Rennes 1 and IRISA, Rennes, France
[4] DGA/MI and IRISA, Rennes, France
[5] DGA/MI and IRMAR, Rennes, France
[6] ANSSI, Paris, France

Abstract. A side-channel analysis of multiplication in $\mathsf{GF}(2^{128})$ has recently been published by Belaïd, Fouque and Gérard at Asiacrypt 2014, with an application to AES-GCM. Using the least significant bit of the Hamming weight of the multiplication result, the authors have shown how to recover the secret multiplier efficiently. However such least significant bit is very sensitive to noise measurement; this implies that, without averaging, their attack can only work for high signal-to-noise ratios ($\mathsf{SNR} > 128$). In this paper we describe a new side-channel attack against the multiplication in $\mathsf{GF}(2^{128})$ that uses the most significant bits of the Hamming weight. We show that much higher values of noise can be then tolerated. For instance with an SNR equal to 8, the key can be recovered using 2^{20} consumption traces with time and memory complexities respectively equal to $2^{51.68}$ and 2^{36}. We moreover show that the new method can be extended to attack the fresh re-keying countermeasure proposed by Medwed, Standaert, Großschädl and Regazzoni at Africacrypt 2010.

Keywords: Side-channel analysis · Galois Field Multiplication · LPN problem

1 Introduction

Side-Channel Attacks. The cornerstone of side-channel analysis (SCA for short) is that information about some key-dependent variable x leaks through *e.g.* the power consumption or the electromagnetic information of the device manipulating x. A side-channel attack classically follows a *divide-and-conquer* approach and the secret is recovered by exhaustively testing the likelihood of every possible value for every secret piece. This *modus operandi* implicitly assumes that x depends on a short portion of the secret (for example only 8 bits if x

T. Güneysu and H. Handschuh (Eds.): CHES 2015, LNCS 9293, pp. 395–415, 2015.
DOI: 10.1007/978-3-662-48324-4_20

corresponds to the output of the AES sbox). It is particularly suited to the context of software implementations where the processing is sequentially split into operations on data whose size depends on the device architecture (*e.g.* 8 bit or even 32 bit for smart cards).

Side-Channel Analysis of Finite-Field Multiplication. At Asiacrypt 2014 [BFG14], Belaïd, Fouque and Gérard consider an attack scenario dedicated to hardware implementations where many operations are performed simultaneously. Following previous works as [MSGR10, MSJ12], they assume that when performing a multiplication $\mathbf{a}\cdot\mathbf{k}$ over $\mathsf{GF}(2^n)$ for some known $\mathbf{a}$, only the Hamming weight of the result $\mathbf{a}\cdot\mathbf{k} \in \mathsf{GF}(2^n)$ is leaking, with some noise; the goal is to recover the secret multiplier $\mathbf{k}$. Formally, after denoting by $\mathcal{N}(0,\sigma)$ the Gaussian distribution with null mean and standard deviation σ and by HW the Hamming weight over $\mathsf{GF}(2^n)$, for a given basis of $\mathsf{GF}(2^n)$, the SCA then amounts to solve the following problem:

Definition 1 (Hidden Multiplier Problem). *Let* $\mathbf{k} \leftarrow \mathsf{GF}(2^n)$. *Let* $\ell \in \mathbb{N}$. *Given a sequence* $(\mathbf{a}_i, \mathcal{L}_i)_{1\leq i\leq \ell}$ *where* $\mathbf{a}_i \leftarrow \mathsf{GF}(2^n)$ *and* $\mathcal{L}_i = \mathsf{HW}(\mathbf{a}_i \cdot \mathbf{k}) + \varepsilon_i$ *where* $\varepsilon_i \leftarrow \mathcal{N}(0,\sigma)$, *recover* $\mathbf{k}$.

The Belaïd-Fouque-Gérard Attack and the LPN Problem. As noted in [BFG14], for $\sigma = 0$ (no noise) the above problem is easy to solve. Namely the least significant bit of the Hamming weight of x is the xor of the bits of x. Hence for known $\mathbf{a}_i$ the least significant bit of $\mathsf{HW}(\mathbf{a}_i \cdot \mathbf{k})$ is a linear function of the bits of the secret $\mathbf{k}$. Therefore every Hamming weight gives a linear equation over the n bits of $\mathbf{k}$ and, if the system of equations has rank n (which happens with good probability), the secret $\mathbf{k}$ can be recovered by solving a linear system. However such least significant bit is very sensitive to the observation noise ε_i. Even for relatively high signal-to-noise ratios (*i.e.,* low σ), this induces a significant error probability for the linear equations. This is all the more damageable that a device is never exactly leaking the Hamming weight of manipulated data, and a modeling (aka epistemic) error therefore adds to the observation noise. The problem of solving a system of noisy linear equations over $\mathsf{GF}(2)$ is known as the Learning Parity with Noise (LPN) problem. New algorithms for solving LPN have recently been proposed [GJL14, BTV15]. The previous best method to solve the LPN problem was the Fouque-Levieil algorithm from [LF06], which is a variant of the algorithm BKW proposed by Blum, Kalai and Wasserman in [BKW00]. According to [BFG14] the Fouque-Levieil algorithm can solve the LPN for $n = 128$ bits with error probability $p = 0.31$ (corresponding to $\mathsf{SNR} = 128$) with 2^{48} acquisitions and 2^{50} complexity (it becomes 2^{334} when $\mathsf{SNR} = 8$). Therefore the Belaïd-Fouque-Gérard (BFG for short) algorithm for solving the Hidden Multiplier Problem is quite efficient for relatively high signal-to-noise ratios ($\mathsf{SNR} > 128$); however it becomes prohibitively inefficient for smaller values (*e.g.*, larger values of σ).

Our New Attack. In this paper we describe a new algorithm for solving the Hidden Multiplier Problem, in which we use several most significant bits of the Hamming weight instead of the single least significant bit; we show that much smaller values of SNR can then be tolerated (SNR $\simeq$ 8), which increases the practicability of the attack. Our technique works as follows. We only keep the observations with small Hamming weight or high Hamming weight. Namely if $\mathsf{HW}(\mathbf{a}_i \cdot \mathbf{k})$ is close to 0, this means that most of the bits of $\mathbf{a}_i \cdot \mathbf{k}$ are equal to 0. This can be written as a system of n equations over the bits of $\mathbf{k}$, all equal to 0, where some of the equations are erroneous. Similarly if the Hamming weight is close to n, we can assume that all n equations are equal to 1, and we obtain again a set of n noisy equations. Hence in both cases we obtain an instance of the LPN problem. For example, if we only keep observations with Hamming weight less than $n/4$ or greater than $3n/4$, we obtain a set of noisy equations with error probability less than 1/4.

To solve the LPN problem we will use BKW style algorithms [BKW00]. The main drawback of these algorithms is the huge samples requirement that makes them unpractical for side-channel attacks. In this paper we use some improvements to reduce the query complexity using Shamir-Schroeppel [SS79] or the variant proposed by Howgrave-Graham and Joux in [HGJ10]. We also take advantage of secret-error switching lemma [Kir11, ACPS09] to further reduce the time complexity.

Since our attack is based on filtering for abnormally low or high Hamming weights, it is much less sensitive to noise in Hamming weight measurement than the BFG attack, which relies on the least significant bit of the Hamming weight. Namely even for small SNR (*i.e.,* close to 8), our filtering remains essentially correct, whereas the information from the least significant bit of the Hamming weight is buried in noise and becomes useless. However, for high SNR, our attack requires a larger amount of observations. Therefore in the latter contexts, the BFG attack stays better.

We also describe an attack when the messages $\mathbf{a}_i$ can be chosen. In that case, the attack becomes much more efficient. We also attack a fresh re-keying scheme proposed in [MSGR10] to defeat side-channel cryptanalysis. Whereas the latter scheme is not vulnerable to the technique used in [BFG14], we demonstrate that our attack enables to recover the secret key very efficiently.

Organization of the Paper. In Sect. 2, we recall the field multiplication for the AES-GCM, the leakage model, the LPN problem and the BKW algorithm. Then, we present our new attack in Sect. 3 and the new algorithmic techniques to reduce the number of queries. In Sect. 4 we describe a new chosen message attack and in Sect. 5 our attack on the fresh re-keying scheme. Finally, in Sect. 6 we present the result of our practical experiments.

2 Preliminaries

2.1 Galois Field Multiplication

For any positive integer n, the finite field of 2^n elements is denoted by $\mathsf{GF}(2^n)$ and the n-dimensional vector space over $\mathsf{GF}(2)$ is denoted by $\mathsf{GF}(2)^n$. Choosing a

basis of $\mathsf{GF}(2^n)$ over $\mathsf{GF}(2)$ enables to represent elements of $\mathsf{GF}(2^n)$ as elements of $\mathsf{GF}(2)^n$ and *vice versa*. In the following, we assume that the same basis is always used to represent elements of $\mathsf{GF}(2^n)$ over $\mathsf{GF}(2)$.

This paper analyses the multiplication in the field $\mathsf{GF}(2^n)$, with a particular focus on $n = 128$, with the representation $\mathsf{GF}(2)[x]/(x^{128} + x^7 + x^2 + x + 1)$ which is used in the AES-GCM protocol. If $\mathbf{a} = (\mathrm{a}_0, \mathrm{a}_1, \cdots, \mathrm{a}_{127})$ and $\mathbf{k} = (\mathrm{k}_0, \mathrm{k}_1, \cdots, \mathrm{k}_{127})$ are two elements of $\mathsf{GF}(2^{128})$ viewed as 128-bit vectors, the multiplication $\mathbf{a} \cdot \mathbf{k}$ can be represented by a matrix/vector product in the following way:

$$\begin{pmatrix} \mathrm{a}_0 & \mathrm{a}_{127} & \cdots & \mathrm{a}_1 \oplus \mathrm{a}_{127} \oplus \mathrm{a}_{126} \\ \mathrm{a}_1 & \mathrm{a}_0 \oplus \mathrm{a}_{127} & \cdots & \mathrm{a}_2 \oplus \mathrm{a}_{123} \oplus \mathrm{a}_1 \oplus \mathrm{a}_{127} \oplus \mathrm{a}_{122} \\ \vdots & \vdots & \ddots & \vdots \\ \mathrm{a}_{127} & \mathrm{a}_{126} & \cdots & \mathrm{a}_0 \oplus \mathrm{a}_{127} \oplus \mathrm{a}_{126} \oplus \mathrm{a}_{121} \end{pmatrix} \cdot \begin{pmatrix} \mathrm{k}_0 \\ \mathrm{k}_1 \\ \vdots \\ \mathrm{k}_{127} \end{pmatrix} = \begin{pmatrix} \mathrm{z}_0 \\ \mathrm{z}_1 \\ \vdots \\ \mathrm{z}_{127} \end{pmatrix}, \quad (1)$$

where the product $\cdot$ is processed over $\mathsf{GF}(2)$.

2.2 Probabilities

In this paper we shall use an upper-case letter, *e.g.* X, to denote a random variable, while the lower-case letter, x, shall denote a value taken by X. The probability of an event ev is denoted by $\Pr(ev)$. The *mean* and the *variance* of a random variable X are respectively denoted by $\mathbb{E}(X)$ and $\mathrm{Var}(X)$ (the *standard deviation* of X is the square root of the variance). A *continuous* random variable X with mean μ and variance σ^2 is said to follow a *normal* (Gaussian) distribution, denoted by $X \sim \mathcal{N}(\mu, \sigma)$, if, $\forall\ x \in \mathbb{R}$, $\Pr[X \leqslant x] = \int_{-\infty}^{x} \phi_{\mu,\sigma}(x)$ where $\phi_{\mu,\sigma}$ is the normal *probability distribution function* (pdf) defined by

$$\phi_{\mu,\sigma}(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}.$$

The other distributions used in this paper are the *uniform distribution* over $\mathsf{GF}(2)^n$, denoted by $\mathcal{U}(\mathsf{GF}(2)^n)$, and the Bernoulli distribution $\mathsf{Ber}(p)$ over $\mathsf{GF}(2)$, with $\Pr[X = 1] = p$, for some $p \in [0, 1]$. We shall also use the Binomial distribution which is defined over $\{0, \ldots, n\}$ by $\Pr[X = j] = \binom{n}{j} p^j (1-p)^{n-j}$ and is denoted by $\mathrm{B}(n, p)$.

2.3 Leakage Model

A common assumption is to consider that a processing on an embedded device leaks a noisy observation of the Hamming weight of the manipulated values. Namely, for such manipulated value $\mathbf{z} \in \mathsf{GF}(2)^n$, it is assumed that the adversary obtains the following observation $\mathcal{L}(\mathbf{z})$:

$$\mathcal{L}(\mathbf{z}) = \mathsf{HW}(\mathbf{z}) + \varepsilon, \quad (2)$$

with an independent noise ε satisfying $\varepsilon \sim \mathcal{N}(0, \sigma)$. In practice, each bit of $\mathbf{z}$ can leak differently. Instead of the Hamming weight, the deterministic part

of the observation can hence be modeled as a multivariate polynomial in the bits of $\mathbf{z}$, where the coefficients are taken in $\mathbb{R}$, see [SLP05, RKSF11, DDP13]. For most of current microprocessor architectures, the latter polynomial is well approximated by a linear combination of the bits of $\mathbf{z}$, leading to generalize (2) with $\mathcal{L}(\mathbf{z}) = \sum_{i=0}^{n-1} \beta_i z_i + \varepsilon$. For simplicity, we will describe our attack under the noisy Hamming Weight leakage model given by (2) but in Sect. 6, we will show that it also works for such generic leakage model.

In the rest of the paper, the level of noise in the observations is quantified with the *signal-to-noise ratio* (SNR for short), that we define as the ratio between the signal variance and the noise variance. This value, which equals $n/(4\sigma^2)$ under Assumption (2), is a useful notion to compare different contexts where the variances of both the signal and the noise are different (*e.g.* with different devices).

As in [BFG14], the main purpose of our attack is to show that the key $\mathbf{k}$ can be recovered with only the observations $\mathcal{L}(\mathbf{k} \cdot \mathbf{a_i})$ for many known $\mathbf{a_i}$'s. Thus, we assume that the attacker has no access to the internal leakage of the field multiplication $\mathbf{k} \cdot \mathbf{a_i}$ and that the n-bit results are stored in n-bit registers, which is the worst case to attack.

2.4 Learning Parities with Noise

As briefly explained in the introduction the problem of recovering a secret $\mathbf{k}$ from noisy observations of $\mathsf{HW}(\mathbf{a} \cdot \mathbf{k})$ relates to the well known LPN problem.

Definition 2 (Learning Parity with Noise (LPN) Problem). *Let* $\mathbf{k} \in \mathsf{GF}(2)^n$ *and* $p \in (0, 1/2)$. *Given a family of* ν *values* $(\mathbf{a}_i)_{0 \leqslant i < \nu}$ *in* $\mathsf{GF}(2)^n$ *and the family of corresponding observations* $(b_i = \langle \mathbf{a}_i, \mathbf{k} \rangle + e_i)_{0 \leqslant i < \nu}$, *where* $\langle \cdot, \cdot \rangle$ *denotes the scalar product* $\in \mathsf{GF}(2)^n$ *and where the* $\mathbf{a}_i$ *are drawn uniformly in* $\mathsf{GF}(2^n)$ *and the* e_i *are generated according to Bernoulli's distribution* $\mathsf{Ber}(p)$ *with parameter* p, *recover* $\mathbf{k}$.

We denote by $\mathsf{LPN}(n, \nu, p)$ an instance of the LPN problem with parameters (n, ν, p). In this paper, the noisy equations $\langle \mathbf{a}_i, \mathbf{k} \rangle + e_i$ will come from the noisy observations of a device performing field (or ring) multiplications in the form $\mathbf{z} = \mathbf{a} \cdot \mathbf{k}$ in $\mathsf{GF}(2^n)$.

2.5 The BKW Algorithm and Its Variants

Blum *et al.* described in [BKW00] a subexponential algorithm for solving the LPN problem: it performs a clever Gaussian elimination using a small number of linear combinations, which reduces the dimension of the problem. Then, Levieil and Fouque proposed a practical improvement in [LF06] for the second phase of the algorithm and Kirchner [Kir11] proposed to switch secret and error [ACPS09] to further improve the method. Later Arora and Ge [AG11] proposed an algebraic approach for specifically structured noise. Recently Guo *et al.* proposed to use error-correcting codes [GJL14].

The BKW Algorithm. Given as input $b_i = \langle \mathbf{a}_i, \mathbf{k} \rangle + e_i$ for known $\mathbf{a}_i$'s, the goal of the BKW algorithm is to find linear combinations of the $\mathbf{a}_i$'s with ℓ terms such that:

$$\mathbf{a}_{i_1} \oplus \cdots \oplus \mathbf{a}_{i_\ell} = \mathbf{u}_j, \tag{3}$$

where $(\mathbf{u}_j)_{1 \leq j < n}$ is the canonical basis, that is $\mathbf{u}_j$ has its j^{th} coordinate equal to 1 and the other coordinates are 0. Then one gets:

$$\langle \mathbf{u}_j, \mathbf{k} \rangle = \mathrm{k}_j = \bigoplus_{r=1}^{\ell} b_{i_r} \oplus \bigoplus_{r=1}^{\ell} e_{i_r}.$$

It is not difficult to evaluate the new bias of the linear combination of equations using the Piling-Up lemma. Letting $\delta = 1 - 2p$, for ℓ variables $e_1, \ldots, e_\ell$ such that $\Pr[e_i = 1] = p = (1 - \delta)/2$, we have $\Pr[e_1 \oplus \cdots \oplus e_\ell = 0] = \frac{1+\delta^\ell}{2}$. This shows that if we sum ℓ error terms e_i with $\Pr[e_i = 1] = (1 - \delta)/2$, the resulting error term e is such that $\Pr[e = 1] = (1 - \delta')/2$ with $\delta' = \delta^\ell$. If ℓ is not too large, then the bias of the error term $\bigoplus_{r=1}^{\ell} e_{i_r}$ is also not too large and with enough such equations and a majority vote one can recover the j^{th} coordinate of $\mathbf{k}$.

Finding linear combinations. To find linear combinations satisfying (3), we first split the $\mathbf{a}_i$'s into a blocks of b bits, where $n = a \cdot b$ (*e.g.* for $n = 128$ we can take $a = 8$ and $b = 16$). Initially we have ν vectors $\mathbf{a}_i$. Consider the rightmost b bits of each $\mathbf{a}_i$, and sort the $\mathbf{a}_i$'s into 2^b classes according to this value. We xor all elements of each class with a single one element of it, and we discard this element. Hence we get at least $\nu - 2^b$ new vectors $\mathbf{a}_i^{(1)}$, whose rightmost b bits are zero; these $\mathbf{a}_i^{(1)}$ are the xor of 2 initial vectors $\mathbf{a}_i$. One can then proceed recursively. For the next block of b bits we get at least $\nu - 2 \cdot 2^b$ vectors $\mathbf{a}_i^{(2)}$ whose rightmost $2b$ bits are zero; they are the xor of 4 initial vectors $\mathbf{a}_i$. Stopping at the last-but-one block, we get at least $\nu - (a-1) \cdot 2^b$ vectors, for which only the first b-bit block is possibly non-zero, and which are the xor of 2^{a-1} initial vectors $\mathbf{a}_i$. Among these $\nu - (a-1) \cdot 2^b$ vectors, we select the ones equal to the basis vectors $\mathbf{u}_j$ and we perform a majority vote. With the xor of $\ell = 2^{a-1}$ vectors, the bias is $(1-2p)^{2^{a-1}}$. Therefore for the majority vote we need roughly $c/(1-2p)^{2^{a-1}}$ such vectors, for some logarithmic factor c [BKW00]. A variant of BKW algorithm is described by Levieil and Fouque in [LF06]: it finds linear combinations similarly, however at the end, it uses a Walsh Transform to recover the last b bits of $\mathbf{k}$ at once.

3 Our New Attack

In this section, we describe our new side-channel attack on the result of the multiplication in $\mathsf{GF}(2^n)$, which benefits from being weakly impacted by the observation noise. As in [BFG14], we aim at recovering the n-bit secret key $\mathbf{k}$ from a sequence of t queries $(\mathbf{a}_i, \mathsf{HW}(\mathbf{k} \cdot \mathbf{a}_i) + \varepsilon_i)_{0 \leq i < t}$ where the $\mathbf{a}_i$ are drawn uniformly in $\mathsf{GF}(2^n)$ and the ε_i are drawn from the Gaussian distribution $\mathcal{N}(0, \sigma)$.

3.1 Overview

The cornerstone of the attack is to filter the collected measurements to keep only the lowest and the highest Hamming weights. Then we assume that for each low (resp. high) Hamming weight, the multiplication result is exactly n bits of zeros (resp. ones). As a consequence, each filtered observation of $\mathbf{z}_i = \mathbf{a}_i \cdot \mathbf{k}$ gives n equations each with some error probability p. In our context, the equations correspond to the row-by-column scalar products in (1) and the binary error associated to the ith equation is denoted by e_i, with $\Pr[e_i = 1] = p$. Therefore given t messages and corresponding measurements, we get an instance of the $\mathsf{LPN}(n, n \cdot t, p)$ problem that we can solve using techniques described in Sect. 3.3. To correctly scale the latter techniques, we need to know the error probability p with good precision. In the next section we show how to compute p from the filtering threshold and the measurement noise σ in (2).

3.2 Filtering

We describe here how we filter the lowest and highest leakage and we compute the error probabilities of our final set of equations. In order to catch the extreme Hamming weight values of the multiplication results, we choose a threshold real value λ and we filter all the observations below $n/2 - \lambda s$ and above $n/2 + \lambda s$, with $s = \sqrt{n}/2$ the standard deviation of the leakage deterministic part (here the Hamming weight). In the first case, we assume that all the bits of the multiplication result are zeros and in the second case we assume that they are all set to one. In both cases, we get n linear equations on the key bits, each having the same error probability p.

We first compute the proportion of filtered acquisitions before focusing on the error probability p. Let $\mathbf{z} = \mathbf{a} \cdot \mathbf{k}$ be the result of a finite field multiplication; since $\mathbf{z} \sim \mathcal{U}(\mathsf{GF}(2)^n)$, we deduce $\mathsf{HW}(\mathbf{z}) \sim \mathrm{B}(n, 1/2)$. Moreover since $\mathcal{L}(\mathbf{z}) = \mathsf{HW}(\mathbf{z}) + \varepsilon$, with $\varepsilon \sim \mathcal{N}(0, \sigma)$, we obtain that the pdf h of $\mathcal{L}(\mathbf{z})$ is defined over $\mathbb{R}$ by:

$$h(x) = 2^{-n} \sum_{y=0}^{n} \binom{n}{y} \phi_{y,\sigma}(x).$$

Since our filtering rejects the observations with leakage $\mathcal{L}(\mathbf{z})$ between $n/2 - \lambda s$ and $n/2 + \lambda s$ for some parameter λ, the proportion of filtered acquisition $F(\lambda)$ is then:

$$\forall \lambda \in \mathbb{R}, \quad F(\lambda) = 1 - 2^{-n} \sum_{y=0}^{n} \binom{n}{y} \int_{n/2-\lambda s}^{n/2+\lambda s} \phi_{y,\sigma}(t) dt. \tag{4}$$

After filtering, our attack consists in assuming that the n bits of $\mathbf{z}$ are all zeros if $\mathcal{L}(\mathbf{z}) < n/2 - \lambda s$, and are all ones if $\mathcal{L}(\mathbf{z}) > n/2 + \lambda s$. Therefore in the first case out of the n equations, $\mathsf{HW}(\mathbf{z})$ equations are erroneous, whereas in the second case $n - \mathsf{HW}(\mathbf{z})$ equations are erroneous. In the first case, this corresponds to an error probability $\mathsf{HW}(\mathbf{z})/n$, while in the second case this corresponds to an

error probability $1 - \mathsf{HW}(\mathbf{z})/n$. On average over filtered observations, we obtain an error probability:

$$p(\lambda) = \frac{1}{F(\lambda)} \sum_{y=0}^{n} \frac{\binom{n}{y}}{2^n} \left(\frac{y}{n} \int_{-\infty}^{n/2-\lambda s} \phi_{y,\sigma}(t)dt + \left(1 - \frac{y}{n}\right) \int_{n/2+\lambda s}^{+\infty} \phi_{y,\sigma}(t)dt \right).$$

This error probability $p(\lambda)$ (or p for short) is a crucial parameter as it gives the error probability in the LPN problem. Our goal is to minimize p in order to minimize the complexity of solving the LPN problem. This can be done by increasing the filtering threshold λ; however a larger λ implies that a larger number of observations must be obtained initially. Therefore a tradeoff must be found between the error probability p in the LPN problem and the proportion $F(\lambda)$ of filtered observations.

The main advantage of our attack is that this error probability p is quite insensitive to the noise σ in the observations, as illustrated in Table 1. For $n = 128$ and for various values of σ, we provide the corresponding filtering threshold λ that leads to a filtering probability $F(\lambda)$, expressed with $\log_2 1/F(\lambda)$; we then give the corresponding error probability p. For example, for $\mathsf{SNR} = 128$, with $\lambda = 6.00$ we get a filtering probability $F(\lambda) = 2^{-30}$, which means that on average 2^{30} observations are required to get $n = 128$ equations for the LPN problem; in that case the error probability for the LPN problem is $p = 0.23$. We see that this error probability does not grow too fast as SNR decreases, as we get $p = 0.25$ for $\mathsf{SNR} = 8$ and $p = 0.34$ for $\mathsf{SNR} = 0.5$.

Study in the General Case. For completeness, we exhibit hereafter the expressions of the probabilities $F(\lambda)$ and $p(\lambda)$ when the leakage satisfies (2) for another function than $\mathsf{HW}(\cdot)$. If we relax the Hamming weight assumption but still assume that the noise is independent, additive and Gaussian, we get the following natural generalization of (2):

$$\mathcal{L}(\mathbf{z}) = \varphi(\mathbf{z}) + \varepsilon,$$

where $\varphi(\mathbf{z}) \doteq \mathbb{E}\left(\mathcal{L}(Z) \mid Z = \mathbf{z}\right)$ and $\varepsilon \sim \mathcal{N}(0, \sigma)$. This leads to the following generalization of (4):

$$\forall \lambda \in \mathbb{R}, \quad F(\lambda) = 1 - \sum_{y \in \mathrm{Im}(\varphi)} \mathbb{P}\left(\varphi(Z) = y\right) \int_{-\lambda s}^{\lambda s} \phi_{y,\sigma}(t + \mu)dt,$$

Table 1. Error probability p and λ w.r.t. the filtering proportion $F(\lambda)$ and the SNR

$\log_2(1/F(\lambda))$	30	25	20	15	10	5	30	25	20	15	10	5
	$\mathsf{SNR} = 128, \sigma = 0.5$						$\mathsf{SNR} = 2, \sigma = 4$					
λ	6.00	5.46	4.85	4.15	3.29	2.16	7.42	6.73	5.97	5.09	4.03	2.64
p	0.23	0.25	0.28	0.31	0.34	0.39	0.28	0.30	0.32	0.34	0.37	0.41
	$\mathsf{SNR} = 8, \sigma = 2$						$\mathsf{SNR} = 0.5, \sigma = 8$					
λ	6.37	5.79	5.14	4.39	3.48	2.28	10.57	9.58	8.48	7.21	5.71	3.73
p	0.25	0.27	0.29	0.32	0.35	0.40	0.34	0.36	0.37	0.39	0.41	0.44

where μ and s respectively denote the mean and the standard deviation of $\varphi(Z)$. Analogously, we get:

$$p(\lambda) = \frac{1}{F(\lambda)} \sum_{y=0}^{n} \frac{\binom{n}{y}}{2^n} \left(\frac{y}{n} \int_{-\infty}^{\lambda s} g_y(t+\mu)dt + \left(1 - \frac{y}{n}\right) \int_{\lambda s}^{+\infty} g_y(t+\mu)dt \right),$$

where for every y, the pdf $g_{\mathcal{L}|\mathsf{HW}(Z)=y}$ is defined by:

$$g_y(\ell) = \binom{n}{y}^{-1} \sum_{z \in \mathsf{HW}^{-1}(y)} \phi_{\varphi(\mathbf{z}),\sigma}(\ell).$$

In the case $\varphi = \mathsf{HW}$ (*i.e.,* when the device leaks perfectly in the Hamming weight model), it can be checked that g_y is simply the pdf of $\mathcal{N}(\mathsf{HW}(y), \sigma)$, otherwise it is a Gaussian mixture. In Sect. 6, we will approximate it by a Gaussian pdf with mean $\mathbb{E}(\mathcal{L}(Z) \mid \mathsf{HW}(Z) = y)$ and standard deviation $\sqrt{\mathrm{Var}(\mathcal{L}(Z) \mid \mathsf{HW}(Z) = y)}$.

3.3 Solving the LPN Problem

Numerous algorithms for solving LPN are known in the literature; a good survey is given by Pietrzak in [Pie12]. They generally require a huge number of LPN equations. However in our context, these equations come from side-channel acquisitions and thus remain in a rather scarce number. A well-known result of Lyubashevsky reduces the sample complexity, but its limitations on the noise render it inapplicable to our problem [Lyu05]. In this section we summarize the ideas we set-up for solving the LPN problem with a reduced number of samples and under reasonable levels of noise.

We take the point of view of an attacker: she has a limited quantity of side-channel information, thus a limited number of initial LPN samples. She also has a limited computing power and (most importantly) memory. She has two goals: firstly she wants to make sure that the attack will indeed be feasible in theory (this depends on the final number of reduced equations), thus she must compute it as exactly as possible (she cannot afford to miss one bit of complexity in the computations). Secondly, she has reasonable but limited resources and wants to make the attack as efficient as possible.

Algorithm Sketch. The main parameter of the algorithm is the initial bias: it determines the number of linear combinations steps we will be able to do before the final bias explodes. We fix it to 3 reductions (8 linear combinations). We look for small-weight linear combinations of initial equations that have their MSB cancelled. There's not enough initial LPN equations to use BKW or LF1 (*cf* Sect. 2.5) algorithms directly (they do not remove enough bits per iteration).

We thus first (rather artificially) square the number ν of LPN samples: for all elements $\mathbf{a}_i$ in the initial set, with error probability p (bias $\delta = 1 - 2p$), we build the set $(\mathbf{a}_{i,j})_{i \neq j} \doteq (\mathbf{a}_i \oplus \mathbf{a}_j)_{i,j}$. We then can do only 2 reductions. However, on the one hand, BKW-like algorithms will still not find enough reduced equations.

On the other hand, exhaustively looking for reduced equations among all linear combinations of at most 4 (corresponding to 2 reductions) amplified equations would not be very efficient. Consequently, we apply two steps of a generalized birthday paradox-like algorithm [Wag02].

Then assume that we obtain w-bits reduced equations. Once enough equations are found (this depends on the final bias of the equations, which is δ^8), we can directly apply a Walsh-Hadamard transform (WHT) to recover the w LSB of the secret if the attacker memory is greater than 2^w w-bits words. If we can only obtain equations reduced to $w' > w$ bits, we can simply guess the $w' - w$ bits of the secret and do a WHT on the last w bits. In this case, the search space can be reduced using the error/secret switching idea at the very beginning of the algorithm.

The algorithm steps as well as its time and space complexities are analyzed in details in [BCF+15]. From a practical perspective, the optimal choice depends on several parameters: number of traces, filtering ratio, level of noise, available memory, computing power. Several trade-offs are thus available to the attacker. The most obvious one is to trade side-channel measurements against computing needs. Using more traces either makes it possible to reduce the bias of the selected equations, or increases their number, reducing the reduction time (birthday paradox phase). In a nutshell, the more traces are available, the better. Given a fixed number of traces (order of magnitude 2^{20} to 2^{24}), the attacker fixes the filtering threshold λ. Increasing λ improves the bias of the selected equations. Thus less reduced equations are required for the WHT to correctly find w bits of the secret. Nonetheless, increasing λ also reduces the number of initial equations and thus makes the birthday paradox part of the algorithm slower. Concerning the reduction phase, it is well known that balancing the two phases of the generalized birthday paradox is the best way to reduce its complexity. Finally doubling the memory makes it possible recover one bit more with the WHT, while slightly more than doubling its time complexity: we fill the table with equations that are 1 bit less reduced, halving the time needed by the birthday paradox phase.

3.4 Comparison with State-of-the Art Attacks

Compared to [BFG14], our new attack performs better except in one scenario when $\mathsf{SNR} = 128$ and the number of available queries is very limited by the context. Indeed, for $\mathsf{SNR} = 128$ the attack in [BFG14] requires only 128 observations to get 128 equations with error probability 0.31 whereas our attack requires 2^{15} observations to achieve the same error probability. In the other contexts (*i.e.,* for higher levels of noise) the attack in [BFG14] faces strong limitations. Concretely, recovering the secret key becomes very hard if the inputs are not chosen. On the contrary, since our attack benefits from being quite insensitive to noise, it stays successful even for higher noise levels.

4 Extension to Chosen Inputs

In this section, we present a key-recovery technique which can be applied when the attacker is able to control the public multiplication operands $\mathbf{a}_i$. It is based on comparing the leakage for related inputs.

4.1 Comparing Leaks

In the so-called *chosen message model*, the attacker chooses ν messages $(\mathbf{a}_i)_{0 \leqslant i < \nu}$ in $\mathsf{GF}(2^n)$ and gets the corresponding leakages $\mathcal{L}(\mathbf{k} \cdot \mathbf{a}_i)$ as defined by Equation (2).

From the underlying associative property of the field $\mathsf{GF}(2^n)$, we remark[1] that the relation $(2 \cdot \mathbf{a}_i) \cdot \mathbf{k} = 2 \cdot (\mathbf{a}_i \cdot \mathbf{k})$ stands for every query $\mathbf{a}_i$. If the most significant bit of $\mathbf{a}_i \cdot \mathbf{k}$ is zero, then the latter relation implies that the bits of $\mathbf{a}_i \cdot k$ are simply shifted when computing $2 \cdot (\mathbf{a}_i \cdot \mathbf{k})$ which results in $\mathsf{HW}((2 \cdot \mathbf{a}_i) \cdot \mathbf{k}) = \mathsf{HW}(\mathbf{a}_i \cdot \mathbf{k})$. However, if the most significant bit of $\mathbf{a}_i \cdot \mathbf{k}$ is one, then the bits are also shifted but the result is summed with the constant value 23, which corresponds to the decimal representation of the binary coefficients of the non-leading monomials of the polynomial $x^{128} + x^7 + x^2 + x + 1$ involved in the representation of the field $\mathsf{GF}(2^{128})$ in AES-GCM. In this case, the Hamming weight values $\mathsf{HW}((2 \cdot \mathbf{a}_i) \cdot \mathbf{k})$ and $\mathsf{HW}(\mathbf{a}_i \cdot \mathbf{k})$ are necessarily different. Indeed, the bits are shifted, the less significant bit is set to one and the bits of $(\mathbf{a}_i \cdot \mathbf{k})$ at positions 0, 1 and 6 are flipped. Thus, the absolute value of the difference between both Hamming Weight values is equal to 3 with probability 1/4 or to 1 with probability 3/4.

Without noise, we can perfectly distinguish whether both Hamming weight values are equal or not, and thus get knowledge of the most significant bit of $\mathbf{a}_i \cdot \mathbf{k}$. Repeating the experiment for every power of two until 2^{128} (*i.e.,* with 128 queries) gives us the knowledge of every bit of the multiplication result and thus the recovery of $\mathbf{k}$. With noise, the recovery is no longer straightforward. To decide whether the noisy Hamming weights are equal or different, we fix a threshold τ depending on the SNR. Namely, if the distance $|\mathcal{L}((2 \cdot \mathbf{a}_i) \cdot \mathbf{k}) - \mathcal{L}(\mathbf{a}_i \cdot \mathbf{k})|$ is greater than τs where s is the signal standard deviation (here the standard deviation of $\mathsf{HW}(Z)$, say $\sqrt{n}/2$), then we decide that $\mathsf{HW}((2 \cdot \mathbf{a}_i) \cdot \mathbf{k}) \neq \mathsf{HW}(\mathbf{a}_i \cdot \mathbf{k})$ and thus that the most significant bit of $(\mathbf{a}_i \cdot \mathbf{k})$ equals one. The type I error probability p_{I} associated to this decision (*i.e.,* the probability of deciding that the Hamming weights are different while they are equal) satisfies:

$$\begin{aligned} p_{\mathrm{I}} &= \mathbb{P}\left[|\mathcal{L}((2 \cdot \mathbf{a}_i) \cdot \mathbf{k}) - \mathcal{L}(\mathbf{a}_i \cdot \mathbf{k})| > \tau s \mid \mathsf{HW}((2 \cdot \mathbf{a}_i) \cdot \mathbf{k}) = \mathsf{HW}(\mathbf{a}_i \cdot \mathbf{k})\right] \\ &= \mathbb{P}\left[|\varepsilon_{i+1} - \varepsilon_i| > \tau s\right] = 1 - \int_{-\tau s}^{\tau s} \phi_{\sigma\sqrt{2}}(u) du, \end{aligned}$$

where we recall that, according to (2), the variable ε_i (resp. ε_{i+1}) corresponds to the noise in the ith (resp. $(i+1)$th) observation $\mathcal{L}(\mathbf{a}_i \cdot \mathbf{k})$ (resp. $\mathcal{L}(\mathbf{a}_{i+1} \cdot \mathbf{k}) = \mathcal{L}((2 \cdot \mathbf{a}_i) \cdot \mathbf{k})$).

[1] We can simply choose $\mathbf{a}_i$ equal to 1.

Similarly, the type II error probability p_{II} (of deciding that the Hamming weight values are equal when they are different) satisfies:

$$p_{\text{II}} = \frac{3}{8}\left(\int_{-\tau s-1}^{\tau s-1} \phi_{\sigma\sqrt{2}}(u)du + \int_{-\tau s+1}^{\tau s+1} \phi_{\sigma\sqrt{2}}(u)du\right) + \frac{1}{8}\left(\int_{-\tau s-3}^{\tau s-3} \phi_{\sigma\sqrt{2}}(u)du + \int_{-\tau s+3}^{\tau s+3} \phi_{\sigma\sqrt{2}}(u)du\right).$$

Since, the key bits are all assumed to be balanced between one and zero, the probability of error p for each key bit is equal to $\frac{1}{2}(p_{\text{I}} + p_{\text{II}})$. Table 2 gives the thresholds τ which minimizes the error probability for different values of standard deviations.[2]

Table 2. Optimal threshold and probability of deciding correctly w.r.t. the SNR

SNR (σ)	128 (0.5)	8 (2.0)	2 (4.0)	0.5 (8.0)
τ	0.094	0.171	0.301	0.536
p	0.003	0.27	0.39	0.46

Comparing to Table 1, the error probabilities in Table 2 are much more advantageous and only 129 queries are required. If the number of queries is not limiting, the traces can be averaged to decrease the noise and thus improve the success rate. Another improvement is to correlate not only two consecutive powers of 2 but also non-consecutive ones (*e.g.*, 2^j and 2^{j+2}). Without noise, we do not get more information but in presence of noise we can improve the probability of deciding correctly.

4.2 Key Recovery

With the method described above, we only get 128 different linear equations in the key bits. Thus, we cannot use an LPN solving algorithm to recover the secret key in presence of errors. However, since we can average the measurements, we can significantly reduce the level of noise and remove the errors almost completely. For instance, with an SNR of 128 (which can also be achieved from an SNR of 2 and 64 repetitions), we get an average of $128 \times 0.003 = 0.384$ errors. Solving the system without error is straightforward when we use the powers of two since we directly have the key bits. Thus, inverting all the second members of the equations one-by-one to remove a single error leads to a global complexity of 2^7 key verifications. This complexity is easily achievable and remains reasonable to recover a 128-bit key.

5 Adaptation to Fresh Re-Keying

The core idea of the fresh re-keying countermeasure originally proposed in [MSGR10] for block cipher algorithm is to create a new *session* key from a

[2] Note that we did not consider so far the bias induced by the recovery of the less significant bits (whose values have been altered by previous squarings) since it is very negligible in practice.

public *nonce* for each new processing of the encryption algorithm. It guaranties that the secret (master) key is never used directly. To allow for the decryption of the ciphertext, the latter one is sent together with the nonce. For soundness, the fresh re-keying must satisfy two properties. First, it must be easy to protect against side-channel attacks. Secondly, it must have a good *diffusion* so that each bit of the new session key depends on a large number of bits of the master key, rendering attacks based on key-hypotheses testing inefficient. To satisfy the first property, [MSGR10] proposes to base the re-keying on linear functions. Efficient techniques are indeed known to secure the latter functions against SCA (*e.g.* higher-order masking has linear complexity for linear functions [ISW03, CGP+12]). To additionally satisfy the second property, [MSGR10] proposes to define the linear functions from *circulant* matrices deduced from the random nonce.

Let $\mathbf{k} \in \mathsf{GF}(2^8)^n$ denote the master key which must be protected and let $\mathbf{a} \in \mathsf{GF}(2^8)^n$ denote the nonce (generated at random). The square matrix whose lines correspond to all the rotations of the byte-coordinates of $\mathbf{a}$ (*e.g.* the i^{th} row corresponds to the vector $\mathbf{a}$ right-rotated i times) is denoted by $\mathsf{circ}(\mathrm{a}_0, \cdots, \mathrm{a}_{n-1})$. It satisfies:

$$\mathsf{circ}(\mathrm{a}_0, \cdots, \mathrm{a}_{n-1}) = \begin{pmatrix} \mathrm{a}_0 & \mathrm{a}_{n-1} & \mathrm{a}_{n-2} & \cdots & \mathrm{a}_1 \\ \mathrm{a}_1 & \mathrm{a}_0 & \mathrm{a}_{n-1} & \cdots & \mathrm{a}_2 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathrm{a}_{n-1} & \mathrm{a}_{n-2} & \mathrm{a}_{n-3} & \cdots & \mathrm{a}_0 \end{pmatrix},$$

and the session key $\mathbf{k}'$ is deduced from $(\mathbf{k}, \mathbf{a})$ as follows:

$$\mathbf{k}' = \mathsf{circ}(\mathrm{a}_0, \cdots, \mathrm{a}_{n-1}) \cdot \mathbf{k}, \tag{5}$$

where $\cdot$ denotes the scalar product in $\mathsf{GF}(2^8)^n$. After denoting the multiplication on $\mathsf{GF}(2^8)$ by $\otimes$, Equation (5) implies in particular that the i^{th} byte of $\mathbf{k}'$ satisfies:

$$\mathrm{k}'_i = \sum_{j=0}^{n-1} \mathrm{a}_{i+j \bmod n} \otimes \mathrm{k}_j.$$

It may be checked that the attack described in Sect. 3.1 applies against the multiplication specified by (5) similarly as for the multiplication in (1). Indeed, the matrix-vector product defined in (5) over $\mathsf{GF}(2^8)$ can be rewritten over $\mathsf{GF}(2)$ expressing each bit of $\mathbf{k}'$ as a linear combination of the bits of $\mathbf{k}$ with coefficients being themselves linear combinations of the bits of $\mathbf{a} \in \mathsf{GF}(2)^{128}$. Eventually, exactly like in previous section, for ν filtered messages the attack leads to an instance of the $\mathsf{LPN}(128, 128\nu, p)$ problem.[3] Actually, looking further in the fresh re-keying protocol, we can improve the attack by taking advantage of the context in which the fresh re-keying is used.

[3] As observed by the authors of [BFG14], the attack in [BFG14] does not apply to the multiplication specified by (5), essentially because of the circulant property of the matrix.

Until now, we have assumed that the multiplication output was stored in a 128-bit register, which essentially corresponds to an hardware implementation and is the worst case from the attacker point of view. If we switch to a software implementation *e.g.* running on a w-bit architecture, then the attacker can now target the manipulation of w-bit sub-parts of the refresh key $\mathbf{k}'$ which puts him in a more favourable context. By moreover assuming that $\mathbf{k}'$ is used as a secret parameter of a block cipher like AES (as proposed in [MSGR10]), then the attacker can exploit information leakage when the byte-coordinates of $\mathbf{k}'$ are manipulated separately. Observing the manipulation of each of the sixteen 8-bit chunks separately gives, for a same filtering ratio, a much lower error probability on the equations that what was achieved in the previous (hardware) context. This can be explained by the fact that exhibiting extreme Hamming weights is obviously much more easier on 8 bits than on 128 bits. For instance, filtering one observation over 2^{10} (i.e., $F(\lambda) = 2^{-10}$) with a SNR equal to 2 results in an error probability of $p = 0.28$ for $n = 128$ and $p = 0.065$ for $n = 8$, that is more than four times less. Table 3 gives the error probability p according to the proportion of filtered acquisitions $F(\lambda)$ for SNR equal to 128, 8, 2 and then 0.5 (as in Table 1) and $n = 8$.

Table 3. Error probability p according to the proportion of filtered acquisitions $F(\lambda)$.

$\log_2(1/F(\lambda))$	10	5	4	3	2	1	10	5	4	3	2	1
	SNR $= 128, \sigma = 0.125$						SNR $= 2, \sigma = 1$					
λ	2.93	2.15	2.02	1.47	1.33	0.71	3.88	2.62	2.28	1.89	1.42	0.83
p	$2.8 \cdot 10^{-19}$	0.09	0.11	0.17	0.21	0.28	$6.5 \cdot 10^{-2}$	0.16	0.19	0.22	0.26	0.32
	SNR $= 8, \sigma = 0.5$						SNR $= 0.5, \sigma = 2$					
λ	3.25	2.26	1.97	1.63	1.24	0.74	5.66	3.73	3.22	2.66	1.99	1.17
p	$5.9 \cdot 10^{-3}$	0.10	0.14	0.18	0.23	0.29	0.17	0.25	0.28	0.30	0.33	0.37

This confirms on different parameters that with much fewer observations, we have smaller error probabilities. Therefore, even for $F(\lambda) = 0.5$ (i.e., we only filter one observation over two), the system can be solved to recover the 128-bit key. Furthermore, it is worth noting that this new attack on an AES using a one-time key allows to recover the master key without observing any leakage in the fresh re-keying algorithm.

By using this trick which consists in observing the leakage of 8-bit session keys in the first round of the AES, we can also mount an attack towards the outlines of the approach proposed in [BFG14] against the AES-GCM multiplication. Since in this case only the first matrix row is involved in the computation, the coefficients of the key bits are different and each observation gives a useful linear equation. Plus, since we observe the leakage on 8-bit data, the noise impacts on the less significant bit of Hamming weight is reduced, which improves the system solving. However, the resulting attack remains much less efficient than our new attack, even in the number of required observations.

6 Practical Experiments

We showed in previous sections how to mount efficient side-channel attacks on finite-field multiplication over 128-bit data in different scenarios according to the

attacker capabilities. In order to verify the truthfullness of our leakage assumptions, we have mounted few of these attacks in practice and made some simulations. In particular, we implemented the AES-GCM and the fresh re-keying protocol on an ATMega328p and measured the leakage using the ChipWhisperer kit [OC14]. We also obtained the 100,000 traces of AES-GCM multiplication from [BFG14] corresponding to EM radiations of an FPGA implementation on the Virtex 5 of a SASEBO board.

We first illustrate the leakage behavior we obtained on the ATMega328p. Then we present experimental confirmations that the attack on AES-GCM with known inputs can actually be mounted. Afterwards, we show how efficient is the attack on fresh re-keying when the attacker can exploit 8-bit leakages of the first round of AES. Eventually, the reader may find in the extended version of this paper [BCF+15] an experiment corresponding to the chosen-message attack presented in Sect. 4 for a 128-bit multiplication implemented on the ATMega328p.

6.1 ATMega328p Leakage Behaviour

Since we are in software on an 8-bit implementation, we simulate a 128-bit leakage by summing the intermediate leakage on 8-bit parts of the result[4]. We randomly generated $100,000$ vectors $\mathbf{a} \in \mathsf{GF}(2)^{128}$ and, for a fixed key $\mathbf{k}$, we measured the leakage during the processing of $\mathbf{z} = \mathbf{a} \cdot \mathbf{k}$ as specified in AES-GCM (see (1)). Each measurement was composed of $4,992$ points among which we detected 16 points of interest by following a T-test approach as *e.g.* described in [GJJR11]. We afterwards verified that these points corresponded to the manipulation of the byte-coordinates $\mathbf{z}[i]$ of $\mathbf{z}$ after the multiplication processing.

For each $i \in [1..16]$, we denote by $g_{\mathrm{Id},i}$ the function $z \mapsto \mathbb{E}\left(\mathcal{L}(\mathbf{z}[i]) \mid \mathbf{z}[i] = z\right)$ and by $g_{\mathrm{HW},i}$ the function $y \mapsto \mathbb{E}\left(\mathcal{L}(\mathbf{z}[i]) \mid \mathsf{HW}(\mathbf{z}[i]) = y\right)$ (the first function corresponds to the mean of the leakage $\mathcal{L}(\mathbf{z}[i])$ knowing $\mathbf{z}[i] = z \in \mathsf{GF}(2)^8$ and the second function corresponds to the mean of the leakage $\mathcal{L}(\mathbf{z}[i])$ knowing $\mathsf{HW}(\mathbf{z}[i]) = y \in [0..8]$). In the top of Fig. 1, we plot for each $i \in [1..16]$ the distribution of our estimations of the values of $g_{\mathrm{Id},i}()$ (left-hand figure) and the distribution of the values of $g_{\mathrm{HW},i}()$. First, it may be observed that all the byte-coordinates, except the first one, leak quite similarly. The average mean and standard deviation of the functions $g_{\mathrm{ID},i}$ are -0.0301 and 0.0051 respectively. They are -0.0291 and 0.0092 for the functions $g_{\mathrm{HW},i}$. While the left-hand figure shows that the distributions of values differ from normal distributions, the right-hand figure exhibits a strong dependency between them and the distribution of the Hamming weight values of $\mathbf{z}[i]$. This shows that our implementation is a good target for our attack which requires that the deterministic part of the leakage monotonously depends on the Hamming weight of the manipulated data.

[4] Our purpose was to test the practical soundness of our theoretical analyses; we hence chose to artificially build a 128-bit leakage. The application of our attack to 8-bit chunks is the purpose of Sect. 6.3 where it is shown that this situation is much more favourable to the attacker.

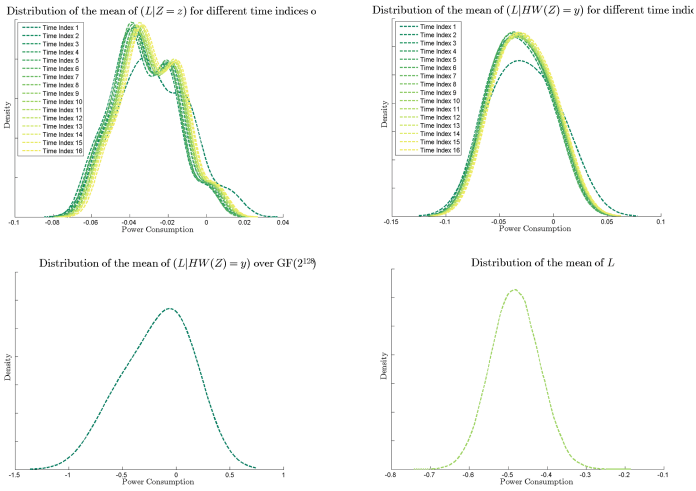

Fig. 1. Behaviour of the leakage w.r.t. the manipulated data Z

Eventually, we plot in the bottom-left figure an estimate (with kernel methods) of the distribution of the values $\mathbb{E}(\mathcal{L}(\mathbf{z})) \mid \mathsf{HW}(\mathbf{z}) = y)$ when y ranges in $[0..128]$ and $\mathcal{L}(\mathbf{z}) \doteq \sum_{i=1}^{16} \mathcal{L}(\mathbf{z}[i])$. Once again, the distribution is a not a perfect binomial one, but the figure shows that the deterministic part of the leakage monotonously depends on the Hamming weight of the manipulated data. The mean and the standard deviation of the plotted distribution are -0.1781 and $0,2392$ respectively. For completeness, we also plot in the bottom-right of Fig. 1 the distribution of the leakage values (after combining the 16 point of interest): the distribution looks very close to a Gaussian one.

6.2 Attacks on AES-GCM with Known Inputs

The aforementioned attack of AES-GCM with known inputs was almost completely performed for 96-bit keys (simulations for more leakage traces) and partially performed for 128-bit keys (the error probabilities were confirmed in practice).

Experiments on Filtering.

ATMega328p (128-bit). In this context, the leakage $\mathcal{L}(\mathbf{z})$ is built by summing the sixteen leakages $\mathcal{L}(\mathbf{z}[i])$, with $i \in [1..16]$. Theoretically, summing the sixteen intermediate Hamming weight values gives us exactly the Hamming weight value of the multiplication result. And summing the sixteen noise of standard deviation σ_8 results in a Gaussian noise of standard deviation $\sigma_{128} = 4 \cdot \sigma_8$. In practice, we get an SNR of 8.21 on the 128-bit simulated leakage. In Table 4 we provide the experimental bounds λ_{exp} and error probabilities p_{exp} corresponding to few levels

Table 4. Experimental and theoretical parameters corresponding to filtering proportion $F(\lambda)$ on the ATmega for 128-bit AES-GCM.

SNR $= 8.21, \sigma = 0.0206$							
$\log_2(1/F(\lambda))$	14	12	10	8	6	4	2
λ_{exp}	4.37	3.96	3.49	3.05	2.54	1.97	1.22
p_{exp}	0.383	0.386	0.393	0.407	0.420	0.434	0.452
λ_{the}	4.27	3.90	3.51	3.08	2.59	2.00	1.24
p_{the}	0.381	0.390	0.399	0.409	0.421	0.435	0.453

of filtering. We also indicate the theoretical estimates λ_{the} and p_{the} obtained by applying Formulas (3.2) and (3.2) to the template we obtained using the same set of traces. As it can be observed, the theoretical estimates are very close to the ones obtained experimentally (which validates our theoretical analysis, even for non Hamming weight model)[5].

Virtex 5 (128-bit). We additionally performed filtering on the traces from [BFG14] obtained from an FPGA implementation of GCM. Hereafter we provide theoretical (p_{the}) and experimental (p_{exp}) error probabilities for different values of the filtering parameter λ (Table 5). It must be noticed that experimental results correspond to expectations. The largest deviation (for $\lambda = 3.847$) is due to the fact that only 20 traces were kept after filtering[6].

Table 5. Error probabilities obtained from real traces.

λ	0.906	1.270	1.645	2.022	2.409	2.794	3.165	3.847
p_{the}	0.442	0.431	0.419	0.407	0.395	0.382	0.369	0.357
p_{exp}	0.441	0.430	0.418	0.405	0.392	0.379	0.370	0.361

ATMega328p (96-bit). As in the 128-bit case, the 96-bit leakage is simulated by summing the twelve intermediate 8-bit leakage of the multiplication result. Table 6 gives the bounds q and the error probabilities p corresponding to some levels of filtering[7].

[5] It must be noticed that a SNR equal to 8.21 in our experiments (with a noise standard deviation 0.0206) corresponds to a noise with standard deviation $\sigma = \sqrt{32/8.21} = 1.97$ in the theoretical Hamming weight model over 128-bit data.

[6] It must be noticed that, surprisingly, we also obtained an SNR equal to 8.21 in FPGA experiments but corresponding to a noise standard deviation of 7.11.

[7] An SNR equal to 8.7073 in our experiments (with a noise standard deviation 0.0173) corresponds to a noise with standard deviation $\sqrt{24/8.7073} = 1.66$ in the theoretical Hamming weight model over 96-bit data.

Table 6. Experimental and theoretical parameters corresponding to filtering proportion $F(\lambda)$ on the ATmega for 96-bit AES-GCM

SNR $= 8.7073, \sigma = 0.0173$						
$\log_2(1/F(\lambda))$	12	10	8	6	4	2
λ_{exp}	4.27	3.80	3.29	2.76	2.14	1.31
p_{exp}	0.377	0.387	0.402	0.414	0.429	0.449

LPN Experiments.

Attack on Simulated Traces (96-bit). We successfully performed our new attack on AES-GCM for a block-size reduced to 96 bits. We generated a 96-bit key $\mathbf{k}$, then generated 2^{20} uniform random $\mathbf{a}_i$. We simulated a leakage corresponding to the one obtained on the ATMega328p (*i.e.,* with the same statistics) and chose λ equal to 3.80 (filtering with probability 2^{-10}, error probability 0.387). This kept 916 relations, the less noisy one having weight 25 (error rate 0.260). We used this relation for secret/error switch. All in all, we got 87840 $\approx 2^{16,42}$ LPN equations. After 6 hours of parallelized generalized birthday computation (32 cores, 200 GB of RAM), we got $\approx 2^{39}$ equations reduced down to 36 bits. After a 36-bit Walsh transform ($\approx$ 2000 seconds, same machine), we recovered the 36 least significant bits of the error that we converted in 36 bits of the secret. This heavy computation corresponds to the most complex part of the attack and validates its success. We can afterwards find the remaining bits by iterating the attack with the knowledge of the recovered bits. This is a matter of minutes: it corresponds to an attack on a 60-bit key, which is much less expensive than the 96-bit case.

Expected Attack Complexities (128-bit). We provide here theoretical complexities for the key-recovery attack on 128-bit secret key. Experiments have been performed on 96-bit secrets and presented in the previous paragraph which confirm the accuracy of our theoretical estimates. We can see in Fig. 2 the evolution of the time complexity as a function of the memory available for the attack. Plots are provided for three different data complexities. We notice that the time/memory trade-off is only exploitable up to one point. This is due to the fact that when lots of memory is available, one may perform a larger Walsh-Hadamard transform to obtain more reduced equations. At some point, the time complexity of this transform will be predominant compared to the birthday paradox step and thus there will be no gain in increasing the Walsh size. A relevant time/memory trade-off for 2^{20} acquisitions is a time complexity of $2^{51.68}$ for 2^{36} bytes in memory (servers with such amount of memory can be bought easily).

6.3 Attack on Fresh Re-Keying

We detail here the attack that aims at recovering the master key from the leakages corresponding to the first round of the AES when the secret key is generated

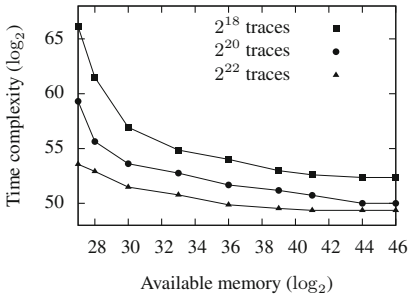

Fig. 2. Estimated complexities of the 128-bit attack ($\mathsf{SNR} = 8.21$).

by the fresh re-keying primitive described in Sect. 5. We present the known-input version of the attack, the chosen-input attack is described in [BCF+15].

Leakage Acquisition. We randomly generated 15,000 vectors $\mathbf{a} \in \mathsf{GF}(2)^{128}$ and 15,000 vectors $\mathbf{b} \in \mathsf{GF}(2)^8$. We then measured the 8-bit leakage during the processing of $\mathsf{Sbox}(\mathbf{z}[0] \oplus \mathbf{b})$ with $\mathbf{z}[0]$ the first byte of the multiplication between $\mathbf{a}$ and $\mathbf{k}$.

Filtering. We filtered the extreme consumption measurements in order to exhibit the extreme Hamming weight values. Table 7 gives the empirical error probabilities according to the proportion of filtering on the 15,000 observations. As explained in Sect. 5, the error probabilities are naturally much lower than for a 128-bit leakage.

Table 7. Error probability p according to the proportion of filtered acquisitions $F(\lambda)$ on the ATMega328p for the fresh re-keying with known inputs

$\log_2(1/F(\lambda))$	9	8	7	6	5	4	3	2	1
$\mathsf{SNR} = 8.6921, \sigma = 0.0165$									
λ	0.555	0.514	0.473	0.432	0.391	0.349	0.288	0.226	0.123
p	0.0	0.013	0.056	0.089	0.11	0.15	0.18	0.22	0.29

Key Recovery. With a sufficient (but still reasonable) filtering, we can directly recover the key by inverting the linear system of equations. For instance, in our experiments, filtering one observation over 2^9 gives $33 \times 8 = 264$ linear equations on the bits of $\mathbf{k}$ without a single error. Thus, inverting the system directly gives us the correct key.

References

[ACPS09] Applebaum, B., Cash, D., Peikert, C., Sahai, A.: Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 595–618. Springer, Heidelberg (2009)

[AG11] Arora, S., Ge, R.: New algorithms for learning in presence of errors. In: Aceto, L., Henzinger, M., Sgall, J. (eds.) ICALP 2011, Part I. LNCS, vol. 6755, pp. 403–415. Springer, Heidelberg (2011)

[BCF+15] Belaïd, S., Coron, J.-S., Fouque, P.-A., Gérard, B., Kammerer, J.-G., Prouff, E.: Improved side-channel analysis of finite-field multiplication. Cryptology ePrint Archive, Report 2015/542, (2015). http://eprint.iacr.org/

[BFG14] Belaïd, S., Fouque, P.-A., Gérard, B.: Side-Channel analysis of multiplications in GF(2128). In: Sarkar, P., Iwata, T. (eds.) ASIACRYPT 2014, Part II. LNCS, vol. 8874, pp. 306–325. Springer, Heidelberg (2014)

[BKW00] Blum, A., Kalai, A., Wasserman, H.: Noise-tolerant learning, the parity problem, and the statistical query model. In: 32nd ACM STOC, pp. 435–440. ACM Press, May 2000

[BTV15] Bogos, S., Tramer, F., Vaudenay, S.: On solving LPN using BKW and variants. Cryptology ePrint Archive, Report 2015/049, (2015). http://eprint.iacr.org/2015/049

[CGP+12] Carlet, C., Goubin, L., Prouff, E., Quisquater, M., Rivain, M.: Higher-Order masking schemes for S-boxes. In: Canteaut, A. (ed.) FSE 2012. LNCS, vol. 7549, pp. 366–384. Springer, Heidelberg (2012)

[CJRT05] Chekuri, C., Jansen, Rolim, K., J.D.P., Trevisan, L. (eds.) Approximation, randomization and combinatorial optimization, algorithms and techniques. In: 8th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems, APPROX 2005 and 9th International Workshop on Randomization and Computation, RANDOM 2005, Berkeley, CA, USA, August 22–24, 2005, Proceedings, vol. 3624 of Lecture Notes in Computer Science. Springer, Heidelberg (2005)

[DDP13] Dabosville, G., Doget, J., Prouff, E.: A new second-order side channel attack based on linear regression. IEEE Trans. Comput. **62**(8), 1629–1640 (2013)

[GJJR11] Goodwill, G., Jun, B., Jaffe, J., Rohatgi, P.: A testing methodology for side-channel resistance validation. In: Workshop NIAT (2011)

[GJL14] Guo, Q., Johansson, T., Löndahl, C.: Solving LPN using covering codes. In: Sarkar, P., Iwata, T. (eds.) ASIACRYPT 2014. LNCS, vol. 8873, pp. 1–20. Springer, Heidelberg (2014)

[HGJ10] Howgrave-Graham, N., Joux, A.: New generic algorithms for hard knapsacks. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 235–256. Springer, Heidelberg (2010)

[ISW03] 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 (2003)

[Kir11] Kirchner, P.: Improved generalized birthday attack. Cryptology ePrint Archive, Report 2011/377, (2011). http://eprint.iacr.org/2011/377

[LF06] Levieil, É., Fouque, P.-A.: An improved LPN algorithm. In: De Prisco, R., Yung, M. (eds.) SCN 2006. LNCS, vol. 4116, pp. 348–359. Springer, Heidelberg (2006)

[Lyu05] Lyubashevsky, V.: The parity problem in the presence of noise, decoding random linear codes, and the subset sum problem. In: Chekuri et al. (eds.) [CJRT05], pp. 378–389 (2005)

[MSGR10] Medwed, M., Standaert, F.-X., Großschädl, J., Regazzoni, F.: Fresh rekeying: security against side-channel and fault attacks for low-cost devices. In: Bernstein, D.J., Lange, T. (eds.) AFRICACRYPT 2010. LNCS, vol. 6055, pp. 279–296. Springer, Heidelberg (2010)

[MSJ12] Medwed, M., Standaert, F.-X., Joux, A.: Towards super-exponential side-channel security with efficient leakage-resilient PRFs. In: Prouff, E., Schaumont, P. (eds.) CHES 2012. LNCS, vol. 7428, pp. 193–212. Springer, Heidelberg (2012)

[OC14] O'Flynn, C., Chen, Z.: Chipwhisperer: an open-source platform for hardware embedded security research. Cryptology ePrint Archive, Report 2014/204 (2014). http://eprint.iacr.org/

[Pie12] Pietrzak, K.: Cryptography from learning parity with noise. In: Bieliková, M., Friedrich, G., Gottlob, G., Katzenbeisser, S., Turán, G. (eds.) SOFSEM 2012. LNCS, vol. 7147, pp. 99–114. Springer, Heidelberg (2012)

[RKSF11] Renauld, M., Kamel, D., Standaert, F.-X., Flandre, D.: Information theoretic and security analysis of a 65-nanometer DDSLL AES S-Box. In: Preneel, B., Takagi, T. (eds.) CHES 2011. LNCS, vol. 6917, pp. 223–239. Springer, Heidelberg (2011)

[SLP05] Schindler, W., Lemke, K., Paar, C.: A Stochastic Model for Differential Side Channel Cryptanalysis. In: Rao, J.R., Sunar, B. (eds.) CHES 2005. LNCS, vol. 3659, pp. 30–46. Springer, Heidelberg (2005)

[SS79] Schroeppel, R., Shamir, A.: A T s^2 = o(2^n) time/space tradeoff for certain np-complete problems. In: 20th Annual Symposium on Foundations of Computer Science, pp. 328–336. IEEE Computer Society, San Juan, Puerto Rico, 29–31 October (1979)

[Wag02] Wagner, D.: A generalized birthday problem. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, pp. 288–303. Springer, Heidelberg (2002)