

On Dual Lattice Attacks Against Small-Secret LWE and Parameter Choices in HELib and SEAL

Martin R. Albrecht^(✉)

Information Security Group, Royal Holloway,
University of London, Egham, Surrey TW20 0EX, UK
`martin.albrecht@royalholloway.ac.uk`

Abstract. We present novel variants of the dual-lattice attack against LWE in the presence of an unusually short secret. These variants are informed by recent progress in BKW-style algorithms for solving LWE. Applying them to parameter sets suggested by the homomorphic encryption libraries HELib and SEAL yields revised security estimates. Our techniques scale the exponent of the dual-lattice attack by a factor of $(2L)/(2L+1)$ when $\log q = \Theta(L \log n)$, when the secret has constant hamming weight h and where L is the maximum depth of supported circuits. They also allow to half the dimension of the lattice under consideration at a multiplicative cost of 2^h operations. Moreover, our techniques yield revised concrete security estimates. For example, both libraries promise 80 bits of security for LWE instances with $n = 1024$ and $\log_2 q \approx 47$, while the techniques described in this work lead to estimated costs of 68 bits (SEAL) and 62 bits (HELib).

1 Introduction

Learning with Errors (LWE), defined in Definition 1, has proven to be a rich source of cryptographic constructions, from public-key encryption and Diffie-Hellman-style key exchange (cf. [Reg09, Pei09, LPR10, DXL12, BCNS15, ADPS16, BCD+16]) to fully homomorphic encryption (cf. [BV11, BGV12, Bra12, FV12, GSW13, CS15]).

Definition 1 (LWE [Reg09]). Let n, q be positive integers, χ be a probability distribution on \mathbb{Z} and \mathbf{s} be a secret vector in \mathbb{Z}_q^n . We denote by $L_{\mathbf{s}, \chi, q}$ 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 pairs $(\mathbf{a}, c) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$ are sampled according to $L_{\mathbf{s}, \chi, q}$ or the uniform distribution on $\mathbb{Z}_q^n \times \mathbb{Z}_q$.

Search-LWE is the problem of recovering \mathbf{s} from $(\mathbf{a}, c) = (\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle + e) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$ sampled according to $L_{\mathbf{s}, \chi, q}$.

This research was supported by EPSRC grants EP/L018543/1 “Multilinear Maps in Cryptography” and EP/P009417/1 “Bit Security of Learning with Errors for Post-Quantum Cryptography and Fully Homomorphic Encryption”.

We may write LWE instances in matrix form (\mathbf{A}, \mathbf{c}) , where rows correspond to samples (\mathbf{a}_i, c_i) . In many instantiations, χ is a discrete Gaussian distribution with standard deviation $\alpha q/\sqrt{2\pi}$. Though, in this work, like in many works on cryptanalysis of LWE, the details of the error distribution do not matter as long as we can bound the size of the error under additions.

The bit-security of concrete LWE instances is a prominent area of current cryptographic research, in particular in light of standardisation initiatives for LWE-based schemes and LWE-based (somewhat) homomorphic encryption being proposed for applications such as computation with medical data [KL15]. See [APS15] for a relatively recent survey of known (classical) attacks.

Applications such as [KL15] are enabled by progress in homomorphic encryption in recent years. The two most well-known homomorphic encryption libraries are HELib and SEAL. HELib [GHS12a, HS14] implements BGV [BGV12]. SEAL v2.0 [LP16] implements FV [Bra12, FV12]. Both schemes fundamentally rely on the security of LWE.

However, results on the expected cost of solving generic LWE instances do not directly translate to LWE instances as used in fully homomorphic encryption (FHE). Firstly, because these instances are typically related to the Ring-LWE assumption [LPR10, LPR13] instead of plain LWE. Secondly, because these instances are typically *small-secret* instances. In particular, they typically sample the secret \mathbf{s} from some distribution \mathcal{B} as defined below. We call such instances \mathcal{B} -secret LWE instances.

Definition 2. *Let n, q be positive integers. We call*

\mathcal{B} *any distribution on \mathbb{Z}_q^n where each component ≤ 1 in absolute value, i.e. $\|\mathbf{s}_{(i)}\| \leq 1$ for $\mathbf{s} \leftarrow_{\mathcal{B}}$.*

\mathcal{B}^+ *the distribution on \mathbb{Z}_q^n where each component is independently sampled uniformly at random from $\{0, 1\}$.*

\mathcal{B}^- *the distribution on \mathbb{Z}_q^n where each component is independently sampled uniformly at random from $\{-1, 0, 1\}$.*

\mathcal{B}_h^+ *the distribution on \mathbb{Z}_q^n where components are sampled independently uniformly at random from $\{0, 1\}$ with the additional guarantee that at most h components are non-zero.*

\mathcal{B}_h^- *the distribution on \mathbb{Z}_q^n where components are sampled independently uniformly at random from $\{-1, 0, 1\}$ with the additional guarantee that at most h components are non-zero.*

Remark 1. In [BLP+13], instances with $\mathbf{s} \leftarrow_{\mathcal{B}^+}$ are referred to as binary-secret; \mathcal{B}^+ is used in [FV12]; \mathcal{B}^- is used in Microsoft’s SEAL v2.0 library¹ and [LN14]; \mathcal{B}_{64}^- is the default choice in HELib, cf. [GHS12b, Appendix C.1.1] and [HS14].

It is an open question how much easier, if any, \mathcal{B} -secret LWE instances are compared to regular LWE instances. On the one hand, designers of FHE schemes

¹ cf. `KeyGenerator::set_poly_coeffs_zero_one_negone()` at <https://sealcrypto.codeplex.com/SourceControl/latest#SEAL/keygenerator.h>.

typically ignore this issue [GHS12a, LN14, CS16]. This could be considered as somewhat justified by a reduction from [ACPS09] showing that an LWE instance with an arbitrary secret can be transformed into an instance with a secret following the noise distribution in polynomial time and at the loss of n samples. Hence, such instances are not easier than instances with a uniformly random secret, assuming sufficiently many samples are available. As a consequence, LWE with a secret following the noise distribution is considered to be in *normal form*. Given that the noise in homomorphic encryption libraries is also typically rather small—SEAL and HELib use standard deviation $\sigma \approx 3.2$ —the distribution \mathcal{B}^- gives rise to LWE instances which could be considered relatively close to normal-form LWE instances. However, considering the actual distributions, not just the standard deviations, it is known that LWE with *error distribution* \mathcal{B} is insecure once sufficiently many samples are available [AG11, ACFP14, KF15].

On the other hand, the best, known reduction from regular LWE to \mathcal{B}^+ -secret LWE has an expansion factor of $\log q$ in the dimension. That is, [BLP+13] gives a reduction from regular LWE in dimension n to LWE with $\mathbf{s} \leftarrow_{\S} \mathcal{B}^+$ in dimension $n \log q$.

In contrast, even for noise with width $\approx \sqrt{n}$ and $\mathbf{s} \leftarrow_{\S} \mathcal{B}^-$ the best known lattice attacks suggest an expansion factor of at most $\log \log n$ [BG14], if at all. Overall, known algorithms do not perform significantly better for \mathcal{B} -secret LWE instances, perhaps reinforcing our confidence in the common approach of simply ignoring the special form of the secret.

One family of algorithms has recently seen considerable progress with regards to \mathcal{B} -secret instances: combinatorial algorithms. Already in [Reg09] it was observed that the BKW algorithm, originally proposed for LPN by Blum, Kalai and Wasserman [BKW00], leads to an algorithm in $2^{\Theta(n)}$ time and space for solving LWE. The algorithm proceeds by splitting the components of the vectors \mathbf{a}_i into blocks of k components. Then, it searches for collisions in the first block in an “elimination table” holding entries for (possibly) all q^k different values for that block. This table is constructed by sampling fresh (\mathbf{a}_i, c_i) pairs from the LWE oracle. By subtracting vectors with colliding components in the first block, a vector of dimension $n - k$ is recovered, applying the same subtraction to the corresponding c_i values, produces an error of size $\sqrt{2}\alpha q$. Repeating the process for consecutive blocks reduces the dimension further at the cost of an increase in the noise by a factor $\sqrt{2}$ at each level. This process either continues until all components of \mathbf{a}_i are eliminated or when there are so few components left that exhaustive search can solve the remaining low-dimensional LWE instance.

A first detailed study of this algorithm when applied to LWE was provided in [ACF+15]. Subsequently, improved variants were proposed, for small secret LWE instances via “lazy modulus switching” [AFFP14], via the application of an FFT in the last step of the algorithm [DTV15], via varying the block size k [KF15] and via rephrasing the problem as the coding-theoretic problem of quantisation [GJS15]. In particular, the works [KF15, GJS15] improve the exploitation of a small secret to the point where these techniques improve the cost of solving instances where the secret is as big as the error, i.e. arbitrary

LWE instances. Yet, combinatorial algorithms do not perform well on FHE-style LWE instances because of their large dimension n to accommodate the large modulus q .

1.1 Our Contribution/Outline

We first review parameter choices in HELib and SEAL as well as known algorithms for solving LWE and related problems in Sect. 2.

Then, we reconsider the dual-lattice attack (or “dual attack” in short) which finds short vectors \mathbf{y} such that $\mathbf{y} \cdot \mathbf{A} \equiv 0 \pmod{q}$ using lattice reduction. In particular, we recast this attack as the lattice-reduction analogue of the BKW algorithm and adapt techniques and lessons learned from BKW-style algorithms. Applying these techniques to parameter sets suggested for HELib and SEAL, we arrive at revised concrete and asymptotic security estimates.

First, in Sect. 3, we recall (the first stage of) BKW as a recursive dimension reduction algorithm for LWE instances. Each step transforms an LWE instance in dimension n to an instance in dimension $n - k$ at the cost of an increase in the noise by a factor of $\sqrt{2}$. This smaller instance is then reduced further by applying BKW again or solved using another algorithm for solving LWE; typically some form of exhaustive search once the dimension is small enough. To achieve this dimension reduction, BKW first produces elimination tables and then makes use of these tables to sample possibly many LWE samples in dimension $n - k$ relatively cheaply. We translate this approach to lattice reduction in the low advantage regime: we perform one expensive lattice reduction step followed by many relatively cheap lattice reductions on rerandomised bases. This essentially reduces the overall solving cost by a factor of m , where m is the number of samples required to distinguish a discrete Gaussian distribution with large standard deviation from uniform modulo q . We note that this approach applies to any LWE instance, i.e. does not rely on an unusually short secret and thus gives cause for a moderate revision of many LWE estimates based on the dual-attack in the low advantage regime. It does, however, rely on the heuristic that these cheap lattice reduction steps produce sufficiently short and random vectors. We give evidence that this heuristic holds.

Second, in Sect. 4, we observe that the normal form of the dual attack—finding short vectors \mathbf{y} such that $\mathbf{y} \cdot \mathbf{A} \equiv \mathbf{x} \pmod{q}$ is short—is a natural analogue of “lazy modulus switching” [AFFP14]. Then, to exploit the unusually small secret, we apply lattice scaling as in [BG14]. The scaling factor is somewhat analogous to picking the target modulus in modulus switching resp. picking the (dimension of the) code for quantisation. This technique applies to any \mathcal{B} -secret LWE instance. For \mathcal{B}_h^- -secret instances, it reduces the cost of the dual attack by a factor of $2L/(2L + 1)$ in the exponent when $\log q = \Theta(L \log n)$ for L the supported depth of FHE circuits and when h is a constant.

Third, in Sect. 5, we focus on $\mathbf{s} \leftarrow_{\mathcal{S}} \mathcal{B}_h^\pm$ and adapt the dual attack to find short vectors which produce zero when multiplied with a *subset* of the columns of \mathbf{A} . This, as in BKW, produces a smaller, easier LWE instance which is then solved using another algorithm. In BKW, these smaller instances typically have

very small dimension (say, 10). Here, we consider instances with dimension of several hundreds. This is enabled by exploiting the sparsity of the secret and by relaxing the conditions on the second step: we recover a solution only with a small probability of success. The basic form of this attack does not rely on the size of the non-zero components (only on the sparsity) and reduces the cost of solving an instance in dimension n to the cost of solving an instance in dimension $n/2$ multiplied by 2^h where h is the hamming weight of the secret (other trade-offs between multiplicative cost increase and dimension reduction are possible and typically optimal). We also give an improved variant when the non-zero components are also small.

In Sect. 6, we put everything together to arrive at our final algorithm SILKE, which combines the techniques outlined above; inheriting their properties. We also give revised security estimates for parameter sets suggested for HELib and SEAL in Table 1. Table 1 highlights that the techniques described in this work can, despite being relatively simple, produce significantly revised concrete security estimates for both SEAL and HELib.

Table 1. Costs of dual attacks on HELib and SEAL. Rows “ $\log_2 q$ ” give bit sizes for the maximal modulus for a given n , for SEAL it is taken from [LN14], for HELib it is chosen such that the expected cost is 2^{80} resp. 2^{128} s according to [GHS12a]. The rows “dual” give the log cost (in operations) of the dual attack according to our lattice-reduction estimates without taking any of our improvements into account; The row “SILKE_{small}” gives the log cost of Algorithm 3 with “sparse” set to false; The rows “SILKE_{sparse}” give the log cost of Algorithm 3 with “sparse” set to true. The “sparse” flag toggles whether the approach described in Sect. 5 is enabled or not in Algorithm 3.

n	1024	2048	4096	8192	16384
SEAL 80-bit					
$\log_2 q$	47.5	95.4	192.0	392.1	799.6
dual	83.1	78.2	73.7	71.1	70.6
SILKE _{small}	68.1	69.0	68.2	68.4	68.8
HELlib 80-bit					
$\log_2 q$	47.0	87.0	167.0	326.0	638.0
dual	85.2	85.2	85.3	84.6	85.5
SILKE _{sparse}	61.3	65.0	67.9	70.2	73.1
HELlib 128-bit					
$\log_2 q$	38.0	70.0	134.0	261.0	511.0
dual	110.7	110.1	109.3	108.8	108.9
SILKE _{sparse}	73.2	77.4	81.2	84.0	86.4

Table 2. Logarithms of algorithm costs in operations mod q when applied to example parameters $n = 2048$, $q \approx 2^{63.4}$, $\alpha \approx 2^{-60.4}$ and $\mathbf{s} \leftarrow_{\S} \mathcal{B}_{64}^-$. The row “base line” gives the log cost of attacks according to our lattice-reduction estimates without taking any of our improvements into account.

Strategy	Dual	Decode	Embed
HElib	188.9	—	—
Base line	124.2	116.6	114.5
Sect. 4	101.0	—	—
Sect. 5	97.1	111.0	110.9
Sect. 6	83.9	—	—

2 Preliminaries

Logarithms are base 2 if not stated otherwise. We write vectors in bold, e.g. \mathbf{a} , and matrices in upper-case bold, e.g. \mathbf{A} . By $\mathbf{a}_{(i)}$ we denote the i -th component of \mathbf{a} , i.e. a scalar. In contrast, \mathbf{a}_i is the i -th element of a list of vectors. We write \mathbf{I}_m for the $m \times m$ identity matrix over whichever base ring is implied from context. We write $\mathbf{0}_{m \times n}$ for the $m \times n$ zero matrix. A lattice is a discrete subgroup of \mathbb{R}^n . It can be represented by a basis \mathbf{B} . We write $\Lambda(\mathbf{B})$ for the lattice generated by the rows of the matrix \mathbf{B} , i.e. all integer-linear combinations of the rows of \mathbf{B} . We write $\Lambda_q(\mathbf{B})$ for the q -ary lattice generated by the rows of the matrix \mathbf{B} over \mathbb{Z}_q , i.e. the lattice spanned by the rows \mathbf{B} and multiples of q . We write $\mathbf{A}_{n:m}$ for the rows $n, \dots, m-1$ of \mathbf{A} . If the starting or end point is omitted it is assumed to be 0 or the number of rows respectively, i.e. we follow Python’s slice notation.

2.1 Rolling Example

Throughout, we are going to use Example 1 below to illustrate the behaviour of the techniques described here. See Table 2 for an overview of complexity estimates for solving this set of parameters using the techniques described in this work.

Example 1. The LWE dimension is $n = 2048$, the modulus is $q \approx 2^{63.4}$, the noise parameter is $\alpha \approx 2^{-60.4}$, i.e. we have a standard deviation of $\sigma \approx 3.2$. We have $\mathbf{s} \leftarrow_{\S} \mathcal{B}_{64}^-$, i.e. only $h = 64$ components of the secret are ± 1 , all other components are zero. This set of parameters is inspired by parameter choices in HElib and produced by calling the function `fhe_params(n=2048,L=2)` of the LWE estimator from [APS15].

2.2 Parameter Choices in HElib

HElib [GHS12a, HS14] uses the cost of the dual attack for solving LWE to establish parameters. The dual strategy reduces the problem of distinguishing LWE from uniform to the SIS problem [Ajt96]:

Definition 3 (SIS). Given $q \in \mathbb{Z}$, a matrix \mathbf{A} , and $t < q$; find \mathbf{y} with $0 < \|\mathbf{y}\| \leq t$ and

$$\mathbf{y} \cdot \mathbf{A} \equiv \mathbf{0} \pmod{q}.$$

Now, given samples \mathbf{A}, \mathbf{c} where either $\mathbf{c} = \mathbf{A} \cdot \mathbf{s} + \mathbf{e}$ or \mathbf{c} uniform, we can distinguish the two cases by finding a short \mathbf{y} which solves SIS on \mathbf{A} and by computing $\langle \mathbf{y}, \mathbf{c} \rangle$. On the one hand, if $\mathbf{c} = \mathbf{A} \cdot \mathbf{s} + \mathbf{e}$, then $\langle \mathbf{y}, \mathbf{c} \rangle = \langle \mathbf{y} \cdot \mathbf{A}, \mathbf{s} \rangle + \langle \mathbf{y}, \mathbf{e} \rangle \equiv \langle \mathbf{y}, \mathbf{e} \rangle \pmod{q}$. If \mathbf{y} is short then $\langle \mathbf{y}, \mathbf{e} \rangle$ is also short. On the other hand, if \mathbf{c} is uniformly random, so is $\langle \mathbf{y}, \mathbf{c} \rangle$.

To pick a target norm for \mathbf{y} , HELib picks $\|\mathbf{y}\| = q$ which allows distinguishing with good probability because q is not too far from q/σ since $\sigma \approx 3.2$ and q is typically rather large. More precisely, we may rely on the following lemma:

Lemma 1 ([LP11]). Given an LWE instance characterised by n, α, q and a vector \mathbf{y} of length $\|\mathbf{y}\|$ such that $\mathbf{y} \cdot \mathbf{A} \equiv 0 \pmod{q}$, the advantage of distinguishing $\langle \mathbf{y}, \mathbf{e} \rangle$ from random is close to

$$\exp(-\pi(\|\mathbf{y}\| \cdot \alpha)^2).$$

To produce a short enough \mathbf{y} , we may call a lattice-reduction algorithm. In particular, we may call the BKZ algorithm with block size β . After performing BKZ- β reduction the first vector in the transformed lattice basis will have norm $\delta_0^m \cdot \det(\Lambda)^{1/m}$ where $\det(\Lambda)$ is the determinant of the lattice under consideration, m its dimension and the root-Hermite factor δ_0 is a constant based on the block size parameter β . Increasing the parameter β leads to a smaller δ_0 but also leads to an increase in run-time; the run-time grows at least exponential in β (see below).

In our case, the expression above simplifies to $\|\mathbf{y}\| \approx \delta_0^m \cdot q^{n/m}$ whp, where n is the LWE dimension and m is the number of samples we consider. The minimum of this expression is attained at $m = \sqrt{\frac{n \log q}{\log \delta_0}}$ [MR09].

Explicitly, we are given a matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$, construct a basis \mathbf{Y} for its left kernel modulo q and then consider the q -ary lattice $\Lambda_q(\mathbf{Y})$ spanned by the rows of \mathbf{Y} . With high probability \mathbf{Y} is an $(m - n) \times m$ matrix and $\Lambda_q(\mathbf{Y})$ has volume q^n . Let \mathbf{L} be a basis for $\Lambda_q(\mathbf{Y})$, $m' = m - n$ and write $\mathbf{Y} = [\mathbf{I}_{m'} | \mathbf{Y}']$ then we have

$$\mathbf{L} = \begin{pmatrix} \mathbf{I}_{m'} & \mathbf{Y}' \\ 0 & q \mathbf{I}_n \end{pmatrix}.$$

In other words, we are attempting to find a short vector \mathbf{y} in the integer row span of \mathbf{L} .

Given a target for the norm of \mathbf{y} and hence for δ_0 , HELib² estimates the cost of lattice reduction by relying on the following formula from [LP11]:

$$\log t_{BKZ}(\delta_0) = \frac{1.8}{\log \delta_0} - 110, \tag{1}$$

² <https://github.com/shaib/HELlib/blob/a5921a08e8b418f154be54f4e39a849e74489319/src/FHEContext.cpp#L22>.

where $t_{BKZ}(\delta_0)$ is the time in seconds it takes to BKZ reduce a basis to achieve root-Hermite factor δ_0 . This estimate is based on experiments with BKZ in the NTL library [Sho01] and extrapolation.

2.3 LP Model

The [LP11] model for estimating the cost of lattice-reduction is not correct.

Firstly, it expresses runtime in seconds instead of units of computation. As Moore's law progresses and more parallelism is introduced, the number of instructions that can be performed in a second increases. Hence, we first must translate Eq. (1) to units of computation. The experiments of Lindner and Peikert were performed on a 2.33 Ghz AMD Opteron machine, so we may assume that about $2.33 \cdot 10^9$ operations can be performed on such a machine in one second and we scale Eq. (1) accordingly.³

Secondly, the LP model does not fit the implementation of BKZ in NTL. The BKZ algorithm internally calls an oracle for solving the shortest vector problem in smaller dimension. The most practically relevant algorithms for realising this oracle are enumeration without preprocessing (Fincke-Pohst) which costs $2^{\Theta(\beta^2)}$ operations, enumeration with recursive preprocessing (Kannan) which costs $\beta^{\Theta(\beta)}$ and sieving which costs $2^{\Theta(\beta)}$. NTL implements enumeration without preprocessing. That is, while it was shown in [Wal15] that BKZ with recursive BKZ pre-processing achieves a run-time of $\text{poly}(n) \cdot \beta^{\Theta(\beta)}$, NTL does not implement the necessary recursive preprocessing with BKZ in smaller dimensions. Hence, it runs in time $\text{poly}(n) \cdot 2^{\Theta(\beta^2)}$ for block size β .

Thirdly, the LP model assumes a linear relation between $1/\log(\delta_0)$ and the log of the running time of BKZ, but from the ‘‘lattice rule-of-thumb’’ ($\delta_0 \approx \beta^{1/(2\beta)}$) and $2^{\Theta(\beta)}$ being the complexity of the best known algorithm for solving the shortest vector problem, we get:

Lemma 2 ([APS15]). *The log of the time complexity achieve a root-Hermite factor δ_0 with BKZ is*

$$\Theta\left(\frac{\log(1/\log \delta_0)}{\log \delta_0}\right)$$

if calling the SVP oracle costs $2^{\Theta(\beta)}$.

To illustrate the difference between Lemma 2 and Eq. (1), consider Regev's original parameters [Reg05] for LWE: $q \approx n^2$, $\alpha q \approx \sqrt{n}$. Then, solving LWE with the dual attack and advantage ϵ requires a log root-Hermite factor $\log \delta_0 = \log^2\left(\alpha \sqrt{\ln(1/\epsilon)/\pi}^{-1}\right)/(4n \log q)$ [APS15]. Picking ϵ such that $\log \sqrt{\ln(1/\epsilon)/\pi} \approx 1$, the log root-Hermite factor becomes $\log \delta_0 = \frac{9 \log n}{32n}$. Plugging this result into Eq. 1, we would estimate that solving LWE for these parameters takes $\log t_{BKZ}(\delta_0) = \frac{32n}{5 \log n} - 110$ s, which is subexponential in n .

³ The number of operations on integers of size $\log q$ depends on q and is not constant. However, constant scaling provides a reasonable approximation for the number of operations for the parameter ranges we are interested in here.

2.4 Parameter Choices in SEAL 2.0

SEAL v2.0 [LP16] largely leaves parameter choices to the user. However, it provides the `ChooserEvaluator::default_parameter_options()` function which returns values from [LN14, Table 2].⁴ This table gives a maximum $\log q$ for 80 bits of security for $n = 1024, 2048, 4096, 8192, 16384$. We reproduce these values for $\log q$ in Table 1. The default standard deviation is $\sigma = 3.19$.

The values of [LN14, Table 2] are based on enumeration costs and the simulator from [CN11, CN12]. Furthermore, to extrapolate from available enumeration costs from [CN12, LN14] assumes calling the SVP oracle in BKZ grows only exponentially with β , i.e. as $2^{0.64\beta-28}$. Note that this is overly optimistic, as [CN12] calls enumeration with recursive preprocessing to realise the SVP oracle inside BKZ, which has a complexity of $\beta^{\Theta(\beta)}$.

Finally, we note that the SEAL v2.0 manual [LP16] cautions the user against relying on the security provided by the list of default parameters.

2.5 Lattice Reduction

We will estimate the cost of lattice reduction using the following assumptions: BKZ- β produces vectors with $\delta_0 \approx \left(\frac{\beta}{2\pi e}(\pi\beta)^{\frac{1}{\beta}}\right)^{\frac{1}{2(\beta-1)}}$ [Che13]. The SVP oracle in BKZ is realised using sieving and sieving in blocksize β costs $t_\beta = 2^{0.292\beta+12.31}$ clock cycles. Here, 0.292β follows from [BDGL16], the additive constant $+12.31$ is based on experiments in [Laa15]. BKZ- β costs $cn \cdot t_\beta$ clock cycles in dimension n for some small constant c based on experiments in [Che13]; cf. [Che13, Figure 4.6]. This corresponds roughly to $2c$ tours of BKZ. We pick $c = 8$ based on our experiments with [FPL16].

This estimate is more optimistic than the estimate in [APS15], which does not yet take [BDGL16] into account and bases the number of SVP oracle calls on theoretical convergence results [HPS11] instead of experimental evidence. On the other hand, this estimate is more pessimistic than [BCD+16] which assumes *one* SVP call to be sufficient in order to protect against future algorithmic developments. While such developments, amortising costs across SVP calls during one BKZ reduction, are plausible, we avoid this assumption here in order not to “oversell” our results. However, we note that our improvements are somewhat oblivious to the underlying lattice-reduction model used. That is, while the concrete estimates for bit-security will vary depending on which estimate is employed, the techniques described here lead to improvements over the plain dual attack regardless of model. For completeness, we give estimated costs in different cost models in Appendix C.

According to the [LP11] estimate, solving Example 1 costs $2^{157.8}$ s or $2^{188.9}$ operations using the standard dual attack. The estimates outlined in this section predict a cost of $2^{124.2}$ operations for the same standard dual attack.

⁴ Note that the most recent version of SEAL now recommends more conservative parameters [LCP16], partly in reaction to this work.

2.6 Related Work

LWE. Besides the dual attack, via BKW or lattice-reduction, there is also the primal attack, which solves the bounded distance decoding (BDD) problem directly. That is, given (\mathbf{A}, \mathbf{c}) with $\mathbf{c} = \mathbf{A} \cdot \mathbf{s} + \mathbf{e}$ or $\mathbf{c} \leftarrow_{\S} \mathcal{U}(\mathbb{Z}_q^m)$ find \mathbf{s}' such that $\|\mathbf{w} - \mathbf{c}\|$ with $\mathbf{w} = \mathbf{A} \cdot \mathbf{s}'$ is minimised. For this, we may employ Kannan's embedding [AFG14] or variants of Babai's nearest planes after lattice reduction [LP11, LN13]. For Example 1 the cost of the latter approach is $2^{116.6}$ operations, i.e. about a factor 190 faster than the dual attack.

Arora & Ge proposed an asymptotically efficient algorithm for solving LWE [AG11], which was later improved in [ACFP14]. However, these algorithms involve large constants in the exponent, ruling them out for parameters typically considered in cryptography. We, hence, do not consider them further in this work.

Small-Secret LWE. As mentioned in [GHS12b], we can transform instances with an unusually short secret into instances where the secret follows the error distribution, but n samples have the old, short secret as noise [ACPS09].

Given a random $m \times n$ matrix $\mathbf{A} \bmod q$ and an m -vector $\mathbf{c} = \mathbf{A} \cdot \mathbf{s} + \mathbf{e} \bmod q$, let \mathbf{A}_0 denotes the first n rows of \mathbf{A} , \mathbf{A}_1 the next n rows, etc., $\mathbf{e}_0, \mathbf{e}_1, \dots$ are the corresponding parts of the error vector and $\mathbf{c}_0, \mathbf{c}_1, \dots$ the corresponding parts of \mathbf{c} . We have $\mathbf{c}_0 = \mathbf{A}_0 \cdot \mathbf{s} + \mathbf{e}_0$ or $\mathbf{A}_0^{-1} \cdot \mathbf{c}_0 = \mathbf{s} + \mathbf{A}_0^{-1} \mathbf{e}_0$. For $i > 0$ we have $\mathbf{c}_i = \mathbf{A}_i \cdot \mathbf{s} + \mathbf{e}_i$, which together with the above gives $\mathbf{A}_i \mathbf{A}_0^{-1} \mathbf{c}_0 - \mathbf{c}_i = \mathbf{A}_i \mathbf{A}_0^{-1} \mathbf{e}_0 - \mathbf{e}_i$. The output of the transformation is $\mathbf{z} = \mathbf{B} \cdot \mathbf{e}_0 + \mathbf{f}$ with $\mathbf{B} = (\mathbf{A}_0^{-1} \mid \mathbf{A}_1 \cdot \mathbf{A}_0^{-1} \mid \dots)$ and $\mathbf{z} = (\mathbf{A}_0^{-1} \mathbf{c}_0 \mid \mathbf{A}_1 \mathbf{A}_0^{-1} \mathbf{c}_1 \mid \dots)$ and $\mathbf{f} = (\mathbf{s} \mid \mathbf{e}_1 \mid \dots)$. For Example 1, this reduces α from $2^{-60.4}$ to $\approx 2^{-60.8}$ and marginally improves the cost of solving.

An explicit variant of this approach is given in [BG14]. Consider the lattice

$$\Lambda = \{\mathbf{v} \in \mathbb{Z}^{n+m} \mid [\mathbf{A} \mid \mathbf{I}_m] \cdot \mathbf{v} \equiv 0 \bmod q\}.$$

It has an unusually short vector $(\mathbf{s} \parallel \mathbf{e})$. When $\|\mathbf{s}\| \ll \|\mathbf{e}\|$, the vector $(\mathbf{s} \parallel \mathbf{e})$ is uneven in length. To balance the two sides, rescale the first part to have the same norm as the second. When $\mathbf{s} \leftarrow_{\S} \mathcal{B}^-$, this scales the volume of the lattice by σ^n . When $\mathbf{s} \leftarrow_{\S} \mathcal{B}^+$, this scales the volume of the lattice by $(2\sigma)^n$ because we can scale by 2σ and then re-balance. When $\mathbf{s} \leftarrow_{\S} \mathcal{B}_h^\pm$, the volume is scaled depending on h . For our rolling example, this approach costs $2^{114.5}$ operations, i.e. is about a factor 830 faster than the dual attack.

Independently and concurrently to this work, a new key-exchange protocol based on sparse secret LWE was proposed in [CKH+16]. A subset of the techniques discussed here are also discussed in [CKH+16], in particular, ignoring components of the secret and using lattice scaling as in [BG14].

Combinatorial. This work combines combinatorial and lattice-reduction techniques. As such, it has some similarities with the hybrid attack on NTRU [HG07]. This attack was recently adapted to LWE in the \mathcal{B} -secret case in [BGPW16] and its complexity revisited in [Wun16].

Rings. Recently, [ABD16] proposed a subfield lattice-attack on the two fully homomorphic encryption schemes YASHE [BLLN13] and LTV [LTV12], showing that NTRU with “overstretched” moduli q is less secure than initially expected. Quickly after, [KF16] pointed out that the presence of subfields is not necessary for attacks to succeed. NTRU can be considered as the homogeneous version of Ring-LWE, but there is currently no indication that these attacks can be translated to the Ring-LWE setting. There is currently no known algorithm which solves Ring-LWE faster than LWE for the parameter choices (ring, error distribution, etc.) typically considered in FHE schemes.

3 Amortising Costs

If the cost of distinguishing LWE from random with probability ε is c , the cost of solving is customarily estimated as at least c/ε [LP11]. More precisely, applying Chernoff bounds, we require about $1/\varepsilon^2$ samples to amplify a decision experiment succeeding with advantage ε to a constant advantage. Hence, e.g. in [APS15], the dual attack is costed as the cost of running BKZ- β to achieve the target δ_0 multiplied by the number of samples required to distinguish with the target advantage, i.e. $\approx c/\varepsilon^2$.

In the case of the dual attack, this cost can be reduced by performing rerandomisation on the already reduced basis. If \mathbf{L} is a basis for the lattice $\Lambda_q(\mathbf{Y})$, we first compute \mathbf{L}' as the output of BKZ- β reduction where β is chosen to achieve the target δ_0 required for some given target advantage. Then, in order to produce sufficiently many relatively short vectors $\mathbf{y}_i \in \Lambda_q(\mathbf{Y})$ we repeatedly multiply \mathbf{L}' by a fresh random sparse unimodular matrix with small entries to produce \mathbf{L}'_i . As a consequence, \mathbf{L}'_i remains somewhat short. Finally, we run BKZ- β' with $\beta' \leq \beta$ on \mathbf{L}'_i and return the smallest non-zero vector as \mathbf{y}_i . See Algorithm 1, where ε_d is chosen following Lemma 1 (see below for the expectation of $\|\mathbf{y}\|$) and m is chosen following [SL12].

That is, similar to BKW, which in a first step produces elimination tables which allow sampling smaller dimensional LWE samples in $\mathcal{O}(n^2)$ operations, we first produce a relatively good basis \mathbf{L}' to allow sampling \mathbf{y}_i relatively efficiently.

To produce the estimates in Table 1, we assume the same rerandomisation strategy as is employed in fpLLL’s implementation [FPL16] of extreme pruning for BKZ 2.0.⁵ This rerandomisation strategy first permutes rows and then adds three existing rows together using ± 1 coefficients, which would increase norms by a factor of $\sqrt{3} < 2$ when all vectors initially have roughly the same norm. For completeness, we reproduce the algorithm in Appendix A. We then run LLL, i.e. we set $\beta' = 2$, and assume that our \mathbf{y}_i have their norms increased by a factor of two, i.e. $E[\|\mathbf{y}_i\|] = 2 \cdot \delta_0^m q^{n/m}$.

Heuristic. We note that, in implementing this strategy, we are losing statistical independence. To maintain statistical independence, we would consider fresh

⁵ <https://github.com/fplll/fplll/blob/b75fe83/fplll/bkz.cpp#L43>.

Data: candidate LWE samples $\mathbf{A}, \mathbf{c} \in \mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^m$
Data: BKZ block sizes $\beta, \beta' \geq 2$
Data: target success probability ε
 $\varepsilon_d \leftarrow \exp(-\pi(\mathbb{E}[\|\mathbf{y}_i\|] \cdot \alpha)^2)$;
 $m \leftarrow \lceil 2 \log(2 - 2\varepsilon) / \log(1 - 4\varepsilon_d^2) \rceil$;
 $\mathbf{L} \leftarrow$ basis for $\{\mathbf{y} \in \mathbb{Z}^m : \mathbf{y} \cdot \mathbf{A} \equiv 0 \pmod{q}\}$;
 $\mathbf{L}' \leftarrow$ BKZ- β reduced basis for \mathbf{L} ;
for $i \leftarrow 0$ **to** $m - 1$ **do**
 $\mathbf{U} \leftarrow_{\S}$ a sparse unimodular matrix with small entries;
 $\mathbf{L}_i \leftarrow \mathbf{U} \cdot \mathbf{L}'$;
 $\mathbf{L}'_i \leftarrow$ BKZ- β' reduced basis for \mathbf{L}_i ;
 $\mathbf{y}_i \leftarrow$ shortest row vector in \mathbf{L}'_i ;
 $\mathbf{e}'_i \leftarrow \langle \mathbf{y}_i, \mathbf{c} \rangle$;
end
if \mathbf{e}'_i follow discrete Gaussian distribution **then**
 return \top ;
else
 return \perp ;
end

Algorithm 1. SILKE₁: Amortising costs in BKW-style SIS strategy for solving LWE

LWE samples and distinguish $\langle \mathbf{y}_i, \mathbf{e}_i \rangle$ from uniform. However, neither HELib nor SEAL provides the attacker with sufficiently many samples to run the algorithm under these conditions. Instead, we are attempting to distinguish $\langle \mathbf{y}_i, \mathbf{e} \rangle$ from uniform. Furthermore, since we are performing only light rerandomisation our distribution could be skewed if our \mathbf{y}_i in $\langle \mathbf{y}_i, \mathbf{e} \rangle$ are not sufficiently random. Just as in BKW-style algorithms [ACF+15] we assume the values $\langle \mathbf{y}_i, \mathbf{e} \rangle$ are distributed closely enough to the target distribution to allow us to ignore this issue.

Experimental Verification. We tested the heuristic assumption of Algorithm 1 by rerandomising a BKZ-60 reduced basis using Algorithm 4 with $d = 3$ followed by LLL reduction several hundred times. In this experiment, we recovered fresh somewhat short vectors in each call, where somewhat short means with a norm at most twice that of the shortest vector of \mathbf{L}' . We give further experimental evidence in Sect. 6.

Finally, we note that this process shares some similarities with random sampling reduction (RSR) [Sch03], where random linear combinations are LLL reduced to produce short vectors. While, here, we are only performing sparse sums and accept *larger* norms, the techniques used to analyse RSR might permit reducing our heuristic to a more standard heuristic assumption.

4 Scaled Normal-Form

The line of research improving the BKW algorithm for small secrets starting with [AFFP14] proceeds from the observation that we do not need to find

$\mathbf{y} \cdot \mathbf{A} \equiv 0 \pmod q$, but if the secret is sufficiently small then any \mathbf{y} such that $\mathbf{y} \cdot \mathbf{A}$ is short suffices, i.e. we seek short vectors (\mathbf{w}, \mathbf{v}) in the lattice

$$\Lambda = \{(\mathbf{y}, \mathbf{x}) \in \mathbb{Z}^m \times \mathbb{Z}^n : \mathbf{y} \cdot \mathbf{A} \equiv \mathbf{x} \pmod q\}.$$

Note that this lattice is the lattice considered in dual attacks on normal form LWE instances (cf. [ADPS15]).⁶ Given a short vector in $(\mathbf{w}, \mathbf{v}) \in \Lambda$, we have

$$\mathbf{w} \cdot \mathbf{c} = \mathbf{w} \cdot (\mathbf{A} \cdot \mathbf{s} + \mathbf{e}) = \langle \mathbf{v}, \mathbf{s} \rangle + \langle \mathbf{w}, \mathbf{e} \rangle.$$

Here, \mathbf{v} corresponds to the noise from “modulus switching” or quantisation in BKW-style algorithms and \mathbf{w} to the multiplicative factor by which the LWE noise increases due to repeated subtractions.

Now, in small secret LWE instances we have $\|\mathbf{s}\| < \|\mathbf{e}\|$. As a consequence, we may permit $\|\mathbf{v}\| > \|\mathbf{w}\|$ such that

$$\|\langle \mathbf{w}, \mathbf{s} \rangle\| \approx \|\langle \mathbf{v}, \mathbf{e} \rangle\|.$$

Hence, we consider the lattice

$$\Lambda_c = \{(\mathbf{y}, \mathbf{x}/c) \in \mathbb{Z}^m \times (1/c \cdot \mathbb{Z})^n : \mathbf{y} \cdot \mathbf{A} \equiv \mathbf{x} \pmod q\}$$

for some constant c , similar to [BG14]. The lattice Λ_c has dimension $m' = m + n$ and whp volume $(q/c)^n$. To construct a basis for Λ_c , assume $\mathbf{A}_{m-n:m}$ has full rank (this holds with high probability for large q). Then $\Lambda_c = \Lambda(\mathbf{L}')$ with

$$\mathbf{L}' = \begin{pmatrix} \frac{1}{c} \mathbf{I}_n & \mathbf{0}_{n \times (m-n)} & \mathbf{A}_{m-n:m}^{-1} \\ & \mathbf{I}_{m-n} & \mathbf{B}' \\ & & q \mathbf{I}_n \end{pmatrix}$$

where $[\mathbf{I}_{m-n} | \mathbf{B}']$ is a basis for the left kernel of $\mathbf{A} \pmod q$.

Remark 2. In our estimates for HELib and SEAL, we typically have $m = n$ and $[\mathbf{I}_{m-n} | \mathbf{B}'] \in \mathbb{Z}^{0 \times n}$.

It remains to establish c . Lattice reduction produces a vector (\mathbf{w}, \mathbf{v}) with

$$\|(\mathbf{w}, \mathbf{v})\| \approx \delta_0^{m'} \cdot (q/c)^{n/m'}, \tag{2}$$

which translates to a noise value

$$e = \mathbf{w} \cdot \mathbf{A} \cdot \mathbf{s} + \langle \mathbf{w}, \mathbf{e} \rangle = \langle c \cdot \mathbf{v}, \mathbf{s} \rangle + \langle \mathbf{w}, \mathbf{e} \rangle$$

and we set

$$c = \frac{\alpha q}{\sqrt{2} \pi h} \equiv \sqrt{m' - n}$$

to equalise the noise contributions of both parts of the above sum.

As a consequence, we arrive at the following lemma, which is attained by combining Eq. (2) with Lemma 1.

⁶ The strategy seems folklore, we were unable to find a canonical reference for it.

Lemma 3. *Let $m' = 2n$ and $c = \frac{\alpha q}{\sqrt{2\pi h}} \cdot \sqrt{m' - n}$. A lattice reduction algorithm achieving δ_0 such that*

$$\log \delta_0 = \frac{\log \left(\frac{2n \log^2 \varepsilon}{\pi \alpha^2 h} \right)}{8n}$$

leads to an algorithm solving decisional LWE with $\mathbf{s} \leftarrow_{\S} \mathcal{B}_h^-$ instance with advantage ε and the same cost.

Remark 3. We focus on $m' = 2n$ in Lemma 3 for ease of exposure. For the instances considered in this work, $m' = 2n$ is a good approximation for m' (see Sect. 6).

For Example 1 we predict at a cost of $2^{107.4}$ operations mod q for solving Decision-LWE when applying this strategy. Amortising costs as suggested in Sect. 3 reduces it further to $2^{101.0}$ operations mod q .

Asymptotic Behaviour. The general dual strategy, without exploiting small secrets, requires

$$\log \delta_0 = \frac{\log \left(-\frac{2 \log \varepsilon}{\alpha^2 q} \right)}{4n}$$

according to [APS15]. For HELib’s choice of $8 = \alpha q$ and $h = 64$ and setting ε constant, this expression simplifies to

$$\log \delta_0 = \frac{\log q + C_d}{4n},$$

for some constant C_d . On the other hand, Lemma 3 simplifies to

$$\log \delta_0 = \frac{\log q + \frac{1}{2} \log n + C_m}{4n}, \tag{3}$$

for some constant $C_m < C_d$.

For a circuit of depth L , BGV requires $\log q = L \log n + \mathcal{O}(L)$ [GHS12b, Appendix C.2]. Applying Lemma 2, we get that

$$\lim_{\kappa \rightarrow \infty} \frac{\text{cost}_m}{\text{cost}_d} = \lim_{n \rightarrow \infty} \frac{\text{cost}_m}{\text{cost}_d} = \frac{2L}{2L + 1},$$

where cost_d is the log cost of the standard dual attack, cost_m is the log cost under Lemma 3 and κ the security parameter. The same analysis applies to any constant h . Finally, when $h = 2/3n$, i.e. $\mathbf{s} \leftarrow_{\S} \mathcal{B}^-$, then the term $1/2 \cdot \log n$ vanishes from (3), but $C_m > C_d$.

5 Sparse Secrets

Recall that BKW-style algorithms consist of two stages or, indeed, sub-algorithms. First, in the reduction stage, combinatorial methods are employed

to transform an LWE instance in dimension n into an instance of dimension $0 \leq n' \leq n$, typically with increased noise level α . This smaller LWE instance is then, in the solving stage, is solved using some form of exhaustive search over the secret.

Taking the same perspective on the dual attack, write $\mathbf{A} = [\mathbf{A}_0 \mid \mathbf{A}_1]$ with $\mathbf{A}_0 \in \mathbb{Z}_q^{m \times (n-k)}$ and $\mathbf{A}_1 \in \mathbb{Z}_q^{m \times k}$ and find a short vector in the lattice

$$\Lambda = \{\mathbf{y} \in \mathbb{Z}^m : \mathbf{y} \cdot \mathbf{A}_0 \equiv 0 \pmod q\}.$$

Each short vector $\mathbf{y} \in \Lambda$ produces a sample for an LWE instance in dimension k and noise rate $\alpha' = \mathbb{E}[\|\mathbf{y}\|] \cdot \alpha$. Setting $k = 0$ recovers the original dual attack. For $k > 0$, we may now apply our favourite algorithm for solving small dimensional, easy LWE instances. Applying exhaustive search implies $\log_2 k < \kappa$ for $\mathbf{s} \leftarrow_{\S} \mathcal{B}^+$ resp. $\log_3 k < \kappa$ for $\mathbf{s} \leftarrow_{\S} \mathcal{B}^-$ when κ is the target level of security.

The case $\mathbf{s} \leftarrow_{\S} \mathcal{B}_h^{\pm}$ permits much larger k by relaxing the conditions we place on solving the k -dimensional instance. Instead of solving with probability one, we solve with some probability p_k and rerun the algorithm in case of failure.

For this, write $\mathbf{A} \cdot \mathbf{P} = [\mathbf{A}_0 \mid \mathbf{A}_1]$ and $\mathbf{s} \cdot \mathbf{P} = [\mathbf{s}_0 \mid \mathbf{s}_1]$ where \mathbf{P} is a random permutation matrix. Now, over the choice of \mathbf{P} there is a good chance that $\mathbf{s}_1 = 0$ and hence that $\mathbf{A}_1 \cdot \mathbf{s}_1 \equiv 0 \pmod q$. That is, the right choice of \mathbf{P} places all non-zero components of \mathbf{s} in the \mathbf{s}_0 part.

In particular, with probability $1-h/n$ a coordinate $\mathbf{s}_{(i)}$ is zero. More generally, picking k components of \mathbf{s} at random will pick only components such that $\mathbf{s}_{(i)} = 0$ with probability

$$p_k = \prod_{i=0}^{k-1} \left(1 - \frac{h}{n-i}\right) = \frac{\binom{n-h}{k}}{\binom{n}{k}} \approx \left(1 - \frac{h}{n}\right)^k.$$

Hence, simply treating $k > 0$ in the solving stage the same as $k = 0$ succeeds with probability p_k . The success probability can be amplified to close to one by repeating the elimination and solving stages $\approx 1/p_k$ times assuming we distinguish with probability close to 1.

It is clear that the same strategy translates to the primal attack by simply dropping random columns before running the algorithm. However, for the dual attack, the following improvement can be applied. Instead of considering only $\mathbf{s}_1 = 0$, perform exhaustive search over those solutions that occur with sufficiently high probability. In particular, over the choice of \mathbf{P} , the probability that \mathbf{s}_1 contains $k-j$ components with $\mathbf{s}_{1,(i)} = 0$ and exactly j components with $\mathbf{s}_{1,(i)} \neq 0$ is

$$p_{k,j} = \frac{\binom{n-h}{k-j} \binom{h}{j}}{\binom{n}{k}},$$

i.e. follows the hypergeometric distribution.

Now, assuming $\mathbf{s} \leftarrow_{\S} \mathcal{B}_h^-$, to check if any of those candidates for \mathbf{s}_1 is correct, we need to compare $\binom{k}{j} \cdot 2^j$ distributions against the uniform distribution mod q .

Thus, after picking a parameter ℓ we arrive at Algorithm 2 with cost:

1. m calls to BKZ- β in dimension $n - k$.
2. $m \cdot \sum_{i=0}^{\ell} \binom{k}{i} \cdot 2^i \cdot i$ additions mod q to evaluate m samples on all possible solutions up to weight ℓ .

Assuming m is chosen such that distinguishing LWE from uniform succeeds with probability close to one, then Algorithm 2 succeeds with probability $\sum_{j=0}^{\ell} p_{k,j}$.

Data: $m \times n$ matrix \mathbf{A} over \mathbb{Z}_q
Data: m vector \mathbf{c} over \mathbb{Z}_q
Data: density parameter $0 \leq \ell \leq 64$
Data: dimension parameter $0 \leq k \leq n$
 $\mathbf{P} \leftarrow_{\S} n \times n$ permutation matrices;
 $[\mathbf{A}_0 \mid \mathbf{A}_1] \leftarrow \mathbf{A} \cdot \mathbf{P}$ with $\mathbf{A}_0 \in \mathbb{Z}_q^{m \times (n-k)}$;
 $\mathbf{L} \leftarrow$ basis for scaled-dual lattice of \mathbf{A}_0 ;
for $i \leftarrow 0$ **to** $m - 1$ **do**
 | $\mathbf{y}_i \leftarrow$ a short vector in the row span of \mathbf{L} ;
 | $e'_i \leftarrow \langle \mathbf{y}_i, \mathbf{c} \rangle$;
end
if e'_i follow discrete Gaussian distribution **then**
 | **return** \top ;
end
foreach \mathbf{s}' in the set of $\sum_{i=0}^{\ell} \binom{k}{i} \cdot 2^i$ candidate solutions **do**
 | **for** $i \leftarrow 0$ **to** $m - 1$ **do**
 | $e''_i = e'_i + \langle \mathbf{y}_i \cdot \mathbf{A}_1, \mathbf{s}' \rangle$;
 | **end**
 | **if** e''_i follow discrete Gaussian distribution **then**
 | **return** \top ;
 | **end**
 | **end**
end
return \perp ;

Algorithm 2. SILKE₂: Sparse secrets in BKW-style SIS strategy for solving LWE.

Asymptotic Behaviour. We arrive at the following simple lemma:

Lemma 4. Let $0 \leq h < n$ and $d > 1$ be constants, $p_{h,d}$ be some constant depending on h and d , $c_{n,\alpha,q}$ be the cost of solving LWE with parameters n, α, q with probability $\geq 1 - 2^{-p_{h,d}^2}$. Then, solving LWE in dimension n with $\mathbf{s} \leftarrow_{\S} \mathcal{B}_h^{\pm}$ costs $\mathcal{O}(c_{n-n/d,\alpha,q})$ operations.

Proof. Observe that $p_{h,d} = \lim_{n \rightarrow \infty} \binom{n-h}{n/d} / \binom{n}{n/d}$ is a constant for any constant $0 \leq h < n$ and $d > 1$. Hence, solving $\mathcal{O}(1/p_{h,d}) = \mathcal{O}(1)$ instances in dimension $n - n/d$ solves the instance in dimension n . \square

Remark 4. Picking $d = 2$ we get $\lim_{n \rightarrow \infty} \binom{n-h}{n/2} / \binom{n}{n/2} = 2^{-h}$ and an overall costs of $\mathcal{O}(2^h \cdot c_{n/2, \alpha, q})$. This improves on exhaustive search, which costs $\mathcal{O}(2^h \cdot \binom{n}{h})$, when $c_{n/2, \alpha, q} \in o\left(\binom{n}{h}\right)$.

6 Combined

Combining the strategies described in this work, we arrive at Algorithm 3 (SILKE). It takes a flag *sparse* which enables the sparse strategy of Algorithm 2. In this case, we enforce that distinguishing LWE from uniform succeeds with probability $1 - 2^{-\kappa}$ when we guessed \mathbf{s}' correctly. Clearly, this parameter can be improved, i.e. this probability reduced, but amplifying the success probability is relatively cheap, so we forego this improvement.

We give an implementation of Algorithm 3 for *sparse* = false in Appendix B. For brevity, we skip the *sparse* = true case. We also tested our implementation on several parameter sets:⁷

1. Considering an LWE instance with $n = 100$ and $q \approx 2^{23}$, $\alpha = 8/q$ and $h = 20$, we first BKZ-50 reduced the basis \mathbf{L} for $c = 16$. This produced a short vector \mathbf{w} such that $|\langle \mathbf{w}, \mathbf{c} \rangle| \approx 2^{15.3}$. Then, running LLL 256 times, we produced short vectors such that $E[|\langle \mathbf{w}_i, \mathbf{c} \rangle|] = 2^{15.7}$ and standard deviation $2^{16.6}$.
2. Considering an LWE instance with $n = 140$ and $q \approx 2^{40}$, $\alpha = 8/q$ and $h = 32$, we first BKZ-70 reduced the basis \mathbf{L} for $c = 1$. This took 64 hours and produced a short vector \mathbf{w} such that $|\langle \mathbf{w}, \mathbf{c} \rangle| \approx 2^{23.7}$, with $E[|\langle \mathbf{w}, \mathbf{c} \rangle|] \approx 2^{25.5}$ conditioned on $|\mathbf{w}|$. Then, running LLL 140 times (each run taking about 50s on average), we produced short vectors such that $E[|\langle \mathbf{w}_i, \mathbf{c} \rangle|] = 2^{26.0}$ and standard deviation $2^{26.4}$ for $\langle \mathbf{w}_i, \mathbf{c} \rangle$.
3. Considering the same LWE instance with $n = 140$ and $q \approx 2^{40}$, $\alpha = 8/q$ and $h = 32$, we first BKZ-70 reduced the basis \mathbf{L} for $c = 16$. This took 65 hours and produced a short vector \mathbf{w} such that $|\langle \mathbf{w}, \mathbf{c} \rangle| \approx 2^{24.7}$ after scaling by c , cf. $E[|\langle \mathbf{w}, \mathbf{c} \rangle|] \approx 2^{24.8}$. Then, running LLL 140 times (each run taking about 50s on average), we produced short vectors such that $E[|\langle \mathbf{w}_i, \mathbf{c} \rangle|] = 2^{25.5}$ and standard deviation $2^{25.9}$ for $\langle \mathbf{w}_i, \mathbf{c} \rangle$.
4. Considering again the same LWE instance with $n = 140$ and $q \approx 2^{40}$, $\alpha = 8/q$ and $h = 32$, we first BKZ-70 reduced the basis \mathbf{L} for $c = 1$. This took 30 hours and produced a short vector \mathbf{w} such that $|\langle \mathbf{w}, \mathbf{c} \rangle| \approx 2^{25.2}$, cf. $E[|\langle \mathbf{w}, \mathbf{c} \rangle|] \approx 2^{25.6}$. Then, running LLL 1024 times (each run taking about 50s on average), we produced 1016 short vectors such that $E[|\langle \mathbf{w}_i, \mathbf{c} \rangle|] = 2^{25.8}$ and standard deviation $2^{26.1}$ for $\langle \mathbf{w}_i, \mathbf{c} \rangle$.

⁷ All experiments on “strombenzin” with Intel(R) Xeon(R) CPU E5-2667 v2 @ 3.30 GHz.

Data: candidate LWE samples $\mathbf{A}, \mathbf{c} \in \mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^m$
Data: BKZ block sizes $\beta, \beta' \geq 2$
Data: target success probability ε
Data: *sparse* flag toggling sparse strategy
Data: scale factor $c \geq 1$
Data: dimension parameter $0 \leq k \leq n$, 0 when *sparse* is set
Data: density parameter $0 \leq \ell \leq k$, 0 when *sparse* is set
// distinguishing advantage per sample from β, β'
 $\varepsilon_d \leftarrow \exp(-\pi(\mathbb{E}[\|\mathbf{y}_i\|] \cdot \alpha)^2)$;
if *sparse* **then**
 $\varepsilon_t \leftarrow 1 - 1/2^\kappa$; *// for security parameter κ*
 $r \leftarrow \max\left(\lceil \log(1 - \varepsilon) / \log(1 - \sum_{j=0}^{\ell} p_{k,j}) \rceil, 1\right)$;
else
 $\varepsilon_t, r \leftarrow \varepsilon, 1$;
end
// required number of samples for majority vote
 $m \leftarrow \lceil 2 \log(2 - 2\varepsilon_t) / \log(1 - 4\varepsilon_d^2) \rceil$;
repeat r *times*
 $\mathbf{P} \leftarrow_{\S} n \times n$ permutation matrices;
 $[\mathbf{A}_0 \mid \mathbf{A}_1] \leftarrow \mathbf{A} \cdot \mathbf{P}$ with $\mathbf{A}_0 \in \mathbb{Z}_q^{m \times (n-k)}$;
 $\mathbf{L} \leftarrow$ basis for $\{(\mathbf{y}, \mathbf{x}/c) \in \mathbb{Z}^m \times (1/c \cdot \mathbb{Z})^n : \mathbf{y} \cdot \mathbf{A}_0 \equiv \mathbf{x} \pmod{q}\}$;
 $\mathbf{L}' \leftarrow$ BKZ- β reduced basis for \mathbf{L} ;
 for $i \leftarrow 0$ **to** $m - 1$ **do**
 $\mathbf{U} \leftarrow_{\S}$ a sparse unimodular matrix with small entries;
 $\mathbf{L}_i \leftarrow \mathbf{U} \cdot \mathbf{L}'$;
 $\mathbf{L}'_i \leftarrow$ BKZ- β' reduced basis for \mathbf{L}_i ;
 $(\mathbf{w}_i, \mathbf{v}_i) \leftarrow$ shortest row vector in \mathbf{L}'_i ;
 $e'_i \leftarrow \langle \mathbf{w}_i, \mathbf{c} \rangle$;
 end
 if e'_i *follow discrete Gaussian distribution* **then**
 return \top ;
 end
 foreach \mathbf{s}' *in the set of $\sum_{i=1}^{\ell} \binom{k}{i} \cdot 2^i$ candidate solutions* **do**
 for $i \leftarrow 0$ **to** $m - 1$ **do**
 $e''_i = e'_i + \langle \mathbf{w}_i \cdot \mathbf{A}_1, \mathbf{s}' \rangle$;
 end
 if e''_i *follow discrete Gaussian distribution* **then**
 return \top ;
 end
 end
 end
return \perp ;

Algorithm 3. SILKE: (Sparse) BKW-style SIS Strategy for solving LWE

5. Considering an LWE instance with $n = 180$ and $q \approx 2^{40}$, $\alpha = 8/q$ and $h = 48$, we first BKZ-70 reduced the basis \mathbf{L} for $c = 8$. This took 198 hours⁸ and produced a short vector \mathbf{w} such that $|\langle \mathbf{w}, \mathbf{c} \rangle| \approx 2^{26.7}$, cf. $E[|\langle \mathbf{w}, \mathbf{c} \rangle|] \approx 2^{25.9}$. Then, running LLL 180 times (each run taking about 500 s on average), we produced short vectors such that $E[|\langle \mathbf{w}_i, \mathbf{c} \rangle|] = 2^{26.6}$ and standard deviation $2^{26.9}$ for $\langle \mathbf{w}_i, \mathbf{c} \rangle$.

All our experiments match our prediction bounding the growth of the norms of our vectors by a factor of two. Note, however, that in the fourth experiment 1 in 128 vectors found with LLL was a duplicate of previously discovered vector, indicating that re-randomisation is not perfect. While the effect of this loss on the running time of the overall algorithm is small, it highlights that further research is required on the interplay of re-randomisation and lattice reduction.

Applying Algorithm 3 to parameter choices from HELib and SEAL, we arrive at the estimates in Table 1. These estimates were produced using the Sage [S+15] code available at <http://bitbucket.org/malb/lwe-estimator> which optimises the parameters c, ℓ, k, β to minimise the overall cost.

For the HELib parameters in Table 1 we chose the sparse strategy. Here, amortising costs as in Sect. 3 did not lead to a significant improvement, which is why we did not use it in these cases. All considered lattices have dimension $< 2n$. Hence, one Ring-LWE sample is sufficient to mount these attacks. Note that this is less than the dual attack as described in [GHS12a] would require (two samples).

For the SEAL parameter choices in Table 1, dimension $n = 1024$ requires two Ring-LWE samples, larger dimensions only require one sample. Here, amortising costs as in Algorithm 1 does lead to a modest improvement and is hence enabled.

Finally, we note that reducing q to $\approx 2^{34}$ resp. $\approx 2^{560}$ leads to an estimated cost of 80 bits for $n = 1024$ resp. $n = 16384$ for $\mathbf{s} \leftarrow_{\S} \mathcal{B}_{64}^-$. For $\mathbf{s} \leftarrow_{\S} \mathcal{B}^-$, $q \approx 2^{40}$ resp. $q \approx 2^{660}$ leads to an estimated cost of 80 bits under the techniques described here. In both cases, we assume $\sigma \approx 3.2$.

Acknowledgements. We thank Kenny Paterson and Adeline Roux-Langlois for helpful comments on an earlier draft of this work. We thank Hao Chen for reporting an error in an earlier version of this work.

⁸ We ran 49 BKZ tours until fpLLL's auto abort triggered. After 16 tours the norm of the then shortest vector was by a factor 1.266 larger than the norm of the shortest vector found after 49 tours.

A Rerandomisation

Data: $n \times m$ matrix \mathbf{L}

Data: density parameter d , default $d = 3$

Result: $\mathbf{U} \cdot \mathbf{L}$ where \mathbf{U} is a sparse, unimodular matrix.

```

for  $i \leftarrow 0$  to  $4 \cdot n - 1$  do
  |  $a \leftarrow_{\S} \{0, n - 1\}$ ;
  |  $b \leftarrow_{\S} \{0, n - 1\} \setminus \{a\}$ ;
  |  $\mathbf{L}_{(b)}, \mathbf{L}_{(a)} \leftarrow \mathbf{L}_{(a)}, \mathbf{L}_{(b)}$  ;
end
for  $a \leftarrow 0$  to  $n - 2$  do
  | for  $i \leftarrow 0$  to  $d - 1$  do
    |  $b \leftarrow_{\S} \{a + 1, n - 1\}$ ;
    |  $s \leftarrow_{\S} \{0, 1\}$ ;
    |  $\mathbf{L}_{(a)} \leftarrow \mathbf{L}_{(a)} + (-1)^s \cdot \mathbf{L}_{(b)}$ ;
  | end
end
return  $\mathbf{L}$ ;

```

Algorithm 4. Rerandomisation strategy in the `fpLLL` library [FPL16].

B Implementation

```

# -*- coding: utf-8 -*-
from sage.all import shuffle, randint, ceil, next_prime, log, cputime, mean, variance,
set_random_seed, sqrt
from copy import copy
from sage.all import GF, ZZ
from sage.all import random_matrix, random_vector, vector, matrix, identity_matrix
from sage.stats.distributions.discrete_gaussian_integer import DiscreteGaussianDistributionIntegerSampler \
as DiscreteGaussian
from estimator.estimator import preprocess_params, stddev

def gen_fhe_instance(n, q, alpha=None, h=None, m=None, seed=None):
    """
    Generate FHE-style LWE instance

    :param n: dimension
    :param q: modulus
    :param alpha: noise rate (default: 8/q)
    :param h: hamming weight of the secret (default: 2/3n)
    :param m: number of samples (default: n)

    """
    if seed is not None:
        set_random_seed(seed)

    q = next_prime(ceil(q)-1, proof=False)
    if alpha is None:
        alpha = ZZ(8)/q

    n, alpha, q = preprocess_params(n, alpha, q)

    stddev = stddev(alpha*q)

    if m is None:
        m = n
    K = GF(q, proof=False)
    A = random_matrix(K, m, n)

    if h is None:
        s = random_vector(ZZ, n, x=-1, y=1)
    else:
        S = [-1, 1]
        s = [S[randint(0, 1)] for i in range(h)]
        s += [0 for _ in range(n-h)]
        shuffle(s)

```

```

    s = vector(ZZ, s)
    c = A*s

    D = DiscreteGaussian(stddev)

    for i in range(m):
        c[i] += D()

    return A, c

def dual_instance0(A):
    """
    Generate dual attack basis.

    :param A: LWE matrix A
    """
    q = A.base_ring().order()
    B0 = A.left_kernel().basis_matrix().change_ring(ZZ)
    m = B0.ncols()
    n = B0.nrows()
    r = m-n
    B1 = matrix(ZZ, r, n).augment(q*identity_matrix(ZZ, r))
    B = B0.stack(B1)
    return B

def dual_instance1(A, scale=1):
    """
    Generate dual attack basis for LWE normal form.

    :param A: LWE matrix A
    """
    q = A.base_ring().order()
    n = A.ncols()
    B = A.matrix_from_rows(range(0, n)).inverse().change_ring(ZZ)
    L = identity_matrix(ZZ, n).augment(B)
    L = L.stack(matrix(ZZ, n, n).augment(q*identity_matrix(ZZ, n)))

    for i in range(0, 2*n):
        for j in range(n, 2*n):
            L[i, j] = scale*L[i, j]

    return L

def balanced_lift(e):
    """
    Lift e mod q to integer such that result is between -q/2 and q/2

    :param e: a value or vector mod q
    """
    from sage.rings.finite_rings.integer_mod import is_IntegerMod

    q = e.base_ring().order()
    if is_IntegerMod(e):
        e = ZZ(e)
        if e > q//2:
            e -= q
        return e
    else:
        return vector(balanced_lift(ee) for ee in e)

def apply_short1(y, A, c, scale=1):
    """
    Compute 'y*A', 'y*c' where y is a vector in the integer row span of
    'dual_instance(A)'

    :param y: (short) vector in scaled dual lattice
    :param A: LWE matrix
    :param c: LWE vector
    """
    m = A.nrows()
    y = vector(ZZ, 1/ZZ(scale) * y[-m:])
    a = balanced_lift(y*A)
    e = balanced_lift(y*c)
    return a, e

def log_mean(X):
    return log(mean([abs(x) for x in X]), 2)

def log_var(X):
    return log(variance(X).sqrt(), 2)

```

```

def silke(A, c, beta, h, m=None, scale=1, float_type="double"):
    """
    :param A:     LWE matrix
    :param c:     LWE vector
    :param beta:  BKW block size
    :param m:     number of samples to consider
    :param scale: scale rhs of lattice by this factor
    """
    from fpylll import BKZ, IntegerMatrix, LLL, GSO
    from fpylll.algorithms.bkz2 import BKZReduction as BKZ2

    if m is None:
        m = A.nrows()

    L = dual_instance1(A, scale=scale)
    L = IntegerMatrix.from_matrix(L)
    L = LLL.reduction(L, flags=LLL.VERBOSE)
    M = GSO.Mat(L, float_type=float_type)
    bkz = BKZ2(M)
    t = 0.0
    param = BKZ.Param(block_size=beta,
                      strategies=BKZ.DEFAULT_STRATEGY,
                      auto_abort=True,
                      max_loops=16,
                      flags=BKZ.VERBOSE|BKZ.AUTO_ABORT|BKZ.MAX_LOOPS)

    bkz(param)
    t += bkz.stats.total_time

    H = copy(L)

    import pickle
    pickle.dump(L, open("L-%d-%d.sobj"%(L.nrows, beta), "wb"))

    E = []
    Y = set()
    V = set()
    y_i = vector(ZZ, tuple(L[0]))
    Y.add(tuple(y_i))
    E.append(apply_short1(y_i, A, c, scale=scale)[1])

    v = L[0].norm()
    v_ = v/sqrt(L.ncols)
    v_r = 3.2*sqrt(L.ncols - A.ncols())*v_/scale
    v_l = sqrt(h)*v_

    fmt = u"\t\t": %5.1fs, \log(\sigma)\": %5.1f, \log(|y|)\": %5.1f, \log(E[\sigma]):\%5.1f}"

    print
    print fmt%(t,
               log(abs(E[-1]), 2),
               log(L[0].norm(), 2),
               log(sqrt(v_r**2 + v_l**2), 2))

    print
    for i in range(m):
        t = cputime()
        M = GSO.Mat(L, float_type=float_type)
        bkz = BKZ2(M)
        t = cputime()
        bkz.randomize_block(0, L.nrows, stats=None, density=3)
        LLL.reduction(L)
        y_i = vector(ZZ, tuple(L[0]))
        l_n = L[0].norm()
        if L[0].norm() > H[0].norm():
            L = copy(H)
        t = cputime(t)

        Y.add(tuple(y_i))
        V.add(y_i.norm())
        E.append(apply_short1(y_i, A, c, scale=scale)[1])
        if len(V) >= 2:
            fmt = u"{\i}\": %4d, \t\t": %5.1fs, \log(|e_i|)\": %5.1f, \log(|y_i|)\": %5.1f,
            fmt += u"\log(\sigma)\": (%5.1f,%5.1f), \log(|y|)\": (%5.1f,%5.1f), |Y|: %5d}"
            print fmt%(i+2, t, log(abs(E[-1]), 2), log(l_n, 2), log_mean(E), log_var(E),
                          log_mean(V), log_var(V), len(Y))

    return E

```

C Alternative Cost Models

See Table 3.

Table 3. Costs of dual attacks on HELib and SEAL in the [LP11] cost model resp. assuming SVP in dimension β costs $2^{0.64\beta-28}$ operations as in [LN14] plugged into the estimator from [APS15]; cf. Table 1.

[LP11]						
n	1024	2048	4096	8192	16384	
SEAL 80-bit						
q	47.5	95.4	192.0	392.1	799.6	
dual	107.9	97.4	88.0	82.0	78.8	
small	80.5	81.1	78.2	76.4	75.9	
HELlib 80-bit						
q	47.0	87.0	167.0	326.0	638.0	
dual	111.5	112.4	111.5	111.2	111.2	
sparse	58.1	62.6	65.4	69.2	71.5	
HELlib 128-bit						
q	38.0	70.0	134.0	261.0	511.0	
dual	162.0	162.1	160.1	159.1	159.5	
sparse	76.3	81.9	85.8	86.2	90.3	
[LN14, APS15], 8 – 16 BKZ tours						
n	1024	2048	4096	8192	16384	
SEAL 80-bit						
q	47.5	95.4	192.0	392.1	799.6	
dual	101.2	91.7	83.1	78.3	76.1	
small	74.5	76.0	74.1	73.5	73.2	
HELlib 80-bit						
q	47.0	87.0	167.0	326.0	638.0	
dual	105.1	107.1	106.8	107.7	108.8	
sparse	54.1	59.1	62.8	65.8	68.9	
HELlib 128-bit						
q	38.0	70.0	134.0	261.0	511.0	
dual	158.4	159.8	158.6	158.3	160.0	
sparse	72.0	77.4	81.4	84.3	87.1	

References

- [ABD16] Albrecht, M., Bai, S., Ducas, L.: A subfield lattice attack on overstretched NTRU assumptions. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016. LNCS, vol. 9814, pp. 153–178. Springer, Heidelberg (2016). doi:[10.1007/978-3-662-53018-4_6](https://doi.org/10.1007/978-3-662-53018-4_6)

- [ACF+15] Albrecht, M.R., Cid, C., Faugère, J.-C., Fitzpatrick, R., Perret, L.: On the complexity of the BKW algorithm on LWE. *Des. Codes Crypt.* **74**, 325–354 (2015)
- [ACFP14] Albrecht, M.R., Cid, C., Faugère, J.-C., Perret, L.: Algebraic algorithms for LWE. *Cryptology ePrint Archive*, Report 2014/1018 (2014). <http://eprint.iacr.org/2014/1018>
- [ACPS09] Applebaum, B., Cash, D., Peikert, C., Sahai, A.: Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In: Halevi, S. (ed.) *CRYPTO 2009*. LNCS, vol. 5677, pp. 595–618. Springer, Heidelberg (2009). doi:[10.1007/978-3-642-03356-8_35](https://doi.org/10.1007/978-3-642-03356-8_35)
- [ADPS15] Alkim, E., Ducas, L., Pöppelmann, T., Schwabe, P.: Post-quantum key exchange - a new hope. *Cryptology ePrint Archive*, Report 2015/1092 (2015). <http://eprint.iacr.org/2015/1092>
- [ADPS16] Alkim, E., Ducas, L., Pöppelmann, T., Schwabe, P.: Post-quantum key exchange - a new hope. In: Holz, T., Savage, S. (eds.) *25th USENIX Security Symposium*, *USENIX Security*, vol. 16, Austin, TX, USA, 10–12 August 2016, pp. 327–343. USENIX Association (2016)
- [AFFP14] Albrecht, M.R., Faugère, J.-C., Fitzpatrick, R., Perret, L.: Lazy modulus switching for the bkW algorithm on LWE. In: Krawczyk, H. (ed.) *PKC 2014*. LNCS, vol. 8383, pp. 429–445. Springer, Heidelberg (2014). doi:[10.1007/978-3-642-54631-0_25](https://doi.org/10.1007/978-3-642-54631-0_25)
- [AFG14] Albrecht, M.R., Fitzpatrick, R., Göpfert, F.: On the efficacy of solving LWE by reduction to unique-SVP. In: Lee, H.-S., Han, D.-G. (eds.) *ICISC 2013*. LNCS, vol. 8565, pp. 293–310. Springer, Cham (2014). doi:[10.1007/978-3-319-12160-4_18](https://doi.org/10.1007/978-3-319-12160-4_18)
- [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). doi:[10.1007/978-3-642-22006-7_34](https://doi.org/10.1007/978-3-642-22006-7_34)
- [Ajt96] Ajtai, M.: Generating hard instances of lattice problems (extended abstract). In: *28th ACM STOC*, pp. 99–108. ACM Press, May 1996
- [APS15] Albrecht, M.R., Player, R., Scott, S.: On the concrete hardness of Learning with Errors. *J. Math. Cryptology* **9**(3), 169–203 (2015)
- [BCD+16] Bos, J.W., Costello, C., Ducas, L., Mironov, I., Naehrig, M., Nikolaenko, V., Raghunathan, A., Stebila, D.: Frodo: take off the ring! practical, quantum-secure key exchange from LWE. In: Weippl, E.R., Katzenbeisser, S., Kruegel, C., Myers, A.C., Halevi, S. (eds.) *ACM CCS 2016*, pp. 1006–1018. ACM Press, October 2016
- [BCNS15] Bos, J.W., Costello, C., Naehrig, M., Stebila, D.: Post-quantum key exchange for the TLS protocol from the ring learning with errors problem. In: *2015 IEEE Symposium on Security and Privacy*, pp. 553–570. IEEE Computer Society Press, May 2015
- [BDGL16] Becker, A., Ducas, L., Gama, N., Laarhoven, T.: New directions in nearest neighbor searching with applications to lattice sieving. In: Krauthgamer, R. (ed.) *27th SODA*, pp. 10–24. ACM-SIAM, January 2016
- [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). doi:[10.1007/978-3-319-08344-5_21](https://doi.org/10.1007/978-3-319-08344-5_21)
- [BGPW16] Buchmann, J., Göpfert, F., Player, R., Wunderer, T.: On the hardness of LWE with binary error: revisiting the hybrid lattice-reduction and meet-in-the-middle attack. In: Pointcheval, D., Nitaj, A., Rachidi, T.

- (eds.) AFRICACRYPT 2016. LNCS, vol. 9646, pp. 24–43. Springer, Cham (2016). doi:[10.1007/978-3-319-31517-1_2](https://doi.org/10.1007/978-3-319-31517-1_2)
- [BGV12] Brakerski, Z., Gentry, C., Vaikuntanathan, V.: (Leveled) fully homomorphic encryption without bootstrapping. In: Goldwasser, S. (ed.), ITCS 2012, pp. 309–325. ACM, January 2012
- [BKW00] Blum, A., Kalai, A., Wasserman, H.: Noise-tolerant learning, the parity problem, and the statistical query model. In: 32nd ACM STOC, pp. 435–440. ACM Press, May 2000
- [BLLN13] Bos, J.W., Lauter, K., Loftus, J., Naehrig, M.: Improved security for a ring-based fully homomorphic encryption scheme. In: Stam, M. (ed.) IMACC 2013. LNCS, vol. 8308, pp. 45–64. Springer, Heidelberg (2013). doi:[10.1007/978-3-642-45239-0_4](https://doi.org/10.1007/978-3-642-45239-0_4)
- [BLP+13] Brakerski, Z., Langlois, A., Peikert, C., Regev, O., Stehlé, D.: Classical hardness of learning with errors. In: Boneh, D., Roughgarden, T., Feigenbaum, J. (eds.) 45th ACM STOC, pp. 575–584. ACM Press, June 2013
- [Bra12] Brakerski, Z.: Fully homomorphic encryption without modulus switching from classical GapSVP. In: In Safavi-Naini and Canetti [SNC12], pp. 868–886
- [BV11] Brakerski, Z., Vaikuntanathan, V.: Efficient fully homomorphic encryption from (standard) LWE. In: Ostrovsky, R. (ed.) 52nd FOCS, pp. 97–106. IEEE Computer Society Press, October 2011
- [Che13] Chen, Y.: Réduction de réseau et sécurité concrète du chiffrement complètement homomorphe. PhD thesis, Paris 7 (2013)
- [CKH+16] Cheon, J.H., Han, K., Kim, J., Lee, C., Son, Y.: A practical post-quantum public-key cryptosystem based on spLWE. In: Hong, S., Park, J.H. (eds.) ICISC 2016. LNCS, vol. 10157, pp. 51–74. Springer, Cham (2017). doi:[10.1007/978-3-319-53177-9_3](https://doi.org/10.1007/978-3-319-53177-9_3)
- [CS15] Cheon, J.H., Stehlé, D.: Fully homomorphic encryption over the integers revisited. In: Oswald and Fischlin [OF15], pp. 513–536
- [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). doi:[10.1007/978-3-642-25385-0_1](https://doi.org/10.1007/978-3-642-25385-0_1)
- [CN12] Chen, Y., Nguyen, P.Q.: BKZ 2.0: Better lattice security estimates (full version) (2012). http://www.di.ens.fr/~ychen/research/Full_BKZ.pdf
- [CS16] Costache, A., Smart, N.P.: Which ring based somewhat homomorphic encryption scheme is best? In: Sako, K. (ed.) CT-RSA 2016. LNCS, vol. 9610, pp. 325–340. Springer, Cham (2016). doi:[10.1007/978-3-319-29485-8_19](https://doi.org/10.1007/978-3-319-29485-8_19)
- [DXL12] Ding, J., Xie, X., Lin, X.: A simple provably secure key exchange scheme based on the learning with errors problem. Cryptology ePrint Archive, Report 2012/688 (2012). <http://eprint.iacr.org/2012/688>
- [DTV15] Duc, A., Tramèr, F., Vaudenay, S.: Better algorithms for LWE and LWR. In: Oswald and Fischlin [OF15], pp. 173–202
- [FPL16] The FPLLL development team. FPLLL 5.0, a lattice reduction library (2016). <https://github.com/fplll/fplll>
- [FV12] Fan, J., Vercauteren, F.: Somewhat practical fully homomorphic encryption. Cryptology ePrint Archive, Report 2012/144 (2012). <http://eprint.iacr.org/2012/144>
- [GHS12a] Gentry, C., Halevi, S., Smart, N.P.: Homomorphic evaluation of the AES Circuit. In: Safavi-Naini and Canetti [SNC12], pages 850–867

- [GHS12b] Gentry, C., Halevi, S., Smart, N.P.: Homomorphic evaluation of the AES circuit. Cryptology ePrint Archive, Report 2012/099 (2012). <http://eprint.iacr.org/2012/099>
- [GJS15] Guo, Q., Johansson, T., Stankovski, P.: Coded-BKW: solving LWE using lattice codes. In: Gennaro and Robshaw [GR15], pp. 23–42
- [GR15] Gennaro, R., Robshaw, M. (eds.): CRYPTO 2015. LNCS, vol. 9215. Springer, Heidelberg (2015)
- [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). doi:[10.1007/978-3-642-40041-4_5](https://doi.org/10.1007/978-3-642-40041-4_5)
- [HG07] Howgrave-Graham, N.: A hybrid lattice-reduction and meet-in-the-middle attack against NTRU. In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 150–169. Springer, Heidelberg (2007). doi:[10.1007/978-3-540-74143-5_9](https://doi.org/10.1007/978-3-540-74143-5_9)
- [HPS11] Hanrot, G., Pujol, X., Stehlé, D.: Analyzing blockwise lattice algorithms using dynamical systems. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 447–464. Springer, Heidelberg (2011). doi:[10.1007/978-3-642-22792-9_25](https://doi.org/10.1007/978-3-642-22792-9_25)
- [HS14] Halevi, S., Shoup, V.: Algorithms in HELib. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014. LNCS, vol. 8616, pp. 554–571. Springer, Heidelberg (2014). doi:[10.1007/978-3-662-44371-2_31](https://doi.org/10.1007/978-3-662-44371-2_31)
- [KF15] Kirchner, P., Fouque, P.-A.: An improved BKW algorithm for LWE with applications to cryptography and lattices. In: Gennaro, R., Robshaw, M. (eds.) CRYPTO 2015. LNCS, vol. 9215, pp. 43–62. Springer, Heidelberg (2015). doi:[10.1007/978-3-662-47989-6_3](https://doi.org/10.1007/978-3-662-47989-6_3)
- [KF16] Kirchner, P., Fouque, P.-A.: Comparison between subfield and straight-forward attacks on NTRU. IACR Cryptology ePrint Archive, 2016: 717 (2016)
- [KL15] Kim, M., Lauter, K.: Private genome analysis through homomorphic encryption. BMC Med. Inform. Decis. Mak. **15**(5), 1–12 (2015)
- [Laa15] Laarhoven, T.: Sieving for shortest vectors in lattices using angular locality-sensitive hashing. In: Gennaro, R., Robshaw, M. (eds.) CRYPTO 2015. LNCS, vol. 9215, pp. 3–22. Springer, Heidelberg (2015). doi:[10.1007/978-3-662-47989-6_1](https://doi.org/10.1007/978-3-662-47989-6_1)
- [LCP16] Laine, K., Chen, H., Player, R.: Simple Encrypted Arithmetic Library - SEAL (v2.1). Technical report, Microsoft Research, MSR-TR-2016-68, September 2016
- [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). doi:[10.1007/978-3-642-36095-4_19](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). doi:[10.1007/978-3-319-06734-6_20](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). doi:[10.1007/978-3-642-19074-2_21](https://doi.org/10.1007/978-3-642-19074-2_21)
- [LP16] Laine, K., Player, R.: Simple Encrypted Arithmetic Library - SEAL (v2.0). Technical report, Microsoft Research, MSR-TR-2016-52, September 2016

- [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). doi:[10.1007/978-3-642-13190-5_1](https://doi.org/10.1007/978-3-642-13190-5_1)
- [LPR13] Lyubashevsky, V., Peikert, C., Regev, O.: A toolkit for ring-LWE cryptography. Cryptology ePrint Archive, Report 2013/293 (2013). <http://eprint.iacr.org/2013/293>
- [LTV12] López-Alt, A., Tromer, E., Vaikuntanathan, V.: On-the-fly multiparty computation on the cloud via multikey fully homomorphic encryption. In: Karloff, H.J., Pitassi, T. (eds.) 44th ACM STOC, pp. 1219–1234. ACM Press, May 2012
- [MR09] Micciancio, D., Regev, O.: Lattice-based cryptography. In: Bernstein, D.J., Buchmann, J., Dahmen, E. (eds.) Post-Quantum Cryptography, Heidelberg, New York, pp. 147–191 (2009)
- [OF15] Oswald, E., Fischlin, M. (eds.): EUROCRYPT 2015. LNCS, vol. 9056. Springer, Heidelberg (2015)
- [Pei09] Peikert, C.: Some recent progress in lattice-based cryptography. In: Reingold, O. (ed.) TCC 2009. LNCS, vol. 5444, pp. 72–72. Springer, Heidelberg (2009). doi:[10.1007/978-3-642-00457-5_5](https://doi.org/10.1007/978-3-642-00457-5_5)
- [Reg05] Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. In: Gabow, H.N., Fagin, R. (eds.) 37th ACM STOC, pp. 84–93. ACM Press, May 2005
- [Reg09] Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. J. ACM **56**(6), 1–40 (2009)
- [S+15] Stein, W., et al.: Sage Mathematics Software Version 7.1. The Sage Development Team (2015). <http://www.sagemath.org>
- [Sch03] Schnorr, C.P.: Lattice reduction by random sampling and birthday methods. In: Alt, H., Habib, M. (eds.) STACS 2003. LNCS, vol. 2607, pp. 145–156. Springer, Heidelberg (2003). doi:[10.1007/3-540-36494-3_14](https://doi.org/10.1007/3-540-36494-3_14)
- [Sho01] Shoup, V.: NTL: A library for doing number theory (2001). <http://www.shoup.net/ntl/>
- [SL12] Sarma, J., Lunawat, P.: IITM-CS6840: Advanced Complexity Theory – Lecture 11: Amplification Lemma (2012). <http://www.cse.iitm.ac.in/~jayalal/teaching/CS6840/2012/lecture11.pdf>
- [SNC12] Safavi-Naini, R., Canetti, R. (eds.): CRYPTO 2012. LNCS, vol. 7417. Springer, Heidelberg (2012)
- [Wal15] Walter, M.: Lattice point enumeration on block reduced bases. In: Lehmann, A., Wolf, S. (eds.) ICITS 2015. LNCS, vol. 9063, pp. 269–282. Springer, Cham (2015). doi:[10.1007/978-3-319-17470-9_16](https://doi.org/10.1007/978-3-319-17470-9_16)
- [Wun16] Wunderer, T.: Revisiting the hybrid attack: Improved analysis and refined security estimates. Cryptology ePrint Archive, Report 2016/733 (2016). <http://eprint.iacr.org/2016/733>