

# Just A Little Bit More

Joop van de Pol<sup>1</sup>, Nigel P. Smart<sup>1</sup>, and Yuval Yarom<sup>2</sup>

<sup>1</sup> Dept. Computer Science, University of Bristol, United Kingdom.  
joop.vandepol@bristol.ac.uk, nigel@cs.bris.ac.uk

<sup>2</sup> School of Computer Science, The University of Adelaide, Australia.  
yval@cs.adelaide.edu.au

**Abstract.** We extend the FLUSH+RELOAD side-channel attack of Bengier et al. to extract a significantly larger number of bits of information per observed signature when using OpenSSL. This means that by observing only 25 signatures, we can recover secret keys of the **secp256k1** curve, used in the Bitcoin protocol, with a probability greater than 50 percent. This is an order of magnitude improvement over the previously best known result.

The new method of attack exploits two points: Unlike previous partial disclosure attacks we utilize all information obtained and not just that in the least significant or most significant bits, this is enabled by a property of the “standard” curves choice of group order which enables extra bits of information to be extracted. Furthermore, whereas previous works require direct information on ephemeral key bits, our attack utilizes the indirect information from the wNAF double and add chain.

## 1 Introduction

The Elliptic Curve Digital Signature Algorithm (ECDSA) is the elliptic curve analogue of the Digital Signature Algorithm (DSA). It has been well known for over a decade that the randomization used within the DSA/ECDSA algorithm makes it susceptible to side-channel attacks. In particular a small leakage of information on the ephemeral secret key utilized in each signature can be combined over a number of signatures to obtain the entire key.

Howgrave-Graham and Smart [14] showed that DSA is vulnerable to such partial ephemeral key exposure and their work was made rigorous by Nguyen and Shparlinski [21], who also extended these results to ECDSA [22]. More specifically, if, for a polynomially bounded number of random messages and ephemeral keys about  $\log^{1/2} q$  least significant bits (LSBs) are known, the secret key  $\alpha$  can be recovered in polynomial time. A similar result holds for a consecutive sequence of the most significant bits (MSBs), with a potential need for an additional leaked bit due to the paucity of information encoded in the most significant bit of the ephemeral key. When an arbitrary sequence of consecutive bits in the ephemeral key is known, about twice as many bits are required. The attack works by constructing a lattice problem from the obtained digital signatures and side-channel information, and then applying lattice reduction techniques such as LLL [16] or BKZ [23] to solve said lattice problem.

Brumley and co-workers employ this lattice attack to recover ECDSA keys using leaked LSBs (in [4]) and leaked MSBs (in [5]). The former uses a cache side-channel to extract the leaked information and the latter exploits a timing side-channel. In both attacks, a fixed number of bits from each signature is used and signatures are used only if the values of these bits are all zero. Signatures in which the value of any of these bits are one are ignored. Consequently, both attacks require more than 2,500 signatures to break a 160-bit private key.

More recently, again using a cache based side-channel, Bengier et al. [2] use the LSBs of the ephemeral key for a wNAF (a.k.a. sliding window algorithm) multiplication technique. By combining a new side-channel called the FLUSH+RELOAD side-channel [26, 27], and a more precise lattice attack strategy, which utilizes all of the leaked LSBs from every signature, Bengier et al. are able to significantly reduce the number of signatures required. In particular they report that the full secret key of a 256-bit system can be recovered with about 200 signatures in a reasonable length of time, and with a reasonable probability of success.

In this work we extend the FLUSH+RELOAD technique of Bengier et al. to reduce the number of required signatures by an order of magnitude. Our methodology abandons the concentration on extraction of bits in just the MSB and LSB positions, and instead focuses on all the information leaked by all the bits of the ephemeral key. In particular we

exploit a property of many of the standardized elliptic curves as used in OpenSSL. Our method, just as in [2], applies the FLUSH+RELOAD side-channel technique to the wNAF elliptic curve point multiplication algorithm in OpenSSL.

**ECDSA Using Standard Elliptic Curves:** The domain parameters for ECDSA are an elliptic curve  $E$  over a field  $\mathbb{F}$ , and a point  $G$  on  $E$ , of order  $q$ . Given a hash function  $h$ , the ECDSA signature of a message  $m$ , with a private key  $0 < \alpha < q$  and public key  $Q = \alpha G$ , is computed by:

- Selecting a random ephemeral key  $0 < k < q$
- Computing  $r = x(kG) \pmod{q}$ , the X coordinate of  $kG$ .
- Computing  $s = k^{-1}(h(m) + \alpha \cdot r) \pmod{q}$ .

The process is repeated if either  $r = 0$  or  $s = 0$ . The pair  $(r, s)$  is the signature.

To increase interoperability, standard bodies have published several sets of domain parameters for ECDSA [1, 7, 20]. The choice of moduli for the fields used in these standard curves is partly motivated by efficiency arguments. For example, all of the moduli in the curves recommended by FIPS [20] are generalised Mersenne primes [24] and many of them are pseudo-Mersenne primes [10]. This choice of moduli facilitates efficient modular arithmetic by avoiding a division operation which may otherwise be required.

A consequence of using pseudo-Mersenne primes as moduli is that, due to Hasse’s Theorem, not only is the finite-field order close to a power of two, but so is the elliptic-curve group order.

That is,  $q$  can be expressed as  $2^n - \varepsilon$ , where  $|\varepsilon| < 2^p$  for some  $p \approx n/2$ . We demonstrate that such curves are more susceptible to partial disclosure of ephemeral keys than was hitherto known. This property increases the amount of information that can be used from partial disclosure and allows for a more effective attack on ECDSA.

**Our Contribution:** We demonstrate that the above property of the standardized curves allows the utilization of far more leaked information, in particular some arbitrary sequences of consecutive leaked bits. In a nutshell, adding or subtracting  $q$  to or from an unknown number is unlikely to change any bits in positions between  $p + 1$  and  $n$ . Based on this observation we are able to use (for wNAF multiplication algorithms) all the information in consecutive bit sequences in positions above  $p + 1$ . Since in many of the standard curves  $p \approx n/2$ , a large amount of information is leaked per signature. (Assuming one can extract the sequence of additions and doubles in an algorithm.) As identified by Ciet and Joye [8] and exploited by Feix et al. [11], the same property also implies that techniques for mitigating side-channel attack, such as the scalar blinding suggested in [4, 18], do not protect bits in positions above  $p + 1$ .

Prior works deal with the case of partial disclosure of consecutive sequences of bits of the ephemeral key. Our work offers two improvements: It demonstrates how to use partial information leaked from the double and add chains of the wNAF scalar multiplication algorithm [13, 19]. In most cases, the double and add chain does not provide direct information on the value of bits. It only identifies sequences of repeating bits without identifying the value of these bits. We show how to use this information to construct a lattice attack on the private key. Secondly, our attack does not depend on the leaked bits being consecutive. We use information leaked through the double and add chain even though it is spread out along the ephemeral key.

By using more leaked information and exploiting the above property of the elliptic curves, our attack only requires a handful of leaked signatures to fully break the private key. Our experiments show that the perfect information leaked on double and add chains of only 13 signatures is sufficient for recovering the 256 bit private key of the **secp256k1** curve with probability greater than 50 percent. For the 521 bit curve **secp521r1**, 40 signatures are required. We further demonstrate that for the **secp256k1** case observing 25 signatures is highly likely to recover 13 perfect double and add chains. Hence, by observing 25 Bitcoin transactions using the same key, an attacker can expect to recover the private key. For most of the paper we discuss the case of perfect side channels which result in perfect double and add chains, then in Section 6 we show how this assumption can be removed in the context of a real FLUSH+RELOAD attack.

## 2 Background

In this section we discuss three basic procedures we will be referring to throughout. Namely the FLUSH+RELOAD side-channel attack technique, wNAF scalar multiplication method and the use of lattices to extract secret keys from triples. The side-channel information we obtain from executing the wNAF algorithm produces instances of the Hidden

Number Problem (HNP) [3]. Since the HNP is traditionally studied via lattice reduction it is therefore not surprising that we are led to lattice reduction in our analysis.

## 2.1 The FLUSH+RELOAD Side-Channel Attack Technique

FLUSH+RELOAD is a recently discovered cache side-channel attack [26, 27]. The attack exploits a weakness in the Intel implementation of the popular X86 architecture, which allows a spy program to monitor other programs’ read or execute access to shared regions of memory. The spy program only requires read access to the monitored memory.

Unlike most cache side-channel attacks, FLUSH+RELOAD uses the Last-Level Cache (LLC), which is the cache level closest to the memory. The LLC is shared by the execution cores in the processor, allowing the attack to operate when the spy and victim processes execute on different cores. Furthermore, as most virtual machine hypervisors (VMMs) actively share memory between co-resident virtual machines, the attack is applicable to virtualized environment and works cross-VM.

```

Input: adrs—the probed address
Output: true if the address was accessed by the victim
begin
  evict(adrs)
  wait_a_bit()
  time ← current_time()
  tmp ← read(adrs)
  readTime ← current_time()-time
  return readTime < threshold
end

```

**Algorithm 1:** FLUSH+RELOAD Algorithm

To monitor access to memory, the spy repeatedly evicts the contents of the monitored memory from the LLC, waits for some time and then measures the time to read the contents of the monitored memory. See Algorithm 1 for a pseudo-code of the attack. FLUSH+RELOAD uses the X86 `clflush` instruction to evict contents from the cache. To measure time the spy uses the `rdtsc` instruction which returns the time since processor reset measured in processor cycles.

As reading from the LLC is much faster than reading from memory, the spy can differentiate between these two cases. If, following the wait, the contents of memory is retrieved from the cache, it indicates that another process has accessed the memory. Thus, by measuring the time to read the contents of memory, the spy can decide whether the victim has accessed the monitored memory since the last time it was evicted.

To implement the attack, the spy needs to share the monitored memory with the victim. For attacks occurring within the same machine, the spy can map files used by the victim into its own address space. Examples of these files include the victim program file, shared libraries or data files that the victim accesses. As all mapped copies of files are shared, this gives the spy access to memory pages accessed by the victim. In virtualized environments, the spy does not have access to the victim’s files. The spy can, however, map copies of the victim files to its own address space, and rely on the VMM to merge the two copies using page de-duplication [15, 25]. It should be pointed that, as the LLC is physically tagged, the virtual address in which the spy maps the files is irrelevant for the attack. Hence, FLUSH+RELOAD is oblivious to address space layout randomization [17].

This sharing only works when the victim does not make private modifications to the contents of the shared pages. Consequently, many FLUSH+RELOAD attacks target executable code pages, monitoring the times the victim executes specific code. The spy typically divides time into fixed width time slots. In each time slot the spy monitors a few memory locations and records the times that these locations were accessed by the victim. By reconstructing a trace of victim access, the spy is able to infer the data the victim is operating on. Prior works used this attack to recover the private key of GnuPG RSA [27] as well as for recovering the ephemeral key used in OpenSSL ECDSA signatures either completely, for curves over binary fields [26], or partially, for curves over prime fields [2].

## 2.2 The wNAF Scalar Multiplication Method

Several algorithms for computing the scalar multiplication  $kG$  have been proposed. One of the suggested methods is to use the *windowed nonadjacent form* (wNAF) representation of the scalar  $k$ , see [13]. In wNAF a number is represented by a sequence of digits  $k_i$ . The value of a digit  $k_i$  is either zero or an odd number  $-2^w < k_i < 2^w$ , with each pair of non-zero digits separated by at least  $w$  zero digits. The value of  $k$  can be calculated from its wNAF representation using  $k = \sum 2^i \cdot k_i$ . See Algorithm 2 for a method to convert a scalar  $k$  into its wNAF representation. We use  $|\cdot|_x$  to denote the reduction modulo  $x$  into the range  $[-x/2, \dots, x/2)$ .

```

Input: Scalar  $k$  and window width  $w$ 
Output:  $k$  in wNAF:  $k_0, k_1, k_2 \dots$ 
begin
   $e \leftarrow k$ 
   $i \leftarrow 0$ 
  while  $e > 0$  do
    if  $e \bmod 2 = 1$  then
       $k_i \leftarrow |e|_{2^{w+1}}$ 
       $e \leftarrow e - k_i$ 
    else
       $k_i \leftarrow 0$ 
    end
     $e \leftarrow e/2$ 
     $i \leftarrow i+1$ 
  end
end

```

**Algorithm 2:** Conversion to Non-Adjacent Form

Let  $\bar{k}_i$  be the value of the variable  $e$  at the start of the  $i^{\text{th}}$  iteration in Algorithm 2. From the algorithm, it is clear that

$$k_i = \begin{cases} 0 & \bar{k}_i \text{ is even} \\ |\bar{k}_i|_{2^{w+1}} & \bar{k}_i \text{ is odd} \end{cases} \quad (1)$$

Furthermore:

$$k = 2^i \cdot \bar{k}_i + \sum_{j < i} 2^j \cdot k_j \quad (2)$$

Let  $m$  and  $m+l$  be the position of two consecutive non-zero wNAF digits, i.e.  $k_m, k_{m+l} \neq 0$  and  $k_{m+i} = 0$  for all  $0 < i < l$ . We now have

$$-2^{m+w} < \sum_{i \leq m} k_i \cdot 2^i < 2^{m+w}, \quad (3)$$

and because  $l > w$ , we get  $-2^{m+l-1} < \sum_{i \leq m+l-1} k_i \cdot 2^i < 2^{m+l-1}$ . Substituting  $m$  for  $m+l$  gives

$$-2^{m-1} < \sum_{i \leq m-1} k_i \cdot 2^i < 2^{m-1} \quad (4)$$

We note that for the minimal  $m$  such that  $k_m \neq 0$  we have  $\sum_{i \leq m-1} k_i \cdot 2^i = 0$ . Hence (4) holds for every  $m$  such that  $k_m \neq 0$ .

Because  $k_m$  is odd, we have  $-(2^w - 1) \leq k_m \leq 2^w - 1$ . Adding  $k_m \cdot 2^m$  to (4) gives a slightly stronger version of (3):

$$-(2^{m+w} - 2^{m-1}) < \sum_{i \leq m} k_i \cdot 2^i < 2^{m+w} - 2^{m-1} \quad (5)$$

One consequence of subtracting negative wNAF components is that the wNAF representation may be one digit longer than the binary representation of the number. For  $n$ -digits binary numbers Möller [19] suggests using  $k_i \leftarrow \lfloor k \rfloor_{2^w}$  when  $i = n - w - 1$  and  $e$  is odd, where  $\lfloor \cdot \rfloor_x$  denotes the reduction modulo  $x$  into the interval  $[0, \dots, x)$ . This avoids extending the wNAF representation in half the cases at the cost of weakening the non-adjacency property of the representation.

### 2.3 Lattice background

Before we describe how to get the necessary information from the side-channel attack, we recall from previous works what kind of information we are looking for. As in previous works [2, 4, 5, 14, 21, 22], the side-channel information is used to construct a lattice basis and the secret key is then retrieved by solving a lattice problem on this lattice. Generally, for a private key  $\alpha$  and a group order  $q$ , in previous works the authors somehow derive triples  $(t_i, u_i, z_i)$  from the side-channel information such that

$$-q/2^{z_i+1} < v_i = |\alpha \cdot t_i - u_i|_q < q/2^{z_i+1}. \quad (6)$$

Note that for arbitrary  $\alpha$  and  $t_i$ , the values of  $v_i$  are uniformly distributed over the interval  $[-q/2, q/2)$ . Hence, each such triple provides about  $z_i$  bits of information about  $\alpha$ . The use of a different  $z_i$  per equation was introduced in [2]. If we take  $d$  such triples we can construct the following lattice basis

$$B = \begin{pmatrix} 2^{z_1+1} \cdot q & & & & \\ & \ddots & & & \\ & & 2^{z_d+1} \cdot q & & \\ 2^{z_1+1} \cdot t_1 & \dots & 2^{z_d+1} \cdot t_d & & 1 \end{pmatrix},$$

whose rows generate the lattice that we use to retrieve the secret key. Now consider the vector  $\mathbf{u} = (2^{z_1+1} \cdot u_1, \dots, 2^{z_d+1} \cdot u_d, 0)$ , which consists of known quantities. Equation (6) implies the existence of integers  $(\lambda_1, \dots, \lambda_d)$  such that for the vectors  $\mathbf{x} = (\lambda_1, \dots, \lambda_d, \alpha)$  and  $\mathbf{y} = (2^{z_1+1} \cdot v_1, \dots, 2^{z_d+1} \cdot v_d, \alpha)$  we have

$$\mathbf{x} \cdot B - \mathbf{u} = \mathbf{y}.$$

Again using Equation (6), we see that the 2-norm of the vector  $\mathbf{y}$  is at most  $\sqrt{d \cdot q^2 + \alpha^2} \approx \sqrt{d+1} \cdot q$ . Because the lattice determinant of  $L(B)$  is  $2^{d+\sum z_i} \cdot q^d$ , the lattice vector  $\mathbf{x} \cdot B$  is heuristically the closest lattice vector to  $\mathbf{u}$ . By solving the Closest Vector Problem (CVP) on input of the basis  $B$  and the target vector  $\mathbf{u}$ , we obtain  $\mathbf{x}$  and hence the secret key  $\alpha$ .

There are two important methods of solving the closest vector problem: using an exact CVP-solver or using the heuristic embedding technique to convert it to a Shortest Vector Problem (SVP). Exact CVP-solvers require exponential time in the lattice rank ( $d+1$  in our case), whereas the SVP instance that follows from the embedding technique can sometimes be solved using approximation methods that run in polynomial time. Because the ranks of the lattices in this work become quite high when attacking a 521 bit key, we mostly focus on using the embedding technique and solving the associated SVP instance in this case.

The embedding technique transforms the previously described basis  $B$  and target vector  $\mathbf{u}$  to a new basis  $B'$ , resulting in a new lattice of dimension one higher than that generated by  $B$ :

$$B' = \begin{pmatrix} B & 0 \\ \mathbf{u} & q \end{pmatrix},$$

Following the same reasoning as above, we can set  $\mathbf{x}' = (\mathbf{x}, -1)$  and obtain the lattice vector  $\mathbf{y}' = \mathbf{x}' \cdot B' = (\mathbf{y}, -q)$ . The 2-norm of  $\mathbf{y}'$  is upper bounded by approximately  $\sqrt{d+2} \cdot q$ , whereas this lattice has determinant  $2^{d+\sum z_i} \cdot q^{(d+1)}$ . Note, however, that this lattice also contains the vector

$$(-t_1, \dots, -t_d, q, 0) \cdot B' = (0, \dots, 0, q, 0)$$

which will most likely be the shortest vector of the lattice. Still, our approximation algorithms for SVP work on bases and it is obvious to see that any basis of the same lattice must contain a vector ending in  $\pm q$ . Thus, it is heuristically likely that the resulting basis contains the short vector  $\mathbf{y}'$ , which reveals  $\alpha$ .

To summarize, we turn the side-channel information into a lattice and claim that, heuristically, finding the secret key is equivalent to solving a CVP instance. Then, we claim that, again heuristically, solving this CVP instance is equivalent to solving an SVP instance using the embedding technique. In Section 5 we will apply the attack to simulated data to see whether these heuristics hold up.

### 3 Using the wNAF Information

Assuming we have a side channel that leaks the double and add chain of the scalar multiplication. We know how to use the leaked LSBs [2]. These leaked LSBs carry, on average, two bits of information.

Given a double and add chain, the positions of the add operations in the chain correspond to the non-zero digits in the wNAF representation of the ephemeral key  $k$ . Roughly speaking, in half the cases the distance between consecutive non-zero digits is  $w + 1$ . In a quarter of the cases it is  $w + 2$  and so on. Hence, the average distance between consecutive non-zero digits is  $w + \sum_i i/2^i = w + 2$ . Since there are  $2^w$  non-zero digits, we expect that the double and add chain carries two bits of information per each non-zero digit position.

Reducing this information to an instance of the HNP presents three challenges:

- The information is not consecutive, but is spread along the scalar.
- Due to the use of negative digits in the wNAF representation, the double and add chain does not provide direct information on the bits of the scalar
- Current techniques lose half the information when the information is not at the beginning or end of the scalar.

As described in [2], the OpenSSL implementation departs slightly from the descriptions of ECDSA in Section 1. As a countermeasure to the Brumley and Tuveri remote timing attack [5], OpenSSL adds  $q$  or  $2 \cdot q$  to the randomly chosen ephemeral key, ensuring that  $k$  is  $n + 1$  bits long. While the attack is only applicable to curves defined over binary fields, the countermeasure is applied to all curves. Consequently, our analysis assumes that  $2^n \leq k < 2^{n+1}$ .

To handle non-consecutive information, we extract a separate HNP instance for each consecutive set of bits, and use these in the lattice. The effect this has on the lattice attack is discussed in Section 4.

To handle the indirect information caused by the negative digits in the wNAF representation we find a linear combination of  $k$  in which we know the values of some consecutive bits, we can use that to build an HNP instance.

Let  $m$  and  $m + l$  be the positions of two consecutive non-zero wNAF digits where  $m + l < n$ . From the definition of the wNAF representation we know that  $k = \overline{k_{m+l}}2^{m+l} + \sum_{i \leq m} k_i 2^i$ . We can now define the following values:

$$a = \frac{\overline{k_{m+l}} - 1}{2}$$

$$c = \sum_{i \leq m} k_i \cdot 2^i + 2^{m+w}$$

By (5) we have

$$2^{m-1} < c < 2^{m+w+1} - 2^{m-1} \quad (7)$$

From (2) we have

$$k - 2^{m+l} + 2^{m+w} = a \cdot 2^{m+l+1} + c$$

where  $0 \leq a < 2^{n-m-l}$  and because  $l \geq w + 1$  there are  $l - w$  consecutive zero bits in  $k - 2^{m+l} + 2^{m+w}$ .

In order to extract this information, we rely on a property of the curve where the group order  $q$  is close to a power of two. More precisely,  $q = 2^n - \varepsilon$  where  $|\varepsilon| < 2^p$  for  $p \approx n/2$ . We note that many of the standard curves have this property.

Let  $K = A \cdot 2^n + C$ , with  $0 \leq A < 2^{L_1}$  and  $2^{p+L_1} \leq C < 2^{L_1+L_2} - 2^{p+L_1}$ , note that this implies  $L_2 > p$ . Because  $q = 2^n - \varepsilon$  we get  $K - A \cdot q = K - A \cdot 2^n + A \cdot \varepsilon = C + A \cdot \varepsilon$ . Now,  $|\varepsilon| < 2^p$ . Consequently,  $0 \leq K - A \cdot q < 2^{L_1+L_2}$  and we get  $|K - 2^{L_1+L_2-1}|_q < 2^{L_1+L_2-1}$ . For  $p + 1 < m < n - l$  we can set

$$L_1 = n - m - l$$

$$L_2 = m + w$$

$$C = c \cdot 2^{n-m-l-1} = c \cdot 2^{L_1-1}$$

$$K = (k - 2^{m+l} + 2^{m+w}) \cdot 2^{n-m-l-1} = (k - 2^{m+l} + 2^{m+w}) \cdot 2^{L_1-1} = a \cdot 2^n + C$$

From (7) we obtain  $2^{L_1+m-2} < C < 2^{L_1+L_2} - 2^{L_1+m-2}$  which, because  $m \geq p - 2$ , becomes  $2^{p+L_1} < C < 2^{L_1+L_2} - 2^{p+L_1}$ . Thus, we have

$$\left| (k - 2^{m+l} + 2^{m+w}) \cdot 2^{n-m-l-1} - 2^{n-l+w-1} \right|_q < 2^{n-l+w-1}$$

Noting that  $k = \alpha \cdot r \cdot s^{-1} + h \cdot s^{-1} \pmod{q}$ , we can define the values

$$\begin{aligned} t &= \lfloor r \cdot s^{-1} \cdot 2^{n-m-l-1} \rfloor_q, \\ u &= \lfloor 2^{n+w-l-1} - (h \cdot s^{-1} + 2^{m+w} - 2^{m+l}) \cdot 2^{n-m-l-1} \rfloor_q, \\ v &= |\alpha \cdot t - u|_q. \end{aligned}$$

$|v| \leq 2^{n-l+w-1} \approx q/2^{l-w+1}$ , which gives us an instance of the HNP which carries  $l - w$  bits of information.

## 4 Heuristic Analysis

Now we know how to derive our triples  $t_i$ ,  $u_i$  and  $z_i$  that are used to construct the lattice. The next obvious question is: How many do we need before we can retrieve the private key  $\alpha$ ? Because the lattice attack relies on several heuristics, it is hard to give a definitive analysis. However, we will give heuristic reasons here, similar to those for past results.

Each triple  $(t_i, u_i, z_i)$  gives us  $z_i$  bits of information. If this triple comes from a pair  $(m, l)$  such that  $p+1 < m < n-l$ , then  $z_i = l - w$ . In Section 3 we know that on average  $l = w + 2$ . Since the positions of the non-zero digits are independent of  $p$ , on average we lose half the distance between non-zero digits, or  $(w+2)/2$  bits, before the first usable triple and after the last usable triple, which leaves us with  $n - 1 - (p+2) - (w+2)$  bits where our triples can be. The average number of triples is now given by  $(n - p - 3 - (w+2))/(w+2)$  and each of these triples gives us  $l - w = 2$  bits on average. Combining this yields  $2 \cdot (n - p - 3 - (w+2))/(w+2) = 2 \cdot (n - p - 3)/(w+2) - 2$  bits per signature. For the **secp256k1** curve we have that  $n = 256$ ,  $p = 129$  and  $w = 3$ , leading to 47.6 bits per signature on average. Our data obtained from perfect side-channels associated to 1001 signatures gives us an average of 47.6 with a 95% confidence interval of  $\pm 0.2664$ . For the **secp521r1** curve, we have that  $n = 521$ ,  $p = 259$  and  $w = 4$ , which suggests 84.33 bits per signature on average. The data average here is 84.1658 with a 95% confidence interval of  $\pm 0.3825$ . See also the  $Z = 1$  cases of Figures 1 and 2, which show the distribution of the bits leaked per signature in the 256-bit and 521-bit cases, respectively.

This formula suggests that on average, six signatures would be enough to break a 256-bit key (assuming a perfect side channel), since  $47.6 \cdot 6 = 285.6 > 256$ . However, in our preliminary experiments the attack did not succeed once when using six or even seven signatures. Even eight or nine signatures gave a minimal success probability. This indicates that something is wrong with the heuristic. In general there are two possible reasons for failure. Either the lattice problem has the correct solution but it was too hard to solve, or the solution to the lattice problem does not correspond to the private key  $\alpha$ . We will now examine these two possibilities and how to deal with them.

### 4.1 Hardness of the lattice problem

Generally, the lattice problem becomes easier when adding more information to the lattice, but it also becomes harder as the rank increases. Since each triple adds information but also increases the rank of the lattice, it is not always clear whether adding more triples will solve the problem or make it worse. Each triple contributes  $z_i$  bits of information, so we would always prefer triples with a higher  $z_i$  value. Therefore, we set a bound  $Z \geq 1$  and only keep those triples that have  $z_i \geq Z$ . However, this decreases the total number of bits of information we obtain per signature. If  $Z$  is small enough, then roughly speaking we only keep a fraction  $2^{1-Z}$  of the triples, but now each triple contributes  $Z + 1$  bits on average. Hence, the new formula of bits per signature becomes

$$2^{1-Z} \cdot (Z + 1) \cdot ((n - p - 3)/(w + 2) - 1).$$

Our data reflects this formula as well as can be seen in Figures 1 and 2 for the 256-bit and the 521-bit cases, respectively. In our experiments we will set an additional bound  $d$  on the number of triples we use in total, which limits the lattice rank to  $d + 1$ . To this end, we sort the triples by  $z_i$  and then pick the first  $d$  triples to construct the lattice. We adopt this approach for our experiments and the results can be found in Section 5.




Fig. 1: Number of signatures against bits per signature in the 256 bit case.




Fig. 2: Number of signatures against bits per signature in the 521 bit case.

## 4.2 Incorrect solutions

The analysis of Nguyen and Shparlinski [22] requires that the  $t_i$  values in the triples are taken uniformly and independently from a distribution that satisfies some conditions. However, it is easy to see that when two triples are taken from the same signature, the values for the  $t_i = \lfloor r \cdot s^{-1} \cdot 2^{n-m_i-l_i-1} \rfloor_q$  and  $t_j = \lfloor r \cdot s^{-1} \cdot 2^{n-m_j-l_j-1} \rfloor_q$  are not even independent, as they differ mod  $q$  by a factor that is a power of 2 less than  $2^n$ .

Recall from Sections 2.3 and 3 how the triples are used and created, respectively. Consider a triple  $(t_{ij}, u_{ij}, z_{ij})$  corresponding to a signature  $(r_i, s_i, h_i)$ . The corresponding  $v_{ij} = |\alpha \cdot t_{ij} - u_{ij}|_q$  satisfies

$$\begin{aligned} |v_{ij}| &= \left| |\alpha \cdot (r_i \cdot s_i^{-1} \cdot 2^{n-m_j-l_j-1}) - 2^{n+w-l_j-1} \right. \\ &\quad \left. + (h_i \cdot s_i^{-1} + 2^{m_j+w} - 2^{m_j+l}) \cdot 2^{n-m_j-l_j-1} \right|_q \\ &\leq q/2^{z_{ij}+1}, \end{aligned}$$

which is equivalent to

$$|v_{ij}| = \left| |(\alpha \cdot r_i + h_i) \cdot s_i^{-1} \cdot 2^{n-m_j-l_j-1} - 2^{n-1}|_q \right| \leq q/2^{z_{ij}+1},$$

where  $p+1 < m_j < n-l_j$  and  $z_{ij} = l-w$ . Now  $(\alpha \cdot r_i + h_i) \cdot s_i^{-1} = k_i \pmod{q}$  and we know that the previous statement holds due to the structure of  $k_i$ , specifically due to its bits  $m_j+w, \dots, m_j+l_j-1$  repeating, with bit  $m_j+l_j$  being different than the preceding bit. But the map  $x \mapsto (x \cdot r_i + h_i) \cdot s_i^{-1}$  is a bijection mod  $q$ , and hence for each  $i$  there will be many numbers  $X$  such that for all  $j$

$$|v_{ij}(X)| = \left| |(X \cdot r_i + h_i) \cdot s_i^{-1} \cdot 2^{n-m_j-l_j-1} - 2^{n-1}|_q \right| \leq q/2^{z_{ij}+1}.$$

Let  $S_i = \{X : |v_{ij}(X)| \leq q/2^{z_{ij}+1} \text{ for all } j\}$ . If we now have that there exists an  $X \in \bigcap_i S_i$  such that

$$X^2 + \sum_{i,j} (2^{z_{ij}} \cdot v_{ij}(X))^2 < \alpha^2 + \sum_{i,j} (2^{z_{ij}} \cdot v_{ij}(\alpha))^2,$$

then it is very unlikely that the lattice algorithm will find  $\alpha$ , because  $X$  corresponds to a better solution to the lattice problem. Note that this problem arises when fewer signatures are used, because this leads to fewer distinct values for  $(r_i, s_i, h_i)$  and hence fewer sets  $S_i$  that need to intersect. This suggests that increasing the number of signatures could increase the success probability.

Assuming that the  $S_i$  are random, we want to determine what is the probability that their intersection is non-empty. First we consider the size of the  $S_i$ . Recall that  $S_i$  consists of all  $X \pmod{q}$  such that  $v_{ij}(X)$  has ‘the same structure as



$k_i'$ . This means that for each triple specified by  $m_j$  and  $l_j$ , the bits  $m_j + w, \dots, m_j + l_j - 1$  repeat, and bit  $m_j + l_j$  is the opposite of the preceding bits. There are approximately  $2^{n-(l_j-w+1)+1}$  numbers mod  $q$  that have this structure. Let  $f_i$  be the number of triples of signature  $i$  and  $g_{ij} = (l_j - w + 1)$  be the number of bits fixed by triple  $j$  of signature  $i$ . Then, because the triples do not overlap and because  $v_{ij}(\cdot)$  is a bijection, we have that

$$\log_2(|S_i|) = n - \sum_{j=1}^{f_i} (1 - g_{ij}) = n - f_i + \sum_{j=1}^{f_i} g_{ij}.$$

Let  $s_i = |S_i|$  and assume that the  $S_i$  are chosen randomly and independently from all the subsets of integers in the range  $[0, \dots, N - 1]$  (of size  $s_i$ ), where  $N = 2^n$ . Consider the following probability

$$p_i = \mathbb{P}(0 \in S_i) = s_i/N,$$

since  $S_i$  is randomly chosen. Now, because the  $S_i$  are also chosen independently, we have

$$\mathbb{P}\left(0 \in \bigcap_i S_i\right) = \prod_i p_i.$$

Finally, since this argument holds for any  $j \in [0, \dots, N - 1]$ , we can apply the union bound to obtain

$$p_{\text{fail}} = \mathbb{P}\left(\bigcup_j \left(j \in \bigcap_i S_i\right)\right) \leq \sum_j \mathbb{P}\left(0 \in \bigcap_i S_i\right) = N \cdot \prod_i p_i. \quad (8)$$

Recall that each signature has  $f_i = 2^{1-Z} \cdot ((n - p - 3)/(w + 2) - 1)$  triples on average and each triple contributes  $Z + 1$  bits on average, which means  $g_{ij} = Z + 2$  on average. If we plug in the numbers  $n = 256$ ,  $p = 129$ ,  $w = 3$  and  $Z = 3$ , we get that  $f_i \approx 6$ ,  $g_{ij} = 5$  and hence  $p_i \approx 2^{-6 \cdot (5-1)} \approx 2^{-24}$  if we assume an average number of triples and bits in each signature. This in turn gives us an upper bound of  $p_{\text{fail}} \leq N/2^{24 \cdot k}$ . If  $k \geq 11$ , this upper bound is less than one, so this clearly suggests that from about eleven signatures and up, we should succeed with some probability, which is indeed the case from our experiments.

Repeating this for  $n = 521$ ,  $p = 259$ ,  $w = 4$  and  $Z = 4$ , we obtain  $f_i \approx 5$ ,  $g_{ij} = 6$  and hence  $p_i \approx 2^{-5 \cdot (6-1)} \approx 2^{-25}$ . Consequently,  $p_{\text{fail}} \leq N/2^{25 \cdot k}$ , which is less than one when  $k \geq 21$ . However, in our experiments we require at least 30 signatures to obtain the secret key with some probability. Thus the above analysis is only approximate as the secret key length increases.

## 5 Results With a Perfect Side-Channel

Subsection 2.3 outlined our (heuristic) approach to obtain the secret key from a number of triples  $(t_i, u_i, z_i)$  using lattices and Section 3 outlined how to generate these triples from the side-channel information. In this section we will look at some experimental results to see if our heuristic assumptions are justified.

As per Section 4, we used the following approach for our experiments. First, we fix a number of signatures  $s$ , a lattice rank  $d$  and a bound  $Z$ . We then take  $s$  signatures at random from our data set and derive all triples such that  $z_i \geq Z$ , sorting them such that the  $z_i$  are in descending order. If we have more than  $d$  triples, we only take the first  $d$  to construct the lattice. Finally we attempt to solve the lattice problem and note the result. All executions were performed in single thread on an Intel Core i7-3770S CPU running at 3.10 GHz.

When solving the CVP instances there are three possible outcomes. We obtain either no solution, the private key or a wrong solution. No solution means that the lattice problem was too hard for the algorithm and constraints we used, but spending more time and using stronger algorithms might still solve it. When a ‘wrong’ solution is obtained, this means that our heuristics failed: the solution vector was not unique, in the sense that there were other lattice vectors within the expected distance from our target vector.

When solving the SVP instance there are only two outcomes. Either we obtain the private key or not. However, in this case it is not as clear whether a wrong solution means that there were other solutions due to the additional heuristics involved. The complete details of our experimental data are given in the Appendix.

### 5.1 256 bit key

For the 256 bit case, we used BKZ with block size 20 from fplll [6] to solve the SVP instances, as well as to pre-process the CVP instances. To solve the CVP, we applied Schnorr-Euchner enumeration [23] using linear pruning [12] and limiting the number of enumerated nodes to  $2^{29}$ .

The CVP approach seems the best, as the lattice rank ( $d + 1$ ) remains quite small. We restrict our triples to  $Z = 3$  to keep the rank small, but a smaller  $Z$  would not improve our results much. See the appendix for details. We observed that failures are mostly caused by ‘wrong’ solutions in this case, rather than the lattice problem being too hard. In all cases we found that using 75 triples gave the best results. Table 2 in the Appendix lists the runtimes and success probabilities of the lattice part of the attack for varying  $s$ . The results are graphically presented in Figures 4 and 5 in the Appendix.

### 5.2 521 bit key

For the 521 bit case, we used BKZ with block size 20 from fplll [6] to solve the SVP instances. Due to the higher lattice ranks in this case, solving the CVP instances proved much less efficient, even when restricting the triples to  $Z = 4$ .

With 30 signatures we get a small probability of success in the lattice attack whereas with 40 signatures we can obtain the secret key in more than half of the cases. It should be noted that as the number of signatures increases, the choice of  $d$  becomes less important, because the number of triples with more information increases. See the Appendix for Table 4 details and Figures 6 and 7 for a graphical representation.

## 6 Results in a Real-Life Attack

So far our discussion was based on the assumption of a perfect side-channel. That is, we assumed that the double-and-add chains are recovered without any errors. Perfect side-channels are, however, very rare. In this section we extend the results to the actual side-channel exposed by the FLUSH+RELOAD technique.

The attack was carried on an HP Elite 8300, running CentOS 6.5. The victim process runs OpenSSL 1.0.1f, compiled to include debugging symbols. These symbols are not used at run-time and do not affect the performance of OpenSSL. We use them because they assist us in finding the addresses to probe by avoiding reverse engineering [9]. The spy uses a time slot of 1,200 cycles ( $0.375\mu s$ ). In each time slot it probes the memory lines containing the last field multiplication within the group add and double functions. (`ec_GFp_simple_add` and `ec_GFp_simpledbl`, respectively.) Memory lines that contain function calls are accessed both before and after the call, reducing the chance of a spy missing the access due to overlap with the probe. Monitoring code close to the end of the function eliminates false positives due to speculative execution. See Yarom and Falkner [27] for a discussion of overlaps and speculative execution.



Fig. 3: FLUSH+RELOAD spy output. Vertical bars indicate time-slot boundaries; ‘A’ and ‘D’ are probes for OpenSSL access to add and double; dashes indicate missed time-slots.

Figure 3 shows an example of the output of the spy when OpenSSL signs using `secp256k1`. The double and three addition operations at the beginning of the captured sequence are the calculation of the pre-computed wNAF digits. Note the repeated capture of the double and add operations due to monitoring a memory line that contains a function call. The actual wNAF multiplication starts closer to the end of the line, with 7 double operations followed by a group addition.

In this example, the attack captures most of the double and add chain. It does, however, miss a few time-slots and consequently a few group operations in the chain. The spy recognises missed time-slots by noting inexplicable gaps in the processor cycle counter. As we do not know which operations are missed, we lose the bit positions of the operations that precede the missed time-slots. We believe that the missed time-slots are due to system activity which suspends the spy.

Occasionally OpenSSL suspends the calculation of the scalar multiplication to perform memory management functions. These suspends confuse our spy program, which assumes that the scalar multiplication terminated. This, in turn, results in a short capture, which cannot be used for the lattice attack.

To test prevalence of capture errors we captured 1,000 scalar multiplications and compared the capture results to the ground truth. 342 of these captures contained missed time-slots. Another 77 captures contains less than 250 group operations and are, therefore, too short. Of the remaining 581 captures, 577 are perfect while only four contain errors that we could not easily filter out.

Recall, from Section 5, that 13 perfectly captured signatures are sufficient for breaking the key of a 256 bits curve with over 50% probability. An attacker using FLUSH+RELOAD to capture 25 signatures can thus expect to be able to filter out 11 that contain obvious errors, leaving 14 that contain no obvious errors. With less than 1% probability that each of these 14 captures contains an error, the probability that more than one of these captures contains an error is also less than 1%. Hence, the attacker only needs to test all the combination of choosing 13 captures out of these 14 to achieve a 50% probability of breaking the signing key.

Several optimisations can be used to improve the figure of 25 signatures. Some missed slots can be recovered and the spy can be improved to correct short captures. Nevertheless, it should be noted that this figure is still an order of magnitude than the previously best known result of 200 signatures [2], where 200 signatures correspond to a 3.5% probability of breaking the signing key, whereas 300 signatures were required to get a success probability greater than 50%.

## Acknowledgements

The authors would like to thank Ben Sach for helpful conversations during the course of this work. The first and second authors work has been supported in part by ERC Advanced Grant ERC-2010-AdG-267188-CRIPTO, by EPSRC via grant EP/I03126X, and by Defense Advanced Research Projects Agency (DARPA) and the Air Force Research Laboratory (AFRL) under agreement number FA8750-11-2-0079<sup>3</sup>.

The third author wishes to thank Dr Katrina Falkner for her advice and support and the Defence Science and Technology Organisation (DSTO) Maritime Division, Australia, who partially funded his work.

## References

1. American National Standards Institute. *ANSI X9.62, Public Key Cryptography for the Financial Services Industry: The Elliptic Curve Digital Signature Algorithm*, 1999.
2. Naomi Benger, Joop van de Pol, Nigel P. Smart, and Yuval Yarom. “Ooh aah... just a little bit”: A small amount of side channel can go a long way. In Lejla Batina and Matthew Robshaw, editors, *Proceedings of the 16th International Workshop on Cryptographic Hardware and Embedded Systems*, volume 8731 of *Lecture Notes in Computer Science*, pages 75–92, Busan, Korea, September 2014. Springer.
3. Dan Boneh and Ramarathnam Venkatesan. Hardness of computing the most significant bits of secret keys in diffie-hellman and related schemes. In *CRYPTO*, volume 1109 of *Lecture Notes in Computer Science*, pages 129–142, 1996.
4. Billy Bob Brumley and Risto M. Hakala. Cache-timing template attacks. In Mitsuru Matsui, editor, *Advances in Cryptology - ASIACRYPT 2009*, volume 5912 of *Lecture Notes in Computer Science*, pages 667–684. Springer-Verlag, 2009.
5. Billy Bob Brumley and Nicola Taveri. Remote timing attacks are still practical. In Vijay Atluri and Claudia Diaz, editors, *Computer Security - ESORICS 2011*, pages 355–371, Leuven, Belgium, September 2011.

<sup>3</sup> The US Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of Defense Advanced Research Projects Agency (DARPA) or the U.S. Government.

6. David Cadé, Xavier Pujol, and Damien Stehlé. FPLLL-4.0.4. <http://perso.ens-lyon.fr/damien.stehle/fplll/>, 2013.
7. Certicom Research. *SEC 2: Recommended Elliptic Curve Domain Parameters, Version 2.0*, January 2010.
8. Mathieu Ciet and Marc Joye. (virtually) free randomization techniques for elliptic curve cryptography. In Sihan Qing, Dieter Gollmann, and Jianying Zhou, editors, *Proceedings of the fifth International Conference on Information and Communications Security*, volume 2836 of *Lecture Notes in Computer Science*, pages 348–359, Huhehaote, China, October 2003. Springer.
9. Teodoro Ciproso and Mark Stamp. Software reverse engineering. In Peter Stavroulakis and Mark Stamp, editors, *Handbook of Information and Communication Security*, chapter 31, pages 659–696. Springer, 2010.
10. Richard E. Crandall. Method and apparatus for public key exchange in a cryptographic system. US Patent 5,159,632, October 1992.
11. Benoit Feix, Mylène Roussellet, and Alexandre Venelli. Side-channel analysis on blinded regular scalar multiplications. In Willi Meier and Debdeep Mukhopadhyay, editors, *Proceedings of the 15th International Conference on Cryptology in India (INDOCRYPT 2014)*, Lecture Notes in Computer Science, New Delhi, India, December 2014. Springer.
12. Nicolas Gama, Phong Q. Nguyen, and Oded Regev. Lattice enumeration using extreme pruning. In *EUROCRYPT*, volume 6110 of *Lecture Notes in Computer Science*, pages 257–278, 2010.
13. Daniel M. Gordon. A survey of fast exponentiation methods. *Journal of Algorithms*, 27(1):129–146, April 1998.
14. Nick Howgrave-Graham and Nigel P. Smart. Lattice attacks on digital signature schemes. *Designs, Codes and Cryptography*, 23(3):283–290, 2001.
15. Avi Kivity, Yaniv Kamay, Dor Laor, Uri Lublin, and Anthony Liguori. kvm: the Linux virtual machine monitor. In *Proceedings of the Linux Symposium*, volume one, pages 225–230, Ottawa, Ontario, Canada, June 2007.
16. Arjen K. Lenstra, Hendrik W. Lenstra, and László Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261(4):515–534, 1982.
17. Lixin Li, James E. Just, and R. Sekar. Address-space randomization for Windows systems. In *Proceedings of the 22nd Annual Computer Security Applications Conference*, pages 329–338, Miami Beach, Florida, United States, December 2006.
18. Bodo Möller. Parallelizable elliptic curve point multiplication method with resistance against side-channel attacks. In Agnes Hui Chan and Virgil D. Gligor, editors, *Proceedings of the fifth International Conference on Information Security*, number 2433 in *Lecture Notes in Computer Science*, pages 402–413, São Paulo, Brazil, September 2002.
19. Bodo Möller. Improved techniques for fast exponentiation. In P. J. Lee and C. H. Lim, editors, *Information Security and Cryptology - ICISC 2002*, number 2587 in *Lecture Notes in Computer Science*, pages 298–312. Springer-Verlag, 2003.
20. National Institute of Standards and Technology. *FIPS PUB 186-4 Digital Signature Standard (DSS)*, 2013.
21. Phong Q. Nguyen and Igor E. Shparlinski. The insecurity of the digital signature algorithm with partially known nonces. *Journal of Cryptology*, 15(3):151–176, June 2002.
22. Phong Q. Nguyen and Igor E. Shparlinski. The insecurity of the elliptic curve digital signature algorithm with partially known nonces. *Designs, Codes and Cryptography*, 30(2):201–217, September 2003.
23. Claus-Peter Schnorr and M. Euchner. Lattice basis reduction: Improved practical algorithms and solving subset sum problems. In *Fundamentals of Computation Theory – FCT 1991*, volume 529 of *Lecture Notes in Computer Science*, pages 68–85. Springer, 1991.
24. Jerome A. Solinas. Generalized Mersenne numbers. Technical Report CORR-39, University of Waterloo, 1999.
25. Carl A. Waldspurger. Memory resource management in VMware ESX Server. In David E. Culler and Peter Druschel, editors, *Proceedings of the Fifth Symposium on Operating Systems Design and Implementation*, pages 181–194, Boston, Massachusetts, United States, December 2002.
26. Yuval Yarom and Naomi Benger. Recovering OpenSSL ECDSA nonces using the FLUSH+RELOAD cache side-channel attack. Cryptology ePrint Archive, Report 2014/140, February 2014. <http://eprint.iacr.org/>.
27. Yuval Yarom and Katrina Falkner. FLUSH+RELOAD: a high resolution, low noise, L3 cache side-channel attack. In *Proceedings of the 23rd USENIX Security Symposium*, pages 719–732, San Diego, California, United States, August 2014.

## A Experimental results

### A.1 256 Bit Keys

$s$	$d$	SVP		CVP	
		Time (s)	$p_{\text{succ}}$ (%)	Time (s)	$p_{\text{succ}}$ (%)
10	60	1.47	0.0	1.56	0.5
10	65	1.42	1.0	1.90	2.5
10	70	1.44	1.5	2.45	4.0
10	75	1.50	1.5	2.25	7.0
11	60	1.28	0.0	1.63	0.5
11	65	1.68	5.0	2.35	6.5
11	70	1.86	2.5	3.15	19.0
11	75	2.05	7.5	4.66	25.0
11	80	2.12	6.0		
12	60	1.27	2.0	1.69	7.0
12	65	1.71	2.5	2.45	10.5
12	70	2.20	7.5	3.99	29.5
12	75	2.57	10.5	7.68	38.5
12	80	2.90	13.0		
12	85	3.12	8.5		
12	90	3.21	15.5		
13	60	1.30	3.5	1.92	8.5
13	65	1.77	6.0	2.79	25.5
13	70	2.39	11.0	4.48	46.5
13	75	3.16	19.0	11.30	54.0
13	80	3.67	18.5		
13	85	3.81	21.5		
13	90	4.37	25.0		

Table 1: Results for  $d$  triples taken from  $s$  signatures with a 256-bit key ( $Z = 3$ )

$s$	Time (s)	$p_{\text{succ}}$ (%)
10	2.25	7.0
11	4.66	25.0
12	7.68	38.5
13	11.30	54.0

Table 2: CVP results for 75 triples taken from  $s$  signatures with a 256-bit key ( $Z = 3$ )

### A.2 521 Bit Keys

$s$	$d$	Time (s)	$p_{\text{succ}}$ (%)
30	130	50.10	4.0
30	135	58.80	3.0
30	140	66.65	3.5
30	145	69.68	2.5
32	130	50.15	6.5
32	135	58.07	6.5
32	140	62.55	4.0
32	145	67.46	5.0
32	150	70.77	9.5
34	130	50.00	15.5
34	135	55.93	10.5
34	140	62.83	16.0
34	145	64.41	14.0
34	150	70.50	16.0
34	155	71.07	11.5
36	130	48.71	24.5
36	135	54.74	21.0
36	140	59.25	22.5
36	145	62.32	29.0
36	150	65.60	29.0
36	155	68.57	24.5
38	130	49.04	38.5
38	135	53.86	36.0
38	140	57.14	38.5
38	145	61.31	42.5
38	150	66.75	36.5
38	155	66.52	36.5
40	130	47.73	53.0
40	135	50.80	49.0
40	140	54.88	52.0
40	145	60.47	47.0
40	150	64.77	53.0
40	155	64.95	52.5

$s$	$d$	Time (s)	$p_{\text{succ}}$ (%)
31	130	48.50	7.5
31	135	59.91	3.5
31	140	67.35	6.0
31	145	69.96	5.5
33	130	49.70	8.0
33	135	56.52	11.5
33	140	60.31	11.5
33	145	66.39	8.5
33	150	70.54	13.5
33	155	75.49	8.5
35	130	49.76	12.0
35	135	55.33	24.5
35	140	59.50	15.5
35	145	65.59	19.5
35	150	66.93	24.0
35	155	69.67	20.0
37	130	48.20	24.0
37	135	54.79	23.5
37	140	58.60	28.0
37	145	60.05	29.0
37	150	63.40	27.5
37	155	69.14	34.5
39	135	50.99	45.5
39	140	58.81	46.0
39	145	57.08	47.5
39	150	62.35	41.5
39	155	64.99	42.5

Table 3: SVP results for  $d$  triples taken from  $s$  signatures with a 521-bit key ( $Z = 4$ )

$s$	$d$	Time (s)	$p_{\text{succ}}$ (%)
30	130	50.10	4.0
31	130	48.50	7.5
32	150	70.77	9.5
33	150	70.54	13.5
34	140	62.83	16.0
35	135	55.33	24.5
36	145	62.32	29.0
37	155	69.14	34.5
38	145	61.31	42.5
39	145	57.08	47.5
40	130	47.73	53.0

Table 4: SVP results for  $d$  triples taken from  $s$  signatures with a 521-bit key ( $Z = 4$ )

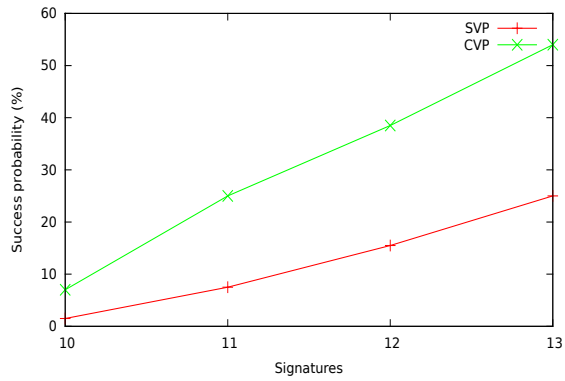


Fig. 4: Success probability per number of signatures against a 256 bit key

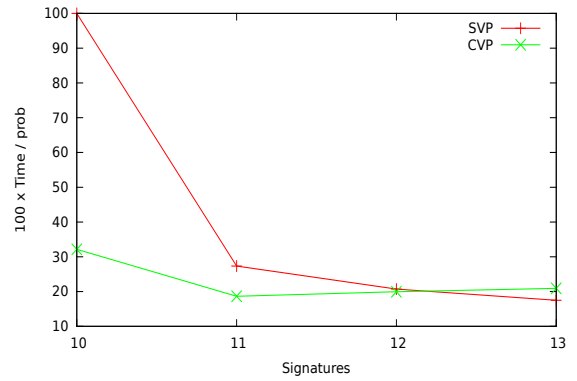


Fig. 5: Expected running time per number of signatures against a 256 bit key




Fig. 6: Success probability per number of signatures against a 521 bit key




Fig. 7: Expected running time per number of signatures against a 521 bit key