A Tale of Three Signatures: Practical Attack of ECDSA with wNAF

Gabrielle De Micheli^{1(✉)}, Rémi Piau^{1,2}, and Cécile Pierrot¹

¹ Université de Lorraine, CNRS, Inria, LORIA, F-54000 Nancy, France

gabrielle.de-micheli@inria.fr

² ENS Rennes, Rennes, France

Abstract. Attacking ECDSA with wNAF implementation for the scalar multiplication first requires some side channel analysis to collect information, then lattice based methods to recover the secret key. In this paper, we reinvestigate the construction of the lattice used in one of these methods, the *Extended Hidden Number Problem* (EHNP). We find the secret key with only 3 signatures, thus reaching a known theoretical bound, whereas best previous methods required at least 4 signatures in practice. Given a specific leakage model, our attack is more efficient than previous attacks, and for most cases, has better probability of success. To obtain such results, we perform a detailed analysis of the parameters used in the attack and introduce a preprocessing method which reduces by a factor up to 7 the total time to recover the secret key for some parameters. We perform an error resilience analysis which has never been done before in the setup of EHNP. Our construction find the secret key with a small amount of erroneous traces, up to 2% of false digits, and 4% with a specific type of error.

Keywords: Public key cryptography · ECDSA · Side channel attack · Windowed non-adjacent form · Lattice techniques

1 Introduction

The Elliptic Curve Digital Signature Algorithm (ECDSA) [13], first proposed in 1992 by Scott Vanstone [26], is a standard public key signature protocol widely deployed. ECDSA is used in the latest library TLS 1.3, email standard OpenPGP and smart cards. It is also implemented in the library OpenSSL, and can be found in cryptocurrencies such as Bitcoin, Ethereum and Ripple. It benefits from a high security based on the hardness of the elliptic curve discrete logarithm problem and a fast signing algorithm due to its small key size. Hence, it is recognized as a standard signature algorithm by institutes such as ISO since 1998, ANSI since 1999, and IEEE and NIST since 2000.

The ECDSA signing algorithm requires scalar multiplications of a point P on an elliptic curve by an ephemeral key k . Since this operation is time-consuming and often the most time-consuming part of the protocol, it is necessary to use

an efficient algorithm. The Non Adjacent Form (NAF) and its windowed variant (wNAF) were introduced as an alternative to the binary representation of the nonce k to reduce the execution time of the scalar multiplication. Indeed, the NAF representation does not allow two non-zero digits to be consecutive, thus reducing the Hamming weight of the representation of the scalar. This improves on the time of execution as the latter is dependent on the number of non-zero digits. The wNAF representation is present in implementations such as in Bitcoin, as well as in the libraries Cryptlib, BouncyCastle and Apple's Common-Crypto. Moreover, until very recently (May 2019), wNAF was present in all three branches of OpenSSL.

However, implementing the scalar multiplication using wNAF representation and no added layer of security makes the protocol vulnerable to side-channel attacks. Side-channel attacks were first introduced about two decades ago by Kocher et al. [14], and have since been used to break many implementations, and in particular some cryptographic primitives such as AES, RSA, and ECDSA. They allow to recover secret information throughout observable leakage. In our case, this leakage corresponds to differences in the execution time of a part of the signing algorithm, observable by monitoring the cache.

For ECDSA, cache side-channel attacks such as FLUSH&RELOAD [28, 29] have been used to recover information about either the sequence of operations used to execute the scalar multiplication, or for example in [8] the modular inversion. For the scalar multiplication, these operations are either a multiplication or an addition depending on the bits of k . This information is usually referred to as a double-and-add chain or the trace of k . A trace is created when a signature is produced by ECDSA and thus we talk about signatures and traces in an equivalent sense. At this point, we ask how many traces need to be collected to successfully recover the secret key. Indeed, from an attacker's perspective, the least traces necessary, the more efficient the attack is. This quantity depends on how much information can be extracted from a single trace and how combining information of multiple traces is used to recover the key. We work on the latter to minimize the number of traces needed.

The nature of the information obtained from the side channel attack allows to determine what kind of method should be carried out to recover the secret key. Attacks on ECDSA are inspired by attacks on a similar cryptosystem, DSA. In 2001, Howgrave-Graham and Smart [12] showed how knowing partial information of the nonce k in DSA can lead to a full secret key recovery. Later, Nguyen and Shparlinski [19] gave a polynomial time algorithm that recovers the secret key in ECDSA as soon as some consecutive bits of the ephemeral key are known. They showed that using the information leaked by the side channel attack, one can recover the secret key by constructing an instance of the Hidden Number Problem (HNP) [4]. The basic structure of the attack algorithm is to construct a lattice which contains the knowledge of consecutive bits of the ephemeral keys, and by solving CVP or SVP, to recover the secret key. This type of attack has been done in [3, 8, 26, 28]. However, these results considered perfect traces, but obtaining traces without any misreadings is very rare. In 2018,

Dall et al. [6] included an error-resilience analysis to their attack: they showed that key recovery with HNP is still possible even in the presence of erroneous traces.

In 2016, Fan, Wang and Cheng [7] used another lattice-based method to attack ECDSA: the Extended Hidden Number Problem (EHNP) [11]. EHNP mostly differs from HNP by the nature of the information given as input. Indeed, the information required to construct an instance of EHNP is not sequences of consecutive bits, but the positions of the non-zero coefficients in any representation of some integers. This model, which we consider in this article as well, is relevant when describing information coming from FLUSH&RELOAD or PRIME&PROBE attacks for example, the latter giving a more generic scenario with no shared data between the attacker and the victim. In [7], the authors are able to extract 105.8 bits of information per signature on average, and thus require in theory only 3 signatures to recover a 256-bit secret key. In practice, they were able to recover the secret key using 4 error-free traces.

In order to optimize an attack on ECDSA various aspects should be considered. By minimizing the number of signatures required in the lattice construction, one minimizes the number of traces needed to be collected during the side-channel attack. Moreover, reducing the time of the lattice part of the attack, and improving the probability of success of key recovery allows to reduce the overall time of the attack. In this paper, we improve on all three of these aspects. Furthermore, we propose the first error-resilience analysis for EHNP and show that key recovery is still possible with erroneous traces too.

Contributions: In this work, we reinvestigate the attack against ECDSA with wNAF representation for the scalar multiplication using EHNP. We focus on the lattice part of the attack, *i.e.*, the exploitation of the information gathered by the side-channel attack. We first assume we obtain a set of error-free traces from a side-channel analysis. We preselect some of these traces to optimize the attack. The main idea of the lattice part is then to use the ECDSA equation and the knowledge gained from the selected traces to construct a set of modular equations which include the secret key as an unknown. These modular equations are then incorporated into a lattice basis similar to the one given in [7], and a short vector in it will contain the necessary information to reconstruct the secret key. We call “experiment” one run of this algorithm. An experiment succeeds if the algorithm recovers the secret key.

A New Preprocessing Method. The idea of selecting good traces beforehand has already been explored in [27]. The authors suggest three rules to select traces that improve the attack on the lattice part. Given a certain (large) amount of traces available, the lattice is usually built with a much smaller subset of these traces. Trying to identify beforehand the traces that would result in a better attack is a clever option. The aim of our new preprocessing—that completely differs from [27]—is to regulate the size of the coefficients in the lattice, and this results in a better lattice reduction time. For instance, with 3 signatures, we were able to reduce the total time of the attack by a factor of 7.

Analyzing the Attack. Several parameters intervene while building and reducing the lattice. We analyze the performance of the attack with respect to these parameters and present the best parameters that optimize either the total time or the probability of success.

First, we focus on the attack time. Note that when talking about the overall time of the attack, we consider the average time of a single experiment multiplied by the number of trials necessary to recover the secret key. We compare¹ our times with the numbers reported in [7, Table 3] with method C. Indeed, methods *A* and *B* in [7] use extra information that comes from choices in the implementation which we choose to ignore as we want our analysis to remain as general as possible. The comparison is justified as we consider the same leakage model, and compare timings when running experiments on similar machines. For 4 signatures, our attack is slightly slower² than timings in [7]. However, when considering more than 4 signatures, our attack is faster. We experiment up to 8 signatures to further improve our overall time. In this case, our attack runs at best in 2 min and 25 s. Timings for 8 signatures are not reported in [7], and the case of 3 signatures was never reached before our work. In Table 1, we compare our times with the fastest times reported by [7]. We choose their fastest times but concerning our results we choose to report experiments which are faster (not the fastest) with, if possible, better probability than theirs.

Table 1. Comparing attack times with [7], for 5000 experiments.

Number of signatures	Our attack		[7]	
	Time	Success (%)	Time	Success (%)
3	39 h	0.2	–	–
4	1 h 17 min	0.5	41 min	1.5
5	8 min 20 s	6.5	18 min	1
6	≈ 5 min	25	18 min	22
7	≈ 3 min	17.5	34 min	24
8	≈ 2 min	29	–	–

The overall time of the attack is also dependent on the success probability of key recovery. From Table 2, one can see that our success probability is higher than [7], except for 7 signatures. They have 68% of success with their best parameters whereas we only reach 45% in this case.

¹ In order to have a fair comparison with our methodology, the times reported in [7] with which we compare ourselves have to be multiplied by the number of trials necessary for their attack succeed, thus increasing their total time by a lot. Using 5 signatures, their best total time would be around 15 h instead of 18 min.

² For 4 signatures, no times are reported without method *A*. Thus, we have no other choice than to compare our times with theirs, using *A*. Yet their time for 4 signatures without *A* should at least be the time they report with it.

For the sake of completeness, we mention that in [21], the authors use HNP to recover the secret key using 13 signatures. Their success probability in this case is around 54% and their overall time is close to 20 s, hence much faster. However, as their leakage model is different, we do not further mention their work.

Finding the Key with Only Three Signatures. Overall, combining a new preprocessing method, a modified lattice construction and a careful choice of parameters allows us to mount an attack which works in practice with only 3 signatures. However, the probability of success in this case is very low. We were able to recover the secret key only once with BKZ-35 over 5000 experiments. This result is difficult to quantify as a probability but we note that finding the key a single time over 5000 experiments is still much better than randomly finding a 256-bit integer. If we assume the probability is around 0.02%, as each trial costs 200 s in average, we can expect to find the secret key after 12 days using a single core. Note that this time can be greatly reduced when parallelizing the process, *i.e.*, each trial can be run on a separate core. On the other hand, if we use our preprocessing method, with 3 signatures we obtain a probability of success of 0.2% and a total time of key recovery of 39 h, thus the factor 7 of improvement mentioned above. Despite the low probability of success, this result remains interesting nonetheless. Indeed, the authors in [7] reported that in practice, the key couldn't be recovered using less than 4 signatures and we improve on their result.

Table 2. Comparing success probability with [7], for 5000 experiments.

Number of signatures	Our attack		[7]	
	Success (%)	Time	Success (%)	Time
3	0.2	39 h	–	–
4	4	25 h 28 min	1.5	41 min
5	20	2 h 42 min	4	36 min
6	40	1 h 4 min	35	1 h 43 min
7	45	2 h 36 min	68	3 h 58 min
8	45	5 h 2 min	–	–

Resilience to Errors. We also investigate the resilience to errors of our attack. Such an analysis has not yet been done in the setup of EHNP. It is important to underline that collecting traces without any errors using any side-channel attack is very hard. Previous works used perfect traces to mount the lattice attack. Thus, it required collecting more traces. As pointed out in [7], more or less twice as many signatures are needed if errors are considered. In practice, this led [7] to gather in average 8 signatures to be able to find the key with 4 perfect traces. We experimentally show that we are still able to recover the secret key even in

the presence of faulty traces. In particular, we find the key using only 4 faulty traces, but with a very low probability of success. As the percentage of incorrect digits in the trace grows, the probability of success decreases and thus more signatures are required to successfully recover the secret key. For instance, if 2% of the digits are wrong among all the digits of a given set of traces, it is still possible to recover the key with 6 signatures. This result is valid if errors are uniformly distributed over the digits. However, we have a better probability to recover the key if errors consist in 0-digit faulty readings, *i.e.*, 0 digits read as non-zero. In other words, the attack could work with a higher percentage of errors, around 4%, if we could ensure from the side channel attack and some preprocessing methods that none of the non-zero digits have been flipped to 0.

Organization: Sect. 2 gives background on ECDSA and the wNAF representation. In Sect. 3, we explain how to transform EHNP into a lattice problem. We explicit the lattice basis and analyze the length of the short vectors found in the reduced basis. In Sect. 4, we introduce our preprocessing method which allows us to reduce the overall time of our attack. In Sect. 5, we give experimental results. Finally, in Sect. 6, we give an error resilience analysis.

2 Preliminaries

2.1 Elliptic Curves Digital Signature Algorithm

The ECDSA algorithm is a variant of the Digital Signature Algorithm, DSA, [17] which uses elliptic curves instead of finite fields. The parameters used in ECDSA are an elliptic curve E over a finite field, a generator G of prime order q and a hash function H . The private key is an integer α such that $1 < \alpha < q - 1$ and the public key is $p_k = [\alpha]G$, the scalar multiplication of G by α .

To sign a message m using the private key α , randomly select an ephemeral key $k \leftarrow_R \mathbb{Z}_q$ and compute $[k]G$. Let r be the x -coordinate of $[k]G$. If $r = 0$, select a new nonce k . Then, compute $s = k^{-1}(H(m) + \alpha r) \bmod q$ and again if $s = 0$, select a new nonce k . Finally, the signature is given by the pair (r, s) .

In order to verify a signature, first check if $r, s \in \mathbb{Z}_q$, otherwise the signature is not valid. Then, compute $v_1 = H(m) \cdot s^{-1} \bmod q$, $v_2 = r \cdot s^{-1} \bmod q$ and $(x, y) = [v_1]G + [v_2]p_k$. Finally, the signature is valid if $x \equiv r \pmod{q}$.

We consider a 128-bit level of security. Hence α, q and k are 256-bit integers.

2.2 WNAF Representation

The ECDSA algorithm requires the computation of $[k]G$, a scalar multiplication. In [10], various methods to compute fast exponentiation are presented. One family of such methods is called window methods and comes from NAF representation. Indeed, the NAF representation does not allow two non-zero digits to be consecutive, thus reducing the Hamming weight of the representation of the scalar. The basic idea of a window method is to consider chunks of w bits

in the representation of the scalar k , compute powers in the window bit by bit, square w times and then multiply by the power in the next window. The window methods can be combined with the NAF representation of k . For any $k \in \mathbb{Z}$, a representation $k = \sum_{j=0}^{\infty} k_j 2^j$ is called a NAF if $k_j \in \{0, \pm 1\}$ and $k_j k_{j+1} = 0$ for all $j \geq 0$. Moreover, every k has a unique NAF representation. The NAF representation minimizes the number of non-zero digits k_j . It is presented in Algorithm 1.

The NAF representation can be combined with a sliding window method to further improve the execution time. For instance, in OpenSSL (up to the latest versions using wNAF 1.1.1b), the window size usually chosen was $w = 3$, which is the value we set for all our experiments. The scalar k is converted into wNAF form using Algorithm 2. The sequence of digits m_i belongs to the set $\{0, \pm 1, \pm 3, \dots, \pm(2^w - 1)\}$. Let k be the sum of its non-zero digits, renamed k_i . More precisely, we get $k = \sum_{j=1}^{\ell} k_j 2^{\lambda_j}$, where ℓ is the number of non-zero digits, and λ_j represents the position of the digit k_j in the wNAF representation.

Example 1. In binary, we can write $23 = 2^4 + 2^2 + 2^1 + 2^0 = (1, 0, 1, 1, 1)$ whereas in NAF-representation, we have $23 = 2^5 - 2^3 - 2^0 = (1, 0, -1, 0, 0, -1)$. With $w = 3$, the wNAF representation gives $23 = 2^4 + 7 \times 2^0 = (1, 0, 0, 0, 7)$.

Input: $k \in \mathbb{Z}^+$
Output: NAF representation of k

```

i = 0;
while k > 0 do
  if k (mod 2) = 1 then
    | ki = 2 - (k (mod 4));
    | k = k - ki;
  else
    | ki = 0;
  end
  k = k/2;
  i = i + 1;
end
return ki-1, ki-2, ..., k1, k0
Algorithm 1: NAF algorithm

```

Input: $k \in \mathbb{Z}^+$, $w \in \mathbb{N}$
Output: $(m_0, m_1, \dots, m_n)$, i.e., k in its wNAF representation

```

i = 0;
while k > 0 do
  if k (mod 2) = 1 then
    | mi = k (mod  $2^{w+1}$ );
    | if mi ≥  $2^w$  then
    |   | mi = mi -  $2^{w+1}$ ;
    | end
    | k = k - mi;
  else
    | mi = 0;
  end
  k = k/2;
  i = i + 1;
end
Algorithm 2: wNAF representation

```

2.3 Lattice Reduction Algorithms

A $\mathbb{Z}$ -lattice is a discrete additive subgroup of $\mathbb{Z}^n$. It is usually specified by a *basis matrix* $B \in \mathbb{Z}^{n \times n}$. The lattice $L(B)$ generated by B consists of all integer

combinations of the row vectors in B . The determinant of a lattice is the absolute value of the determinant of a basis matrix: $\det L(B) = |\det B|$. The discreteness property ensures that there is a vector v_1 reaching the minimum non-zero value for the euclidean norm. Let us write $\|v_1\|_2 = \lambda_1$. Let λ_i be the i^{th} successive minimum of the lattice. The LLL algorithm [15] takes as an input a lattice basis, and returns in polynomial time in the lattice dimension n a reduced lattice basis whose vectors b_i satisfy the worst-case approximation bound $\|b_i\|_2 \leq 2^{(n-1)/2} \lambda_i$. In practice, for random lattices, LLL obtains approximation factors such that $\|b_1\|_2 \leq 1.02^n \lambda_1$ as noted by Nguyen and Stehlé [18]. Moreover, for random lattices, the Gaussian heuristic implies that $\lambda_1 \approx \sqrt{n/(2\pi e)} \det(L)^{1/n}$.

The BKZ algorithm [22, 24] is exponential in some given block-size β and polynomial in the lattice dimension n . It outputs a reduced lattice basis whose vectors b_i satisfy the approximation $\|b_i\|_2 \leq i \gamma_\beta^{(n-i)/(k-1)} \lambda_i$ [23], where γ_β is the Hermite constant. In practice, Chen and Nguyen [5] observed that BKZ returns vectors such that $b_1 \leq (1 + \epsilon_\beta)^n \lambda_1$ where ϵ_β depends on the block-size β . For random lattices, they get $1 + \epsilon_\beta = 1.01$ for a block-size $\beta = 85$.

3 Attacking ECDSA Using Lattices

Using some side-channel attack, one can recover information about the wNAF representation of the nonce k . In particular, it allows us to know the positions of the non-zero coefficients in the representation of k . However, the value of these coefficients are unknown. This information can be used in the setup of the Extended Hidden Number Problem (EHNP) to recover the secret key. For many messages m , we use ECDSA to produce signatures (r, s) and each run of the signing algorithm produces a nonce k . We assume we have the corresponding trace of the nonce, that is, the equivalent of the double-and-add chain of kG using wNAF. The goal of the attack is to recover the secret α while optimizing either the number of signatures required or the total time of the attack.

3.1 The Extended Hidden Number Problem

The Hidden Number Problem (HNP) allows to recover a secret element $\alpha \in \mathbb{Z}_q$ if some information about the most significant bits of random multiples of $\alpha \pmod{q}$ are known for some prime q . Boneh and Venkatesan show how to recover α in polynomial time with probability greater than $1/2$. In [11], the authors extend the HNP and present a polynomial time algorithm for solving the instances of this extended problem. The Extended Hidden Number Problem is defined as follows. Given u congruences of the form

$$a_i \alpha + \sum_{j=1}^{\ell_i} b_{i,j} k_{i,j} \equiv c_i \pmod{q}, \quad (1)$$

where the secret α and $0 \leq k_{i,j} \leq 2^{n_{ij}}$ are unknown, and the values n_{ij} , a_i , $b_{i,j}$, c_i , ℓ_i are known for $1 \leq i \leq u$ (see [11], Definition 3), one has to recover α in polynomial time. The EHNP can then be transformed into a lattice problem and one recovers the secret α by solving a short vector problem in a given lattice.

3.2 Using EHNP to Attack ECDSA

From the ECDSA algorithm, we know that given a message m , the algorithm outputs a signature (r, s) such that

$$\alpha r = sk - H(m) \pmod{q}. \quad (2)$$

The value $H(m)$ is just some hash of the message m . We consider a set of u signature pairs (r_i, s_i) with corresponding message m_i that satisfy Eq. (2). For each signature pair, we have a nonce k . Using the wNAF representation of k , we write $k = \sum_{j=1}^{\ell} k_j 2^{\lambda_j}$, with $k_j \in \{\pm 1, \pm 3, \dots, \pm(2^w - 1)\}$ and the choice of w depends on the implementation. Note that the coefficients k_j are unknown, however, the positions λ_j are supposed to be known via some side-channel leakage. It is then possible to represent the ephemeral key k as the sum of a known part, and an unknown part. As the value of k_j is odd, one can write $k_j = 2k'_j + 1$, where $-2^{w-1} \leq k'_j \leq 2^{w-1} - 1$. Using the same notations as in [7], set $d_j = k'_j + 2^{w-1}$, where $0 \leq d_j \leq 2^w - 1$. In the rest of the paper, we will denote by μ_j the window-size of d_j . Note that here, $\mu_j = w$ but this window-size will be modified later. This allows to rewrite the value of k as

$$k = \sum_{j=1}^{\ell} k_j 2^{\lambda_j} = \bar{k} + \sum_{j=1}^{\ell} d_j 2^{\lambda_j + 1}, \quad (3)$$

with $\bar{k} = \sum_{j=1}^{\ell} 2^{\lambda_j} - \sum_{j=1}^{\ell} 2^{\lambda_j + w}$. The expression of $\bar{k}$ represents the known part of k . By substituting k in Eq. (3), we get a system of modular equations:

$$\alpha r_i - \sum_{j=1}^{\ell_i} 2^{\lambda_{i,j} + 1} s_i d_{i,j} - (s_i \bar{k}_i - H(m_i)) \equiv 0 \pmod{q} \quad (4)$$

where the unknowns are α and the $d_{i,j}$. The known values are ℓ_i , which is the number of non-zero digits in k for the i^{th} sample, $\lambda_{i,j}$, which is the position of the j^{th} non-zero digit in k for the i^{th} sample and $\bar{k}$ defined above. Equation (4) is then used as input to EHNP, following the method explained in [11]. The problem of finding the secret key is then reduced to solving the short vector problem in a given lattice presented in the following section.

3.3 Constructing the Lattice

Before giving the lattice basis construction, we redefine Eq. (4) to reduce the number of unknown variables in the system. This will allow us to construct a lattice of smaller dimension. Again, we use the same notations as in [7].

Eliminating One Variable. One straightforward way to reduce the lattice dimension is to eliminate a variable from the system. In this case, one can eliminate α

from Eq. (4). Let E_i denote the i^{th} equation of the system. Then by computing $r_1 E_i - r_i E_1$, we get the following new modular equations

$$\sum_{j=1}^{\ell_1} \underbrace{(2^{\lambda_{1,j}+1} s_1 r_i)}_{:=\tau_{j,i}} d_{1,j} + \sum_{j=1}^{\ell_i} \underbrace{(-2^{\lambda_{i,j}+1} s_i r_1)}_{:=\sigma_{i,j}} d_{i,j} - \underbrace{r_1(s_i \bar{k}_i - H(m_i)) + r_i(s_1 \bar{k}_1 - H(m_1))}_{:=\gamma_i} \equiv 0 \pmod{q}. \quad (5)$$

Using the same notations as in [7], we define $\tau_{j,i} = 2^{\lambda_{1,j}+1} s_1 r_i$, $\sigma_{i,j} = -2^{\lambda_{i,j}+1} s_i r_1$ and $\gamma_i = r_1(s_i \bar{k}_i - H(m_i)) + r_i(s_1 \bar{k}_1 - H(m_1))$ for $2 \leq i \leq u$, $1 \leq j \leq \ell_i$. Even if α is eliminated from the equations, if we recover some $d_{i,j}$ values from a short vector in the lattice, we can recover α using any equation in the modular system (4). We now use Eq. (5) to construct the lattice basis.

From a Modular System to a Lattice Basis. Let $\mathcal{L}$ be the lattice constructed for the attack, and we have $\mathcal{L} = \mathcal{L}(\mathcal{B})$ where the lattice basis $\mathcal{B}$ is given below. Let $m = \max_{i,j} \mu_{ij}$ for $1 \leq j \leq \ell_i$ and $2 \leq i \leq u$. We set a scaling factor $\Delta \in \mathbb{N}$ to be defined later. The lattice basis is given by

$$\mathcal{B} = \begin{pmatrix} \text{Eq (5), } i=2 \dots \text{Eq (5), } i=u & & & & \\ \Delta 2^m q & 0 & 0 & 0 & \\ 0 & \ddots & \vdots & & \\ 0 & \dots & \Delta 2^m q & 0 & \\ \Delta 2^m \tau_{1,2} & \dots & \Delta 2^m \tau_{1,u} & 2^{m-\mu_{1,1}} & \\ \vdots & & \vdots & 0 & \ddots \\ \Delta 2^m \tau_{\ell_1,2} & \dots & \Delta 2^m \tau_{\ell_1,u} & 2^{m-\mu_{1,\ell_1}} & \\ \Delta 2^m \sigma_{2,1} & 0 & 0 & 2^{m-\mu_{2,1}} & \\ \vdots & & \vdots & & \ddots \\ \Delta 2^m \sigma_{2,\ell_2} & & \vdots & 2^{m-\mu_{2,\ell_2}} & \\ 0 & \ddots & 0 & \vdots & \ddots \\ \vdots & & \Delta 2^m \sigma_{u,1} & & 2^{m-\mu_{u,1}} \\ \vdots & & \vdots & & \ddots \\ 0 & 0 & \Delta 2^m \sigma_{u,\ell_u} & 0 & 2^{m-\mu_{u,\ell_u}} \\ \Delta 2^m \gamma_2 & \dots & \Delta 2^m \gamma_u & 2^{m-1} & \dots & 2^{m-1} & 2^{m-1} \end{pmatrix}$$

Let $n = (u-1) + T + 1 = T + u$, with $T = \sum_{i=1}^u \ell_i$, be the dimension of the lattice. The $u-1$ first columns correspond to Eq. (5) for $2 \leq i \leq u$. Each of the remaining columns, except the last one, corresponds to a d_{ij} , and contains coefficients that allow to regulate the size of the d_{ij} . The determinant of $\mathcal{L}$ is given by $\det \mathcal{L} = q^{u-1} (\Delta 2^m)^{u-1} 2^{\sum_{i,j} (m-\mu_{i,j})} 2^{m-1}$.

The lattice is built such that there exists $w \in \mathcal{L}$ which contains the unknowns $d_{i,j}$. To find it, we know there exists some values $t_2, t_2, \dots, t_u$ such that if $v = (t_2, \dots, t_u, d_{1,1}, \dots, d_{u,\ell_u}, -1)$, we get $w = v\mathcal{B}$, and

$$w = (0, \dots, 0, d_{1,1} 2^{m-\mu_{1,1}} - 2^{m-1}, \dots, d_{u,\ell_u} 2^{m-\mu_{u,\ell_u}} - 2^{m-1}, -2^{m-1}).$$

If we are able to find w in the lattice, then we can reconstruct the secret key α . In order to find w , we estimate its norm and make sure w appears in the reduced basis. After reducing the basis, we look for vectors of the correct shape, *i.e.*, with sufficiently many zeros at the beginning and the correct last coefficient, and attempt to recover α for each of these.

How the Size of Δ Affects the Norms of the Short Vectors. In order to find the vector w in the lattice, we reduce $\mathcal{B}$ using LLL or BKZ. For w to appear in the reduced basis, one should at least set Δ such that

$$\|w\|_2 \leq (1.02)^n (\det L)^{1/n}. \quad (6)$$

The vector w we expect to find has norm $\|w\|_2 \leq 2^{m-1} \sqrt{T+1}$. From Eq. (6), one can deduce the value of Δ needed to find w in the reduced lattice:

$$\Delta \geq \frac{(T+1)^{(T+u)/(2(u-1))} 2^{\frac{1+\sum \mu_{i,j} - (u+T)}{u-1}}}{q(1.02)^{\frac{(T+u)^2}{u-1}}} := \Delta_{th}.$$

In our experiments, the average value of ℓ_i for $1 \leq i \leq u$ is $\tilde{\ell} = 26$, and thus $T = u \times \tilde{\ell}$ on average. Moreover, the average value of μ_{ij} is 7 and so on average $\sum \mu_{ij} = 7 \times u \times \tilde{\ell}$. Hence, if we compute Δ_{th} for $u = 3, \dots, 8$, with these values, we obtain $\Delta_{th} \ll 1$, which does not help us to set this parameter. In practice, we verify that $\Delta = 1$ allows us to recover the secret key. In Sect. 5, we vary the size of Δ to see whether a slightly larger value affects the probability of success.

Too Many Small Vectors. While running BKZ on $\mathcal{B}$, we note that for some specific sets of parameters the reduced basis contains some undesired short vectors, *i.e.*, vectors that are shorter than w . This can be explained by looking at two consecutive rows in the lattice basis given above, say the j^{th} row and the $(j+1)^{th}$ row. For example, one can look at rows which correspond to the $\sigma_{i,j}$ values but the same argument is valid for the rows concerning the $\tau_{j,i}$. From the definitions of the σ values we have $\sigma_{i,j+1} = -2^{\lambda_{i,j+1}+1} \cdot s_i r_1 = -2^{\lambda_{i,j+1}+1} \cdot (\frac{\sigma_{i,j}}{-2^{\lambda_{i,j}+1}})$. So $\sigma_{i,j+1} = 2^{\lambda_{i,j+1}-\lambda_{i,j}} \cdot \sigma_{i,j}$. Thus the linear combination given by the $(j+1)^{th}$ row minus $2^{\lambda_{i,j+1}-\lambda_{i,j}}$ times the j^{th} row gives a vector

$$(0, \dots, 0, -2^{\lambda_{i,j+1}-\lambda_{i,j}+m-\mu_{i,j}}, 2^{m-\mu_{i,j+1}}, 0, \dots, 0). \quad (7)$$

Yet, this vector is expected to have smaller norm than w . Some experimental observations are detailed in Sect. 5.

Differences with the Lattice Construction Given in [7]. Let $\mathcal{B}'$ be the lattice basis constructed in [7]. Our basis $\mathcal{B}$ is a rescaled version of $\mathcal{B}'$ such that $\mathcal{B} = 2^m \Delta \mathcal{B}'$. This rescaling allows us to ensure that all the coefficients in our lattice basis are integer values. Note that [7] have a value δ in their construction which corresponds to $1/\Delta$. In this work, we give a precise analysis of the value of Δ , both theoretically and experimentally in Sect. 5, which is missing in [7].

4 Improving the Lattice Attack

4.1 Reducing the Lattice Dimension: The Merging Technique

In [7], the authors present another way to further reduce the lattice dimension, which they call the merging technique. It aims at reducing the lattice dimension by reducing the number of non-zero digits of k . The lattice dimension depends on the value $T = \sum_{i=1}^u \ell_i$, and thus reducing T reduces the dimension. To understand the attack, it suffices to know that after merging, we obtain some new values ℓ' corresponding to the new number of non-zero digits and λ'_j the position of these digits for $1 \leq j \leq \ell'$. After merging, one can rewrite $k = \bar{k} + \sum_{j=1}^{\ell'} d'_j 2^{\lambda'_j+1}$, where the new d'_j have a new window size which we denote μ_j , i.e., $0 \leq d'_j \leq 2^{\mu_j} - 1$.

We present our merging algorithm based on Algorithm 3 given in [7]. Our algorithm modifies directly the sequence $\{\lambda_j\}_{j=1}^{\ell}$, whereas [7] work on the double-and-add chains. This helped us avoid some implementation issues such as an index overrun present in Algorithm 3 [7], line 7. To facilitate the ease of reading of (our) Algorithm 3, we work with dynamic tables. Let us first recall various known methods we use in the algorithm: *push_back*(e) inserts an element e at the end of the table, *at*(i) outputs the element at index i , and *last*() returns the last element of the table. We consider tables of integers indexed in $[0; S - 1]$, where S is the size of the table.

Input: v_λ , a table of size n with the positions of non-zero digits in the trace sorted in increasing order and $n \geq 1$, a window size w .

Output: $v_{\lambda'}$, a table of size $n' \leq n$ containing the merged λ values and table v_μ of same size n' , with the values of the window size μ_i .

Initialisation

$i \leftarrow 1$;

$v_{\lambda'} \leftarrow$ empty array;

$v_\mu \leftarrow$ empty array;

Processing

$v_{\lambda'}.push_back(v_\lambda.at(0))$;

while $i < n$ **do**

$dist \leftarrow v_\lambda.at(i) - v_\lambda.at(i - 1)$;

if $dist > w + 1$ **then**

$v_\mu.push_back(v_\lambda.at(i - 1) - v_{\lambda'}.last() + w)$;

$v_{\lambda'}.push_back(v_\lambda.at(i))$;

end

$i \leftarrow i + 1$;

end

$v_\mu.push_back(v_\lambda.at(n) - v_{\lambda'}.last() + w)$;

return $(v_{\lambda'}, v_\mu)$

Algorithm 3: Merging algorithm

A useful example of the merging technique is given in [7]. From 3 to 8 signatures the approximate dimension of the lattices using the elimination and merging techniques are the following: 80, 110, 135, 160, 190 and 215. Each new lattice dimension is roughly 54% of the dimension of the lattice before applying these techniques, for the same number of signatures. For instance, with 8 signatures we would have a lattice of dimension 400 on average, far too large to be easily reduced. For the traces we consider, after merging the mean of the ℓ_i is 26, the minimum being 17 and the maximum 37 with a standard deviation of 3. One could further reduce the lattice dimension by preprocessing traces with small ℓ_i . However, the standard deviation being small, the difference in the reduction times should not be affected too much.

4.2 Preprocessing the Traces

The two main pieces of information we can extract and use in our attack are first the number of non-zero digits in the wNAF representation of the nonce k , denoted ℓ and the weight of each non-zero digit, denoted μ_j for $1 \leq j \leq \ell$. Let $\mathcal{T}$ be the set of traces we obtained from the side-channel leakage representing the wNAF representation of the nonce k used while producing an ECDSA signature. We consider the subset $S_a = \{t \in \mathcal{T} \mid \max_j \mu_j \leq a, 1 \leq j \leq \ell\}$. We choose to preselect traces in a subset S_a for small values of a . The idea behind this preprocessing is to regulate the size of the coefficients in the lattice. Indeed, when selecting u traces for the attack, by upper-bounding $m = \max_{i,j} \mu_{i,j}$ for $2 \leq i \leq u$, we force the coefficients to remain smaller than when taking traces at random.

We work with a set $\mathcal{T}$ of 2000 traces such that for all traces $11 \leq \max_j \mu_j \leq 67$. The proportion of signatures corresponding to the different preprocessing subsets we consider in our experiments are: 2% for S_{11} , 18% for S_{15} and 44% for S_{19} . The effect of preprocessing on the total time is explained in Sect. 5.

5 Performance Analysis

Traces from the Real World. We work with the elliptic curve `secp256k1` but none of the techniques introduced here are limited to this specific elliptic curve. We consider traces from a FLUSH&RELOAD attack, executed through hyper-threading, as it can virtually recover the most amount of information.³

To the best of our knowledge, the only information we can recover are the positions of the non-zero digits. We are not able to determine the sign or the value of the digits in the wNAF representation. In [7], the authors exploit the

³ In practice, measurements done during the cache attack depend on the noise in the execution environment, the threat model and the target leaky implementation. For instance, FLUSH&RELOAD ran from another core would be noisy. PRIME&PROBE would give the same information, with a more generic scenario. In an SGX scenario, it would recover the largest amount of information but in a user/user threat model it would be too noisy to lead to practical key recovery.

fact that the length of the binary string of k is fixed in implementations such as OpenSSL, and thus more information can be recovered by comparing this length to the length of the double-and-add chain. In particular, they were able to recover the MSB of k , and in some cases the sign of the second MSB. We do not consider this extra information as we want our analysis to remain general.

We report calculations ran on error-free traces where we evaluate the total time necessary to recover the secret key and the probability of success of the attack. Our experiments have two possible outputs: either we reconstruct the secret key α and thus consider the experiment a success, or we do not recover the secret key, and the experiment fails. In order to compute the success probability and the average time of one reduction, we run 5000 experiments for some specific sets of parameters using either Sage’s default BKZ implementation [25] or a more recent implementation of the latest sieving strategies, the General Sieve Kernel (G6K) [1]. The experiments were ran using the cluster Grid’5000 on a single core of an Intel Xeon Gold 6130. The total time is the average time of a single reduction multiplied by the number of trials necessary to recover the key. The number of trials necessary to recover the secret key corresponds the number of experiments ran until we have a success for a given set of parameters. For a fixed number of signatures, we either optimize the total time or the success probability. We report numbers in Tables 3, 4 when using BKZ.⁴

Table 3. Fastest key recovery with respect to the number of signatures.

Number of signatures	Total time	Parameters			Probability of success (%)
		BKZ	Preprocessing	Δ	
3	39 h	35	S_{11}	$\approx 2^3$	0.2
4	1 h 17	25	S_{15}	$\approx 2^3$	0.5
5	8 min 20	25	S_{19}	$\approx 2^3$	6.5
6	3 min 55	20	S_{all}	$\approx 2^3$	7
7	2 min 43	20	S_{all}	$\approx 2^3$	17.5
8	2 min 25	20	S_{all}	$\approx 2^3$	29

Comments on G6K: We do not report the full experiments ran with G6K since using this implementation does not lead to the fastest total time of our attack: around 2 min using 8 signatures for BKZ and at best 5 min for G6K.

⁴ In [7], the authors use an Intel Core i7-3770 CPU running at 3.40 GHz on a single core. In order for the time comparison to be meaningful, we ran experiments with a machine of comparable performance to estimate the timings of a single reduction. As we obtained similar timings with an older machine than used in [7], the variations we find when comparing ourselves to them solely come from the lattice construction and the reduction algorithm being used rather than hardware differences.

Table 4. Highest probability of success with respect to the number of signatures.

Number of signatures	Probability of success (%)	Parameters			Total time
		BKZ	Preprocessing	Δ	
3	0.2	35	S_{11}	$\approx 2^3$	39 h
4	4	35	S_{all}	$\approx 2^3$	25 h 28
5	20	35	S_{all}	$\approx 2^3$	2 h 42
6	40	35	S_{all}	$\approx 2^3$	1 h 04
7	45	35	S_{all}	$\approx 2^3$	2 h 36
8	45	35	S_{all}	$\approx 2^3$	5 h 02

However, G6K allows to reduce lattices with much higher block-sizes than BKZ. For comparable probabilities of success, G6K is faster. Considering the highest probability achieved, on one hand, BKZ-35 leads to a probability of success of 45%, and a single reduction takes 133 min. On the other hand, to reach around the same probability of success with G6K, we increase the block-size to 80, and a single reduction is only around 45 min on average. This is an improvement by a factor of 3 in the reduction time.

Only 3 Signatures. Using $\Delta \approx 2^3$ and no preprocessing, we recovered the secret key using 3 signatures with BKZ-35 only once and three times with BKZ-40. When using pre-processing S_{11} , BKZ-35 and $\Delta \approx 2^3$, the probability of success went up to 0.2%. Since all the probabilities remain much less than 1% an extensive analysis would have taken too long. Thus, in the rest of the section, the number of signatures only varies between 4 and 8. However, we want to emphasize that it is precisely this detailed analysis on a slightly higher number of signatures that allowed us to understand the impact of the parameters on the performance of the attack and resulted in finding the right ones allowing to mount the attack with 3 signatures.

Varying the Bitsize of Δ . In Fig. 1, we analyze the total time of key recovery as a function of the bitsize of Δ . We fix the block-size of BKZ to 25 and take traces without any preprocessing. We are able to recover the secret key by setting $\Delta = 1$, which is the lowest theoretical value one can choose. However, we observed a slight increase in the probability of success by taking a larger Δ . Without any surprise, we note that the total time to recover the secret key increases with the bitsize of Δ as the coefficients in the lattice basis become larger.

Analyzing the Effect of Preprocessing. We analyze the influence of our preprocessing method on the attack time. We fix BKZ block-size to 25. The effect of preprocessing is influenced by the bitsize of Δ and we give here an analyze for $\Delta \approx 2^{25}$ since the effect is more noticeable.

Fig. 1. Analyzing the overall time to recover the secret key as a function of the bitsize of Δ . We report the numbers BKZ-25 and no preprocessing. The optimal value for Δ is around 2^3 except for $u = 8$ where it is 2^5 .

The effect of preprocessing is difficult to predict since its behavior varies depending on the parameters, having both positive and negative effects. On the one hand, we reduce the size of all the coefficients in the lattice, thus reducing the reduction time. On the other hand, we generate more potential small vectors⁵ with norms smaller than the norm of w . For this reason, the probability of success of the attack decreases since the vector w is more likely to be a linear combination of vectors already in the reduced basis. For example, with 7 signatures we find in average w to be the third or fourth vector in the reduced basis without preprocessing, whereas with S_{11} it is more likely to appear in position 40.

The positive effect of preprocessing is most noticeable for $u = 4$ and $u = 5$, as shown in Fig. 2. For instance, using S_{15} and $u = 4$ lowers the overall time by a factor up to 5.7. For $u = 5$, we gain a factor close to 3 by using either S_{15} or S_{19} . For $u > 5$, using preprocessed traces is less impactful. For large Δ such as $\Delta \approx 2^{25}$, we still note some lower overall times when using S_{15} and S_{19} , up to a factor 2. When the bitsize gets smaller, reducing the size of the coefficients in the lattice is less impactful.

Balancing the Block-size of BKZ. Finally, we vary the block-size in the BKZ algorithm. We fix $\Delta \approx 2^3$ and use no preprocessing. We plot the results in Fig. 3 for 6 and 7 signatures. For other values of u , the plot is very similar and we omit them in Fig. 3. Without any surprise, we see that as we increase the block-size, the probability of success increases, however the reduction time increases significantly as well. This explains the results shown in Table 3 and Table 4: to reach the best probability of success one needs to increase the block-size in BKZ (we did not try any block-size greater than 40), but to get the fastest key recovery attack, the block-size is chosen between 20 and 25, except for 3 signatures where the probability of success is too low with these parameters.

⁵ In the sense of vectors exhibited in (7).

Fig. 2. Overall time to recover the secret key as a function of the preprocessing subset for 4 and 5 traces. The other parameters are fixed: $\Delta \approx 2^{25}$ and BKZ-25.

Fig. 3. Analyzing the number of trials to recover the secret key and the reduction time of the lattice as a function of the block-size of BKZ. We consider the cases where $u = 6$ and $u = 7$. The dotted lines correspond to the number of trials, and the continued lines to the reduction time in seconds.

6 Error Resilience Analysis

It is not unexpected to have errors in the traces collected during side-channel attacks. Obtaining error-free traces requires some amount of work on the signal processing side. Prior to [6], the presence of errors in traces was either ignored or preprocessing was done on the traces until an error-free sample was found, see [2, 9]. In [6], it is shown the lattice attack still successfully recovers the secret key even when traces contain errors. An error in the setup given in [6] corresponds to an incorrect bound on the size of the values being collected. In our setup, a trace without errors corresponds to a trace where every single coefficient in the wNAF representation of k has been identified correctly as either non-zero or not. The probability of having an error in our setup is thus much higher. Side-channel attacks without any errors are very rare. Both [21] and [6] give some analysis of the attacks FLUSH&RELOAD and PRIME&PROBE in real life scenarios.

In [7], the results presented in the paper assume the FLUSH&RELOAD is implemented perfectly, without any error. In particular, to obtain 4 perfect traces and be able to run their experiment and find the key, one would need to have in average 8 traces from FLUSH&RELOAD – the probability to conduct to a perfect reading of the traces being 56% as pointed out in [21]. In our work, we show that it is possible to recover the secret key using only 4, even erroneous, traces. However, the probability of success is very low.

Recall that an error in our case corresponds to a flipped digit in the trace of k . Table 5 shows the attack success probability in the presence of errors. We ran used BKZ-25 and $\Delta \approx 2^3$ with traces taken from S_{all} . We average over 5000 experiments. We write $\ll 1$ when the attack succeeded less than five times over 5000 experiments, thus making it difficult to evaluate the probability of success.

Table 5. Error analysis using BKZ-25, $\Delta \approx 2^3$ and S_{all} .

Number of signatures	Probability of success (%)				
	0 error	5 errors	10 errors	20 errors	30 errors
4	0.28	$\ll 1$	0	0	0
5	4.58	0.86	0.18	$\ll 1$	0
6	19.52	5.26	1.26	0.14	$\ll 1$
7	33.54	10.82	3.42	0.32	$\ll 1$
8	35.14	13.26	4.70	0.58	$\ll 1$

The attack works up to a resilience to 2% of errors. Indeed, for $u = 6$, we recovered the secret key with 30 errors, *i.e.*, 30 flipped digits over 6×257 digits.

Different Types of Errors. There exists two possible types of errors. In the first case, a coefficient which is zero is evaluated as a non-zero coefficient. In theory, this only adds a new variable to the system, *i.e.*, the number ℓ of non-zero coefficients is overestimated. This does not affect the probability of success much. Indeed, we just have an overly-constrained system. We can see in Fig. 4 that the probability of success of the attack indeed decreases slowly as we add errors of this form. With errors only of this form, we were able to recover the secret key up to nearly 4% of errors, (for instance with $u = 6$, using BKZ-35). The other type of errors consists of a non-zero coefficients which is misread as a zero coefficient. In this case, we lose information necessary for the key recovery and thus this type of error affects the probability of success far more importantly as can also be seen in Fig. 4. In this setup, we were not able to recover the secret key when more than 3 errors of this type appear in the set of traces considered.

Strategy. If the signal processing method is hesitant between a non-zero digit or 0, we would recommend to favor putting a non-zero instead of 0 to increase the chance of having an error of type $0 \rightarrow$ non-zero, for which the attack is a lot more tolerant.

Fig. 4. Probability of success for key recovery with various types of errors when using $u = 8$, BKZ-25, $\Delta \approx 2^3$, and no preprocessing.

7 Conclusion and Countermeasures

In the last decades, most implementations of ECDSA have been the target of microarchitectural attacks, and thus existing implementations have either been replaced by more robust algorithms, or layers of security have been added.

For example, one way of minimizing leakage from the scalar multiplication is to use the Montgomery ladder scalar-by-point multiplication [16], much more resilient to side-channel attacks due to the regularity of the operations. However, this does not entirely remove the risk of leakage [28]. Additional countermeasures are necessary.

When looking at common countermeasures, many implementations use blinding or masking techniques [20], for example in BouncyCastle implementation of ECDSA. The former consists in blinding the data before doing any operations, and masking techniques randomize all the data-dependent operations by applying random transformations, thus making any leakage inexploitable.

However, it is important to keep in mind these lattices attacks as they can be applied at any level of an implementation that leaks the correct information.

Acknowledgement. We would like to thank Nadia Heninger for discussions about possible lattice constructions, Medhi Tibouchi for answering our side-channel questions, Alenka Zajic and Milos Prvulovic for providing us with traces from OpenSSL that allowed us to confirm our results on a deployed implementation, Daniel Genkin for pointing us towards the Extended Hidden Number Problem, and Pierrick Gaudry for his precious support and reading. Experiments presented in this paper were carried out using the Grid’5000 testbed, supported by a scientific interest group hosted by Inria and including CNRS, RENATER and several universities as well as other organizations.

References

1. Albrecht, M.R., Ducas, L., Herold, G., Kirshanova, E., Postlethwaite, E.W., Stevens, M.: The general sieve kernel and new records in lattice reduction. *Cryptology ePrint Archive, Report 2019/089* (2019). <https://eprint.iacr.org/2019/089>

2. Angel, J., Rahul, R., Ashokkumar, C., Menezes, B.: DSA signing key recovery with noisy side channels and variable error rates. In: Patra, A., Smart, N.P. (eds.) INDOCRYPT 2017. LNCS, vol. 10698, pp. 147–165. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-71667-1_8
3. Bengier, N., van de Pol, J., Smart, N.P., Yarom, Y.: “Ooh Aah... Just a Little Bit” : a small amount of side channel can go a long way. In: Batina, L., Robshaw, M. (eds.) CHES 2014. LNCS, vol. 8731, pp. 75–92. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-662-44709-3_5
4. Boneh, D., Venkatesan, R.: Hardness of computing the most significant bits of secret keys in diffie-hellman and related schemes. In: Kobitz, N. (ed.) CRYPTO 1996. LNCS, vol. 1109, pp. 129–142. Springer, Heidelberg (1996). https://doi.org/10.1007/3-540-68697-5_11
5. Chen, Y., Nguyen, P.Q.: BKZ 2.0: better lattice security estimates. In: Lee, D.H., Wang, X. (eds.) ASIACRYPT 2011. LNCS, vol. 7073, pp. 1–20. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-25385-0_1
6. Dall, F., et al.: Cachequote: Efficiently recovering long-term secrets of SGX EPID via cache attacks (2018)
7. Fan, S., Wang, W., Cheng, Q.: Attacking OpenSSL implementation of ECDSA with a few signatures. In: Weippl, E.R., Katzenbeisser, S., Kruegel, C., Myers, A.C., Halevi, S. (eds.) ACM CCS 2016, pp. 1505–1515. ACM Press, October 2016
8. García, C.P., Brumley, B.B.: Constant-time callees with variable-time callers. In: Kirda, E., Ristenpart, T. (eds.) USENIX Security 2017, pp. 83–98. USENIX Association, August 2017
9. Genkin, D., Pachmanov, L., Pipman, I., Tromer, E., Yarom, Y.: ECDSA key extraction from mobile devices via nonintrusive physical side channels. In: Weippl, E.R., Katzenbeisser, S., Kruegel, C., Myers, A.C., Halevi, S. (eds.) ACM CCS 2016, pp. 1626–1638. ACM Press, October 2016
10. Gordon, D.M.: A survey of fast exponentiation methods. *J. Algorithms* **27**(1), 129–146 (1998)
11. Hlaváč, M., Rosa, T.: Extended hidden number problem and its cryptanalytic applications. In: Biham, E., Youssef, A.M. (eds.) SAC 2006. LNCS, vol. 4356, pp. 114–133. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-74462-7_9
12. Howgrave-Graham, N.A., Smart, N.P.: Lattice attacks on digital signature schemes. *Des. Codes Cryptol.* **23**(3), 283–290 (2001)
13. Johnson, D., Menezes, A., Vanstone, S.: The elliptic curve digital signature algorithm (ECDSA). *Int. J. Inf. Secur.* **1**(1), 36–63 (2001)
14. Kocher, P., Jaffe, J., Jun, B.: Differential power analysis. In: Wiener, M. (ed.) CRYPTO 1999. LNCS, vol. 1666, pp. 388–397. Springer, Heidelberg (1999). https://doi.org/10.1007/3-540-48405-1_25
15. Lenstra, A.K., Lenstra, H.W., Lovász, L.: Factoring polynomials with rational coefficients. *Math. Ann.* **261**(4), 515–534 (1982)
16. Montgomery, P.L.: Speeding the pollard and elliptic curve methods of factorization. *Math. Comput.* **48**(177), 243–243 (1987)
17. National Institute of Standards and Technology: Digital Signature Standard (DSS) (2013)
18. Nguyen, P.Q., Stehlé, D.: LLL on the average. In: Hess, F., Pauli, S., Pohst, M. (eds.) ANTS 2006. LNCS, vol. 4076, pp. 238–256. Springer, Heidelberg (2006). https://doi.org/10.1007/11792086_18
19. Nguyen, P.Q., Shparlinski, I.E.: The insecurity of the elliptic curve digital signature algorithm with partially known nonces. *Des. Codes Cryptol.* **30**(2), 201–217 (2003)

20. Osvik, D.A., Shamir, A., Tromer, E.: Cache attacks and countermeasures: the case of AES. In: Pointcheval, D. (ed.) CT-RSA 2006. LNCS, vol. 3860, pp. 1–20. Springer, Heidelberg (2006). https://doi.org/10.1007/11605805_1
21. van de Pol, J., Smart, N.P., Yarom, Y.: Just a little bit more. In: Nyberg, K. (ed.) CT-RSA 2015. LNCS, vol. 9048, pp. 3–21. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-16715-2_1
22. Schnorr, C.P.: A hierarchy of polynomial time lattice basis reduction algorithms. Theoret. Comput. Sci. **53**(2–3), 201–224 (1987)
23. Schnorr, C.P.: Block reduced lattice bases and successive minima. Comb. Probab. Comput. **3**, 507–522 (1994)
24. Schnorr, C.P., Euchner, M.: Lattice basis reduction: Improved practical algorithms and solving subset sum problems. Math. Program. **66**(2), 181–199 (1994)
25. The FPLLL development team: FPLLL, a lattice reduction library (2016)
26. Vanstone, S.: Responses to NIST’s proposals (1992)
27. Wang, W., Fan, S.: Attacking OpenSSL ECDSA with a small amount of side-channel information. Sci. China Inf. Sci. **61**(3), 032105 (2017)
28. Yarom, Y., Benger, N.: Recovering OpenSSL ECDSA nonces using the FLUSH+RELOAD cache side-channel attack. IACR Cryptol. ePrint Archive **2014**, 140 (2014)
29. Yarom, Y., Falkner, K.: FLUSH+RELOAD: A high resolution, low noise, L3 cache side-channel attack. In: Proceedings of the 23rd USENIX Conference on Security Symposium, SEC 2014, Berkeley, CA, USA, pp. 719–732. USENIX Association (2014)