

Masking ring-LWE

Oscar Reparaz¹ · Sujoy Sinha Roy¹ · Ruan de Clercq¹ · Frederik Vercauteren¹ · Ingrid Verbauwhede¹

Received: 30 November 2015 / Accepted: 26 February 2016 / Published online: 16 March 2016
© Springer-Verlag Berlin Heidelberg 2016

Abstract In this paper, we propose a masking scheme to protect ring-LWE decryption from first-order side-channel attacks. In an unprotected ring-LWE decryption, the recovered plaintext is computed by first performing polynomial arithmetic on the secret key and then decoding the result. We mask the polynomial operations by arithmetically splitting the secret key polynomial into two random shares; the final decoding operation is performed using a new bespoke masked decoder. The outputs of our masked ring-LWE decryption are Boolean shares suitable for derivation of a symmetric key. Thus, the masking scheme keeps all intermediates, including the recovered plaintext, in the masked domain. We have implemented the masking scheme on both hardware and software. On a Xilinx Virtex-II FPGA, the masked ring-LWE processor requires around 2000 LUTs, a 20 % increase in the area with respect to the unprotected architecture. A masked decryption operation takes 7478 cycles, which is only a factor $2.6\times$ larger than the unprotected decryption. On a 32-bit ARM Cortex-M4F processor, the masked software implementation costs around $5.2\times$ more cycles than the unprotected implementation.

Keywords Lattice cryptography · Ring-LWE · DPA · Masking

1 Introduction

Once the quantum computer is built, Shor's algorithm will make most current public-key cryptographic algorithms obsolete. In particular, public-key cryptosystems that rely on number-theoretic hardness assumptions such as integer factorization (RSA) or discrete logarithms, either in \mathbb{Z}_p^* (Diffie–Hellman) or in elliptic curves over finite fields, will be insecure. On the bright side, there is an entire branch of post-quantum cryptography that is believed to resist mathematical attacks running on quantum computers.

There are three main branches of post-quantum cryptosystems: based on codes, on multivariate quadratic equations or on lattices [2]. Lattice-based cryptographic constructions, founded on the learning with errors (LWE) problem [23] and its ring variant known as ring-LWE problem [17], have become a versatile tool for designing asymmetric encryption schemes [17], digital signatures [10] and homomorphic encryption schemes [4, 11]. Several hardware and software implementations of such schemes have appeared in the literature. So far, the reported implementations have focused mainly on efficient implementation strategies, and very little research work has appeared in the area of side channel security of the lattice-based schemes.

It comes as no surprise that implementations of post-quantum algorithms are vulnerable to side-channel attacks. Side-channel attacks, as introduced by Kocher [15], exploit timing, power consumption or the electromagnetic emanation from a device executing a cryptographic implementation to extract secrets, such as cryptographic keys. A particularly powerful side-channel technique is differential power analysis (DPA), introduced by Kocher et al. [16]. In a typical DPA attack, the adversary measures the instantaneous power consumption of a device, places hypotheses on subkeys and applies statistical tests to confirm or reject the hypotheses.

This journal version is based on a paper that appeared at the CHES 2015 conference [26]. Sections 6, 7.3 and 8.2 carry substantial differences.

✉ Oscar Reparaz
oscar.reparaz@esat.kuleuven.be

¹ Kasteelpark Arenberg 10 bus 2452, Leuven, Belgium

DPA attacks can be surprisingly easy to mount even with low-end equipment, and hence it is important to protect against them.

There are plenty of countermeasures against DPA. Most notably, masking [7, 14] is provably sound and popular in industry. Masking effectively randomizes the computation of the cryptographic algorithm by splitting each intermediate into several shares, in such a way that each share is independent of any secret. This property is preserved through the entire computation. Thus, observing any single intermediate (for example, by a side channel, be it known or unknown) reveals nothing about the secret. However, there are not many masking schemes specifically designed for post-quantum cryptography. In [5], Brenner et al. present a masked FPGA implementation of the post-quantum pseudo-random function SPRING.

In the rest of the paper, we focus on protecting the ring-LWE decryption operation against side-channel attacks with masking. The decryption algorithm is considerably exposed to DPA attacks since it repeatedly uses long-term private keys. In contrast, the encryption or key-generation procedures use ephemeral secrets only [27].

Our contribution In this paper, we present a compact masked implementation of the ring-LWE decryption function. The masking countermeasure adds a small overhead when compared with the other previous approaches, thanks to a bespoke probabilistic masked decoder designed specifically for our implementation. We implemented the masked ring-LWE decryption on a Virtex-II FPGA and on an ARM Cortex-M4F processor, and tested the side-channel security with practical experiments.

Organization The paper is structured as follows: we provide a brief mathematical background of the ring-LWE encryption scheme in Sect. 2 and describe a high-level overview of the proposed masked ring-LWE decryption in Sect. 3. Next, in Sect. 4 we construct the masked decoder. In Sect. 5, we describe our hardware implementation, and in the next Sect. 6 we describe our software implementation. We analyze the error rates of the decryption operation in Sect. 7. We dedicate Sect. 8 for the side-channel evaluation of both hardware and software implementations and draw conclusions in the last section.

2 Preliminaries

Notation The Latin letters r, c_i indicate polynomials. When we want to explicitly access a coefficient of the polynomial, we write $r[i]$. Multiplication of polynomials is written as $r * c_1$. Coefficient-wise multiplication is denoted as $r \cdot c_1$. The letter m denotes a string of message bits, and q is an integer. Letters with prime x' or double prime x'' represent shares

of variable x . Depending on the context, these shares are split either arithmetically $x = x' + x'' \pmod{q}$ or Boolean $x = x' + x'' \pmod{2}$. A polynomial r is shared into (r', r'') by additively sharing each of its coefficients $r[i]$ such that $r = r' + r''$.

Ring-LWE For completeness, we give in this section a description of the three major algorithms of the ring-LWE public-key cryptosystem [17]: key generation, encryption and decryption.

The ring-LWE encryption scheme works with polynomials in a ring $R_q = \mathbb{Z}_q[x]/(f(x))$, where $f(x)$ is an irreducible polynomial of degree n . During the key generation, encryption and decryption operations, polynomial arithmetic such as polynomial addition, subtraction and multiplication are performed. In addition, the key-generation and encryption operations require sampling of error polynomials from an error distribution (typically a discrete Gaussian.)

The ring-LWE encryption scheme is described in this way:

- In the key generation phase, two error polynomials r_1 and r_2 are sampled from the discrete Gaussian distribution. The secret key is the polynomial r_2 and the public key is the polynomial $p = r_1 - g * r_2$. After key generation, there is no use of the polynomial r_1 . The polynomial g is globally known.
- In the encryption operation of a binary message vector m of length n , the message is first lifted to a ring element $\bar{m} \in R_q$ by multiplying the message bits by $q/2$. The ciphertext is computed as a pair of polynomials (c_1, c_2) where $c_1 = g * e_1 + e_2$ and $c_2 = p * e_1 + e_3 + \bar{m} \in R_q$. The encryption operation requires generation of three error polynomials e_1, e_2 and e_3 .
- The decryption operation uses the private key r_2 to compute the message as $m = \text{th}(c_1 * r_2 + c_2)$. The decoding function th is a simple threshold decoder that is applied coefficient-wise and is defined as

$$\text{th}(x) = \begin{cases} 0 & \text{if } x \in (0, q/4) \cup (3q/4, q) \\ 1 & \text{if } x \in (q/4, 3q/4) \end{cases}. \quad (1)$$

Efficiency improvements To achieve an efficient implementation of the encryption scheme, the irreducible polynomial $f(x)$ is taken as $x^n + 1$ where n is a power of two, and the modulus q is chosen as a prime number satisfying $q \equiv 1 \pmod{2n}$ [20, 28]. In this setting, polynomial multiplications can be efficiently performed in $O(n \log n)$ time using the number theoretic transform (NTT).

Following [28], we keep the ciphertext polynomials c_1 and c_2 in the NTT domain to reduce the computation cost of the decryption operation. The decryption operation thus computes the decrypted message as

$$m = \text{th}(\text{INTT}(\tilde{c}_1 \cdot \tilde{r}_2 + \tilde{c}_2)). \quad (2)$$

Here, the symbol \tilde{r} represents the NTT of a polynomial r , and $\text{INTT}(\cdot)$ represents the inverse NTT operation. The multiplication of $\tilde{c}_1 \cdot \tilde{r}_2$ is thus performed coefficient-wise (as well as the addition of $\tilde{c}_1 \cdot \tilde{r}_2 + \tilde{c}_2$.) For convenience, we drop the tildes in the rest of the paper and work with c_1, c_2 and r_2 in the NTT domain. We furthermore refer to \tilde{r}_2 simply as r . (We recall that the INTT is a linear transformation applied to the n coefficients of $a = r \cdot c_1 + c_2$.) The decoding function th applies a threshold function to each coefficient of a as defined in Eq. (1) to output n recovered message bits.

3 High-level overview

In this section, we give a high-level view of the masked ring-LWE implementation. The most natural way to split the computation of the decryption as Eq. (2) is to split the secret polynomial r additively into two shares r' and r'' , such that $r[i] = r'[i] + r''[i] \pmod{q}$ for all i . The n coefficients of r' are chosen uniformly at random in \mathbb{Z}_q in each execution of the decryption.

The bulk of the computation from Eq. (2) is amenable to this splitting, since by linearity of the multiplication and INTT operation, we have that $\text{INTT}(r \cdot c_2 + c_1) = \text{INTT}(r' \cdot c_2 + c_1) + \text{INTT}(r'' \cdot c_2)$. Thus, we can split almost the entire computation from Eq. (2) into two branches, as drawn in Fig. 1. The first branch computes on r' to determine the polynomial

$$a' = \text{INTT}(r' \cdot c_1 + c_2) \tag{3}$$

and the second branch operates on r'' to determine

$$a'' = \text{INTT}(r'' \cdot c_1). \tag{4}$$

The advantage of such a high-level masking is that the operations of Eqs. (3) and (4) can be performed on an arithmetic processor without any particular protection against DPA. (This is because any intermediate appearing in either branch is independent of the secret r . This situation is very similar to, for example, base point blinding in elliptic curve scalar multiplication.) We can reuse an existing ring-LWE processor for these operations and leverage the numerous optimizations carried out for this block [9,20,28].

The final threshold $\text{th}(\cdot)$ operation of Eq. (2) is obviously non-linear in the base field \mathbb{F}_q and hence cannot be independently applied to each branch (Eqs. (3) and (4)). There are generic approaches to mask arbitrary functions. For instance, in [5], an approach based on masked tables was used. However, these generic approaches are usually quite expensive in terms of area or randomness. In the following Sect. 4, we pursue another direction. We design a bespoke masked decoder that results in a compact implementation.




Fig. 1 General data flow of the masked ring-LWE decryption. r' and r'' are the arithmetic shares of the private key r ; c_1 and c_2 are the input unmasked ciphertext; m' and m'' are the Boolean shares of the recovered plaintext

4 Masked decoder

In this section, we describe a compact, probabilistic masked decoder. In the sequel, a denotes a single coefficient and (a', a'') its shares such that $a' + a'' = a \pmod{q}$. The decoder computes the function $\text{th}(a)$ from the shares (a', a'') . We also drop the symbol \pmod{q} when obvious.

First crack The key idea of the efficient masked decoder is that we do not need to know the exact values of the shares a' and a'' of a coefficient a to compute $\text{th}(a)$. For example, if $0 < a' < q/4$ and $q/4 < a'' < q/2$, then $a = a' + a''$ is bounded by $q/4 < a < 3q/4$ and thus $\text{th}(a) = 1$, that is, we learnt $\text{th}(a)$ from only a few most significant bits from a' and a'' . We can use this idea to substantially simplify the complexity of the masked th function.

4.1 Rules

Figure 2, left, illustrates the situation from the last paragraph. In this case, $0 < a' < q/4$ and $q/4 < a'' < q/2$, so obviously a can range only from $q/4$ to $3q/4$, and hence $\text{th}(a) = 1$. Analogously to this rule, we can formulate 3 other rules:

- If $q/2 < a' < 3q/4$ and $3q/4 < a'' < q$ then $q/4 < a < 3q/4$ and thus $\text{th}(a) = 1$.
- If $q/4 < a' < q/2$ and $q/2 < a'' < 3q/4$ then a belongs to $(0, q/4) \cup (3q/4, q)$ and thus $\text{th}(a) = 0$ (quadrants I and IV, left half of the circle).
- If $3q/4 < a' < q$ and $0 < a'' < q/4$ then a belongs to $(0, q/4) \cup (3q/4, q)$ and thus $\text{th}(a) = 0$.

There are 4 other rules that result from interchanging a' with a'' in the above expressions. (This follows straight from the symmetry of the additive splitting.) Essentially, with the only information of the *quadrant* of each share a' and a'' , we can, in half of the cases, deduce the output of $\text{th}(a)$. (For the simplicity of the explanation, we obviated what happens in the boundaries of the quadrant intervals. Similar conclusions hold when including them.)




Fig. 2 Idea for the masked decoder. Elements in \mathbb{Z}_q are shown in a circle. Adding two elements translates into adding their respective angles. Left case $0 < a' < q/4, q/4 < a'' < q/2$, and therefore $\text{th}(a) = 1$.

Center and right case $0 < a' < q/4, 0 < a'' < q/4$, which does not allow to infer $\text{th}(a)$

What if no rule is hit? In roughly half of the cases, we can apply one of the eight rules previously described to deduce the value of $\text{th}(a)$. However, in the other half of the cases, none of the rules applies. A representative case of this event is shown in Fig. 2, center and right. In both cases, $0 < a' < q/4$ and $0 < a'' < q/4$. This situation is not covered by any of the eight rules previously described. We see that in the center sub figure $\text{th}(a) = 0$, while in the right sub figure $\text{th}(a) = 1$, so in this case the quadrants of each share a' and a'' do not allow us to infer $\text{th}(a)$.

The solution in this case is to refresh the splitting (a', a'') , that is, update $a' \leftarrow a' + \Delta_1$ and $a'' \leftarrow a'' - \Delta_1$ for certain Δ_1 . (This refreshing¹ naturally preserves the unshared value $a = a' + a''$.) After the refreshing, the eight rules can be checked again. If still no rule applies, the process is repeated with a different refreshing value Δ_i . Note that in each iteration of the step, roughly half of the possible values of $(a', a'') \in \mathbb{Z}_q \times \mathbb{Z}_q$ are successfully decoded, and thus the amount of pairs (a', a'') that do not get decoded shrinks exponentially with the number of iterations. In our implementation, $N = 16$ iterations produces a satisfactory result. This will be studied in detail in Sect. 7.1.

Optimal cooked values for Δ_i One can determine a sequence of Δ_i values that maximizes the number of pairs successfully decoded after N iterations. We performed a first-order search for such a sequence of Δ_i values. Each Δ_i maximizes the number of successfully decoded pairs after $i - 1$ iterations. For $q = 7681$, the sequence of Δ_i shown in Appendix A was found.

Architecture The hardware architecture for the masked decoder follows from the previous working principle description. Our implementation is shown in Fig. 3. From left to




Fig. 3 Data flow for the masked decoder

right, we see the first refreshing step by the constants Δ_i . The constants Δ_i vary from iteration to iteration. After the refreshing step, the quadrant function is applied to each share a', a'' . This quadrant function outputs x if a belongs to the x -th quadrant, and thus the output consists of 2 bits. These blocks are essentially 13-bit comparators, and thus relatively inexpensive in logic.² The subsequent rule checking on (q', q'') is performed by a masked table lookup that is described in the following section. The whole process is repeated $N = 16$ iterations, and this number of iterations stays fixed even if the decoding is successful after the few first iterations.

4.2 Masked table lookup

The final step in the masked decoder is a masked table lookup. This table implements the rules described in Sect. 4.1, and essentially maps the output of each quadrant q'_i and q''_i (2 bits each) after the i -th iteration ($i \in [1, N]$) to a (Boolean) masked output bit value (m'_i, m''_i) . In our specific implementation, we have other inputs: the result of the decoding from the previous iteration (m'_{i-1}, m''_{i-1}) and an extra randomness bit r (fresh at each of the N iterations for each of the n coefficients).

¹ We use here the term “refresh” to refer to the process of modifying the masked representation (a', a'') of a without modifying the unshared value a , but, contrary to other contexts in the literature, we do not imply that we are pumping new randomness in the new representation.

² Note that in the special case, q is a prime close to a power of two and the construction of the quadrant block can be further simplified.

This is a well-studied problem that arises in other situations (for instance, when masking the sbox lookup in a typical block cipher) and there are plenty of approaches here to implement such masked table lookup.

Hardware In hardware, we opted for the approach of masked tables as in [30]. We set $m'_i \leftarrow r$ and we compute $m''_i \leftarrow f(r, q'_i, q''_i, m'_{i-1}, m''_{i-1})$. The function f essentially bypasses the previous decoded value when no rule applies to q'_i, q''_i by setting the output m''_i to $r + m'_{i-1} + m''_{i-1}$ (refreshing the content of the output registers). If a rule applies to q'_i, q''_i , it sets the output m''_i accordingly. By doing this, we can register always the output of this table and no control logic to enable such output register is needed (it is implicitly integrated into this masked table.) This is the reason why the table sees also the previous decoded value m'_{i-1} and m''_{i-1} .

The usual precautions are applied when implementing f . For our target FPGA platform, we carefully split the 7-bit input to 1-bit output function f into a balanced tree of 4-bit input LUTs, in such a way that any intermediate input or output of LUTs does not leak in the first order. Note that here we are assuming that each LUT is an atomic operation. If stronger security guarantees are needed, other approaches to implement such function f should be followed. When implemented in an ASIC, it may be preferable to store this masked table in ROM (since the contents of the table are immutable and the size is small).

The output of this table is (Boolean) masked, and thus no unmasked value lives within the implementation. This is suited for consumption of a masked AES module (say) after some preprocessing as will be detailed later. We stress that we use masked tables on the *output* of the quadrants. This is the key for our reduced area requirements, as will be explained in Sect. 5.

Software For the software implementation of the masked table lookup, we base our approach on the previous hardware description. We first write an unmasked decoder in a (software) bitsliced way, and then apply the method of [1] to provide “gate-level” masking to the bitsliced software implementation. More details are given in Sect. 6.

5 Hardware implementation

We implemented the fully masked ring-LWE decryption system with the proof-of-concept parameter set $(n, q, s) = (256, 7681, 11.32)$ first introduced in [13], corresponding to a medium-term security level. Note that these concrete choice of parameters is not meant to be deployed. The target platform is a Xilinx Virtex-II xc2vp7 FPGA. The HDL files were synthesized within Xilinx ISE v8.2 with optimiza-

tion settings set to balanced and KEEP HIERARCHY flag when appropriate to prevent optimization of security-critical components. We base our arithmetic processor on the design from [28].

5.1 Area

In our case, a single arithmetic coprocessor performs serially the computations of Eq. (3) and then that of Eq. (4). This incurs in a very slight area overhead (only the control microcode is slightly modified, plus the masked decoder), at the obvious cost of an increased execution time. In comparison to the unprotected version, our protected decryption scheme consumes more memory as now we store two shares r' and r'' of the secret polynomial r , and the two output polynomials a' and a'' from the two INTT operations.

In Table 1, we can see that the proposed masking of the ring-LWE architecture incurs an additional area overhead of only 301 LUTs and 129 FFs in comparison to the unprotected version. This additional area cost is mostly due to a pair of masked decoders. Due to its low area overhead, we chose to keep two masked decoders in parallel, decoding two coefficients simultaneously. (This nicely fits with the memory organization of the arithmetic coprocessor, since it fits two 13-bit coefficients in each memory word.) Thus, we use two addition and subtraction circuits for the refreshing with Δ_i (accounting for 160 LUTs) and two masked tables (90 LUTs in total).

We note that we could straightforwardly reduce the additional area cost by reusing the 13-bit addition and subtraction circuits present in the arithmetic coprocessor. Since during a decoding operation, the arithmetic coprocessor remains idle, reusing of the addition and subtraction circuits does not cause any increase in the cycle count. For simplicity, we did not implement this approach.

5.2 Cycle count

The cycle count for our approach is decomposed in the computation of Eqs. (3) and (4) and the masked decoder. Equation 3 takes 2840 cycles (one unprotected ring-LWE decryption), while Eq. (4) takes 2590 cycles, slightly less

Table 1 Performance and comparison on Xilinx Virtex-II xc2vp7 FPGA

	LUTs/FFs/DSPs	f_{\max} (MHz)	Cycles
Unprotected	1713/830/1	120	2.8k
Protected	2014/959/1	100	7.5k

Note that these results are not directly comparable with [28], since the latter were obtained from a more advanced Virtex-6 FPGA, which has 6-bit input LUTs and superior routing mechanisms in comparison to our target FPGA

than Eq. (3) since there is no addition present in the second branch.

The two-way parallel masked decoder takes $\frac{1}{2} \times n \times N + \epsilon$ cycles to decode all the coefficients into message bits. In our case with $n = 256$, $N = 16$ the masked decoder takes 2048 cycles. Thus in total, a masked decryption operation requires 7478 cycles. The arithmetic coprocessor and the masked decoder run in constant time and constant flow.

5.3 Comparison with an elliptic-curve cryptosystem

We compare our protected decryption scheme with the unprotected high-speed elliptic curve scalar multiplier architecture proposed by Rebeiro et al. in [22]. The architecture for the field $\text{GF}(2^{233})$ consumes 23,147 LUTs and computes an unprotected scalar multiplication in 12.5 μs on a more advanced Virtex-4 FPGA. Thus, the scalar multiplier has an area \times time product of approximately 289,337. Our protected ring-LWE decryption (for a similar security) achieves an area \times time product of approximately 151,452 on a Virtex-2 FPGA, thus achieving at least 1.9 times better figure of merit.

5.4 Trade-offs

The previous figures are subject to trade-offs. If smaller latency is desired instead of a compact implementation, two coprocessors can perform the two computations of Eqs. (3) and (4) in parallel. Trade-offs also apply to the masked decoder, and the parallelization could be extended easily to reduce latency in this stage. Since the BRAMs present in the Xilinx FPGAs support reading of multiple consecutive words, we could keep more pairs of masked decoders in parallel and reduce the number of cycles. Another alternative is to keep the masked decoder in pipeline with the polynomial arithmetic block. Such type of setting is suitable for systems where many decryption operations are performed in a chain. While the masked decoder works on the coefficients of a previous computation, the polynomial arithmetic unit processes new ciphertexts. Since the masked decoder is faster than the polynomial arithmetic unit, the cycle count of the masked decoder is not an overhead in such type of setting. But of course, in this situation we could not reuse the arithmetic circuitry of the arithmetic coprocessor for the refreshing operation of the masked decoder.

5.5 Maximum frequency

We note that the arithmetic coprocessor is a very optimized unit with a complex pipeline organization. We thus insert two pipeline stages in the masked decoder to match the maximum frequency of the whole system to that of the arithmetic coprocessor. In this way, the design can run up to almost 100 MHz. The critical path is inside the arithmetic multiplier.

6 Software implementation

We wrote a software implementation of the complete system for an ARM Cortex-M4F with the same parameter set as the previous section. This Cortex-M4F is a popular and powerful embedded platform. It has a 32-bit word size, 13 general-purpose registers, its instruction set supports single-cycle 32-bit multiplications and 16-bit SIMD arithmetic.

6.1 Arithmetic operations

Our implementation for the arithmetic part (the two branches from Eqs. (3) and (4)) follows the lines of de Clercq et al. [9]. We remind here the key ideas of the software implementation. Each coefficient requires 13 bits of storage for $q = 7681$, and we therefore store two coefficients in every processor word. We use the negative-wrapped NTT along with computational optimizations from [28] to implement the polynomial multiplication. We can reduce the number of memory accesses, pointer operations, and loop overhead by 50 % by performing a twofold unrolling of the inner loop of the NTT transformation. The expensive calculation of twiddle factors can be avoided by storing precomputed twiddle factors and inverse twiddle factors in a lookup table. The code is constant time and constant flow (SPA resistant.). Since each branch operates on only one share, no special protection against DPA is required.

6.2 Masked decoder

Quadrants The quadrant operation is implemented in a constant-time and constant-flow way. It relies on arithmetic subtraction to perform successive comparisons against $q/4$, $q/2$ and $3q/4$. From these comparisons, the quadrant result is constructed by bitmasks. As in the previous paragraph, since each quadrant operates on a single share, no further DPA protection is required.

Table lookup The table lookup is the most sensitive part since it sees both shares q' and q'' . We mask the table lookup following [1]. This approach takes as input an unprotected software bitsliced implementation written as a straight-line sequence of XOR and AND instructions. Then, the input data are shared in a Boolean fashion and the instructions are replaced by its secure equivalent. The masked XOR operation is very easy to derive; the masked AND instruction is more elaborate due to the non-linearity of the operation. The dataflow for the AND instruction is represented in Fig. 4. It is essentially Trichina's masked AND gate.

We wrote the unmasked function that applies the rules of Sect. 4 (including output feedback) in a bitsliced fashion. We then used espresso [29] and MisII (part of Octtools) for logic minimization and synthesis into XOR and AND "gates"




Fig. 4 Trichina AND gate. This masked computes the unshared function $c = ab$. Each variable a is shared into two shares a_1, a_2 . This is the construction we use for our secure AND instruction in our software implementation

Table 2 Timings for major operations in software

Operation	k Cycles
Equation (3)	43
Inverse NTT transform	39
Masked decoder	168

= instructions. We then substituted the XOR and AND instructions for its secure equivalents. We perform 32 table lookups (for 32 different coefficients) concurrently, and the decoder always performs 16 iterations. This part (a series of XOR and AND) was prototyped in C and the assembly output carefully inspected.

6.3 Timings

In Table 2, we can see an overview of the time required for each major operation. Note that while the arithmetic part is heavily optimized, we did not focus on achieving the fastest implementation in the masked decoder implementation. The most expensive part of arithmetic computation is the inverse NTT, requiring 39k cycles. The computation of Eq. (3) takes around 43k cycles. The masked decoder takes around 168k cycles. (Most of the time goes to computing the quadrant functions. An assembler version for these functions would greatly benefit the overall timing.) The overhead in cycles for the masked version is around 5.8 times more cycles.

7 Discussion

7.1 Error rates

Cryptosystems based on ring-LWE are inherently probabilistic. This means that there is a non-zero probability that the recovered plaintext after ring-LWE decryption is not exactly the plaintext before encryption. In our case, due to the probabilistic nature of our masked decoder approach, there is a

second source of noise. Since the number of iterations of the masked decoder is finite, there are some pair values (a', a'') that will not get decoded within the fixed finite number of iterations. In this section, we first explain the error rate of the probabilistic decoding in isolation, and then we switch to the global system error rate and point out strategies to mitigate it.

Errors due to the probabilistic decoding In this section, we assume that the plaintext bit is 1 and the unmasked input a to the masked decoder is in $(q/4, 3q/4)$. The additional error due to the probabilistic masked decoder is the probability p_e that (a', a'') does not get successfully decoded. Let us write $p_s = 1 - p_e$.

This probability p_s is influenced by two distributions. We have that

$$p_s = \sum \Pr[\text{successful decode}|a] \cdot \Pr[a], \tag{5}$$

where the sum is taken over $a \in (q/4, 3q/4)$. On the one hand, $\Pr[\text{successful decode}|a]$ is the probability that the decoder successfully decodes a . On the other, $\Pr[a]$ is the probability with which a takes various values in $(q/4, 3q/4)$.

The distribution of the decoder success probability $\Pr[\text{successful decode}|a]$ as a function of the unshared input value a to the decoder can be easily computed by averaging over all possible pairs (a', a'') , such that $a' + a'' = a$. Since for any given value of a , its shares a' or a'' are (individually) equiprobable, we compute $\Pr[\text{successful decode}|a]$ as $\Pr[\text{successful decode}|a] = \frac{1}{q} \sum_{a'+a''=a} \Pr[\text{successful decode of } (a', a'')]$.

The distribution $\Pr[\text{successful decode}|a]$ is shown in Fig. 5. We see that the decoder performs best when $a \approx q/2$, in which case all possible inputs get decoded correctly. Only when the input value a approaches $q/4$ or $3q/4$, the performance degrades. When using a larger number of iterations $N = 16$, this effect is less pronounced when compared to $N = 2$ iterations, as Fig. 5 shows.

On the other hand, it is easy to see that not all unshared inputs a to the decoder are equally likely. By the construction of the ring-LWE decryption function, the unshared input to the decoder a is either centered around $q/2$ (resp. 0) when the message bit is 1 (resp. 0). This distribution $\Pr[a]$ is plotted in Fig. 6.

These two observations combined produce a nice interaction between the *prior* distribution $\Pr[a]$ of a (given by the ring-LWE decryption) and the success distribution of the masked decoder $\Pr[\text{successful decode}|a]$ as in Eq. (5). Namely, values of a that are difficult to decode (those with low $\Pr[\text{successful decode}|a]$) are quite unlikely to appear as input to the masked decoder (their $\Pr[a]$ is also low). This positive interaction keeps the global error rate of the system quite low. This is precisely quantified in the next paragraph.




Fig. 5 Empirical success distribution for the masked decoder




Fig. 6 Distribution of a when plaintext is 1

Table 3 Global error rates with the probabilistic decoder

Iterations	$p_g [\times 10^{-5}]$	p_g/p_{baseline}
$N = 2$	332.24	91.41
3	178.44	49.09
4	25.36	6.97
5	20.77	5.71
6	16.22	4.46
8	6.97	1.91
16	4.32	1.19
24	4.06	1.11
30	3.87	1.06

Global error rate and number of iterations We performed simulations to estimate the global error rate and determine the required number of iterations N in our design. Over 10^6 bits, the average error per bit using a deterministic decoder was $p_{\text{baseline}} = 3.634375 \times 10^{-5}$. This is a baseline error intrinsic to the ring-LWE construction. When we plug in the probabilistic decoder, the global, end-to-end, error rate per bit p_g increases. (We have $p_g = p_{\text{baseline}} + p_e$.) In Table 3, we can find the global error rate for different values of the number




Fig. 7 Evolution of the ratio p_g/p_{baseline} as the number of iterations N grows

of iterations N of the decoding. At $N = 3$, for instance, the error rate is $p_g = 1.7844 \times 10^{-3}$, which is ≈ 49 times larger than p_{baseline} . As already hinted, the error rate quickly decreases with N (roughly exponentially, as can be seen in Fig. 7). In our design, we set $N = 16$ (we iterate 16 times per coefficient) as a balanced trade-off between cycle count and error rate. The impact of the masked probabilistic decoder on the global error rate is quite low, adding less than 20 % to the intrinsic error rate when compared to a deterministic decoder, as it can be seen in Table 3. We note that one could generalize the masked decoder to trade area for less number of iterations. For details, see Appendix C.

7.2 Comparison with other decoding strategies

We are only aware of a similar masked decoder, the one presented in [5]. There, the authors resort to a generic masking method, namely masked tables, to perform the decoding. Translating the ideas of [5] in our context, we would need two tables of 2^{13} bits (one of them random). For a smaller group \mathbb{Z}_d with $d = 257$, the authors report an utilization of 1331 slices on a Virtex 6 FPGA. The results in slices are not directly comparable with ours, we point out that the size of the masked table following the approach of [5] grows linearly in the group size q , while for our solution the size of the masked table stays constant (independent of q) and the quadrant blocks grow only logarithmically in q . The cycle count, however, is larger in our solution. The critical observation of our masked decoder is that we can *compress* the input coefficient shares a' and a'' to a mere two-bit per share (the output of each quadrant) and then perform the decoding based on the information of the two quadrants (4 bits).

7.3 Post-processing

Albeit the computation from Eq. (2) is commonly referred as the “ring-LWE decryption”, the decryption process should

include a post-processing on the recovered message m . This post-processing consists of error correction and padding verification.

Linear codes with masking One approach to deal with the probabilistic nature of the ring-LWE decryption system is to use forward error correcting codes (FEC). The message prior to encryption is encoded using a FEC and the resulting composite is ring-LWE encrypted. The output of the ring-LWE decryption should be corrected for errors, preferably in the masked domain. For syndrome decoding of linear codes, this can easily be done by masking the syndrome table. A clever choice of the linear code (for example, perfect codes) can allow very easy masked implementation. (The only perfect linear codes are repetition, Hamming and Golay codes.)

Padding schemes As presented, the ring-LWE system is malleable. CCA security can be achieved with a padding mechanism. The Fujisaki and Okamoto [12] padding scheme is known to work with ring-LWE [19]. This padding scheme makes use of standard symmetric cryptographic constructions whose masked implementations are well studied. We point out that key-encapsulation mechanisms may result in a more compact and simpler implementation.

We remind that the Fujisaki–Okamoto padding scheme requires a negligible decryption error rate for honestly generated ciphertexts,³ as explained by Peikert [19]. Thus, the designer must ensure that the global error rate due to the intrinsic noise of ring-LWE and the probabilistic decoder is negligible. This can be achieved with FEC as previously described.

A formal generic analysis to choose an FEC code that sets the error rate to, say, 2^{-80} or 2^{-128} , is not straightforward. The analysis is greatly simplified if one chooses (n, p, s) parameters such that there is no error contribution due to those parameters and at the same time a required bit security level is maintained. We leave this for a future work.

7.4 Extension to higher-order security

We point out that the approach laid out in Sect. 3 scales quite well with the security order. To achieve security at level $d + 1$, one would need to split the computation of Eq. (2) into d branches analogously to Eq. (3). The masked decoder can follow the same principles with the appropriate modifications. The complexity of this decoder obviously grows. Generic approaches to perform this computation have been discussed in [3, 8, 24].

³ We would like to thank the anonymous reviewer for bringing this important issue to our attention.

8 Evaluation

In this section, we evaluate both the hardware and the software implementations described above.

We provide a very advantageous setting for the adversary: we assume that the evaluator knows the details about the implementation (for example, pipeline stages and register allocation). In addition, we assume that while guessing a subkey, the adversary knows the rest of the key. These assumptions allow to comfortably place predictions on intermediates arbitrarily deep into the computation. While this may represent a very powerful attacker and somewhat unrealistic, the algebraic structure of such cryptosystem may help the attacker to predict deep intermediates with relatively low effort. In the Appendix B, we describe an attack on half-masked ring-LWE decryption that uses these ideas. This stresses the necessity of masking the decoding function entirely.

The evaluation methodology to test if the masking is sound is as follows. We first proceed with the first-order key-recovery attacks when the randomness source (PRNG) is switched off. We demonstrate that in that situation the attacks are successful, indicating that the setup and procedure are sound. Then we switch on the PRNG and repeat the attacks. If the masking is sound, the first-order attacks shall not succeed. In addition, we perform second-order attacks to confirm that the previous first-order analyses were carried out with enough traces.

We modeled the power consumption as the Hamming distance between two consecutive values held in a register, and used Pearson's correlation coefficient to compare predictions with measurements [6].

8.1 Hardware implementation evaluation

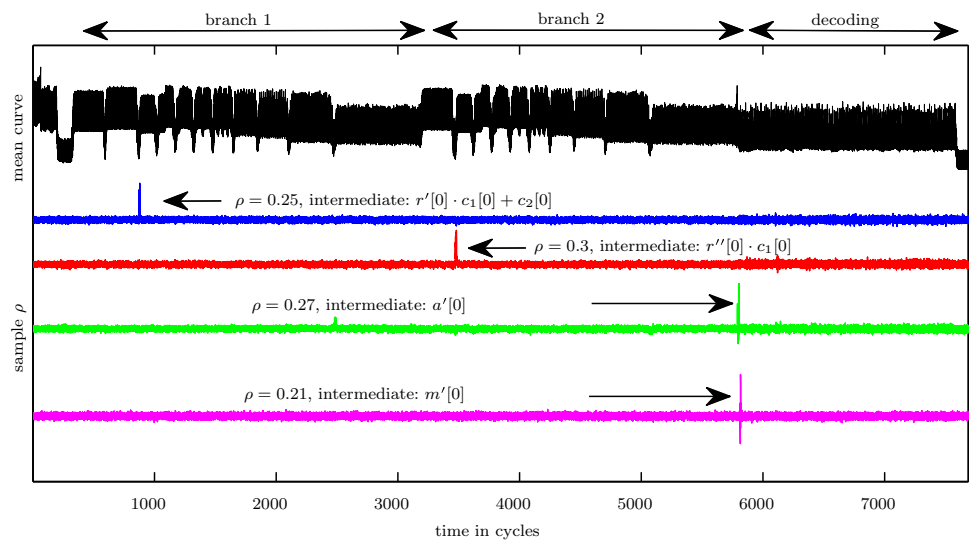
Measurement setup We implemented the full design on a SASEBO G board. The design was clocked at 18.75 MHz and the power consumption was sampled at 500 MS/s. This platform is of very low noise.

We test 4 different points which covers all the relevant parts of the computation. The targets are the first 13-bit coefficient of $r' \cdot c_1 + c_2$, the first 13-bit coefficient of $r'' \cdot c_1$, the first input coefficient to the shared decoder and the first output bit.

PRNG off We first begin the experiments when the PRNG is off. That is, the sharing of r into r' and r'' on each execution is deterministic. This would not happen in practice, as an active PRNG would randomize the representation of r in each execution. In our setting, this would mean that the masking is switched off.

In Fig. 8 we draw the result of correlating against the 4 intermediates with 10,000 traces. On top, we draw a mean

Fig. 8 Hardware implementation. PRNG off. On top, black, one power consumption trace. The different computational stages can be distinguished: first branch, second branch and decoding. Next, in blue, the correlation trace for the value is $r'[0] \cdot c_1[0] + c_2[0]$. The correlation achieves a maximum value of $\rho = 0.25$. Below, in red, correlation for $r'' \cdot c_1$ (max $\rho \approx 0.3$); in green correlation for the input of the masked decoder $a'[0]$. At the bottom, correlation with one message bit $m'[0]$



trace for orientation. The correlation values are, from top to bottom, 0.25, 0.3, 0.27 and 0.21, respectively. This means that the attacks are successful, and confirms the soundness of our setting. In Fig. 9, we can see the evolution of the




Fig. 9 Hardware implementation. PRNG off. Evolution of the correlation coefficient as the number of traces increases for the intermediates $r'[0] \cdot c_1[0] + c_2[0]$ (left) and $r''[0] \cdot c_1[0]$ (right). Correct subkey guess in red, all other guesses in green. A 99.99 % confidence interval for $\rho = 0$ is plotted in black discontinuous line. We can see that starting from hundred measurements the attacks are successful

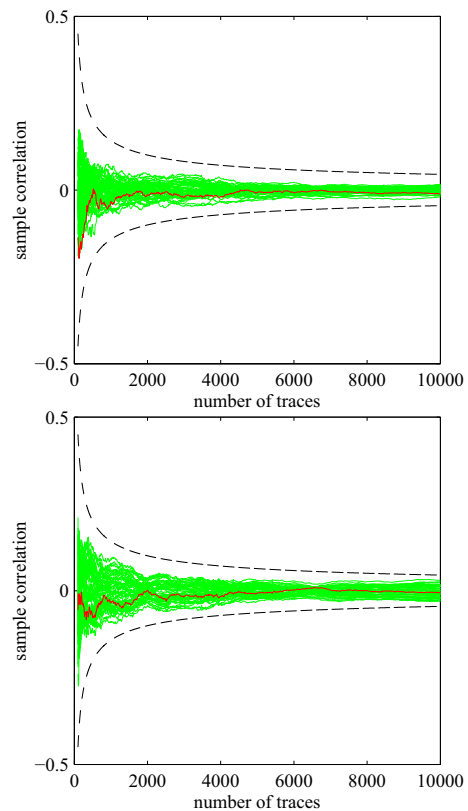



Fig. 10 Hardware implementation. Analogous to Fig. 9, but with PRNG on. The correct subkey is no longer identifiable. This is expected and means that the masking is effective

correlation coefficient as the number of traces increases for the first two intermediates. We can see that starting from 100 traces, the attack is successful. A similar behavior was observed for other intermediates.

PRNG on In Fig. 10, we draw the result of the previous analysis when the masks are switched on. This corresponds to the

Fig. 11 Hardware implementation. Correlation traces for intermediates within the shared decoder. On top, a power measurement trace showing the last 15 decodings. Below, correlation traces. The first two (masks and masked values) assume that the adversary knows the masks. The third one, in light blue, is a first-order attack without knowing the attack, and is unsuccessful. In contrast, the second-order attack against the same intermediate is successful, as the traces in magenta and yellow show



situation that an adversary would face in reality. We can see that the correct key guess is no longer distinguishable, even when using 10,000 traces. We repeated the same experiments for other intermediates and other intermediate positions with identical results.

Second-order attacks To confirm that we used enough traces in our previous analyses, we perform here second-order attacks on the masked implementation with the PRNG on. We will focus on the masked decoder. In Fig. 11, we draw on the top a mean curve in the region of 7400–7700 cycles, corresponding to the end of the masked decoding. We target one output bit of the decoding: $m[254]$.

In Fig. 11, we first begin by correlating against masks and masked values. This is a test scenario, since for this attack we need to know the masks, something that would not happen in a real deployment. Correlations with masks or masked value yield high correlation as expected ($\rho = 0.32$ and $\rho = 0.34$, respectively). In contrast, when correlating against the unshared value (in light blue), the correlation coefficient does not traverse the confidence interval for $\rho = 0$. This indicates that the masking is effective. We can repeat the same attack against centered and squared traces [7, 21]. This is effectively a second-order attack and is expected to work. It is shown in magenta in Fig. 11, and we can see that the attack succeeds. Using the centered absolute value to pre-process traces also works as expected, as shown in yellow.

In Fig. 12, we can see the evolution as a function of the number of traces. We can see that starting from ≈ 2000 measurements, this second-order attack is successful. This confirms that the first-order attacks from above were carried out with enough traces, since a second-order attack is already successful starting from ≈ 2000 measurements.

We remark that the relatively low number of traces required for the second-order attack is due to the very friendly

scenario for the evaluator. The platform is of low noise and no other countermeasure except masking was implemented. In practice, masking needs a source of noise to be effective,




Fig. 12 Hardware implementation. Top correlation as the number of traces increases for the first-order attack (PRNG on), around clock cycle 7560. Bottom correlation for the second-order attack with masks on. The attack begins to be successful with 2000 measurements





Fig. 13 Hardware implementation. Crosscorrelation trace. The x and y axes represent time, flowing from the *upper left* hand side corner to the *lower right*. The entire figure spans 7500 cycles (as Fig. 8). It is possible to distinguish the two branch computations (including its components) and the decoding. *Colors* enhanced to improve contrast

and consequently the higher-order attacks would be harder to mount, requiring more traces [7] and more computation [25] (Fig. 13).

8.2 Software implementation evaluation

Measurement setup We deployed the masked software implementation on a 32-bit ARM STM32F407VG Cortex-M4. The MCU operates at 168 MHz and has 192 kB of SRAM. We take contactless power measurements from a decoupling capacitor in the power loop with a Langer LF2-5 H-field probe and 20 dB amplification. This laboratory setup is of very low noise. DPA on an unprotected byte-oriented AES suc-

Fig. 14 Software implementation. PRNG off. On *top*, *black* one EM consumption trace. The selected region comprises one masked bitsliced table lookup. *Below* different correlation traces for various intermediates. Correlation for different shares of the same intermediate are drawn in the *same color*. The 99.99 % confidence interval for $\rho = 0$ is drawn in *light gray*



ceeds with 20 traces. We focus the evaluation on the most challenging part: the masked decoding operation.

Masks off Figure 14 shows the successful correlations when the adversary knows the secret PRNG seed. This serves to confirm that our setup is sound. We selected many different intermediates within the table lookup operation and used 20k traces to produce a good-looking picture. The maximum absolute value for the correlation against the correct key hypothesis is around $|\rho| \approx 0.71$. In Fig. 15, top, we see the evolution of sample correlation coefficient as the number of curves at timesample 1390. We can see that starting from less than 100 traces, the attack is successful, since the correct subkey stands out from all other competing key hypotheses.

Masks on When the PRNG output is unknown, first-order attacks are expected not to work. This is the case in our implementation. In Fig. 15, middle, the evolution of the correlation coefficient is plotted at the same timesample 1390. The correct subkey is indistinguishable among competing ones. Similar observations apply to the entire timespan.

Second-order attacks We also performed second-order attacks. Note that our implementation does not claim second-order security. One can see from Fig. 15, bottom, that second-order attacks begin to work from a couple of hundred measurements. This means that the previous analyses were carried out with enough number of measurements (up to 20k measurements). Similar observations apply here: our software setting is very friendly toward the evaluator, since there is no additional noise present in the measurements. In reality, one would always implement masking along with a source of noise to be effective.




Fig. 15 Software implementation. Evolution of Pearson’s correlation coefficient with the number of traces for different attacks at timesample 1390. On *top* (successful) first-order attack with PRNG off. *Middle* (unsuccessful) first-order attack with PRNG on. *Bottom* (successful) second-order attack with PRNG on. Correct subkey in *red*, and incorrect in *green*. We also plot the 99.99 % confidence interval for $\rho = 0$ in *dashed line*

8.3 Horizontal DPA attacks

During the decoder operation, the input coefficients are refreshed $N - 1 = 15$ times with publicly known offsets Δ_i . The device thus handles consecutively the values $a', a' + \Delta_1, \dots, a' + \Delta_1 + \dots + \Delta_{15}$. This may enable a horizontal DPA attack [18] during the operation: the adversary may collect a single trace, split it into 16 chunks and then perform a DPA on these 16 chunks to recover the mask a' . Once the masks from all traces are discovered, a first-order, vertical DPA applies.

There are two factors that mitigate this threat. First, we note that the adversary is given a very limited number of traces to recover each mask (namely, $N = 16$). Secondly, this attack can be easily prevented by shuffling the public coefficients Δ_i . This randomizes the order of execution of each refreshing with Δ_i , and thus the exposure to horizontal DPA attacks is minimized.

9 Conclusion

In this paper, we described a practical side-channel protected implementation of the lattice-based ring-LWE asymmetric decryption. Our solution is based on the sound principles of masking and incurs in a manageable overhead (in cycles and

area). A key component of our solution is a bespoke masked decoder. Our implementation performs the entire ring-LWE decryption computation in the masked domain.

Acknowledgments The authors would like to thank the CHES 2015 reviewers for their valuable comments. This work has been supported in part by the European Commission through the ICT programme under contracts H2020-ICT-645622 PQCRYPTO, H2020-ICT-644209 HEAT and FP7-ICT-2013-10-SEP-210076296 PRACTICE, by the Research Council KU Leuven TENSE (GOA/11/007); by the Flemish Government FWO G.0550.12N, G.00130.13N and G.0876.14N; and by the Hercules Foundation AKUL/11/19. Oscar Reparaz is funded by a PhD fellowship of the Fund for Scientific Research-Flanders (FWO). Sujoy Sinha Roy was supported by Erasmus Mundus PhD Scholarship.

Appendix A: Optimal values of Δ_i for $q = 7681$

- $\Delta(i) = (960, 1440, 480, 1680, 240, 720, 1200, 1800,$ (6)
- $120, 360, 600, 840, 1080, 1320, 1560, 1860,$ (7)
- $60, 180, 300, 420, 540, 660, 780, 900, 1020,$ (8)
- $1140, 1260, 1380, 1500, 1620, 1740, 1890,$ (9)
- $30, 90, 150, 210, 270, 330, 390, 450, 510,$ (10)
- $570, 630, 690, 750, 810, 870, 930, 990, 1050,$ (11)
- $1110, 1170, 1230)$ (12)

These values were found by exhaustive first-order search. The value Δ_i is chosen so that it maximizes the number of pairs that get decoded after i iterations.

Appendix B: Attack on half-masked variant

In this section, we analyze the security of a masked ring-LWE variant where the intermediates just before decoding are unmasked, and the decoding is performed in the unmasked domain. This alternative is definitely cheaper than full masking.

In the following, we provide evidence to show that this clearly does not provide enough security in our case.

(A seemingly similar situation appears in [5]. However, there are important differences—namely, it is not possible to choose ciphertexts. In the following, we are not analyzing the variant of [5], but only the half-masked ring-LWE.)

A common argument is that after key diffusion is complete, prediction of the intermediates is not possible and hence standard DPA attacks to the half-masked ring-LWE do not apply. We will see that this is not strictly true, if the attacker can choose ciphertexts.

Assume that the coefficients of the polynomial $a = \text{INTT}(r \cdot c_1 + c_2)$ appear unmasked in the implementation. Let the adversary collect measurements with the chosen cipher-

text. The ciphertext c_1 has the following structure: all the coefficients are fixed except $c_1[0]$ that is randomly varying. The ciphertext c_2 has the same structure. Then observe that due to linearity of the INTT operation, $a[0]$ can be written as $a[0] = \alpha(r[0] \cdot c_1[0] + c_2[0]) + \beta$, where

- α is a public constant determined by the INTT transformation.
- β is a secret constant that is a function of the other (unknown) key coefficients $r[1], \dots, r[255]$. Note that by construction, β is constant within the set of collected traces.

Thus, an attacker can perform a DPA attack targeting the intermediate $a[0]$ and placing predictions on $(r[0], \beta)$. The adversary recovers $r[0]$ and proceeds to recover other key coefficients. We have verified this attack in simulations, even when using $\text{th}(a[i])$ as intermediate.

(It may seem that the high number of hypotheses, 2^{26} , may produce a cumbersome attack. However, one can apply techniques of partial correlation [6] to alleviate the computational effort of DPA on large word sizes [31]. We have experimented that in practice it makes sense to first recover $r[0]$ (this is easier due to larger non-linearity of the modular multiplication) and then β (which may be harder due to the low non-linearity of the modular addition), splitting the 2^{26} effort in two 2^{13} steps.)

Appendix C: Generalization of the decoding scheme

The probability of not hitting any rule can be reduced by increasing the number of rules, i.e., by splitting the domain of decoding into more than four sections. For example, in Table 4, the rules are shown for the case when the decoding domain is split into eight sections or octant. As seen from the table, the probability of not hitting a rule has reduced to 1/4. Hence to meet a same decryption failure rate, an

Table 4 The rules for octant decoding

a'	a''							
	I_1	I_2	I_3	I_4	I_5	I_6	I_7	I_8
$I_1 = (0, q/8)$	✓	×	✓	✓	✓	×	✓	✓
$I_2 = (q/8, 2q/8)$	×	✓	✓	✓	×	✓	✓	✓
$I_3 = (2q/8, 3q/8)$	✓	✓	✓	×	✓	✓	✓	×
$I_4 = (3q/8, 4q/8)$	✓	✓	×	✓	✓	✓	×	✓
$I_5 = (4q/8, 5q/8)$	✓	×	✓	✓	✓	×	✓	✓
$I_6 = (5q/8, 6q/8)$	×	✓	✓	✓	×	✓	✓	✓
$I_7 = (6q/8, 7q/8)$	✓	✓	✓	×	✓	✓	✓	×
$I_8 = (7q/8, 8q/8)$	✓	✓	×	✓	✓	✓	×	✓

The cases where no rule is hit are marked with ×

octant decoder needs almost half the number of iterations as required by a quad decoder. However, there are overheads associated with an octant decoder when it is compared to a quad decoder: the number of comparisons to locate the position of a coefficient in the octant chart doubles and the sizes of the tables quadruple.

References

- Balasch, J., Gierlichs, B., Reparaz, O., Verbauwhede, I.: Dpa, bit-slicing and masking at 1 ghz. In: Güneysu, Handschuh [15], pp. 599–619
- Bernstein, D.J., Buchmann, J., Dahmen, E.: Post Quantum Cryptography, 1st edn. Springer, Berlin (2008)
- Bilgin, B., Gierlichs, B., Nikova, S., Nikov, V., Rijmen, V.: Higher-order threshold implementations. In: ASIACRYPT, Volume 8874 of LNCS, pp. 326–343. Springer, Berlin (2014)
- Bos, J.W., Lauter, K., Loftus, J., Naehrig, M.: Improved security for a ring-based fully homomorphic encryption scheme. In: Cryptography and Coding, Volume 8308 of LNCS, pp. 45–64. Springer, Berlin (2013)
- Brenner, H., Gaspar, L., Leurent, G., Rosen, A., Standaert, F.-X.: FPGA implementations of SPRING—and their countermeasures against side-channel attacks. In: CHES, Volume 8731 of LNCS, pp. 414–432. Springer, Berlin (2014)
- Brier, E., Clavier, C., Olivier, F.: Correlation power analysis with a leakage model. In: CHES, Volume 3156 of LNCS, pp. 16–29. Springer, Berlin (2004)
- Chari, S., Jutla, C.S., Rao, J.R., Rohatgi, P.: Towards sound approaches to counteract power-analysis attacks. In: CRYPTO, Volume 1666 of LNCS, pp. 398–412. Springer, Berlin (1999)
- Coron, J.S.: Higher order masking of look-up tables. In: EUROCRYPT, Volume 8441 of LNCS, pp. 441–458. Springer, Berlin (2014)
- de Clercq, R., Roy, S.S., Vercauteren, F., Verbauwhede, I.: Efficient software implementation of ring-LWE encryption. In: Nebel, W., Atienza, D. (eds.) Proceedings of the 2015 Design, Automation and Test in Europe Conference and Exhibition, DATE 2015, Grenoble, France, March 9–13, 2015, pp. 339–344. ACM (2015)
- Ducas, L., Durmus, A., Lepoint, T., Lyubashevsky, V.: Lattice signatures and bimodal gaussians. In: CRYPTO, Volume 8042 of LNCS, pp. 40–56. Springer, Berlin (2013)
- Fan, J., Vercauteren, F.: Somewhat practical fully homomorphic encryption. Cryptology ePrint Archive, Report 2012/144 (2012). <http://eprint.iacr.org/>
- Fujisaki, E., Okamoto, T.: Secure integration of asymmetric and symmetric encryption schemes. J. Cryptol. **26**(1), 80–101 (2013)
- Göttert, N., Feller, T., Schneider, M., Buchmann, J., Huss, S.: On the design of hardware building blocks for modern lattice-based encryption schemes. In: CHES, Volume 7428 of LNCS, pp. 512–529. Springer, Berlin (2012)
- Goubin, L., Patarin, J.: DES and differential power analysis the duplication method. In: CHES, Volume 1717 of LNCS, pp. 158–172. Springer, Berlin (1999)
- Kocher, P.: Timing attacks on implementations of Diffie–Hellman, RSA, DSS, and other systems. In: CRYPTO, Volume 1109 of LNCS, pp. 104–113. Springer, Berlin (1996)
- Kocher, P., Jaffe, J., Jun, B.: Differential power analysis. In: CRYPTO, Volume 1666 of LNCS, pp. 388–397. Springer, Berlin (1999)
- Lyubashevsky, V., Peikert, C., Regev, O.: On ideal lattices and learning with errors over rings. In: EUROCRYPT, Volume 6110

- of LNCS, pp. 1–23. Springer, 2010. Full Version available at Cryptology ePrint Archive, Report 2012/230
18. Pan, J., den Hartog, J.I., Lu, J.: You cannot hide behind the mask: power analysis on a provably secure s-box implementation. In: Information Security Applications, Volume 5932 of LNCS, pp. 178–192. Springer, Berlin (2009)
 19. Peikert, C.: Lattice cryptography for the internet. In: Post-Quantum Cryptography—6th International Workshop, PQCrypto 2014, Waterloo, ON, Canada, October 1–3, 2014. Proceedings, pp. 197–219 (2014)
 20. Pöppelmann, T., Güneysu, T.: Towards practical lattice-based public-key encryption on reconfigurable hardware. In: Selected Areas in Cryptography—SAC 2013, Volume 8282 of LNCS, pp. 68–85. Springer, Berlin (2014)
 21. Prouff, E., Rivain, M., Bevan, R.: Statistical analysis of second order differential power analysis. *IEEE Trans. Comput.* **58**(6), 799–811 (2009)
 22. Rebeiro, C., Roy, S.S., Mukhopadhyay, D.: Pushing the limits of high-speed $GF(2^m)$ elliptic curve scalar multiplication on fpgas. In: CHES, Volume 7428 of LNCS, pp. 494–511. Springer, Berlin (2012)
 23. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. In: Proceedings of the Thirty-seventh Annual ACM Symposium on Theory of Computing, STOC '05, pp. 84–93, New York, NY, USA, 2005. ACM
 24. Reparaz, O., Bilgin, B., Nikova, S., Gierlichs, B., Verbauwhede, I.: Consolidating masking schemes. In: CRYPTO, Volume 9215 of LNCS, pp. 764–783. Springer, Berlin (2015)
 25. Reparaz, O., Gierlichs, B., Verbauwhede, I.: Selecting time samples for multivariate DPA attacks. In: CHES, Volume 7428 of LNCS, pp. 155–174. Springer, Berlin (2012)
 26. Reparaz, O., Roy, S.S., Vercauteren, F., Verbauwhede, I.: A masked ring-lwe implementation. In: Güneysu and Handschuh [15], pp. 683–702 (2015)
 27. Roy, S.S., Reparaz, O., Vercauteren, F., Verbauwhede, I.: Compact and side channel secure discrete gaussian sampling. In: IACR Cryptology ePrint Archive, vol. 2014, p. 591 (2014)
 28. Roy, S.S., Vercauteren, F., Mentens, N., Chen, D.D., Verbauwhede, I.: Compact ring-lwe cryptoprocessor. In: CHES, Volume 8731 of LNCS, pp. 371–391. Springer, Berlin (2014)
 29. Rudell, R.L.: Multiple-valued logic minimization for pla synthesis. Technical report, DTIC Document (1986)
 30. E.V. Trichina. Table lookup operation on masked data, 2013. US Patent 8,422,668
 31. Tunstall, M., Hanley, N., McEvoy, R.P., Whelan, C., Murphy, C.C., Marnane, W.P.: Correlation power analysis of large word sizes. In: IET Irish Signals and Systems Conference (ISSC) 2007, September 2007. Available at <http://www.cs.bris.ac.uk/home/tunstall/papers/THMWMM.pdf>