Horizontal Side-Channel Attacks and Countermeasures on the ISW Masking Scheme

Alberto Battistello, Jean-Sébastien Coron, Emmanuel Prouff, Rina Zeitoun

► To cite this version:

Alberto Battistello, Jean-Sébastien Coron, Emmanuel Prouff, Rina Zeitoun. Horizontal Side-Channel Attacks and Countermeasures on the ISW Masking Scheme. 18th Conference on Cryptographic Hardware and Embedded Systems (CHES 2016), Aug 2016, Santa Barbara, CA, United States. pp.23 - 39, 10.1007/978-3-662-53140-2_2 . hal-01399577

HAL Id: hal-01399577

<https://hal.science/hal-01399577>

Submitted on 19 Nov 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Horizontal Side-Channel Attacks and Countermeasures on the ISW Masking Scheme

Alberto Battistello¹, Jean-Sébastien Coron², Emmanuel Prouff^{3*}, and Rina Zeitoun¹

¹ Oberthur Technologies, France
{a.battistello,r.zeitoun}@oberthur.com

² University of Luxembourg
jean-sebastien.coron@uni.lu

³ Sorbonne Universités, UPMC Univ Paris 06, CNRS, INRIA,
Laboratoire d'Informatique de Paris 6 (LIP6), Équipe PolSys, 4 place Jussieu, 75252
Paris Cedex 05, France

Abstract. A common countermeasure against side-channel attacks consists in using the masking scheme originally introduced by Ishai, Sahai and Wagner (ISW) at Crypto 2003, and further generalized by Rivain and Prouff at CHES 2010. The countermeasure is provably secure in the probing model, and it was showed by Duc, Dziembowski and Faust at Eurocrypt 2014 that the proof can be extended to the more realistic noisy leakage model. However the extension only applies if the leakage noise σ increases at least linearly with the masking order n , which is not necessarily possible in practice.

In this paper we investigate the security of an implementation when the previous condition is not satisfied, for example when the masking order n increases for a constant noise σ . We exhibit two (template) horizontal side-channel attacks against the Rivain-Prouff's secure multiplication scheme and we analyze their efficiency thanks to several simulations and experiments.

Eventually, we describe a variant of Rivain-Prouff's multiplication that is still provably secure in the original ISW model, and also heuristically secure against our new attacks.

1 Introduction

Side-channel analysis is a class of cryptanalytic attacks that exploit the physical environment of a cryptosystem to recover some *leakage* about its secrets. To secure implementations against this threat, security developers usually apply techniques inspired from *secret sharing* [Bla79,Sha79] or *multi-party computation* [CCD88]. The idea is to randomly split a secret into several shares such that the adversary needs all of them to reconstruct the secret. For these schemes, the

* Part of this work has been done at Safran Identity and Security, and while the author was at ANSSI, France.

number of shares n in which the key-dependent data are split plays the role of a security parameter.

A common countermeasure against side-channel attacks consists in using the masking scheme originally introduced by Ishai, Sahai and Wagner (ISW) [ISW03]. The countermeasure achieves provable security in the so-called *probing security model* [ISW03], in which the adversary can recover a limited number of intermediate variables of the computation. This model has been argued to be practically relevant to address so-called *higher-order* side-channel attacks and it has been the basis of several efficient schemes to protect block ciphers.

More recently, it has been shown in [DDF14] that the probing security of an implementation actually implies its security in the more realistic *noisy leakage model* introduced in [PR13]. More precisely, if an implementation obtained by applying the compiler in [ISW03] is secure at order n in the probing model, then [DFS15, Theorem 3] shows that the success probability of distinguishing the correct key among $|\mathbb{K}|$ candidates is bounded above by $|\mathbb{K}| \cdot 2^{-n/9}$ if the leakage L_i on each intermediate variable X_i satisfies:

$$I(X_i; L_i) \leq 2 \cdot (|\mathbb{K}| \cdot (28n + 16))^{-2} ,$$

where $I(\cdot; \cdot)$ denotes the *mutual information* and where the index i ranges from 1 to the total number of intermediate variables.

In this paper we investigate what happens when the above condition is not satisfied. Since the above mutual information $I(X_i; L_i)$ can be approximated by $k/(8\sigma^2)$ in the Hamming weight model in $\mathbb{F}_{2^k}$, where σ is the noise in the measurement (see the full version of this paper [BCPZ16]), this amounts to investigating the security of Ishai-Sahai-Wagner's (ISW) implementations when the number of shares n satisfies:

$$n > c \cdot \sigma$$

As already observed in previous works [VGS14, CFG⁺10], the fact that the same share (or more generally several data depending on the same sensitive value) is manipulated several times may open the door to new attacks which are not taken into account in the probing model. Those attacks, sometimes called *horizontal* [CFG⁺10] or (*Template*) *algebraic* [ORSW12, VGS14] exploit the algebraic dependency between several intermediate results to discriminate key hypotheses.

In this paper, we exhibit two (horizontal) side channel attacks against the ISW multiplication algorithm. These attacks show that the use of this algorithm (and its extension proposed by Rivain and Prouff in [RP10]) may introduce a weakness with respect to horizontal side channel attacks if the sharing order n is such that $n > c \cdot \sigma^2$, where σ is the measurement noise. While the first attack is too costly (even for low noise contexts) to make it applicable in practice, the second attack, which essentially iterates the first one until achieving a satisfying likelihood, shows very good performances. For instance, when the leakages are simulated by noisy Hamming weights computed over $\mathbb{F}_{2^8}$ with $\sigma = 1$, it recovers all the shares of a 21-sharing. We also confirm the practicality of our attack with a real life experiment on a development platform embedding the **ATMega328**

processor (see the full version of this paper [BCPZ16]). Actually, in this context where the leakages are multivariate and not univariate as in our theoretical analyses and simulations, the attack appears to be more efficient than expected and recovers all the shares of a n -sharing when $n \geq 40$.

Eventually, we describe a variant of Rivain-Prouff’s multiplication that is still provably secure in the original ISW model, and also heuristically secure against our new attacks. Our new countermeasure is similar to the countermeasure in [FRR⁺10], in that it can be divided in two steps: a “matrix” step in which starting from the input shares x_i and y_j , one obtains a matrix $x_i \cdot y_j$ with n^2 elements, and a “compression” step in which one uses some randomness to get back to a n -sharing c_i . Assuming a leak-free component, the countermeasure in [FRR⁺10] is proven secure in the noisy leakage model, in which the leakage function reveals all the bits of the internal state of the circuit, perturbed by independent binomial noise. Our countermeasure does not use any leak-free component, but is only heuristically secure in the noisy leakage model (see Section 8.2 for our security analysis).

2 Preliminaries

For two positive integers n and d , a (n, d) -sharing of a variable x defined over some finite field $\mathbb{F}_{2^k}$ is a random vector $(x_1, x_2, \dots, x_n)$ over $\mathbb{F}_{2^k}$ such that $x = \sum_{i=1}^n x_i$ holds (*completeness equality*) and any tuple of $d-1$ shares x_i is a uniform random vector over $(\mathbb{F}_{2^k})^{d-1}$. If $n = d$, the terminology simplifies to n -sharing. An algorithm with domain $(\mathbb{F}_{2^k})^n$ is said to be $(n-1)$ th-order secure in the probing model if on input an n -sharing $(x_1, x_2, \dots, x_n)$ of some variable x , it admits no tuple of $n-1$ or fewer intermediate variables that depends on x .

We refer to the full version of this paper [BCPZ16] for the definitions of Signal to Noise Ratio (SNR), Gaussian distribution, entropy and differential entropy.

3 Secure Multiplication Schemes

In this section, we recall the secure multiplication scheme over $\mathbb{F}_2$ introduced in [ISW03] and its extension to any field $\mathbb{F}_{2^k}$ proposed in [RP10].

Ishai-Sahai-Wagner’s Scheme [ISW03]. Let x^* and y^* be binary values from $\mathbb{F}_2$ and let $(x_i)_{1 \leq i \leq n}$ and $(y_i)_{1 \leq i \leq n}$ be n -sharings of x^* and y^* respectively. To securely compute a sharing of $c = x^* \cdot y^*$ from $(x_i)_{1 \leq i \leq n}$ and $(y_i)_{1 \leq i \leq n}$, the ISW method works as follows:

1. For every $1 \leq i < j \leq n$, pick up a random bit $r_{i,j}$.
2. For every $1 \leq i < j \leq n$, compute $r_{j,i} = (r_{i,j} + x_i \cdot y_j) + x_j \cdot y_i$.
3. For every $1 \leq i \leq n$, compute $c_i = x_i \cdot y_i + \sum_{j \neq i} r_{i,j}$.

The above multiplication scheme achieves security at order $\lfloor n/2 \rfloor$ in the probing security model [ISW03].

The Rivain-Prouff Scheme. The ISW countermeasure was extended to $\mathbb{F}_{2^k}$ by Rivain and Prouff in [RP10]. As showed in [BBD⁺15], the **SecMult** algorithm below is secure in the ISW probing model against t probes for $n \geq t + 1$ shares; the authors also show that with some additional mask refreshing, the Rivain-Prouff countermeasure for the full AES can be made secure with $n \geq t + 1$ shares.

Algorithm 1 SecMult

Input: the n -sharings $(x_i)_{i \in [1..n]}$ and $(y_j)_{j \in [1..n]}$ of x^* and y^* respectively
Output: the n -sharing $(c_i)_{i \in [1..n]}$ of $x^* \cdot y^*$

```

1: for  $i = 1$  to  $n$  do
2:   for  $j = i + 1$  to  $n$  do
3:      $r_{i,j} \leftarrow \mathbb{F}_{2^k}$ 
4:      $r_{j,i} \leftarrow (r_{i,j} + x_i \cdot y_j) + x_j \cdot y_i$ 
5:   end for
6: end for
7: for  $i = 1$  to  $n$  do
8:    $c_i \leftarrow x_i \cdot y_i$ 
9:   for  $j = 1$  to  $n$ ,  $j \neq i$  do  $c_i \leftarrow c_i + r_{i,j}$ 
10: end for
11: return  $(c_1, c_1, \dots, c_n)$ 

```

In Algorithm 1, one can check that each share x_i or y_j is manipulated n times, whereas each product $x_i y_j$ is manipulated a single time. This gives a total of $3n^2$ manipulations that can be observed through side channels.

4 Horizontal DPA Attack

4.1 Problem description.

Let $(x_i)_{i \in [1..n]}$ and $(y_i)_{i \in [1..n]}$ be respectively the n -sharings of x^* and y^* (namely, we have $x^* = x_1 + \dots + x_n$ and $y^* = y_1 + \dots + y_n$). We assume that an adversary gets, during the processing of Algorithm 1, a single observation of each of the following random variables for $1 \leq i, j \leq n$:

$$L_i = \varphi(x_i) + B_i \tag{1}$$

$$L'_j = \varphi(y_j) + B'_j \tag{2}$$

$$L''_{ij} = \varphi(x_i \cdot y_j) + B''_{ij} \tag{3}$$

where φ is an unknown function which depends on the device architecture, where B_i, B'_j are Gaussian noise of standard deviation $\sigma/\sqrt{n}$, and B''_{ij} is Gaussian noise with standard deviation σ . Namely we assume that each x_i and y_j is processed n times, so by averaging the standard deviation is divided by a factor $\sqrt{n}$, which gives $\sigma/\sqrt{n}$ if we assume that the initial noise standard deviation is σ .

The random variables associated to the i th share x_i and the j th share y_j are respectively denoted by X_i and Y_j . Our goal is to recover the secret variable x^* (and/or y^*).

4.2 Complexity Lower Bound: Entropy Analysis of Noisy Hamming Weight Leakage

For simplicity, we first restrict ourselves to a leakage function φ equal to the Hamming weight of the variable being manipulated. In that case, the mutual information $I(X; L)$ between the Hamming weight of a uniform random variable X defined over $\mathbb{F}_{2^k}$ and a noisy observation L of this Hamming weight can be approximated as:

$$I(X; L) \simeq \frac{k}{8\sigma^2} , \quad (4)$$

if the noise being modeled by a Gaussian random variable has standard deviation σ . This approximation, whose derivation is given in the full version of this paper [BCPZ16], is only true for large σ .

To recover a total of $2n$ shares (n shares of x^* and y^* respectively) from $3n^2$ Hamming weight leakages (namely each manipulation leaks according to (1)-(3) with $\varphi = \text{HW}$), the total amount of information to be recovered is $2n \cdot k$ if we assume that the shares are i.i.d. with uniform distribution over $\mathbb{F}_{2^k}$. Therefore, since we have a total of $3n^2$ observations during the execution of Algorithm 1, we obtain from (4) that the noise standard deviation σ and the sharing order n must satisfy the following inequality for a side channel attack to be feasible:

$$3 \cdot n^2 \cdot \frac{k}{8\sigma^2} > 2n \cdot k . \quad (5)$$

We obtain an equality of the form $n > c \cdot \sigma^2$ for some constant c , as in a classical (vertical) side channel attack trying to recover x^* from n observations of intermediate variables depending on x^* [CJRR99]. This analogy between horizontal and vertical attacks has already been noticed in previous papers like [CFG⁺10] or [BJPW13]. Note that in principle the constant c is independent of the field degree k (which has also been observed in previous papers, see for instance [SVO⁺10]).

4.3 Attack With Perfect Hamming Weight Observations

In the full version of this paper [BCPZ16], we consider the particular case of perfect Hamming weight measurements (no noise), using a maximum likelihood approach. We show that even with perfect observations of the Hamming weight, depending on the finite-field representation, we are not always guaranteed to recover the secret variable x^* ; however for the finite field representation used in AES the attack enables to recover the secret x^* for a large enough number of observations.

4.4 Maximum Likelihood Attack: Theoretical Attack with the full ISW State

For most field representations and leakage functions, the maximum likelihood approach used in the previous section recovers the i -th share of x^* from an observation of L_i and an observation of (L'_j, L''_{ij}) for every $j \in [1..n]$. It extends straightforwardly to noisy scenarios and we shall detail this extension in Section 5.1. However, the disadvantage of this approach is that it recovers each share separately, before rebuilding x^* and y^* from them. From a pure information theoretic point of view this is suboptimal since (1) the final purpose is not to recover all the shares perfectly but only the shared values and (2) only $3n$ observations are used to recover each share whereas the full tuple of $3n^2$ observations brings more information. Actually, the most efficient attack in terms of leakage exploitation consists in using the joint distribution of $(L_i, L'_j, L''_{ij})_{i,j \in [1..n]}$ to distinguish the correct hypothesis about $x^* = x_1 + x_2 + \dots + x_n$ and $y^* = y_1 + y_2 + \dots + y_n$.

As already observed in Section 3, during the processing of Algorithm 1, the adversary may get a tuple $(\ell_{ij})_{j \in [1..n]}$ (resp. $(\ell'_{ij})_{i \in [1..n]}$) of n observations for each L_i (resp. each L'_j) and one observation ℓ''_{ij} for each L''_{ij} . The full tuple of observations $(\ell_{ij}, \ell'_{ij}, \ell''_{ij})_{i,j}$ is denoted by ℓ , and we denote by $\mathbf{L}$ the corresponding random variable¹. Then, to recover (x^*, y^*) from ℓ , the *maximum likelihood approach* starts by estimating the pdfs $f_{\mathbf{L}|X^*=x^*, Y^*=y^*}$ for every possible (x^*, y^*) , and then estimates the following vector of distinguisher values for every hypothesis (x, y) :

$$d_{\text{ML}}^*(\ell) \doteq (f_{\mathbf{L}|(X^*, Y^*)}(\ell, (x, y)))_{(x,y) \in \mathbb{F}_{2^k}^2} \quad (6)$$

The pair (x, y) maximizing the above probability is eventually chosen.

At a first glance, the estimation of the pdfs $f_{\mathbf{L}|X^*=x^*, Y^*=y^*}$ seems to be challenging. However, it can be deduced from the estimations of the pdfs associated to the manipulations of the shares. Indeed, after denoting by $p_{x,y}$ each probability value in the right-hand side of (6), and by using the law of total probability together with the fact that the noises are independent, we get:

$$2^{2kn} \cdot p_{x,y} = \sum_{\substack{x_1, \dots, x_n \in \mathbb{F}_{2^k} \\ x = x_1 + \dots + x_n}} \sum_{\substack{y_1, \dots, y_n \in \mathbb{F}_{2^k} \\ y = y_1 + \dots + y_n}} \prod_{i,j=1}^n f_{L_i|X_i}(\ell_{ij}, x_i) \cdot f_{L'_j|Y_j}(\ell'_{ij}, y_j) \cdot f_{L''_{ij}|X_i Y_j}(\ell''_{ij}, x_i y_j) .$$

Unfortunately, even if the equation above shows how to deduce the pdfs $f_{\mathbf{L}|(X^*, Y^*)}(\cdot, (x^*, y^*))$ from characterizations of the shares' manipulations, a direct processing of the probability has complexity $O(2^{2nk})$. By representing the sum over the x_i 's as a sequence of convolution products, and thanks to Walsh

¹ In (1)-(3), it is assumed that the observations $(\ell_{ij})_{j \in [1..n]}$ and $(\ell'_{ij})_{i \in [1..n]}$ are averaged to build a single observation with noise divided by $\sqrt{n}$. This assumption is not done here in order to stay as general as possible.

transforms processing, the complexity can be easily reduced to $O(n2^{n(k+1)})$. The latter complexity stays however too high, even for small values of n and k , which led us to look at alternatives to this attack.

5 First Attack: Maximum Likelihood Attack on a Single Matrix Row

5.1 Attack Description

In this section, we explain how to recover each share x_i of x^* separately, by observing the processing of Algorithm 1. Applying this attack against all the shares leads to the full recovery of the sensitive value x^* with some success probability, which is essentially the product of the success probabilities of the attack on each share separately.

Given a share x_i , the attack consists in collecting the leakages on $(y_j, x_i \cdot y_j)$ for every $j \in [1..n]$. Therefore the attack is essentially a horizontal version of the classical (vertical) second-order side-channel attack, where each share x_i is multiplicatively masked over $\mathbb{F}_{2^k}$ by a random y_j for $j \in [1..n]$.

The most efficient attack to maximize the amount of information recovered on X_i from a tuple of observations $\ell \doteq \ell_i, (\ell'_j, \ell''_{ij})_{j \in [1..n]} \leftrightarrow \mathbf{L} \doteq L_i, (L'_j, L''_{ij})_{j \in [1..n]}$ consists in applying a *maximum likelihood approach* [CJRR99, GHR15], which amounts to computing the following vector of distinguisher values:

$$d_{\text{ML}}(\ell) \doteq (f_{\mathbf{L}|X_i}(\ell, \hat{x}_i))_{\hat{x}_i \in \mathbb{F}_{2^k}} \quad (7)$$

and in choosing the candidate $\hat{x}_i$ which maximizes the probability. We refer to the full version of this paper [BCPZ16] for the derivation of each score $f_{\mathbf{L}|X_i}(\ell, \hat{x}_i)$ in (7); we obtain:

$$f_{(L'_j, L''_{ij})|X_i}((\ell'_j, \ell''_{ij}), \hat{x}_i) = \sum_{y \in \mathbb{F}_{2^k}} f_{(L'_j, L''_{ij})|(X_i, Y_j)}((\ell'_j, \ell''_{ij}), (\hat{x}_i, y)) \cdot \text{pr}_{Y_j}(y) \quad , \quad (8)$$

and similarly:

$$f_{(L_i, L''_{ij})|Y_j}((\ell_i, \ell''_{ij}), \hat{y}_j) = \sum_{x \in \mathbb{F}_{2^k}} f_{(L_i, L''_{ij})|(X_i, Y_j)}((\ell_i, \ell''_{ij}), (x, \hat{y}_j)) \cdot \text{p}_{X_i}(x) \quad . \quad (9)$$

5.2 Complexity Analysis

As mentioned previously, given a share x_i , the attack consists in collecting the leakages on $(y_j, x_i \cdot y_j)$ for every $j \in [1..n]$. Therefore the attack is essentially an horizontal version of the classical (vertical) second-order side-channel attack. In principle the number n of leakage samples needed to recover x_i with good probability (aka the attack complexity) should consequently be $n = \mathcal{O}(\sigma^4)$ [CJRR99, GHR15, SVO⁺10]. This holds when multiplying two leakages both with noise having σ as standard deviation. However here the leakage on y_j has a

noise with a standard deviation $\sigma/\sqrt{n}$ instead of σ (thanks to the averaging step). Therefore the noise of the product becomes $\sigma^2/\sqrt{n}$ (instead of σ^2), which gives after averaging with n measurements a standard deviation of σ^2/n , and therefore an attack complexity satisfying $n = \mathcal{O}(\sigma^2)$, as in a classical first-order side-channel attack.

5.3 Numerical Experiments

The attack presented in Sect. 5.1 has been implemented against each share x_i of a value x , with the leakages being simulated according to (1)-(3) with $\varphi = \text{HW}$. For the noise standard deviation σ and the sharing order n , different values have been tested to enlighten the relation between these two parameters. We stated that an attack succeeds iff the totality of the n shares x_i have been recovered, which leads to the full recovery of x^* . We recall that, since the shares x_i are manipulated n times, measurements for the leakages L_i and L'_j have noise standard deviations $\sigma/\sqrt{n}$ instead of σ . For efficiency reasons, we have chosen to work in the finite field $\mathbb{F}_{2^4}$ (namely $k = 4$ in previous analyses).

For various noise standard deviations σ with $\text{SNR} = k(2\sigma)^{-2}$ (*i.e.* $\text{SNR} = \sigma^{-2}$ for $k = 4$), Table 1 gives the average minimum number n of shares required for the attack to succeed with probability strictly greater than 0.5 (the averaging being computed over 300 attack iterations). The attack complexity $n = \mathcal{O}(\sigma^2)$ argued in Sect. 5.2 is confirmed by the trend of these numerical experiments. Undeniably, this efficiency is quickly too poor for practical applications where n is small (clearly lower than 10) and the SNR is high (smaller than 1).

σ (SNR)	0 ($+\infty$)	0.2 (25)	0.4 (6.25)	0.6 (2.77)	0.8 (1.56)	1 (1)
n	12	14	30	73	160	284

Table 1: First attack: number of shares n as a function of the noise σ to succeed with probability > 0.5

6 Second Attack: Iterative Attack

6.1 Attack Description

From the discussions in Sect. 4.4, and in view of the poor efficiency of the previous attack, we investigated another strategy which targets all the shares simultaneously. Essentially, the core idea of our second attack described below is to apply several attacks recursively on the x_i 's and y_j 's, and to refine step by step the likelihood of each candidate for the tuple of shares. Namely, we start by applying the attack described in Section 5.1 in order to compute, for every i , a likelihood probability for each hypothesis $X_i = x$ (x ranging over $\mathbb{F}_{2^k}$); then we

apply the same attack in order to compute, for every j , a likelihood probability for each hypothesis $Y_j = y$ (y ranging over $\mathbb{F}_{2^k}$) with the single difference that the probability $p_{X_i}(x)$ in (9) is replaced by the likelihood probability which was just computed². Then, one reiterates the attack to refine the likelihood probabilities $(p_{X_i}(x))_{x \in \mathbb{F}_{2^k}}$, by evaluating (8) with the uniform distribution $p_{Y_j}(y)$ being replaced by the likelihood probability $\text{new-}p_{Y_j}(y)$ which has been previously computed. The scheme is afterwards repeated until the maximum taken by the pdfs of each share X_i and Y_j is greater than some threshold β . In order to have better results, we perform the whole attack a second time, by starting with the computation of the likelihood probability for each hypothesis $Y_j = y$ instead of starting by $X_i = x$.

We give the formal description of the attack processing in Algorithm 2 (in order to have the complete attack, one should perform the while loop a second time, by rather starting with the computation of $\text{new-}p_{Y_j}(y)$ instead of $\text{new-}p_{X_i}(x)$).

6.2 Numerical Experiments

The iterative attack described in Algorithm 2 has been tested against leakages simulations defined exactly as in Section 5.3. As previously we stated that an attack succeeds if the totality of the n shares x_i have been recovered, which leads to the full recovery of x^* . For various noise standard deviations σ with $\text{SNR} = k(2\sigma)^{-2}$, Table 2 gives the average minimum number of shares n required for the attack to succeed with probability strictly greater than 0.5 (the averaging being computed over 300 attack iterations). The first row corresponds to $k = 4$, and the second row to $k = 8$ (the corresponding SNRs are $\text{SNR}_4 = \sigma^{-2}$ and $\text{SNR}_8 = (\sqrt{2}\sigma^2)^{-1}$). Numerical experiments yield greatly improved results in comparison to those obtained by running the basic attack. Namely, in $\mathbb{F}_{2^4}$, for a noise $\sigma = 0$, the number of shares required is 2, while 12 shares were needed for the basic attack, and the improvement is even more confirmed with a growing σ : for a noise $\sigma = 1$, the number of shares required is 25, while 284 shares were needed for the basic attack. It can also be observed that the results for shares in $\mathbb{F}_{2^4}$ and $\mathbb{F}_{2^8}$ are relatively close, the number of shares being most likely slightly smaller for shares in $\mathbb{F}_{2^4}$ than in $\mathbb{F}_{2^8}$. This observation is in-line with the lower bound in (5), where the cardinality 2^k of the finite field plays no role.

σ ($\text{SNR}_4, \text{SNR}_8$)	0 ($+\infty, +\infty$)	0.2 (25, 17.67)	0.4 (6.25, 4.41)	0.6 (2.77, 1.96)	0.8 (1.56, 1.10)	1 (1, 0.7071)
n (for $\mathbb{F}_{2^4}$)	2	2	3	6	13	25
n (for $\mathbb{F}_{2^8}$)	5	6	8	11	16	21

Table 2: Iterative attack: number of shares n as a function of the noise σ to succeed with probability > 0.5 in $\mathbb{F}_{2^4}$ (first row) and in $\mathbb{F}_{2^8}$ (second row).

² Actually to get the probability of $X_i \mid \mathbf{L}$ instead of $\mathbf{L} \mid X_i$, Bayes' Formula is applied which explains the division by the sum of probabilities in the lines 14 and 19 in Algorithm 2.

Algorithm 2 Iterative Maximum Likelihood Attack

Input: a threshold β , an observation ℓ_i of each L_i , an observation ℓ'_j of each L'_j and one observation ℓ''_{ij} of each L''_{ij} (the random variables being defined as in (1)-(3))

Output: a n -tuple of pdfs $(p_{X_i})_i$ (resp. $(p_{Y_i})_i$) such that, for every $i \in [1..n]$, at least one $\hat{x}_i$ (resp. $\hat{y}_i$) satisfies $p_{X_i}(\hat{x}_i) \geq \beta$ (resp. $p_{Y_i}(\hat{y}_i) \geq \beta$)

```

1: for  $i = 1$  to  $n$  do
2:   for  $x \in \mathbb{F}_{2^k}$  do    # Initialize the likelihood of each candidate for  $X_i$ 
3:      $p_{X_i}(x) = f_{L_i|X_i}(\ell_i, x)$ 
4:   end for
5:   for  $y \in \mathbb{F}_{2^k}$  do    # Initialize the likelihood of each candidate for  $Y_i$ 
6:      $p_{Y_i}(y) = f_{L'_i|Y_i}(\ell'_i, y)$ 
7:      $\text{new-}p_{Y_i}(y) = p_{Y_i}(y)$ 
8:   end for
9: end for

10: while  $\text{end} \neq n$  do
11:    $\text{end} \leftarrow 0$ 
12:   for  $i = 1$  to  $n$  do
13:     for  $x \in \mathbb{F}_{2^k}$  do # Compute/Update the likelihood of each candidate for  $X_i$ 
14:        $\text{new-}p_{X_i}(x) = 2^{-(2n+1)k} \frac{p_{X_i}(x)}{\sum_{x' \in \mathbb{F}_{2^k}} p_{X_i}(x')} \prod_{j=1}^n \sum_{y \in \mathbb{F}_{2^k}} \frac{\text{new-}p_{Y_j}(y)}{\sum_{y' \in \mathbb{F}_{2^k}} \text{new-}p_{Y_j}(y')} \cdot$ 
          $f_{L''_{ij}|X_i Y_j}(\ell''_{ij}, x \cdot y)$ 
15:     end for
16:   end for
17:   for  $i = 1$  to  $n$  do
18:     for  $y \in \mathbb{F}_{2^k}$  do # Compute/Update the likelihood of each candidate for  $Y_i$ 
19:        $\text{new-}p_{Y_i}(y) = 2^{-(2n+1)k} \frac{p_{Y_i}(y)}{\sum_{y' \in \mathbb{F}_{2^k}} p_{Y_i}(y')} \prod_{j=1}^n \sum_{x \in \mathbb{F}_{2^k}} \frac{\text{new-}p_{X_j}(x)}{\sum_{x' \in \mathbb{F}_{2^k}} \text{new-}p_{X_j}(x')} \cdot$ 
          $f_{L''_{ij}|X_i Y_j}(\ell''_{ij}, x \cdot y)$ 
20:     end for
21:   end for
22:   for  $i = 1$  to  $n$  do
23:     if  $\max_x(\text{new-}p_{X_i}(x)) \geq \beta$  and  $\max_y(\text{new-}p_{Y_i}(y)) \geq \beta$  then
24:        $\text{end} \leftarrow \text{end} + 1$ 
25:     end if
26:   end for
27: end while

```

7 Practical Results

In the full version of this paper [BCPZ16], we describe the result of practical experiments of our attack against a development platform embedding the ATmega328 processor.

8 A Countermeasure against the previous Attacks

8.1 Description

In the following, we describe a countermeasure against the previous attack against the Rivain-Prouff algorithm. More precisely, we describe a variant of Algorithm 1, called RefSecMult, to compute an n -sharing of $c = x^* \cdot y^*$ from $(x_i)_{i \in [1..n]}$ and $(y_i)_{i \in [1..n]}$. Our new algorithm is still provably secure in the original ISW probing model, and heuristically secure against the horizontal side-channel attacks described in the previous sections.

Algorithm 3 RefSecMult

Input: n -sharings $(x_i)_{i \in [1..n]}$ and $(y_j)_{j \in [1..n]}$ of x^* and y^* respectively

Output: an n -sharing $(c_i)_{i \in [1..n]}$ of $x^* \cdot y^*$

```

1:  $M_{ij} \leftarrow \text{MatMult}((x_1, \dots, x_n), (y_1, \dots, y_n))$ 
2: for  $i = 1$  to  $n$  do
3:   for  $j = i + 1$  to  $n$  do
4:      $r_{i,j} \xleftarrow{\$} \mathbb{F}_{2^k}$ 
5:      $r_{j,i} \leftarrow (r_{i,j} + M_{ij}) + M_{ji}$ 
6:   end for
7: end for
8: for  $i = 1$  to  $n$  do
9:    $c_i \leftarrow M_{ii}$ 
10:  for  $j = 1$  to  $n$ ,  $j \neq i$  do  $c_i \leftarrow c_i + r_{i,j}$ 
11: end for
12: return  $(c_1, c_1, \dots, c_n)$ 
```

As observed in [FRR⁺10], the ISW and Rivain-Prouff countermeasures can be divided in two steps: a “matrix” step in which starting from the input shares x_i and y_j , one obtains a matrix $x_i \cdot y_j$ with n^2 elements, and a “compression” step in which one uses some randomness to get back to a n -sharing c_i . Namely the matrix elements $(x_i \cdot y_j)_{1 \leq i, j \leq n}$ form a n^2 -sharing of $x^* \cdot y^*$:

$$x^* \cdot y^* = \left(\sum_{i=1}^n x_i \right) \cdot \left(\sum_{j=1}^n y_j \right) = \sum_{1 \leq i, j \leq n} x_i \cdot y_j \quad (10)$$

and the goal of the compression step is to securely go from such n^2 -sharing of $x^* \cdot y^*$ to a n -sharing of $x^* \cdot y^*$.

Our new countermeasure (Algorithm 3) uses the same compression step as Rivain-Prouff, but with a different matrix step, called **MatMult** (Algorithm 4), so that the shares x_i and y_j are not used multiple times (as when computing the matrix elements $x_i \cdot y_j$ in Rivain-Prouff). Eventually the **MatMult** algorithm outputs a matrix $(M_{ij})_{1 \leq i, j \leq n}$ which is still a n^2 -sharing of $x^* \cdot y^*$, as in (10); therefore using the same compression step as Rivain-Prouff, Algorithm 3 outputs a n -sharing of $x^* \cdot y^*$, as required.

Algorithm 4 MatMult

Input: the n -sharings $(x_i)_{i \in [1..n]}$ and $(y_j)_{j \in [1..n]}$ of x^* and y^* respectively

Output: the n^2 -sharing $(M_{ij})_{i \in [1..n], j \in [1..n]}$ of $x^* \cdot y^*$

```

1: if  $n = 1$  then
2:    $M \leftarrow [x_1 \cdot y_1]$ 
3: else
4:    $\mathbf{X}^{(1)} \leftarrow (x_1, \dots, x_{n/2}), \mathbf{X}^{(2)} \leftarrow (x_{n/2+1}, \dots, x_n)$ 
5:    $\mathbf{Y}^{(1)} \leftarrow (y_1, \dots, y_{n/2}), \mathbf{Y}^{(2)} \leftarrow (y_{n/2+1}, \dots, y_n)$ 
6:    $\mathbf{M}^{(1,1)} \leftarrow \text{MatMult}(\mathbf{X}^{(1)}, \mathbf{Y}^{(1)})$ 
7:    $\mathbf{X}^{(1)} \leftarrow \text{RefreshMasks}(\mathbf{X}^{(1)}), \mathbf{Y}^{(1)} \leftarrow \text{RefreshMasks}(\mathbf{Y}^{(1)})$ 
8:    $\mathbf{M}^{(1,2)} \leftarrow \text{MatMult}(\mathbf{X}^{(1)}, \mathbf{Y}^{(2)})$ 
9:    $\mathbf{M}^{(2,1)} \leftarrow \text{MatMult}(\mathbf{X}^{(2)}, \mathbf{Y}^{(1)})$ 
10:   $\mathbf{X}^{(2)} \leftarrow \text{RefreshMasks}(\mathbf{X}^{(2)}), \mathbf{Y}^{(2)} \leftarrow \text{RefreshMasks}(\mathbf{Y}^{(2)})$ 
11:   $\mathbf{M}^{(2,2)} \leftarrow \text{MatMult}(\mathbf{X}^{(2)}, \mathbf{Y}^{(2)})$ 
12:   $M \leftarrow \begin{bmatrix} \mathbf{M}^{(1,1)} & \mathbf{M}^{(1,2)} \\ \mathbf{M}^{(2,1)} & \mathbf{M}^{(2,2)} \end{bmatrix}$ 
13: end if
14: return  $M$ 

```

As illustrated in Fig. 1, the **MatMult** algorithm is recursive and computes the $n \times n$ matrix in four sub-matrix blocs. This is done by splitting the input shares x_i and y_j in two parts, namely $\mathbf{X}^{(1)} = (x_1, \dots, x_{n/2})$ and $\mathbf{X}^{(2)} = (x_{n/2+1}, \dots, x_n)$, and similarly $\mathbf{Y}^{(1)} = (y_1, \dots, y_{n/2})$ and $\mathbf{Y}^{(2)} = (y_{n/2+1}, \dots, y_n)$, and recursively processing the four sub-matrix blocs corresponding to $\mathbf{X}^{(u)} \times \mathbf{Y}^{(v)}$ for $1 \leq u, v \leq 2$. To prevent the same share x_i from being used twice, each input block $\mathbf{X}^{(u)}$ and $\mathbf{Y}^{(v)}$ is refreshed before being used a second time, using a mask refreshing algorithm. An example of such mask refreshing, hereafter called **RefreshMasks**, can for instance be found in [DDF14]; see Algorithm 5. Since the mask refreshing does not modify the xor of the input $n/2$ -vectors $\mathbf{X}^{(u)}$ and $\mathbf{Y}^{(v)}$, each sub-matrix block $\mathbf{M}^{(u,v)}$ is still a $n^2/4$ -sharing of $(\oplus \mathbf{X}^{(u)}) \cdot (\oplus \mathbf{Y}^{(v)})$, and therefore the output matrix M is still a n^2 -sharing of $x^* \cdot y^*$, as required. Note that without the **RefreshMasks**, we would have $M_{ij} = x_i \cdot y_j$ as in Rivain-Prouff.

Since the **RefreshMask** algorithm has complexity $\mathcal{O}(n^2)$, it is easy to see that the complexity of our **RefSecMult** algorithm is $\mathcal{O}(n^2 \log n)$ (instead of $\mathcal{O}(n^2)$ for the original Rivain-Prouff countermeasure in Alg. 1). Therefore for a circuit of

Algorithm 5 RefreshMasks

Input: $a_1, \dots, a_n$ **Output:** $c_1, \dots, c_n$ such that $\sum_{i=1}^n c_i = \sum_{i=1}^n a_i$

```
1: For  $i = 1$  to  $n$  do  $c_i \leftarrow a_i$ 
2: for  $i = 1$  to  $n$  do do
3:   for  $j = i + 1$  to  $n$  do do
4:      $r \leftarrow \{0, 1\}^k$ 
5:      $c_i \leftarrow c_i + r$ 
6:      $c_j \leftarrow c_j + r$ 
7:   end for
8: end for
9: return  $c_1, \dots, c_n$ 
```

size $|C|$ the complexity is $\mathcal{O}(|C| \cdot n^2 \log n)$, instead of $\mathcal{O}(|C| \cdot n^2)$ for Rivain-Prouff. The following lemma shows the soundness of our RefSecMult countermeasure.

Lemma 1 (Soundness of RefSecMult). *The RefSecMult algorithm, on input n -sharings $(x_i)_{i \in [1..n]}$ and $(y_j)_{j \in [1..n]}$ of x^* and y^* respectively, outputs an n -sharing $(c_i)_{i \in [1..n]}$ of $x^* \cdot y^*$.*

Proof. We prove recursively that the MatMult algorithm, taking as input n -sharings $(x_i)_{i \in [1..n]}$ and $(y_j)_{j \in [1..n]}$ of x^* and y^* respectively, outputs an n^2 -sharing M_{ij} of $x^* \cdot y^*$. The lemma for RefSecMult will follow, since as in Rivain-Prouff the lines 2 to 12 of Alg. 3 transform a n^2 -sharing M_{ij} of $x^* \cdot y^*$ into a n -sharing of $x^* \cdot y^*$.

The property clearly holds for $n = 1$. Assuming that it holds for $n/2$, since the RefreshMasks does not change the xor of the input $n/2$ -vectors $\mathbf{X}^{(u)}$ and $\mathbf{Y}^{(v)}$, each sub-matrix block $\mathbf{M}^{(u,v)}$ is still an $n^2/4$ -sharing of $(\oplus \mathbf{X}^{(u)}) \cdot (\oplus \mathbf{Y}^{(v)})$, and therefore the output matrix $\mathbf{M}$ is still an n^2 -sharing of $x^* \cdot y^*$, as required. This proves the lemma. $\square$

Remark 1. The description of our countermeasure requires that n is a power of two, but it is easy to modify the countermeasure to handle any value of n . Namely in Algorithm 4, for odd n it suffices to split the inputs x_i and y_j in two parts of size $(n-1)/2$ and $(n+1)/2$ respectively, instead of $n/2$.

8.2 Security Analysis.

Proven security in the ISW probing model. We prove that our RefSecMult algorithm achieves at least the same level of security of Rivain-Prouff, namely it is secure in the ISW probing model against t probes for $n \geq t+1$ shares. For this we use the refined security model against probing attacks recently introduced in [BBD⁺15], called t -SNI security. This stronger definition of t -SNI security enables to prove that a gadget can be used in a full construction with $n \geq t+1$ shares, instead of $n \geq 2t+1$ for the weaker definition of t -NI security

Fig. 1: The recursive MatMult algorithm, where R represents the RefreshMasks Algorithm, and $\otimes$ represents a recursive call to the MatMult algorithm.

(corresponding to the original ISW security proof). The authors of [BBD⁺15] show that the ISW (and Rivain-Prouff) multiplication gadget does satisfy this stronger t -SNI security definition. They also show that with some additional mask refreshing satisfying the t -SNI property (such as RefreshMasks), the Rivain-Prouff countermeasure for the full AES can be made secure with $n \geq t + 1$ shares.

The following lemma shows that our RefSecMult countermeasure achieves the t -SNI property; we provide the proof in Appendix A. The proof is essentially the same as in [BBD⁺15] for the Rivain-Prouff countermeasure; namely the compression step is the same, and for the matrix step, in the simulation we can assume that all the randoms in RefreshMasks are given to the adversary. The t -SNI security implies that our RefSecMult algorithm is also composable, with $n \geq t + 1$ shares.

Lemma 2 (t -SNI of RefSecMult). *Let $(x_i)_{1 \leq i \leq n}$ and $(y_i)_{1 \leq i \leq n}$ be the input shares of the SecMult operation, and let $(c_i)_{1 \leq i \leq n}$ be the output shares. For any set of t_1 intermediate variables and any subset $|\mathcal{O}| \leq t_2$ of output shares such that $t_1 + t_2 < n$, there exists two subsets I and J of indices with $|I| \leq t_1$ and $|J| \leq t_1$, such that those t_1 intermediate variables as well as the output shares $c_{|\mathcal{O}|}$ can be perfectly simulated from $x_{|I|}$ and $y_{|J|}$.*

Heuristic security against horizontal-DPA attacks. We stress that the previous lemma only proves the security of our countermeasure against t probes for $n \geq t + 1$, so it does not prove that our countermeasure is secure against the horizontal-DPA attacks described in the previous sections, since such attacks use information about n^2 intermediate variables instead of only $n - 1$.

As illustrated in Fig. 1, the main difference between the new RefSecMult algorithm and the original SecMult algorithm (Alg. 1) is that we keep refreshing

the x_i shares and the y_j shares blockwise between the processing of the finite field multiplications $x_i \cdot y_j$. Therefore, as opposed to what happens in SecMult, we never have the same x_i being multiplied by all y_j 's for $1 \leq j \leq n$. Therefore an attacker cannot accumulate information about a specific share x_i , which heuristically prevents the attacks described in this paper.

Acknowledgments. We are very grateful to the anonymous CHES reviewers for pointing a flaw in a previous version of our countermeasure in Section 8.

References

- [BBD⁺15] Gilles Barthe, Sonia Belaïd, François Dupressoir, Pierre-Alain Fouque, and Benjamin Grégoire. Compositional verification of higher-order masking: Application to a verifying masking compiler. Cryptology ePrint Archive, Report 2015/506, 2015. <http://eprint.iacr.org/>.
- [BCPZ16] Alberto Battistello, Jean-Sebastien Coron, Emmanuel Prouff, and Rina Zeitoun. Horizontal side-channel attacks and countermeasures on the ISW masking scheme. Cryptology ePrint Archive, Report 2016/540, 2016. <http://eprint.iacr.org/>. Full version of this paper.
- [BJPW13] Aurélie Bauer, Éliane Jaulmes, Emmanuel Prouff, and Justine Wild. Horizontal and vertical side-channel attacks against secure RSA implementations. In Ed Dawson, editor, *Topics in Cryptology - CT-RSA 2013 - The Cryptographers' Track at the RSA Conference 2013, San Francisco, CA, USA, February 25-March 1, 2013. Proceedings*, volume 7779 of *Lecture Notes in Computer Science*, pages 1–17. Springer, 2013.
- [Bla79] G.R. Blakely. Safeguarding cryptographic keys. In *National Comp. Conf.*, volume 48, pages 313–317, New York, June 1979. AFIPS Press.
- [CCD88] David Chaum, Claude Crépeau, and Ivan Damgård. Multiparty unconditionally secure protocols (extended abstract). In Janos Simon, editor, *Proceedings of the 20th Annual ACM Symposium on Theory of Computing, May 2-4, 1988, Chicago, Illinois, USA*, pages 11–19. ACM, 1988.
- [CFG⁺10] Christophe Clavier, Benoit Feix, Georges Gagnerot, Mylène Roussellet, and Vincent Verneuil. Horizontal correlation analysis on exponentiation. In Miguel Soriano, Sihon Qing, and Javier López, editors, *Information and Communications Security - 12th International Conference, ICICS 2010, Barcelona, Spain, December 15-17, 2010. Proceedings*, volume 6476 of *Lecture Notes in Computer Science*, pages 46–61. Springer, 2010.
- [CJRR99] Suresh Chari, Charanjit S. Jutla, Josyula R. Rao, and Pankaj Rohatgi. Towards sound approaches to counteract power-analysis attacks. In *Advances in Cryptology - CRYPTO '99, 19th Annual International Cryptology Conference, Santa Barbara, California, USA, August 15-19, 1999, Proceedings*, pages 398–412, 1999.
- [DDF14] Alexandre Duc, Stefan Dziembowski, and Sebastian Faust. Unifying leakage models: From probing attacks to noisy leakage. In *Advances in Cryptology - EUROCRYPT 2014 - 33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Copenhagen, Denmark, May 11-15, 2014. Proceedings*, pages 423–440, 2014.

- [DFS15] Alexandre Duc, Sebastian Faust, and François-Xavier Standaert. Making masking security proofs concrete - or how to evaluate the security of any leaking device. In *Advances in Cryptology - EUROCRYPT 2015 - 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Sofia, Bulgaria, April 26-30, 2015, Proceedings, Part I*, pages 401–429, 2015.
- [FRR⁺10] Sebastian Faust, Tal Rabin, Leonid Reyzin, Eran Tromer, and Vinod Vaikuntanathan. Protecting circuits from leakage: the computationally-bounded and noisy cases. In *Advances in Cryptology - EUROCRYPT 2010, 29th Annual International Conference on the Theory and Applications of Cryptographic Techniques, French Riviera, May 30 - June 3, 2010. Proceedings*, pages 135–156, 2010.
- [GHR15] Sylvain Guilley, Annelie Heuser, and Olivier Rioul. A key to success - success exponents for side-channel distinguishers. In Alex Biryukov and Vipul Goyal, editors, *Progress in Cryptology - INDOCRYPT 2015 - 16th International Conference on Cryptology in India, Bangalore, India, December 6-9, 2015, Proceedings*, volume 9462 of *Lecture Notes in Computer Science*, pages 270–290. Springer, 2015.
- [ISW03] Yuval Ishai, Amit Sahai, and David Wagner. Private circuits: Securing hardware against probing attacks. In *Advances in Cryptology - CRYPTO 2003, 23rd Annual International Cryptology Conference, Santa Barbara, California, USA, August 17-21, 2003, Proceedings*, pages 463–481, 2003.
- [ORSW12] Yossef Oren, Mathieu Renauld, François-Xavier Standaert, and Avishai Wool. Algebraic side-channel attacks beyond the hamming weight leakage model. In Emmanuel Prouff and Patrick Schaumont, editors, *Cryptographic Hardware and Embedded Systems - CHES 2012 - 14th International Workshop, Leuven, Belgium, September 9-12, 2012. Proceedings*, volume 7428 of *Lecture Notes in Computer Science*, pages 140–154. Springer, 2012.
- [PR13] Emmanuel Prouff and Matthieu Rivain. Masking against side-channel attacks: A formal security proof. In *Advances in Cryptology - EUROCRYPT 2013, 32nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Athens, Greece, May 26-30, 2013. Proceedings*, pages 142–159, 2013.
- [RP10] Matthieu Rivain and Emmanuel Prouff. Provably secure higher-order masking of AES. In *Cryptographic Hardware and Embedded Systems, CHES 2010, 12th International Workshop, Santa Barbara, CA, USA, August 17-20, 2010. Proceedings*, pages 413–427, 2010.
- [Sha79] Adi Shamir. How to Share a Secret. *Communications of the ACM*, 22(11):612–613, November 1979.
- [SVO⁺10] François-Xavier Standaert, Nicolas Veyrat-Charvillon, Elisabeth Oswald, Benedikt Gierlichs, Marcel Medwed, Markus Kasper, and Stefan Mangard. The world is not enough: Another look on second-order DPA. In Masayuki Abe, editor, *Advances in Cryptology - ASIACRYPT 2010 - 16th International Conference on the Theory and Application of Cryptology and Information Security, Singapore, December 5-9, 2010. Proceedings*, volume 6477 of *Lecture Notes in Computer Science*, pages 112–129. Springer, 2010.
- [VGS14] Nicolas Veyrat-Charvillon, Benoit Gérard, and François-Xavier Standaert. Soft analytical side-channel attacks. In Palash Sarkar and Tetsu Iwata, editors, *Advances in Cryptology - ASIACRYPT 2014 - 20th International Conference on the Theory and Application of Cryptology and Information*

A Proof of Lemma 2

Our proof is essentially the same as in [BBD⁺15]. We construct two sets I and J corresponding to the input shares of x^* and y^* respectively. We denote by M_{ij} the result of the subroutine $\text{MatMult}((x_1, \dots, x_n), (y_1, \dots, y_n))$. From the definition of MatMult and RefreshMasks , it is easy to see that each M_{ij} can be perfectly simulated from x_i and y_j ; more generally any internal variable within MatMult can be perfectly simulated from x_i and/or y_j for some i and j ; for this it suffices to simulate the randoms in RefreshMasks exactly as they are generated in RefreshMasks .

We divide the internal probes in 4 groups. The four groups are processed separately and sequentially, that is we start with all probes in Group 1, and finish with all probes in Group 4.

- Group 1: If M_{ii} is probed, add i to I and J .
- Group 2: If $r_{i,j}$ or $c_{i,j}$ is probed (for any $i \neq j$), add i to I and J .

Note that after the processing of Group 1 and 2 probes, we have $I = J$; we denote by U the common value of I and J after the processing of Group 1 and 2 probes.

- Group 3: If $M_{ij} \oplus r_{i,j}$ is probed: if $i \in U$ or $j \in U$, add $\{i, j\}$ to both I and J .
- Group 4: If M_{ij} is probed (for any $i \neq j$), then add i to I and j to J . If some probe in MatMult requires the knowledge of x_i and/or y_j , add i to I and/or j to J .

We have $|I| \leq t_1$ and $|J| \leq t_1$, since for every probe we add at most one index in I and J . The simulation of probed variables in groups 1 and 4 is straightforward. Note that for $i < j$, the variable r_{ij} is used in all partial sums c_{ik} for $k \geq j$; moreover r_{ij} is used in $r_{ij} \oplus M_{ij}$, which is used in r_{ji} , which is used in all partial sums c_{jk} for $k \geq i$. Therefore if $i \notin U$, then r_{ij} is not probed and does not enter in the computation of any probed c_{ik} ; symmetrically if $j \notin U$, then r_{ji} is not probed and does not enter in the computation of any probed c_{jk} .

For any pair $i < j$, we can now distinguish 4 cases:

- Case 1: $\{i, j\} \in U$. In that case, we can perfectly simulate all variables r_{ij} , M_{ij} , $M_{ij} \oplus r_{i,j}$, M_{ji} and r_{ji} . In particular, we let $r_{ij} \leftarrow \mathbb{F}_{2^k}$, as in the real circuit.
- Case 2: $i \in U$ and $j \notin U$. In that case we simulate $r_{ij} \leftarrow \mathbb{F}_{2^k}$, as in the real circuit. If $M_{ij} \oplus r_{i,j}$ is probed (Group 3), we can also simulate it since $i \in U$ and $j \in J$ by definition of the processing of Group 3 variables.

- Case 3: $i \notin U$ and $j \in U$. In that case r_{ij} has not been probed, nor any variable c_{ik} , since otherwise $i \in U$. Therefore r_{ij} is not used in the computation of any probed variable (except r_{ji} , and possibly $M_{ij} \oplus r_{i,j}$). Therefore we can simulate $r_{ji} \leftarrow \mathbb{F}_{2^k}$; moreover if $M_{ij} \oplus r_{ij}$ is probed, we can perfectly simulate it using $M_{ij} \oplus r_{ij} = M_{ji} \oplus r_{ji}$, since $j \in U$ and $i \in J$ by definition of the processing of Group 3 variables.
- Case 4: $i \notin U$ and $j \notin U$. If $M_{ij} \oplus r_{i,j}$ is probed, since r_{ij} is not probed and does not enter into the computation of any other probed variable, we can perfectly simulate such probe with a random value.

From cases 1, 2 and 3, we obtain that for any $i \neq j$, we can perfectly simulate any variable r_{ij} such that $i \in U$. This implies that we can also perfectly simulate all partial sums c_{ik} when $i \in U$, including the output variables c_i . Finally, all probed variables are perfectly simulated.

We now consider the simulation of the output variables c_i . We must show how to simulate c_i for all $i \in \mathcal{O}$, where $\mathcal{O}$ is an arbitrary subset of $[1, n]$ such that $t_1 + |\mathcal{O}| < n$. For $i \in U$, such variables are already perfectly simulated, as explained above. We now consider the output variables c_i with $i \notin U$. We construct a subset of indices V as follows: for any probed Group 3 variable $M_{ij} \oplus r_{ij}$ such that $i \notin U$ and $j \notin U$ (this corresponds to Case 4), we put j in V if $i \in \mathcal{O}$, otherwise we put i in V . Since we have only considered Group 3 probes, we must have $|U| + |V| \leq t_1$, which implies $|U| + |V| + |\mathcal{O}| < n$. Therefore there exists an index $j^* \in [1, n]$ such that $j^* \notin U \cup V \cup \mathcal{O}$. For any $i \in \mathcal{O}$, we can write:

$$c_i = M_{ii} \oplus \bigoplus_{j \neq i} r_{ij} = r_{i,j^*} \oplus \left(M_{ii} \oplus \bigoplus_{j \neq i, j^*} r_{ij} \right)$$

We claim that neither r_{i,j^*} nor $r_{j^*,i}$ do enter into the computation of any probed variable or other $c_{i'}$ for $i' \in \mathcal{O}$. Namely $i \notin U$ so neither r_{i,j^*} nor any partial sum c_{ik} was probed; similarly $j^* \notin U$ so neither $r_{j^*,i}$ nor any partial sum $c_{j^*,k}$ was probed, and the output c_{j^*} does not have to be simulated since by definition $j^* \notin \mathcal{O}$. Finally if $i < j^*$ then $M_{i,j^*} \oplus r_{i,j^*}$ was not probed since otherwise $j^* \in V$ (since $i \in \mathcal{O}$); similarly if $j^* < i$ then $M_{j^*,i} \oplus r_{j^*,i}$ was not probed since otherwise we would have $j^* \in V$ since $j^* \notin \mathcal{O}$. Therefore, since neither r_{i,j^*} nor $r_{j^*,i}$ are used elsewhere, we can perfectly simulate c_i by generating a random value. This proves the Lemma.