



Universiteit
Leiden
The Netherlands

Lattice cryptography: isomorphisms & dual attacks

Pulles, L.N.

Citation

Pulles, L. N. (2026, September 23). *Lattice cryptography: isomorphisms & dual attacks*. Retrieved from <https://hdl.handle.net/1887/4309918>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/4309918>

Note: To cite this publication please use the final published version (if applicable).

Chapter 4

Accurate Score Prediction for Dual-Sieve Attacks

This chapter is based on joint work with L. Ducas published in the Journal of Cryptology [DP26], which combines published work [DP23c] with new material [DP23b].

We propose a heuristic for the output of a lattice sieve that seems weaker than the independent score heuristic. When recovering part of the secret in the Dual-Sieve attack, we use this heuristic to devise predictions for the score distribution associated to candidates: for correct candidates with noise drawn from any radial distribution, we predict scores using a central limit heuristic; for incorrect candidates, we predict scores by approximating the Voronoi cell by a ball.

In addition, we compare the predicted score distributions to experimental data, and observe that these predictions are qualitatively and quantitatively quite accurate. This allows one to accurately estimate the number of false-positives and false-negatives encountered in the Dual-Sieve attack, and it opens the door to a sound analysis thereof. One can now, for example, explore a way of mitigating the many false-positives.

4.1 Introduction

The previous chapter has revealed that the **independent score heuristic**, as used in the analysis of the Dual-Sieve attack [LW21, GJ21, MAT22], leads to an incorrect prediction for the score distribution of both **BDD** and uniform targets. In particular, the experiments in Section 3.6 suggest that this heuristic overestimates the success probability of distinguishing between a **BDD** target and a uniform target. Subsequently, by parametrising a Dual-Sieve attack according to this flawed heuristic, the attack will most likely fail.

For assessing the security of LWE-based cryptosystems, such as ML-KEM and ML-DSA, it is required to know precisely what the concrete cost is of running a state-of-the-art primal *or* dual attack. For this, we will need a new, reliable model for the Dual-Sieve attack, which possibly requires a newly devised heuristic.

In Section 3.5, the **independent score heuristic** contradicts in three regimes with other lemmata or well-tested heuristics, and thus seems to be the reason for the mispredictions in Section 3.6. These mispredictions indicate that there may be some correlation between the inner products $\langle \mathbf{w}, \mathbf{t} \rangle \pmod{1}$, while computing the score of a target $\mathbf{t}$.

The situation may be resolved by identifying the ‘confounding variable’ that causes dependencies between these random variables. In this case, the confounding variable seems to be the target $\mathbf{t}$, or more precisely its norm $\|\mathbf{t}\|$. For example, scaling the target by a factor α , multiplies *all* the inner products by α . Instead, by fixing the confounding variable $\mathbf{t}$, one may more reasonably hope the inner products $\langle \mathbf{w}, \mathbf{t} \rangle \pmod{1}$ become independent again. Fixing the target $\mathbf{t}$ leads very naturally to studying the score distribution of $\mathbf{t}$, where we now consider the randomness of the lattice and the set of short dual vectors.

4.1.1 Contributions

We propose new models for the score functions of **BDD** targets and uniform targets in Section 4.3. In Section 4.4, we then verify the predictions given by these models in an experimental setting, and observe that the predictions are quite accurate.

Accurate score distribution models (Section 4.3)

First, we will model the set of short dual vectors obtained by lattice sieving [BDGL16], using Heuristic 2.114 as being uniformly from some ball. Such a model is more accurate compared to that of existing literature, which approximates the distribution of lattice sieve vectors by i.i.d. Gaussians [MAT22, Assumption 4.4]. Although our choice somewhat complicates the Fourier analysis, it is still feasible thanks to the well-known Bessel functions, which we introduce in Section 4.2.

For a fixed target $\mathbf{t}$, we derive a novel *analytic expression* for the expectation value and variance of the score distribution. Note that in this case, each individual score has the same individual score distribution, so we can resort to Heuristic 2.3—which seems reasonable as there are exponentially many dual vectors—to reason that the sum of individual scores is normally distributed.

While the case of a particular target is of limited interest in our context of search-BDD, fortunately, we can easily extend the analysis to any other radial distribution, such as Gaussian or uniform in a ball.

Although it is notoriously hard to determine the exact shape of the Voronoi cell [Laa21], one can roughly approximate the Voronoi cell by a ball of same volume. This then approximates the uniform target distribution by a radial distribution, which allows us to predict the score distribution for uniform targets using the above prediction.

Experimental validation (Section 4.4)

The overall approach remains heuristic, but the previously identified issues have now been sidestepped, and we further confirm those score distribution predictions with extensive experiments. All the predictions that are made throughout this work have been extensively tested by large experiments to verify whether the proposed heuristics are reasonable in the context of solving decision-BDD.

Figures 4.1 and 4.2 show the predicted score distributions for BDD targets drawn uniformly from a sphere or ball, and from a Gaussian, using a lattice of dimension 90. In addition, the predicted score distribution for uniform targets is shown in Figure 4.3. This figure contains data of experiments in dimensions 40, 50, 60 and 70, each based on 2^{48} samples.

Finally, Figure 4.4 shows the score distribution for uniform targets when using a larger saturation radius inside lattice sieving. According to

our prediction, the ‘waterfall-floor’ phenomenon can be observed with fewer samples, compared to the default saturation radius of $r_{\text{sat}} = \sqrt{4/3}$. The experiments confirm this behaviour.

Open data

All the written code and the obtained data to generate the figures, can be found in the GitHub repository <https://github.com/ludopulles/AccurateScorePredictionDualSieveAttacks>. Experiments are written in Python and use the G6K and FPyLLL libraries [ADH⁺19, FPL25], and there is a binding to some C code to accelerate the FFT.

4.1.2 Related work

There are several related works [BW25, PS24, CDMT24], which are concurrent works with [DP23b].

The work of Bashiri and Wiemers [BW25] proposes an analysis of the variance of the score of the BDD target and the uniform targets. They use a significantly different approach, and obtain predictions compatible with the experimental data from Section 3.6.3. This does not directly allow us to conclude on false-negative nor false-positive rates, as it does not fully characterise the distribution.

The work of Pouly and Shen [PS24] proposes a provable variant of the dual attack using discrete Gaussian sampling in place of sieving vectors. Notably, their provable regime does not intersect the heuristic contradictory regime in Section 3.5.3. In fact, and quite interestingly, both regimes are separated by a constant factor of 2 on the length of the BDD target.

The work of Carrier, Debris-Alazard, Meyer-Hilfiger and Tillich, while mostly focusing on the case of statistical decoding for codes, also includes a short section on the case of lattices [CDMT24, Sec. 8]. Here, they try to predict, also using Bessel functions, the floor phenomenon for uniform targets modulo the lattice. They use a somewhat comparable but not identical reasoning and derivations as that of our Section 4.3.4.

4.2 Bessel function

The class of Bessel functions is defined as follows, and will be useful for the expected score of errors from a ball later on. The following definition

is based on [Wat1922, §3.3 (4)], which considers the more general setting where α and t may be complex.

Definition 4.1 (Bessel function). For any $\alpha > -\frac{1}{2}$, the *Bessel function (of the first kind) of order α* is given by

$$J_\alpha(t) = \frac{(t/2)^\alpha}{\sqrt{\pi} \cdot \Gamma(\alpha + \frac{1}{2})} \int_{-1}^1 e^{its} (1 - s^2)^{\alpha - \frac{1}{2}} ds,$$

for $t > 0$.

Observe that $J_\alpha(t)$ is real because $I(-s) = \overline{I(s)}$ where $I(s)$ denotes the integrand in the definition, cf. [AS64, 9.1.20].

It is well known that the Bessel function of some order α has an infinite number of positive roots. Let us denote with $j_{\alpha,n}$, the n th positive root of the Bessel function of order α .

Lemma 4.2 ([Wat1922, §15.81]). *The first positive root of the Bessel function of order α satisfies*

$$j_{\alpha,1} = \alpha + 1.855757 \dots \alpha^{1/3} + O(\alpha^{-1/3}),$$

as $\alpha \rightarrow \infty$, where $1.855757 \dots$ can be computed numerically up to arbitrary precision. In addition, we have $J_\alpha(x) > 0$ for all $0 < x < j_{\alpha,1}$.

For convenience later on, we define the following function.

Definition 4.3. The real-valued function ξ has domain $\mathbb{R}_{>-1/2} \times \mathbb{R}_{>0}$ and is given by

$$\xi(\alpha, x) = \frac{\Gamma(\alpha + 1)}{(\pi x)^\alpha} J_\alpha(2\pi x) , \quad (4.1)$$

for any (α, x) in its domain.

The image of the function $x \mapsto \xi(\alpha, x)$ is contained in $[-1, 1]$.

Remark 4.4. By using [AS64, 9.1.10], and the crude approximation $\Gamma(\alpha + k + 1)/\Gamma(\alpha + 1) = (\alpha + 1) \cdot (\alpha + 2) \cdot \dots \cdot (\alpha + k) \approx (\alpha + 1)^k$, we get the following approximation for small, positive x :

$$\begin{aligned} \xi(\alpha, x) &= \sum_{k=0}^{\infty} \frac{(-\pi^2 x^2)^k}{k!(\alpha + 1)(\alpha + 2) \dots (\alpha + k)} \\ &\approx \sum_{k=0}^{\infty} \frac{(-\pi^2 x^2)^k}{k!(\alpha + 1)^k} = e^{-\frac{\pi^2 x^2}{\alpha + 1}} . \end{aligned}$$

However, this approximation is not needed in computations, since ξ can be computed easily in many programming languages.

Remark 4.5. In high dimensions, one may get some numerical errors when computing ξ directly based on Equation (4.1). These issues are circumvented by computing ξ using following identity

$$\xi(\alpha, x) = {}_0F_1\left(; \alpha + 1; -\pi^2 x^2 \right) ,$$

where ${}_0F_1$ is the confluent hypergeometric function [AS64, 9.1.69]. For example, in `Python` one can use the function `hyp0f1` from the `mpmath` package, which implements ${}_0F_1$.

4.2.1 Relation to the Fourier transform

Bessel functions occur very naturally in the context of Fourier transforms, in particular as the Fourier transform of the indicator function of a ball.

Lemma 4.6. *The Fourier transform of $f = \mathbf{1}_{r\mathcal{B}^n}$ is given by*

$$\widehat{f}(\mathbf{x}) = \left(\frac{r}{\|\mathbf{x}\|} \right)^{\frac{n}{2}} \cdot J_{n/2}(2\pi r \|\mathbf{x}\|) = \text{Vol}_n(r\mathcal{B}^n) \cdot \xi\left(\frac{n}{2}, r\|\mathbf{x}\|\right) ,$$

where $\mathbf{x} \in \mathbb{R}^n$.

The following proof is based on [Fis13].

Proof. First, since f is radial, $\widehat{f}$ is also radial. Hence, let us do the proof only for $\mathbf{x} = (s, 0, \dots, 0)$. Then,

$$\begin{aligned} \widehat{f}(\mathbf{x}) &= \int_{r\mathcal{B}^n} e^{-2\pi i s y_1} d\mathbf{y} = \int_{-r}^r e^{-2\pi i s t} \text{Vol}_{n-1}\left(\sqrt{r^2 - t^2} \mathcal{B}^{n-1}\right) dt \\ &= \frac{r^n \pi^{\frac{n-1}{2}}}{\Gamma(\frac{n}{2} + \frac{1}{2})} \int_{-1}^1 e^{2\pi i r s u} (1 - u^2)^{\frac{n-1}{2}} du = \left(\frac{r}{s}\right)^{n/2} J_{\frac{n}{2}}(2\pi r s) . \end{aligned}$$

For the last equality, note that we have

$$\begin{aligned} \left(\frac{r}{\|\mathbf{x}\|} \right)^{\frac{n}{2}} J_{\frac{n}{2}}(2\pi r \|\mathbf{x}\|) &= \left(\frac{r}{\|\mathbf{x}\|} \right)^{\frac{n}{2}} \frac{(\pi r \|\mathbf{x}\|)^{\frac{n}{2}}}{\Gamma(\frac{n}{2} + 1)} \xi\left(\frac{n}{2}, r\|\mathbf{x}\|\right) \\ &= \text{Vol}_n(r\mathcal{B}^n) \xi\left(\frac{n}{2}, r\|\mathbf{x}\|\right) . \quad \square \end{aligned}$$

4.3 Score distribution models

In this section, we will gradually develop new predictions for the score distributions of both **BDD** targets and uniform targets, which hopefully match closely experimental data from Section 3.6.

In Section 4.3.1, we derive analytic expressions for the mean and variance of the individual score function $f_{\mathbf{w}}(\mathbf{t})$. In Section 4.3.2, we derive expressions for the exact mean and variance of the score distribution for a fixed **BDD** target based on Heuristic 2.114. Here, we present a new heuristic that says the score distribution for a single target $\mathbf{t}$ is Gaussian and only depends on $d(\mathbf{t}, \Lambda)$, when the probability is taken over the lattice. Using this heuristic and by integrating over the radius, we predict in Section 4.3.3 the score distribution for a radial target distribution, such as uniform in a ball and Gaussian. Moreover, by approximating the Voronoi cell by a ball of volume $\det(\Lambda)$, we can also predict the score distribution for an error uniform modulo Λ , which is done in Section 4.3.4.

Throughout this section, let us assume Λ has full rank and $\det(\Lambda) = 1$ without loss of generality, because the ambient space can be restricted to $\text{span}(\Lambda)$, see Proposition 2.59, and any lattice Λ can be rescaled to one of unit-volume.

We strongly advise the reader to think carefully for which lattice Λ to use the predictions of this section. Namely, in many of the recent dual attacks, it is a projected sublattice of Λ_{LWE} having rank- β_{sieve} , see Remark 3.2.

4.3.1 Individual score function

First, let us fix an arbitrary vector $\mathbf{t} \in \mathbb{R}^n$, and look at the score function $f_{\mathbf{w}}(\mathbf{t}) = \cos(2\pi \langle \mathbf{w}, \mathbf{t} \rangle)$, where $\mathbf{w}$ is sampled uniformly from a sphere. In this case, the mean and variance are expressible in terms of Bessel functions, and thus in terms of ξ from Definition 4.3.

Lemma 4.7. *Given an arbitrary vector $\mathbf{t} \in \mathbb{R}^n$ and $r \in \mathbb{R}_{>0}$, when the random variable $\mathbf{w}$ follows the uniform distribution on the $(n-1)$ -*

dimensional sphere of radius r , we have

$$\begin{aligned} \mathbb{E}_{\mathbf{w} \leftarrow \mathcal{U}(r\mathcal{S}^{n-1})}[f_{\mathbf{w}}(\mathbf{t})] &= \xi\left(\frac{n}{2} - 1, r\|\mathbf{t}\|\right), \quad \text{and} \\ \mathbb{V}_{\mathbf{w} \leftarrow \mathcal{U}(r\mathcal{S}^{n-1})}[f_{\mathbf{w}}(\mathbf{t})] &= \frac{1}{2} + \frac{1}{2}\xi\left(\frac{n}{2} - 1, 2r\|\mathbf{t}\|\right) - \xi\left(\frac{n}{2} - 1, r\|\mathbf{t}\|\right)^2. \end{aligned}$$

Proof. The expectation value follows directly from a classic result from Fourier Analysis, see e.g. [GS64, p. 198] or [SW71, p. 154], as we have

$$\begin{aligned} \mathbb{E}_{\mathbf{w} \leftarrow \mathcal{U}(r\mathcal{S}^{n-1})}[f_{\mathbf{w}}(\mathbf{t})] &= \frac{1}{\text{Vol}_n(r\mathcal{S}^{n-1})} \int_{r\mathcal{S}^{n-1}} e^{2\pi i \cdot \langle \mathbf{w}, \mathbf{t} \rangle} d\mathbf{w} \\ &= \frac{\Gamma(\frac{n}{2}) \cdot J_{\frac{n}{2}-1}(2\pi r\|\mathbf{t}\|)}{(\pi r\|\mathbf{t}\|)^{\frac{n}{2}-1}}. \end{aligned}$$

The trigonometric identity $f_{\mathbf{w}}(\mathbf{t})^2 = \frac{1}{2} + \frac{1}{2} \cos(4\pi \langle \mathbf{w}, \mathbf{t} \rangle) = \frac{1}{2} + \frac{1}{2} f_{\mathbf{w}}(2 \cdot \mathbf{t})$ allows us to derive a similar expression for the variance, by observing

$$\begin{aligned} \mathbb{V}_{\mathbf{w} \leftarrow \mathcal{U}(r\mathcal{S}^{n-1})}[f_{\mathbf{w}}(\mathbf{t})] &= \mathbb{E}_{\mathbf{w}}[f_{\mathbf{w}}(\mathbf{t})^2] - \mathbb{E}_{\mathbf{w}}[f_{\mathbf{w}}(\mathbf{t})]^2 \\ &= \frac{1}{2} + \frac{1}{2} \mathbb{E}_{\mathbf{w}}[f_{\mathbf{w}}(2 \cdot \mathbf{t})] - \mathbb{E}_{\mathbf{w}}[f_{\mathbf{w}}(\mathbf{t})]^2, \end{aligned}$$

and reusing the result for the expectation value. $\square$

Remark 4.8. The individual score function has been studied before, for example by Laarhoven and Walter [LW21, Section 4.1]. However, they approximate the distribution for $\mathbf{w}$ by $\mathcal{N}_n(0, r^2/n)$, yielding only an approximation for the expectation value and variance, whereas the above lemma has an exact expression for both. Still, if one is inclined to work with a crude approximation, by Remark 4.4, the expectation in the above lemma is for small values r approximated by $\exp\left(-\frac{2\pi^2\|\mathbf{t}\|^2 r^2}{n}\right)$, which matches the approximated expectation in [LW21].

As we will reuse the mean and variance for the ball, let us introduce the following notation.

Definition 4.9. For $n \in \mathbb{Z}_{\geq 1}$, $r_p, r_d \in \mathbb{R}_{>0}$ let us denote

$$\begin{aligned} E_{n,r_d}(r_p) &= \xi\left(\frac{n}{2}, r_d \cdot r_p\right), \quad \text{and} \\ V_{n,r_d}(r_p) &= \frac{1}{2} + \frac{1}{2} E_{n,r_d}(2 \cdot r_p) - (E_{n,r_d}(r_p))^2. \end{aligned}$$

The result in Lemma 4.7 can be translated easily to a uniform distribution on a ball, because sampling uniformly from the n -dimensional ball can be done by first sampling from a $(n + 1)$ -dimensional sphere and then dropping the last two coordinates [BGMN05, Corollary 4].

Corollary 4.10. *Given an arbitrary vector $\mathbf{t} \in \mathbb{R}^n$ and $r \in \mathbb{R}_{>0}$, when the random variable $\mathbf{w}$ follows the uniform distribution on the n -dimensional ball of radius r , we have*

$$\mathbb{E}_{\mathbf{w} \leftarrow \mathcal{U}(r\mathcal{B}^n)}[f_{\mathbf{w}}(\mathbf{t})] = E_{n,r}(\|\mathbf{t}\|), \quad \text{and} \quad \mathbb{V}_{\mathbf{w} \leftarrow \mathcal{U}(r\mathcal{B}^n)}[f_{\mathbf{w}}(\mathbf{t})] = V_{n,r}(\|\mathbf{t}\|) .$$

4.3.2 Total score function

The next step is to determine the expectation value and variance of the total score $f_{\mathcal{W}}(\mathbf{t})$, in the case that the set $\mathcal{W}$ is obtained by lattice sieving. It is now necessary to take into account the randomness with which the lattice Λ is sampled, because the lattice has influence on the set of dual vectors $\mathcal{W} \subseteq r_d\mathcal{B}^n \cap \Lambda^\vee$.

In this section, let us fix a target $\mathbf{t} \in \mathbb{R}^n$ initially. We will study the score distribution for this target $\mathbf{t}$ by considering the randomness of the lattice and the set of dual vectors $\mathcal{W}$.

Upon first reading the rest of this section, the reader is encouraged to think that the set of dual vectors is $\mathcal{W} = r_d\mathcal{B}^n \cap \Lambda^\vee$, or rather: half of them to break the negation symmetry (recall Section 2.6). To argue about the dual vectors we will make use of Heuristic 2.114, which can be used for more general saturation parameters. When sieving in practice, a saturation ratio $f_{\text{sat}} < 1$ is used to obtain enough dual vectors in the ball, and it is very natural to believe that different saturation parameters do not change the score distribution significantly, other than having an effect on N , the expected size of $\mathcal{W}$. Although all the shortest vectors are usually found, most of the missing $1 - f_{\text{sat}}$ portion of vectors is around the boundary. Still, we believe this does not noticeably affect the score distribution for reasonably large f_{sat} .

By using Heuristic 2.114, the $\mathbf{w} \in \mathcal{W}$ are i.i.d. samples and match the distribution in Corollary 4.10, too. Because heuristically the scores $(f_{\mathbf{w}}(\mathbf{t}))_{\mathbf{w} \in \mathcal{W}}$ are i.i.d., we expect that the total score is Gaussian for large $|\mathcal{W}|$ by Heuristic 2.3. This argument results in the following heuristic.

Heuristic 4.11. *Fix a target $\mathbf{t} \in \mathbb{R}^n$ of norm $\|\mathbf{t}\| \in (0, \text{GH}(n))$. The score distribution $f_{\mathcal{W}}(\mathbf{t})$ is Gaussian with mean $N \cdot E_{n,r_d}(\|\mathbf{t}\|)$ and variance*

$N \cdot V_{n,r_d}(\|\mathbf{t}\|)$ when sampling Λ and obtaining N dual vectors $\mathcal{W} \leftarrow \text{SIEVE}(\Lambda^\vee, r_{\text{sat}}, f_{\text{sat}})$ from a sieve algorithm, where E_{n,r_d} and V_{n,r_d} are defined in Definition 4.9 and $r_d = r_{\text{sat}} \text{GH}(n)$. In particular, the CDF of the score distribution is as follows:

$$\Pr_{\Lambda, \mathcal{W}}[f_{\mathcal{W}}(\mathbf{t}) \leq x] = \frac{1}{2} + \frac{1}{2} \operatorname{erf} \left(\frac{x - N \cdot E_{n,r_d}(\|\mathbf{t}\|)}{\sqrt{2N \cdot V_{n,r_d}(\|\mathbf{t}\|)}} \right). \quad (4.2)$$

Heuristic Justification. First, based on Heuristic 2.114, we assume, over the randomness of Λ and $\mathcal{W}$, that $\mathcal{W}$ is a set of i.i.d. $\mathbf{w} \leftarrow \mathcal{U}(r_d \mathcal{B}^n)$ of size N . Then, Corollary 4.10 proves that each individual score $(f_{\mathbf{w}})_{\mathbf{w} \in \mathcal{W}}$ has mean $E_{n,r_d}(\|\mathbf{t}\|)$ and variance $V_{n,r_d}(\|\mathbf{t}\|)$.

Because the $(f_{\mathbf{w}}(\mathbf{t}))_{\mathbf{w} \in \mathcal{W}}$ are i.i.d. individual scores, we can use Heuristic 2.3 here. In this case, the total score distribution will have mean $N \cdot E_{n,r_d}(\|\mathbf{t}\|)$ and variance $N \cdot V_{n,r_d}(\|\mathbf{t}\|)$. $\triangle$

Remark 4.12. The heuristic should not be used for targets $\mathbf{t}$ of norm above $\text{GH}(n)$. Given a random lattice, one expects there exists $\mathbf{t}' \in \mathbf{t} + \Lambda$ of norm $\|\mathbf{t}'\| \leq \text{GH}(n)$ by the Gaussian heuristic. Taking $\mathbf{t}' \in \mathbf{t} + \Lambda$ of minimal norm we have $f_{\mathcal{W}}(\mathbf{t}) = f_{\mathcal{W}}(\mathbf{t}')$, but using Heuristic 2.114 and Corollary 4.10 yields a *different* expected score of $E_{n,r_d}(\|\mathbf{t}'\|)$ for $\mathbf{t}'$, which is much higher than $E_{n,r_d}(\|\mathbf{t}\|) \approx 0$ if $\|\mathbf{t}'\| < 0.5 \text{GH}(n)$, say. Hence, use of Heuristic 2.114, which forgets that $\mathcal{W}$ are dual vectors, results in contradictory predictions in the case $\|\mathbf{t}\| > \text{GH}(n)$.

On the other hand, given a target $\mathbf{t}$ of norm $\alpha \text{GH}(n)$ for a constant $\alpha < 1$, we expect $\|\mathbf{t}\| = d(\mathbf{t}, \Lambda)$ holds for most random lattices as $n \rightarrow \infty$ by Heuristic Claim 2.48, avoiding this issue with Heuristic 2.114.

Interestingly, the expected score $f_{\mathcal{W}}(\mathbf{t})$ is very similar to the score when one takes $\mathcal{W} = r_d \mathcal{B}^n \cap \Lambda^\vee$. Without any heuristics, Theorem 2.75 and Lemma 4.6 yield the following identity:

$$\begin{aligned} f_{\mathcal{W}}(\mathbf{t}) &= \sum_{\mathbf{w} \in \Lambda^\vee} \mathbf{1}_{r_d \mathcal{B}^n}(\mathbf{w}) e^{2\pi i \langle \mathbf{w}, \mathbf{t} \rangle} \\ &= \text{Vol}_n(r_d \mathcal{B}^n) \cdot \sum_{\mathbf{v} \in \Lambda} \xi \left(\frac{n}{2}, r_d \|\mathbf{v} + \mathbf{t}\| \right). \end{aligned} \quad (4.3)$$

For $r_p \ll \text{GH}(n)$, the vector $\mathbf{v} = \mathbf{0}$ gives the main contribution of $\xi(\frac{n}{2}, r_d r_p)$ to the summation in Equation (4.3), because $\xi(n/2, -)$ is a rapidly decaying function.

At this point, one could argue that Heuristic 4.11 is not better at predicting a score distribution than the **independent score heuristic** because it still uses Heuristic 2.3. However, observe that Heuristic 4.11 uses Heuristic 2.3 for **i.i.d.** $f_{\mathbf{w}}(\mathbf{t})$, after first fixing the target $\mathbf{t}$ independent of the lattice.

On the other hand, usage of the **independent score heuristic** leads to a prediction that uses Heuristic 2.3 for the individual scores $f_{\mathbf{w}}(\mathbf{t})$ *irrespective of the target distribution for $\mathbf{t}$* , which ignores the dependency of its norm on the mean score. Heuristic 4.11, being restricted to a fixed target norm, thus seems reasonable. However, large experiments are ultimately needed to gain confidence in the heuristic, and these can be found in Section 4.4.2.

Lastly, the following lemma contains an upper bound on the norm of a **BDD** target such that it is distinguishable from targets uniform modulo the lattice.

Lemma 4.13. *If $r_p \cdot r_d < \frac{n}{4\pi}$, then $E_{n,r_d}(r_p) > 0$.*

Proof. The first zero of the function $x \mapsto \xi(\frac{n}{2}, x)$ is at

$$x = \frac{1}{2\pi} j_{n/2,1} > \frac{n}{4\pi} > r_p r_d ,$$

where we use Lemma 4.2 in the first inequality. $\square$

Using the approximation $\text{GH}(n) \approx \sqrt{\frac{n}{2\pi e}}$, note the condition on $r_p r_d$ translates to roughly $r_p r_d < \frac{e}{2} \text{GH}(n)^2$ for large n .

Note that above lemma is sharp, i.e., if we take $r_p r_d = \alpha n$ for some constant $\alpha > \frac{1}{4\pi}$, then $r_p r_d = \frac{n}{4\pi} + \left(\alpha - \frac{1}{4\pi}\right)n$. Lemma 4.2 implies there exists some $N \in \mathbb{N}$ such that for all $n \geq N$ we have $r_p r_d > \frac{1}{2\pi} j_{n/2,1}$, so there will be some $n \geq N$ giving a negative expectation score.

4.3.3 Radial error distributions

The predictions for a particular error can now be extended to any error distribution χ that is radial, i.e., the **PDF** is the same at two points of the same norm. In particular, given the radial distribution χ , let us write $f(r)$ to denote the probability (density) of a sample $\mathbf{e} \leftarrow \chi$ having norm r , i.e.,

$$f(r) = \frac{d}{dr} \Pr_{\mathbf{t} \leftarrow \chi} [\|\mathbf{t}\| \leq r].$$

We refer to $f(r)$ as the *norm distribution* of χ .

Using Heuristic 4.11 we arrive at the following claim.

Heuristic Claim 4.14. *Let χ be an error distribution with norm distribution $f(r)$, such that the support of χ is mostly contained in $\text{GH}(n)\mathcal{B}^n$. The score distribution $f_{\mathcal{W}}(\mathbf{t})$ for an error vector $\mathbf{t} \leftarrow \chi$ has the following CDF:*

$$\Pr_{\Lambda, \mathbf{t} \leftarrow \chi} [f_{\mathcal{W}}(\mathbf{t}) \leq x] = \frac{1}{2} + \frac{1}{2} \int_0^\infty \text{erf} \left(\frac{x - N \cdot E_{n, r_d}(r)}{\sqrt{2N \cdot V_{n, r_d}(r)}} \right) f(r) \, dr, \quad (4.4)$$

where the probability is taken over randomness from sampling Λ and obtaining N dual vectors $\mathcal{W} \leftarrow \text{SIEVE}(\Lambda^\vee, r_{\text{sat}}, f_{\text{sat}})$ from a sieve algorithm.

There are now two radial distributions of particular interest: Gaussian and the uniform distribution on the ball.

Gaussian error

The norm distribution of samples from $\mathcal{N}_n(0, 1)$, is the χ -distribution of order n , which has a PDF given by:

$$f(r; n) = \frac{r^{n-1} \exp\left(-\frac{r^2}{2}\right)}{2^{\frac{n}{2}-1} \Gamma\left(\frac{n}{2}\right)} \quad (r \in \mathbb{R}_{>0}).$$

Note that one should not get confused with the (perhaps more well-known) χ^2 -distribution, which considers the *squared norm*.

Now instantiating Heuristic Claim 4.14 for $\mathcal{N}_n(0, \sigma^2)$, the Gaussian error distribution of parameter $\sigma > 0$, yields the following claim.

Heuristic Claim 4.15. *Let $\sigma \in (0, \text{GH}(n)/\sqrt{n})$, and let $f(x; n)$ be the PDF of the χ -distribution. The score distribution $f_{\mathcal{W}}(\mathbf{t})$ for an error $\mathbf{t} \leftarrow \mathcal{N}_n(0, \sigma^2)$ has the following CDF:*

$$\Pr_{\substack{\Lambda, \mathcal{W}, \\ \mathbf{t} \leftarrow \mathcal{N}_n(0, \sigma^2)}} [f_{\mathcal{W}}(\mathbf{t}) \leq x] = \frac{1}{2} + \frac{1}{2} \int_0^\infty \text{erf} \left(\frac{x - N \cdot E_{n, r_d}(r)}{\sqrt{2N \cdot V_{n, r_d}(r)}} \right) f\left(\frac{r}{\sigma}; n\right) \frac{dr}{\sigma}, \quad (4.5)$$

where the probability is taken over randomness from sampling Λ and obtaining N dual vectors $\mathcal{W} \leftarrow \text{SIEVE}(\Lambda^\vee, r_{\text{sat}}, f_{\text{sat}})$ from a sieve algorithm.

Note that the χ -distribution is very concentrated around $\sigma\sqrt{n}$, so the best numerical approximations are obtained when the numerical integration is giving special attention to the region around $\sigma\sqrt{n}$ (see Section 4.4.1).

Error uniform from a ball

Consider the distribution $U(r_p \mathcal{B}^n)$ for some $r_p > 0$. In this case,

$$\Pr_{\mathbf{t} \leftarrow U(r_p \mathcal{B}^n)}[\|\mathbf{t}\| \leq r] = r^n / r_p^n$$

for any $r \in [0, r_p]$, while this equals 1 for $r \geq r_p$. Strictly speaking, the **CDF** of the norm distribution is not differentiable at $r = r_p$, so $f(r_p)$ is not defined. Still, one can mitigate this issue by integrating from 0 to r_p in Equation (4.4), yielding the following claim.

Heuristic Claim 4.16. *Let $r_p \in (0, \text{GH}(n))$. The score distribution $f_{\mathcal{W}}(\mathbf{t})$ for an error $\mathbf{t} \leftarrow U(r_p \mathcal{B}^n)$ has the following **CDF**:*

$$\Pr_{\substack{\Lambda, \mathcal{W}, \\ \mathbf{t} \leftarrow U(r_p \mathcal{B}^n)}}[f_{\mathcal{W}}(\mathbf{t}) \leq x] = \frac{1}{2} + \frac{1}{2} \int_0^{r_p} \text{erf}\left(\frac{x - N \cdot E_{n,r_d}(r)}{\sqrt{2N \cdot V_{n,r_d}(r)}}\right) \frac{nr^{n-1} dr}{r_p^n}, \quad (4.6)$$

where the probability is taken over randomness from sampling Λ and obtaining N dual vectors $\mathcal{W} \leftarrow \text{SIEVE}(\Lambda^\vee, r_{\text{sat}}, f_{\text{sat}})$ from a sieve algorithm.

4.3.4 Error uniform modulo lattice

Given uniform targets $\mathbf{t} \leftarrow U(\mathbb{R}^n / \Lambda)$, the inner product $\langle \mathbf{w}, \mathbf{t} \rangle \pmod{1}$ follows the uniform distribution on $\mathbb{R}/\mathbb{Z}$, for any (nonzero) dual vector $\mathbf{w} \in \Lambda^\vee$. Since experiments in Section 3.6 show that the total score function $f_{\mathcal{W}}(\mathbf{t})$ is not Gaussian, we cannot use Heuristic 2.3 for the individual scores $(f_{\mathbf{w}}(\mathbf{t}))_{\mathbf{w} \in \mathcal{W}}$ as it would imply $f_{\mathcal{W}}(\mathbf{t})$ is Gaussian, which is contradictory to experiments. In particular, Section 3.5.2 reveals that use of Heuristic 2.3 would underestimate the probability on an exceptionally high score, because it would not consider the probability of a uniform target landing close to Λ , to which a distinguisher assigns a high score.

In this section, we want to derive predictions for uniform targets, without using the **independent score heuristic**. The above motivation

shows that having a high score is driven by a target close to the lattice, so it seems important to know how the distance $d(\mathbf{t}, \Lambda)$ from a uniform target to a lattice is distributed to predict the score distribution.

First, because the Voronoi cell tiles the space, assume targets are sampled from $U(\mathcal{V}(\Lambda))$. The score for such targets would ideally be predicted by using Heuristic Claim 4.14, but $U(\mathcal{V}(\Lambda))$ is only a radial distribution in dimension $n = 1$. Instead, because the dual vectors already have rotational invariance when using Heuristic 2.114, we could approximate the score distribution by still using Heuristic Claim 4.14 with norm distribution $f(r) = dF(r)/dr$, where

$$F(r) = \Pr_{\mathbf{t} \leftarrow U(\mathbb{R}^n/\Lambda)} [d(\mathbf{t}, \Lambda) \leq r] = \text{Vol}_n(\mathcal{V}(\Lambda) \cap r\mathcal{B}^n). \quad (4.7)$$

However, already computing the Voronoi cell is considered a difficult task [MV10]. Note that we have $\frac{1}{2}\lambda_1\mathcal{B}^n \subseteq \mathcal{V}(\Lambda) \subseteq \mu(\Lambda)\mathcal{B}^n$, where $\mu(\Lambda)$ is the covering radius of the lattice. Hence, $F(r) = \text{Vol}_n(r\mathcal{B}^n)$ for $r \leq \frac{1}{2}\lambda_1$. For radii $r \in (\frac{1}{2}\lambda_1, \mu(\Lambda))$, there is no easy expression for $F(r)$ because the ball of radius r is not necessarily contained in the Voronoi cell. However, we still have the upper bound,

$$F(r) \leq \text{Vol}_n(r\mathcal{B}^n), \quad (4.8)$$

which is not much smaller than what is derived in Heuristic Claim 3.13 for $r = (1 - \varepsilon)\text{GH}(n)$ with $\varepsilon > 0$. The upper bound in Equation (4.8) can thus be seen as a first order approximation of $F(r)$. The following heuristic implies $F(r)$ equals the upper bound in Equation (4.8).

Heuristic 4.17. *Let $\Lambda \subset \mathbb{R}^n$ be a random full-rank unit-volume lattice. Then,*

$$\mathcal{V}(\Lambda) = \text{GH}(n) \cdot \mathcal{B}^n.$$

Using Heuristic 4.17 and Heuristic Claim 4.16, directly results in the following claim.

Heuristic Claim 4.18. *The score distribution $f_{\mathcal{W}}(\mathbf{t})$ for an error vector $\mathbf{t} \leftarrow U(\mathbb{R}^n/\Lambda)$ has the following CDF:*

$$\Pr_{\substack{\Lambda, \mathcal{W}, \\ \mathbf{t} \leftarrow U(\mathbb{R}^n/\Lambda)}} [f_{\mathcal{W}}(\mathbf{t}) \leq x] = \frac{1}{2} + \frac{1}{2} \int_0^{\text{GH}(n)} \text{erf} \left(\frac{x - N \cdot E_{n,r_d}(r)}{\sqrt{2N \cdot V_{n,r_d}(r)}} \right) \frac{nr^{n-1} dr}{\text{GH}(n)^n}, \quad (4.9)$$

where the probability is taken over randomness from sampling Λ and obtaining N dual vectors $\mathcal{W} \leftarrow \text{SIEVE}(\Lambda^\vee, r_{\text{sat}}, f_{\text{sat}})$ from a sieve algorithm.

This heuristic claim predicts the waterfall and floor shapes in the score distribution for uniform targets.

Waterfall shape: Suppose x is close to zero. If r is much smaller than $\text{GH}(n)$, the integrand is negligible because the expected score $N \cdot E_{n,r_d}(r)$ is much higher than x . On the other hand, for $r \approx \text{GH}(n)$, we have $N \cdot E_{n,r_d}(r) \approx 0$ and $N \cdot V_{n,r_d}(r) \approx N/2$. By using these approximations for $r > (1 - \varepsilon) \text{GH}(n)$ given some constant $\varepsilon > 0$, we may approximate the integral in Equation (4.9) by $(1 - e^{-n/\varepsilon}) \cdot \text{erf}(x/\sqrt{N})$. This approximation approaches $\text{erf}(x/\sqrt{N})$ as n tends to infinity. Hence, the survival function has a waterfall shape around $x = 0$.

Floor shape: Let us consider the survival function of $f_{\mathcal{W}}(\mathbf{t})$:

$$\Pr_{\substack{\Lambda, \mathcal{W}, \\ \mathbf{t} \leftarrow \mathbf{U}(\mathbb{R}^n/\Lambda)}} [f_{\mathcal{W}}(\mathbf{t}) > x] = \frac{1}{2} \int_0^{\text{GH}(n)} \text{erfc} \left(\frac{x - N \cdot E_{n,r_d}(r)}{\sqrt{2N \cdot V_{n,r_d}(r)}} \right) \frac{nr^{n-1} dr}{\text{GH}(n)^n}.$$

For large x , say $c\sqrt{N}$ for some $c > 10$, the biggest contribution to the integral comes from $r \approx r_0$ where r_0 is a solution of the equation $N \cdot E_{n,r_d}(r_0) = x$, because here erfc evaluates to ≈ 1 . However, for $r \gg r_0$, the integrand drops to zero as erfc decreases much quicker than r^{n-1} , while for $r \ll r_0$, the integrand also drops to zero because r^{n-1} is simply too small.

4.4 Experiments

In this section, we provide further substantiation of the concrete predictions made in Section 4.3, in particular the predictions in Equations (4.2), (4.5), (4.6) and (4.9), with experimental support. By running experiments, we can verify whether the used heuristics lead to conclusions that precisely match practice.

Similar to Section 3.6, we take $\mathcal{W}$ to be the full output of a lattice sieve on $\Lambda^\vee$ in the experiments. We will compare the score distribution for three possible BDD target distributions with their respective prediction from Section 4.3: uniform from a sphere, uniform from a ball and Gaussian. In addition, we compare the score distribution for uniform targets with the predicted CDF of Equation (4.9).

4.4.1 Implementation details

In the experiments, we used the **G6K** software [ADH⁺19] to obtain the dual vectors. We used Python on a high-level together with a binding to some C/C++ code to speed up **BDD** sampling.

BDD targets

The script `bdd_sample.py` samples the three **BDD** score distributions by first sampling a random q -ary lattice from a matrix of dimension $n \times n/2$ (n is only even), then running a lattice sieve to acquire many dual vectors, and finally computing scores for many samples. To make the implementation easier, the sampled q -ary lattice is used as the dual lattice $\Lambda^\vee$, because we do not work with primal lattice vectors. The prime used is $q = 3329$.

For the experiments, we sample **BDD** targets from a distribution with norms concentrated around $f_{\text{GH}} \cdot \text{GH}(n)/\sqrt{q}$ for some parameter $f_{\text{GH}} \in (0, 1]$, where Λ is the primal lattice. Note that the **Gaussian heuristic** predicts $\lambda_1(\Lambda) \approx \text{GH}(n)/\sqrt{q}$ holds, because the dual lattice has determinant $q^{n/2}$.

More specifically, for a Gaussian sample, we simply sample $\mathbf{x} \leftarrow \mathcal{N}_n(0, \sigma^2)$ with $\sigma = \frac{f_{\text{GH}} \cdot \text{GH}(n)}{\sqrt{n \cdot q}}$. Note that we need a factor of $1/\sqrt{n}$ such that $\|\mathbf{x}\|$ is concentrated around the target length $f_{\text{GH}} \cdot \text{GH}(n)/\sqrt{q}$. For a sample $\mathbf{y}$ uniform on the sphere, we reuse the same Gaussian sample $\mathbf{x}$, and compute

$$\mathbf{y} = \frac{f_{\text{GH}} \cdot \text{GH}(n)}{\sqrt{q}} \cdot \frac{\mathbf{x}}{\|\mathbf{x}\|},$$

which is distributed uniformly on the sphere of radius $f_{\text{GH}} \cdot \text{GH}(n)/\sqrt{q}$. For a sample uniform in the ball, we take a sample uniformly from the $(n+1)$ -dimensional sphere in $\mathbb{R}^{n+2}$ and drop the last two coordinates, making use of [BGMN05, Corollary 4]. We only generated 2 more Gaussian samples, and reused the n from $\mathbf{x}$ here. The script `bdd_sample.py` ran to collect 100 000 samples.

The script `bdd_predict.py` then uses this data to plot the experimental data, the prediction from Section 4.3 and the heuristic prediction. The predictions from Section 4.3 required evaluating the Bessel function and an integration, for which we used `hyp0f1` (see Remark 4.5) and `quad` respectively, both in the Python package `mpmath`.

For the integration, the Gaussian prediction was numerically more accurate after splitting the interval $(0, \infty)$ into two at the expected length. For more details, see ‘Highly variable functions’ in `quad`’s [documentation](#). Predictions for uniform samples from the ball of radius r were obtained by integrating from $0.001r$ to r to prevent the singularity around 0.

Uniform targets

For targets uniform in the lattice, the scripts `unif_sample.py` and `unif_predict.py` were used similar to the **BDD** case. One should pay special attention to the integration errors caused by numerical integration when calculating the predictions of Heuristic Claim 4.18. Namely, the numerical integration should give an answer with a small *relative error*, because the output can be very small, as can be seen in Figure 4.3. Thus, for numerical integration we used the function `quad` from the Python package `SciPy` with parameters `epsabs=0` and `epsrel=0.001`.

4.4.2 BDD targets

The experimental results for **BDD** targets can be found in Figure 4.1 and Figure 4.2. The former samples **BDD** targets at expected distance of $0.7 \text{ GH}(n)$ from the lattice, while the latter samples at distance $0.5 \text{ GH}(n)$ from the lattice. Here, a saturation ratio of $f_{\text{sat}} = 0.99$ and saturation radius of $r_{\text{sat}} = \sqrt{4/3}$ was used, so $N = \lceil \frac{1}{2} f_{\text{sat}} \cdot r_{\text{sat}}^{90} \rceil = 207419$ dual vectors were used.

It is clear that the predictions made in Sections 4.3.2 and 4.3.3 give accurate estimates on the score distribution for **BDD** targets that are from a sphere, from a ball or Gaussian. Moreover, these experiments show that Heuristic 4.11 appears to be a reasonable heuristic to use in this regime.

4.4.3 Uniform targets

Figure 4.3 compares experimental data with the prediction of the score distribution for uniform targets. Here, the lattice sieve used a saturation ratio of $f_{\text{sat}} = 0.9$ and saturation radius of $r_{\text{sat}} = \sqrt{4/3} \approx 1.1547$. The experimental data for uniform targets is acquired independently of that in Section 3.6.2, which used $2N$ dual vectors. For each of the four experimental curves of Figure 4.3, we ran the script `code/unif_sample.py`, which took slightly more than one day on a server with 40 physical cores.




Figure 4.1: CDF of the score for three BDD target distributions in dimension $n = 90$. The prior prediction uses the *independent score heuristic*, while ‘Prediction (this work)’ is based on Sections 4.3.2 and 4.3.3. Each experimental curve contains 10^5 samples, which are obtained with `code/bdd_sample.py` and are listed in `data/predictions_n90_ghf0.7.csv` of the auxiliary files.




Figure 4.2: CDF of the score for three BDD target distributions in dimension $n = 90$. The prior prediction uses the *independent score heuristic*, while ‘Prediction (this work)’ is based on Sections 4.3.2 and 4.3.3. Each experimental curve contains 10^5 samples, which are obtained with `code/bdd_sample.py` and are listed in `data/predictions_n90_ghf0.5.csv` of the auxiliary files.

Note that in dimension 50, 60 and 70, the uniform prediction seems to be very close to the experimental data. The right tail of the experimental data depends on extremely rare events (happening once in 2^{48} trials), so a few of the most-right data points can change a little bit in another run. We believe if more samples were taken, the experimental data would get closer to the prediction, but then the experiment would require a runtime of multiple days on our server.

In dimension 40, it seems that the experiments give scores higher than expected by our prediction. However, the sieve provides only $N = \lceil \frac{1}{2} f_{\text{sat}} \cdot r_{\text{sat}}^{40} \rceil = 142$ dual vectors in this dimension, which may be too small for the heuristic prediction to be accurate, as Heuristic 2.3 might not apply with few dual vectors, but also the lengths of the dual vectors are small. Therefore, we believe that the predictions become more accurate when there are more dual vectors, which happens, e.g., in higher dimensions, or when sieving with a larger saturation radius.

To test the robustness of our prediction for the floor phenomenon, we also ran experiments with a lattice sieve using a different saturation radius r_{sat} , which produces more (and therefore a bit longer) dual vectors than using the standard saturation radius of $\sqrt{4/3} \approx 1.1547$. In these scenarios, the floor phenomenon is observable already at a larger failure probability. Thus, fewer samples are needed to observe the floor phenomenon in the experimental data. This eases the computational power required to get the experiments, but the analysis should of course still hold in a situation where the saturation radius is larger than $\sqrt{4/3}$. This motivates Figure 4.4, which shows the predictions for a larger saturation radius. In these cases the floor phenomenon can be seen after a decent computation time, and it shows that the new predictions match the experimental data very accurately.

4.5 Conclusion

After identifying the norm of the target $\|\mathbf{t}\|$ as a confounding variable in the individual scores, we propose alternative heuristics in which this confounding variable is fixed. This leads us to predict the score distributions for both BDD targets and uniform targets in the Dual-Sieve Attack.

The newly derived heuristic predictions are tested in extensive experiments. Specifically, Heuristic Claim 4.15 and Heuristic Claim 4.18 experimentally show accurate predictions regarding both BDD and uniform




Figure 4.3: Binary logarithm ($\log_2$) of the survival function (SF) of the score distribution for uniform targets in various dimensions n . Each experimental curve contains 2^{48} samples, and ran a sieve with saturation ratio $f_{\text{sat}} = 0.9$ and saturation radius $r_{\text{sat}} = \sqrt{4/3}$. For $n = 40$, samples scoring below 60 are in buckets.




Figure 4.4: Binary logarithm ($\log_2$) of the survival function (SF) of the score distribution for uniform targets in various dimensions n , with a different *saturation radius*. Each experimental curve contains 2^{40} samples, and ran a sieve with saturation ratio $f_{\text{sat}} = 0.9$.

targets, compared to experimental data. In conclusion, the experiments strongly suggest that Heuristics 2.114, 4.11 and 4.17 may be used confidently in an analysis of the Dual-Sieve (with or without FFT) attack against BDD.

4.5.1 Future work

The effectiveness of the state-of-the-art dual attack remains as future work. In particular, note that existing dual attacks, e.g. [GJ21], will most likely require a different parametrisation to achieve the best runtime, while having a constant success probability, when the analysis will use the new model of Section 4.3. Moreover, there will be other aspects to the costing of the attack that may still require some attention, as highlighted in [DP23a, App. A].

On another note, theoretically it would be interesting to predict the scores for uniform targets from Section 4.3.4: instead of relying on a ball approximation of the Voronoi cell, it could be more satisfactory to adapt the Poisson model of [MT23, Assumption 8] to the lattice setting. One expects on average to have $\text{Vol}_n(r\mathcal{B}^n)/\det(\Lambda)$ many lattice points at distance at most r from a target $\mathbf{t}$ sampled uniformly modulo a lattice. Whereas the code setting has a Poisson process for each weight $w = 0, 1, \dots, n$, the situation is a bit more complex for lattices, as there is a continuum of processes to consider.