On the Efficacy of Solving LWE by Reduction to Unique-SVP

Martin R. Albrecht[1], Robert Fitzpatrick[2], and Florian Göpfert[3]

[1] Technical University of Denmark, Denmark
[2] ISG, Royal Holloway, University of London, Egham, United Kingdom
[3] CASED, TU Darmstadt, Darmstadt, Germany
maroa@dtu.dk, robert.fitzpatrick.2010@live.rhul.ac.uk, fgoepfert@cdc.informatik.tu-darmstadt.de

Abstract. We present a study of the concrete complexity of solving instances of the unique shortest vector problem (uSVP). In particular, we study the complexity of solving the Learning with Errors (LWE) problem by reducing the Bounded-Distance Decoding (BDD) problem to uSVP and attempting to solve such instances using the 'embedding' approach. We experimentally derive a model for the success of the approach, compare to alternative methods and demonstrate that for the LWE instances considered in this work, reducing to uSVP and solving via embedding compares favorably to other approaches.

1 Introduction

The Learning with Errors (LWE) problem is a generalisation to large moduli of the Learning Parity with Noise (LPN) problem. Since introduction by Regev [20], it has proved a remarkably flexible base for building cryptosystems. For example, Gentry, Peikert and Vaikuntanathan presented in [11] LWE-based constructions of identity-based encryption and many recent (fully) homomorphic encryption constructions are related to LWE [10, 7, 2]. Besides the flexibility of LWE, the main reason for the popularity of this problem is the convincing theoretical arguments underlying its hardness, namely a reduction from worst-case lattice problems such as GapSVP and SIVP to average-case LWE.

Definition 1 (LWE [20]). *Let n, q be positive integers, χ be a probability distribution on $\mathbb{Z}$ and $\boldsymbol{s}$ be a secret vector following the uniform distribution on $\mathbb{Z}_q^n$. We denote by $L_{\boldsymbol{s},\chi}$ the probability distribution on $\mathbb{Z}_q^n \times \mathbb{Z}_q$ obtained by choosing $\boldsymbol{a}$ from the uniform distribution on $\mathbb{Z}_q^n$, choosing $e \in \mathbb{Z}$ according to χ and returning $(\boldsymbol{a}, c) = (\boldsymbol{a}, \langle \boldsymbol{a}, \boldsymbol{s} \rangle + e) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$. The LWE problem is then, given a set of samples, to determine whether they originated from a $L_{\boldsymbol{s},\chi}$ oracle for some $\boldsymbol{s}$ or whether they follow the uniform distribution on $\mathbb{Z}_q^n \times \mathbb{Z}_q$.*

The modulus is typically taken to be polynomial in n and χ is the discrete Gaussian on $D_{\mathbb{Z},\alpha \cdot q}$ on $\mathbb{Z}$ with mean 0 and standard deviation $\sigma = \alpha \cdot q/\sqrt{2\pi}$ for some α. For these choices it was shown in [21, 6] that if $\alpha q > 2\sqrt{n}$ then (worst-case) GapSVP-$\tilde{\mathcal{O}}(n/\alpha)$ reduces to (average-case) LWE.

However, while the asymptotic hardness of LWE is well-understood, current understanding of the concrete hardness of solving particular instances of LWE leaves much to be desired. In this work, we examine the applicability of Kannan's embedding technique [13] to LWE and present the results of experiments using LLL and BKZ. While the embedding approach has been successfully employed in several past works, the approach remains somewhat mysterious with our current understanding of the efficacy of the approach being comparatively poor.

1.1 Related Work

In [16] Liu et. al. investigate similar questions, though their work lacks an experimental component which, in our opinion, form an indispensable part of any such work, given the current state of knowledge regarding the concrete complexity of unique-SVP. The current understanding of how a particular gap is related to the success of a particular reduction algorithm in disclosing a shortest vector, is poor. In [9] the results of a number of experiments were reported in which the authors examined the success of a number of algorithms in disclosing a shortest vector when (at least) a good approximation to the gap was known (though not in bounded-distance decoding/LWE cases). A simple model was proposed as a criterion for the success of a particular algorithm and particular class of lattices, with 'explaining the uSVP phenomenon' being posed as an open question.

1.2 Contribution & Organisation

We provide some background in Section 2 and discuss the embedding gap in Section 3.1. In Section 4 we apply the embedding approach to lattices derived from LWE instances and discuss the limits of this approach. Finally, in Section 5 we discuss the limits of the embedding approach and compare our results with results from the literature.

2 Background and Notation

2.1 Lattices and Discrete Gaussians

A *lattice* Λ in $\mathbb{R}^m$ is a discrete additive subgroup. We view a lattice as being generated by a (non-unique) basis $\mathbb{B} = \{\boldsymbol{b}_0, \ldots, \boldsymbol{b}_{n-1}\} \subset \mathbb{Z}^m$ of linearly-independent integer vectors:

$$\Lambda = \mathcal{L}(\mathbf{B}) = \mathbb{Z}^m \cdot \mathbf{B} = \left\{ \sum_{i=0}^{m-1} x_i \cdot \boldsymbol{b}_i : x_i \in \mathbb{Z} \right\}$$

The *rank* of the lattice $\mathcal{L}(\mathbf{B})$ is defined to be the rank of the matrix $\mathbf{B}$ with rows consisting of the basis vectors. If $m = n$ we say that $\mathcal{L}(\mathbf{B})$ is *full-rank*. We are only concerned with such lattices in this work and henceforth assume that the lattices we deal with are full-rank. In addition, in this work we are only concerned with q-ary lattices which are those such that $q\mathbb{Z}^m \subseteq \Lambda \subseteq \mathbb{Z}^m$. Note that every q-ary lattice is full-rank.

Given a lattice Λ we denote by $\lambda_i(\Lambda)$ the i-th minimum of Λ:

$$\lambda_i(\Lambda) := \inf \left\{ r \mid \dim(\operatorname{span}(\Lambda \cap \bar{\mathcal{B}}_m(\mathbf{0}, r))) \geq i \right\}$$

where $\bar{\mathcal{B}}_m(\mathbf{0}, r)$ denotes the closed, zero-centered m-dimensional (Euclidean) ball of radius r.

The determinant $\det(\mathcal{L})$ of a full-rank lattice is the absolute value of the determinant of any basis of the lattice.

Throughout, we work exclusively in the Euclidean norm unless otherwise stated and omit norm subscripts i.e. $\|\boldsymbol{x}\| = \|\boldsymbol{x}\|_2$. Given a point $\boldsymbol{t} \in \mathbb{R}^m$ and a lattice Λ, we define the minimum distance from $\boldsymbol{t}$ to the lattice by $\operatorname{dist}(\boldsymbol{t}, \Lambda) = \min \{\|\boldsymbol{t} - \boldsymbol{x}\| \mid \boldsymbol{x} \in \Lambda\}$.

The following are some computational problems on lattices which will be of relevance to our discussion

- ζ-Approx SVP (ζ-Approx-SVP), $\zeta \geq 1$: Given a lattice $\mathcal{L}$, find a vector $\boldsymbol{v} \in \mathcal{L}$ such that $0 < \|\boldsymbol{v}\| \leq \zeta \cdot \lambda_1(\mathcal{L})$
- κ-Hermite SVP (κ-HSVP), $\kappa \geq 1$: Given a lattice $\mathcal{L}$, find a vector $\boldsymbol{v} \in \mathcal{L}$ such that $0 < \|\boldsymbol{v}\| \leq \kappa \cdot \det(\mathcal{L})^{1/n}$
- η-Bounded Distance Decoding (BDD$_\eta$) ($\eta \leq 1/2$): Given a lattice $\mathcal{L}$ and a vector $\boldsymbol{t}$ such that $dist(\boldsymbol{t}, \mathcal{L}) < \eta\lambda_1(\mathcal{L})$, output the lattice vector $\boldsymbol{y}$ closest to $\boldsymbol{t}$
- ϱ-Unique Shortest Vector Problem (ϱ-uSVP): Given a lattice $\mathcal{L}$ such that $\lambda_2(\mathcal{L}) > \varrho \cdot \lambda_1(\mathcal{L})$, find the shortest non-zero lattice vector in $\mathcal{L}$

Now, if we have an algorithm which can solve Hermite-SVP with approximation factor κ, we can use this algorithm linearly many times [17] to solve Approx-SVP with approximation factor κ^2. Hence, we can use our κ-HSVP algorithm to solve uSVP instances in which the gap is at least κ^2. Similarly, if we have a Hermite root factor δ_0 characterising our Hermite-SVP algorithm, we can solve uSVP instances of gap δ_0^{2m}. However, one of the conclusions from [9], as we discuss later, is that, as with the gulf between the theoretical and practical performance of lattice reduction algorithms, we can generally solve uSVP instances with much smaller gap. More specifically, the results of [9] indicate that, while an exponential gap is still required to solve uSVP, the size of the gap only needs to grow on the order of the Hermite factor rather than its square, as indicated by the theoretical (worst-case) results. To the best of our knowledge, this behaviour remains unexplained and the practical performance of lattice reduction algorithms on lattices possessing a λ_2/λ_1 gap remains somewhat mysterious.

We always start counting from zero and denote vectors and matrices in lower-case and upper-case bold, respectively. We always assume that a lattice is generated by row combinations and that, when treating a collection of LWE samples as a 'matrix-LWE' sample, we assume this takes the form $\boldsymbol{b} = \mathbf{A}^T\boldsymbol{s} + \boldsymbol{e}$. Given a random variable X, E[X] denotes the expected value of X.

All experiments were carried out using the NTL implementation of BKZ and all LWE instances were generated using the LWE instance generator [3]. For simplicity, we assume throughout that enough LWE samples are exposed by a cryptosystem to allow the employment of the embedding technique in the optimal lattice dimension.

The discrete Gaussian distribution with parameter s, denoted $D_{\Lambda,s}$, over a lattice Λ is defined to be the probability distribution with support Λ which, for each $\boldsymbol{x} \in \Lambda$, assigns probability proportional to $\exp(-\pi\|\boldsymbol{x}\|^2/s^2)$. When we refer to the value s in this work with regard to LWE instantiations, we mean a discrete Gaussian with parameter s over the integer lattice. In an abuse of notation, we also use $D_{\Lambda,s}$ to denote a random variable following this distribution.

The following tail bound on discrete Gaussians from [5] is needed:

Lemma 1. *[[5]] Let $c \geq 1$ and $C = c \cdot \exp((1-c^2)/2) < 1$. Then for any real $s > 0$ and any integer $n \geq 1$ we have*

$$\Pr[\|D_{\mathbb{Z}^n,s}\| \geq c \cdot \frac{s\sqrt{n}}{\sqrt{2\pi}}] \leq C^n.$$

2.2 The Concrete Complexity of BKZ and BKZ 2.0

A central difficulty which arises in all works which require the use of 'strong' lattice reduction (by which we mean BKZ and improved variants) is the prediction of the concrete complexity of such algorithms. In [8] the authors present a study of 'BKZ 2.0', the amalgamation of three folklore techniques to improve the performance of BKZ: pruned enumeration; pre-processing of local blocks and early termination. While no implementations of such algorithms are publicly available, the authors of [8] present a simulator to predict the behaviour of out-of-reach BKZ computations. In [14] a model for the running time of BKZ 2.0 is proposed by running a limited set of experiments using the standard NTL implementation of BKZ and then adjusting the extrapolated running times by a certain factor to try and account for the improved running times promised by BKZ 2.0. The model arrived at is

$$\log_2 T_{sec} = 1.8/\log_2 \delta_0 - 110$$

In a recent work [15], the authors re-visit this model and compare the predictions to the BKZ 2.0 simulator of [8] in a few cases. In the cases examined in [15], the running-time predictions obtained by the use of the BKZ 2.0 simulator are quite close to those obtained by the model of Lindner and Peikert.

However, based on the data-points provided in [15] and converting these to the same metric as in the Lindner-Peikert model, the function

$$\log_2 T_{sec}^{\mathrm{BKZ2.0}} = 0.009/\log_2^2 \delta_0 - 27$$

provides a close approximation to the running-time output of the simulator for these particular cases.

This is a non-linear approximation and hence naturally grows faster than the approximation in [14]. However, given the greater sophistication of the latter 'BKZ 2.0' extrapolations derived from the simulator of [8], we expect this model to provide more accurate approximations of running times than the model of [14].

We note that a BKZ logarithmic running-time model which is non-linear in $\log_2 \delta_0$ appears more intuitive than a linear model. While, in practise, the root Hermite factors achievable through the use of BKZ with a particular blocksize β are much better than their best provable upper bounds, the root factor achievable appears to behave similarly to the upper bounds as a function of β. Namely, the best proven upper bounds on the root Hermite factor are of the form $\sqrt{\gamma_\beta}^{1/(\beta-1)}$, where γ_β denotes the best known upper bound on the Hermite constant for lattices of dimension β. Now since, asymptotically, γ_β grows linearly in β, if we assume that the root Hermite factor achievable in practise displays asymptotic behaviour similar to that of the best-known upper bound, then the root Hermite factor achievable as a function of β, denoted $\delta_0(\beta)$, is such that $\delta_0(\beta) \in \Omega(1/\beta)$. Since the running time of BKZ appears to be doubly-exponential in β, we can derive that $\log T_{sec}$ is non-linear in $1/\log \delta_0$, as is borne out by the results in [15]. In Section 4, we employ both models for completeness and comparison.

2.3 Alternative Algorithms for Solving LWE

Several previous works examine algorithms for solving LWE instances. The main methods for solving LWE consist of

1. Using combinatorial methods or lattice reduction to find a short (scaled) dual-lattice vector, permitting to distinguish LWE samples from uniform [19, 1]
2. Employing lattice reduction on the primal lattice then employing Babai's nearest-plane algorithm or a decoding variant thereof [14, 15]
3. Reduce the problem to noise-free non-linear polynomial system solving as proposed by Arora and Ge [4].

While, asymptotically, combinatorial methods for finding short (scaled) dual-lattice vectors are most efficient even for moderate parameter sizes [1], the exponential space requirements of these algorithms imply that lattice-based methods are more suitable for attacking practical instantiations of LWE. The algorithm due to Arora and Ge, at present, is largely of theoretical interest and impractical in its current form.

2.4 Concrete Hardness of uSVP

It is folklore that the presence of a significant gap between the first and second minima of a lattice makes finding the a shortest non-zero vector somewhat easier than would otherwise be the case, with an exponential gap allowing a shortest non-zero vector to be disclosed by application of LLL. However, in cases with sub-exponential gap, the success of lattice reduction algorithms in disclosing shortest non-zero vectors is poorly understood with a brief investigation in [9] being (to the best of our knowledge) the only practical investigation of such effects.

In [9] it was posited that given a lattice-reduction algorithm which we assume to be characterised by a root-Hermite factor δ_0 and a (full-rank) m-dimensional lattice Λ, the algorithm will be successful in disclosing a shortest non-zero vector with high probability when $\lambda_2(\Lambda)/\lambda_1(\Lambda) \geq \tau \cdot \delta_0^m$, where τ was taken to be a constant depending both on the nature of the lattices examined and also on the lattice reduction algorithm applied. In [9] values of τ ranging between 0.18 and 0.48 were experimentally-derived for various classes of lattices (though not LWE-derived lattices) and algorithms. However, the phrase 'with high probability' was not elaborated on in [9] and thus it is unclear as to whether a fixed threshold was used throughout the experiments in [9] or a variable threshold.

3 The Embedding Approach

In this section we outline and examine our application of Kannan's embedding technique, the resulting λ_2/λ_1-gap distributions and the resulting implications for the success of the approach.

3.1 Construction of Embedding Lattices

We consider the problem of being given an m-dimensional LWE-derived lattice basis $\mathbf{A}$ derived from LWE samples as follows. We take a sample matrix $\mathbf{A}' \in \mathbb{Z}_q^{n\times m}$ and calculate the reduced echelon form $\mathbf{A}'' \in \mathbb{Z}_q^{n\times m}$. For the right permutation matrix $\mathbf{P} \in \mathbb{Z}^{m \times m}$, we obtain the form $\mathbf{A}'' \cdot \mathbf{P} = \left(\mathbf{I}\,\overline{\mathbf{A}}\right)$ with $\mathbf{I} \in \mathbb{Z}_q^{n \times n}$ and $\overline{\mathbf{A}} \in \mathbb{Z}_q^{(m-n)\times(m-n)}$. If we interpret this matrix as a matrix over $\mathbb{Z}$, extend it with $q\mathbf{I} \in \mathbb{Z}^{(m-n)\times(m-n)}$ and define

$$\mathbf{A} = \begin{pmatrix} \mathbf{I} & \overline{\mathbf{A}} \\ \mathbf{0} & q\mathbf{I} \end{pmatrix} \mathbf{P}^{-1},$$

$\mathbf{A}$ is a basis of the lattice $\{\mathbf{v} \in \mathbb{Z}^m \mid \exists \mathbf{x} \in \mathbb{Z}_q n : \mathbf{x}\mathbf{A} = \mathbf{v} \mod q\}$. Now, given this $\mathbf{A}$ and a target vector $\boldsymbol{t} \in \mathbb{Z}_q^m$ and attempting to solve the LWE instance by reducing the embedding lattice basis

$$\mathbf{B}_{(\mathbf{A},\boldsymbol{t},t)} := \begin{pmatrix} \mathbf{A} & \mathbf{0} \\ \boldsymbol{t} & t \end{pmatrix}$$

where $t > 0$ is an embedding factor to be determined. We then define $\Lambda_e := \mathcal{L}(\mathbf{B})$. Note that, with overwhelming probability, $\det(\Lambda_{\boldsymbol{e}}) = t \cdot q^{m-n}$, i.e. $\mathbf{A}'$ has full rank over $\mathbb{Z}_q$.

It is well-known [18] that $1/(2\gamma)$-BDD can be reduced to solving γ-USVP by setting the embedding factor $t \geq \operatorname{dist}(\boldsymbol{t}, \mathcal{L}(\mathbf{A}))$. In practise, however, employing a smaller embedding factor generally allows us to create a unique-SVP instance with larger λ_2/λ_1 gap than by setting $t = \operatorname{dist}(\boldsymbol{t}, \mathcal{L}(\mathbf{A}))$. However, by setting $t < \operatorname{dist}(\boldsymbol{t}, \mathcal{L}(\mathbf{A}))$, with non-zero probability there exists a vector $\boldsymbol{v} \in \Lambda$ such that $\|\boldsymbol{v} + c \cdot [\boldsymbol{t} \quad t]\| < \|[\boldsymbol{e} \quad t]\|$ where $c \in \mathbb{Z}$ and, in general, if $t < \operatorname{dist}(\boldsymbol{t}, \mathcal{L}(\mathbf{A}))$, we will have $\lambda_2(\Lambda_e) < \lambda_1(\mathcal{L}(\mathbf{A}))$. Thus when we reduce t, quantification of the resulting λ_2/λ_1 gap becomes difficult.

To the best of our knowledge, no good model exists to determine the distribution of the lattice gap when taking an embedding factor smaller than $\|\boldsymbol{e}\|$ – an attempt is made in [16] but fails for reasons we outline in Appendix **??**. To attempt circumvention of such difficulties, we conduct experiments on LWE-derived uSVP lattices, examining firstly the λ_2/λ_1 gap required for success when we set $t = \lceil \operatorname{dist}(\boldsymbol{t}, \mathcal{L}(\mathbf{A})) \rceil$ (where we know λ_2/λ_1) and then for the case $t = 1$, under the assumption that the 'necessary gap' is unlikely to change, allowing us to derive analogous models.

3.2 On the Determination of τ when $t = \lceil \|\boldsymbol{e}\| \rceil$

As mentioned in 2.4, we employ the simple model of Gama and Nguyen for predicting the success of a particular basis-reduction algorithm in recovering a shortest non-zero vector, namely that there exist values of τ such that, for a given probability, basis-reduction algorithm and lattice class, the basis-reduction algorithm finds a shortest non-zero vector with probability greater or equal than the given probability over the random choice of lattices in the class whenever $\lambda_2(\Lambda)/\lambda_1(\Lambda) \geq \tau \cdot \delta_0^m$ where Λ represents a random choice of lattice in the class with dimension m. Thus, if we are able to sample such lattices randomly, determining a particular value of τ requires us to know (at least approximately) the λ_2/λ_1 gap of the lattices we are dealing with.

In the q-ary lattices we consider (i.e. lattices of the form $\mathcal{L}(\mathbf{A})$), unfortunately, there is no known good bound (in the Euclidean norm) on the first minimum when $m < 5n \log_2 q$. The case of $m \geq 5n \log_2 q$ is dealt with in [22]. In [16], the authors attempt to prove a probabilistic lower-bound for smaller m in the Euclidean norm, however the arguments employed in the attempted proof are inadequate – we revisit these in Appendix **??**. For the case of random lattices (in the sense of [12]), it is known that with overwhelming probability the minima of such an n-dimensional lattice are all asymptotically close to the Gaussian heuristic i.e.

$$\frac{\lambda_i(\Lambda)}{\operatorname{vol}(\Lambda)^{1/n}} \approx \frac{\Gamma(1 + n/2)^{1/n}}{\sqrt{\pi}}.$$

Now the q-ary lattices (e.g. $\mathcal{L}(\mathbf{A})$) widely employed in lattice-based cryptography are not random in this sense, being instead 'Ajtai' or LWE lattices, endowed with the worst-to-average-case properties. However, in all cases, it appears that the Gaussian heuristic appears to hold exceedingly well for such lattices (at least for the first minimum and with the added property that we always have vectors of norm q within the lattice), thus we assume throughout that the first minimum of such lattices is lower-bounded by the Gaussian heuristic with overwhelming probability.

For the first minimum of the embedded lattices, we only deal with this explicitly in the 'known-λ_1' case where we take this to be $\lambda_1(\Lambda_e) = \sqrt{2} \cdot \|\boldsymbol{e}\|$.

Then we can state the following lemma

Lemma 2. *Let* $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$*, let* $s > 0$ *and let* $c > 1$*. Let* $\boldsymbol{e}$ *be drawn from* $D_{\mathbb{Z}^m,s}$*. Under the assumption that* $\lambda_1(\Lambda(\mathbf{A})) \geq \mathrm{GH}_{q,n,m}$[1] *and that the rows of* $\mathbf{A}$ *are linearly-independent over* $\mathbb{Z}_q$*, we can create an embedding lattice* Λ_e *with* λ_2/λ_1*-gap greater than*

$$\frac{\min\left\{q, \frac{q^{1-\frac{n}{m}}\Gamma(1+\frac{m}{2})^{\frac{1}{m}}}{\sqrt{\pi}}\right\}}{\frac{cs\sqrt{m}}{\sqrt{\pi}}} \approx \frac{\min\left\{q, q^{1-\frac{n}{m}}\sqrt{\frac{m}{2\pi e}}\right\}}{\frac{cs\sqrt{m}}{\sqrt{\pi}}}$$

with probabillity greater than $1 - (c \cdot \exp((1-c^2)/2))^m$*.*

Proof. Omitted

We wish to obtain the value of m for which we can expect to gain the largest gaps (again, probabilistically).

Corollary 1. *Under the assumptions stated in Lemma 2 and for a fixed value of* c $(c > 1)$*, we can construct embedding lattices with the largest possible gap when*

$$q = \frac{q^{1-\frac{n}{m}}\Gamma(1+\frac{m}{2})^{\frac{1}{m}}}{\sqrt{\pi}}.$$

Proof. We assume that the approximation is close enough such that the maximum occurs for the same value of m. Consider the functions:

- $f_0(m) = \frac{\sqrt{\pi}q^{1-\frac{n}{m}}\sqrt{\frac{m}{2\pi e}}}{cs\sqrt{m}}$ where $c > 1$, $s > 0$.
- $f_1(m) = \frac{q\sqrt{\pi}}{cs\sqrt{m}}$ where $c, m > 1$ and $s > 0$.

Then $f_1(m)$ is clearly monotonically-decreasing and $f_0(m)$ has the form $f_0(m) = d \cdot q^{1-\frac{n}{m}}$, where d is a positive constant, hence is clearly monotonically-increasing under the conditions given. □

Thus, in our experiments, it appears valid to derive values of τ by assuming the Gaussian heuristic holds and that the (Euclidean) norm of the noise vector is equal to the expected value.

3.3 On the Determination of $\boldsymbol{\tau}$ when $\boldsymbol{t} < \lceil\|\boldsymbol{e}\|\rceil$

However, as mentioned, the employment of an embedding factor smaller than the norm of the noise vector $\boldsymbol{e}$ generally leads to a modest decrease in the size of the second minimum of the resulting lattice. In all cases observed, however, this decrease in the second minimum is less than the corresponding decrease in the first minimum (as a result of making the target vector shorter), leading to a more effective attack. However, quantification of the resulting gap is not simple - we know of no efficient method for determining the distribution of the λ_2/λ_1 gap under such conditions.

In an attempt to circumvent the lack of knowledge of the distribution of the λ_2/λ_1 gap when we take an embedding factor t such that $t < \|\boldsymbol{e}\|$, we assume that (for the same probabilistic success of a given basis-reduction algorithm) the same size of gap is required as in the case where we take $t = \lceil\|\boldsymbol{e}\|\rceil$ and then derive a modified value for τ. That is, we assume that the basis-reduction algorithm is in some sense oblivious to the embedding factor, with the size of the gap being the 'deciding' factor. While this is a somewhat arbitrary assumption, we believe it to be reasonable and intuitive. We denote the value of τ when $t = \lceil\|\boldsymbol{e}\|\rceil$ by $\tau_{\|\boldsymbol{e}\|}$ and the analogous value of τ when $t = 1$

[1] We employ the notation $\mathrm{GH}_{q,n,m}$ to denote the application of the Gaussian heuristic to an LWE lattice formed from m LWE samples of dimension n, with modulus q.

by τ_1. Given a particular value of n and knowing $\tau_{\|\boldsymbol{e}\|}$, we hence know (approximately) the gap required, denoted by $g_{\|\boldsymbol{e}\|}$ and hence a corresponding minimum lattice dimension which we denote by $m_{\|\boldsymbol{e}\|}$. Then, denoting by m_1 the minimum lattice dimension in the case $t = 1$ and assuming that the minimum required gap, denoted g_1 in the second case is the same, we can write

$$\tau_1 = \min\left\{\tau_{\|\boldsymbol{e}\|} \cdot \delta_0^{(m_{\|\boldsymbol{e}\|} - m_1)}, 1\right\}.$$

However, for easier and more intuitive comparison, we wish to express τ values for the case $t = 1$ when using the gaps from the $t = \lceil \|\boldsymbol{e}\| \rceil$ cases. For this comparison, we simply use the λ_2/λ_1 gaps from the case $t = \lceil \|\boldsymbol{e}\| \rceil$ and plug in the minimum dimension values from the case $t = 1$. We denote these 'illustrative' values of τ by τ'.

4 Application to LWE and comparisons

We now examine in more detail the model of [9] when applied to such unique-SVP instances. One difficulty with this model is that, while Gama and Nguyen state that success will occur with 'high probability', this probability is not explained. In the cases examined in this work, it appears to be often impossible to exceed a success probability and forms, in our opinion, an interesting subject for future work. For instance, Figure 1 demonstrates success probabilities for LLL for the case of Regev's parameterisation with $n \in \{35, 40, 45\}$ ($t = \|\boldsymbol{e}\|$) and increasing values of m, with between 50 and 100 cases being run for each value of m.

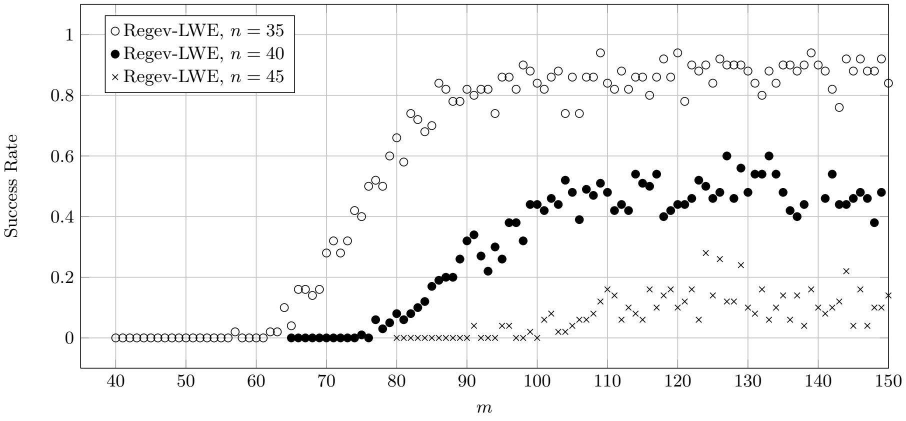

Fig. 1. Experimental Success Rates, Regev-LWE, LLL, $n \in \{35, 40, 45\}$, $t = \|\boldsymbol{e}\|$

We treat only the LWE parameterisations proposed by Regev [21] and Lindner/Peikert [14] and view each family of LWE instances as being parameterised by a value of n, from which values of s and q are derived. We then wish to examine the conditions under which applying the embedding approach yields a basis in which the target vector is present (though not necessarily the shortest vector in this reduced basis).

As in [9], our experiments indicate that the target vector lies in the reduced basis with some (fixed) probability whenever the gap is large enough such that

$$\frac{\lambda_2(\Lambda_m)}{\lambda_1(\Lambda_m)} \geq \tau \cdot \delta_0^m$$

where τ is some real constant such that $0 < \tau \leq 1$ depending on the desired probability level, the 'nature' of the lattices considered and the basis-reduction algorithm used. Our experiments proceed by fixing values of n to obtain corresponding LWE parameterisations then generating instances with increasing values of m – using [3] – until finding the minimum such value that recovery of the target vector is possible with the desired probability. We denote such values of m by $m_{\min}(n)$. In the $t = \lceil \|e\| \rceil$ case, plugging this value $m_{\min}(n)$ in $\frac{\lambda_2(\Lambda_m)}{\lambda_1(\Lambda_m)} = \tau \cdot \delta_0^m$ for m where we use Lemma 2 then recovers $\tau_{\|e\|}$. From this value and experimental data for $t = 1$ we can then derive

$\tau_1 = \min\left\{\tau_{\|\boldsymbol{e}\|} \cdot \delta_0^{(m_{\|\boldsymbol{e}\|} - m_1)}, 1\right\}$ and τ' by solving $\frac{\lambda_2(\Lambda_m)}{\lambda_1(\Lambda_m)} = \tau' \cdot \delta_0^{m_1}$. Throughout, the experimental data points indicate the minimum lattice dimension for which the lattice basis reduction algorithm succeeds in recovering the target vector with success rate 10%.

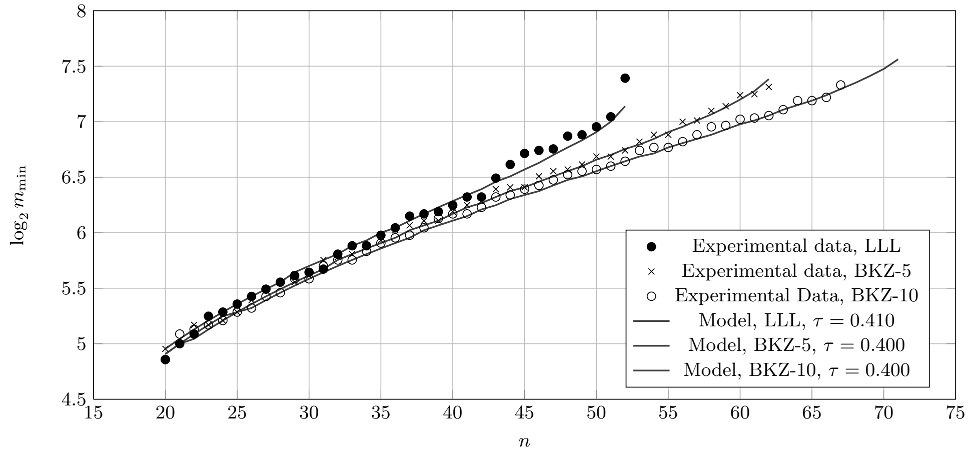

Fig. 2. Minimum lattice dimension, Regev-LWE, success rate 10%, $t = \|\boldsymbol{e}\|$.

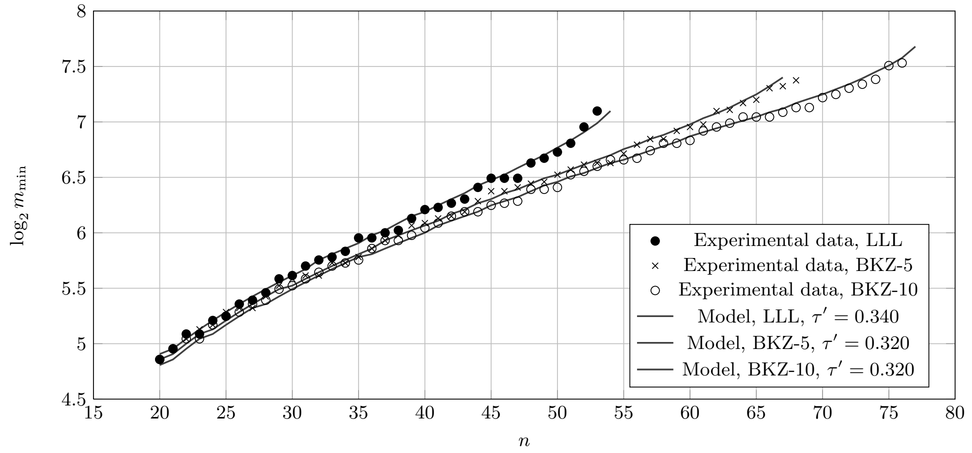

Fig. 3. Minimum lattice dimension, Regev-LWE, success rate 10%, $t = 1$.

In all experiments carried out, we artificially force that every $\|\boldsymbol{e}\|$ takes value $\approx \mathrm{E}[\|\boldsymbol{e}\|]$. This allows us to gain a good estimate of the λ_2/λ_1 gap in the $t = \|\boldsymbol{e}\|$ case. In addition, for the $m_{\min}$ calculations, we used the experimentally-derived root Hermite factors (see Appendix A) with linear interpolation.

4.1 Regev's Parameters

We firstly examine the case of Regev's original parameters as proposed in [20]. We take $q \approx n^2$ and set $\alpha = 1/(\sqrt{n} \cdot \log_2^2 n)$, $s = \alpha q$. Figure 2 illustrates the predicted feasible regions when $t = \lceil\|\boldsymbol{e}\|\rceil$. Similarly, Figure 3 gives analogous plots in the case $t = 1$, using the 'illustrative' values of τ' mentioned in Section 3.3. Figure 4 gives the $m_{\|\boldsymbol{e}\|}/m_1$ ratio for LLL and BKZ-5, illustrating the greater efficiency of using $t = 1$.

Based on the results as displayed above, we obtain parameters for embedding factors of $\lceil\|\boldsymbol{e}\|\rceil$ and 1, given in Table 1. We note that, while using an embedding factor $t = 1$ is most efficient, obtaining $\tau_1 > \tau_{\|\boldsymbol{e}\|}$ possibly seems

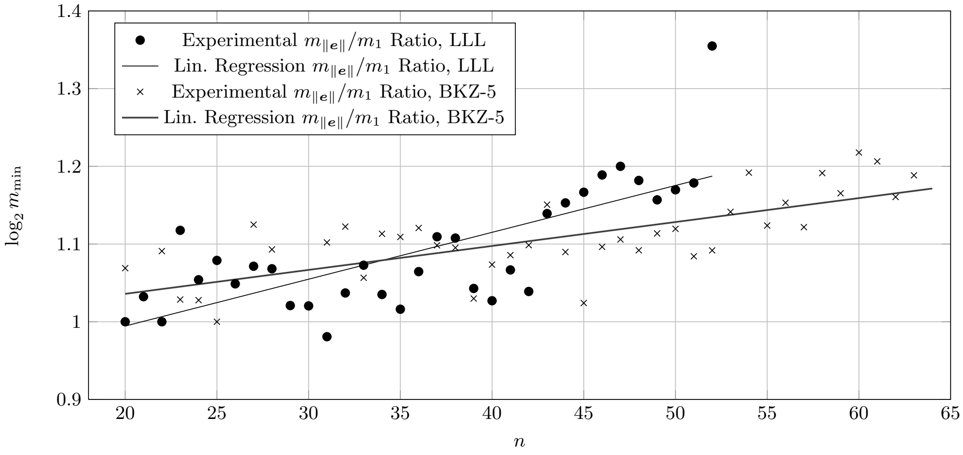

Fig. 4. $m_{\|e\|}/m_1$ Ratios, Regev-LWE, LLL and BKZ-5

counter-intuitive. However, the assumption of a fixed gap required for success while success occurs for a smaller value of m indeed leads to a larger value for τ_1.

	LLL	BKZ-5	BKZ-10
τ $(t = \|e\|)$	0.410	0.400	0.400
τ $(t = 1)$	0.467	0.464	0.444
τ' $(t = 1)$	0.340	0.320	0.320

Table 1. Parameters for finding e with success rate 10%, Regev's parameterisation.

4.2 Lindner and Peikert's Parameters

In [14], parameters for an improved LWE-based cryptosystem were proposed. For more details on this variant, the reader is refered to [14]. For our purposes, the principal difference from the Regev-LWE case is the smaller moduli employed by Lindner and Peikert. As in the Regev-LWE case, we choose a series of values for n and generate parameters accordingly, then apply LLL, BKZ-5 and BKZ-10 to solve such instances as far as is possible. Specifically, Table 2 gives a selection of the parameters considered as produced by [3].

n	20	30	40	50	60	70	80
q	2053	2053	2053	2053	2053	2053	2053
s	9.026	8.566	8.225	7.953	7.728	7.536	7.369

Table 2. Selected Lindner/Peikert LWE Parameters

We proceed similarly to the Regev-LWE case, with minimum lattice dimensions being given in Figure 5 for the $t = \|e\|$ case and in Figure 6 for the $t = 1$ case. Table 3 gives the derived values of τ for Lindner and Peikert's parameterisation.

We note that the values of τ derived seem consistent and do not vary widely between parameterisations. Of course, the value of τ may be expected to change when using 'stronger' algorithms than BKZ-10 or BKZ-20, however our limited experiments, and the results reported in [9] appear to indicate that the use of 'stronger' basis reduction algorithms leads to modest decreases in the values of τ. Thus, when we project these results in Section 5.1, we use the experimentally-derived τ values and thus expect the resulting complexity predictions to be somewhat conservative.

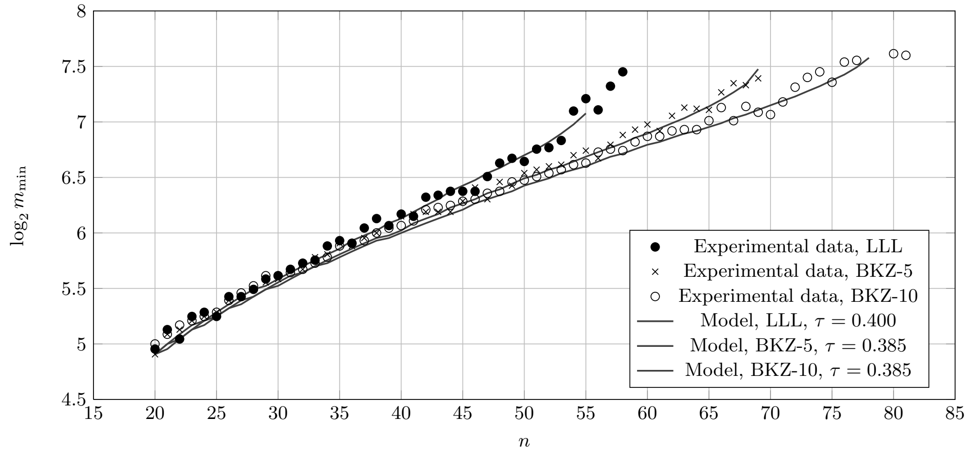

Fig. 5. Minimum lattice dimension, Lindner/Peikert Parameterisation, success rate 10%, $t = \|\boldsymbol{e}\|$

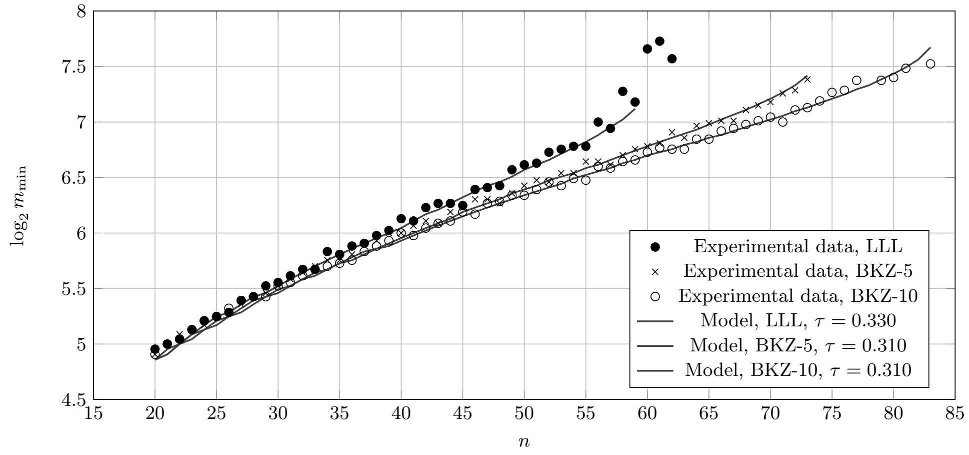

Fig. 6. Minimum lattice dimension, Lindner/Peikert Parameterisation, success rate 10%, $t = 1$

	LLL	BKZ-5	BKZ-10
τ $(t = \|\boldsymbol{e}\|)$	0.400	0.385	0.385
τ $(t = 1)$	0.435	0.431	0.439
τ' $(t = 1)$	0.330	0.310	0.310

Table 3. Parameters for finding $\boldsymbol{e}$ with success rate 10%, Lindner and Peikert's parameters

5 Limits of the Embedding Approach

Using the above model, we can derive an estimation of the limits of applicability of the embedding approach. Given a values (δ_0, τ), we can define the maximum value of n for which we can recover the target vector using the embedding approach to be

$$n_{\max} := \max\left\{ n \colon \exists m \quad \text{s.t.} \frac{\lambda_2(\Lambda_e(n,m))}{\lambda_1(\Lambda_e(n,m))} = \tau \cdot \delta_0^m \right\}$$

The goal is to determine the values of $n_{\max}$. Lemma 2 shows that we can construct a gap of size (under the assumption that we use basis-reduction algorithms with δ_0 small enough that $q^{1-(n/m)}\sqrt{m/(2\pi e)} < q$)

$$\frac{\lambda_2}{\lambda_1} \approx \frac{q^{1-\frac{n}{m}}\sqrt{\frac{1}{2e}}}{cs}.$$

If we want to solve an LWE instance with secret-dimension n, we have to find m such that

$$\frac{q^{1-\frac{n}{m}}\sqrt{\frac{1}{2e}}}{cs \cdot \tau \cdot \delta_0^m} \geq 1.$$

In order to determine the optimal m, we want to maximize the function

$$f_n(m) = \frac{q^{1-\frac{n}{m}}\sqrt{\frac{1}{2e}}}{c \cdot s \cdot \tau \cdot \delta_0^m}.$$

the first derivative of which is zero only when

$$\frac{n \log q}{m^2} = \log \delta_0,$$

and therefore $m = \sqrt{\frac{n \log q}{\log \delta_0}}$ is the optimal sub-dimension. In other words, we expect the attack to succeed if

$$\frac{q^{\left(1-\frac{n}{\sqrt{\frac{n \log q}{\log \delta_0}}}\right)}\sqrt{\frac{1}{2e}}}{c \cdot s \cdot \tau \cdot \delta_0^{\sqrt{\frac{n \log q}{\log \delta_0}}}} \geq 1$$

Thus, we only need to consider the optimal sub-dimension to ascertain whether we can expect the attack to succeed (with the given probability). Since, in our experiments we force $\|\boldsymbol{e}\| \approx \mathrm{E}[\|\boldsymbol{e}\|]$, we increase the value of c to cover all but the upper-tail of the distribution of $\|\boldsymbol{e}\|$. We can then state the following:

Assumption 1 *Given a fixed LWE parameterisation and a given value of τ (derived as above using $\|\boldsymbol{e}\| \approx \mathrm{E}[\|\boldsymbol{e}\|]$ instances and also corresponding to a fixed δ_0) corresponding to a fixed success rate p_s, we can solve general instances from the parameterisation with secret-dimension n with a particular value of m with probability*

$$p_c \geq p_s \cdot \left(1 - (c \cdot \exp((1-c^2)/2))^m\right) \tag{1}$$

if

$$\frac{q^{\left(1-\frac{n}{\sqrt{\frac{n \log q}{\log \delta_0}}}\right)}\sqrt{\frac{1}{2e}}}{c \cdot s \cdot \tau \cdot \delta_0^{\sqrt{\frac{n \log q}{\log \delta_0}}}} \geq 1 \tag{2}$$

We note that this assumption follows immediately from the above discussion and Lemma 1. Thus, given a target success probability, we attempt to satisfy conditions 1 and 2.

n	δ_0	$\log_2(\mathrm{sec}) = 1.8/\log_2 \delta_0 - 110$	$\log_2(\mathrm{sec}) = 0.009/\log_2^2 \delta_0 - 27$
64	1.0159	negl.	negl.
96	1.0111	negl.	negl.
128	1.0085	37.41	33.36
160	1.0069	71.44	64.45
192	1.0058	105.74	102.29
224	1.0050	140.16	146.83
256	1.0045	167.88	187.50
288	1.0040	202.54	244.37
320	1.0036	237.20	307.85

Table 4. Estimated cost of finding $\boldsymbol{e}$ with success rate 0.099, Regev's parameters.

5.1 Comparisons

We briefly compare the application of BKZ in both the embedding approach and the short dual-lattice vector distinguishing approach. For all embedding approach predictions, we take success probability slightly lower than 0.1, employing Assumption 1 - we choose c such that condition 1 holds for $p_c \geq 0.099$. While the dual-lattice distinguishing approach is not the best-known attack (the best practical attacks being that in [14] or modified versions [15]), it is easy to analyse in comparison to reduction-then-decode algorithms. We consider the application of BKZ in both situations. In the distinguishing approach, we can choose a desired distinguishing advantage ϵ and set $\gamma = q/s \cdot \sqrt{\ln(1/\epsilon)/\pi}$, from which we can compute a required Hermite root factor of $\delta_0 = 2^{\log_2^2(\gamma)/(4n\log_2 q)}$. So, for instance, with $n = 128$, we require $\delta_0 \approx 1.0077$ to gain a distinguishing advantage of ≈ 0.099, i.e. significantly worse than the 1.0085 required for the embedding attack. In Table 5 we give comparable estimated costs for distinguishing between LWE samples and uniformly random samples using the approach of Micciancio and Regev.

n	δ_0	$\log_2(\text{sec}) = 1.8/\log_2\delta_0 - 110$	$\log_2(\text{sec}) = 0.009/\log_2^2\delta_0 - 27$
64	1.0144	negl.	negl.
96	1.0099	negl.	negl.
128	1.0077	53.15	46.93
160	1.0063	89.99	84.10
192	1.0053	126.44	128.29
224	1.0046	162.56	179.35
256	1.0040	198.39	237.18
288	1.0036	234.00	301.70
320	1.0033	269.38	372.81

Table 5. Estimated cost of solving decision-LWE, advantage ~ 0.099, Regev's parameters, dual-lattice distinguisher

However, we note that the expression of Lindner and Peikert for the advantage of the dual-lattice distinguishing approach gives an upper-bound on the advantage obtained through the use of a specific algorithm. While the approximation is close overall, in the high-advantage regime the model is somewhat optimistic in estimating the advantage obtainable.

More rigourous comparison to the dual-lattice distinguishing attack is difficult, however, since the optimal strategy for said attack is to run a large number of low-advantage attacks and we can only analyse the embedding approach for high-advantages due to the (current) practical component of the analysis. We also note that if the embedding approach is used with $t = \|\boldsymbol{e}\|$ and fails, we can extract the resulting reduced basis of the lattice Λ and can then we proceed to run enumeration/decoding procedures, a strategy worthy of further investigation.

In conclusion, we provide evidence that the model of Gama and Nguyen is applicable to the solution of unique-SVP instances constructed from LWE instances and experimentally derive the constants which embody the performance of the approach. Based on the models used and assumptions made, we show that the embedding approach outperforms the dual-lattice distinguishing approach of Micciancio and Regev (in the high-advantage regime).

Open Questions We view a more in-depth comparison of the efficiency of the embedding technique and enumeration techniques as a pressing research question. The practical behaviour of lattice-reduction algorithms on unique-SVP instances remains mysterious, with (to the best of our knowledge) no recent progress in explaining the phenomena observed.

References

1. Martin R. Albrecht, Carlos Cid, Jean-Charles Faugère, Robert Fitzpatrick, and Ludovic Perret. On the complexity of the BKW algorithm on LWE. *Designs, Codes and Cryptography*, pages 1–30, 2013.
2. Martin R. Albrecht, Pooya Farshim, Jean-Charles Faugère, and Ludovic Perret. Polly Cracker, revisited. In *Advances in Cryptology – ASIACRYPT 2011*, volume 7073 of *Lecture Notes in Computer Science*, pages 179–196, Berlin, Heidelberg,

New York, 2011. Springer Verlag. full version available as Cryptology ePrint Archive, Report 2011/289, 2011 `http://eprint.iacr.org/`.

3. Martin R. Albrecht, Robert Fitzpatrick, Daniel Cabracas, Florian Göpfert, and Michael Schneider. A generator for LWE and Ring-LWE instances, 2013. available at `http://www.iacr.org/news/files/2013-04-29lwe-generator.pdf`.
4. Sanjeev Arora and Rong Ge. New algorithms for learning in presence of errors. In Luca Aceto, Monika Henzinger, and Jiri Sgall, editors, *ICALP*, volume 6755 of *Lecture Notes in Computer Science*, pages 403–415, Berlin, Heidelberg, New York, 2011. Springer Verlag.
5. Wojciech Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 296(1):625–635, 1993.
6. Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of Learning with Errors. to appear STOC 2013, 2013.
7. Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In Rafail Ostrovsky, editor, *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011*, pages 97–106. IEEE, 2011.
8. Yuanmi Chen and Phong Q. Nguyen. BKZ 2.0: better lattice security estimates. In *Advances in Cryptology - ASIACRYPT 2011*, volume 7073 of *Lecture Notes in Computer Science*, pages 1–20, Berlin, Heidelberg, 2011. Springer Verlag.
9. Nicolas Gama and Phong Q. Nguyen. Predicting lattice reduction. In Nigel P. Smart, editor, *EUROCRYPT*, volume 4965 of *Lecture Notes in Computer Science*, pages 31–51. Springer, 2008.
10. Craig Gentry. *A fully homomorphic encryption scheme*. PhD thesis, Stanford University, 2009. Available at `http://crypto.stanford.edu/craig`.
11. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC 08: Proceedings of the 40th annual ACM symposium on Theory of computing*, pages 197–206. ACM, 2008.
12. Daniel Goldstein and Andrew Mayer. On the equidistribution of Hecke points. *Forum Mathematicum*, 15:165–189, 2003.
13. Ravi Kannan. Minkowski's convex body theorem and integer programming. *Mathematics of operations research*, 12(3):415–440, 1987.
14. Richard Lindner and Chris Peikert. Better key sizes (and attacks) for LWE-based encryption. In *Topics in Cryptology – CT-RSA 2011*, volume 6558 of *Lecture Notes in Computer Science*, pages 319–339, Berlin, Heidelberg, New York, 2011. Springer Verlag.
15. Mingjie Liu and Phong Q. Nguyen. Solving BDD by enumeration: An update. In Ed Dawson, editor, *CT-RSA*, volume 7779 of *Lecture Notes in Computer Science*, pages 293–309. Springer, 2013.
16. Mingjie Liu, Xiaoyun Wang, Guangwu Xu, and Xuexin Zheng. Shortest lattice vectors in the presence of gaps. Cryptology ePrint Archive, Report 2011/139, 2011. `http://eprint.iacr.org/`. Last accessed 4th March 2012.
17. László Lovász. *An algorithmic theory of numbers, graphs, and convexity*. CBMS-NSF regional conference series in applied mathematics. Philadelphia, Pa. Society for Industrial and Applied Mathematics, 1986.
18. Vadim Lyubashevsky and Daniele Micciancio. On bounded distance decoding, unique shortest vectors, and the minimum distance problem. In Shai Halevi, editor, *CRYPTO*, volume 5677 of *Lecture Notes in Computer Science*, pages 577–594. Springer, 2009.
19. Daniele Micciancio and Oded Regev. Lattice-based cryptography. In Daniel J. Bernstein, Johannes Buchmann, and Erik Dahmen, editors, *Post-Quantum Cryptography*, pages 147–191. Springer Verlag, Berlin, Heidelberg, New York, 2009.
20. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *STOC*, pages 84–93. ACM, 2005.
21. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6), 2009.
22. Damien Stehlé, Ron Steinfeld, Keisuke Tanaka, and Keita Xagawa. Efficient public key encryption based on ideal lattices. In Mitsuru Matsui, editor, *Advances in Cryptology - ASIACRYPT 2009*, volume 5912 of *Lecture Notes in Computer Science*, pages 617–635, Berlin, Heidelberg, New York, 2009. Springer Verlag.

A Root-Hermite Factors for LWE-derived Lattices

It is a generally-accepted heuristic that the norms of shortest lattice vectors found by lattice basis reduction algorithms can be approximated by

$$\|\boldsymbol{b}_1\| \approx \det(\mathcal{L})^{1/m} \cdot \delta_0(m)^m$$

where $\delta_0(m)$ rapidly converges to a constant, denoted δ_0, as m grows. The following tables give experimentally-derived root-Hermite factors for LLL and some BKZ algorithms as applied to the LWE-derived lattices studied in this work – all root-Hermite factors being obtained for the minimum dimension in which the given algorithm solves the LWE-n instance with probability 0.1.

n	20	25	30	35	40	45	50
δ_0	1.0151	1.0169	1.0178	1.0182	1.0192	1.0204	1.0204

Table 6. Root Hermite Factors, LLL, Regev's Parameters

n	20	25	30	35	40	45	50	55	60
δ_0	1.0138	1.0146	1.0147	1.0147	1.0148	1.0157	1.0161	1.0159	1.0160

Table 7. Root Hermite Factors, BKZ-5, Regev's Parameters

n	20	25	30	35	40	45	50	55	60	65
δ_0	1.0121	1.0129	1.0136	1.0139	1.0138	1.0141	1.0145	1.0145	1.0146	1.0143

Table 8. Root Hermite Factors, BKZ-10, Regev's Parameters