One Fault is All it Needs: Breaking Higher-Order Masking with Persistent Fault Analysis

Jingyu Pan†, Fan Zhang and Kui Ren
Zhejiang University, China
State Key Laboratory of Cryptology, China
Email: {joeypan,fanzhang,kuiren}@zju.edu.cn

Shivam Bhasin
Temasek Laboratories,
Nanyang Technological University, Singapore
Email: sbhasin@ntu.edu.sg

***Abstract*—Persistent fault analysis (PFA) was proposed at CHES 2018 as a novel fault analysis technique. It was shown to completely defeat standard redundancy based countermeasure against fault analysis. In this work, we investigate the security of masking schemes against PFA. We show that with only one fault injection, masking countermeasures can be broken at any masking order. The study is performed on publicly available implementations of masking.**

***Index Terms*—fault attacks, masking, persistent**

I. Introduction

Fault attacks [1] are a type of physical attacks which considers an active attacker capable of disturbing the operation of the target system. Fault attacks have been powerful against standard cryptographic schemes, such as AES, RSA, etc.

Most fault attacks assume a *transient* fault model, where the injected disturbance or fault is temporary, and ideally it affects only one instance of the target function call (eg. one encryption). Some attacks also consider a *permanent* fault model which affects all calls to the target function. Such faults often arise from physical defects in the device. Recently at CHES 2018 [2], a new fault model was highlighted which remains between transient and permanent, called as *persistent* fault. Unlike transient fault, it affects several calls of the target function, however, *persistent* fault is not permanent, and disappears with a device reboot.

A specific fault analysis technique to exploit persistent fault on block ciphers was also developed and called as *Persistent Fault Attack (PFA) [2]*. PFA was shown to break fault countermeasures based on module redundancy and comparison. Masking [3] is one of the most-studied countermeasures against side-channel attacks.

In this work, we investigate the security of some popular masking schemes against PFA. Publicly available implementations are used for the analysis. Our results show that masking countermeasures can be easily broken with PFA. We highlight the main advantage of PFA over other fault attacks. PFA needs only one fault injection as compared to typically one fault per encryption for other fault analysis technique. With one fault injection and multiple encryptions with same fault, PFA can break masking countermeasure at any order. This reduces the practical effort that an attacker should bare to minimum.

The rest of the paper is organized as follows. Section 2 recalls principles of PFA and masking. Section 3 applies PFA on general masking construction. Case study on security of public implementation of masking schemes against PFA is described in Section 4 and final conclusions are drawn in Section 5.

This work was supported in part by the National Natural Science Foundation of China under the grants 61472357,61571063, and by the the Open Fund of State Key Laboratory of Cryptology, China.

† This research was conducted in parts during author's summer visit to Temasek Laboratories, NTU, Singapore.

II. Background

This section recalls general background concepts about PFA and masking.

A. Persistent Fault Analysis (PFA)

PFA was recently introduced as a novel fault analysis technique in CHES 2018 [2]. Unlike other fault attacks which rely on transient or permanent fault model, PFA exploits persistent fault model. As stated earlier, under persistent fault model, the fault affects several consecutive encryptions. The fault, typically, alters a stored constant (like one or more entry in Sbox look-up) in the target algorithm.

For better comprehension, let us take an example of PRESENT cipher where a random nibble fault alters one Sbox element v to v'. In absence of fault, all elements in 4×4 Sbox including v, v' have an expectation of $\mathbb{E}(v) = \mathbb{E}(v') = \frac{1}{16}$. If a persistent fault is injected to change v to v', $\mathbb{E}(v) = 0, \mathbb{E}(v') = \frac{2}{16}$, while all other elements still hold the expectation $\frac{1}{16}$. If in a certain Sbox call in any round, the original output is v, it will be replaced by v', leading to faulty ciphertext. Some encryptions will still be correct as they won't access the element v of the Sbox during the whole encryption. This difference can be detected statistically over a big set of ciphertexts, just by observing the distribution of each nibble, leaking information of the whole last round key k. Fig. 1 illustrates PFA. A fault is injected into the Sbox and turns an element of Sbox from $v = 10$ into $v' = 12$. v, v' are not required to be known to the attacker and can be brute forced. The following statistical tools can be used for key recovery:

1) t_{min}: find the missing value in Sbox table. Then $k = t_{min} \oplus v$;

2) $t \neq t_{min}$: find other values t where $t \neq t_{min}$ and eliminate candidates for k;
3) t_{max}: find the value with with maximal probability $k = t_{max} \oplus v'$.

The attacker needs enough number of ciphertexts to confidently distinguish distribution of t_{min} or t_{max} from others. The minimum number of ciphertexts N can be computed by the classical coupon collector's problem where it needs $N = (2^b - 1) \times (\sum_{i=1}^{(2^b-1)} \frac{1}{i})$, where b is the bit width of x. For PRESENT ($b = 4$) $N \approx 50$, and for AES (b=8) $N \approx 1560$. More details on PFA and its application on redundancy based fault countermeasures can be found in [2].

B. Masking

Masking [3] is the most studied countermeasure against side-channel attacks. The key idea behind masking is to mask the side-channel activity of a sensitive intermediate value in a cryptographic algorithm by mixing it with a random value. Each encryption call requires fresh randomness to totally remove dependency between sensitive value and side-channel activity. Randomness are sometimes updated several times within an encryption to avoid sophisticated attacks like higher-order attacks. Theoretically, masking does not prevent against fault attacks, however, due to randomness involved, the fault analysis can be complicated.

C. Related Works

Masking has come under the scanner of fault attacks in few previous work. Boscher and Handschuh [4] showed that masking does not protect against classical differential fault attacks. While the analysis was a bit more restrictive in terms of the fault model and the number of faults that are required, the key recovery was possible with increased attack effort. A new kind of fault analysis called fault sensitivity analysis (FSA) was shown to break masking by Li et al [5]. FSA used some side-channel information with fault attack to achieve the goal, again with increased effort as compared to unprotected implementation. FSA was further combined with collision attack to enhance its power leading to stronger attack on several countermeasures including masking and threshold implementation [6]. Use of randomness was recommended as a fault countermeasure prerequisite by Lomne et al [7]. Recently in CHES 2018, a special class of fault attack called statistical ineffective fault attack (SIFA [8]) were used to target and break masking countermeasure at any masking order. SIFA requires several ineffective fault injection to statistically determine the key. In this work, we assess the security of several public implementations of masking countermeasure under PFA. As shown later, PFA on masking requires only one fault injection and breaks masking at any order d.

III. PFA on Masking

A. Fault Model

We follow the general PFA threat model that is:

1) The adversary can inject the persistent fault in some cipher constant (or look-up tables) before the encryption process. The (serialized) implementation of block cipher uses one look-up table for all words (bytes or nibbles) and all rounds.
2) The adversary is able to collect multiple ciphertext outputs with random plaintext (not known).

As required in PFA, fault injection to disturb memory content has been practically demonstrated in range of devices including micro-controllers, FPGA and ASIC [9], [10]. Persistent fault on modern CPU using rowhammer was presented in [2]. In the following, we analyze masking schemes under the said threat model.

B. General Idea

Block ciphers are composed of repetitions of a round function. In PFA, we are mainly concerned about the final round since it's directly related to the ciphertext. The last round of cipher with basic boolean masking can be written as follows:

$$c = (L(S'(x \oplus m) \oplus m^{'}) \oplus k) \oplus L(m^{'}) \tag{1}$$

where c denotes the ciphertext, L denotes some linear function (typically permutations), x denotes the last round input, m and m' denote penultimate and last round masks, respectively. k denotes the round key and $S^{'}(x)$ denotes the masked Sbox which can be calculated as $S'(x) = S(x \oplus m)$. Note that higher order masking can also be included in this analysis, where m can be calculated as $m = m_1 \oplus m_2 \oplus \cdots \oplus m_d$ with d as the masking order.

In our attack model against masking block ciphers, we assume the original (unmasked) Sbox is stored for look-up and a persistent fault is injected. The analysis scheme remains generic as illustrated in the previous section. For each Sbox call in the encryption, ideally a fresh set of masks are drawn and a new masked Sbox S' is computed. This is popularly known as the re-computation method.

If faulty value x' is injected to the i^{th} element of S where the original value $S(i) = x \neq x'$, it leads to the faulty element in the correspondingly calculated masked Sbox where $S'(i \oplus m) = x' \oplus m'$. Consequently, the $x \oplus m$ element is missing in the S' and the $x' \oplus m$ element is doubled. With this knowledge, the adversary can deduce that $c_* = L(x \oplus m') \oplus L(m') \oplus k = L(x) \oplus k$ will not appear in the output ciphertexts. Similarly, $c'_* = L(x') \oplus k$ will be doubled. Since the computation of c_*, c'_* does not depend on either m or m', the attack is equivalent to attacking an unmasked implementation. Even for d order masking, m and m' can be written as the combination of d mask, which eventually gets cancelled out to compute the ciphertext, making the complexity constant even when increasing order d.

IV. Case Studies: Breaking Public Implementation of Masking Schemes with Single Fault

We target a few public implementations of masking in this section. The key advantage of PFA is that it requires only

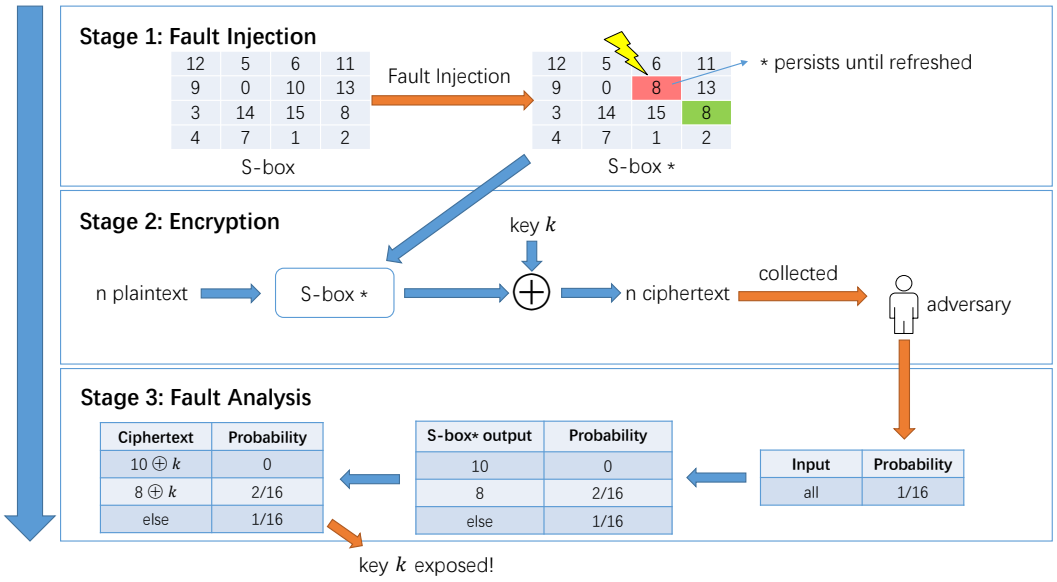

Fig. 1. Overview of Persistent Fault Attack.

one fault injection and multiple encryptions, thus limiting the practical effort of injecting the faults. The required fault model is described before and several works have been practically validated in a range of devices. In the following, we focus on developing the analysis technique with simulations under compatible fault models.

A. Bytewise Masking AES

We apply PFA to the public implementation of bytewise masking available at [11]. It is a typical implementation that follows the general idea illustrated in the previous section. In this case, 6 randomly-generated masks denoted by $m, m^{'}, m_1, m_2, m_3, m_4$ are involved in each encryption, where $m_i, 1 \leq i \leq 4$ correspond to 4 rows of AES, respectively. For the MixColumns operation $MC(col_1, col_2, col_3, col_4)$, 4 output-masks have to be calculated in advance accordingly, denoted by $m_1^{'}, m_2^{'}, m_3^{'}, m_4^{'}$, such that $(m_1^{'}, m_2^{'}, m_3^{'}, m_4^{'}) = MC(m_1, m_2, m_3, m_4)$.

When all 10 masks are generated, a masked AES Sbox denoted by S' is pre-calculated prior to the encryption as:

$$S^{'}_{m,m'}(x) = S(x \oplus m) \oplus m^{'} \tag{2}$$

where S denotes the original AES Sbox in which a persistent fault will be injected, and m and m' are the generated masks. With a persistent fault in S, every S' will contain a fault, irrespective of mask values. One single fault would be enough to reveal the key with the statistical method.

The algorithm of this bytewise masking AES is shown in Algorithm 1. The operations directly affected by the persistent fault are shown in red. Here we apply the t_{min} strategy. We use the available code for our analysis. We injected one persistent fault in S, by randomly changing one Sbox element. The attack was repeated 100 times and the average of all results

Algorithm 1: Bytewise Masking AES

Input: plaintext $p = (p_1, p_2, p_3, p_4)$, where $p_i, 1 \leq i \leq 4$ represent the i^{th} column vector of p, key k
Output: ciphertext c

$rk \leftarrow KeySchedule(k)$
$(m, m^{'}, m_1, m_2, m_3, m_4) \leftarrow^{\$} (\mathbb{F}_{2^8}, \dots)$
$(m_1^{'}, m_2^{'}, m_3^{'}, m_4^{'}) \leftarrow MixColumn(m_1, m_2, m_3, m_4)$
$S' \leftarrow GetMaskedSbox(S, m, m^{'})$ // Eq (2)
$x \leftarrow p \oplus (m, \dots) \oplus rk[0]$
for $i = 1; i < 10; i++$ **do**
  $x \leftarrow S^{'}(x)$
  $x \leftarrow ShiftRows(x)$
  $x \leftarrow x \oplus (m_1, m_2, m_3, m_4) \oplus (m^{'}, \dots)$
  $x \leftarrow MixColumn(x)$
  $x \leftarrow x \oplus rk[i] \oplus (m_1^{'}, m_2^{'}, m_3^{'}, m_4^{'}) \oplus (m, \dots)$
end
$x \leftarrow S^{'}(x)$
$x \leftarrow ShiftRows(x)$
$c \leftarrow x \oplus rk[10] \oplus (m^{'}, \dots)$

is computed. By coupon collector's problem the minimum number of ciphertext required are ≈ 1560. In the experiments we found that with 1500 ciphertexts the attacker has on average less than 2 key byte candidates to test and a unique key with little over 2000 ciphertexts. The analysis remains exactly the same to recover all the bytes independently from same set of ciphertext, thus revealing the last round key and eventually the master key.

B. Coron's Higher-order Masking of Look-up Tables [12]

In Eurocrypt 2014, Coron presented a method to securely compute look-up tables in a block cipher, secure at any order d [12]. This scheme is an ideal target for PFA as it uses look-up tables by design, which is vulnerable to persistent faults. We

target the publicly available implementation of AES protected with this scheme, provided by the author [13].

The key feature of Coron's countermeasure [12] is table recomputation. It uses independent masks with additional refresh of the masks between every successive shift of the input. One can view every line u of the randomized table as a n-dimensional vector of elements in $\{0,1\}^k$, and for all inputs $u \in \{0,1\}^k$:

$$T(u) = (s_{u,1}, s_{u,2}, \ldots, s_{u,n})$$

where initially each vector $T(u)$ is a n-boolean sharing of the value $S(u \oplus x_1)$. The vectors $T(u)$ of the randomized table are then progressively shifted for all $u \in \{0,1\}^k$, first by x_2 and so on until x_{n-1}. Eventually the evaluation of $T(x_n)$ gives a vector of n output shares that corresponds to $S(x)$.

To refresh the masks between successive shifts one can generate a random n-sharing of 0, that is $a_1, \ldots, a_n \in \{0,1\}^k$ such that $\bigoplus_{i=1}^{n} a_i = 0$, and XOR the vector $T(u)$ with $(a_1, \ldots, a_n)$, independently for every u. More concretely, we can use the **RefreshMasks** procedure in Algorithm 2 from [3], which gives a masking of y as $y = y_1 \oplus \cdots \oplus y_n$ by XORing both y_1 and y_i with $r_i \xleftarrow{\$} \mathbb{F}_{2^k}$, in an iterative manner from $i = 2$ to n, where the original value of y_1 is y. The full description of the procedure of Coron's higher order masking of look-up tables is provided in Algorithm 3.[1]

Algorithm 2: RefreshMasks

Input: shares $(x_i)_i$ satisfying $\bigoplus_i x_i = x$
Output: shares $(x_i^{'})_i$ satisfying $\bigoplus_i x_i^{'} = x$
$(z_0^{'}, z_1^{'}, \ldots, z_d^{'}) \leftarrow (z_0, z_1, \ldots, z_d)$
for $i = 1; i < d+1; i++$ **do**
  $r_i \xleftarrow{\$} \mathbb{F}_{2^k}$
  $z_0^{'} \leftarrow z_0^{'} \oplus r_i$
  $z_i \leftarrow z_i \oplus r_i$
end

Algorithm 3: Coron's Masked Computation of $y = S(x)$

Input: shares $x_1, \ldots, x_n$ such that $\bigoplus_i x_i = x$
Output: shares $y_1, \ldots, y_n$ such that $\bigoplus_i y_i = y = S(x)$
for *all* $u \in \mathbb{F}_{2^k}$ **do**
  $T(u) \leftarrow (S(u), 0, \ldots, 0) \in (\mathbb{F}_{2^k})^n$ // $(T(u)) = S(u)$
end
for $i = 1$ *to* $n-1$ **do**
  for *all* $u \in \mathbb{F}_{2^k}$ **do**
    for $j = 1$ *to* n **do**
      $T^{'}(u)[j] \leftarrow T(u \oplus x_i)[j]$ // $T^{'}(u) \leftarrow T(u \oplus x_i)$
    end
  end
  for *all* $u \in \mathbb{F}_{2^k}$ **do**
    $T(u) \leftarrow RefreshMasks(T^{'}(u))$
    // $\bigoplus(T(u)) = S(u \oplus x_1 \oplus \cdots \oplus x_i)$
  end
end
// $\bigoplus(T(u)) = S(u \oplus x_1 \oplus \cdots \oplus x_{n-1})$ for all $u \in \mathbb{F}_{2^k}$
$(y_1, \ldots, y_n) \leftarrow RefreshMasks(T(x_n))$ // $\bigoplus(T(x_n)) = S(x)$

[1]For simplicity, we assume both the input and output of $S(x)$ are words of k bits.

Algorithm 3 uses two temporary tables T and T' in RAM. Both are generated on the basis of the look-up table $S : \{0,1\}^k \rightarrow \{0,1\}^k$. We show that, however, with as few as one single faulty element in table S, the following masking provides no protection against PFA. The operation marked in red in Algorithm 3 denotes the one directly involving injected persistent fault. It results in a faulty table $S^{'}$, which is same as table S but one element.

The attack is performed on AES implementation available at [13], which follows Algorithm 3. For each attack, a single fault is injected into S, and PFA is applied for $d = 1$. The masking offers no resistance against PFA as it reduces to the generic case presented in Section III-B, where the key recovery remains independent of the mask. This results in the attack similar to unprotected AES with key recovery with around 2000 ciphertexts. The increase in masking order d has no impact on the attack because the combination of d different masks can be reduced to a single equivalent mask as $m = m_1 \oplus m_2 \oplus \cdots \oplus m_d$.

Next, we target other masking schemes which do not directly use the Sbox and thus making the analysis more complicated, yet possible.

C. Rivain and Prouff's Masking [3]

In CHES 2010, Rivain and Prouff [3] proposed an efficient method to mask the AES Sbox processing at any order. Specifically, the authors use the algebraic structure of the AES Sbox, which is the composition of an affine function over $\mathbb{F}_2^8$ with the power function $x \mapsto x^{254}$ over $\mathbb{F}_{256}$, and they showed that it can be expressed as a sequence of operations involving a few linear functions over $\mathbb{F}_2^8$, which is easy to mask, and four multiplications over $\mathbb{F}_{256}$. If this computation is performed completely on the fly without any look-up tables, PFA does not apply in principle.

Now, we look at the public implementation of this scheme available at [13]. Let's focus on the Sbox masking part, where component affine transformation is realized through table look-up [13]. The additive part of the affine transformation is $0x63$, thus it can be checked that:

$$Af(x_0) \oplus \cdots \oplus Af(x_d) = \begin{cases} Af(x) & \text{if } d \text{ is even,} \\ Af(x) \oplus 0x63 & \text{if } d \text{ is odd,} \end{cases} \quad (3)$$

where $x = x_0 \oplus x_1 \oplus \cdots \oplus x_d$, for a d order masking. The vulnerable table look-up is highlighted in red in Algorithm 4, which we target by PFA.

Algorithm 4: Rivain and Prouff's secure AES Sbox

Input: shares x_i satisfying $\bigoplus_i x_i = x$
Output: shares y_i satisfying $\bigoplus_i y_i = y = S(x)$
$(y_0, \ldots, y_d,) \leftarrow Exp254(x_0, \ldots, x_d)$
for $i = 0; i \leq d; i++$ **do**
  $y_i \leftarrow Af(y_i)$
end
if $d \, mod \, 2 = 1$ **then**
  $y_0 \leftarrow y_0 \oplus 0x63$
end

However, we need to update the strategy of PFA to target this implementation. Recall that the main idea of PFA is to make a distinct disturbance, which is predictable or observable for the adversary, on the distribution of the output. The previous cases are ideally vulnerable to PFA since the output of the target function (Sbox) is linearly dependent on one single look-up of a permutation table, thus it's rather easy to produce distinguishable and predictable faulty outputs with one single fault. When multiple look-up operations are involved in the target function, as in the case of Rivain-Prouff's Sbox, we show that the output is still distinguishable and predictable with one random fault injection for any masking order, to allow PFA.

Consider a random variable $r(v, v^*, \delta) \in \{0,1\}^b, b \in \mathbb{N}^+$ whose probability is

$$Pr(r = k) = \begin{cases} \frac{1}{2^b} + \delta & k = v^*, \\ \frac{1}{2^b} - \delta & k = v, \\ \frac{1}{2^b} & \text{else}, \end{cases} \tag{4}$$

where $v, v^* \in \{0,1\}^b$ and $0 < \delta \leq \frac{1}{2^b}$. Therefore for independent $r_0(v, v^*, \delta)$ and $r_1(v, v^*, \Delta)$, we have

$$Pr(r_0 \oplus r_1 = k) = \begin{cases} \frac{1}{2^b} + 2\delta\Delta & k = 0, \\ \frac{1}{2^b} - 2\delta\Delta & k = v \oplus v^*, \\ \frac{1}{2^b} & \text{else}. \end{cases} \tag{5}$$

So $r_0(v, v^*, \delta) \oplus r_1(v, v^*, \Delta)$ is equivalent to $r(v \oplus v^*, 0, 2\delta\Delta)$. Similarly we can show that $r_0(v, v^*, \delta) \oplus r_2(v \oplus v^*, 0, \Delta)$ is equivalent to $r(v, v^*, 2\delta\Delta)$.

With one persistent random fault injection into the Af table, when the random input x is under uniform distribution, the output of the faulted table $Af'(x)$ is equivalent to the random variable r above as $r(v, v^*, \frac{1}{2^8})$, where v denotes the original value of the element where the fault is injected, and v^* denotes the faulty value.

For masking order $d = 1$, by Equation (5), we have

$$Pr(Af'(x_0) \oplus Af'(x_1) = k) = \begin{cases} \frac{1}{2^8} + 2 \times (\frac{1}{256})^2 & k = 0, \\ \frac{1}{2^8} - 2 \times (\frac{1}{256})^2 & k = v \oplus v^*, \\ \frac{1}{2^8} & \text{else}, \end{cases} \tag{6}$$

which is equivalent to $r(v \oplus v^*, 0, 2 \times (\frac{1}{256})^2)$. The bias is much lower as compared to previous cases, requiring more samples for the attack.

For any odd masking order d, we can decompose $\bigoplus_{i=0}^{d} Af'(x_i) = \bigoplus_{i=0}^{\frac{d-1}{2}} (Af'(2i) \oplus Af'(2i+1))$ to $\frac{d+1}{2}$ pairs of independent outputs of Af'. Each pair is equivalent to $r(v \oplus v^*, 0, 2 \times (\frac{1}{256})^2)$. By applying Equation 5 $\frac{d+1}{2}$ times, we have $\bigoplus_{i=0}^{d} Af'(x_i)$ is equivalent to $r(v \oplus v^*, 0, 2^d \times (\frac{1}{256})^{d+1}) = r(v \oplus v^*, 0, 2^{-7d-8})$. For any even masking order d, we consider it as a combination of the $d-1$ order masking and $Af'(x_d)$, whose probability should be the same with $r_{(d-1)}(v \oplus v^*, 0, 2^{d-1} \times (\frac{1}{256})^d) \oplus r_d(v, v^*, \frac{1}{2^8})$ which is equivalent to $r(v, v^*, 2^d \times (\frac{1}{256})^{d+1}) = r(v, v^*, 2^{-7d-8})$. In Fig. 2, we apply this strategy to the public implementation of [13], where key k can be extracted with both t_{max} and t_{min} strategy, when d is odd.

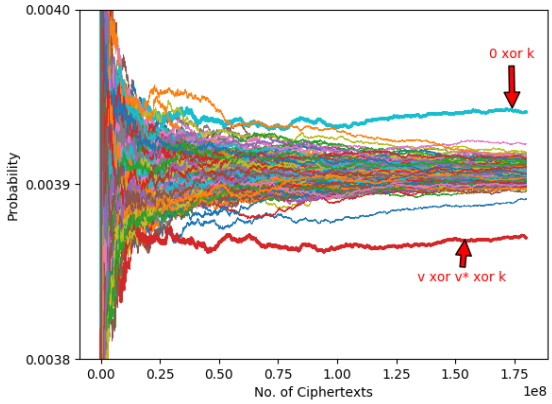

Fig. 2. Key Extraction for Rivain and Prouff scheme [13] with $d = 1$.

However, since $\delta = 2^{-7d-8}$, it decreases exponentially as masking order d increases, and thus more ciphertexts are required to perform PFA. In order to make an estimation of the number of ciphertexts required with higher masking order d, we study the case of AES. For each ciphertext byte, it has the probability of $\frac{1}{256}$ of appearing, so with n ciphertexts, the total number c of its appearance obeys binomial distribution as $c \sim \mathcal{B}(n, p)$, where $p = \frac{1}{256}$. Therefore the variance of $\frac{c}{n}$ is $\frac{p(1-p)}{n}$, and by central limit theorem, $\frac{c}{n}$ approximately follows normal distribution $\mathcal{N}(p, \frac{p(1-p)}{n})$. To perform PFA successfully, we need $\frac{\sqrt{\frac{p(1-p)}{n}}}{2^{-7d-8}} \propto$ constant. Therefore we have $n \propto 2^{14d}$, which means n grows exponentially as d increases.

D. Software Threshold [14]

Sasdrich *et al.* [14] extended the widely used threshold implementation (TI [15]) for software targets. They use PRESENT cipher as a case study, showing a first-order secure implementation. Interested readers can refer to [14] for details on software TI implementation of PRESENT. As public source code is not available, we implemented it in *C* language.

We implemented Algorithm 5. It uses a look-up table:

$$T(x^i, x^j) = A''(f_{\mathcal{Q}_{12}}(A(x^i), A(x^j))$$

which is composed of 256 elements of 4 bits. Targeting at T is not optimal, as each element stands a much less chance of being accessed in the process of encryption. Instead, we target at the smaller look-up table A''' : 8FDACB9E43160752 which is an affine permutation of 4-bit elements and already marked in red in Algorithm 5.

Intuitively, one single fault seems insufficient for PFA since each access of the faulted table is relevant to only one share of all three. However, we can use the same model with the Rivain-Prouff's AES Sbox to estimate the probability distribution of the final output of threshold implementation.

For example, a faulty value 0 is injected into the first element of table A''', whose original value is 8. This injects

Algorithm 5: First-Order Threshold Implementation of PRESENT

Input: $\bar{x} = (x^1, x^2, x^3)$: shared plaintext
k: cipher key
Output: $\bar{y} = (y^1, y^2, y^3)$: shared ciphertext

$rk \leftarrow KeySchedule(k)$
for $i = 1; i \leq 31; i++$ **do**
  $x^1 \leftarrow x^1 \oplus rk[i]$
  $t^3 \leftarrow T(x^1, x^2)$
  $t^2 \leftarrow T(x^3, x^1)$
  $t^1 \leftarrow T(x^2, x^3)$
  $t^3 \leftarrow A'''(t^3)$
  $t^2 \leftarrow A'''(t^2)$
  $t^1 \leftarrow A'''(t^1)$
  $x^3 \leftarrow T(t^1, t^2)$
  $x^2 \leftarrow T(t^3, t^1)$
  $x^1 \leftarrow T(t^2, t^3)$
  $x^1 \leftarrow P(x^1)$
  $x^2 \leftarrow P(x^2)$
  $x^3 \leftarrow P(x^3)$
end
$y^1 \leftarrow x^1 \oplus rk[32]$
$y^2 \leftarrow x^2$
$y^3 \leftarrow x^3$

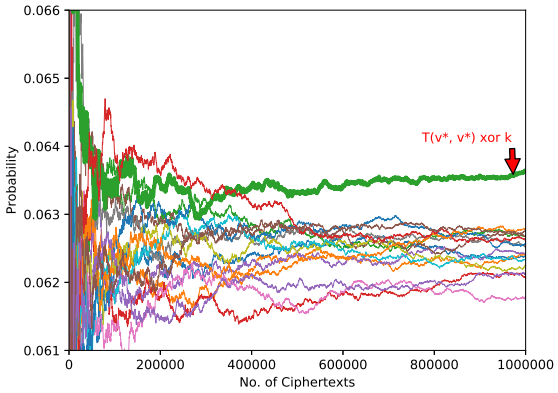

Fig. 3. Key Extraction with t_{max} strategy on Software TI [14].

a bias in the input of T (see Algorithm 5). While an input of 8 will never arrive, input 0 is doubled. In this condition, the probability distribution of the outputs of function T is biased as well. Let T' denote the biased T. The truth table of T' shows probability of 6 being the output $\frac{9}{256}$, and the probability of 12 is $\frac{23}{256}$, while all the others have probabilities that are much closer or equal to $\frac{16}{256}$.

We can use the same analysis model in the Rivain-Prouff's case and calculate the probability distribution of $T'(x_0, x^1) \oplus T'(x_2, x^0) \oplus T'(x_1, x^2)$. Note that for any fault injection with a random fault f, $T(v^*, v^*)$ will have maximal probability to appear at output of T'. Correspondingly, with v denoting the original value where the fault is injected, $T(v, v)$ will always be the one with minimal probability. Therefore, either t_{max} or t_{min} strategy can be applied to extract the key k. In Fig. 3, we show how t_{max} strategy can be applied to recover the key with less than $400K$ ciphertexts.

V. Conclusions

PFA was recently introduced as a novel fault attack. In this work, we show that one persistent fault is enough to break masking at any masking order d. This is validated on public implementations. To conclude, the main advantage of PFA over other fault analysis is that, PFA needs only one fault injection which could last for multiple encryptions, bringing the practical effort of injecting a fault to bare minimum. While avoiding usage of look-up tables completely can prevent PFA, it cannot be a practical solution. This motivates research for novel countermeasures against PFA. Application of PFA to fault detection enhanced masking or other combined countermeasure is another interesting direction.

References

[1] H. Bar-El, H. Choukri, D. Naccache, M. Tunstall, and C. Whelan, "The sorcerer's apprentice guide to fault attacks," *Proceedings of the IEEE*, vol. 94, no. 2, pp. 370–382, 2006.

[2] F. Zhang, X. Lou, X. Zhao, B. Shivam, W. He, R. Ding, S. Qureshi, and K. Ren, "Persistent Fault Analysis on Block Ciphers," in *IACR Transactions on Cryptographic Hardware and Embedded Systems*, vol. 2018, no. 3, 2018, pp. 150–172.

[3] M. Rivain and E. Prouff, "Provably secure higher-order masking of aes," in *International Workshop on Cryptographic Hardware and Embedded Systems*. Springer, 2010, pp. 413–427.

[4] A. Boscher and H. Handschuh, "Masking does not protect against differential fault attacks," in *Fault Diagnosis and Tolerance in Cryptography, 2008. FDTC'08. 5th Workshop on*. IEEE, 2008, pp. 35–40.

[5] Y. Li, K. Sakiyama, S. Gomisawa, T. Fukunaga, J. Takahashi, and K. Ohta, "Fault sensitivity analysis," in *International Workshop on Cryptographic Hardware and Embedded Systems*. Springer, 2010, pp. 320–334.

[6] A. Moradi, O. Mischke, C. Paar, Y. Li, K. Ohta, and K. Sakiyama, "On the power of fault sensitivity analysis and collision side-channel attacks in a combined setting," in *International Workshop on Cryptographic Hardware and Embedded Systems*. Springer, 2011, pp. 292–311.

[7] V. Lomné, T. Roche, and A. Thillard, "On the need of randomness in fault attack countermeasures-application to aes," in *Fault Diagnosis and Tolerance in Cryptography (FDTC), 2012 Workshop on*. IEEE, 2012, pp. 85–94.

[8] C. Dobraunig, M. Eichlseder, H. Gross, S. Mangard, F. Mendel, and R. Primas, "Statistical ineffective fault attacks on masked aes with fault countermeasures," Cryptology ePrint Archive, Report 2018/357, 2018. https://eprint. iacr. org/2018/357, Tech. Rep.

[9] B. Selmke, S. Brummer, J. Heyszl, and G. Sigl, "Precise laser fault injections into 90 nm and 45 nm sram-cells," in *International Conference on Smart Card Research and Advanced Applications*. Springer, 2015, pp. 193–205.

[10] C. Roscian, A. Sarafianos, J.-M. Dutertre, and A. Tria, "Fault model analysis of laser-induced faults in sram memory cells," in *Fault Diagnosis and Tolerance in Cryptography (FDTC), 2013 Workshop on*. IEEE, 2013, pp. 89–98.

[11] "Masked-AES-Implementation." [Online]. Available: https://github.com/Secure-Embedded-Systems/Masked-AES-Implementation

[12] J.-S. Coron, "Higher order masking of look-up tables," in *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 2014, pp. 441–458.

[13] "Higher Order Countermeasures for AES and DES." [Online]. Available: https://github.com/coron/htable

[14] P. Sasdrich, R. Bock, and A. Moradi, "Threshold implementation in software," in *International Workshop on Constructive Side-Channel Analysis and Secure Design*. Springer, 2018, pp. 227–244.

[15] S. Nikova, V. Rijmen, and M. Schläffer, "Secure Hardware Implementation of Nonlinear Functions," *J. Cryptol.*, pp. 292–321, 2011.

[16] Robust-AES. [Online]. Available: https://github.com/vernamlab/Robust-AES

[17] A. Fernandez-rubio, "Efficient side-channel resistant mpc-based software implementation of the AES," 2017.

[18] S. G. Michael Ben-Or and A. Wigderson, "Completeness theorems for non-cryptographic fault-tolerant distributed computation," in *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing*. ACM, 1988, pp. 1–10.

[19] M. O. R. Rosario Gennaro and T. Rabin, "Simpli
ed vss and fast-track multiparty computations with applications to threshold cryptography," in *Proceedings of the Seventeenth Annual ACM Symposium on Principles of Distributed Computing*. ACM, 1998, pp. 101–110.

[20] T. Roche and E. Prouff, "Higher-order glitch free implementation of the AES using secure multi-party computation protocols," *Journal of Crptographic Engineering*, vol. 2, no. 2, pp. 111–127, 2012.

APPENDIX

Robust-AES is an AES-128 implementation using Shamir's Secret Sharing and Secure Multiparty Computation (SMC), claiming to be resistant to any d^{th} order side-channel attacks and partially resistant to fault injection attacks. The public implementation we target is available at [16].

$$\begin{bmatrix} P(\alpha_0) \\ P(\alpha_1) \\ \vdots \\ P(\alpha_d) \\ \vdots \\ P(\alpha_{n-1}) \end{bmatrix} = \begin{bmatrix} 1 & \alpha_0 & \alpha_0^2 & \cdots & \alpha_0^{n-1} \\ 1 & \alpha_1 & \alpha_1^2 & \cdots & \alpha_1^{n-1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \alpha_d & \alpha_d^2 & \cdots & \alpha_d^{n-1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \alpha_{n-1} & \alpha_{n-1}^2 & \cdots & \alpha_{n-1}^{n-1} \end{bmatrix} \begin{bmatrix} s \\ \alpha_0 \\ \alpha_1 \\ \vdots \\ \alpha_d \\ \vdots \\ \alpha_{n-1} \end{bmatrix} \tag{7}$$

Shamir's Secret Sharing masking scheme can be described by Equation 7, where s is a sensitive variable that are split by a d-degree polynomial $P(x) = s + a_1 x + \cdots + a_d x^d$, and the coefficients a_i ($1 \leq i \leq d$) are randomly selected and are meant to remain secret. And the number n of shares $P(\alpha_0), P(\alpha_1), \cdots, P(\alpha_{n-1})$ must be larger than d (at least $d+1$) in order to reconstruct the coefficients of the polynomial. The evaluation points $\alpha_0, \alpha_1, \cdots, \alpha_{n-1}$ can be considered public here. Note that to reconstruct the secret value s, only the first row (denoted as $\lambda_0, \lambda_1, \cdots, \lambda_{n-1}$) of the inverse of the Vandermonde matrix in Equation 7 is required such as $s = \sum_{i=0}^{n-1} P(\alpha_i)\lambda_i$.

Robust-AES made Secure Multiparty Computation on the shares of Shamir's Secret Sharing, such that it causes the same effect on the secret value. Computations of linear operations are quite straightforward here [17]. Consider two secrets s_1, s_2 previously generated by two different polynomials $f(x)$ and $g(x)$ respectively. The addition of the two shares $h(\alpha_i) = f(\alpha_i) + g(\alpha_i)$ of all parties returns the same result as that of the sum of the two shares $s_1 + s_2$. It can be easily checked that such linear property applies to affine transformations such as $k(\alpha_i) = c_1 f(\alpha_i) + c_2$ as well.

However, multiplication of two masked secrets involves a more complicated process and is fundamental to higher order non-linear functions. By [18], [19] and [20], a description of the multiplication algorithm is defined as below:

1) Each player I_i computes $h(\alpha_i)g(\alpha_i)$ locally,
2) I_i generates a degree d polynomial $\mathcal{Q}_i(x)$ such that $\mathcal{Q}_I(0) = h(\alpha_i)$ and sends the value $\mathcal{Q}_i(\alpha_j)$ to player I_j.
3) I_i computes $\mathbf{Q}(\alpha_i) = \sum_{j=0}^{n-1} \lambda_j \mathcal{Q}_j(\alpha_j)$ where $\lambda_0, \cdots, \lambda_{n-1}$ represent the first row of the inverse of Vandermonde matrix in Equation 7.
4) The family $\mathbf{Q}(\alpha_i)_{i=0,1,\cdots,n-1}$ presents a shared representation of s_1, s_2.

In our target public implementation [16], such multiplication is implemented as the sum of two independent table look-ups. We target this operation since it's the very last non-linear operation in the last round of Robust-AES and involves accessing constants stored in memory, which makes it fulfill the conditions of PFA naturally and the coherent analysis strategy can be applied here for key extraction.

Recall that in Rivain-Prouff's AES, each look-up of Af in Algorithm 4 with one fault bias represents a random variable $r(v, v^*, \delta)$ whose probability is described as Equation (4), where $\delta = \frac{1}{2^b}$ and $b = 8$ since only one of all 2^b elements is tampered.

Similarly, in Robust-AES, multiplication is based on two independent pre-calculated tables, each of which is composed of $16 \times 256 = 4096$ bytes. Therefore $\delta = \frac{1}{4096}$ for each table of SMC multiplication with a single fault bias. For masking order d, knowledge of $d + 1$ shares is required for secret reconstruction. In [16], however, more than $d + 1$ shares are involved in encryption, yet for simplicity, we assume $d + 1$ in the following analysis, which meets the requirement of the minimal number of shares. For each share, the multiplication operation requires $2 \times (d+1)$ look-ups at the final stage, $d+1$ look-ups in either target table respectively. And the secret is a weighted sum of all required $d + 1$ shares in Galois Field, where the weights are public. Therefore, the output of Robust-AES has $\delta' = \delta^{2\times(d+1)^2}$. By comparison to Rivain-Prouff's AES, we estimate that, for $d = 1$, the time complexity is about 2^{204}. And the relationship between the number of ciphertexts required n and the protection order d is $n \propto 2^{48\times(d+1)^2}$.