A Tale of Three Signatures: practical attack of ECDSA with wNAF

Gabrielle de Micheli, Rémi E. Piau, Cécile Pierrot

A Tale of Three Signatures:
practical attack of ECDSA with wNAF

Gabrielle De Micheli[1], Rémi Piau[1,2], and Cécile Pierrot[1]

[1] Université de Lorraine, CNRS, Inria, LORIA, F-54000 Nancy, France
[2] ENS 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) [16], first proposed in 1992 by Scott Vanstone [31], 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 [18], 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 [33, 35] have been used to recover information about either the sequence of operations used to execute the scalar multiplication, or for example in [10] 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 [15] showed how knowing partial information of the nonce k in DSA can lead to a full secret key recovery. Later, Nguyen and Shparlinski [23] 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 epheremal keys, and by solving CVP or SVP, to recover the secret key. This type of attack has been done in [3, 30, 32, 10]. 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) [13]. 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 [32]. 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 [32] — 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 compare[3] 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 slower[4] 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 minutes and 25 seconds. 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	Our attack		[7]	
signatures	Time	Success (%)	Time	Success (%)
3	39 hours	0.2	−	−
4	1 hour 17 minutes	0.5	41 minutes	1.5
5	8 minutes 20 seconds	6.5	18 minutes	1
6	≈ 5 minutes	25	18 minutes	22
7	≈ 3 minutes	17.5	34 minutes	24
8	≈ 2 minutes	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.

For the sake of completeness, we mention that in [25], 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 seconds, hence much

---

[3] 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 hours instead of 18 minutes.

[4] 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.

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 seconds 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 hours, 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 | Our attack | | [7] | |
signatures	Success (%)	Time	Success (%)	Time
3	0.2	39 hours	–	–
4	4	25 hours 28 minutes	1.5	41 minutes
5	20	2 hours 42 minutes	4	36 minutes
6	40	1 hour 4 minutes	35	1 hour 43 minutes
7	45	2 hours 36 minutes	68	3 hours 58 minutes
8	45	5 hours 2 minutes	–	–

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.

Looking at Coppersmith. Finally, as the EHNP setup consists of a system of modular equations for which we look for integer roots, we investigate the use of Coppersmith's method for finding small roots of integer polynomials, as an alternative to EHNP. Albeit our attempts to apply Coppersmith's method were not successful as some bound on the unknowns is not satisfied, we briefly sketch our ideas hoping it could lead to further improvements.

Organization: In Section 2, we introduce some background on ECDSA and the wNAF representation. Moreover, we give details on lattices and well known reduction algorithms. In Section 3, we explain how the Extended Hidden Number problem can be transformed into a lattice problem. We explicit the lattice basis and give some analysis on the length of the short vectors found in the reduced basis. In Section 4, we introduce our preprocessing method which allows us to reduce the overall time of our attack. In Section 5, we give experimental results which shows the performance of our attack as a function of the various parameters that are being considered. In Section 6, we analyze the resilience of our attack to erroneous traces. Finally in Section 7, we describe our attempt to use Coppersmith's method instead of EHNP.

2   Preliminaries

2.1   Elliptic Curves Digital Signature Algorithm

The Elliptic Curves Digital Signature Algorithm is a variant of the Digital Signature Algorithm, DSA, [21] which uses elliptic curves instead of finite fields. The parameters used in the ECDSA algorithm 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 $x \equiv r \pmod{q}$.

Remark 1. In this paper, we consider a 128 bit level of security and thus α, q and k are all 256-bit integers.

2.2 wNAF representation

The ECDSA algorithm presented above requires the computation of $[k]G$ which corresponds to a scalar multiplication. In [12], 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, as we will explain now. 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.

> **Input** : $k \in \mathbb{Z}^+$
> **Output:** NAF representation of k
> $i = 0$;
> **while** $k > 0$ **do**
>   **if** $k \pmod 2 = 1$ **then**
> $k_i = 2 - (k \pmod 4)$;
> $k = k - k_i$;
>   **else**
> $k_i = 0$;
>   **end**
> $k = k/2$;
> $i = i + 1$;
> **end**
> **return** $k_{i-1}, k_{i-2}, \ldots, k_1, k_0$

Algorithm 1: NAF algorithm

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 for example), the window size usually chosen was $w = 3$. The scalar k is converted into wNAF form using Algorithm 2. Note that in Algorithm 2, the sequence of digits m_i belongs to the set $\{0, \pm 1, \pm 3, \ldots, \pm (2^w -$

Input: $k \in \mathbb{Z}^+$, $w \in \mathbb{N}$
Output: $(m_0, m_1, \ldots, m_n)$, *i.e.*, k in its wNAF representation
$i = 0$;
while $k > 0$ **do**
$\quad$ **if** $k \pmod 2 = 1$ **then**
$\quad\quad$ $m_i = k \pmod{2^{w+1}}$;
$\quad\quad$ **if** $m_i \geq 2^w$ **then**
$\quad\quad\quad$ $m_i = m_i - 2^{w+1}$;
$\quad\quad$ **end**
$\quad\quad$ $k = k - m_i$;
$\quad$ **else**
$\quad\quad$ $m_i = 0$;
$\quad$ **end**
$\quad$ $k = k/2$;
$\quad$ $i = i + 1$;
end

Algorithm 2: wNAF representation

1)}. We can rewrite k as a sum of its non-zero digits, which we rename 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

$$23 = 2^5 - 2^3 - 2^0 = (1, 0, -1, 0, 0, -1).$$

Using a window size $w = 3$, the wNAF representation gives

$$23 = 2^4 + 7 \times 2^0 = (1, 0, 0, 0, 7).$$

There exists a modified wNAF representation used in OpenSSL for example. In the non-modified wNAF representation, at most one of any $w + 1$ consecutive digits is non-zero and in the modified version, this also stands with the exception that the most significant digit may be only $w - 1$ zeros away from that next non-zero digit.

2.3 Lattice reduction algorithms

A lattice is a discrete additive subgroup of $\mathbb{R}^n$. It is usually specified by giving 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 some $\lambda_1 > 0$ such that the length of one of the shortest non-zero vectors v_1 in the lattice satisfies $||v_1|| = \lambda_1$. Let λ_i be the i^{th} successive minimum of the lattice. The LLL algorithm [19] 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 $b_1 \leq 1.02^n \lambda_1$ as noted by Nguyen and Stehlé [22]. Moreover, for random lattices, we note that the Gaussian heuristic implies that

$$\lambda_1 \approx \sqrt{n/(2\pi e)} \det(L)^{1/n}. \tag{1}$$

The BKZ algorithm [26, 28] 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 $||b_i||_2 \leq i\gamma_\beta^{(n-i)/(k-1)}\lambda_i$ [27], 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), introduced in 1996 [4], 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 [14], 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

$$a_i\alpha + \sum_{j=1}^{\ell_i} b_{i,j}k_{i,j} \equiv c_i \pmod{q}, \tag{2}$$

where the secret α and $0 \leqslant k_{i,j} \leqslant 2^{\eta_{ij}}$ are unknown, and the values η_{ij}, a_i, $b_{i,j}$, c_i, ℓ_i are all known for $1 \leqslant i \leqslant u$ (see [14], Definition 3), one has to recover α in polynomial time. Similarly to the HNP, the EHNP can be transformed into a lattice problem and one can recover 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

$$\alpha r = sk - H(m) \pmod{q}. \tag{3}$$

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 satisfies Equation (3). 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 $k_j \in \{\pm 1, \pm 3, \ldots, \pm(2^w - 1)\}$ 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 $-2^{w-1} \leqslant k'_j \leqslant 2^{w-1} - 1$. Using the same notations as in [8], 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

$$k = \sum_{j=1}^{\ell} k_j 2^{\lambda_j} = \bar{k} + \sum_{j=1}^{\ell} d_j 2^{\lambda_j + 1}, \tag{4}$$

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 Equation (4), we obtain a system of modular equations of the form

$$\alpha r_i - \sum_{j=1}^{\ell_i} 2^{\lambda_{i,j} + 1} s_i d_{i,j} - (s_i \bar{k}_i - H(m_i)) \equiv 0 \pmod{q} \tag{5}$$

where the unknowns are α and the $d_{i,j}$. The known values are the 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. We can then use Equation (5) as input to the Extended Hidden Number Problem, following the method explained in [14]. The problem of finding the secret key is then reduced to solving the short vector problem in a given lattice which we give in the following section.

3.3 Constructing the lattice

Before giving the lattice basis construction, we redefine Equation (5) 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 [8].

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 Equation (5). 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

$$\sum_{j=1}^{\ell_1} \underbrace{(2^{\lambda_{1,j}+1} s_1 r_i)}_{:=\tau_{j,i}} d_{1,j} + \sum_{j=1}^{\ell_i} \underbrace{(-2^{\lambda_{i,j}+1} s_i r_1)}_{:=\sigma_{i,j}} d_{i,j}$$
$$- \underbrace{r_1(s_i \bar{k}_i - H(m_i)) + r_i(s_1 \bar{k}_1 - H(m_1))}_{:=\gamma_i} \equiv 0 \pmod{q}.$$
$$(6)$$

Again, using the same notations as in [8], 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 $2 \leqslant i \leqslant u$, $1 \leqslant j \leqslant \ell_i$. Even if α is eliminated from the equations, if we are able to recover some $d_{i,j}$ values from a short vector in the lattice, we can recover α using any equation in the modular system (5). We will now use Equation (6) 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 $1 \leqslant j \leqslant \ell_i$ and $2 \leqslant i \leqslant u$. We set a scaling factor $\Delta \in \mathbb{N}$ to be defined later. The lattice basis is given by

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 Equation (6) 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

$$\det \mathcal{L} = q^{u-1} \left(\Delta 2^m\right)^{u-1} 2^{\sum_{i,j}(m-\mu_{i,j})} 2^{m-1}.$$

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

$$w = v\mathcal{B}, \tag{7}$$

and

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

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

$$||w||_2 \leqslant (1.02)^n (\det L)^{1/n}. \tag{8}$$

The vector w we expect to find has norm $||w||_2 \leqslant 2^{m-1}\sqrt{T+1}$. From Equation (8), one can deduce the value of Δ needed to find w in the reduced lattice, which is given by the expression

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

In our experiments, the average value of ℓ_i for $1 \leqslant i \leqslant u$ 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 $u = 3, \ldots, 8$, with these values, we obtain $\Delta_{th} \ll 1$, which does not help us to set this parameter.

In practice, we verify that setting $\Delta = 1$ allows us to recover the secret key. In our experiments, we vary the bitsize of Δ to see whether a slightly larger value affects the probability of success. This comment will be adressed in Section 5.

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

$$\begin{aligned}
\sigma_{i,j+1} &= -2^{\lambda_{i,j+1}+1} \cdot s_i r_1 \\
&= -2^{\lambda_{i,j+1}+1} \cdot \left(\frac{\sigma_{i,j}}{-2^{\lambda_{i,j+1}+1}} \right) \\
&= 2^{\lambda_{i,j+1}-\lambda_{i,j}} \cdot \sigma_{i,j}
\end{aligned}$$

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

$$(0, \cdots, 0, -2^{\lambda_{i,j+1}-\lambda_{i,j}+m-\mu_{i,j}}, 2^{m-\mu_{i,j+1}}, 0, \cdots, 0). \tag{9}$$

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

Remark 2. It would be of interest to understand how one can modify the lattice construction to always find w as the shortest vector of the reduced basis. Indeed, by reducing the number of vectors shorter than w we expect to increase the probability of success of our attack. This would lower the chances of w being a linear combination of short vectors and thus not appearing in the reduced basis.

Differences with the lattice construction given in [8]. Let $\mathcal{B}'$ be the lattice basis constructed in [8]. 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 [8] have a value δ in their construction which corresonds to $1/\Delta$. In this work, we give a precise analysis of the value of Δ, both theoretically and experimentally in Section 5, which is missing in [8].

4   Improving the lattice attack

4.1   Reducing the lattice dimension: the merging technique

In [8], 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 $1 \leqslant j \leqslant \ell'$. 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.*, $0 \leqslant d'_j \leqslant 2^{\mu_j} - 1$.

We present here our merging algorithm based on Algorithm 3 given in [8]. Our algorithm modifies directly the sequence $\{\lambda_j\}_{j=1}^{\ell}$, whereas [8] work on the double-and-add chains. This helped us avoid some implementation issues such as an index outrun present in Algorithm 3 [8], line 7. To facilitate the ease of

reading of (our) Algorithm 3, we work with dynamic tables. To do so, we 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 $n \geqslant 1$, a window size w.
Output: $v_{\lambda'}$, a table of size $n' \leqslant n$ 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 [8]. We give in Table 3 the approximate dimension of the lattices we obtain using the elimination and merging techniques. 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.

Table 3: Average dimensions of the lattices after merging.

Number of signatures	Average dimension
3	80
4	110
5	135
6	160
7	190
8	215

Remark 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 $1 \leqslant j \leqslant \ell$. 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 $S_a = \{t \in \mathcal{T} \,|\, \max_j \mu_j \leqslant a, 1 \leqslant j \leqslant \ell\}$. 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 $2 \leqslant i \leqslant u$, we force the coefficients to remain smaller than when taking traces at random.

In practice, we work with a set $\mathcal{T}$ of 2000 traces such that $\min_{t \in \mathcal{T}} \max_j \mu_j = 11$ and $\max_{t \in \mathcal{T}} \max_j \mu_j = 67$. We consider the sets S_{11}, S_{15} and S_{19} in our experiments. In Table 4, we give the proportion of signatures corresponding to the different preprocessing subsets.

Table 4: Proportion of preprocessing subsets.

Preprocessing	Proportion (%)
S_{11}	2
S_{15}	18
S_{19}	44

The effect of preprocessing on the total time of the attack is explained in Section 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 hyperthreading, as it can virtually recover the most amount of information.[5]

---

[5] 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.

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 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 [29] 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 5, 6 when using BKZ. [6]

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 minutes using 8 signatures for BKZ and at best 5 minutes for G6K.

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 minutes. 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 minutes on average. This is an improvement by a factor of 3 in the reduction time.

Experimentally, we vary the parameters that are considered in the attack: the bitsize of Δ, the preprocessing subset and the block-size used in BKZ.

Only 3 signatures. Using $\Delta \approx 2^3$ and no preprocessing, we were able to recover the secret key using 3 signatures with BKZ-35 only once and three times with

---

[6] In [7], the authors use an Intel Core i7-3770 CPU running at 3.40GHz 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 5: Fastest key recovery with respect to the number of signatures.

| Number of | Total | Parameters | | | Probability of |
signatures	time	BKZ	Preprocessing	Δ	success (%)
3	39 hours	35	S_{11}	$\approx 2^3$	0.2
4	1 hour 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

Table 6: Highest probability of success with respect to the number of signatures.

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

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 been too much time consuming to do. For this reason, in the rest of this section, the number of signatures only vary 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 Figure 1, we analyze the total time to recover the secret key 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. Details of the experiments are given in Appendix A.

Analyzing the effect of preprocessing. We also analyze the influence of our pre-processing 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. We report results for $\Delta \approx 2^3$ in Appendix B. We still gain time using the preprocessing but less than with $\Delta \approx 2^{25}$.
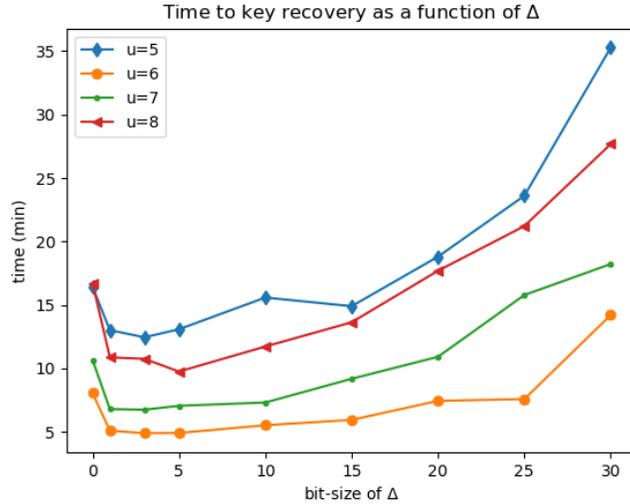
Fig. 1: Analyzing the overall time to recover the secret key as a function of the bitsize of Δ. We report the numbers BKZ-25 and no preprocessing. The optimal value for Δ is around 2^3 except for $u = 8$ where it is 2^5.

The effect of preprocessing is difficult to predict since its behavior varies a lot 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 vectors[7] 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 on average.

The positive effect of preprocessing is most noticeable for $u = 4$ and $u = 5$, as shown in Figure 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 impacful. Details are given in Appendix B.

---

[7] In the sense of vectors exhibited in (9).
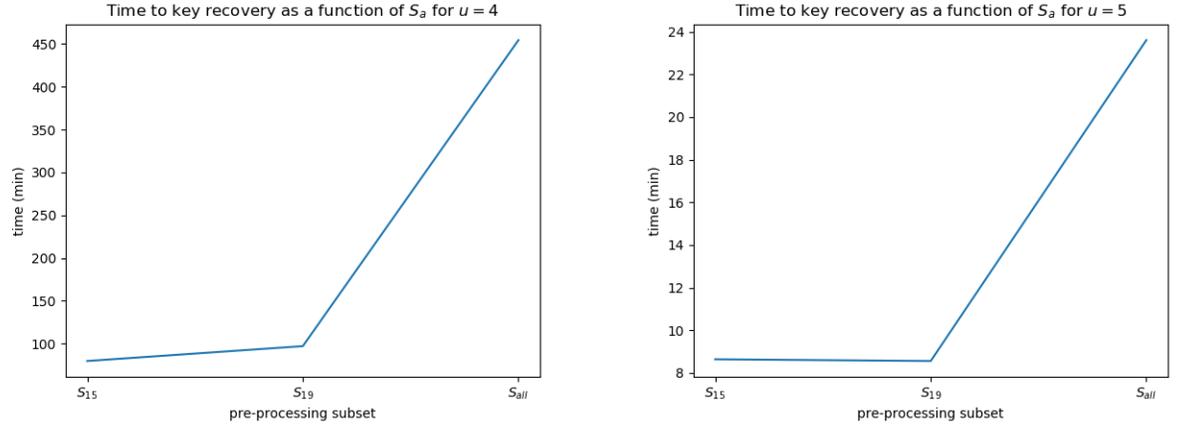
Fig. 2: Analyzing the overall time to recover the secret key as a function of the preprocessing subset for 4 and 5 traces. The other parameters are fixed: $\Delta \approx 2^{25}$ and BKZ-25.

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 Figure 3 for 6 and 7 signatures. For other values of u, the plot is very similar and we omit them in Figure 3 for ease of lecture. 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 5 and Table 6: 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. Details are given in Appendix C.

6   Error resilience analysis

It is not unexpected to have errors in the traces collected during the side-channel attack. Obtaining a set of 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 [11, 2]. In [6], it is shown that the lattice attack still successfully recovers the secret key even when some 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
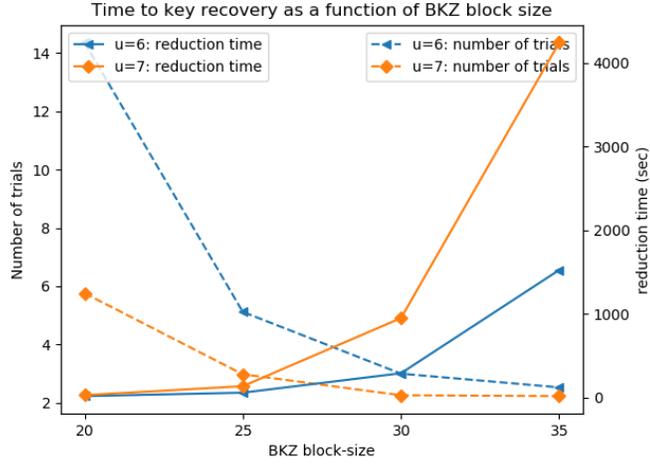
Fig. 3: Analyzing the number of trials to recover the secret key and the reduction time of the lattice as a function of the block-size of BKZ. We consider the cases where $u = 6$ and $u = 7$. The dotted lines correspond to the number of trials, and the continued lines to the reduction time in seconds.

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 [25] and [6] give some analysis of the attacks FLUSH&RELOAD and PRIME&PROBE in real life scenarios.

In [8], 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 [25]. 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. The following Table 7 shows the probability of success of the attack in the presence of errors. We ran experiments for BKZ-25 using $\Delta \approx 2^3$ and 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.

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

Table 7: Error analysis using BKZ-25, $\Delta \approx 2^3$ and S_{all}.

Number of	Probability of success (%)				
signatures	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	$\ll 1$

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 Figure 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, see Table 10 in Appendix D.

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 Figure 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. More details on the probabilities of success of these two types of errors can be seen in Appendix D.

If the signal processing method is hesitant between a 1 or 0 digit, we would recommend to favor putting 1 instead of 0 to increase the chance of having an error of type $0 \to 1$, for which the attack is a lot more tolerant.

7   An attempt at using Coppersmith's methods

Given that the setup of the Extended Hidden Number Problem gives a system of modular equations with the unknowns $(\alpha, d_{1,1}, \cdots, d_{u,l_u})$, it is natural to ask whether this system can be solved using Coppersmith's method for finding small modular roots of integer polynomials. Admittedly, α is of the order of magnitude of q so not so small, but by small root we mean that it is sufficient to know a bound on each variable.

Coppersmith's methods in the case of bivariate polynomials can be expressed as the following theorem [9, Theorem 19.2.1]. It states that a small modular root of a bivariate polynomial can be detected as an integer root of other polynomials.

Theorem 1. *Let $F(x_1, x_2) \in \mathbb{Z}[x_1, x_2]$ be a polynomial of total degree n. Let $X_1, X_2, q \in \mathbb{N}$ be such that $X_1 X_2 < q^{1/n - \epsilon}$ for some $0 < \epsilon < 1/n$. Then*
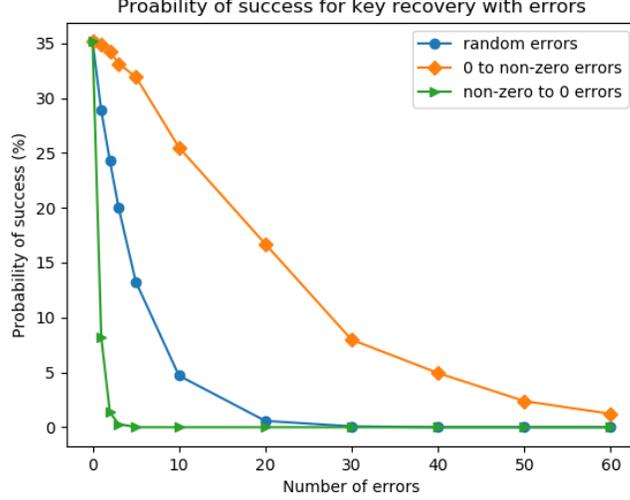
Fig. 4: Probability of success for key recovery with various types of errors when using $u = 8$, BKZ-25, $\Delta \approx 2^3$, and no preprocessing.

one can compute in time polynomial in $\log(q)$ *and* $1/\epsilon > n$ *two polynomials* $F_1(x_1, x_2)$ *and* $F_2(x_1, x_2) \in \mathbb{Z}[x_1, x_2]$ *such that for all* $(x_1^{(0)}, x_2^{(0)}) \in \mathbb{Z}^2$ *with* $|x_1^{(0)}| < X_1, |x_2^{(0)}| < X_2$ *and* $F(x_1^{(0)}, x_2^{(0)}) \equiv 0 \pmod{q}$, *one has* $F_1(x_1^{(0)}, x_2^{(0)}) = F_2(x_1^{(0)}, x_2^{(0)}) = 0$ *over* $\mathbb{Z}$.

Theorem 1 is generalized to m variables in [17]. Let $|x_i^{(0)}| < X_i$ for $i = 1, \cdots, m$. In this setting, the condition $X_1 \cdot X_2 < q^{1/n-\epsilon}$ for some $0 < \epsilon < 1/n$ is replaced by $X_1 \cdot X_2 \cdots X_m < q^{1/n-\epsilon}$.

In our setup. We consider the system of $u - 1$ modular equations after elimination, with $T = \sum_{i=2}^{u} \ell_i$ variables. This allow us to have one less unknown and to avoid having to recover α which would be our largest variable as it is of the same order of magnitude of q whereas $d_{i,j} < 2^{\mu_{i,j}}$.

We have the following equations

$$F_i(d_{1,1}, \ldots, d_{1,\ell_1}, \ldots d_{u,1}, \ldots, d_{u,\ell_u}) = \sum_{j=1}^{\ell_1} \tau_{j,i} d_{1,j} + \sum_{j=1}^{\ell_i} \sigma_{i,j} d_{i,j} - \gamma_i \equiv 0 \pmod{q}$$

for $2 \leqslant i \leqslant u$, and where τ_{ji}, σ_{ij} and γ_i are defined as in Section 3.3. This system has $u - 1$ equations and T unknowns. Note that F_i is a linear polynomial and its total degree is $n = 1$.

Let D be a bound on the unknowns d_{ij}, *i.e.*, $|d_{ij}| < D$. The condition in the theorem requires that

$$D^T < q^{1-\epsilon}$$

which means $D < q^{(1-\epsilon)/T}$. When $\epsilon \to 1$, we get that $D < 1$, and when $\epsilon \to 0$, we have $D < q^{1/T}$. If we consider the attack scenario where the number of signatures $u \in [3, 8]$, the value of T grows with u. For $u = 3$, the value T is of the order of 150 on average. This results in the condition $D \leqslant 3$. But restricting the bound on the d_{ij} to 3 at best seems too restrictive for the key recovery to be successful. Indeed, it means that the algorithm will miss all the solutions with at least one $3 < d_{i,j} < 2^{\mu_{i,j}}$.

Remark 4. We also considered the equations without elimination, *i.e.*, keeping the variable α. However, this resulted in an even stronger condition of D ($D < 1$) due to the size of α.

Remark 5. The theorem mentioned above is one of the many variations of Coppersmith's method. The proof of the theorem relies on the construction of a lattice whose coefficients correspond to the coefficients of the polynomials for which we want to find a modular root. The idea is to use LLL on this lattice to construct new polynomials with small coefficients, small enough so that the expected modular root is in truth a root of these new polynomials over the integers.

We have tested various construction for the lattice basis. In particular, we give our lattice construction in the elimination case in Appendix E. However, none of our constructions have allowed us to successfully recover the secret key.

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 [20], much more resilient to side-channel attacks due to the regularity of the operations. However, this does not entirely remove the risk of leakage[34]. Additional countermeasures are necessary.

When looking at common countermeasures, many implementations use blinding or masking techniques [24], 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

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: INDOCRYPT. Lecture Notes in Computer Science, vol. 10698, pp. 147–165 (2017)
3. Benger, 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 (Sep 2014).
4. Boneh, D., Venkatesan, R.: Hardness of computing the most significant bits of secret keys in Diffie-Hellman and related schemes. In: Koblitz, N. (ed.) CRYPTO'96. LNCS, vol. 1109, pp. 129–142. Springer, Heidelberg (Aug 1996).
5. Chen, Y., Nguyen, P.Q.: BKZ 2.0: Better lattice security estimates. In: ASIACRYPT. Lecture Notes in Computer Science, vol. 7073, pp. 1–20. Springer (2011)
6. Dall, F., De Micheli, G., Eisenbarth, T., Genkin, D., Heninger, N., Moghimi, A., Yarom, Y.: 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 (Oct 2016).
8. Fan, S., Wang, W., Cheng, Q.: Attacking OpenSSL implementation of ECDSA with a few signatures. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security. pp. 1505–1515. CCS '16 (2016)
9. Galbraith, S.D.: Mathematics of Public Key Cryptography. Cambridge University Press, New York, NY, USA, 1st edn. (2012)
10. 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 (Aug 2017)
11. Genkin, D., Pachmanov, L., Pipman, I., Tromer, E., Yarom, Y.: ECDSA key extraction from mobile devices via nonintrusive physical side channels. In: CCS. pp. 1626–1638 (2016)
12. Gordon, D.M.: A survey of fast exponentiation methods. Journal of Algorithms **27**(1), 129–146 (Apr 1998)

13. Hlavác, 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 (Aug 2007).
14. Hlaváč, M., Rosa, T.: Extended hidden number problem and its cryptanalytic applications. In: Biham, E., Youssef, A.M. (eds.) Selected Areas in Cryptography. pp. 114–133. Springer Berlin Heidelberg, Berlin, Heidelberg (2007)
15. Howgrave-Graham, N.A., Smart, N.P.: Lattice attacks on digital signature schemes. Designs, Codes and Cryptography **23**(3), 283–290 (2001)
16. Johnson, D., Menezes, A., Vanstone, S.: The elliptic curve digital signature algorithm (ECDSA). International Journal of Information Security **1**(1), 36–63 (2001)
17. Jutla, C.S.: On finding small solutions of modular multivariate polynomial equations. In: Nyberg, K. (ed.) Advances in Cryptology — EUROCRYPT'98. pp. 158–170. Springer Berlin Heidelberg, Berlin, Heidelberg (1998)
18. Kocher, P., Jaffe, J., Jun, B.: Differential power analysis. In: Wiener, M. (ed.) Advances in Cryptology — CRYPTO' 99. pp. 388–397. Springer Berlin Heidelberg, Berlin, Heidelberg (1999)
19. Lenstra, A.K., Lenstra, H.W., Lovász, L.: Factoring polynomials with rational coefficients. Mathematische Annalen **261**(4), 515–534 (1982)
20. Montgomery, P.L.: Speeding the pollard and elliptic curve methods of factorization. Mathematics of Computation **48**(177), 243–243 (jan 1987)
21. National Institute of Standards and Technology: Digital Signature Standard (DSS) (2013)
22. Nguyen, P., Stehlé, D.: LLL on the average. In: Proceedings of the 7th International Conference on Algorithmic Number Theory. pp. 238–256. ANTS'06, Springer-Verlag, Berlin, Heidelberg (2006)
23. Nguyen, P.Q., Shparlinski, I.E.: The insecurity of the elliptic curve digital signature algorithm with partially known nonces. Designs, Codes and Cryptography **30**(2), 201–217 (Sep 2003)
24. 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 (Feb 2006).
25. Van de Pol, J., Smart, N.P., Yarom, Y.: Just a little bit more. In: Nyberg, K. (ed.) Topics in Cryptology — CT-RSA 2015. pp. 3–21. Springer International Publishing, Cham (2015)
26. Schnorr, C.P.: A hierarchy of polynomial time lattice basis reduction algorithms. Theoretical Computer Science **53**(2-3), 201–224 (1987)
27. Schnorr, C.P.: Block reduced lattice bases and successive minima. Combinatorics, Probability & Computing **3**, 507–522 (1994)
28. Schnorr, C.P., Euchner, M.: Lattice basis reduction: Improved practical algorithms and solving subset sum problems. Mathematical Programming **66**(2), 181–199 (1994)
29. The FPLLL development team: FPLLL, a lattice reduction library (2016)
30. 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, Heidelberg (Apr 2015).
31. Vanstone, S.: Responses to NIST's proposals (1992)
32. Wang, W., Fan, S.: Attacking OpenSSL ECDSA with a small amount of side-channel information. Science China Information Sciences **61**(3), 032105 (2017)
33. Yarom, Y., Benger, N.: Recovering OpenSSL ECDSA nonces using the FLUSH+RELOAD cache side-channel attack. IACR Cryptology ePrint Archive **2014**, 140 (2014)

34. Yarom, Y., Benger, N.: Recovering openssl ecdsa nonces using the flush+reload cache side-channel attack. IACR Cryptology ePrint Archive **2014**, 140 (2014)

35. 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. pp. 719–732. SEC'14, USENIX Association, Berkeley, CA, USA (2014)

A   Bitsize of Δ effect over the key recovery total time

We analyze the effect of the bitsize of Δ. We fix BKZ-25 and use no preprocessing. We average over 5000 experiments.

Parameters				Results		
u	BKZ-β	Preprocessing	Δ bitsize	Probability of success (%)	Time of one experiment (sec)	Total time to key recovery (min)
4	25	S_{all}	0	0.14	31	375
4	25	S_{all}	1	0.16	31	330
4	25	S_{all}	3	0.28	32	**191**
4	25	S_{all}	5	0.22	30	234
4	25	S_{all}	10	0.24	33	228
4	25	S_{all}	15	0.16	39	411
4	25	S_{all}	20	0.20	45	379
4	25	S_{all}	25	0.20	54	454
4	25	S_{all}	30	0.10	31	515
5	25	S_{all}	0	3.74	37	16
5	25	S_{all}	1	4.60	36	13
5	25	S_{all}	3	4.58	34	**12**
5	25	S_{all}	5	4.38	34	13
5	25	S_{all}	10	3.92	36	15
5	25	S_{all}	15	4.62	41	15
5	25	S_{all}	20	4.60	52	19
5	25	S_{all}	25	4.52	64	23
5	25	S_{all}	30	4.18	88	35

Parameters				Results		
u	BKZ-β	Preprocessing	Δ bitsize	Probability of success (%)	Time of one experiment (sec)	Total time to key recovery (min)
6	25	S_{all}	0	15.94	77	8
6	25	S_{all}	1	19.96	61	5
6	25	S_{all}	3	19.52	57	**5**
6	25	S_{all}	5	20.10	59	5
6	25	S_{all}	10	19.04	63	5
6	25	S_{all}	15	20.34	72	6
6	25	S_{all}	20	20.58	92	7
6	25	S_{all}	25	20.02	91	7
6	25	S_{all}	30	19.26	164	14
7	25	S_{all}	0	28.86	185	10
7	25	S_{all}	1	33.00	134	7
7	25	S_{all}	3	33.54	136	**6**
7	25	S_{all}	5	33.69	142	7
7	25	S_{all}	10	33.99	149	7
7	25	S_{all}	15	33.81	186	9
7	25	S_{all}	20	34.94	229	11
7	25	S_{all}	25	31.68	300	15
7	25	S_{all}	30	32.08	351	18
8	25	S_{all}	0	32.12	322	16
8	25	S_{all}	1	36.40	237	101
8	25	S_{all}	3	35.14	227	10
8	25	S_{all}	5	36.00	211	**9**
8	25	S_{all}	10	34.86	245	11
8	25	S_{all}	15	36.18	296	13
8	25	S_{all}	20	35.48	376	17
8	25	S_{all}	25	36.12	460	21
8	25	S_{all}	30	34.54	573	27

B   Preprocessing effect over the key recovery total time

We analyze the effect of the preprocessing. We fix BKZ-25 and $\Delta \approx 2^3, 2^{25}$. We average over 5000 experiments.

Parameters				Results		
u	BKZ-β	Preprocessing	Δ bitsize	Probability of success (%)	Time of one experiment (sec)	Total time to key recovery (min)
4	25	S_{11}	25	0.20	9	**79**
4	25	S_{15}	25	0.52	24	79
4	25	S_{19}	25	0.50	29	97
4	25	S_{all}	25	0.20	54	454
5	25	S_{11}	25	1.70	17	17
5	25	S_{15}	25	5.74	29	8
5	25	S_{19}	25	6.28	32	**8**
5	25	S_{all}	25	4.52	64	23
6	25	S_{11}	25	3.64	38	17
6	25	S_{15}	25	22.12	77	5
6	25	S_{19}	25	25.12	77	**5**
6	25	S_{all}	25	20.02	91	7

Parameters				Results		
u	BKZ-β	Preprocessing	Δ bitsize	Probability of success (%)	Time of one experiment (sec)	Total time to key recovery (min)
7	25	S_{11}	25	3.40	55	27
7	25	S_{15}	25	26.20	151	9
7	25	S_{19}	25	43.90	162	**7**
7	25	S_{all}	25	31.68	300	15
8	25	S_{11}	25	4.50	85	31
8	25	S_{15}	25	32.50	237	12
8	25	S_{19}	25	43.90	267	**10**
8	25	S_{all}	25	36.12	460	21

Parameters				Results		
u	BKZ-β	Preprocessing	Δ bitsize	Probability of success (%)	Time of one experiment (sec)	Total time to key recovery (min)
4	25	S_{11}	3	0.18	9	89
4	25	S_{15}	3	0.52	24	**77**
4	25	S_{19}	3	0.38	29	130
4	25	S_{all}	3	0.28	32	191
5	25	S_{11}	3	1.18	19	27
5	25	S_{15}	3	5.90	30	8
5	25	S_{19}	3	6.50	32	**8**
5	25	S_{all}	3	4.58	34	12
6	25	S_{11}	3	4.04	40	16
6	25	S_{15}	3	20.36	78	6
6	25	S_{19}	3	24.76	72	**4**
6	25	S_{all}	3	19.52	57	5
7	25	S_{11}	3	4.15	60	24
7	25	S_{15}	3	27.00	158	9
7	25	S_{19}	3	35.25	173	8
7	25	S_{all}	3	33.54	135	**6**
8	25	S_{11}	3	4.40	88	33
8	25	S_{15}	3	35.20	249	11
8	25	S_{19}	3	40.70	268	11
8	25	S_{all}	3	35.14	227	**10**

C   BKZ block-size effect over the key recovery total time

We analyze the effect of the BKZ block-size. We set $\Delta \approx 2^3$ and use no preprocessing. We average over 5000 experiments.

Parameters				Results		
u	BKZ-β	Preprocessing	Δ bitsize	Probability of success (%)	Time of one experiment (sec)	Total time to key recovery (min)
4	20	S_{all}	3	0	5	0
4	25	S_{all}	3	0.28	32	**191**
4	30	S_{all}	3	1.30	302	387
4	35	S_{all}	3	4.10	3763	1529

Parameters				Results		
u	BKZ-β	Preprocessing	Δ bitsize	Probability of success (%)	Time of one experiment (sec)	Total time to key recovery (min)
5	20	S_{all}	3	0.82	9	19
5	25	S_{all}	3	4.58	34	**12**
5	30	S_{all}	3	11.60	225	32
5	35	S_{all}	3	20.18	1964	162
6	20	S_{all}	3	6.96	16	**4**
6	25	S_{all}	3	19.52	57	5
6	30	S_{all}	3	32.96	290	14
6	35	S_{all}	3	39.52	1525	64
7	20	S_{all}	3	17.35	28	**2**
7	25	S_{all}	3	33.54	136	6
7	30	S_{all}	3	44.20	950	35
7	35	S_{all}	3	44.80	4245	158
8	20	S_{all}	3	29.40	43	**2**
8	25	S_{all}	3	35.14	227	10
8	30	S_{all}	3	46.66	1894	68
8	35	S_{all}	3	44.70	8119	302

D   Analysis of errors

We analyze the effect of two possible kind of errors on the probability of success of our attack, using BKZ-25, $\Delta \approx 2^3$ and no preprocessing. We average over 5000 experiments. We write $\ll 1$ when the attack succeeded less than five times over 5000 experiments.

Table 8: Error $0 \rightarrow 1$ analysis using BKZ-25, $\Delta \approx 2^3$ and S_{all}.

Number of signatures	Probability of success (%)								
	0 errors	1 error	5 errors	10 errors	20 errors	30 errors	40 errors	50 errors	60 errors
4	0.28	0.18	0.10	$\ll 1$	0	0	0	0	0
5	4.58	3.82	2.70	1.06	0.32	$\ll 1$	0	0	0
6	19.52	10.79	13.88	7.90	2.94	0.86	0.36	0.10	$\ll 1$
7	33.54	31.06	26.04	18.36	9.24	4.54	?	1.02	0.50
8	35.14	34.92	31.94	25.50	16.70	7.96	4.94	2.48	1.22

Table 9: Error $1 \to 0$ analysis using BKZ-25, $\Delta \approx 2^3$ and S_{all}.

Number of	Probability of success (%)			
signatures	0 errors	1 error	2 errors	3 errors
4	0.28	0	0	0
5	4.58	0.36	$\ll 1$	0
6	19.52	2.70	0.36	$\ll 1$
7	33.54	5.54	1.00	0.12
8	35.14	8.20	1.36	0.30

When considering many errors, the probability of success can be increased by augmenting the block-size in the BKZ algorithm, as can be seen in Table 10.

Table 10: Errors $0 \to 1$ analysis with $\Delta \approx 2^3$, S_{all} and increasing block-size.

Number of	Probability of success (%)															
signatures	30 errors				40 errors				50 errors				60 errors			
	25	30	35	40	25	30	35	40	25	30	35	40	25	30	35	40
5	$\ll 1$	0.24	0.35	0.75	0	$\ll 1$	$\ll 1$	0.42	0	0	0	0	0	0	0	0
6	0.86	2.48	3.58	3.97	0.36	0.90	1.18	2.28	0.10	0.36	0.58	0.94	$\ll 1$	$\ll 1$	0.12	0.12
7	4.54	6.44	7.32	8.73	1.80	3.54	3.48	4.58	1.02	1.26	1.84	3.26	0.50	0.62	1.20	1.43
8	7.96	10.46	11.78	10.98	4.94	6.12	6.73	7.12	2.48	3.26	3.78	4.64	1.22	1.84	1.89	2.18

E   Lattice construction for Coppersmith's methods

We consider $u - 1$ equations given after elimination. We construct the following lattice basis

$$\begin{array}{ccccccccccccc}
d_{2,1} & d_{2,2} & \ldots & d_{2,\ell_2} & \ldots d_{u,1} \ldots & d_{u,\ell_u} & d_{1,1} & d_{1,2} & \ldots & d_{1,\ell_1} & 1 \\
\end{array}$$

$$\left(\begin{array}{ccccccccccc}
D & \sigma_{2,2}\sigma_{2,1}^{-1}D & \ldots & \sigma_{2,\ell_2}\sigma_{2,1}^{-1}D & 0 \quad 0 \quad 0 & 0 & \tau_{1,2}\sigma_{2,1}^{-1}D & \tau_{2,2}\sigma_{2,1}^{-1}D & \ldots & \tau_{\ell_1,2}\sigma_{2,1}^{-1}D & -\gamma_1\sigma_{2,1}^{-1} \\
& Dq & & & & & & & & & \\
& & \ddots & & & & & & & & \\
& & & Dq & & & & & & & \\
& & \ddots & D & \ldots \sigma_{u,\ell_u}\sigma_{u,1}^{-1}D & \tau_{1,u}\sigma_{u,1}^{-1}D & \tau_{2,u}\sigma_{u,1}^{-1}D & \ldots & \tau_{\ell_u,u}\sigma_{u,1}^{-1}D & -\gamma_u\sigma_{u,1}^{-1} \\
& & & Dq & & & & & & & \\
& & & & \ddots & & & & & & \\
& & & & & \ddots & & & & & \\
& & & & & & \ddots & & & & \\
& & & & & & & \ddots & & & \\
& & & & & & & & Dq & & \\
& & & & & & & & & & q \\
\end{array}\right)\begin{array}{l} Eq_2 \\ \\ \vdots \\ \\ Eq_u \\ \\ \\ \\ \\ \\ \\ \end{array}$$

The dimension of this lattice is $\dim L = T + 1$ and the determinant is given by

$$\det L = D^T q^{T-u+2}.$$

Coppersmith's method uses LLL to produce polynomials with integer roots equal to those from the inital modular equations and to do so, the lattice basis must satisfy $(1.02)^n (\det L)^{1/n} < q$, where $n = \dim \mathcal{L}$. This implies we need the condition

$$D < \left(\frac{q^{u-1}}{1.02^{(T+1)^2}}\right)^{1/T}.$$

Numerically, we get $D < 1$ for $u \in [3, 8]$.