



NTRU Fatigue: How Stretched is Overstretched?

Léo Ducas^(✉) and Wessel van Woerden^(✉)

Cryptology Group, CWI, Amsterdam, The Netherlands
`{leo.ducas, wvw}@cwi.nl`

Abstract. Until recently lattice reduction attacks on NTRU lattices were thought to behave similar as on (ring-)LWE lattices with the same parameters. However several works (Albrecht-Bai-Ducas 2016, Kirchner-Fouque 2017) showed a significant gap for large moduli q , the so-called overstretched regime of NTRU.

With the NTRU scheme being a finalist to the NIST PQC competition it is important to understand —both asymptotically and concretely— where the fatigue point lies exactly, *i.e.* at which q the overstretched regime begins. Unfortunately the analysis by Kirchner and Fouque is based on an impossibility argument, which only results in an asymptotic upper bound on the fatigue point. It also does not really *explain* how lattice reduction actually recovers secret-key information.

We propose a new analysis that asymptotically improves on that of Kirchner and Fouque, narrowing down the fatigue point for ternary NTRU from $q \leq n^{2.783+o(1)}$ to $q = n^{2.484+o(1)}$, and finally explaining the mechanism behind this phenomenon. We push this analysis further to a concrete one, settling the fatigue point at $q \approx 0.004 \cdot n^{2.484}$, and allowing precise hardness predictions in the overstretched regime. These predictions are backed by extensive experiments.

1 Introduction

1.1 Context

One should certainly recognize that in the field of lattice-based cryptography the NTRU cryptosystem of Hoffstein, Pipher and Silverman [HPS98, CDH+20] was particularly ahead of its time. After two decades spent basing cryptography [Ajt99, Reg05, SSTX09] on the worst-case hardness of lattice problems and concretising this theory into practical cryptosystems for standardisation [PAA+19, SAB+20, DKR+20], it is quite remarkable to see these constructions landing not so far away from the original design of NTRU (q -ary lattices, module structure over similar polynomial rings). In fact, it was even discovered a posteriori, that, up to the choice of parameters, the NTRU scheme itself can also be supported by worst-case hardness [SS11].

Regarding cryptanalysis, it was only recently discovered that the security of NTRU is in fact more subtle than the problem of finding a single unusually

short vector in a lattice. The first dent in this status quo came in 2016, from two concurrent works of Albrecht et al., and Cheon et al. [ABD16, C JL16], which exploits the specific algebraic structure of the NTRU lattice to improve upon pure lattice reduction attacks¹. This approach was shown to be applicable when the modulus q is large enough (say, super-polynomial), a regime coined “overstretched”.

Shortly thereafter Kirchner and Fouque [KF17] showed that this improved complexity does *not* require any algebraic structure, and is instead rooted in the purely geometrical fact that the NTRU lattice contains an unusually dense sublattice of large dimension, *i.e.* a sublattice of small determinant.² They also go further in their analysis, and conclude that moduli q as small as $n^{2.783+o(1)}$ already belong to the overstretched regime —for random ternary secrets. In particular, for q larger than this bound, the security of NTRU is significantly less than that of Learning With Errors [Reg04] and of its Ring variant [SSTX09, LPR13] using similar parameters.³

However, it is not so clear from the analysis of Kirchner and Fouque whether this asymptotic quantity $n^{2.783+o(1)}$ is an estimate or merely an upper bound on the *fatigue point*, that is the value of q separating the standard regime from the overstretched regime. Their analysis is based on a lemma of Pataki and Tural [PT08], that constraints the shape of lattice basis in terms of the volume of their sublattices. While it allows to conclude that the dense sublattice must be discovered after reducing the lattice basis beyond these constraints, it does not really explain *how* lattice reduction ends up discovering the dense sublattice, nor does it exclude that the discovery could happen earlier.

So far, it has been generally considered that only advanced schemes —requiring very large q — such as NTRU-based Homomorphic Encryption [BLLN13] or candidate cryptographic multi-linear maps [GGH13] could be affected by this overstretched regime. Yet, because the analysis of Kirchner and Fouque is only asymptotic, and because it may only provide an upper bound on the fatigue point, there is at the moment little documented evidence that the overstretched regime may not in fact extend further down, maybe down to the NTRU encryption scheme itself [HPS98, CDH+20]! Admittedly, this seems like a far fetched concern: asymptotically this scheme chooses $q = O(n)$, with a hidden constant between 4 and 5 in practice. However, this scheme being now a finalist of the NIST standardisation process for post-quantum cryptography, it appears rather imperious to refine our understanding of the phenomenon, and to finally close this pending question.

We found further motivation to go down this rabbit hole by measuring the concrete value of fatigue point experimentally. Until now, all documented

¹ Though the idea had been inconclusively considered already in 2002 by Gentry, Jonsson, Nguyen Stern and Szydlo as reported in [GS02, Sect. 6].

² Note that one may associate a short vector to a dense sublattice of dimension 1.

³ In fact, the presence of n rotations of the secret key already implies a minor security degradation compared to (Ring)-LWE already in the standard regime [MS01, DDGR20].

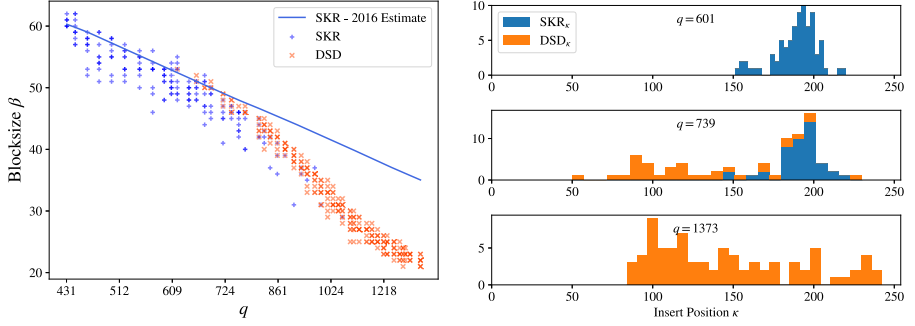


Fig. 1. Progressive BKZ with 8 tours per blocksize on matrix NTRU instances with parameters $n = 127, \sigma^2 = \frac{2}{3}$ for several moduli q . **Left:** the first blocksize β at which Progressive BKZ detects the Secret Key Recovery (SKR $_{\kappa}$) or Dense Sublattice Discovery (DSD $_{\kappa}$) event. We did 10 runs per modulus q . For the 2016-estimates, we use the geometric series assumption (GSA) for the shape of the basis and a probabilistic model for the discovery of the secret vector (see Sect. 2.4). **Right:** the positions κ at which a secret key or dense sublattice vector are detected over 80 runs per modulus.

experiments on the overstretched regime [ABD16, KF17, LW20] have focused on rather large values of q , and only used weak lattice reduction (LLL [LLL82], BKZ with blocksize 20): their goal was to demonstrate the claimed general behaviour when parameters are far in the overstretched regime. On the contrary, we focus our attention to the fatigue point for this preliminary experiment. That is, we ran strong reduction (progressive-BKZ [Sch87, AWHT16] up to blocksize 60) until a vector related to the secret key appeared for a range of moduli q . We distinguished the standard regime from the overstretched regime by classifying according to which event occurs first

- Secret Key Recovery (SKR $_{\kappa}$): a vector as short as a secret key vector is inserted in the basis at any given position κ .
- Dense Sublattice Discovery (DSD $_{\kappa}$): a vector strictly longer than the secret key but belonging to the dense sublattice generated by the secret key is inserted in the basis at any given position κ .

The result (Fig. 1) is rather striking: for $n = 127$, we start seeing a deviation from the standard regime for q as small as 700, while a naive interpretation of the prediction by Kirchner and Fouque [KF17] would suggest a fatigue point at $q \approx n^{2.783} \approx 700\,000$. We can conclude either that the asymptotic bound is not tight, or that the hidden asymptotic term (the $o(1)$ in $n^{2.783+o(1)}$) is significantly negative in practice. In any case, the bound of Kirchner and Fouque does not seem to provide accurate concrete predictions.

Remark. At this point, we should clarify why the DSD event should essentially be considered a successful attack. First, for q not too much larger than the fatigue point, an SKR event typically quickly follows after the DSD event; what happens

is that DSD events cascade, until the full dense sublattice has been extracted: the first half of the reduced basis precisely generates the dense sublattice. Lattice reduction will happen independently on each half of the basis, meaning that the dimension of the search space for the secret key has effectively been halved, and therefore making the problem much easier.

However, as q increases, DSD becomes easier and easier, to the point that it becomes even easier than secret key recovery within the dense sublattice. In other terms, there is a *superstretched* regime for larger q , where DSD does not directly lead to SKR.

Nevertheless, we argue —essentially rephrasing [ABD16]— that the DSD event is typically sufficient for an attack. First, the dense sublattice vector discovered is of length significantly lower than q ; in an FHE scheme such as [BLLN13] it is sufficient to decrypt fresh ciphertexts.⁴ Secondly, in the case of cyclotomic or circulant NTRU, it is possible to recover the secret key from the dense sublattice by other means than pure lattice reduction; in particular the recent line of work on the principal ideal-SVP [EHKS14, CDPR16, BEF+17] showed that this can be done classically in sub-exponential time $\exp(\tilde{O}(\sqrt{n}))$ and quantumly in polynomial time.

1.2 Our Work

Having identified precisely what event distinguishes the standard regime of NTRU from its overstretched regime, we may now proceed to a refined analysis, and determine precisely both the fatigue point and the precise cost⁵ of attacks in the overstretched regime. Our refined analysis diverges from the one of Kirchner and Fouque [KF17] on the following points:

1. we exploit the fact that BKZ runs SVP on large blocks ($\beta \geq 2$) not only to deduce the shape of the basis, but also to actually discover dense sublattice vectors,
2. we do not solely focus on the behaviour at position $\kappa = n - \beta + 1$ out of $d = 2n$ dimensions, but instead predict the most relevant position,
3. we propose an average-case analysis of volumes of the relevant lattices and sublattices, leading to a concrete prediction rather than a worst-case bound,
4. we also validate our intermediate and final predictions quantitatively with extensive experiments.

We note that contributions 1 and 2 alone already give us an important asymptotic result: the fatigue point of NTRU is indeed lower than predicted by Kirchner and Fouque, namely, it should happen at $q = n^{2.484+o(1)}$ instead of $n^{2.783+o(1)}$.

⁴ The secret key being shorter is only required to deal with ciphertexts obtained by homomorphic computation.

⁵ In this work, we only measure cost of lattice reduction in terms of the required BKZ blocksize; the computational cost of BKZ is essentially an orthogonal question.

Furthermore, our concrete average case analysis also differentiates the *circulant* version of NTRU [HPS98] from its *matrix* version [CG05,GGH+19]. We note minor deviations in the concrete analysis of volumes of relevant sublattices, that on average slightly favours the attacker in the matrix case, but also shows a larger variance in the concrete hardness of the circulant case.

In summary: we achieve an explicative and predictive model for the fatigue of NTRU, with concrete predictions confirmed in practice. In particular, the fatigue point is estimated to be at $q \approx 0.004 \cdot n^{2.484}$ for $n > 100$. All our artefacts for experiments and predictions are open-source and can be accessed at <https://github.com/WvanWoerden/NTRUFatigue>. These are based on the FPLLL and FPyLLL libraries [dt21a,dt21b].

Impact. We wish to clarify that this work *does not* contradict the concrete security of the NTRU candidate to the NIST competition [CDH+20]; on the contrary, we close a pending question regarding a potential vulnerability.

Limitation: the Lucky-Lifts. During our experiments, we also noted rare occurrence of DSD events that qualitatively differ from what we expected. Namely, the vector from the dense sublattice was found at positions κ quite larger than what was predicted by our model. More remarkable, these vectors were extremely unbalanced: their $2n - \kappa$ last (Gram-Schmidt) coordinates were much smaller than the κ first coordinates. We call these DSD events lucky-lifts (DSD-LL), while the one we model and mostly observe are called after the Pataki-Tural Lemma (DSD-PT). Despite those two phenomena being very distinct, they nevertheless occurred for the same BKZ blocksize β , at least in the range of parameters we could experiment with.

It could very well be that these rare DSD-LL events are just artefacts of the modest parameters of our experiments and that these events vanish as the dimension grows. Yet, as they seem of a very different nature, a definitive conclusion would require a dedicated study.

1.3 Organisation

We introduce some preliminaries, the NTRU lattice, and the state-of-the-art estimates in Sect. 2. In Sect. 3 we introduce our new DSD-PT estimate and give an asymptotic analysis. In Sect. 4 we give an average-case analysis to construct a concrete estimator. In the final Sect. 5 we compare our estimate with experiments.

2 Preliminaries

2.1 Notation and Distributions

All vectors and matrices are denoted by bold lower and upper case letters respectively. All vectors are column-vectors and we write $\mathbf{B} = [\mathbf{b}_0, \dots, \mathbf{b}_{n-1}]$ for a matrix where the i -th column vector is \mathbf{b}_i . If a matrix $\mathbf{B} \in \mathbb{R}^{d \times n}$ has full rank n we denote by $\mathcal{L}(\mathbf{B}) := \{\mathbf{B}\mathbf{x} : \mathbf{x} \in \mathbb{Z}^n\}$ the lattice spanned by the columns of \mathbf{B} . We call a lattice vector $\mathbf{v} \in \mathcal{L}$ primitive if it is not a strict integer multiple of another lattice vector. For a basis \mathbf{B} and $i \in \{0, \dots, d-1\}$ we define π_i as the orthogonal projection away from $\mathbf{b}_0, \dots, \mathbf{b}_{i-1}$, and the Gram-Schmidt vectors as $\mathbf{b}_0^*, \dots, \mathbf{b}_{d-1}^*$ where $\mathbf{b}_i^* := \pi_i(\mathbf{b}_i)$. We write $\mathbf{B}_{[l:r]}$ for the matrix $[\pi_l(\mathbf{b}_l), \dots, \pi_l(\mathbf{b}_{r-1})]$, and denote the projected⁶ sublattice $\mathcal{L}(\mathbf{B}_{[l:r]})$ as $\mathcal{L}_{[l:r]}$ when the basis is clear from the context. We denote the Euclidean norm of a vector \mathbf{v} by $\|\mathbf{v}\|$ and the volume of a lattice by $\text{vol}(\mathcal{L}(\mathbf{B})) := \prod_{i=0}^{n-1} \|\mathbf{b}_i^*\|$. We write $\lambda_1(\mathcal{L}) := \min_{\mathbf{v} \in \mathcal{L} \setminus \{0\}} \|\mathbf{v}\|$ for the first minimum of a lattice \mathcal{L} . For a lattice \mathcal{L}

we denote the dual lattice as $\mathcal{L}^* := \{\mathbf{w} \in \text{span}(\mathcal{L}) : \langle \mathbf{w}, \mathbf{v} \rangle \in \mathbb{Z} \text{ for all } \mathbf{v} \in \mathcal{L}\}$. We use ‘claim’ to refer to an informal statement based on heuristics.

We denote the continuous centered Gaussian (normal) distribution with variance σ^2 by χ_{σ^2} . We denote the unit sphere over k coordinates as \mathcal{S}^{k-1} and call the uniform distribution over \mathcal{S}^{k-1} the spherical distribution. We write \mathcal{B}_1^d for the d -dimensional unit ball. We write the chi-square distribution with k degrees of freedom as $\chi_{k, \sigma^2}^2 := \sum_{i=1}^k X_i^2$, where X_1, \dots, X_k are independently distributed as χ_{σ^2} . The chi-square distribution has expectation $k\sigma^2$, but for our concrete estimates we consider the log-expectation.

Lemma 2.1. *Let X be distributed as χ_{k, σ^2}^2 , then*

$$\mathbb{E}[\ln(X)] = \ln(2\sigma^2) + \psi(k/2),$$

where $\psi(x) := \Gamma'(x)/\Gamma(x)$ is the digamma function.

2.2 NTRU and Lattice Attacks

We start with the historical definition of NTRU.

Definition 2.2 (NTRU). *Let n be prime, q a positive integer and let $\mathbf{f}, \mathbf{g} \in (\mathbb{Z}/q\mathbb{Z})[X]$ be polynomials of degree n with small coefficients sampled from some distribution χ under the condition that \mathbf{f} is invertible in $\mathcal{R}_q := (\mathbb{Z}/q\mathbb{Z})[X]/(X^n - 1)$. The pair (\mathbf{f}, \mathbf{g}) forms the secret key, and the public key is defined as $\mathbf{h} := \mathbf{g}/\mathbf{f} \bmod \mathcal{R}_q$. The NTRU problem is to recover any rotation $(X^i \mathbf{f}, X^i \mathbf{g})$ of the secret key from \mathbf{h} .*

For *NTRUencrypt* [HPS98, CDH+20] \mathbf{f} and \mathbf{g} have ternary coefficients, with a fixed number of about $n/3$ of each value in $\{-1, 0, 1\}$. For our analysis we

⁶ When $l = 0$, no projection is applied, and $\mathcal{L}_{[0:r]}$ is simply a sublattice of \mathcal{L} .

consider the case where each coefficient is sampled from a discrete Gaussian over \mathbb{Z} with some variance $\sigma^2 > 0$. For simplicity the ternary case is treated as a discrete Gaussian with variance $\sigma^2 = \frac{2}{3}$.

More generally we consider a matrix description of NTRU where the polynomials are replaced by matrices $\mathbf{F}, \mathbf{G}, \mathbf{H} \in \mathbb{Z}^{n \times n}$ such that $\mathbf{H} := \mathbf{G} \cdot \mathbf{F}^{-1} \bmod q$ [CG05, GGH+19]. Variants of NTRU, e.g. based on different algebraic rings [BBC+20], can be encoded in the structure of the matrices. For example, the original problem can be encoded by setting $\mathbf{F}_{i,j} := f_{(i+j \bmod n)}$ where $\mathbf{f} = \sum_{i=0}^{n-1} f_i X^i$, for each polynomial respectively. We call the original variant *circulant NTRU*, based on the resulting shape of the matrices \mathbf{F}, \mathbf{G} , and we treat \mathbf{f}, \mathbf{g} as n -dimensional vectors. We also consider the variant, called *matrix NTRU*, where the matrices \mathbf{F}, \mathbf{G} have no extra structure and the coefficients are independently sampled from a discrete Gaussian.

To reduce the NTRU problem to a lattice problem we define the *NTRU lattice*, which contains a particularly *dense* sublattice generated by the secret key.

Definition 2.3. Let $(n, q, \mathbf{F}, \mathbf{G}, \mathbf{H})$ be an NTRU instance. We define the NTRU lattice as

$$\mathcal{L}^{\mathbf{H},q} := \begin{pmatrix} q\mathbf{I}_n & \mathbf{H} \\ \mathbf{0} & \mathbf{I}_n \end{pmatrix} \cdot \mathbb{Z}^{2n},$$

and its (secret) dense sublattice of rank n by:

$$\mathcal{L}^{\mathbf{G}\mathbf{F}} := \mathbf{B}^{\mathbf{G}\mathbf{F}} \cdot \mathbb{Z}^n \subset \mathcal{L}^{\mathbf{H},q}, \text{ where } \mathbf{B}^{\mathbf{G}\mathbf{F}} := \begin{pmatrix} \mathbf{G} \\ \mathbf{F} \end{pmatrix}.$$

Solving the NTRU problem is equivalent to recovering the dense sublattice basis $\mathbf{B}^{\mathbf{G}\mathbf{F}} = [\mathbf{G}; \mathbf{F}]$ up to some permutation of the columns. For uniformity of notation we will denote such a column by $(\mathbf{g}|\mathbf{f})$. These column vectors have a length of about $\|(\mathbf{g}|\mathbf{f})\| \approx \sqrt{2n\sigma^2}$, which for common parameters is much shorter than the expected minimal length $\lambda_1(\mathcal{L}^{\mathbf{H},q}) \approx \sqrt{nq/(\pi e)}$ of the full lattice $\mathcal{L}^{\mathbf{H},q}$ for a truly uniform random $\mathbf{H} \in (\mathbb{Z}/q\mathbb{Z})^{n \times n}$. To recover the secret key we thus have to find these exceptionally short vectors in the full lattice $\mathcal{L}^{\mathbf{H},q}$.

In [CS97] Coppersmith and Shamir showed that we can slightly relax the problem as any small vector from the dense sublattice $\mathcal{L}^{\mathbf{G}\mathbf{F}}$ is enough to decode a message. We therefore focus our analysis on the recovery of elements from $\mathcal{L}^{\mathbf{G}\mathbf{F}}$, and not (directly) on the full secret basis $\mathbf{B}^{\mathbf{G}\mathbf{F}}$. To recover short vectors we resort to lattice reduction.

2.3 Lattice Reduction

Any lattice $\mathcal{L} = \mathcal{L}(\mathbf{B})$ with basis $\mathbf{B} \in \mathbb{R}^{d \times d}$ has (for $d > 1$) an infinite number of other bases $\mathbf{B} \cdot \mathbf{U}$ with $\mathbf{U} \in \text{GL}_d(\mathbb{Z})$. The goal of lattice reduction is to find a *good* basis: the basis vectors are preferably short and somewhat orthogonal. Looking at the Gram-Schmidt vectors $\mathbf{b}_0^*, \dots, \mathbf{b}_{d-1}^*$ we have the invariant $\prod_{i=0}^{d-1} \|\mathbf{b}_i^*\| =$

$\det(\mathbf{B}) = \text{vol}(\mathcal{L})$ which is independent of the basis. Therefore decreasing the length of the first basis vector $\mathbf{b}_0 = \mathbf{b}_0^*$ forces some of the other Gram-Schmidt vectors to increase in length. We call these lengths $(\|\mathbf{b}_i^*\|)_{i=0,\dots,d-1}$ the *profile* of a basis \mathbf{B} . A good basis has a well balanced profile; in particular one that does not decrease too fast.

The most famous lattice reduction algorithm is the polynomial time LLL algorithm, which gives some guarantees on the slope of an LLL-reduced basis. We consider a generalisation, namely the BKZ algorithm, that gives a flatter slope, but at a higher cost. A basis is BKZ reduced with blocksize β if \mathbf{b}_κ^* is a shortest vector of the projected sublattice $\mathcal{L}_{[\kappa:\min(\kappa+\beta,d))}$ at each position κ . LLL-reduction corresponds to the case that $\beta = 2$.

Definition 2.4 (BKZ). A basis $\mathbf{B} = [\mathbf{b}_0, \dots, \mathbf{b}_{d-1}]$ is called BKZ- β reduced if

$$\|\mathbf{b}_\kappa^*\| = \lambda_1(\mathcal{L}_{[\kappa:\min(\kappa+\beta,d))}) \text{ for all } \kappa = 0, \dots, d-1.$$

A BKZ-reduced basis has several provable bounds on the slope of the profile. In the context of cryptanalysis we are more interested in the average-case behaviour and thus we fall back on heuristics to describe the shape of a BKZ-reduced profile. The most commonly used heuristic for lattices is the Gaussian Heuristic, that states that for a measurable volume \mathcal{V} the number of lattice points $|\mathcal{L} \cap \mathcal{V}|$ approximately equals $\text{vol}(\mathcal{V})/\text{vol}(\mathcal{L})$. Applying this to a ball allows to estimate the first minimum of a lattice.

Heuristic 2.5. Let \mathcal{L} be a d -dimensional lattice with volume $\text{vol}(\mathcal{L})$. The expectation of the first minimum $\lambda_1(\mathcal{L})$ under the Gaussian Heuristic is given by

$$\text{gh}(\mathcal{L}) := \frac{\text{vol}(\mathcal{L})^{1/d}}{\text{vol}(\mathcal{B}_1)^{1/d}} \approx \sqrt{d/(2\pi e)} \cdot \text{vol}(\mathcal{L})^{1/d}.$$

We also denote $\text{gh}(d) \approx \sqrt{d/(2\pi e)}$ for the expected first minimum of a d -dimensional lattice with volume 1.

Applying the above heuristic to the value of $\|\mathbf{b}_\kappa^*\| = \lambda_1(\mathcal{L}_{[\kappa:\min(\kappa+\beta,d))})$ at each position κ gives us relations between the Gram-Schmidt lengths $\|\mathbf{b}_0^*\|, \dots, \|\mathbf{b}_{d-1}^*\|$. Solving these relations for $\beta \ll d$ shows that $\|\mathbf{b}_\kappa^*\| / \|\mathbf{b}_{\kappa+1}^*\| \approx \alpha_\beta$ for some constant α_β only depending on β . So heuristically the profile forms a geometric series. This is made more precise by the Geometric Series Assumption.

Heuristic 2.6 (Geometric Series Assumption (GSA)). Let \mathbf{B} be a BKZ- β reduced basis, then the profile satisfies

$$\ln(\|\mathbf{b}_i^*\|) = \frac{d-1-2i}{2} \cdot \ln(\alpha_\beta) + \frac{\ln(\det(\mathbf{B}))}{d},$$

where $\alpha_\beta = \text{gh}(\beta)^{2/(\beta-1)}$.

The GSA is reasonably precise for say $\beta \geq 50$ and a not too large blocksize $\beta \ll d$ compared to the lattice dimension.

The BKZ algorithm (see Algorithm 1) computes a BKZ-reduced basis from any other basis. The algorithm greedily attempts to satisfy the BKZ condition at each position by computing a shortest vector in each *block* $\mathbf{B}_{[\kappa:\min(\kappa+\beta, d)]}$, and replacing the basis vector \mathbf{b}_κ accordingly. This makes the basis BKZ- β reduced at position κ , but might invalidate the condition at other positions. Applying this once to all positions $\kappa = 0, \dots, d - 2$ is called a *tour*. The BKZ algorithms repeats such tours until the basis remains unchanged and is thus BKZ-reduced.

Algorithm 1: The BKZ algorithm.

Data: A lattice basis \mathbf{B} , blocksize β .

while \mathbf{B} is not BKZ- β reduced **do**

for $\kappa = 0, \dots, d - 2$ **do** // A single BKZ- β tour

$\mathbf{w} \leftarrow$ a shortest vector in $\mathcal{L}(\mathbf{B}_{[\kappa:\min(\kappa+\beta, d)]})$;

Lift \mathbf{w} to a full vector $\mathbf{v} \in \mathcal{L}(\mathbf{B}_{[0:\min(\kappa+\beta, d)]})$ s.t. $\pi_\kappa(\mathbf{v}) = \mathbf{w}$;

Insert \mathbf{v} in \mathbf{B} at position κ and use LLL to resolve linear dependencies;

The number of tours is polynomially bounded, and in practice not much improvement is attained after say a few dozen tours. The cost of BKZ is thus mainly dominated by the exponential (in β) cost of finding a shortest vector in a β -dimensional lattice. *Progressive* BKZ reduces this cost in practice, where instead of running many tours of BKZ- β , one runs only a few tours for increasing $\beta' = 2, 3, \dots, \beta$.

For our experiments we also added a hook to BKZ, using secret key information, to detect if a vector \mathbf{v} is part of the dense sublattice \mathcal{L}^{GF} and to abort early if this is the case.

While the Geometric Series Assumption gives a good first order estimate of the basis profile after BKZ-reduction, it is known to be inaccurate in small dimensions or when the dimension is only a small multiple of the blocksize. Additionally it does not account for the slower convergence when running progressive BKZ with only a few tours. To resolve this problem [CN11] introduced a BKZ simulator based on the Gaussian Heuristic, that was later refined in [YD17, BSW18]. These allow for accurate and efficient predictions of the profile shape for *random* lattices, even for progressive BKZ with a limited number of tours.

Behaviour on q -ary lattices. While by now the behaviour of BKZ on random lattices is reasonably understood, this is less the case for q -ary lattices (for certain parameters) such as the NTRU lattice $\mathcal{L}^{H,q}$.

Definition 2.7 (q-ary lattices). A lattice \mathcal{L} of dimension d is said to be q -ary if for some $q > 0$ we have

$$q\mathbb{Z}^d \subset \mathcal{L} \subset \mathbb{Z}^d.$$

Note that the first n basis vectors of $\mathcal{L}^{\mathbf{H},q}$ are orthogonal q -vectors $(q, 0, \dots, 0)$, $(0, q, 0, \dots, 0), \dots$, and so the initial basis profile starts with $\|\mathbf{b}_0^*\| = \dots = \|\mathbf{b}_{n-1}^*\| = q$. Additionally after projecting away from these q -vectors, the remaining basis vectors are again orthogonal with length 1, and thus we have $\|\mathbf{b}_n^*\| = \dots = \|\mathbf{b}_{d-1}^*\| = 1$. Note that in the BKZ algorithm the length of \mathbf{b}_0 can not increase, and is thus always at most q . Also \mathbf{b}_1 can not increase in length if \mathbf{b}_0 remains unchanged, and so on. For dual-BKZ or the self-dual LLL the profile lengths can not drop below 1 anywhere by the same reasoning. Still LLL and BKZ guarantee that the profile slope in the middle is not too steep. So after LLL reduction the profile must be flat at the start and end, and have a sloped part in the middle, we call this a Z-shape [AD21]. Because BKZ is not self-dual we do not have any guarantee that the last profile elements do not drop below 1, however we could for example run BKZ only on an appropriate middle context $\mathcal{L}_{[n-m:n+m]}$ to force this behaviour. With this description one would expect the middle part to follow the GSA, leading to an alternative heuristic for q -ary lattices.

Heuristic 2.8 (ZGSA). *Let \mathbf{B} be a basis of a $2n$ -dimensional q -ary lattice \mathcal{L} with n q -vectors. After BKZ- β reduction the profile has the following shape:*

$$\|\mathbf{b}_i^*\| = \begin{cases} q & \text{if } i \leq n - m, \\ \sqrt{q} \cdot \alpha_\beta^{\frac{2n-1-2i}{2}}, & \text{if } n - m < i < n + m - 1, \\ 1, & \text{if } i \geq n + m - 1, \end{cases}$$

where $\alpha_\beta = \text{gh}(\beta)^{2/(\beta-1)}$, and $m = \frac{1}{2} + \frac{\ln(q)}{2\ln(\alpha_\beta)}$.

Again this gives us a good first order estimate. Asymptotically setting $\beta = \mathcal{B} \cdot n$ and $q = n^{\mathcal{Q}}$, we obtain $\ln(\alpha_\beta) = \frac{\ln(n)}{\mathcal{B} \cdot n} + O(n^{-1})$, and $m = \frac{1}{2} \mathcal{Q} \mathcal{B} \cdot n + O\left(\frac{n}{\ln(n)}\right)$.

2.4 Estimates

The main question of our work is to better understand how BKZ recovers the dense sublattice $\mathcal{L}^{\mathbf{GF}}$ from an NTRU lattice $\mathcal{L}^{\mathbf{H},q}$. Several works exist that give estimates on the blocksize β for which BKZ successfully recovers the secret key (\mathbf{g}, \mathbf{f}) , or more generally a vector from the dense sublattice. We discuss the state-of-the-art estimates, one known as the *2016 Estimate* [ADPS16] with further refinements [DDGR20, PV21], and one by Kirchner and Fouque [KF17].

While the 2016 Estimate already gives a clear explanation *how* BKZ recovers a suitable vector, the Kirchner and Fouque estimate is only based on an impossibility result. To be more precise about what we mean with recovery we define the following two events.

Definition 2.9 (BKZ Events). *For a BKZ run on an NTRU lattice \mathcal{L} with dense sublattice $\mathcal{L}^{\mathbf{GF}}$ we define two events:*

1. **Secret Key Recovery (SKR):** *The first time one the secret keys $(\mathbf{g}|\mathbf{f})$ is inserted.*

2. Dense Sublattice Discovery (DSD): The first time a dense lattice vector $\mathbf{v} \in \mathcal{L}^{GF}$ strictly longer than the secret key(s) is inserted.

We further specify SKR_κ and DSD_κ when the insertion takes place at position κ in the basis.

2016 Estimate [ADPS16] for SKR. The 2016 Estimate is aimed at the more general problem of detecting an unusually short vector in a lattice. To obtain an estimate for the NTRU problem, and more specifically the SKR event, we apply it to the unusually short vector $(\mathbf{g}|\mathbf{f}) \in \mathcal{L}^{H,q}$.

Claim 2.10 (SKR – 2016 Estimate). Let \mathcal{L} be a lattice of dimension d and let $\mathbf{v} \in \mathcal{L}$ be a unusually short vector $\|\mathbf{v}\| \ll \text{gh}(\mathcal{L})$. Then under the Geometric Series Assumption BKZ recovers \mathbf{v} if

$$\sqrt{\beta/d} \cdot \|\mathbf{v}\| < \sqrt{\alpha_\beta}^{2\beta-d-1} \cdot \text{vol}(\mathcal{L})^{1/d},$$

where $\alpha_\beta = \text{gh}(\beta)^{2/(\beta-1)}$.

The left hand side of the inequality is an estimate for $\|\pi_{d-\beta}(\mathbf{v})\|$, while the right hand size is the expected norm of $\mathbf{b}_{d-\beta}^*$ under the GSA. When the inequality is satisfied we expect that the shortest vector in $\mathcal{L}_{[d-\beta:d]}$ is in fact (a projection of) the unusually short vector, and thus it is inserted by BKZ at position $d - \beta$.

For q -ary lattices we can easily change the estimate to make use of the ZGSA instead, although for successful blocksizes $\mathbf{b}_{d-\beta}^*$ will not lie on the flat tail-part, and thus this will not change anything. Additionally for q -ary lattices it can be beneficial to apply the estimate not to the full lattice but on some projected sublattice $\mathcal{L}_{[i:d]}$ for $i \leq n$; the left hand side of the equation is expected to remain unchanged, while the right hand side might decrease as $\text{vol}(\mathcal{L})$ loses a factor q^i . Note that we do not necessarily have to explicitly let BKZ act on this projected sublattice, as BKZ already does this naturally.

Asymptotics. Consider the NTRU lattice $\mathcal{L}^{H,q}$ and suppose that $q = \Theta(n^{\mathcal{Q}})$, $\|\mathbf{v}\| = \|(\mathbf{g}, \mathbf{f})\| = \Theta(n^{\mathcal{S}})$ and $\beta = (\mathcal{B} + o(1))n$. Applying the 2016 Estimate the right hand side of the inequality is minimised when only keeping $k = \min((\sqrt{2\mathcal{B}\mathcal{Q}} - 1)n, n)$ of the q -vectors, so by applying the estimate to the projected sublattice $\mathcal{L}_{[n-k:2n]}^{H,q}$. For $\mathcal{S} \geq 1$ we have $k = n$, and solving the equation gives $\mathcal{B} = \frac{2}{\mathcal{Q}+2-2\mathcal{S}}$. For $\mathcal{S} < 1$ we have $k = (\sqrt{2\mathcal{B}\mathcal{Q}} - 1)n$, and solving gives $\mathcal{B} = \frac{2\mathcal{Q}}{(\mathcal{Q}+1-\mathcal{S})^2}$. Note in particular that in terms of q we require a blocksize of $\beta = \hat{\Theta}(n/\ln(q))$.

Refinements. The 2016 Estimate gives a clear explanation on how and where the secret vector is recovered. This also allows to further refine the estimate and give concrete predictions. For example by using a BKZ-simulator instead of the GSA, and by accounting for the probability that after the projection $\|\pi_{d-\beta}(\mathbf{v})\|$ has been found, it is successfully lifted to the full vector \mathbf{v} . Also instead of working

with the expected length of the projection, we can directly model the probability distribution under the assumption that \mathbf{v} is distributed as a Gaussian vector. Such refinements were applied in [DDGR20, PV21], and the resulting concrete predictions match with experiments to recover an unusually short vector. In this work, we use the (Z)GSA for the basis shape, but adjusting the slope to account for the speed of convergence using experimentally determined values. However, we do use the advanced probabilistic model for the detection and lifting of the short vector.

For NTRU there is not just a single unusually short vector, but there are $n = d/2$ of them, which makes it more likely that at least one of them is recovered. Because the refined concrete estimator already works with a probability distribution, we can easily take multiple vectors into account. The resulting predictions for the SKR event match the experiments reasonably well for smallish q as can be seen in Fig. 1. For large q , the so-called *overstretched* regime, the estimate is however too pessimistic.

Kirchner–Fouque Estimate [KF17] for DSD. In 2016 Albrecht, Bai and Lucas [ABD16] showed that for very large values of q one can mount an algebraic *subfield attack* on the cyclotomic NTRU problem with sub-exponential or even polynomial complexity. This allowed them to break several homomorphic encryption schemes that relied on NTRU in the overstretched regime.

However soon after, Kirchner–Fouque [KF17] showed that this elaborate algebraic attack was unnecessary: (dual-)BKZ already behaves much better in this regime than the 2016 Estimate predicts, leading to the same asymptotic improvements. The key idea behind their analysis is that in the overstretched regime the NTRU lattice $\mathcal{L}^{H,q}$ contains an exceptionally dense sublattice \mathcal{L}^{GF} of low volume. This gives a constraint on the basis profile via the following lemma by Pataki and Tural.

Lemma 2.11 (Pataki and Tural [PT08]). *Let \mathcal{L} be a d -dimensional lattice with basis $\mathbf{b}_0, \dots, \mathbf{b}_{d-1}$. For any k -dimensional sublattice $\mathcal{L}' \subset \mathcal{L}$ we have*

$$\text{vol}(\mathcal{L}') \geq \min_J \prod_{j \in J} \|\mathbf{b}_j^*\|,$$

where J ranges over the k -size subsets of $\{0, \dots, d-1\}$.

Applying Lemma 2.11 to the n -dimensional sublattice $\mathcal{L}^{GF} \subset \mathcal{L}^{H,q}$, and assuming a non-increasing profile, we obtain an upper bound on the volume of $\mathcal{L}_{[n:2n]}^{H,q}$. Assuming the ZGSA the latter volume increases when running BKZ- β for increasing block sizes, eventually contradicting the upper bound. This allows us to detect if a q -ary lattice is in fact an NTRU lattice, but additionally Kirchner–Fouque argue that BKZ must *somehow* have detected the dense sublattice after this point. Based on this impossibility argument they introduced the following estimate.

Claim 2.12 (DSD – Kirchner–Fouque Estimate). *Let $\mathcal{L}^{H,q}$ be an NTRU lattice of dimension $2n$, with dense sublattice $\mathcal{L}^{GF} \subset \mathcal{L}^{H,q}$. Under the Z-shape*

Geometric Series Assumption BKZ- β triggers the DSD event if

$$\text{vol}(\mathcal{L}^{GF}) < q^{\frac{m-1}{2}} \cdot \alpha_\beta^{-\frac{1}{2}(m-1)^2},$$

where $\alpha_\beta = \text{gh}(\beta)^{2/(\beta-1)}$, and $m = \frac{1}{2} + \frac{\ln(q)}{2\ln(\alpha_\beta)}$.

To apply this estimate we can bound $\text{vol}(\mathcal{L}^{GF})$ using the Hadamard inequality by $\|(\mathbf{g}|\mathbf{f})\|^n$. As a first approximation this is reasonably tight because the secret basis \mathbf{B}^{GF} is close to orthogonal.

Asymptotics. Consider the NTRU lattice $\mathcal{L}^{H,q}$ and suppose that $q = \Theta(n^{\mathcal{Q}})$, $\|(\mathbf{g}, \mathbf{f})\| = \Theta(n^{\mathcal{S}})$ and $\beta = (\mathcal{B} + o(1))n$. We apply the Kirchner–Fouque Estimate using that $m \approx \frac{\mathcal{B}\mathcal{Q}}{2}n$ and $\alpha_\beta \approx (\mathcal{B}n)^{1/(\mathcal{B}n)}$. The left hand side of the inequality is bounded by $n^{n\mathcal{S}+o(n)}$ and the right hand side equals $n^{\frac{\mathcal{B}\mathcal{Q}^2}{8}n+o(n)}$; solving gives $\mathcal{B} \geq \frac{8\mathcal{S}}{\mathcal{Q}^2}$. Note that in terms of q we require a blocksize of $\beta = \tilde{\Theta}(n/\ln^2(q))$, improving upon the 2016 Estimate by a factor $\ln(q)$. So for large enough q the Kirchner–Fouque Estimate predicts a lower successful blocksize than the 2016 Estimate. We call the value of q for which BKZ starts to behave better than predicted by the 2016 Estimate the *fatigue point*. For the common situation that $\mathcal{S} = \frac{1}{2}$, e.g. when each secret coefficient has standard deviation $\sigma = \Theta(1)$, the Kirchner–Fouque Estimate predicts that the fatigue point lies at some $q \leq n^{2.783+o(1)}$.

3 A New Estimate

3.1 Preliminary Experiments

Both the 2016 Estimate and the Kirchner–Fouque Estimate analyse an event that leads to successful recovery of a vector of the dense NTRU sublattice. This only gives an upper bound on the hardness; a different event leading to the recovery might happen at a lower blocksize. Additionally the Kirchner–Fouque Estimate is only based on an impossibility result and gives no explanation as to how BKZ actually recovers a vector from the dense sublattice. In order to derive a tight estimate we first run experiments to track down at which point a dense sublattice vector is actually found during the BKZ tours, i.e. when the DSD_κ event is triggered and at what position. Then we model this event in order to hopefully derive a tight estimate.

We run progressive BKZ on NTRU lattices $\mathcal{L}^{H,q}$ for fixed parameters $n = 127$, $\sigma^2 = \frac{2}{3}$, and several moduli q . For each BKZ insertion at position κ we check if the inserted vector belongs to the dense sublattice \mathcal{L}^{GF} , and thereby if the SKR_κ or DSD_κ event takes place, after which we stop.

The results are shown in Fig. 1. We take a closer look at the observed SKR_κ and DSD_κ events and where they are triggered. We can group our observations in three typical circumstances.

- **SKR-2016**. The SKR_κ event is mostly triggered for small values of q , and this mostly happens at the position $\kappa = 2n - \beta$, so in the last block $[2n - \beta : 2n)$, or slightly earlier. This coincides exactly with the $\text{SKR}_{2n-\beta}$ event as predicted by the 2016 Estimate [ADPS16, AGVW17].
- **DSD-PT**. The DSD_κ event is mostly triggered at positions $\kappa = n + k - \beta$ for $0 < k \ll n$. The inserted dense vector \mathbf{v} is often significantly longer than the secret key but still shorter than the q -vectors. On closer inspection the projected length $\|\pi_{n+k-\beta}(\mathbf{v})\|$ is close to the expected length $\sqrt{\frac{\beta}{n+k}} \|\mathbf{v}\|$ for all instances, more specifically the length of \mathbf{v} is well balanced over the Gram-Schmidt directions $\mathbf{b}_0^*, \dots, \mathbf{b}_{n+k-1}^*$. We name these events after the Pataki–Tural Lemma (DSD-PT).
- **DSD-LL**. For a few instances the DSD_κ event is triggered at large positions κ , up to $2n - \beta$. The inserted dense vector \mathbf{v} is again significantly longer than the secret key, but it has an unexpectedly short projection $\pi_\kappa(\mathbf{v})$ on the BKZ block $[\kappa : \kappa + \beta)$. We call these events lucky-lifts (DSD-LL).

The DSD-LL event could potentially be explained by the relatively large amount of shortish vectors in the close to orthogonal dense sublattice \mathcal{L}^{GF} compared to what one would expect based on the Gaussian Heuristic. These many vectors might compensate for the low probability event that: (1) such a long vector has such a short projection, and (2) the projected vector is correctly lifted by Babai’s nearest plane algorithm (thus a *lucky lift*). The DSD-LL event remains rare for all parameters we used in our experiments, and the successful blocksizes do not seem to deviate from the DSD-PT events. Although we think this circumstance deserves further analysis we therefore base our estimate on the more common DSD-PT event.

For the DSD-PT event the projected length $\|\pi_{n+k-\beta}(\mathbf{v})\|$ is close to $\sqrt{\frac{\beta}{n+k}} \|\mathbf{v}\|$, and thus the inserted dense vector \mathbf{v} must in fact be (close to) a shortest vector of the intersected sublattice $\mathcal{L}_{[0:n+k)}^{H,q} \cap \mathcal{L}^{GF}$. If not, the shortest vector would typically have an even smaller projection and would thus be inserted instead. For ease of analysis we therefore assume that \mathbf{v} is a shortest vector of $\mathcal{L}_{[0:n+k)}^{H,q} \cap \mathcal{L}^{GF}$. In short our new estimate can be described as follows.

Claim 3.1 (DSD-PT estimate). *A tour of $\text{BKZ-}\beta$ triggers the DSD event if*

$$\pi_{n+k-\beta}(\mathbf{v}) < \|\mathbf{b}_{n+k-\beta}^*\|,$$

where \mathbf{v} is a shortest vector of $\mathcal{L}_{[0:n+k)}^{H,q} \cap \mathcal{L}^{GF}$ for some $0 < k \leq n$.

3.2 Asymptotic Analysis

We denote the intersected sublattice by $\mathcal{L}_{\cap[0:r)}^{GF} := \mathcal{L}_{[0:r)}^{H,q} \cap \mathcal{L}^{GF}$. To directly apply Claim 3.1 we are interested in the length of \mathbf{v} , and thus the value of $\lambda_1(\mathcal{L}_{\cap[0:n+k)}^{GF})$. We break down the analysis into several steps. In order to obtain a bound on

the first minimum we first compute a bound on the volume of the intersection $\mathcal{L}_{\cap[0:n+k]}^{\mathbf{GF}}$ in terms of the basis profile and the volume of $\mathcal{L}^{\mathbf{GF}}$. Together with the GSA and a simple bound for $\text{vol}(\mathcal{L}^{\mathbf{GF}})$ we can then apply Minkowski's bound on the first minimum. By optimising $\kappa = n + k - \beta$ we obtain our new asymptotic estimate.

Intersection. To understand the behaviour of the volume of the intersected lattice we first need a small technical Lemma.

Lemma 3.2 ([DDGR20]). *Given a lattice \mathcal{L} with volume $\text{vol}(\mathcal{L})$, and a primitive vector \mathbf{v} with respect to \mathcal{L}^* . Let \mathbf{v}^\perp denote the subspace orthogonal to \mathbf{v} . Then $\mathcal{L} \cap \mathbf{v}^\perp$ is a lattice with volume $\text{vol}(\mathcal{L} \cap \mathbf{v}^\perp) = \|\mathbf{v}\| \cdot \text{vol}(\mathcal{L})$.*

The following Lemma generalises the Pataki–Tural Lemma on which the estimate of Kirchner–Fouque is based. More specifically the Pataki–Tural Lemma only considers the case where the intersection is always trivial ($s = 0$).

Lemma 3.3 (Generalisation of [PT08]). *Let \mathcal{L} be a d -dimensional lattice with basis $\mathbf{b}_0, \dots, \mathbf{b}_{d-1}$, and consider the sublattice $\mathcal{L}_{[0:s]}$. For any n -dimensional sublattice $\mathcal{L}' \subset \mathcal{L}$ we have*

$$\text{vol}(\mathcal{L}_{[0:s]} \cap \mathcal{L}') \leq \text{vol}(\mathcal{L}') \cdot \left(\min_J \prod_{j \in J} \|\mathbf{b}_j^*\| \right)^{-1},$$

where $k := \dim(\mathcal{L}_{[0:s]} \cap \mathcal{L}')$ and J ranges over the $(n - k)$ -size subsets of $\{s, \dots, d - 1\}$.

Proof. We write $\mathcal{L}'_{\cap[0:r]} := \mathcal{L}_{[0:r]} \cap \mathcal{L}'$. For $j = k, \dots, n$ we define $s_j \in \{s, \dots, d\}$ as the maximal index such that $\dim(\mathcal{L}'_{\cap[0:s_j]}) = j$, i.e. we obtain the following strict chain of sublattices:

$$\mathcal{L}'_{\cap[0:s]} = \mathcal{L}'_{\cap[0:s_k]} \subsetneq \mathcal{L}'_{\cap[0:s_{k+1}]} \subsetneq \dots \subsetneq \mathcal{L}'_{\cap[0:s_n]} = \mathcal{L}'.$$

Fix $j \in \{k, \dots, n - 1\}$. Because the basis vectors $\mathbf{b}_0, \dots, \mathbf{b}_{d-1}$ are linearly independent we have that $\mathcal{L}'_{\cap[0:s_{(j+1)}]} = \mathcal{L}'_{\cap[0:s_j+1]}$. This allows us to focus on the volume decrease from index $s_j + 1$ to s_j , for which we know that

$$\mathcal{L}'_{\cap[0:s_j]} = \mathcal{L}'_{\cap[0:s_j+1]} \cap (\mathbf{b}_{s_j}^*)^\perp,$$

where $(\mathbf{b}_{s_j}^*)^\perp$ denotes the subspace orthogonal to $\mathbf{b}_{s_j}^*$. The corresponding dual basis of $\mathbf{b}_0, \dots, \mathbf{b}_{s_j}$ contains a dual vector $\mathbf{d} \in \mathcal{L}_{[0:s_j+1]}^*$ of length $\|\mathbf{b}_{s_j}^*\|^{-1}$ with $\text{span}(\mathbf{d}) = \text{span}(\mathbf{b}_{s_j}^*)$. Let π be the orthogonal projection onto $\text{span}(\mathcal{L}'_{\cap[0:s_j+1]})$, then $\pi(\mathbf{d}) \in (\mathcal{L}'_{\cap[0:s_j+1]})^*$. Let $m \in \mathbb{Z}_{\geq 1}$ be such that $\pi(\mathbf{d})/m$ is primitive w.r.t. $(\mathcal{L}'_{\cap[0:s_j+1]})^*$, then by Lemma 3.2 we obtain:

$$\text{vol}(\mathcal{L}'_{\cap[0:s_j]}) = \text{vol}(\mathcal{L}'_{\cap[0:s_j+1]}) \cdot \|\pi(\mathbf{d})/m\| \leq \text{vol}(\mathcal{L}'_{\cap[0:s_j+1]}) \cdot \|\mathbf{b}_{s_j}^*\|^{-1}.$$

We conclude the proof by chaining the above inequality for $j = k, \dots, n - 1$. \square

Before recovering a dense lattice vector we heuristically assume that there is no special relation between the current lattice basis and the dense sublattice. More specific we can consider that the span of $\mathbf{b}_0, \dots, \mathbf{b}_{n-1}$ and that of \mathcal{L}^{GF} behave like random n -dimensional subspaces, and thus they have a trivial intersection with high probability in the $2n$ -dimensional space. As a direct result we have that $\dim(\mathcal{L}_{\cap[0:n+k]}^{GF}) = k$ for $k = 0, \dots, n$. Applying this to Lemma 3.3 we obtain the following corollary.

Corollary 3.4. *Let $\mathcal{L}^{H,q}$ be an NTRU lattice with dense sublattice \mathcal{L}^{GF} of dimension n , if $\dim(\mathcal{L}_{\cap[0:n+k]}^{GF}) = k$ for some $k \geq 0$, then*

$$\text{vol}(\mathcal{L}_{\cap[0:n+k]}^{GF}) \leq \text{vol}(\mathcal{L}^{GF}) \cdot \left(\prod_{j=n+k}^{d-1} \|\mathbf{b}_j^*\| \right)^{-1}.$$

Note that Corollary 3.4 already shows that the new estimate can not be worse than the Kirchner–Fouque Estimate. Namely if the Kirchner–Fouque Estimate is triggered, then for intersection dimension $k = 1$ the right hand side is smaller than $\|\mathbf{b}_n^*\|$. Assuming a non-decreasing profile we then have $\lambda_1(\mathcal{L}_{\cap[0:n+1]}^{GF}) = \text{vol}(\mathcal{L}_{\cap[0:n+1]}^{GF}) \leq \|\mathbf{b}_n^*\| \leq \|\mathbf{b}_{n+1-\beta}^*\|$, which implies that BKZ- β would find a dense sublattice vector in this block (or earlier).

Volume Dense Sublattice. To use Corollary 3.4 we also need to bound the volume of the dense sublattice \mathcal{L}^{GF} . Because the secret basis is close to orthogonal the Hadamard Inequality $\text{vol}(\mathcal{L}^{GF}) \leq \|(\mathbf{g}|\mathbf{f})\|^n$ is sufficient as a first order approximation.

Conclusion. To obtain a heuristic asymptotic estimate we will assume that before finding a dense lattice vector the basis follows the ZGSA shape.

Claim 3.5. *The BKZ algorithm with blocksize $\beta = \mathcal{B}n$ applied to an NTRU instance with parameters $q = \Theta(n^{\mathcal{Q}})$, $\|(\mathbf{g}|\mathbf{f})\| = O(n^{\mathcal{S}})$ triggers the DSD event if*

$$\mathcal{B} = \frac{8\mathcal{S}}{\mathcal{Q}^2 + 1} + o(1).$$

Justification. By the Hadamard Inequality we have $\ln(\text{vol}(\mathcal{L}^{GF})) \leq \mathcal{S}n \ln(n) + O(n)$. Let $k := \mathcal{K}n > 0$. Heuristically we expect that $\dim(\mathcal{L}_{\cap[0:n+k]}^{H,q} \cap \mathcal{L}^{GF}) = k$, and thus by Corollary 3.4 and by assuming the ZGSA we obtain a bound on the volume of the intersected sublattice:

$$\begin{aligned} \ln(\text{vol}(\mathcal{L}_{\cap[0:n+k]}^{H,q} \cap \mathcal{L}^{GF})) &\leq \mathcal{S}n \ln(n) - \frac{1}{2} \sum_{i=n+k}^{n+m-1} \left(Q + \frac{2n-1-2i}{\mathcal{B}n} \right) \ln(n) + O(n) \\ &= \mathcal{S}n \ln(n) - \frac{(\mathcal{B}\mathcal{Q} - 2\mathcal{K})^2}{8\mathcal{B}} n \ln(n) + O(n) \end{aligned}$$

By Minkowski's bound we bound the first minimum using the above volume

$$\begin{aligned} \ln(\lambda_1(\mathcal{L}_{[0:n+k]}^{H,q} \cap \mathcal{L}^{GF})) &\leq \frac{1}{2} \ln(\mathcal{K}n) + \frac{\ln(\text{vol}(\mathcal{L}_{[0:n+k]}^{H,q} \cap \mathcal{L}^{GF}))}{\mathcal{K}n} + O(1) \\ &\leq \left(-\frac{(\mathcal{B}\mathcal{Q} - 2\mathcal{K})^2}{8\mathcal{B}\mathcal{K}} + \frac{\mathcal{S}}{\mathcal{K}} + \frac{1}{2} \right) \ln(n) + O(1). \end{aligned}$$

After projecting onto the block $[n+k-\beta : n+k]$ the above short vector does not increase in length.⁷ BKZ detects the projected dense lattice vector in this block if the length is less than $\|\mathbf{b}_{n+k-\beta}^*\| = (\frac{1}{2}\mathcal{Q} + \frac{\mathcal{B}-\mathcal{K}}{\mathcal{B}}) \ln(n) + O(1)$. Solving for \mathcal{B} shows that this is the case when

$$\mathcal{B} \geq \frac{2\sqrt{(2\mathcal{S} - \mathcal{K})^2 + \mathcal{K}^2\mathcal{Q}^2} + 2(2\mathcal{S} - \mathcal{K})}{\mathcal{Q}^2}.$$

When $\mathcal{K} = \frac{4\mathcal{S}}{\mathcal{Q}^2+1}$ the right hand side is minimised and we obtain that BKZ detects the projected dense lattice vector when $\mathcal{B} \geq \frac{8\mathcal{S}}{\mathcal{Q}^2+1}$, which concludes the claim. This routine computation can be verified symbolically via our sage notebook `claim3_5.ipynb`. \triangle

Our new estimate gives an asymptotic improvement over the Kirchner–Fouque Estimate ($\frac{8\mathcal{S}}{\mathcal{Q}^2}$). Asymptotically the optimal position is at $\kappa = n+k-\beta \approx n - \frac{1}{2}\beta$. Interestingly, if we do not optimize k and only consider $k = O(1)$ we obtain the same asymptotic estimate as Kirchner–Fouque, which again emphasizes that we generalised their analysis.

For the fatigue point we compare the relative blocksize of $\frac{8\mathcal{S}}{\mathcal{Q}^2+1}$ to that of the 2016 Estimate given by $\frac{2\mathcal{Q}}{(\mathcal{Q}+1-\mathcal{S})^2}$ for $\mathcal{S} < 1$ and by $\frac{2}{\mathcal{Q}+2-2\mathcal{S}}$ for $\mathcal{S} \geq 1$. For ternary secrets ($\mathcal{S} = \frac{1}{2}$) this narrows down the fatigue point from $q \leq n^{2.783+o(1)}$ to $q = n^{2.484+o(1)}$ compared to the Kirchner–Fouque Estimate. This is still far above the (sub)linear parameters used for NTRU encryption schemes, and thus asymptotically we can close the pending question if these parameters fall in the weaker overstretched regime or not. In practice however we do observe fatigue points that are significantly lower than the naive value of $q = n^{2.484}$, which motivates a concrete analysis with concrete predictions (Fig. 2).

4 Concrete Analysis

In this section we consider a concrete analysis of our new DSD-PT estimate, based on simple heuristics, to better predict the behaviour in practice, and to show that our analysis matches experiments and is thus likely to be tight. The first order asymptotics shown in Sect. 3.2 will remain unchanged, but the differences are significant for practical parameters. Again we split the analysis

⁷ One may also be concerned that the short vector would collapse to $\mathbf{0}$ after projection onto the block $[n+k-\beta : n+k]$, but this becomes increasingly unlikely as the dimension β of the block grows.




Fig. 2. Comparison of asymptotic estimates and new fatigue point for $n \rightarrow \infty$ when the secret key coefficients have standard deviation $\sigma = \Theta(1)$.

into several steps, but now derive heuristic expectations instead of loose upper bounds.

We assume that lattice vectors we encounter follow the Gaussian heuristic, and thus in particular that vectors are spherically distributed after normalisation. When projecting such vectors to a lower dimension they become shorter. The following Lemma shows how much shorter we expect them to become.

Lemma 4.1. *Let $\mathbf{x} \in \mathcal{S}^{d-1}$ follow a spherical distribution, and let $\pi_V : \mathbb{R}^d \rightarrow V$ be a projection to some k -dimensional subspace $V \subset \mathbb{R}^d$, then*

$$\mathbb{E}[\ln(\|\pi_V(\mathbf{x})\|)] = \frac{1}{2}(\psi(k/2) - \psi(d/2)).$$

Proof. Let X_0, \dots, X_{d-1} be standard normal random variables, then the vector $\mathbf{x} = (x_0, \dots, x_{d-1})$, with $x_j = X_j / \sqrt{\sum_{i=0}^{d-1} X_i^2}$, is spherically distributed. Without loss of generality we can assume that π_V projects onto the first k -coordinates. Then we conclude by Lemma 2.1 that

$$\begin{aligned} \mathbb{E}[\ln(\|\pi_V(\mathbf{x})\|)] &= \frac{1}{2} \mathbb{E} \left[\ln \left(\frac{\sum_{i=0}^{k-1} X_i^2}{\sum_{i=0}^{d-1} X_i^2} \right) \right] \\ &= \frac{1}{2} \mathbb{E} \left[\ln \left(\sum_{i=0}^{k-1} X_i^2 \right) \right] - \frac{1}{2} \mathbb{E} \left[\ln \left(\sum_{i=0}^{d-1} X_i^2 \right) \right] \\ &= \frac{1}{2}(\psi(k/2) - \psi(d/2)). \end{aligned}$$

□

4.1 Intersection

We start by giving a concrete average-case estimate for the intersection volumes. Assuming that projections behave as random we obtain the following concrete estimate.

Claim 4.2. *Let \mathcal{L} be a $2n$ -dimensional NTRU lattice with dense sublattice \mathcal{L}^{GF} , before the DSD event is triggered we have for $k = 1, \dots, n$ that $\dim(\mathcal{L}_{\cap[0:n+k]}^{GF}) = k$, and*

$$\begin{aligned} \mathbb{E}[\ln \text{vol}(\mathcal{L}_{\cap[0:n+k]}^{GF})] &= \ln \text{vol}(\mathcal{L}^{GF}) - \left(\sum_{j=n+k}^{2n-1} \ln \|\mathbf{b}_j^*\| \right) \\ &\quad + \sum_{l=k+1}^n \psi\left(\frac{l}{2}\right) - \psi\left(\frac{n+l}{2}\right) + \frac{\zeta'(l)}{\zeta(l)}, \end{aligned}$$

where $\zeta(l) := \sum_{m=1}^{\infty} \frac{1}{m^l}$ is the Riemann zeta function and $\zeta'(l) := \sum_{m=1}^{\infty} \frac{\ln(m)}{m^l}$ its derivative.

Justification. We follow the proof of Lemma 3.3. It is tight except for the length decrease from $\|\mathbf{d}\|$ to the projected and primitive vector $\|\pi(\mathbf{d})\|/m$. Note that when obtaining $\ln \text{vol}(\mathcal{L}_{\cap[0:n+l-1]}^{GF})$ from $\ln \text{vol}(\mathcal{L}_{\cap[0:n+l]}^{GF})$ for some $l = k + 1, \dots, n$, the dual vector \mathbf{d} lives in a $(n+l)$ -dimensional space and is projected to an l -dimensional space. Heuristically we assume that the normalisation of \mathbf{d} is spherically distributed (or that π projects to a random l -dimensional subspace). By Lemma 4.1 the log-expected decrease in length from this projection then equals

$$\mathbb{E}[\ln(\pi(\|\mathbf{d}\|)) - \ln(\|\mathbf{d}\|)] = \psi\left(\frac{l}{2}\right) - \psi\left(\frac{n+l}{2}\right).$$

To conclude we also have to include the primitivity of $\pi(\mathbf{d})$ and thus the log-expectation of $m \geq 1$ such that $\pi(\mathbf{d})/m$ is primitive. For any basis $\mathbf{d}_0, \dots, \mathbf{d}_{l-1}$ and $\pi(\mathbf{d}) = \sum_{i=0}^{l-1} x_i \mathbf{d}_i$, we have $m = \gcd(x_0, \dots, x_{l-1})$. Heuristically we assume that the absolute coefficients $|x_0|, \dots, |x_{l-1}|$ are random integers in the interval $\{1, \dots, B\}$ and we let $B \rightarrow \infty$. For $l \geq 2$ we have (see e.g. [DE+04])

$$\mathbb{P}_{\mathbf{x} \in \{1, \dots, B\}^l} [\gcd(x_0, \dots, x_{l-1}) = m] = \frac{1}{\zeta(l)} \cdot \frac{1}{m^l} + O(\ln(B)/(Bm^{l-1})),$$

where the Riemann zeta function $\zeta(l) = \sum_{m=1}^{\infty} \frac{1}{m^l}$ is just the normalisation factor. From this we conclude that

$$\begin{aligned} \lim_{B \rightarrow \infty} \mathbb{E}_{\mathbf{x} \in \{1, \dots, B\}^l} [\ln \gcd(x_0, \dots, x_{l-1})] &= \lim_{B \rightarrow \infty} \frac{1}{\zeta(l)} \sum_{m=1}^B \left[\frac{\ln(m)}{m^l} + O\left(\frac{\ln(m) \ln(B)}{Bm^{l-1}}\right) \right] \\ &= -\frac{\zeta'(l)}{\zeta(l)} \end{aligned}$$

for $l \geq k + 1 \geq 2$.

△




Fig. 3. Experimental values of $\ln \text{vol}(\mathcal{L}_{[0:n+k]}^{GF})$ versus Claim 4.2 for circulant and matrix NTRU respectively. For each variant we used 256 LLL reduced NTRU lattices with parameters $q = 257, n = 79, \sigma^2 = \frac{2}{3}$ and computed the intersection for each k .

Validation. To validate Claim 4.2 we computed the actual intersection volumes $\text{vol}(\mathcal{L}_{\cap[0:n+k]}^{GF})$ for LLL reduced NTRU instances. We observed here, and also in further experiments, that the dimension assumption $\dim(\mathcal{L}_{\cap[0:n+k]}^{GF}) = k$ holds before we get close to triggering the DSD event. Figure 3 shows that our prediction perfectly matches the experiments for matrix NTRU. For circulant NTRU we see both that the expectation is slightly off and that the variance is much higher. The higher variance can be explained from the fact that the projections are very much dependent due to the circulant structure; in fact a closer inspection shows that for k close to n the differences with the prediction are highly correlated. We were not able to explain the error in the predicted expectation, but it seems to be caused by the circulant structure in combination with the Z-shape: the error decreased and eventually disappeared for large values of q and σ , for which the Z-shape disappeared (and before the DSD event was triggered). A maximal log-error of 2.5 is reached at $k = 1$. Note that a log-error of ϵ on $\text{vol}(\mathcal{L}_{\cap[0:n+k]}^{GF})$ translate into a factor of $e^{\epsilon/k}$ on the predicted length for the shortest vector. Except for very small k , this error appears benign.

4.2 Dense Sublattice

In this section we give a concrete estimate for the expected volume of the dense NTRU sublattice \mathcal{L}^{GF} . Directly from the construction we obtain a basis $[\mathbf{G}|\mathbf{F}]$ of \mathcal{L}^{GF} , with \mathbf{F} invertible. We consider two cases, that of regular NTRU where \mathbf{F} and \mathbf{G} are circulant matrices, and that of matrix NTRU, where all entries are independently sampled. For both constructions the entries are sampled from independent discrete Gaussians over \mathbb{Z} , with some standard deviation $\sigma > 0$. As the only heuristic we assume that the individual entries in fact follow a *continuous* Gaussian instead of the discrete one.

Matrix NTRU. We start with matrix NTRU, where we heuristically assume that all $2n \times n$ coefficients of the basis $[\mathbf{G}|\mathbf{F}]$ are sampled according to independent continuous Gaussians with standard deviation σ . Under this heuristic we can derive an exact expression for the expected log-volume of the dense sublattice.

Lemma 4.3. *Let $[\mathbf{G}|\mathbf{F}]$ be a basis of the lattice $\mathcal{L}^{\mathbf{G}\mathbf{F}}$ where all sampled entries are i.i.d. continuous Gaussians with standard deviation $\sigma > 0$, then*

$$\mathbb{E}[\ln(\text{vol}(\mathcal{L}^{\mathbf{G}\mathbf{F}}))] = \frac{1}{2}n (\ln(2\sigma^2) + \psi(n)) + \sum_{i=0}^{n-1} \left[\psi\left(\frac{2n-i}{2}\right) - \psi(n) \right].$$

Proof. By Lemma 2.1 the log-expectation of the norm of each basis element equals $(\ln(2\sigma^2) + \psi(n))/2$. Note that the i -th Gram-Schmidt vector \mathbf{b}_i^* is obtained after projecting the i -th basis vector orthogonally away from an i -dimensional subspace, and thus onto a $2n - i$ dimensional subspace. However after normalisation the basis vectors follow a spherical distribution and thus by Lemma 4.1 we have

$$\mathbb{E}[\ln \|\mathbf{b}_i^*\|] = (\ln(2\sigma^2) + \psi(n))/2 + \psi\left(\frac{2n-i}{2}\right) - \psi(n).$$

We conclude by noting that $\mathbb{E}[\ln(\text{vol}(\mathcal{L}^{\mathbf{G}\mathbf{F}}))] = \sum_{i=0}^{n-1} \mathbb{E}[\ln \|\mathbf{b}_i^*\|]$. □




Fig. 4. Experimental values of $\ln(\text{vol}(\mathcal{L}^{\mathbf{G}\mathbf{F}}))$ versus Lemma 4.3 for matrix NTRU with discrete Gaussians and variance $\sigma^2 = \frac{2}{3}$. For each parameter n we generated 512 instances.

Circulant NTRU. For circulant NTRU both \mathbf{G} and \mathbf{F} in the basis $[\mathbf{G}|\mathbf{F}]$ are circulant matrices. Again we replace discrete with continuous Gaussians. The eigenvalues and eigenvectors of a circulant matrix are well known and we use this to obtain an exact expression for the expected volume of the dense sublattice.

Lemma 4.4. *Let $[\mathbf{G}|\mathbf{F}]$ be a basis of the lattice $\mathcal{L}^{\mathbf{G}\mathbf{F}}$ where \mathbf{G}, \mathbf{F} are circulant and all sampled entries are i.i.d. continuous Gaussians with standard deviation $\sigma > 0$, then*

$$\mathbb{E}[\ln(\text{vol}(\mathcal{L}^{\mathbf{G}\mathbf{F}}))] = \frac{1}{2}n (\ln(2n\sigma^2) + \psi(1)) + \frac{1}{2}(n-1)(1 - \ln(2)).$$

Proof. For $n \times n$ circulant matrices \mathbf{G}, \mathbf{F} the eigenvectors are identical and given by $\mathbf{v}_j := (1, \omega^j, \omega^{2j}, \dots, \omega^{(n-1)j})$ for $j = 0, \dots, n-1$, where $\omega := e^{2\pi i/n} \in \mathbb{C}$ is a primitive n -th root of unity. Suppose that the circulant matrix \mathbf{G} is generated by the vector $\mathbf{c} = (c_0, \dots, c_{n-1})$, then the corresponding eigenvalues are given by the DFT coefficients of \mathbf{c} , namely $\lambda_j := c_0 + c_{n-1}\omega^j + \dots + c_1\omega^{(n-1)j}$. We have that $\lambda_0 = \sum_{j=0}^{n-1} c_j$, and thus λ_0 follows a Gaussian distribution with variance $n\sigma^2$, and in particular $\lambda_0^2 \sim \chi_{1,n\sigma^2}^2$. Additionally for $j = 1, \dots, n-1$ we can write $\lambda_j = X + i \cdot Y \in \mathbb{C}$ where $X, Y \in \mathbb{R}$ are both linear combinations of the c_i 's and thus (X, Y) follows a jointly Gaussian distribution. A simple computation shows that X and Y both have variance $n\sigma^2/2$ and that they are uncorrelated, which for Gaussians implies that they are independent [PS89, p. 212]. So $|\lambda_j|^2 = X^2 + Y^2 \sim \chi_{2,n\sigma^2/2}^2$. Note that all circulant matrices have the same eigenvectors and thus the squared singular values of the concatenation of two circulant matrices are the sum of the squared absolute eigenvalues. So $[\mathbf{G}|\mathbf{F}]$ has one squared singular value s_0^2 distributed as $\chi_{1,n\sigma^2}^2 + \chi_{1,n\sigma^2}^2 = \chi_{2,n\sigma^2}^2$, and $n-1$ squared singular values s_1^2, \dots, s_{n-1}^2 distributed as $\chi_{2,n\sigma^2/2}^2 + \chi_{2,n\sigma^2/2}^2 = \chi_{4,n\sigma^2/2}^2$. By Lemma 2.1 they have a log-expectation of

$$\mathbb{E}[\ln s_0^2] = \ln(2n\sigma^2) + \psi(1), \text{ and } \mathbb{E}[\ln s_j^2] = \ln(n\sigma^2) + \psi(2)$$

for $j = 1, \dots, n-1$. We conclude by noting that $\mathbb{E}[\ln(\text{vol}(\mathcal{L}^{\mathbf{G}\mathbf{F}}))] = \frac{1}{2} \sum_{i=0}^{n-1} \ln(s_i^2)$. \square

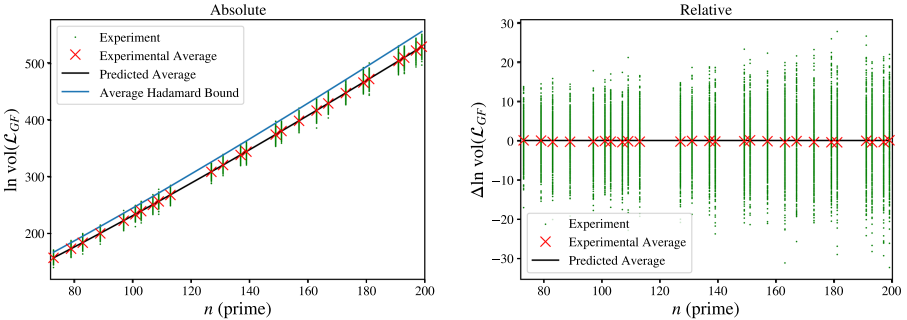


Fig. 5. Experimental values of $\ln(\text{vol}(\mathcal{L}^{\mathbf{G}\mathbf{F}}))$ versus Lemma 4.4 for circulant NTRU with discrete Gaussians and variance $\sigma^2 = \frac{2}{3}$. For each parameter n we generated 512 instances.

Validation. To validate our concrete estimate for $\text{vol}(\mathcal{L}^{GF})$ we generated the NTRU sublattice for several dimensions and computed its volume. We sample the secret coefficients following a discrete Gaussian with variance $\sigma^2 = \frac{2}{3}$ and ran experiments for both matrix NTRU and circulant NTRU. In Figs. 4 and 5 we see that the predictions from Lemmas 4.3 and 4.4 perfectly fit the observed volumes in all dimensions. We do note that the variance is quite significant for the circulant case, but it can be fully explained by the computed eigenvalue distributions in the proof of Lemma 4.4.

4.3 Further Refinements

We discuss some further refinements, some of which were already successfully applied to the 2016 Estimate [AGVW17, DDGR20, PV21].

Gaussian Heuristic. For our asymptotic analysis we used Minkowski’s bound to estimate the length $\lambda_1(\mathcal{L}_{\cap[0:n+k]}^{GF})$ in terms of the volume $\text{vol}(\mathcal{L}_{\cap[0:n+k]}^{GF})$. A natural way to obtain a concrete estimate for the expected minimal length is by assuming that the intersection $\mathcal{L}_{\cap[0:n+k]}^{GF}$ follows the Gaussian Heuristic and thus for our prediction we assume that

$$\lambda_1(\mathcal{L}_{\cap[0:n+k]}^{GF}) = \text{gh}(\mathcal{L}_{\cap[0:n+k]}^{GF}) \approx \sqrt{k/(2\pi e)} \cdot \text{vol}(\mathcal{L}_{\cap[0:n+k]}^{GF})^{1/k}.$$

We should however be careful with this assumption, as in fact it is false for $k = n$. E.g. the above predicts that $\lambda_1(\mathcal{L}^{GF}) \approx \sqrt{n/(2\pi e)} \cdot \sqrt{2n\sigma^2}$, while we know that $\lambda_1(\mathcal{L}^{GF}) = \|(\mathbf{g}, \mathbf{f})\| \approx \sqrt{2n\sigma^2}$, a factor $\Theta(\sqrt{n})$ shorter than predicted. The reason for this is that the dense sublattice is up to rotation and scaling very similar to the orthogonal lattice \mathbb{Z}^n , precisely the lattice for which it is well known that the Gaussian Heuristic is false. For small $k \ll n$ we do observe that the intersected lattice $\mathcal{L}_{\cap[0:s]}^{GF}$ follows the Gaussian Heuristic; the orthogonal structure seems to be broken by the intersection. However we do not have a clear idea how large k can become before the orthogonal structure returns and the minimal length stops following the prediction from the Gaussian Heuristic. We think this behaviour deserves some further investigation, e.g. if the transition is very sudden or not, and we leave it as an open problem. This near-orthogonality of $\mathcal{L}_{\cap[0:s]}^{GF}$ may be critical to model the DSD-LL events.

Probabilities. So far we have only considered expectations of volumes and projections. While this is enough to give a rough concrete estimate we want to be more precise. Success probabilities can accumulate up over multiple BKZ blocks and (progressive) tours, possibly leading to success at much lower block sizes than the rough estimate. We continue using the expected values for the volume of the dense sublattice and the intersection volumes to obtain the expected length $\lambda_1(\mathcal{L}_{\cap[0:s]}^{GF})$ of the dense sublattice vector via the Gaussian Heuristic. However we then model the short dense sublattice vector $\mathbf{v} \in \mathcal{L}_{\cap[0:s]}^{GF}$ as an s -dimensional Gaussian vector with the same expected length; allowing us to compute the

exact probability that $\|\pi_{s-\beta}(\mathbf{v})\| \leq \|\mathbf{b}_{s-\beta}^*\|$ using the CDF of the chi-square distribution with β degrees of freedom.

Up to now we have ignored the probability that after $\pi_{s-\beta}(\mathbf{v})$ is inserted, it is also correctly lifted to the full vector \mathbf{v} by later BKZ tours. While this almost always happens for higher block sizes, it is not so likely for lower block sizes, and ignoring this leads to overly optimistic predictions. For BKZ- β to successfully lift or *eventually* pull the vector \mathbf{v} to the front it should also satisfy $\|\pi_i(\mathbf{v})\| \leq \|\mathbf{b}_i^*\|$ for all $i = s - 2\beta + 1, s - 3\beta + 2, \dots$. These conditions are not independent which makes them hard to compute exactly. We simplify the computation by only considering the dependence for consecutive positions $i, i - \beta + 1$ as done in [DDGR20]. We iteratively run our estimator for progressive $\beta = 2, 3, \dots$ and take account of all probabilities assuming that all tours behave completely independently. Our new concrete estimate will be the expected successful block size. Additionally this allows us to combine both the (probabilistic) SKR 2016 Estimate and our new DSD-PT estimate in a single estimator. With some more administration we can also predict the distribution of the successful location κ , and predict the probability that the SKR event happens before the DSD event.

BKZ Shape for Low Block sizes. While the formulas for the (Z)GSA slope α_β and the expected first minimum $\text{gh}(\beta)$ convert to the experimental values for large block sizes of say $\beta \geq 50$, they are not as accurate for small β . As expected the convergence is worse for progressive BKZ when we only use a few tours of each block size. We ran some experiment on random low dimensional q -ary lattices to obtain practical estimates for $\text{gh}(\beta)$ with $\beta \leq 50$. Earlier works about the 2016 Estimate resorted to BKZ simulators to predict the BKZ shape, which account for the number of tours and also the special shape of the head and tail that do not perfectly follow the GSA shape. Together with the earlier mentioned refinements this resulted in very precise predictions [DDGR20, PV21]. However how BKZ acts on a Z-shaped basis is much less understood [AD21] and as of yet there are no accurate BKZ simulators. Understanding the behaviour and creating an accurate simulator would be very interesting, but is out of the scope of this work. We continue using the ZGSA, but we resort to experimental values for α_β obtained by running BKZ on random q -ary lattices for large q . To remain consistent we also do not use a simulator for the GSA shape, and accept the small discrepancy between the predictions and practical experiments.

5 Experimental Verification

In this section we experimentally confirm our predictions. Further detailed experimental data and discussion is given in the eprint version⁸.

5.1 Successful Blocksize

We start with comparing our concrete predictions to the preliminary experiment from Sect. 3.1. We ran progressive BKZ with 8 tours on matrix NTRU instances

⁸ Section 5.2 in <https://eprint.iacr.org/2021/999>.




Fig. 6. Experiment versus prediction for progressive BKZ with 8 tours on matrix NTRU instances with parameters $n = 127, \sigma^2 = \frac{2}{3}$ for several moduli q . We did 10 runs per modulus q .

with parameters $n = 127, \sigma^2 = \frac{2}{3}$ for several moduli q . In Fig. 6 we show the blocksizes at which the SKR or DSD event is first detected, and compare them to our concrete estimator. We ran the estimator three times for each modulus q : only accounting for SKR, only accounting for DSD-PT, and accounting for both. Note that the combined estimate can be strictly lower than both the first two because the probabilities to succeed accumulate over both events. We calibrated the values of α_β by running the same BKZ routine on $(2 \cdot 127)$ -dimensional q -ary lattices with $q \approx 2^{20}$.

We observe that the experiments match the estimates reasonably well, with an average blocksize error of less than 2 for the DSD events and less than 3 for the SKR events. We shortly discuss potential sources of the small errors error.

- We do not actually run the classical BKZ algorithm, but the BKZ 2.0 algorithm as it is more feasible to run for large blocksizes. One part of the latter algorithm is that in each BKZ block $[\kappa : \kappa + \beta)$ the last $\beta - 1$ vectors are randomised before finding a short projected vector. This temporarily breaks the GSA shape and results a small ‘bump’ in the profile that is pushed to the right during a tour. On average we measured at the SKR events a log-increase of 0.048 on the value of $\|\mathbf{b}_\kappa^*\|$ compared to the GSA (while the rest of the basis matches very closely). Although anecdotal, adjusting our estimator with this offset of 0.048 resulted in very close predictions for the SKR events.
- For small blocksizes $\beta \leq 30$ we see that our DSD-PT estimate is slightly pessimistic compared to the experiments. However the successful profile slope α_β (computed from the profile at the moment of detection) does closely match

the predicted slope $\alpha_{\beta_{pred}}$, pointing to a wrong calibration of the slope parameter for very low blocksizes. Note that the non-flat part of the Z-shape in our experiments has size less than the $2 \cdot 127$ dimensional lattice used for calibration, which plausibly explain why the slope converges more quickly than expected.

5.2 Fatigue Point

Our concrete estimator follows the experiments reasonably well and thus we can use it to estimate the concrete fatigue point for dimensions that are not feasible in practice. To verify our estimate of the fatigue point we also did some experiments in dimensions that are still feasible. For this we ran a *soft* binary search, only decreasing the interval length by $3/4$ so as not view a probabilistic result as a definitive answer. More specifically, starting with a range of $[q_{min}, q_{max}]$ we ran an experiment for a prime $q \approx (q_{min} + q_{max})/2$. If it succeeds with an SKR event we update q_{min} to $(q_{min} + q)/2 + 1$, if it succeeds with a DSD event we update q_{max} to $(q_{max} + q)/2 - 1$. We repeat this until the interval does not contain any prime and we return $(q_{min} + q_{max})/2$ as a rough estimate of the fatigue point. We averaged this over 20 experiments for each parameter n . We chose for matrix NTRU because of the lower variance in the hardness of these instances.

We compared this to our prediction. Because the estimator accounts for probabilities of events, we can predict for which value of q about 50% of the instances succeeds with a DSD event. Because it would be unreasonable to calibrate the




Fig. 7. Concrete fatigue point versus asymptotics using progressive BKZ with 8 tours on matrix NTRU instances with variance $\sigma^2 = \frac{2}{3}$. The 0.5 percentile line shows for which q we estimate that the DSD event is triggered before the SKR event for about 50% of the instances.




Fig. 8. Experiment versus prediction for progressive BKZ with 8 tours on matrix NTRU instances with parameters $n = 113, \sigma^2 = \frac{2}{3}$ for several moduli q . We did 100 runs per modulus q and the plot shows the ratio of these runs succeeding with a DSD event (before an SKR event).

low blocksize slope values α_β for each dimension we reused those of the $2 \cdot 127$ dimensional q -ary lattice from an earlier experiment. This might make the estimates a bit less precise for $n \ll 127$, and $n \gg 127$ if the successful blocksize is small around the fatigue point.

The results are shown in Fig. 7 and plotted against $Cn^{2.484}$ for several constants C . Remarkably the experiments and concrete predictions closely follow the asymptotics already for reasonably small values of n . A loglog-linear regression of the 50% DSD-PT estimate over all primes $199, \dots, 499$ gives $0.0034 \cdot n^{2.506}$. Restricting the exponent to 2.484 gives $0.0038 \cdot n^{2.484}$ with a log-standard deviation of only 0.006.

The experimental average appears slightly higher than the estimator prediction for 50% DSD - 50% SKR. The main reason for this seems to be that the estimator is slightly pessimistic for detecting the SKR event, as already observed and explained in Sect. 5.1. Another small detail is that the binary search is slightly biased to higher values of q because at each iteration we pick the *next* prime after $(q_{\min} + q_{\max})/2$.

5.3 Zoom on the Fatigue Point: A Smooth Probabilistic Transition

We take a closer look at the transition from the non-overstretched to the overstretched regime. For this we ran several experiments on matrix NTRU instances with parameters $n = 113, \sigma^2 = \frac{2}{3}$ for several moduli q , with 100 runs each. We compare the DSD success ratio with our probabilistic concrete estimate. The results are shown in Fig. 8. Just as in Fig. 7 we see a shift between the experiment and prediction, which can again be explained by our SKR estimator being too pessimistic. Note however that while the discrepancy looks significant in this zoomed plot, it only emphasises a small error of about 2 block sizes between the experiments and our predictions. Ignoring this shift the shape of the predicted transition matches the experiments very well.

Acknowledgements. We would like to thank Martin Albrecht and Paul Kirchner for valuable comments and discussions. The research of L. Ducas was supported by the European Union’s H2020 Programme under PROMETHEUS project (grant 780701) and the ERC-StG-ARTICULATE project (no. 947821). W. van Woerden is funded by the ERC-ADG-ALGSTRONGCRYPTO project (no. 740972).

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). https://doi.org/10.1007/978-3-662-53018-4_6
- [AD21] Albrecht, M., Ducas, L.: Lattice attacks on NTRU and LWE: a history of refinements. Cryptology ePrint Archive, Report 2021/799 (2021). <https://eprint.iacr.org/2021/799>
- [ADPS16] Alkim, E., Ducas, L., Pöppelmann, T., Schwabe, P.: Post-quantum key exchange - a new hope. In: USENIX Security 2016, pp. 327–343. USENIX Association (2016)
- [AGVW17] Albrecht, M.R., Göpfert, F., Virdia, F., Wunderer, T.: Revisiting the expected cost of solving uSVP and applications to LWE. In: Takagi, T., Peyrin, T. (eds.) ASIACRYPT 2017. LNCS, vol. 10624, pp. 297–322. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-70694-8_11
- [Ajt99] Ajtai, M.: Generating hard instances of the short basis problem. In: Wiedermann, J., van Emde Boas, P., Nielsen, M. (eds.) ICALP 1999. LNCS, vol. 1644, pp. 1–9. Springer, Heidelberg (1999). https://doi.org/10.1007/3-540-48523-6_1
- [AWHT16] Aono, Y., Wang, Y., Hayashi, T., Takagi, T.: Improved progressive BKZ algorithms and their precise cost estimation by sharp simulator. In: Fischlin, M., Coron, J.-S. (eds.) EUROCRYPT 2016. LNCS, vol. 9665, pp. 789–819. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49890-3_30
- [BBC+20] Bernstein, D.J., et al.: NTRU Prime. Technical report, National Institute of Standards and Technology (2020)
- [BEF+17] Biase, J.-F., Espitau, T., Fouque, P.-A., Gélina, A., Kirchner, P.: Computing generator in cyclotomic integer rings. In: Coron, J.-S., Nielsen, J.B. (eds.) EUROCRYPT 2017. LNCS, vol. 10210, pp. 60–88. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-56620-7_3
- [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). https://doi.org/10.1007/978-3-642-45239-0_4
- [BSW18] Bai, S., Stehlé, D., Wen, W.: Measuring, simulating and exploiting the head concavity phenomenon in BKZ. In: Peyrin, T., Galbraith, S. (eds.) ASIACRYPT 2018. LNCS, vol. 11272, pp. 369–404. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-03326-2_13
- [CDH+20] Chen, C., et al.: NTRU. Technical report, National Institute of Standards and Technology (2020)

- [CDPR16] Cramer, R., Ducas, L., Peikert, C., Regev, O.: Recovering short generators of principal ideals in cyclotomic rings. In: Fischlin, M., Coron, J.-S. (eds.) EUROCRYPT 2016. LNCS, vol. 9666, pp. 559–585. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49896-5_20
- [CG05] Coglianese, M., Goi, B.-M.: MaTRU: a new NTRU-based cryptosystem. In: Maitra, S., Veni Madhavan, C.E., Venkatesan, R. (eds.) INDOCRYPT 2005. LNCS, vol. 3797, pp. 232–243. Springer, Heidelberg (2005). https://doi.org/10.1007/11596219_19
- [CJL16] Cheon, J.H., Jeong, J., Lee, C.: An algorithm for NTRU problems and cryptanalysis of the ggh multilinear map without a low-level encoding of zero. LMS J. Comput. Math. 19(A), 255–266 (2016)
- [CN11] Chen, Y., Nguyen, P.Q.: BKZ 2.0: better lattice security estimates. In: Lee, D.H., Wang, X. (eds.) ASIACRYPT 2011. LNCS, vol. 7073, pp. 1–20. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-25385-0_1
- [CS97] Coppersmith, D., Shamir, A.: Lattice attacks on NTRU. In: Fumy, W. (ed.) EUROCRYPT 1997. LNCS, vol. 1233, pp. 52–61. Springer, Heidelberg (1997). https://doi.org/10.1007/3-540-69053-0_5
- [DDGR20] Dachman-Soled, D., Ducas, L., Gong, H., Rossi, M.: LWE with side information: attacks and concrete security estimation. In: Micciancio, D., Ristenpart, T. (eds.) CRYPTO 2020. LNCS, vol. 12171, pp. 329–358. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-56880-1_12
- [DE+04] Diaconis, P., Erdős, P., et al.: On the distribution of the greatest common divisor. In: A festschrift for Herman Rubin, pp. 56–61. Institute of Mathematical Statistics (2004)
- [DKR+20] D’Anvers, J.: SABER. Technical report, National Institute of Standards and Technology (2020)
- [dt21a] The FPLLL development team. FPLLL, a lattice reduction library, Version: 5.4.1 (2021). <https://github.com/fplll/fplll>
- [dt21b] The FPLLL development team. fpylll, a Python wrapper for the fplll lattice reduction library, Version: 0.5.6 (2021). <https://github.com/fplll/fpylll>
- [EHKS14] Eisenträger, K., Hallgren, S., Kitaev, A., Song, F.: A quantum algorithm for computing the unit group of an arbitrary degree number field. In: 46th ACM STOC, pp. 293–302. ACM Press (2014)
- [GGH13] Garg, S., Gentry, C., Halevi, S.: Candidate multilinear maps from ideal lattices. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 1–17. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-38348-9_1
- [GGH+19] Genise, N., Gentry, C., Halevi, S., Li, B., Micciancio, D.: Homomorphic encryption for finite automata. In: Galbraith, S.D., Moriai, S. (eds.) ASIACRYPT 2019. LNCS, vol. 11922, pp. 473–502. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-34621-8_17
- [GS02] Gentry, C., Szydło, M.: Cryptanalysis of the revised NTRU signature scheme. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 299–320. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-46035-7_20
- [HPS98] Hoffstein, J., Pipher, J., Silverman, J.H.: NTRU: a ring-based public key cryptosystem. In: Buhler, J.P. (ed.) ANTS 1998. LNCS, vol. 1423, pp. 267–288. Springer, Heidelberg (1998). <https://doi.org/10.1007/BFb0054868>

- [KF17] Kirchner, P., Fouque, P.-A.: Revisiting lattice attacks on overstretched NTRU parameters. In: Coron, J.-S., Nielsen, J.B. (eds.) EUROCRYPT 2017. LNCS, vol. 10210, pp. 3–26. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-56620-7_1
- [LLL82] Lenstra, A.K., Lenstra, H.W., Jr., Lovász, L.: Factoring polynomials with rational coefficients. *Mathematische Annalen* **261**(4), 515–534 (1982)
- [LPR13] Lyubashevsky, V., Peikert, C., Regev, O.: On ideal lattices and learning with errors over rings. *J. ACM* **60**(6), 43:1–43:35 (2013). Preliminary version in Eurocrypt 2010
- [LW20] Lee, C., Wallet, A.: Lattice analysis on MinTRU problem. Cryptology ePrint Archive, Report 2020/230 (2020). <https://eprint.iacr.org/2020/230>
- [MS01] May, A., Silverman, J.H.: Dimension reduction methods for convolution modular lattices. In: Silverman, J.H. (ed.) CaLC 2001. LNCS, vol. 2146, pp. 110–125. Springer, Heidelberg (2001). https://doi.org/10.1007/3-540-44670-2_10
- [PAA+19] Poppelmann, T., et al.: NewHope. Technical report, National Institute of Standards and Technology (2019)
- [PS89] Papoulis, A., Saunders, H.: Probability, random variables and stochastic processes (1989)
- [PT08] Pataki, G., Tural, M.: On sublattice determinants in reduced bases. arXiv preprint [arXiv:0804.4014](https://arxiv.org/abs/0804.4014) (2008)
- [PV21] Postlethwaite, E.W., Virdia, F.: On the success probability of solving unique SVP via BKZ. In: Garay, J.A. (ed.) PKC 2021. LNCS, vol. 12710, pp. 68–98. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-75245-3_4
- [Reg04] Regev, O.: Quantum computation and lattice problems. *SIAM J. Comput.* **33**(3), 738–760 (2004). Preliminary version in FOCS 2002
- [Reg05] Regev, O., et al.: On lattices, learning with errors, random linear codes, and cryptography. In: 37th ACM STOC, pp. 84–93. ACM Press (2005)
- [SAB+20] Schwabe, P., et al.: CRYSTALS-KYBER. Technical report, National Institute of Standards and Technology (2020)
- [Sch87] Schnorr, C.-P.: A hierarchy of polynomial time lattice basis reduction algorithms. *Theor. Comput. Sci.* **53**, 201–224 (1987)
- [SS11] Stehlé, D., Steinfeld, R.: Making NTRU as secure as worst-case problems over ideal lattices. In: Paterson, K.G. (ed.) EUROCRYPT 2011. LNCS, vol. 6632, pp. 27–47. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-20465-4_4
- [SSTX09] Stehlé, D., Steinfeld, R., Tanaka, K., Xagawa, K.: Efficient public key encryption based on ideal lattices. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 617–635. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-10366-7_36
- [YD17] Yu, Y., Ducas, L.: Second order statistical behavior of LLL and BKZ. In: Adams, C., Camenisch, J. (eds.) SAC 2017. LNCS, vol. 10719, pp. 3–22. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-72565-9_1