

Blind-Folded: Simple Power Analysis Attacks using Data with a Single Trace and no Training

Xunyu Hu, Quentin L. Meunier and Emmanuelle Encrenaz

Sorbonne Université, LIP6, CNRS, 4 Place Jussieu 75005 Paris

{xunyu.hu,quentin.meunier,emmanuelle.encrenaz}@lip6.fr

Abstract. Side-Channel Attacks target the recovery of key material in cryptographic implementations by measuring physical quantities such as power consumption during the execution of a program. Simple Power Attacks consist in deducing secret information from a trace using a single or a few samples, as opposed to differential attacks which require many traces. Software cryptographic implementations now all contain a data-independent execution path, but often do not consider variations in power consumption associated to data. In this work, we show that a technique commonly used to select a value from different possible values in a control-independent way leads to significant power differences depending on the value selected. This difference is actually so important that a single sample can be considered for attacking one condition, and no training on other traces is required. We exploit this finding to propose the first single-trace attack without any knowledge gained on previous executions, using trace folding. We target the two modular exponentiation implementations in Libcrypt, getting respectively 100% and 99.98% of correct bits in average on 30 executions using 2,048-bit exponents. We also use this technique to attack the scalar multiplication in ECDSA, successfully recovering all secret nonces on 1,000 executions. Finally, the insights we gained from this work allow us to show that a proposed counter-measure from the literature for performing the safe loading of precomputed operands in the context of windowed implementations can be attacked as well.

Keywords: Simple Power Attack · Modular Exponentiation · ECDSA · Constant-Time Implementation · Side-Channel Attacks

1 Introduction

Side-Channel Attacks (SCA) target the recovery of key material in cryptographic implementations by measuring physical quantities such as power consumption during the execution of the program.

There are two major deterministic causes which influence the power consumption: the instruction executed by a program and the data manipulated by these instructions. Instructions have a major role in the power dissipated. In fact, looking at a power trace, it is possible to determine which path of instructions the program has followed, leading to the so called Simple Power Attacks (SPA). For this reason, sensitive programs such as cryptographic primitives are now always designed to have an instruction execution trace which is always the same, and in particular independent from the inputs.

Data, on the other hand, have a slighter effect on power consumption, and exploiting this power difference often requires to capture a lot of traces and perform differential attacks known as Differential Power Analysis attacks (DPA) such as the Correlation Power Analysis attack (CPA). In this article, we show that it is possible to use data power consumption to perform SPA using single sample values, taking advantage of the variation

in the Hamming Weight of the data. These attacks allow to recover some secret values based on a single trace. This article makes three contributions:

- We identified that a common way of implementing constant times cryptographic algorithms leads to a significant weakness. More precisely, the masking technique used to assign one of two values to a variable in a control-independent way leads to significant power differences depending on the chosen value. We additionally identified an inherent weakness specific to elliptic curve cryptography, which is due to the point at infinity.
- We show on two widespread cryptographic implementations how to exploit this data-related power consumption to perform a SPA: on the modular exponentiation in `Libgcrypt` and on a secure elliptic curve implementation for performing a signature. The attack requires a single trace to recover directly the full secret, without any training or profiling phase. The only prerequisite is the identification of the loop performing the computation of the masking technique.
- Using the weaknesses identified in the proposed attacks, we show how to break a counter-measure proposed by Saito *et al.* [SIUH22] to safely load a precomputed operand in a windowed modular exponentiation.

The rest of the article is organized as follows: section 2 describes some background related to modular exponentiation and its usage in RSA, and to ECDSA. These two algorithms are used in the following as attack targets. Section 3 presents some works related to power attacks on cryptographic implementations, and highlights differences of our work with previous works, especially regarding the attack hypotheses. Section 4 presents our attack on the modular exponentiation, and section 5 our attack on ECDSA. In section 6, we show that the proposed operand loading process from [SIUH22] does not prevent the exponent recovery. Finally, section 7 concludes and gives some perspectives for future works.

2 Background

2.1 Modular Exponentiation

2.1.1 Usage in RSA

Modular exponentiation is a critical operation in a cryptographic context, as it is often applied to secret data. More precisely, the RSA asymmetric encryption scheme uses modular exponentiation with the secret key as exponent in two contexts: for deciphering a message, and for signing a message.

In the textbook RSA scheme, a private key is made of values (p, q, d) , a public key is (e, N) : p and q are two large prime numbers and $N = pq$, $d = e^{-1} \bmod (p-1)(q-1)$ where e is a small value and e has to be relatively prime with $(p-1)(q-1)$. Signature and verification of a message m are expressed as:

$$\begin{aligned} s &= m^d \bmod N \\ m &= s^e \bmod N \end{aligned}$$

in which s is the signature. RSA is often implemented using the so-called CRT-RSA scheme for efficiency reasons. In this scheme, the owner of the private key computes:

$$\begin{aligned} d_p &= d \bmod (p-1) \\ d_q &= d \bmod (q-1) \\ q_p &= q^{-1} \bmod p \end{aligned}$$

The signing operates as follows:

$$\begin{aligned} s_p &= m^{d_p} \pmod{p} \\ s_q &= m^{d_q} \pmod{q} \\ s &= (s_p - s_q).q_p.q + s_q \pmod{N} \end{aligned}$$

We can notice that the secret exponent d is no longer used as exponent; however, recovering d_p and d_q from the corresponding exponentiations still allows to recover d : we can then compute $s'_p = m^{d_p} \pmod{N}$, and since we have $s = m^d \pmod{N}$, we can recover:

$$\begin{aligned} p &= \gcd(s - s'_p, N) \\ q &= N/p \\ d &= e^{-1} \pmod{(p-1)(q-1)} \end{aligned}$$

Using the notation $a \equiv b \pmod{n}$ for the congruence of a and b modulo n :

$$\begin{aligned} \text{Since } s &= m^d \pmod{N} \\ \text{we have } s &\equiv m^d \pmod{p} \\ \exists k \text{ such that } s &\equiv m^{d_p+k(p-1)} \pmod{p} \\ &\equiv m^{d_p} \cdot (m^{(p-1)})^k \pmod{p} \\ \text{Euler's Theorem} &\equiv m^{d_p} \cdot (1)^k \pmod{p} \\ &\equiv m^{d_p} \pmod{p} \\ \exists a \text{ such that } s &= m^{d_p} + a.p \end{aligned} \tag{1}$$

$$\begin{aligned} \text{Similarly } s'_p &= m^{d_p} \pmod{N} \\ s'_p &\equiv m^{d_p} \pmod{p} \\ \exists b \text{ such that } s'_p &= m^{d_p} + b.p \end{aligned} \tag{2}$$

$$\begin{aligned} (1) - (2) &\implies s - s'_p = m^{d_p} + a.p - m^{d_p} - b.p \\ s - s'_p &= (a - b).p \end{aligned}$$

So $(s - s'_p)$ and N have a common factor p which is their greatest common divisor since $N = pq$ and p and q are prime numbers.

In `Libgcrypt`'s RSA implementation, a random protection has been added: a random number r is generated for each calculation and the expression $s_p = m^{d_p+r(p-1)} \pmod{p}$ is computed instead of $s_p = m^{d_p} \pmod{p}$. This does not change the value of s_p , but the exponent of the operation is $d_p + r(p-1)$, which means that the "secret" value we recover is actually $d_p + r(p-1)$, and the same for d_q . However, as detailed in Vergnaud's article [Ver20], when we have the full value of $d_p + r(p-1)$, we can use RSA's own properties to find p , and the same for q . Even if we only find a part of $d_p + r(p-1)$, we can still find the complete key using lattice, as described in [MV19] (section 6).

2.1.2 Implementations

Traditional modular exponentiation implementations follow the algorithm shown in Algorithm 1. In this algorithm, the exponent is traversed bit by bit; when the bit is 0, the current result is squared, and when it is 1, the current result is squared then multiplied. Obviously, recovering the sequence of operations, e.g. via a cache attack [YF14, LGS⁺16] or power traces, allows to recover the exponent.

Algorithm 1 Traditional modular exponentiation. Bold variables indicate large integers.

Input: Base c , exponent $d = (d_{k-1} \dots d_0)_2$, modulus N
Output: $r = c^d \bmod N$

```

function MODULAREXPONENTIATION( $c, d, N$ )
   $r \leftarrow 1$ 
  for  $i$  from  $k - 1$  to  $0$  do
     $r \leftarrow r \times r \bmod N$  # Squaring
    if  $d_i = 1$  then
       $r \leftarrow r \times c \bmod N$  # Multiplication
  return  $r$ 

```

Algorithm 2 Windowed modular exponentiation. Bold variables indicate large integers.

Input: Base c , exponent $d = (d_{k-1} \dots d_0)_2$, modulus N , window size w
Output: $r = c^d \bmod N$

```

function MODULAREXPONENTIATION( $c, d, N, w$ )
   $c_0 \leftarrow 1$ 
  for  $i$  from  $1$  to  $2^{w-1} - 1$  do
     $c_i \leftarrow c_{i-1} \times c \bmod N$  # Precomputing the  $2^{w-1}$  first powers in a table
   $r \leftarrow 1$ 
   $z \leftarrow k - 1$ 
  while  $z \geq 0$  do
     $y \leftarrow \max(z - w + 1, 0)$ 
     $u \leftarrow (d_z \dots d_y)_2$ 
    for  $i$  from  $1$  to  $z - y + 1$  do
       $r \leftarrow r \times r \bmod N$  # Squaring
     $r \leftarrow r \times c_u \bmod N$  # Multiplication
  return  $r$ 

```

To circumvent this problem, recent implementations always perform the multiplication and store the result in a different variable, choosing at the end of the iteration which result to keep.

Another approach, which is computationally more efficient consists in precomputing 2^w values of the base at the beginning, and then in traversing the exponent bits by groups of w , as shown in Algorithm 2. We can see that in this version of the algorithm, the computations are independant from the exponent value. However, since the c_u loaded depends on it, it is possible to use cache timing information to infer what the possible values for u are. To avoid this issue, some side-channel resistant implementations load all the elements c_u , keeping as result the result of the load corresponding to the correct window value.

However, as we will show in section 4, both the multiply always version and the windowed version making all loads are subject to SPA when keeping or discarding the result of a useful or useless operation.

2.2 ECDSA

2.2.1 Principle

Elliptic Curve Digital Signature Algorithm (ECDSA) [JMV01] is a public-key digital signature algorithm that is a variant of DSA. It uses elliptic curve cryptography which relies on an elliptic curve defined over some finite field of integers \mathbb{F}_p , $Curve(a, b, p, G, n) = x^3 + ax + b \bmod p$ where p is a large prime, G is a base point on the curve that can generate a subgroup of a large prime-order n .

The public parameters of ECDSA scheme include a description of $Curve(a, b, p, G,$

n). The private key d is a random positive integer less than n , the corresponding public key is set as the point $H = d * G$. The $*$ here is the multiplication between a scalar and a point of the curve, and $d * G$ is defined by doing d times the elliptic curves's addition of the point G .

A signature of a message m is $Sign(z, d) = (r, s)$, where $z = hash(m)$ for some hash function; $r = xcoord(P)$, where $P = k * G$ with k a random nonce greater than 0 and less than n generated for each signature; $s = k^{-1}(z + dr) \pmod n$ where k^{-1} is the inverse of $k \pmod n$. One can then verify the signature from the public key H by computing $u = s^{-1}z \pmod n$, $v = s^{-1}r \pmod n$, and checking whether $xcoord(u * G + v * H) = r$, meaning $s^{-1}z * G + s^{-1}r * H \pmod n = k * G$.

We can see that once we find any nonce k , we can easily compute the private key $d = r^{-1}(sk - z) \pmod n$.

2.2.2 Implementations

Traditional implementations of ECC's multiplication rely on traversing the bits of the scalar, either one by one or using a window of bits for indexing a table containing precomputed multiples of the base point.

Algorithm 3 Fixed-Window ECC Multiplication. Bold variables indicate large integers, and variables in capital indicate EC points.

```

Input: Base Point  $G$ , scalar  $k = (k_{n-1} \dots k_0)_2$ , window size  $w$ 
Output: Point  $P = k * G$ 
function ECCMULTIPLICATION( $G, k, w$ )
   $C_0 \leftarrow O$            #  $O$  represents the point at infinity, neutral element for the addition
  for  $i$  from 1 to  $2^{w-1}$  do
     $C_i \leftarrow C_{i-1} + G$            # Precomputing the first  $2^{w-1}$  multiples in a table
   $P \leftarrow O$ 
   $z \leftarrow n - 1$ 
  while  $z \geq 0$  do
     $y \leftarrow \max(z - w + 1, 0)$ 
     $u \leftarrow (k_z \dots k_y)_2$ 
    for  $i$  from 1 to  $z - y + 1$  do
       $P \leftarrow P + P$            # Computing  $2^{z-y+1} * P$ 
     $P \leftarrow P + C_u$            # Computing  $(2^{z-y+1} + u) * P$ 
  return  $P$ 

```

Algorithm 3 shows the fixed-window ECC multiplication. As for the modular exponentiation, the lookup can be made constant time and indistinguishable from the loaded value by loading each element in the precomputed table and using a mask technique to only keep the correct value.

Yet again, as we will show in section 5, the recombination of all the loaded value can be efficiently exploited to recover the secret nonce in a single trace without prior training.

3 Related Works and Discussion

Several recent works have targeted the recovery of the secret exponent in RSA and the recovery of the secret nonce in ECDSA, often limited to the recovery of a part of the secret bits.

In 2017, Bernstein *et al.* [BBG⁺17] proposed an attack to recover the full RSA exponent of a non constant-time sliding window implementation of modular exponentiation, based on the fact non constant-time implementations seek to reduce the number of multiplication

by adjusting the window depending on the location of the zeros. Following this work, constant-time sliding window exponentiation was proposed.

In 2019, Weissbart *et al.* [WPB19] presented an attack of ECDSA using different techniques of various accuracy for recovering the secret nonce. The full secret recovery is achieved with a Convolutional Neural Network (CNN), which must be trained with around 500 traces for the complete recovery. This work has been extended in 2020 [WCPB20] in which two different implementations of the scalar multiplication are studied: a baseline using power consumption, and a protected implementation using electromagnetic emission (EM). While the authors manage to break the baseline implementation with a single trace, they make use of a CNN trained with 6400 labeled traces, while not being able to explain the cause of the leakage.

In 2020, Lee *et al.* [LH20a] consider the detection of dummy operations as a multi-label classification problem and propose a deep learning method based on CNN to solve it, using power measures. They target an AES software implementation with dummy loads countermeasure, and show that their method performs well compared to their previously proposed method [LH20b]. However, their multi-label CNN is designed for attacking this specific AES implementation, and the applicability to other cryptographic implementations is unclear. Also using deep learning, Lei *et al.* presented a profiling attack based on VGGNet, a small and deep neural network, performed on a smart-card CRT-RSA implementation, using security countermeasures including masking and time jittering [LLQ⁺20]. An ad-hoc method is applied to extract samples points from traces to perform an effective deep learning profiling attack.

Finally, two works are closely related to the contributions we make in this article. First, Monjur *et al.* describe an attack on ECDSA using a fixed window via EM measures [MYW⁺21]. The principle of their attack is similar to ours, but they target the copy and result of an exclusive-or of the masked data, which is less discriminant than the mask computation we propose, and therefore need to consider several samples. Their attack thus needs a profiling comprising a dozen traces, while giving some incorrect bits.

Second, Saito *et al.* propose a single-trace[SIUH22] attack on the GnuMP implementation of the exponentiation in RSA based on EM measures. They use a deep-learning technique to detect dummy loads, but the approach requires a significant training phase (more than 61 millions traces mentioned in the article), while not giving much insight on the reason of the leakage. A counter-measure is presented in the article, which we think does not prevent the exponent recovery. We will explain how to attack this counter-measure in section 6.

In contrary to most of these works which rely on profiling and learning techniques, the attack we present in this article does not need any specific training. The base idea is very simple, and consists in comparing the power consumption of the two different cases in the recombination phase in constant time implementations. As it turns out, the mask computation, which only depends on data, is sufficient to distinguish clearly the condition using a single power measure, even though the attack was carried out on a cheap device limited to a single sample per cycle.

4 Attacking RSA Modular Exponentiation with a SPA

This section presents our Simple Power Analysis attack on modular exponentiation. The attack targets the recovery of the exponent bits, either bit by bit for the traditional version, or slice by slice for the windowed version. Note that we are not interested here in the recovery of the exponent should any error occur, and such algorithms have already been proposed, e.g. in [SIUH22]. Besides, the attack has a very high precision and typically produce no error.

We target implementations of the `Libcrypt` library (v.1.10.3 released on the 14th

November 2023), compiled for Arm v7 and executed on an Arm Cortex-M4 integrated on a STM32F3 board. The board used is the ChipWhispererPro board from NewAE, measuring 4 samples per cycle for short traces, and one sample per cycle using streaming mode for long traces [OC14].

4.1 Characterisation: Conditional Assignment using Masking

A common way to perform a conditional assignment is to use the so-called “masking” technique, consisting in having two words: one containing only '0' bits, and one containing only '1' bits, depending on the condition. A way to achieve this is presented in algorithm 4. It uses the fact that in this implementation, a large integer is represented as an array of limbs, each limb being encoded on a machine word – 32 bits in our case.

Algorithm 4 Control-independant conditional assignment using masking, used for example in `Libgcrypt`. Bold variables indicate large integers.

<pre> 1: function SETCOND(<i>v</i>, <i>u</i>, <i>c</i>) 2: msk0 ← <i>wzero</i> − <i>c</i> 3: msk1 ← <i>c</i> − <i>wone</i> 4: for <i>i</i> from 0 to <i>nlimbs</i> − 1 do 5: <i>v</i>_{<i>i</i>} ← (msk0 & <i>u</i>_{<i>i</i>}) (msk1 & <i>v</i>_{<i>i</i>}) </pre>	<pre> ▷ if <i>c</i> = 0 <i>v</i> remains unchanged, else <i>v</i> gets <i>u</i> ▷ <i>wzero</i> is a word containing only '0' bits ▷ <i>wone</i> is a word containing the value 1 (0b0...01) ▷ <i>nlimbs</i> is the number of limbs ▷ <i>v</i>_{<i>i</i>} and <i>u</i>_{<i>i</i>} are the <i>i</i>th limbs of <i>v</i> and <i>u</i> </pre>
--	--

We first tried to determine whether the computation of `msk0` and `msk1` leads to significant power consumption differences depending on the value of `c`. For this characterisation, we used the implementation of the `SETCOND` function of the library. Extracting only the masks computation lines 2 and 3 (and not the recombination line 4), we captured 1000 traces with a random value for the condition `c`.




Figure 1: Power consumption of a part of the `SETCOND` function in `Libgcrypt` with 4 samples per cycle. Blue curves correspond to a condition value of 0, while red curves correspond to a condition value of 1.

The results of the captures are shown in Figure 1. On this figure, recorded traces with a condition value of 0 are displayed in blue, while recorded traces with a condition value of 1 are displayed in red. We can clearly see that there are some sample points for which the power consumption distributions for the two values of `c` are disjoint. Thus, measuring the power consumption at one of these this sample points only allows to determine the value of `c`, using a simple threshold to discriminate.

4.2 Attacking the full exponentiation

We used the two versions of the modular exponentiation present in the `Libgcrypt` library: the traditional version based on the traversal of the exponent bits, and the second one using a sliding window. In the traditional version, for each bit of the exponent, the condition passed to the `SETCOND` function is the bit value, as shown in Algorithm 5.

Algorithm 5 Simplified iteration of the main loop of the traditional exponentiation function in `Libgcrypt`. Bold variables indicate large integers.

```

  ▷ rp contains the current result
1: xp ← rp × rp                                     ▷ Squaring
2: xp, rp ← rp, xp                                 ▷ Swapping rp and xp
3: xp ← rp × b                                       ▷ Multiplication
4: rp ← SETCOND(rp, xp, ei)                       ▷ Keeping correct result depending on exponent bit

```

Algorithm 6 Simplified iteration of the main internal loop of the sliding window exponentiation function in `Libgcrypt`, loading all precomputed values. Bold variables indicate large integers.

```

1: rp ← result;                                       ▷ rp contains the current result
2: for i from j to 0 do                               ▷ j is the number of squares
3:   for k from 0 to ( $2^w - 1$ ) do                   ▷ w = 5 is the window size
4:     u ← precomp[k]                               ▷ precomp[k] contains  $base^{2 \times k + 1}$  [pointer copy]
5:     SETCOND(w, u, k = e0)                       ▷ w ← precomp[e0] if k = e0, e0 being the slice
6:                                           value derived from the window value
7:   SETCOND(w, rp, i ≠ 0)                             ▷ w ← rp if i ≠ 0, w ← precomp[e0] if i = 0
8:   xp ← MULMOD(rp, w, mp)                       ▷ Square if i ≠ 0, Multiply if i = 0,
9:                                           mp is the modulus
10:  rp ← xp

```

The implementation for the windowed version is shown in Algorithm 6. In order to minimize the number of multiplications, the considered slice value for the exponent is not directly the window value: the implementation first discards the trailing zeros and processes them in the next iteration. Then, the trailing 1 is removed for the index computation: the obtained slice value **e0** is thus comprised between 0 and 15, and is used for indexing the `precomp` table which contains only odd powers between 1 and 31 (i.e. `precomp[e0] = $base^{2 \times e0 + 1}$`). In order to implement constant time, all elements in the table `precomp` are loaded using iteration variable **k**, and using as condition for the `SETCOND` function that **k** equals **e0**.

As an example, if the window value is 10100_2 , the two trailing zeros are first removed, giving 101_2 . Then, the slice value **e0** is obtained by removing the trailing 1, i.e. $e0 = 10_2 = 2$. The function thus copies `precomp[2] = $base^{2 \times 2 + 1} = base^{10_2}$` into **w**, ignoring other loaded values.

Attacks on both the traditional and the windowed versions comprise the following parts: **1)** Locating the Regions of Interest (ROI) corresponding to where the `SETCOND` function is called; **2)** Determining the Points of Interest (POI) from these ROI; **3)** Determining the operations from the power consumption at the POI; **4)** Reconstructing the key from the sequence of operations.

4.3 Locating the Regions of Interest

The ROI correspond to where the `SETCOND` function is called. Examples of such ROI for the windowed version are shown in Figure 2. For this version, each ROI contains 17

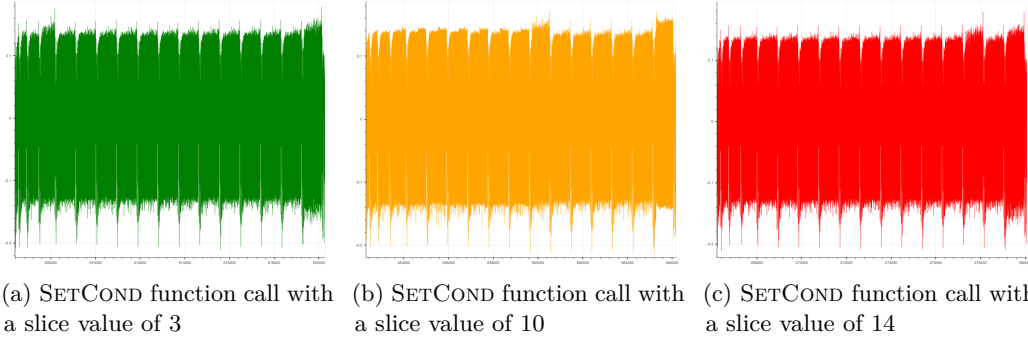


Figure 2: Power consumption for the windowed version where the function SETCOND is called, for three different slice values

sections, and each section contains a POI. The first 16 sections represent the calls to the SETCOND function in the loop and the last one represents the call to the SETCOND function outside of the loop (lines 3-7 in Algorithm 6). As for the traditional version, ROI contain only one section since they comprise a single call to the SETCOND function.

Algorithm 7 Finding all ROI in the trace using a pattern

```

1: function FINDROI(pattern, trace)
2:   roi ← []
3:   for i from 0 to (nb_samples - len(pattern) - 1) do
4:     diff ← 0
5:     for j from 0 to (len(pattern) - 1) do
6:       diff ← diff + abs(pattern[j] - trace[i + j])
7:     if (diff < threshold) then
8:       roi.append(i)
9:   return roi

```

For both versions, the offsets between two consecutive ROI are not the same, as some operations depend on the data values. For the windowed version, it is due to the different number of limbs and the way Libgcrypt cuts the exponent into slices during the computation; while for the traditional version, it is due to some possible required memory allocations depending on the already allocated variable sizes.

In order to locate the different ROI, we use a very simple approach which consists in using one ROI as a power pattern for finding the others. Locating all ROI is then achieved by simply computing, for each subtrace of the same size, the sum of the absolute difference between the subtrace and the pattern, and keeping the ROI if this difference is below some threshold, as illustrated in Algorithm 7. The threshold value is typically obtained empirically: we can manually find a few ROI in the trace, compute this difference and define a threshold. This threshold needs to be larger than the maximum difference observed and smaller than the difference for the other non-ROI subtraces. Finding such a threshold is not complicated since the difference between the results of a ROI and a non-ROI subtrace is very significant. Besides, if we miss some ROI due to the threshold being too low, we can deduce the location of these ROI from the number of cycles between neighbouring ROI we found: if this number is significantly higher than usual, it means that some ROI have been missed, and we can then adapt the threshold consequently.

4.4 Locating the Points of Interest

Focusing on Figure 2, we can observe that for the windowed version, the section for which the iteration number is equal to the slice value (condition $k = e0$ in Algorithm 6) has a different power profile compared to the other, the top part being slightly higher. However, this difference, related to the reading of the precomputed elements, is harder to use than the difference related to the masks' computation that happens at the beginning of each section. This is why we decided to focus on the latter and to use these samples points as POI as they are the most obvious and stable ones.

Since we have only one trace, the approach we use is to “fold” the trace, by overlaying all the ROI found in it. By doing so, the sample points at which the masks' computation occur will exhibit two sets of traces clearly separated, depending on the condition value.

Figure 3 illustrates the overlaid ROI for both versions. In this figure, we notice that two possible areas can be chosen, corresponding to the computation of `msk0` and `msk1`, represented by squared dotted areas. In order to determine the threshold for each POI in the traditional version, we compute the greatest power difference for all ordered power measures at the corresponding sample, and take the middle of the two corresponding samples, as illustrated in Algorithm 8. As each ROI contains a single POI at the same position, locating all the POI is straightforward.

For the windowed version, overlaying the ROI yields 17 distinct POI, and a threshold can be computed for each. Actually, as long as the size of the *base* remains the same, no matter how the exponent changes, the location of the different POI inside the ROI will remain the same. In fact, the width of the sections inside a ROI first increase – as shown in Figure 2 – because when the result of the precomputed power is smaller than the modulus, `Libgcrypt` encodes it only on the required number of limbs. Therefore, the copy time changes according to this number. In our case, we sign a hash from the `hash224` function, which is 224-bit long i.e. 7 limbs. Even when the condition is 0, if the value is not really copied, the number of iterations still depends on the number of limbs as shown lines 4-5 in Algorithm 4. The loaded precomputed results grows until reaching the modulus size, after what it remains the same size. Thus, the POI are always located at the same position in a ROI for a given *base* size, and their location can easily be found. Indeed, each section of a ROI has two parts: one containing the mask generation and one containing the reading of the limbs. In our case, we can observe that the part containing the mask generation takes 89 cycles while the reading takes 13 cycles per limb. Using these values, we can easily calculate the location of all POI in each ROI as long as we know one POI location of one ROI from the trace.

Locating the POI automatically could alternatively be done by looking for the sample in each section of the ROI for which the greatest power difference for all ordered power measures is maximum. However, it still requires to identify the section, therefore bringing little advantage.

Algorithm 8 Finding the POI threshold

```

1: function FINDPOITHRESHOLDS(pm_list)                                ▷ A list of power measures
2:   opm ← sort(pm_list)                                              ▷ Ordered power measures
3:   power_diff ← [(opm[i + 1] - opm[i]) for i in range(len(opm) - 1)]
4:   idx ← argmax(power_diff)                                         ▷ Index of the maximum value
5:   threshold ← opm[idx] + (power_diff[idx] / 2)
6:   return threshold

```




Figure 3: Power consumption at the POI corresponding to the mask computation in the SETCOND function for one trace and for both versions. For the windowed version, the displayed POI are those corresponding to the first of the 17 sections. Traces in orange correspond to a condition of 1, while traces in green correspond to a condition of 0. Dashed areas indicate where the masks are computed.

Algorithm 9 Recovering the sequence of squares and multiplications from the POI

```

1: function RECOVEROPERATIONS(poi, trace)           ▷ Returns the sequence of operations
2:   j ← 0
3:   e0_list ← []                                     ▷ List of exponent slice values
4:   nb_squares ← []  ▷ nb_squares[i] is the number of squares preceding the ith multiply
5:   for i from 0 to len(poi) - 1 do
6:     if trace[poi[i][16]] > threshold16 then           ▷ Multiply detected
7:       for k from 0 to 15 do                               ▷ precomp[k]
8:         if trace[poi[i][k]] < thresholdk then       ▷ k is equal to e0
9:           e0_list.append(k)                               ▷ k contains the exponent slice value
10:          nb_squares.append(j)                            ▷ j squares
11:          j ← 0
12:          break
13:     else
14:       j ← j + 1
15:   return e0_list, nb_squares

```

4.5 Determining the operations from the POI

For the windowed version, each ROI contains 17 POI, corresponding to the 17 calls to the SETCOND function, as shown in Algorithm 6. The last POI indicates if the operation is a Square ($i \neq 0$) or a Multiply ($i = 0$). Therefore, the recovery shown in Algorithm 9 starts by checking the power value at this POI in order to determine the operation (line 6). If the operation is a Multiply, we then check each of the first 16 POI, until we find the section for which the condition is true, meaning that the section number is equal to the exponent slice value (line 8). We then store this slice value and the corresponding number of preceding squares.

For the traditional version, the recovery of the sequence of operations is trivial, as each POI directly gives the operation (Square or Multiply) depending on the condition value, which is directly the exponent bit value.

Algorithm 10 Exponent recovery from the sequence of operations

```

1: function RECOVEREXP(e0_list, nb_squares)
2:   exp  $\leftarrow$  [1]
3:   for i from 0 to (len(e0_list) - 1) do
4:     exp.extend([0]  $\times$  nb_squares[i])            $\triangleright$  Concatenate nb_squares[i] '0' to the right
5:     b  $\leftarrow$  bin( $2 \times$  e0_list[i] + 1)        $\triangleright$  Binary representation
6:     for j from 0 to (len(b) - 1) do
7:       exp[len(exp) - 1 - j]  $\leftarrow$  b[len(b) - 1 - j]
8:   return exp

```

4.6 Reconstructing the secret exponent

The reconstruction of the full secret exponent for the windowed version is shown in Algorithm 10, which traverses the list of exponent slice values and first concatenates the corresponding number of '0' bits (line 4), which is equal to the number of squares, and then adds the exponent value $2 \times e0_list[i] + 1$ (lines 5-7). As an example for the first iteration, if $e0_list[0] = 3$ and $nb_squares[0] = 5$, we first have $exp = (1)_2$, then we concatenate 5 '0' bits to the right ($exp = (100000)_2$), finally the exponent value $(2 \times 3 + 1) = 7 = (111)_2$ is added ($exp = (100111)_2$).

4.7 Experimental Results

In this section, we show how using the power consumption of the mask computation, we can recover very accurately the full exponent with a single trace without any prior knowledge from other traces. For a given target, the only required preliminary step is to identify one ROI in the trace, which can be more or less time-consuming depending on the attack hypotheses – specifically whether the code can be modified or not. For these experiments, we thus suppose that for each trace, we have located manually one ROI, and then we use it as a pattern for the detection of all the other ROI. In practice however, the detection of this ROI of each trace was done using a common pattern in order to speedup the process, since the ROI shape does not change much between two different traces.

The attack achieved in this section comprises the following parts for each trace: **1)** Select one ROI as pattern; **2)** Locate all ROI using the pattern; **3)** Overlay all ROI found in the trace, in order to obtain the folded trace; **4)** For the traditional (resp. windowed) version, locate in the folded trace (resp. in each section of the folded trace), a sample point for which all trace portions split very clearly into two sets, and use it as POI as shown in Figure 3; determine the POI thresholds; **5)** Recover the operations from the power consumption at the POI, and the exponent from the operations.

The attack presented only uses basic knowledge on the traces. However, by making a more detailed analysis, it is possible to improve the attack execution time. For the windowed version, using a 2,048-bit exponent, the full exponentiation takes around 530 millions cycles. Taking as pattern a complete ROI, detecting all the ROI approximately takes two hours. Obviously, the shorter the pattern, the shorter the time required to find all ROI, and we do not necessarily need to use the entire ROI as a pattern if we can find a part of the ROI which is sufficiently different from other parts of the trace. We have thus found two good pattern choices: a 277-sample pattern corresponding to the first section of a ROI, and a 50-sample pattern corresponding to the mask computation part of the 16th section. However, the shorter the pattern, the more sensitive it is to variations in data. Therefore, we need two versions for short patterns, one for each condition value, for matching reliably both cases. In this setup, a ROI is detected when any pattern matches. On the other hand, as mentioned earlier, even though the number of cycles between two consecutive ROI vary, it is comprised in a given range. As a consequence, we can skip a

certain number of samples directly after finding a ROI.

Table 1: Attack results for both the traditional and windowed versions, using 2,048-bit exponents

Version	Traces	Total bits	Pattern samples	Pattern versions	Missed bits	Time/trace	Using as POI	Incorrect bits
Traditional	30	61,440	30	2	0	0.6 s	msk0	14
							msk1	0
Windowed	30	61,440	13711	1	0	5992 s	msk1	0
			227	2	5	77 s		
			50	2	7	11 s		


Experiments were made on 30 traces using 2048-bit exponents for each version. We used as POI both the computation of `msk0` and `msk1` for the traditional version, but only the computation of `msk1` for the windowed version, as the computation of `msk0` is not so clear in the last section. Thus in total, with patterns of different lengths, we obtain 4×30 independent ROI recoveries and 3×30 independent bits recoveries. The results are presented in Table 1. For the traditional version, no bits are missed with the 30-sample pattern corresponding to the ROI, and all bit values are correctly found using the computation of `msk1`, while only 14 bits (in 11 different traces) are incorrectly recovered using the computation of `msk0`. For the windowed version, no bits were missed with a pattern corresponding to a complete ROI (13,711 samples), 5 bits (in 5 different traces) were missed with the 227-sample pattern and 7 bits (in 7 different traces) were missed with the 50-sample pattern. In any case, no incorrect bit value is recovered from a POI, showing the high accuracy of this attack. Besides, the selected short patterns and the skip operation allow to greatly reduce the attack time for the windowed version, from almost two hours for the full pattern without the skip operation, to respectively 77s and 11s.

These results underly the fact that securing the exponentiation against side-channel attacks, even simple ones, is a challenging task. They suggest that it is hopeless to secure the exponentiation with techniques based on constant time or code tricks, but that counter-measures should instead focus on mathematical aspects, or include randomness.

4.8 Limitations

The two hypotheses required for the attack to work are the following: 1) all the ROI should overlay; 2) the power consumption for the mask computation should be significantly different depending on the condition value. While we do not know if these hypotheses will hold for any architecture, it is likely that they should be true often: the first one because the compiled code should follow the source code structure, which does not contain any conditional branch inside the ROI; the second one because the mask computation will always produce a different number of bit flips on the data path depending on the condition value, these bit-flips constituting the most important part in power consumption.

Regarding obtained errors, an example of incorrect bit detection for the traditional version is shown in Figure 4a. We noticed that the incorrect bits always appear in the first three bits of each trace. We can see in Figure 4a that for these three bits, the power is shifted above the others – a phenomenon we already noticed when looking for the ROI. After investigation, it appears that it is because in these iterations, the modulus is not done due to the result size. Therefore, as the preceding operations differ from the rest, they have a different influence on the power consumption from the rest on the `SETCOND` function call. While it is possible to determine a specific threshold for these iterations, we consider that it involves too much knowledge on the trace. Plus, since there are only three of them, it is easy to manually confirm their values.



(a) Incorrect recovered bit from POI (in red) for the traditional version (b) Missed ROI (red) compared to other ROI for the 50-sample pattern for the windowed version

Figure 4: Examples of errors encountered by the attack

We also investigated the missed ROI for the windowed version and noticed that they all occurred at the end of the trace. For these ROI, a small part of the trace seems shifted, leading to a higher difference with the pattern, as illustrated in Figure 4b. While we did not manage to understand the cause of this difference, its impact on the attack remains limited and as we mentioned earlier, all the missed ROI are easy to locate.

Finally, in order to be more robust in the recovered bit values, a slightly more complex approach could be used, considering for example two or more consecutive power values and compute their difference in order to deduce the shape of the correct curve. However, as these cases remained very rare, we did not consider the need to take them into account.

5 Attacking ECDSA with a SPA

5.1 Characterisation and Setup

For ECDSA, we used the `bearssl` library [Por23], which includes some counter measures against side-channel attacks. It is compiled with `gcc 10.2.1` using optimisation level `02`, and using the curve `p256m31`. The ECC multiplication in this library is constant time and uses a fixed 4-bit window over the scalar value, performing the loads of all the values in the table (external loop over `i`), and keeping only the correct value using a masking technique, as shown in algorithm 11. Each value is a point of the curve (which is the multiplication of G by a scalar between 1 and 15), encoded as several 32-bit limbs in order to store the 256-bit coordinates: 9 limbs for the x coordinate and 9 limbs for the y coordinate¹. The internal loop over `j` reads and masks these 18 limbs.

We recorded traces consisting of the 64 iterations, each iteration processing 4 bits of the scalar value. Such a trace requires approximately 6,000,000 processor cycles to complete. As for RSA, we recorded the samples using the ChipWhisperer Pro device, with a STM32F3 target board comprising a Arm Cortex-M4. The number of cycles and samples in a trace imposes here again to use the capture device in streaming mode, which limits the number of samples to a single sample per cycle.

¹A point coordinate is encoded on 9 32-bit words, using 30 bits for the first 8 words, and 16 bits for the last

Algorithm 11 Pseudo-code of the LOOKUPGWIN function in the `bearssl` library

```

▷ idx contains the window value of the scalar bits
▷ Gwin[i][j] contains word j of  $(i+1)*G$ 
▷ EQ(x, y) returns 1 if  $x = y$ , 0 otherwise
1: function LOOKUPGWIN(idx)                                     ▷ returns  $idx * G$ 
2:    $P \leftarrow 0$ 
3:   for i from 0 to 14 do
4:      $mask \leftarrow -EQ(idx, i + 1)$ 
5:     for j from 0 to 17 do
6:        $P[j] \leftarrow mask \& Gwin[i][j]$ 
7:   return  $P$ 

```

5.2 Attacking the masked loads

We first targeted the loads in the precomputed values table, and more specifically the computation of the `mask`. We manually located in the trace the power corresponding to the LOOKUPGWIN function calls, which constitutes our ROI. Since all the ROI are separated with a constant number of samples, the identification of the 64 ROI is straightforward. In Figure 5, we can clearly see that there are 15 sections on each subfigure (again separated by a fixed number of samples), while zooming in reveals that each section containing 18 smaller sections. These are exactly the `for i` and `for j` loops in LOOKUPGWIN.

We can already observe in Figure 5 that when the 4-bit value is equal to $i + 1$ ($mask = -1$), the power consumption of the whole iteration is significantly different. This is because the program actually keeps the precomputed value only in this case, and thus the result of the bitwise AND operation will have a average Hamming Weight value of 15, instead of 0.




Figure 5: Power consumption of one LOOKUPGWIN function call for different window values.

However, as in the case of the exponentiation, the power consumption difference resulting from mask computation is way clearer. By overlaying different ROI, we can observe a sample at the beginning of each section (`for i` iteration) whose value highly depends on the condition value. This sample in each section, that we use as POI, corresponds to the operation: $mask \leftarrow -EQ(idx, i + 1)$.

Figure 6 illustrates the power consumption for all overlaid ROI, at the samples corresponding to the mask computation. In this figure, the top curves correspond to POI in the first section: the dark blue curves when the condition is true which means the window value is 1; the dark orange curves when the condition is false and the window value is different from 0; and the dark green curves when the condition is false and the window value is 0, meaning the condition will be false for all other sections. It is somewhat surprising to see that some sample points allow to discriminate very clearly between a 0 and a non-zero window value; yet regardless, we did not use this information in our attack,




Figure 6: Power consumption associated to the mask computation in ECDSA. The different colors indicate different sections and condition values.

focusing only on the condition result at each iteration. The bottom curves correspond to sections 1 to 14 (from the second to the 15th): the blue curves when the condition is true, i.e. when the iteration matches the window value, the pink curves when the condition is false but has already been true in one of the previous sections of current ROI, the orange curves when the condition is false and has not yet been true in one of the previous sections of current ROI, and the green curves (mixed with the orange ones) correspond to false conditions when the window value is 0, because no condition of any section in this ROI will be true when the value is 0.

From this figure, we can observe that the power consumption in the first section is significantly different from the other sections, for which the measures all coincide. Therefore, we decided to use two thresholds (determined with the technique described in Algorithm 8): one for the first section and one for the others. As the number of iterations is only 64, this allows to have a threshold even for window values which are never reached. As for the window value 1, if it is never reached, all curves at the POI will overlay and the assumption must be made that this value is never obtained.

Algorithm 12 ECDSA Digits recovery using the power measures

```

1: function BITSRECOVERY(trace, poi)
2:   cons1 = [], cons2 = [], secret = ""
3:   for j from 0 to 63 do
4:     cons1.append(trace[poi[j][0]])
5:     for i from 1 to 14 do
6:       cons2.append(trace[poi[j][i]])
7:   threshold2 = find_threshold(cons2)
8:   threshold1 = find_threshold(cons1)
9:   for j from 0 to 63 do
10:    if cons1[j] > threshold1 then
11:      secret += '1'
12:    else
13:      consj = cons2[j × 14 : (j + 1) × 14] ▷ Sublist with the POI of the jth ROI only
14:      if max(consj) > threshold2 - ε then
15:        secret += str(consj.argmax(consj) + 2) ▷ hex character encoding 4 bits
16:      else
17:        secret += '0'
18:   return secret

```

Putting apart the first iteration, we thus decided to look whether one of the measure at the POI was greater than the threshold: if not, the value 0 is deduced, otherwise, the window value taken is the index corresponding to the maximum of the observed values

(Algorithm 12). Finally, since the discrimination using the threshold must only be done between a 0 value and a non-0 value (sets of blue and light green/light orange curves), the threshold can optionally be adapted for the attack to be more robust. In practice, we used an ε value of 0.01, from the observations made in Figure 6.

We captured 1,000 traces and ran the attack on each of these traces: on the $1,000 \times 64$ window values, we recovered correctly all the scalar bits of the nonce k , using Algorithm 12. These results are summarized in table 2.

Table 2: Results for the ECDSA attack for 1,000 traces.

Nb. Traces	Nb. of digits	Nb. of secret nonces	Percentage of recovered digits	Percentage of recovered secret nonces
1,000	64,000	1,000	100%	100%

5.3 Attacking the 0 values




Figure 7: Power consumption for the reads and writes of machine words representing a point of the curve. Red curves indicate the point at infinity, while blue curves indicate another point. The purple line indicates the sample used for the attack.

As seen in section 2.2, the considered implementation iteratively loads precomputed values from a table. Each value consists of several limbs in order to reconstruct the 256-bit coordinates. Due to the nature of the algorithm, when the window value is 0, the result of its multiplication with the base point is the point at infinity, encoded as limbs with value 0 only. Thus, when loading and writing iteratively all of the words of the resulting point, the Hamming Weight of each word will be 0, while it will be 15 on average for other points. This difference in Hamming Weight is actually sufficient to identify when the point at infinity is used. This can actually lead to an attack, since the knowledge of the “zero” values in k on a few traces is sufficient to recover the secret [GRV17, HGS01]. Note that this attack is independent from the masking technique, and targets the copy of the resulting point at the end of the LOOKUPGWIN function.

Here again, by overlaying the ROI corresponding to the copy, two sets of traces appear clearly, as shown in Figure 7. This figure shows that there is a clear difference in power consumption between a window value of 0 (red curves) and a window value different from 0 (blue curves), due to the difference in Hamming Weights.

Once again, we achieved this attack without the knowledge of previous executions, and the thresholds were determined using only the trace itself. We captured 1000 traces and for each trace and each iteration, we guessed if the loaded point was the point at infinity

Table 3: Results for the ECDSA attack consisting in determining the 0 values.

Number of traces	Number of digits	Number of 0 digits incorrectly found	Number of non-0 digits incorrectly found	Attack Accuracy
1,000	64,000	1	0	99.998%

or not based on the power consumption of the selected sample. Similarly to the previous attack, for traces containing no zeros, we used a criteria on the found maximum difference in the ordered power values, and concluded that all windowed values were different from 0. The results, summarized in table 3 show that all digits were correctly identified as zero or non-zero except one.

6 Attacking a Counter-Measure from the Litterature

In their 2022 article, Saito *et al.* [SIUH22] propose the solution recalled in Algorithm 13 to safely load the window value, using two random masks. They conclude that it will mitigate the two vulnerabilities exploited in their proposed attack.

Algorithm 13 Proposed Operand Loading Process from [SIUH22]

Inputs: Precomputed table $(c^0, c^1, \dots, c^{2^v-1})$, window size w , window value b
Output: Multiplication operand $s = c^b$

- 1: **function** LOADOPERAND($(c^0, c^1, \dots, c^{2^v-1})$, w , b)
- 2: $r_{value} \leftarrow \text{GenerateRandom}_l()$ ▷ Generate a l -bit random mask
- 3: $s \leftarrow c^0 \oplus c^1 \oplus \dots \oplus c^{2^v-1} \oplus r_{value}$
- 4: $r_{order} \leftarrow \text{GenerateRandom}_w()$ ▷ Generate a w -bit random mask
- 5: $b_{masked} \leftarrow b \oplus r_{order}$
- 6: **for** i from 0 to 2^{w-1} **do**
- 7: $mask \leftarrow \text{MAKEBITMASK}(i, b_{masked})$ ▷ MAKEBITMASK returns 1...1 if $i = b_{masked}$, else 0
- 8: $s \leftarrow s \oplus (c^{i \oplus r_{order}} \& \neg mask \mid r_{value} \& mask)$
- 9: **return** s


Using the experience acquired with our previous attacks, we show how the exponent value can still be recovered. The attack works by identifying two distinct iterations: the iteration it_0 for which $i = b_{masked}$, and the iteration it_1 for which $i \oplus r_{order} = 0$. We thus have $it_0 = b_{masked}$ and $it_1 = r_{order}$. Using the fact that $b_{masked} = b \oplus r_{order}$, we have $b = it_0 \oplus it_1$.

Identifying $i = b_{masked}$. As we have seen in sections 4 and 5, determining the value of a condition when the latter is used to create a word with only zeros or ones can be done very accurately. Therefore, the MAKEBITMASK function is very likely to exhibit this iteration.

Identifying $i \oplus r_{order} = 0$. In this iteration, the element loaded in the precomputed table is $c^0 = 1$. As we have seen in section 4, in Libgcrypt the precomputed element are encoded on the minimum required number of limbs. Hence in this case, it is sufficient to count the number of iteration during the copy. However, the counter-measure suggests that all precomputed elements are stored on the same number of limbs (64 limbs for 2,048 bits). However, as we have seen in section 5, it is possible to identify precisely when words containing the value 0 are loaded as opposed to words containing roughly 15 or 16 bits with the value 1. In this case, we can then simply count the number of words with value 0 loaded during the copy: the maximum will be reached for c^0 . Note that this also works with a sliding window when the first element is c^1 which is the base c and is a hash with

Table 4

Number of traces	# Matches for both deduced values of r_{order}	# Correct deduced values for r_{order}
1000	239	232



(a) Matching deduced r_{order} values (b) Non-matching deduced r_{order} values (c) Non-matching deduced r_{order} values

Figure 8: Power consumption for the 16 iterations of the micro-benchmarks. The red curve corresponds to the iteration for which $i \oplus r_{\text{order}} = 0$, the green curve for which $i \oplus r_{\text{order}} = 15$

a smaller size, as it is the case in section 4; in our case, the number of "0 words" would then be $64 - 7 = 57$.

Finally, in case the first element is c^1 and the size of c is the same as the modulus, this technique cannot be employed. However, although less accurate, it is possible to look directly at the computation $i \oplus r_{\text{order}}$ and hope that the Hamming Weight of the result gives enough information: it should be minimal with good probability. The found iteration can also be combined with the found iteration for which $i \oplus r_{\text{order}} = 2^{w-1}$ (yielding maximal power), in order to strengthen the confidence in the result: if the two deduced values for r_{order} do not match, we discard the result for this window value.

We designed a micro-benchmark to evaluate the feasibility of this technique using a 4-bit random and made the computation $i \oplus r_{\text{order}}$ for 16 iterations, looking for the maximum and minimum power values, deducing the values of r_{order} and checking their match. We report the results in table 4, which shows the practicability of this approach, considering the least favorable hypotheses: even for this very simple heuristic, out of 1,000 traces, we had a match in 239 cases, 232 of which led to a correct value for r_{order} .

Finally, Figure 8 shows different cases for which the deduced r_{order} values match and do not match: they match when the curve corresponding to $i \oplus r_{\text{order}} = 0$ (red curve) has the maximum power consumption, and when the curve corresponding to $r_{\text{order}} = 15$ (green curve) has the minimum power consumption. Even when the values do not match, we observe that the red and green curves are close to the top and bottom curves respectively.

7 Conclusion

In this article, we showed that power differences associated to variations in data values can be exploited for practical Simple Power Attacks. In particular, we showed that the masking technique often used in constant time implementations to keep one among two possible values leaks its condition in a way that a single sample allows to recover it without error. We illustrated this attacks on the two implementations of the modular exponentiation in `Libgcrypt` and on the scalar ECC multiplication of the `bearssl` library. We additionally

showed that the load of memory words containing either the value 0 or a random value can be clearly identified and also used as an attack vector. Finally, we explained how a proposed counter-measure from the literature for performing secure loads of precomputed values can be attacked. Future work includes the applicability of the proposed attack concepts to other cryptographic domains, as well as designing secure implementations taking into account the presented findings.

References

- [BBG⁺17] Daniel J Bernstein, Joachim Breitner, Daniel Genkin, Leon Groot Bruinderink, Nadia Heninger, Tanja Lange, Christine van Vredendaal, and Yuval Yarom. Sliding right into disaster: Left-to-right sliding windows leak. In *Cryptographic Hardware and Embedded Systems—CHES 2017: 19th International Conference, Taipei, Taiwan, September 25–28, 2017, Proceedings*, pages 555–576. Springer, 2017.
- [GRV17] Dahmun Goudarzi, Matthieu Rivain, and Damien Vergnaud. Lattice attacks against elliptic-curve signatures with blinded scalar multiplication. In *Selected Areas in Cryptography—SAC 2016: 23rd International Conference, St. John’s, NL, Canada, August 10–12, 2016, Revised Selected Papers 23*, pages 120–139. Springer, 2017.
- [HGS01] Nick A Howgrave-Graham and Nigel P. Smart. Lattice attacks on digital signature schemes. *Designs, Codes and Cryptography*, 23:283–290, 2001.
- [JMV01] Don Johnson, Alfred Menezes, and Scott Vanstone. The elliptic curve digital signature algorithm (ecdsa). *International journal of information security*, 1:36–63, 2001.
- [LGS⁺16] Moritz Lipp, Daniel Gruss, Raphael Spreitzer, Clémentine Maurice, and Stefan Mangard. Armageddon: Cache attacks on mobile devices. In *25th USENIX Security Symposium (USENIX Security 16)*, pages 549–564, 2016.
- [LH20a] JongHyeok Lee and Dong-Guk Han. Dlddo: deep learning to detect dummy operations. In *International Conference on Information Security Applications*, pages 73–85. Springer, 2020.
- [LH20b] JongHyeok Lee and Dong-Guk Han. Security analysis on dummy based side-channel countermeasures—case study: Aes with dummy and shuffling. *Applied Soft Computing*, 93:106352, 2020.
- [LLQ⁺20] Qi Lei, Chao Li, Kexin Qiao, Zhe Ma, and Bo Yang. Vgg-based side channel attack on rsa implementation. In *2020 IEEE 19th International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom)*, pages 1157–1161. IEEE, 2020.
- [MV19] Thierry Mefenza and Damien Vergnaud. Cryptanalysis of server-aided rsa protocols with private-key splitting. *The Computer Journal*, 62(8):1194–1213, 2019.
- [MYW⁺21] Alam Monjur, Baki Yilmaz, Frank Werner, Niels Samwel, Alenka Zajic, Daniel Genkin, Yuval Yarom, and Milos Prvulovic. Nonce@ once: a single-trace em side channel attack on several constant-time elliptic curve implementations in mobile platforms. In *IEEE European Symposium on Security and Privacy*, 2021.

- [OC14] Colin O’Flynn and Zhizhang David Chen. Chipwhisperer: An open-source platform for hardware embedded security research. In *COSADE*, 2014.
- [Por23] Thomas Pornin. Bearssl implementation of the ssl/tls protocol. <https://bearssl.org/>, 2023.
- [SIUH22] Kotaro Saito, Akira Ito, Rei Ueno, and Naofumi Homma. One truth prevails: A deep-learning based single-trace power analysis on rsa-crt with windowed exponentiation. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, pages 490–526, 2022.
- [Ver20] Damien Vergnaud. Comment on “efficient and secure outsourcing scheme for rsa decryption in internet of things”. *IEEE Internet of Things Journal*, 7(11):11327–11329, 2020.
- [WCPB20] Léo Weissbart, Łukasz Chmielewski, Stjepan Picek, and Lejla Batina. Systematic side-channel analysis of curve25519 with machine learning. *Journal of Hardware and Systems Security*, 4:314–328, 2020.
- [WPB19] Leo Weissbart, Stjepan Picek, and Lejla Batina. One trace is all it takes: Machine learning-based side-channel attack on eddsa. In *Security, Privacy, and Applied Cryptography Engineering: 9th International Conference, SPACE 2019, Gandhinagar, India, December 3–7, 2019, Proceedings 9*, pages 86–105. Springer, 2019.
- [YF14] Yuval Yarom and Katrina Falkner. Flush+reload: A high resolution, low noise, l3 cache side-channel attack. In *23rd USENIX security symposium (USENIX security 14)*, pages 719–732, 2014.