

A Note on Ring-LWE Security in the Case of Fully Homomorphic Encryption

Guillaume Bonnoron¹ and Caroline Fontaine²(✉)

¹ Chair of Naval Cyber Defense – Ecole Navale, CC600,
29240 Brest Cedex 9, France

`guillaume.bonnoron@ecole-navale.fr`

² CNRS/Lab-STICC – IMT Atlantique, CS 83818, 29238 Brest Cedex 3, France
`caroline.fontaine@imt-atlantique.fr`

Abstract. Evaluating the practical security of Ring-LWE based cryptography has attracted lots of efforts recently. Indeed, some differences from the standard LWE problem enable new attacks. In this paper we discuss the security of Ring-LWE as found in Fully Homomorphic Encryption (FHE) schemes. These FHE schemes require parameters of very special shapes, that an attacker might use to its advantage. First we present the specificities of this case and recall state-of-the-art attacks, then we derive a new special-purpose attack. Our experiments show that this attack has unexpected performance and confirm that we need to study the security of special parameters sets carefully.

Keywords: Homomorphic encryption · Concrete security
Implementation · Ring-LWE

1 Introduction

The Learning With Errors over Rings (Ring-LWE) problem has been introduced by [LPR10] as a ring variant of the Learning With Errors due to Regev [Reg05]. Both problems enjoy security reductions to hard lattice problems (SIVP for LWE and SVP in ideal lattices for Ring-LWE), so they serve as hardness grounds for many cryptographic constructions, among others homomorphic encryption. See the survey from Peikert [Pei15] for an extensive retrospective. Today, the paramount question, that still stands in the way to practical use, concerns the security of concrete instances of these problems. Namely, how shall one choose parameters for these problems to meet a security level objective, say 80 bits of security?

Homomorphic encryption (HE) is a type of encryption that allows to compute over encrypted data. The result, once decrypted, equals that of the same computation done over the plain data. HE enables many new applications because one no longer needs to trust the computing entity, e.g. cloud service providers. Fully homomorphic encryption (FHE) was first achieved by the

Funded and supported by Ecole navale, IMT Atlantique, Thales and Naval Group.

© Springer International Publishing AG 2017

A. Patra and N. P. Smart (Eds.): INDOCRYPT 2017, LNCS 10698, pp. 27–43, 2017.

https://doi.org/10.1007/978-3-319-71667-1_2

ground-breaking works of Gentry [Gen09] and Aguilar et al. [AMGH10]. Since then, huge efforts have been spent and have led to numerous scheme proposals, e.g. [vDGHV10, BGV12, FV12, GSW13, DS15, DM15, CCGH16]. Some of the most efficient ones are now implemented, for early users: HELib [Hal], SEAL [DGBL+15], FV-NFLlib [Cry]. These efficient schemes are known secure under the assumption that Ring-LWE is intractable. For an application, the parameters need to be chosen to guarantee the expected security level. This choice of parameters is also constrained so that the scheme itself works.

Related work: This need for a better understanding of the Ring-LWE security, *in practice*, has already driven some studies. For example, Albrecht et al. [APS15] offer a complete overview of the known attacks against LWE. They also give estimates of their costs against some LWE-based cryptographic schemes, not only FHE. We recall them briefly below for our later discussion. Another more recent line of work has been developed against Ring-LWE specifically, taking advantage of the underlying ring structure, see [Pei16] for a summary and guidelines to draw immune parameters.

Today, the estimates we can use to choose parameters are from the works of Lindner and Peikert [LP11] or Van de Pol and Smart [vdPS13]. Also Albrecht et al. maintain an LWE security estimator based on models of state-of-the-art attacks¹. However none of them include special-purpose attacks for FHE settings.

Our work: It seems to us that a focus on Ring-LWE-based FHE scheme is needed. As we see later, in order to keep correctness in the homomorphic schemes, an application designer needs to pick very special parameters. It seemed unclear to us what concrete advantage an attacker might have in such cases. Our contribution aims at filling this gap. After reviewing state-of-the-art attacks, we derive a new one, specially designed for this case, and present experimental results.

Roadmap: Introducing notation and definitions in Sect. 2, we review the state-of-the-art of the attacks against Ring-LWE in Sect. 3. In Sect. 4 we present our new attack in details, its performance in Sect. 5 and draw some conclusions in Sect. 6.

2 Preliminaries

2.1 Notation

For a positive integer $q > 0$ we note \mathbb{Z}_q the set of elements $\{0, \dots, q - 1\}$.

For $n > 0$ integer, the matrix \mathbf{I}_n refers to the identity matrix of size n . Capital bold letters are used for matrices, e.g. \mathbf{B} , and small ones for (row) vectors, e.g. \mathbf{v} . We similarly write \mathbf{B} for a matrix or for the ordered family of (row) vectors $\mathbf{B} = (\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_n)$ using bold subscripts. For a vector \mathbf{v} , we refer to its components with italic subscripts v_i .

¹ <https://bitbucket.org/malb/lwe-estimator/overview>.

When dealing with a polynomial \mathbf{p} of degree $n - 1$, we use the coefficient embedding and assimilate it as a coordinate vector: $\mathbf{p} = (\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_n)$. We will work with polynomials in polynomial ring $R = \mathbb{Z}[x]/(f(x))$ for some $f \in \mathbb{Z}[x]$. We denote R_q the set of polynomials in R with coefficients in \mathbb{Z}_q .

2.2 Lattices

General definitions. In general, a lattice \mathcal{L} of dimension n is a discrete additive subgroup of \mathbb{R}^n . *Integer* lattices are discrete additive subgroups of \mathbb{Z}^n . In this paper we only work with the integer lattices and simply call them lattices.

Lattices (of size n) are usually represented by a basis \mathbf{B} , a set of n independent integer vectors $(\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_n)$ of size n whose integer linear combinations generate the lattice.

$$\mathcal{L}(\mathbf{B}) = \left\{ \sum_{i=1}^n v_i \mathbf{b}_i : v_i \in \mathbb{Z} \right\} = \{ \mathbf{B}^T \mathbf{v} : \mathbf{v} \in \mathbb{Z}^n \} = \mathbf{B} \mathbb{Z}^n$$

In our lattices, \mathbf{B} is always a square integer matrix, $\mathbf{B} \in \mathbb{Z}^{n \times n}$. For most of the discussion we restrict to this full rank definition and will make it explicit when working with greater generating families.

An invariant of a lattice is its determinant $\det(L)$. It is defined as the absolute value of the determinant of any of its bases.

$$\det(L) = |\det(\mathbf{B})|$$

Gram-Schmidt Orthogonalization (GSO). We refer several times to the GSO of a basis. This algorithm takes the matrix to orthogonalize and outputs the resulting matrix \mathbf{B}^* and a matrix μ (lower triangular with 1 on the diagonal) such that $\mathbf{B} = \mu \times \mathbf{B}^*$. We construct \mathbf{B}^* so that its vectors verify:

- $\mathbf{b}_1^* = \mathbf{b}_1$
- \mathbf{b}_i^* is the projection of \mathbf{b}_i orthogonally to the subspace generated by the $i - 1$ first vectors of \mathbf{B} .

It works in polynomial time in the size of the matrix.

q-ary lattices. When studying LWE and Ring-LWE problems, the lattices we mostly work with are called *q-ary*, because they are defined *modulo* some integer q (not necessarily prime). These lattices are defined as follows:

$$\begin{aligned} \mathcal{L}_q(\mathbf{B}) &= \left\{ \sum_{i=1}^n v_i \mathbf{b}_i \bmod q : v_i \in \mathbb{Z} \right\} \\ &= \{ \mathbf{B}^T \mathbf{v} \bmod q : \mathbf{v} \in \mathbb{Z}^n \} \end{aligned}$$

Hence, since we are working modulo q , we can equivalently consider that all the components of \mathbf{B} and \mathbf{v} are in \mathbb{Z}_q .

Lattice reduction. As we discuss below, working with lattices is better when we have *nice* bases. Therefore we have lattice reduction algorithms that, given a basis, make it nicer according to some criteria. There is a wealth of studies on this computational problem.

We usually consider the following criteria and say that a basis is:

- η -size-reduced if $\forall i < j, |(\mathbf{b}_j | \mathbf{b}_i^*)| \leq \eta \cdot \|\mathbf{b}_i^*\|^2$
- δ -LLL-reduced if it is size-reduced and $\forall i, \delta \|\mathbf{b}_i^*\|^2 \leq \left(\|\mathbf{b}_{i+1}^*\|^2 + \frac{(\mathbf{b}_{i+1} | \mathbf{b}_i^*)^2}{\|\mathbf{b}_i^*\|^2} \right)$
- β -BKZ-reduced if it is LLL-reduced and for all j , \mathbf{b}_j^* is the shortest vector of the sublattice spanned by $(\mathbf{b}_j, \dots, \mathbf{b}_k)$ with $k = \min(j + \beta - 1, n)$

We have algorithms to achieve such reductions:

- The celebrated LLL algorithm from Lenstra et al. [LLL82], running in polynomial time of the dimension and the size of the elements. Improvements have been proposed since then and are available in current implementations [NS05].
- The Blockwise Korkine-Zolotarev algorithm [SE94, CN11] achieves BKZ reduction. It makes a polynomial (in n) call to a SVP oracle in a sublattice of size β . It behaves roughly as a sub-exponential in the basis quality [GN08b] and uses LLL as a sub-routine.
- Slide Reduction [GN08a] is another block algorithm, in the spirit of BKZ, but simpler to express and analyze, whose performance approaches that of BKZ [MW15].

The quantity we usually use to measure the reduction quality, independently of the algorithm, is the *root Hermite factor* γ . It is defined as: $\|\mathbf{b}_1\| = \gamma^n \det(L)^{1/n}$ where \mathbf{b}_1 is the smallest vector of the basis we qualify. The higher the quality, the smaller it gets.

2.3 Learning with Errors and Ring Variant

We recall here the definitions of the Learning With Errors problem [Reg05] and the Ring-LWE variant [LPR10]. Both exist in a *search* version and a *decision* version.

Definition LWE. Let n, q be positive integers, χ a probability distribution on \mathbb{Z} of standard deviation σ and \mathbf{s} a secret random vector in \mathbb{Z}_q^n . We denote by $L_{\mathbf{s}, \chi}$ the probability distribution on $\mathbb{Z}_q^n \times \mathbb{Z}_q$ obtained by choosing $\mathbf{a} \in \mathbb{Z}_q^n$ uniformly at random, choosing $e \in \mathbb{Z}$ according to χ and considering it in \mathbb{Z}_q , and returning $(\mathbf{a}, c) = (\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle + e) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$.

Decision-LWE is the problem of deciding whether given pairs (\mathbf{a}, c) are sampled according to $L_{\mathbf{s}, \chi}$ or the uniform distribution on $\mathbb{Z}_q^n \times \mathbb{Z}_q$.

Search-LWE is the problem of recovering \mathbf{s} from pairs (\mathbf{a}, c) sampled from $L_{\mathbf{s}, \chi}$.

Definition Ring-LWE. Let R be a ring of degree n over \mathbb{Z} (usually $R = \mathbb{Z}[x]/(f(x))$ for some cyclotomic polynomial $f(x)$). Let q be a positive integer, χ a probability distribution on R of standard deviation σ and \mathbf{s} a secret random element in R_q . We denote by $L_{\mathbf{s},\chi}$ the probability distribution on $R_q \times R_q$ obtained by choosing $\mathbf{a} \in R_q$ uniformly at random, choosing $\mathbf{e} \in R$ according to χ and considering it in R_q , and returning $(\mathbf{a}, \mathbf{c}) = (\mathbf{a}, [\mathbf{a} \cdot \mathbf{s} + \mathbf{e}]_q) \in R_q \times R_q$.

Decision-Ring-LWE is the problem of deciding whether given pairs (\mathbf{a}, \mathbf{c}) are sampled according to $L_{\mathbf{s},\chi}$ or the uniform distribution on $R_q \times R_q$.

Search-Ring-LWE is the problem to recovering \mathbf{s} from pairs (\mathbf{a}, \mathbf{c}) sampled from $L_{\mathbf{s},\chi}$.

The hardnesses of (Ring-)LWE problems depend on the three variables n , σ and q . The hardness reductions presented in the introductory papers stands when $\sigma > 2\sqrt{n}$. Besides that, we aim to establish a link between a choice of parameters and the provided security level, in the context of FHE.

2.4 Ring-LWE Based FHE Schemes

Since the works of Gentry [Gen09] and Aguilar et al. [AMGH10], lots of homomorphic encryption schemes have been proposed. They can be divided into two families: those based on integers and those based on lattices. LWE and Ring-LWE serve as building ground for the latter family. Generally, those based on Ring-LWE derive from an equivalent scheme on LWE and are more efficient in terms of space and/or time. The most common Ring-LWE based schemes are: [BGV12], [FV12] and SHIELD [KGV15].

For expository purpose, we recall here only elements of the [FV12] scheme. The discussion remains valid for all schemes as well. The interested reader should refer to the original papers for extensive details about the schemes. In our present study, we are interested in the elements that will be the attack target, namely the public key. In [FV12] the key generation process goes as follows:

1. **FV.ParameterChoice**(λ): choose (n, σ, q) to guarantee of level of security λ and set $R = \mathbb{Z}[x]/(\Phi_n(x))$
2. **FV.KeyGen**(n, σ, q): sample $\mathbf{s} \leftarrow R_2$, $\mathbf{a} \leftarrow R_q$, $\mathbf{e} \leftarrow \chi$ and output

$$\mathbf{sk} = \mathbf{s} \text{ and } \mathbf{pk} = ([-(\mathbf{a} \cdot \mathbf{s} + \mathbf{e})]_q, \mathbf{a})$$

Security. Let aside the sign of the first element, the public key \mathbf{pk} is exactly a Ring-LWE pair as described above. The objective of the attacker is to solve *search-Ring-LWE* (i.e. find \mathbf{s}) when given access to this public key and the public parameters $(n, \sigma, q$ and $R)$.

Correctness. The condition for the scheme to be correct (i.e. decryption yields a result consistent with the computation) is [FV12, Eq. 6]:

$$4 \cdot \beta(\epsilon) \cdot \delta_R^{L_{\max}} \cdot (\delta_R + 1.25)^{L_{\max}+1} \cdot t^{L_{\max}-1} < \frac{q}{\sigma}$$

where

- L_{\max} is the maximum multiplicative depth before bootstrapping is needed.
- t is the size of the plaintext space (R_t).
- $\delta_R = \max\{\|\mathbf{a} \cdot \mathbf{b}\| / (\|\mathbf{a}\| \cdot \|\mathbf{b}\|) : \mathbf{a}, \mathbf{b} \in R\}$, e.g. $\delta_{\mathbb{Z}[x]/(x^n+1)} = n$.
- $\beta(\epsilon)$ is such that the error samples are bounded by $B = \beta(\epsilon) \cdot \sigma$ with probability $1 - \epsilon$, e.g. $\beta(2^{-64}) \approx 9.2$.

As a result of this constraint, we can see that q will have to be very large, in front of σ , to conserve correctness with interesting depth, say $L = 9$. To prevent q from being enormous, many authors tend to take σ very tiny, e.g. 3.2 or 8. However, this completely violates the bound $\sigma > 2\sqrt{n}$ required for the hardness reductions to hold [LPR10].

3 Existing Attacks

In January 2015, Albrecht et al. [APS15] aimed at assessing the concrete hardness of LWE and provided an excellent survey of the state-of-art approaches as of this date. These methods apply also to Ring-LWE. We recall here, from a high level perspective, the different families of methods, together with more recent attacks when such exist. The sections of their paper contains extensive details and are therefore mentioned here as reference for the interested reader. Then we briefly present the attacks targeting Ring-LWE specifically and summarized in the recent survey [Pei16].

Bruteforce on s . Attempting an exhaustive search is always possible, yet rarely efficient to solve the problem if the parameter choice is sound. For completeness they express the time complexity of such attack, in general [APS15, Sect. 5.1] and in the case where $\|\mathbf{s}\| \leq 1$ [APS15, Sect. 6.1]. There are better performing methods, as follow.

Distinction. In [APS15, Sect. 5.3], they present a way to distinguish LWE samples from uniformly random samples, as stated by the Decision-LWE problem. However this does not directly recover the secret.

Arora-Ge Attack [AG11]. A purely algebraic attack [APS15, Sects. 5.6 and 6.5], this one consists in constructing from LWE samples a polynomial whose root is the secret. It is particularly efficient in cases where $\|\mathbf{e}\|$ is very small, for example binary errors. For this case however, we are outside of the domain of validity of the hardness reduction from [Reg05] which requires $\sigma > 2\sqrt{n}$.

Blum-Kalai-Wasserman [BKW03]. This combinatorial method works like the Gauss elimination procedure [APS15, Sects. 5.2 and 6.4]. It requests many samples, searches collisions between parts of these samples and gradually creates linear combinations where more and more components equal 0. At the end, the combination allows to deduce the secret. This method has been improved a lot since its original presentation. To date, the best is from Guo et al. [GJS15]. Due to the number of samples required, we do not consider this attack further. Indeed we prefer to assume availability of only the public key and see what can be done.

Decoding. In LWE samples, Gaussian errors can be considered bounded, with high probability, by a sufficiently large bound. An approach is then to use decoding algorithms to remove the error. This addresses the LWE problem as a *Bounded-Distance Decoding (BDD)* instance in a specially crafted lattice. The BDD problem asks to recover the lattice point closest to a given one (not in the lattice), with the promise that the target point is close to the lattice [APS15, Sect. 5.4].

Decoding with Embedding. Another technique to cope with the difficulties of direct decoding is to embed the lattice into another. The latter has higher dimension and more properties that enable some optimizations. Two embeddings are presented: one in the general case [APS15, Sect. 5.5] by Kannan [Kan87] and one in the specific case of $\|\mathbf{s}\| \leq 1$ [APS15, Sect. 6.3] from [BG14].

Using the Ring Structure. Several recent works [EHL14, ELOS15, CLS15, CIV16] use the underlying ring structure to attack Ring-LWE. Peikert describe a unified framework that encompasses all these attacks and sort them into two classes: reduction to errorless LWE and reduction modulo an ideal for which decision-Ring-LWE is tractable [Pei16]. Consequently, we know some rings to be vulnerable and others to be immune.

4 Our New Attack

Out of these surveys, we conclude to keep the decoding attacks. Algebraic attacks are not usable within the bounds of hardness reductions [Reg05, LPR10] and combinatorial attacks require too many samples, we let them aside. In the case of [FV12], the secret is small, $\|\mathbf{s}\| \leq 1$, so when it comes to embedding, the solution [BG14] is best.

4.1 Bai-Galbraith Embedding Improved

We detail here a slightly improved version of this embedding.

Let $(\mathbf{A}, \mathbf{b} = \mathbf{A}^T \mathbf{s} + \mathbf{e} \bmod q)$ be a LWE instance for some n, q and σ . As we derive the matrix \mathbf{A} from a single Ring-LWE sample (the key), we fix $m = n$, $\mathbf{A} \in \mathbb{Z}^{n \times n}$. Write $\mathbf{A}' = (\mathbf{A} | \mathbf{I}_n)$, being a $n \times 2n$ matrix. We have the following equality

$$\mathbf{b} = \mathbf{A}' \begin{pmatrix} \mathbf{s} \\ \mathbf{e} \end{pmatrix} \bmod q$$

where $\begin{pmatrix} \mathbf{s} \\ \mathbf{e} \end{pmatrix}$ is a short vector with respect to q .

Clearly for $\mathbf{w} = \begin{pmatrix} \mathbf{0} \\ \mathbf{b} \end{pmatrix}$ we have $\mathbf{A}' \mathbf{w} = \mathbf{b}$. We then try to approximate this vector \mathbf{w} by a point \mathbf{v}_0 from the lattice \mathcal{L}' defined by

$$\mathcal{L}' = \{ \mathbf{v} \in \mathbb{Z}^{2n} : \mathbf{A}' \mathbf{v} = \mathbf{0} \bmod q \}$$

For such a vector \mathbf{v}_0 , we have $\mathbf{v}_0 \approx \mathbf{w}$ so $\mathbf{w} - \mathbf{v}_0$ is a short vector (with respect to q) and

$$\begin{aligned} \mathbf{A}'(\mathbf{w} - \mathbf{v}_0) \bmod q &= \mathbf{A}'\mathbf{w} - \mathbf{A}'\mathbf{v}_0 \bmod q \\ &= \mathbf{b} - 0 \bmod q \\ &= \mathbf{b} \end{aligned}$$

Consequently we have $\mathbf{w} - \mathbf{v}_0 = \begin{pmatrix} \mathbf{s} \\ \mathbf{e} \end{pmatrix}$. Hence, finding \mathbf{v}_0 in \mathcal{L}' allows us to recover \mathbf{s} (and also \mathbf{e}). We can use BDD algorithms for this, if we have a basis \mathbf{B} for \mathcal{L}' . We obtain it by embedding as follows:

$$\mathbf{B}^T = \begin{pmatrix} \mathbf{I}_n & \mathbf{0} \\ -\mathbf{A} & q\mathbf{I}_n \end{pmatrix} \in \mathbb{Z}^{2n \times 2n}$$

It can be verified that the columns of this matrix are linearly independent and each of them satisfies the definition of \mathcal{L}' , so \mathbf{B} is a basis of \mathcal{L}' .

This slightly differs from what Bai and Galbraith did. They introduced the matrix

$$\mathbf{M} = \begin{pmatrix} \mathbf{I}_n & \Big| & \\ -\mathbf{A} & \Big| & q\mathbf{I}_{2n} \end{pmatrix} \in \mathbb{Z}^{2n \times 3n}$$

and compute its column Hermite Normal Form to end up with a full rank matrix generating the lattice. Our trick avoids this HNF computation and yields a basis for \mathcal{L}' more efficiently².

After the embedding part, we leave untouched the rescaling operation from [BG14]. Since $\|\mathbf{s}\| \leq 1$ whereas the standard deviation for \mathbf{e} is σ (which is not small as discussed above), the components of the difference $\mathbf{w} - \mathbf{v}_0$ are unbalanced. Therefore we multiply the n first components of the basis vectors. It makes the lattice reduction easier, the difference more balanced and still short (with respect to q). Consequently the next stage of the attack works in the lattice whose basis is

$$\mathbf{B}^T = \begin{pmatrix} \sigma\mathbf{I}_n & \mathbf{0} \\ -\mathbf{A} & q\mathbf{I}_n \end{pmatrix} \in \mathbb{Z}^{2n \times 2n}$$

4.2 Lattice Reduction

Once the embedding and rescaling are done, the next step is to find the closest point \mathbf{v}_0 to \mathbf{w} in \mathcal{L}' . The existing algorithms for BDD require a somehow reduced basis of the lattice, otherwise they are completely impracticable. Since our basis \mathbf{B} in its present form verifies none of the criteria for reduction given at the beginning Sect. 2.2, we need to reduce it. We detail here the different options and our strategy.

Lattice reduction has been studied for long since it is very useful in cryptanalysis, integer programming, to cite only a few. The first and celebrated work of

² Computing the HNF of a matrix is not an intense computation, but can be avoided. See [SL96] for a complexity analysis.

Lenstra et al. [LLL82] presented a polynomial time algorithm for lattice reduction, whose output quality suffices for many applications. Since then, several blockwise algorithms have been introduced, the best performing are BKZ [CN11] and Slide Reduction [GN08a]. Both achieve better output quality than LLL but have worse time complexities. Since BDD becomes quicker with a better reduced basis, we need to find a trade-off between the time spent on the reduction step and on the BDD step, so that the overall attack time is minimized. Therefore, an attacker might attempt to use LLL first and if the BDD step takes too much time, try a stronger reduction algorithm. Same discussion apply when the algorithm is fixed, the attacker can tweak its parameters and achieve different output qualities at different computational costs.

The attacker may also look for specialised lattice reduction algorithm since the previous ones work for any lattices and do not take advantage of any structure in them. To the best of our knowledge, the only specialized algorithm that fits our case is the variant from Gama et al. [GHGN06]. The lattices it reduces are *symplectic*. Denoting $J_{2n} = \begin{pmatrix} \mathbf{0} & \mathbf{I}_n \\ -\mathbf{I}_n & \mathbf{0} \end{pmatrix} \in \mathbb{Z}^{2n \times 2n}$, a lattice with basis \mathbf{B} is symplectic if and only if $\mathbf{B}^T J_{2n} \mathbf{B} = J_{2n}$. The bases of our case are indeed symplectic, as can be verified. So this variant of LLL would be interesting to try. The authors reported a speed-up factor of nearly 10 when compared to reference implementation of classical LLL. However the code is not available, so we were not able to use it in our tests.

In the next section we present the details of our experiments on algorithms and parameters. Our conclusion is that for our case, the BDD step is successful even when we only perform an LLL reduction with quite weak parameters. This fact is absolutely no general conclusion. Here we are in a specific case, our basis \mathbf{B} has several properties (integer, upper triangular, blockwise upper triangular with scaled identity on the diagonal, etc.) and so do our parameters n , σ and q .

4.3 Pruned Enumeration for BDD

Finally, to find \mathbf{v}_0 in \mathcal{L}' close to \mathbf{w} with the reduced basis, we use Liu-Nguyen pruned enumeration [LN13]. To our knowledge pruned enumeration is the best performing algorithms in practice, see [HPS11] for a description of the other candidates.

This method of pruned enumeration is an adaptation to BDD of the extreme pruning technique introduced by Gama et al. [GNR10]. This algorithm enumerates lattice points that are close enough to the target, with a bound provided as a *pruning function*. Since the enumeration works by adding one component at a time to the current partial solution, we may define different bounds for each positions. Unlike previous work [Bab86, LP11], Liu-Nguyen pruned enumeration evaluates the heuristic on the *projections* of the current candidate and not on its *components*.

This concludes the algorithmic description of our attack. We turn now to implementation aspects and the results we get.

5 Implementation and Results

In this section we present the implementation details of our new attack, described in the previous section. Since we were able to use it successfully against many [FV12] keys, we thoroughly analyse its performance and give out some conclusions.

5.1 Implementation Details

In order to implement our attack, we use the following existing codes:

- Lepoint’s implementation [Lep14] of [FV12],
- Stehlé’s `fp111` [dt16] for lattice reduction algorithms,
- Shoup’s NTL [Sho15] for additional lattice operations.

Then, we implemented the embedding/rescaling, the pruned enumeration algorithm from Liu-Nguyen (following their pseudo-code [LN13]) and created the glue between the different libraries to lead the attack from beginning to end.

Lepoint’s implementation. In [LN14], in order to compare [FV12] and YASHE, Lepoint implemented the code needed to use both schemes and perform homomorphic operations. We use his constructor method with light modification. It allows us to create a public key $\mathbf{pk} = ([-(\mathbf{a} \cdot \mathbf{s} + \mathbf{e})]_q, \mathbf{a})$, given n , σ and q . The attack aims to recover the secret when given the `FVKey` object as input.

From [FV12] to R -LWE lattice. From the \mathbf{a} in the public key, we construct a lattice basis \mathbf{A} so that $[\mathbf{a} \cdot \mathbf{s} + \mathbf{e}]_q$ (polynomial operation) equals $\mathbf{A}\mathbf{s} + \mathbf{e} \bmod q$ (matrix/vector operation). For instance, with $R = \mathbb{Z}[x]/(x^n + 1)$ we have:

$$\mathbf{A}^T = \begin{pmatrix} a_1 & a_2 & a_3 & \cdots & a_n \\ -a_n & a_1 & a_2 & \cdots & a_{n-1} \\ -a_{n-1} & -a_n & a_1 & \cdots & a_{n-2} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ -a_2 & -a_3 & -a_4 & \cdots & a_1 \end{pmatrix}$$

Embedding and rescaling. Then we embed this matrix \mathbf{A} into \mathbf{B} following Sect. 4.1. Rescaling is also done here. \mathbf{B} is ready for the next stage.

Lattice reduction. For this step we rely on the different routines from `fp111` [dt16]. Both LLL and BKZ algorithms are available and can be tweaked conveniently. We detail below the settings we experimented.

BDD enumeration. Finally, we implemented the pruned enumeration algorithm from Liu-Nguyen to solve BDD in the reduced lattice \mathcal{L}' . Later, we discovered an undocumented implementation of enumeration in `fp111`, that we included in our experiments.

5.2 Attack Settings – Early Benchmarks

Before launching a full-scale attack, we run it on small cases. It allows us to explore the choices we have, for algorithms and their parameters, for the different stages.

Lattice Reduction Tweaking. State-of-the-art results [APS15] tell us that strong reductions like BKZ is needed to make enumeration possible. Yet after a few toy examples with BKZ as reduction algorithm, we realized that most of the attack time was spent in the lattice reduction stage. So we tried different settings for BKZ and LLL on cases where $n \leq 100$.

Beginning with BKZ, we could confirm [CN11], an intermediate blocksize $\beta = 20$ leads to moderate running times. Whereas smaller values like 5 or greater like 40 lead to prohibitive running times. Even with such blocksize, the attack time is not balanced between reduction and enumeration, so we rapidly shifted to study LLL.

LLL algorithm is governed by two parameters δ and η . We recall the size reduction condition and the Lovász condition from Sect. 2.2:

$$\begin{aligned} \forall i < j, \quad |(\mathbf{b}_j | \mathbf{b}_i^*)| &\leq \eta \cdot \|\mathbf{b}_i^*\|^2 && \text{with } 1/2 \leq \eta \leq \sqrt{\delta} \\ \forall i, \quad \delta \|\mathbf{b}_i^*\|^2 &\leq \left(\|\mathbf{b}_{i+1}^*\|^2 + \frac{(\mathbf{b}_{i+1} | \mathbf{b}_i^*)^2}{\|\mathbf{b}_i^*\|^2} \right) && \text{with } 1/4 \leq \delta \leq 1 \end{aligned}$$

The default values in `fp111` are $(\delta, \eta) = (0.99, 0.51)$.

Optimising δ . First, we tried to decrease δ to loosen the Lovász condition, however with $(0.75, 0.51)$ we gained only limited speedups and only for $n \geq 100$, and it had dramatic effect on the observed success rate of the overall attack. The enumeration stage fails due to a lattice basis of insufficient quality.

Optimising η . Then we restored δ and experimented with an increased $\eta \in [0.60, 0.98]$, to ease the size-reduction condition.

As we can see in Fig. 1, loosening the size reduction condition decreases the running time of LLL. We observe a greater gain (in time) between 0.60 and 0.75 than between 0.75 and 0.95. In addition, with $\eta = 0.95$ we get a less successful attack, whereas for η equal to 0.51 or 0.75 the observed success rates are similar.

Observed root Hermite factor. We report in Table 1 the quality we get with reduction parameters $\delta = 0.99$ and $\eta = 0.71$. As a reminder, the best proved γ for LLL is $\gamma = 1.0754$ and Gama and Nguyen observed $\gamma = 1.0219$ for random lattices of similar dimensions [GN08b]. In our case, we observe that LLL reduces lattices to a very good quality, even with weak reduction parameters δ and η .

Enumeration Behavior. For the enumeration step, we use the technique of pruned enumeration from [LN13]. Its only setting is the pruning function (R_1^2, \dots, R_{2n}^2) . As we know precisely the expected distance between \mathbf{w} and the

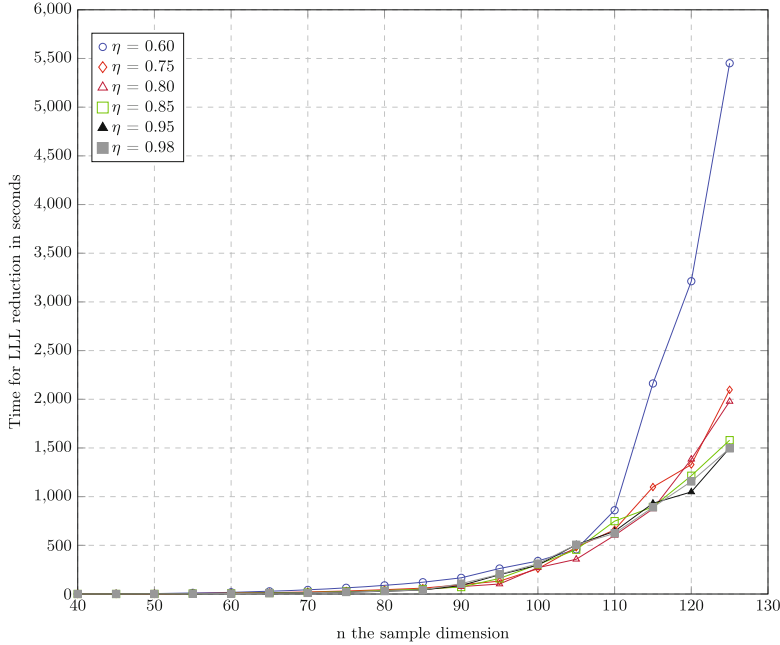


Fig. 1. Execution time in seconds in term of n for different η

Table 1. Observed reduction quality for different lattice parameters

γ	$\log_2 q = 47$	$\log_2 q = 95$
$n = 50$	1.00238	1.00129
$n = 200$	1.00300	1.00147

lattice point \mathbf{v}_0 , $\mathbf{w} - \mathbf{v}_0 = (\sigma_{\mathbf{e}}^s)$, we can set the bound R_{2n} to the expected norm of $(\sigma_{\mathbf{e}}^s)$. For a Gaussian error distribution we have with high probability $\|\mathbf{e}\|^2 \leq n \times (3\sigma^2)$, so $R_{2n}^2 = n\sigma^2 + n(3\sigma)^2 = 10n\sigma^2$.

In their paper, Gama et al. [GNR10] introduce three pruning functions: linear, step and piecewise linear, whose they study the resulting time complexity and success probability. They also mention another pruning function obtained by numerical optimization. In our case, we start with a linear pruning function defined by $R_k^2 = (k/2n)R_{2n}^2$.

One of the surprising finding of our experiments is that the enumeration terminates with the first candidate solutions, in nearly every cases. Recall that enumeration for BDD begins with getting a vector *quite close* to the target, and then, enumerates around, towards the one which minimizes the distance. So this first enumerated vector is equal to that outputed by Babai's Nearest Plane algorithm [Bab86]. It so happens that this solution falls below our pruning bound for most of our test cases, and leads to a successful key recovery. Both our implementation of BDD and `fp111`'s `ClosestVector` routine show this behavior.

A qualitative explanation for this unexpected costless enumeration is that the error norm (depending on σ) is very small compared to the Gram-Schmidt vectors, whose norms highly depend on q . Consequently, the simple rounding from Babai algorithm seems enough to get directly to the closest vector, component by component. This result stands also for other FHE schemes, in all which we have the same properties for σ and q .

5.3 Attacking Higher Dimensions

With these first conclusions we decided to go for a greater range of values for n and q . In terms of possible circuit depth evaluation, this range covers up to $L = 13$. We set the attack to an LLL reduction with $\delta = 0.99$ and $\eta \in \{0.51, 0.61, 0.71\}$ and an enumeration limited to Nearest Plane, and launched it for a few weeks.

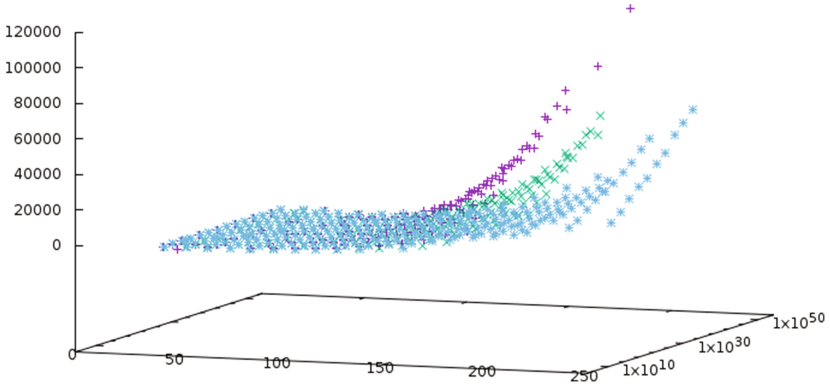


Fig. 2. Execution time in seconds in terms of n and q for different η

With n up to 250 we see in Fig. 2 that with $\eta = 0.71$ the attack takes fairly less time than 0.51, roughly 5 times less and still finishes successfully. This motivated us to keep only the version $\eta = 0.71$ and continue to higher dimensions.

In the end we were able to successfully break [FV12] keys with $n = 250$ and $\log q = 46.8$ in 10 h, or $n = 320$ and $\log q = 68.7$ in little less than 28 h, see Fig. 3. Our result compares favorably with the work of Laine and Lauter [LL15], who were able to recover a key in dimension 350 in 3.5 days, yet with a less generic attack working for q very large (2^{52}) and σ very small (3.2).

Such parameters are not considered secure, even prior to our work. One would take much greater n for instance. Yet they serve as good examples to understand the performance of our special-purpose attack. When composed with an LLL reduction and an enumeration with Nearest Plane, our proposal requires only a polynomial number of operations (in the parameters of the lattice) to complete. Moreover, the estimator from [APS15], which takes into account only generic attacks, including the latest [Alb17], predicts one month of computation to break such key.

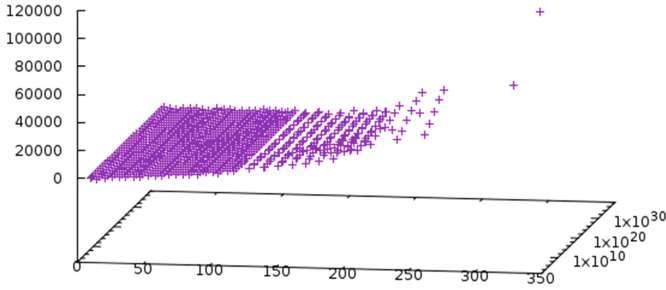


Fig. 3. Execution time in seconds in terms of n and q for $\eta = 0.71$

6 Conclusion

In this work, we aimed at assessing the practical security of FHE schemes based on Ring-LWE. To us, this is a topic that needs more focus, since the requirements of correctness in FHE lead to very special shape of Ring-LWE parameters. After reviewing state-of-the-art attacks, we presented a new special-purpose attack for the case at hands. Our experiments show that such attack has unexpected performance: with only a polynomial number of steps, it successfully breaks keys with parameters beyond toy sizes.

Our main results, on lattice reduction and enumeration in FHE cases, confirm our opinion that attacker may have unexpected advantage in special situations. Our result does not contradict the security reductions of Ring-LWE, but when picking practical parameters, it is really important to consider such results. The discussion about sizing parameters to guarantee a security level objective is far from being closed. We hope we raise interest for this kind of work which is of great importance to move forward with FHE implementation.

Acknowledgements. This work has been supported by the Chair of Naval Cyber Defense, funded by Ecole Navale, IMT-Atlantique, Thales and Naval Group.

References

- [AG11] Arora, S., Ge, R.: New algorithms for learning in presence of errors. In: Aceto, L., Henzinger, M., Sgall, J. (eds.) ICALP 2011. LNCS, vol. 6755, pp. 403–415. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-22006-7_34
- [Alb17] Albrecht, M.R.: On dual lattice attacks against small-secret LWE and parameter choices in HELib and SEAL. In: Coron, J.-S., Nielsen, J.B. (eds.) EUROCRYPT 2017. LNCS, vol. 10211, pp. 103–129. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-56614-6_4
- [AMGH10] Melchor, C.A., Gaborit, P., Herranz, J.: Additively homomorphic encryption with d -operand multiplications. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 138–154. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-14623-7_8

- [APS15] Albrecht, M.R., Player, R., Scott, S.: On the concrete hardness of Learning with Errors. Cryptology ePrint Archive, Report 2015/046 (2015)
- [Bab86] Babai, L.: On lovász' lattice reduction and the nearest lattice point problem. *Combinatorica* **6**(1), 1–13 (1986)
- [BG14] Bai, S., Galbraith, S.D.: Lattice decoding attacks on binary LWE. In: Susilo, W., Mu, Y. (eds.) ACISP 2014. LNCS, vol. 8544, pp. 322–337. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-08344-5_21
- [BGV12] Brakerski, Z., Gentry, C., Vaikuntanathan, V.: (Leveled) fully homomorphic encryption without bootstrapping. In: Proceedings of the 3rd Innovations in Theoretical Computer Science Conference, pp. 309–325. ACM (2012)
- [BKW03] Blum, A., Kalai, A., Wasserman, H.: Noise-tolerant learning, the parity problem, and the statistical query model. *J. ACM (JACM)* **50**(4), 506–519 (2003)
- [CGGI16] Chillotti, I., Gama, N., Georgieva, M., Izabachène, M.: Faster fully homomorphic encryption: bootstrapping in less than 0.1 Seconds. In: Cheon, J.H., Takagi, T. (eds.) ASIACRYPT 2016. LNCS, vol. 10031, pp. 3–33. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53887-6_1
- [CIV16] Castryck, W., Iliashenko, I., Vercauteren, F.: Provably weak instances of ring-LWE revisited. In: Fischlin, M., Coron, J.-S. (eds.) EUROCRYPT 2016. LNCS, vol. 9665, pp. 147–167. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49890-3_6
- [CLS15] Chen, H., Lauter, K., Stange, K.E.: Attacks on Search RLWE. Cryptology ePrint Archive, Report 2015/971 (2015). <http://eprint.iacr.org/>
- [CN11] Chen, Y., Nguyen, P.Q.: BKZ 2.0: better lattice security estimates. In: Lee, D.H., Wang, X. (eds.) ASIACRYPT 2011. LNCS, vol. 7073, pp. 1–20. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-25385-0_1
- [Cry] CryptoExperts. FV-NFLlib. <https://github.com/CryptoExperts/FV-NFLlib>
- [DGBL+15] Dowlin, N., Gilad-Bachrach, R., Laine, K., Lauter, K., Naehrig, M., Wernsing, J.: Manual for using homomorphic encryption for bioinformatics. Technical report MSR-TR-2015-87, November 2015
- [DM15] Ducas, L., Micciancio, D.: FHEW: bootstrapping homomorphic encryption in less than a second. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015. LNCS, vol. 9056, pp. 617–640. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46800-5_24
- [DS15] Doröz, Y., Sunar, B.: Flattening NTRU for Evaluation Key Free Homomorphic Encryption (2015). <http://eprint.iacr.org/>
- [dt16] The FPLLL development team. fpLLL, a lattice reduction library (2016). <https://github.com/fplll/fplll>
- [EHL14] Eisenträger, K., Hallgren, S., Lauter, K.: Weak instances of PLWE. In: Joux, A., Youssef, A. (eds.) SAC 2014. LNCS, vol. 8781, pp. 183–194. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-13051-4_11
- [ELOS15] Elias, Y., Lauter, K.E., Ozman, E., Stange, K.E.: Provably weak instances of ring-LWE. In: Gennaro, R., Robshaw, M. (eds.) CRYPTO 2015. LNCS, vol. 9215, pp. 63–92. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-47989-6_4
- [FV12] Fan, J., Vercauteren, F.: Somewhat practical fully homomorphic encryption. IACR Cryptology ePrint Archive 2012:144 (2012)

- [Gen09] Gentry, C.: Fully homomorphic encryption using ideal lattices. In: STOC, vol. 9, pp. 169–178 (2009)
- [GHGN06] Gama, N., Howgrave-Graham, N., Nguyen, P.Q.: Symplectic lattice reduction and NTRU. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 233–253. Springer, Heidelberg (2006). https://doi.org/10.1007/11761679_15
- [GJS15] Guo, Q., Johansson, T., Stankovski, P.: Coded-BKW: solving LWE using lattice codes. In: Gennaro, R., Robshaw, M. (eds.) CRYPTO 2015. LNCS, vol. 9215, pp. 23–42. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-47989-6_2
- [GN08a] Gama, N., Nguyen, P.Q.: Finding short lattice vectors within mordell’s inequality. In: Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing, pp. 207–216. ACM (2008)
- [GN08b] Gama, N., Nguyen, P.Q.: Predicting lattice reduction. In: Smart, N. (ed.) EUROCRYPT 2008. LNCS, vol. 4965, pp. 31–51. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-78967-3_3
- [GNR10] Gama, N., Nguyen, P.Q., Regev, O.: Lattice enumeration using extreme pruning. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 257–278. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-13190-5_13
- [GSW13] Gentry, C., Sahai, A., Waters, B.: Homomorphic encryption from learning with errors: conceptually-simpler, asymptotically-faster, attribute-based. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013. LNCS, vol. 8042, pp. 75–92. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40041-4_5
- [Hal] Halevi, S.: Helib. <https://github.com/sha1h/HElib>
- [HPS11] Hanrot, G., Pujol, X., Stehlé, D.: Algorithms for the shortest and closest lattice vector problems. In: Chee, Y.M., Guo, Z., Ling, S., Shao, F., Tang, Y., Wang, H., Xing, C. (eds.) IWCC 2011. LNCS, vol. 6639, pp. 159–190. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-20901-7_10
- [Kan87] Kannan, R.: Minkowski’s convex body theorem and integer programming. Math. Oper. Res. **12**(3), 415–440 (1987)
- [KGV15] Khedr, A., Gulak, G., Vaikuntanathan, V.: SHIELD: scalable homomorphic implementation of encrypted data-classifiers. IEEE Trans. Comput. **65**(9), 2848–2858 (2015)
- [Lep14] Lepoint, T.: A proof-of-concept implementation of the homomorphic evaluation of SIMON using FV and YASHE leveled homomorphic cryptosystems (2014). Accessed 18 Aug 2015
- [LL15] Laine, K., Lauter, K.: Key Recovery for LWE in Polynomial Time. Cryptology ePrint Archive, Report 2015/176 (2015)
- [LLL82] Lenstra, A.K., Lenstra, H.W., Lovász, L.: Factoring polynomials with rational coefficients. Math. Ann. **261**(4), 515–534 (1982)
- [LN13] Liu, M., Nguyen, P.Q.: Solving BDD by enumeration: an update. In: Dawson, E. (ed.) CT-RSA 2013. LNCS, vol. 7779, pp. 293–309. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-36095-4_19
- [LN14] Lepoint, T., Naehrig, M.: A comparison of the homomorphic encryption schemes FV and YASHE. In: Pointcheval, D., Vergnaud, D. (eds.) AFRICACRYPT 2014. LNCS, vol. 8469, pp. 318–335. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-06734-6_20

- [LP11] Lindner, R., Peikert, C.: Better key sizes (and Attacks) for LWE-based encryption. In: Kiayias, A. (ed.) CT-RSA 2011. LNCS, vol. 6558, pp. 319–339. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-19074-2_21
- [LPR10] Lyubashevsky, V., Peikert, C., Regev, O.: On ideal lattices and learning with errors over rings. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 1–23. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-13190-5_1
- [MW15] Micciancio, D., Walter, M.: Practical, predictable lattice basis reduction. Cryptology ePrint Archive, Report 2015/1123 (2015). <http://eprint.iacr.org/>
- [NS05] Nguên, P.Q., Stehlé, D.: Floating-point LLL revisited. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 215–233. Springer, Heidelberg (2005). https://doi.org/10.1007/11426639_13
- [Pei15] Peikert, C.: A decade of lattice cryptography. Cryptology ePrint Archive, Report 2015/939 (2015). <http://eprint.iacr.org/>
- [Pei16] Peikert, C.: How (Not) to Instantiate Ring-LWE. Cryptology ePrint Archive, Report 2016/351 (2016)
- [Reg05] 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, pp. 84–93. ACM (2005)
- [SE94] Schnorr, C.-P., Euchner, M.: Lattice basis reduction: improved practical algorithms and solving subset sum problems. Math. Program. **66**(1–3), 181–199 (1994)
- [Sho15] Shoup, V.: NTL - A Library for doing Number Theory (2015). Accessed 18 Aug 2015
- [SL96] Storjohann, A., Labahn, G.: Asymptotically fast computation of hermite normal forms of integer matrices. In: Proceedings of the 1996 International Symposium on Symbolic and Algebraic Computation, pp 259–266. ACM (1996)
- [vDGHV10] van Dijk, M., Gentry, C., Halevi, S., Vaikuntanathan, V.: Fully homomorphic encryption over the integers. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 24–43. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-13190-5_2
- [vdPS13] van de Pol, J., Smart, N.P.: Estimating key sizes for high dimensional lattice-based systems. In: Stam, M. (ed.) IMACC 2013. LNCS, vol. 8308, pp. 290–303. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-45239-0_17

Progress in Cryptology – INDOCRYPT 2017

18th International Conference on Cryptology in India,
Chennai, India, December 10–13, 2017, Proceedings

Patra, A.; Smart, N.P. (Eds.)

2017, XII, 403 p. 68 illus., Softcover

ISBN: 978-3-319-71666-4