



LWE with Side Information: Attacks and Concrete Security Estimation

Dana Dachman-Soled^{1(✉)}, Léo Ducas^{2(✉)}, Huijing Gong^{1(✉)},
and Mélissa Rossi^{3,4,5,6(✉)}

¹ University of Maryland, College Park, USA
`danadach@ece.umd.edu`, `gong@cs.umd.edu`

² CWI, Amsterdam, The Netherlands
`l.ducas@cwi.nl`

³ ANSSI, Paris, France

⁴ ENS Paris, CNRS, PSL University, Paris, France
`melissa.rossi@ens.fr`

⁵ Thales, Gennevilliers, France

⁶ Inria, Paris, France

Abstract. We propose a framework for cryptanalysis of lattice-based schemes, when side information—in the form of “hints”—about the secret and/or error is available. Our framework generalizes the so-called primal lattice reduction attack, and allows the progressive integration of hints before running a final lattice reduction step. Our techniques for integrating hints include sparsifying the lattice, projecting onto and intersecting with hyperplanes, and/or altering the distribution of the secret vector. Our main contribution is to propose a toolbox and a methodology to integrate such hints into lattice reduction attacks and to predict the performance of those lattice attacks with side information.

While initially designed for side-channel information, our framework can also be used in other cases: exploiting decryption failures, or simply exploiting constraints imposed by certain schemes (LAC, Round5, NTRU).

We implement a Sage 9.0 toolkit to actually mount such attacks with hints when computationally feasible, and to predict their performances on larger instances. We provide several end-to-end application examples, such as an improvement of a single trace attack on Frodo by Bos et al. (SAC 2018). In particular, our work can estimate security loss even given very little side information, leading to a smooth measurement/-computation trade-off for side-channel attacks.

The research of L. Ducas and M. Rossi was supported by the European Union’s H2020 Programme under PROMETHEUS project (grant 780701). The research of M. Rossi was also supported by ANRT under the programs CIFRE N 2016/1583. It was also supported by the French Programme d’Investissement d’Avenir under national project RISQ P14158. The research of D. Dachman-Soled and H. Gong was supported in part by NSF grants #CNS-1933033, #CNS-1840893, #CNS-1453045 (CAREER), by a research partnership award from Cisco and by financial assistance award 70NANB15H328 and 70NANB19H126 from the U.S. Department of Commerce, National Institute of Standards and Technology.

$$\text{LWE/BDD} \xrightarrow{\text{Kannan}} \text{uSVP}_{A'} \xrightarrow{\text{Sec 3.4}} \text{Lattice reduction}$$

Fig. 1. Primal attack without hints (prior art).

Keywords: LWE · NTRU · Lattice reduction · Cryptanalysis · Side-channels analysis · Decryption failures

1 Introduction

A large effort is currently underway to replace standardized public key cryptosystems, which are quantum-insecure, with newly developed “post-quantum” cryptosystems, conjectured to be secure against quantum attack. Lattice-based cryptography has been widely recognized as a foremost candidate for practical, post-quantum security and accordingly, a large effort has been made to develop and analyze lattice-based cryptosystems. The ongoing standardization process and anticipated deployment of lattice-based cryptography raises an important question: How resilient are lattices to side-channel attacks or other forms of side information? While there are numerous works addressing this question for specific cryptosystems (See [2, 9, 17, 18, 32, 33] for side channel attacks targeting lattice-based NIST candidates), these works use rather ad-hoc methods to reconstruct the secret key, requiring new techniques and algorithms to be developed for each setting. For example, the work of [9] uses brute-force methods for a portion of the attack, while [7] exploits linear regression techniques. Moreover, ad-hoc methods do not allow (1) to take advantage of decades worth of research and (2) optimization of standard lattice attacks. Second, most of the side-channel attacks from prior work consider substantial amounts of information leakage and show that it leads to feasible recovery of the entire key, whereas one may be interested in more precise tradeoffs in terms of information leakage versus concrete security of the scheme. The above motivates the focus of this work: Can one integrate side information into a standard lattice attack, and if so, by how much does the information reduce the cost of this attack? Given that side-channel resistance is the next step toward the technological readiness of lattice-based cryptography, and that we expect numerous works in this growing area, we believe that a general framework and a prediction software are in order.

Contributions. First, we propose a framework that generalizes the so-called primal lattice reduction attack, and allows the progressive integration of “hints” (i.e. side information that takes one of several forms) before running the final lattice reduction step. This contribution is summarized in Figs. 1 and 2 and developed in Sect. 3.

Second, we implement a Sage 9.0 toolkit to actually mount such attacks with hints when computationally feasible, and to predict their performance on larger instances. Our predictions are validated by extensive experiments. Our tool and

these experiments are described in Sect. 5. Our toolkit is open-source, available at: <https://github.com/lducas/leaky-LWE-Estimator>.

Third, we demonstrate the usefulness of our framework and tool via three example applications. Our main example (Sect. 6.1) revisits the side channel information obtained from the first side-channel attack of [9] against Frodo. In that article, it was concluded that a divide-and-conquer side-channel template attack would not lead to a meaningful attack using standard combinatorial search for reconstruction of the secret. Our technique allows to integrate this side-channel information into lattice attacks, and to predict the exact security drop. For example, the CCS2 parameter set very conservatively aims for 128-bits of post-quantum security (or 448 “bikz” as defined in Sect. 3.4); but after the leakage of [9] we predict that its security drops to 29 “bikz”, i.e. that it can be broken with BKZ-29, a computation that should be more than feasible, but would require a dedicated re-implementation of our framework.

Interestingly, we note that our framework is not only useful in the side-channel scenario; we are for example also able to model decryption failures as hints fitting our framework. This allows us to reproduce some predictions from [14]. This is discussed in Sect. 6.2.

Perhaps more surprisingly, we also find a novel improvement to attack a few schemes (LAC [25], Round5 [16], NTRU [35]) without any side-channel or oracle queries. Indeed, such schemes use ternary distribution for secrets, with a prescribed numbers of 1 and -1 : this hint fits our framework, and lead to a (very) minor improvement, discussed in Sect. 6.3.

Lastly, our framework also encompasses and streamlines existing tweaks of the primal attack: the choice of ignoring certain LWE equations to optimize the volume-dimension trade-off, as well as the re-centering [30] and isotropization [12, 19] accounting for potential a-priori distortions of the secret. It also implicitly solves the question of the optimal choice of the coefficient for Kannan’s Embedding from the Bounded Distance Decoding problem (BDD) to the unique Shortest Vector Problem (uSVP) [21] (See Remark 22).

As a side contribution, we also propose in the full version of our paper [13] a refined method to estimate the required blocksize to solve an LWE/BDD/uSVP instance. This refinement was motivated by the inaccuracy of the standard method from the literature [3, 4] in experimentally reachable blocksizes, which was making the validation of our contribution difficult. While experimentally much more accurate, this new methodology certainly deserves further scrutiny.

Technical overview. Our work is based on a generalization of the Bounded Distance Decoding problem (BDD) to a Distorted version (DBDD), which allows to account for the potentially non-spherical covariance of the secret vector to be found.

Each hint will affect the lattice itself, the mean and/or the covariance parameter of the DBDD instance, making the problem easier (see Fig. 2). At last, we make the distribution spherical again by applying a well-chosen linear transformation, reverting to a spherical BDD instance before running the attack. Thanks to the hints, this new instance will be easier than the initial one. Let us assume




Fig. 2. The primal attack with hints (our work).

that \mathbf{v} , l , k and σ are parameters known by the attacker. Our framework can handle four types of hints on the secret \mathbf{s} or on the lattice Λ .

- Perfect hints: $\langle \mathbf{s}, \mathbf{v} \rangle = l$ *intersect the lattice with an hyperplane.*
- Modular hints: $\langle \mathbf{s}, \mathbf{v} \rangle = l \pmod k$ *sparsify the lattice.*
- Approximate hints: $\langle \mathbf{s}, \mathbf{v} \rangle = l + \epsilon_\sigma$ *decrease the covariance of the secret.*
- Short vector hints: $\mathbf{v} \in \Lambda$ *project orthogonally to \mathbf{v} .*

While the first three hints are clear wins for the performance of lattice attacks, the last one is a trade-off between the dimension and the volume of the lattice. This last type of hint is in fact meant to generalize the standard trick consisting of ‘ignoring’ certain LWE equations; ignoring such an equation can be interpreted geometrically as such a projection orthogonally to a so-called q -vector.

All the transformations of the lattice above can be computed in polynomial time. However, computing with general distribution in large dimension is not possible; we restrict our study to the case of Gaussian distributions of arbitrary covariance, for which such computations are also poly-time.

Some of these transformations remain quite expensive, in particular because they involve rational numbers with very large denominators, and it remains rather impractical to run them on cryptographic-grade instances. Fortunately, up to a necessary hypothesis of primitivity of the vector \mathbf{v} (with respect to either Λ or its dual depending on the type of hint), we can also predict the effect of each hint on the lattice parameters, and therefore run faster predictions of the attack cost.

From Leaks to Hints. At first, it may not be so clear that the types of hints above are so useful in realistic applications, in particular since they need to be linear on the secret. Of course our framework can handle rather trivial hints such as the perfect leak of a secret coefficient $\mathbf{s}_i = l$. Slightly less trivial is the case where the only the low-order bits leaks, a hint of the form $\mathbf{s}_i = l \pmod 2$.

We note that most of the computations done during an LWE decryption are linear: leaking any intermediate register during a matrix vector product leads to a hint of the same form (possibly $\pmod q$). Similarly, the leak of a NTT coefficient of a secret in a Ring/Module variant can also be viewed as such.

Admittedly, such ideal leaks of a full register are not the typical scenario and leaks are typically not linear on the content of the register. However, such non-linearities can be handled by approximate hints. For instance, let \mathbf{s}_0 be a secret coefficient (represented by a signed 16-bits integer), whose a priori distribution is supported by $\{-5, \dots, 5\}$. Consider the case where we learn the Hamming weight of \mathbf{s}_0 , say $H(\mathbf{s}_0) = 2$. Then, we can narrow down the possibilities to $\mathbf{s}_0 \in \{3, 5\}$. This leads to two hints:

- a modular hint: $\mathbf{s}_0 = 1 \pmod 2$,
- an approximate hint: $\mathbf{s}_0 = 4 + \epsilon_1$, where ϵ_1 has variance 1.

While closer to a realistic scenario, the above example remains rather simplified. A detailed example of how realistic leaks can be integrated as hint will be given in Sect. 6.1, based on the leakage data from [9].

2 Preliminaries

2.1 Linear Algebra

We use bold lower case letters to denote vectors, and bold upper case letters to denote matrices. We use row notations for vectors, and start indexing from 0. Let \mathbf{I}_d denote the d -dimensional identity matrix. Let $\langle \cdot, \cdot \rangle$ denote the inner product of two vectors of the same size. Let us introduce the row span of a matrix (denoted $\text{Span}(\cdot)$) as the subspace generated by all \mathbb{R} -linear combinations of the rows of its input.

Definition 1 (Positive Semidefinite). *A $n \times n$ symmetric real matrix \mathbf{M} is positive semidefinite if scalar $\mathbf{xMx}^T \geq 0$ for all $\mathbf{x} \in \mathbb{R}^n$; if so we write $\mathbf{M} \geq 0$. Given two $n \times n$ real matrix \mathbf{A} and \mathbf{B} , we note $\mathbf{A} \geq \mathbf{B}$ if $\mathbf{A} - \mathbf{B}$ is positive semidefinite.*

Definition 2. *A matrix \mathbf{M} is a square root of Σ , denoted $\sqrt{\Sigma}$, if*

$$\mathbf{M}^T \cdot \mathbf{M} = \Sigma,$$

Our techniques involve keeping track of the covariance matrix Σ of the secret and error vectors as hints are progressively integrated. The covariance matrix may become singular during this process and will not have an inverse. Therefore, in the following we introduce some degenerate notions for the inverse and the determinant of a square matrix. Essentially, we restrict these notions to the row span of their input. For $\mathbf{X} \in \mathbb{R}^{d \times k}$ (with any $d, k \in \mathbb{N}$), we will denote $\Pi_{\mathbf{X}}$ the orthogonal projection matrix onto $\text{Span}(\mathbf{X})$. More formally, let \mathbf{Y} be a maximal set of independent row-vectors of \mathbf{X} ; the orthogonal projection matrix is given by $\Pi_{\mathbf{X}} = \mathbf{Y}^T \cdot (\mathbf{Y} \cdot \mathbf{Y}^T)^{-1} \cdot \mathbf{Y}$. Its complement (the projection orthogonally to $\text{Span}(\mathbf{X})$) is denoted $\Pi_{\mathbf{X}}^\perp := \mathbf{I}_d - \Pi_{\mathbf{X}}$. We naturally extend the notation Π_F and Π_F^\perp to subspaces $F \subset \mathbb{R}^d$. By definition, the projection matrices satisfy $\Pi_F^2 = \Pi_F$, $\Pi_F^T = \Pi_F$ and $\Pi_F \cdot \Pi_F^\perp = \Pi_F^\perp \cdot \Pi_F = \mathbf{0}$.

Definition 3 (Restricted inverse and determinant). Let Σ be a symmetric matrix. We define a restricted inverse denoted Σ^\sim as

$$\Sigma^\sim := (\Sigma + \Pi_\Sigma^\perp)^{-1} - \Pi_\Sigma^\perp.$$

It satisfies $\text{Span}(\Sigma^\sim) = \text{Span}(\Sigma)$ and $\Sigma \cdot \Sigma^\sim = \Pi_\Sigma$.

We also denote $\text{rdet}(\Sigma)$ as the restricted determinant defined as follows.

$$\text{rdet}(\Sigma) := \det(\Sigma + \Pi_\Sigma^\perp).$$

The idea behind Definition 3 is to provide an (artificial) invertibility property to the input Σ by adding the missing orthogonal part and to remove it afterwards. For example, if $\Sigma = \begin{bmatrix} \mathbf{A} & 0 \\ 0 & 0 \end{bmatrix}$ where \mathbf{A} is invertible,

$$\Sigma^\sim = \left(\begin{bmatrix} \mathbf{A} & 0 \\ 0 & 0 \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} \right)^{-1} - \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} \mathbf{A}^{-1} & 0 \\ 0 & 0 \end{bmatrix} \text{ and } \text{rdet } \Sigma = \det(\mathbf{A}).$$

2.2 Statistics

Random variables, i.e. variables whose values depend on outcomes of a random phenomenon, are denoted in lowercase calligraphic letters e.g. a, b, e . Random vectors are denoted in uppercase calligraphic letters e.g. C, X, Z .

Before hints are integrated, we will assume that the secret and error vectors follow a multidimensional normal (Gaussian) distribution. Hints will typically correspond to learning a (noisy, modular or perfect) linear equation on the secret. We must then consider the altered distribution on the secret, conditioned on this information. Fortunately, this will also be a multidimensional normal distribution with an altered covariance and mean. In the following, we present the precise formulae for the covariance and mean of these conditional distributions.

Definition 4 (Multidimensional normal distribution). Let $d \in \mathbb{Z}$, for $\mu \in \mathbb{Z}^d$ and Σ being a symmetric matrix of dimension $d \times d$, we denote by $D_{\Sigma, \mu}^d$ the multidimensional normal distribution supported by $\mu + \text{Span}(\Sigma)$ by the following

$$\mathbf{x} \mapsto \frac{1}{\sqrt{(2\pi)^{\text{rank}(\Sigma)} \cdot \text{rdet}(\Sigma)}} \exp\left(-\frac{1}{2}(\mathbf{x} - \mu) \cdot \Sigma^\sim \cdot (\mathbf{x} - \mu)^T\right).$$

The following states how a normal distribution is altered under linear transformation.

Lemma 5. Suppose X has a $D_{\Sigma, \mu}^d$ distribution. Let \mathbf{A} be a $n \times d$ matrix. Then $X\mathbf{A}^T$ has a $D_{\mathbf{A}\Sigma\mathbf{A}^T, \mu\mathbf{A}^T}^n$ distribution.

Lemma 6 shows the altered distribution of a normal random variable conditioned on its noisy linear transformation value, following from [24, Equations (6) and (7)].

Lemma 6 (Conditional distribution $X|XA^T + b$ from [24]). Suppose that $X \in \mathbb{Z}^d$ has a $D_{\Sigma, \mu}^d$ distribution, and $b \in \mathbb{Z}^n$ has a $D_{\Sigma_b, 0}^n$ distribution. Let us fix \mathbf{A} as a $n \times d$ matrix and $\mathbf{z} \in \mathbb{Z}^n$. The conditional distribution of $X \mid (XA^T + b = \mathbf{z})$ is $D_{\Sigma', \mu'}^d$, where

$$\begin{aligned} \mu' &= \mu + (\mathbf{z} - \mu\mathbf{A}^T)(\mathbf{A}\Sigma\mathbf{A}^T + \Sigma_b)^{-1}\mathbf{A}\Sigma \\ \Sigma' &= \Sigma - \Sigma\mathbf{A}^T(\mathbf{A}\Sigma\mathbf{A}^T + \Sigma_b)^{-1}\mathbf{A}\Sigma. \end{aligned}$$

Corollary 7 (Conditional distribution $X|\langle X, \mathbf{v} \rangle + e$). Suppose that $X \in \mathbb{Z}^d$ has a $D_{\Sigma, \mu}^d$ distribution and e has a $D_{\sigma_e^2, 0}^1$ distribution. Let us fix $\mathbf{v} \in \mathbb{R}^d$ as a nonzero vector and $z \in \mathbb{Z}$. We define the following scalars:

$$y = \langle X, \mathbf{v} \rangle + e, \quad \mu_2 = \langle \mathbf{v}, \mu \rangle \quad \text{and} \quad \sigma_2 = \mathbf{v}\Sigma\mathbf{v}^T + \sigma_e^2$$

If $\sigma_2 \neq 0$, the conditional distribution of $X \mid (y = z)$ is $D_{\Sigma', \mu'}^d$, where

$$\mu' = \mu + \frac{(z - \mu_2)}{\sigma_2}\mathbf{v}\Sigma, \quad \Sigma' = \Sigma - \frac{\Sigma\mathbf{v}^T\mathbf{v}\Sigma}{\sigma_2}. \tag{1}$$

If $\sigma_2 = 0$, the conditional distribution of $X \mid (y = z)$ is $D_{\Sigma, \mu}^d$.

Remark 8. We note that Corollary 7 is also useful to describe for $X|\langle X, \mathbf{v} \rangle$ by letting $\sigma_e = 0$.

2.3 Lattices

A *lattice*, denoted as Λ , is a discrete additive subgroup of \mathbb{R}^m , which is generated as the set of all linear integer combinations of n ($m \geq n$) linearly independent basis vectors $\{\mathbf{b}_j\} \subset \mathbb{R}^m$, namely,

$$\Lambda := \left\{ \sum_j z_j \mathbf{b}_j : z_j \in \mathbb{Z} \right\},$$

We say that m is the *dimension* of Λ and n is its rank. A lattice is *full rank* if $n = m$. A matrix \mathbf{B} having the basis vectors as rows is called a *basis*. The *volume* of a lattice Λ is defined as $\text{Vol}(\Lambda) := \sqrt{\det(\mathbf{B}\mathbf{B}^T)}$. The *dual lattice* of Λ in \mathbb{R}^n is defined as follows.

$$\Lambda^* := \{ \mathbf{y} \in \text{Span}(\mathbf{B}) \mid \forall \mathbf{x} \in \Lambda, \langle \mathbf{x}, \mathbf{y} \rangle \in \mathbb{Z} \}.$$

Note that, $(\Lambda^*)^* = \Lambda$, and $\text{Vol}(\Lambda^*) = 1/\text{Vol}(\Lambda)$.

Lemma 9 ([26, Proposition 1.3.4]). *Let Λ be a lattice and let F be a subspace of \mathbb{R}^n . If $\Lambda \cap F$ is a lattice, then the dual of $\Lambda \cap F$ is the orthogonal projection onto F of the dual of Λ . In other words, each element of Λ^* is multiplied by the projection matrix Π_F :*

$$(\Lambda \cap F)^* = \Lambda^* \cdot \Pi_F.$$

Definition 10 (Primitive vectors). *A set of vector $\mathbf{y}_1, \dots, \mathbf{y}_k \in \Lambda$ is said primitive with respect to Λ if $\Lambda \cap \text{Span}(\mathbf{y}_1, \dots, \mathbf{y}_k)$ is equal to the lattice generated by $\mathbf{y}_1, \dots, \mathbf{y}_k$. Equivalently, it is primitive if it can be extended to a basis of Λ . If $k = 1$, \mathbf{y}_1 , this is equivalent to $\mathbf{y}_1/i \notin \Lambda$ for any integer $i \geq 2$.*

To predict the hardness of the lattice reduction on altered instances, we must compute the volume of the final transformed lattice. We devise a highly efficient way to do this, by observing that each time a hint is integrated, we can update the volume of the transformed lattice, given only the volume of the previous lattice and information about the current hint (under mild restrictions on the form of the hint). Lemmas 11 and 12 are proved in the full version of our paper [13].

Lemma 11 (Volume of a lattice slice). *Given a lattice Λ with volume $\text{Vol}(\Lambda)$, and a primitive vector \mathbf{v} with respect to Λ^* . Let \mathbf{v}^\perp denote subspace orthogonal to \mathbf{v} . Then $\Lambda \cap \mathbf{v}^\perp$ is a lattice with volume $\text{Vol}(\Lambda \cap \mathbf{v}^\perp) = \|\mathbf{v}\| \cdot \text{Vol}(\Lambda)$.*

Lemma 12 (Volume of a sparsified lattice). *Let Λ be a lattice, $\mathbf{v} \in \Lambda^*$ be a primitive vector of Λ^* , and $k > 0$ be an integer. Let $\Lambda' = \{\mathbf{x} \in \Lambda \mid \langle \mathbf{x}, \mathbf{v} \rangle = 0 \pmod k\}$ be a sublattice of Λ . Then $\text{Vol}(\Lambda') = k \cdot \text{Vol}(\Lambda)$.*

Fact 13 (Volume of a projected lattice). *Let Λ be a lattice, $\mathbf{v} \in \Lambda$ be a primitive vector of Λ . Let $\Lambda' = \Lambda \cdot \Pi_{\mathbf{v}}^\perp$ be a sublattice of Λ . Then $\text{Vol}(\Lambda') = \text{Vol}(\Lambda) / \|\mathbf{v}\|$. More generally, if \mathbf{V} is a primitive set of vectors of Λ , then $\Lambda' = \Lambda \cdot \Pi_{\mathbf{V}}^\perp$ has volume $\text{Vol}(\Lambda') = \text{Vol}(\Lambda) / \sqrt{\det(\mathbf{V}\mathbf{V}^T)}$.*

Fact 14 (Lattice volume under linear transformations). *Let Λ be a lattice in \mathbb{R}^n , and $\mathbf{M} \in \mathbb{R}^{n \times n}$ a matrix such that $\ker \mathbf{M} = \text{Span}(\Lambda)^\perp$. Then we have $\text{Vol}(\Lambda \cdot \mathbf{M}) = \text{rdet}(\mathbf{M}) \text{Vol}(\Lambda)$.*

3 Distorted Bounded Distance Decoding

3.1 Definition

We first recall the definition of the (search) LWE problem, in its short-secret variant which is the most relevant to practical LWE-based encryption.

Definition 15 (Search LWE problem with short secrets). *Let n, m and q be positive integers, and let χ be a distribution over \mathbb{Z} . The search LWE problem (with short secrets) for parameters (n, m, q, χ) is:*

Given the pair $(\mathbf{A} \in \mathbb{Z}_q^{m \times n}, \mathbf{b} = \mathbf{z}\mathbf{A}^T + \mathbf{e} \in \mathbb{Z}_q^m)$ where:

1. $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ is sampled uniformly at random,
2. $\mathbf{z} \leftarrow \chi^n$, and $\mathbf{e} \leftarrow \chi^m$ are sampled with independent and identically distributed coefficients following the distribution χ .

Find \mathbf{z} .

The primal attack (See for example [3]) against (search)-LWE proceeds by viewing the LWE instance as an instance of a Bounded Distance Decoding (BDD) problem, converting it to a uSVP instance (via Kannan’s embedding [21]), and finally applying a lattice reduction algorithm to solve the uSVP instance. The central tool of our framework is a generalization of BDD that accounts for potential distortion in the distribution of the secret noise vector that is to be recovered.

Definition 16 (Distorted Bounded Distance Decoding problem). *Let $\Lambda \subset \mathbb{R}^d$ be a lattice, $\Sigma \in \mathbb{R}^{d \times d}$ be a symmetric matrix and $\mu \in \text{Span}(\Lambda) \subset \mathbb{R}^d$ such that*

$$\text{Span}(\Sigma) \subsetneq \text{Span}(\Sigma + \mu^T \cdot \mu) = \text{Span}(\Lambda). \tag{2}$$

The Distorted Bounded Distance Decoding problem $DBDD_{\Lambda, \mu, \Sigma}$ is the following problem:

Given μ, Σ and a basis of Λ .

Find the unique vector $\mathbf{x} \in \Lambda \cap E(\mu, \Sigma)$

where $E(\mu, \Sigma)$ denotes the ellipsoid

$$E(\mu, \Sigma) := \{\mathbf{x} \in \mu + \text{Span}(\Sigma) \mid (\mathbf{x} - \mu) \cdot \Sigma^{-1} \cdot (\mathbf{x} - \mu)^T \leq \text{rank}(\Sigma)\}.$$

We will refer to the triple $\mathcal{I} = (\Lambda, \mu, \Sigma)$ as the instance of the $DBDD_{\Lambda, \mu, \Sigma}$ problem.

Intuitively, Definition 16 corresponds to knowing that the secret vector \mathbf{x} to be recovered follows a distribution of variance Σ and average μ . The quantity $(\mathbf{x} - \mu) \cdot \Sigma^{-1} \cdot (\mathbf{x} - \mu)^T$ can be interpreted as a non-canonical Euclidean squared distance $\|\mathbf{x} - \mu\|_{\Sigma}^2$, and the expected value of such a distance for a Gaussian \mathbf{x} of variance Σ and average μ is $\text{rank}(\Sigma)$. One can argue that, for such a Gaussian, there is a constant probability that $\|\mathbf{x} - \mu\|_{\Sigma}^2$ is slightly greater than $\text{rank}(\Sigma)$. Since we are interested in the average behavior of our attack, we ignore this benign technical detail. In fact, we will typically interpret DBDD as the promise that the secret follows a Gaussian distribution of center μ and covariance Σ .

The ellipsoid can be seen as an affine transformation (that we call “distortion”) of the centered hyperball of radius $\text{rank}(\Sigma)$. Let us introduce a notation for the hyperball; for any $d \in \mathbb{N}$

$$B_d := \{\mathbf{x} \in \mathbb{R}^d \mid \|\mathbf{x}\|_2 \leq d\}. \tag{3}$$

One can thus write using Definition 2:

$$E(\mu, \Sigma) = B_{\text{rank}(\Sigma)} \cdot \sqrt{\Sigma} + \mu. \tag{4}$$




Fig. 3. Graphical intuition of DBDD, BDD and uSVP in dimension two: the problem consists in finding a nonzero element of Λ in the colored zone. The identity hyperball is larger for uSVP to represent the fact that, during the reduction, the uSVP lattice has one dimension more than for BDD.

From the Span inclusion in Eq. (2), one can deduce that the condition is equivalent to requiring $\mu \notin \text{Span}(\Sigma)$ and $\text{rank}(\Sigma + \mu^T \cdot \mu) = \text{rank}(\Sigma) + 1 = \text{rank}(\Lambda)$. This technical detail is necessary for embedding it properly into a uSVP instance (See later in Sect. 3.3).

Particular cases of Definition 16. Let us temporarily ignore the condition in Eq. (2) to study some particular cases. As shown in Fig. 3, when $\Sigma = \mathbf{I}_d$, $\text{DBDD}_{\Lambda, \mu, \mathbf{I}_d}$ is BDD instance. Indeed, the ellipsoid becomes a shifted hyperball $E(\mu, \mathbf{I}_d) = \{x \in \mu + \mathbb{R}^{d \times d} \mid \|x - \mu\|_2 \leq d\} = B_d + \mu$. If in addition $\mu = 0$, $\text{DBDD}_{\Lambda, 0, \mathbf{I}_d}$ becomes a uSVP instance on Λ .

3.2 Embedding LWE into DBDD

In the typical primal attack framework (Fig. 1), one directly views LWE as a BDD instance of the same dimension. For our purposes, however, it will be useful to apply Kannan’s Embedding at this stage and therefore increase the dimension of the lattice by 1. While it could be delayed to the last stage of our attack, this extra fixed coefficient 1 will be particularly convenient when we integrate hints (see Remark 22 in Sect. 4). It should be noted that no information is lost through this transformation, since the parameters μ and Σ allow us to encode the knowledge that the solution we are looking for has its last coefficient set to 1 and nothing else. In more details, the solution $s := (e, z)$ of an LWE instance is extended to

$$\bar{s} := (e, z, 1) \tag{5}$$

which is a short vector in the lattice $\Lambda = \{(x, y, w) \mid x + yA^T - bw = 0 \pmod q\}$. A basis of this lattice is given by the row vectors of

$$\begin{bmatrix} q\mathbf{I}_m & 0 & 0 \\ \mathbf{A}^T & -\mathbf{I}_n & 0 \\ \mathbf{b} & 0 & 1 \end{bmatrix}.$$

Denoting μ_χ and σ_χ^2 the average and variance of the LWE distribution χ (See Definition 15), we can convert this LWE instance to a $\text{DBDD}_{\Lambda, \mu, \Sigma}$ instance with

$\boldsymbol{\mu} = [\mu_\chi \cdots \mu_\chi \ 1]$ and $\boldsymbol{\Sigma} = \begin{bmatrix} \sigma_\chi^2 \mathbf{I}_{m+n} & 0 \\ 0 & 0 \end{bmatrix}$. The lattice Λ is of full rank in \mathbb{R}^d where $d := m + n + 1$, and its volume is q^m . Note that the rank of $\boldsymbol{\Sigma}$ is only $d - 1$: the ellipsoid has one less dimension than the lattice. It then validates the requirement of Eq. (2).

Remark 17. Typically, Kannan’s embedding from BDD to uSVP leaves the bottom right matrix coefficient as a free parameter, say c , to be chosen optimally. The optimal value is the one maximizing

$$\frac{\|(\mathbf{z}; c)\|}{\det(\Lambda)^{1/d}} = \frac{(m + n)\sigma_\chi + c}{(c \cdot q^m)^{1/d}},$$

namely, $c = \sigma_\chi$ according to the arithmetic-geometric mean inequality. Some prior works [3, 5] instead chose $c = 1$. While this is benign since σ_χ is typically not too far from 1, it remains a sub-optimal choice. Looking ahead, in our DBDD framework, this choice becomes irrelevant thanks to the *isotropization* step introduced in the next section; we can therefore choose $c = 1$ without worsening the attack.

3.3 Converting DBDD to uSVP

In this Section, we explain how a DBDD instance $(\Lambda, \boldsymbol{\mu}, \boldsymbol{\Sigma})$ is converted into a uSVP one. Two modifications are necessary. First, we need to homogenize the problem. Let us show that the ellipsoid in Definition 16 is contained in a larger centered ellipsoid (with one more dimension) as follows:

$$E(\boldsymbol{\mu}, \boldsymbol{\Sigma}) \subset E(\mathbf{0}, \boldsymbol{\Sigma} + \boldsymbol{\mu}^T \cdot \boldsymbol{\mu}). \tag{6}$$

Using Eq. (4), one can write

$$E(\boldsymbol{\mu}, \boldsymbol{\Sigma}) = B_{\text{rank}(\boldsymbol{\Sigma})} \cdot \sqrt{\boldsymbol{\Sigma}} + \boldsymbol{\mu} \subset B_{\text{rank}(\boldsymbol{\Sigma})} \cdot \sqrt{\boldsymbol{\Sigma}} \pm \boldsymbol{\mu},$$

where $B_{\text{rank}(\boldsymbol{\Sigma})}$ is defined in Eq. (3). And, with Eq. (2), one can deduce $\text{rank}(\boldsymbol{\Sigma} + \boldsymbol{\mu}^T \cdot \boldsymbol{\mu}) = \text{rank}(\boldsymbol{\Sigma}) + 1$, then:

$$B_{\text{rank}(\boldsymbol{\Sigma})} \cdot \sqrt{\boldsymbol{\Sigma}} \pm \boldsymbol{\mu} \subset B_{\text{rank}(\boldsymbol{\Sigma})+1} \cdot \begin{bmatrix} \sqrt{\boldsymbol{\Sigma}} \\ \boldsymbol{\mu} \end{bmatrix}.$$

We apply Definition 2 which confirms the inclusion of Eq. (6):

$$E(\boldsymbol{\mu}, \boldsymbol{\Sigma}) \subset B_{\text{rank}(\boldsymbol{\Sigma})+1} \cdot \begin{bmatrix} \sqrt{\boldsymbol{\Sigma}} \\ \boldsymbol{\mu} \end{bmatrix} = E(\mathbf{0}, \boldsymbol{\Sigma} + \boldsymbol{\mu}^T \cdot \boldsymbol{\mu}).$$

Thus, we can homogenize and transform the instance into a centered one with $\boldsymbol{\Sigma}' := \boldsymbol{\Sigma} + \boldsymbol{\mu}^T \cdot \boldsymbol{\mu}$.

Secondly, to get an isotropic distribution (i.e. with all its eigenvalues being 1), one can just multiply every element of the lattice with the pseudoinverse

of $\sqrt{\Sigma'}$. We get a new covariance matrix $\Sigma'' = \sqrt{\Sigma'} \sim \Sigma' \cdot \sqrt{\Sigma'} \sim T = \Pi_{\Sigma'} \cdot \Pi_{\Sigma'}^T$. And with orthogonal projection properties (see Sect. 2.1), $\Sigma'' = \Pi_{\Sigma'} = \Pi_A$, the last equality coming from Eq. (2).

In summary, one must make by the two following changes:

$$\begin{aligned} \text{homogenize: } (A, \mu, \Sigma) &\mapsto (A, \mathbf{0}, \Sigma' := \Sigma + \mu^T \cdot \mu) \\ \text{isotropize: } (A, \mathbf{0}, \Sigma') &\mapsto (A \cdot \mathbf{M}, \mathbf{0}, \Pi_A) \end{aligned}$$

where $\mathbf{M} := (\sqrt{\Sigma'}) \sim$. From the solution \mathbf{x} to the $\text{uSVP}_{A \cdot \mathbf{M}}$ problem, one can derive $\mathbf{x}' = \mathbf{x} \mathbf{M} \sim$ the solution to the $\text{DBDD}_{A, \mu, \Sigma}$ problem.

Remark 18. One may note that we could solve a DBDD instance without isotropization simply by including the ellipsoid in a larger ball, and directly apply lattice reduction before the second step. This leads, however, to less efficient attacks. One may also note that the first homogenization step “forgets” some information about the secret’s distribution. This, however, is inherent to the conversion to a unique-SVP problem which is geometrically homogeneous, and is already present in the original primal attack.

3.4 Security Estimates of uSVP: Bikz versus Bits

The attack on a uSVP instance consists of applying BKZ- β on the uSVP lattice A for an appropriate block size parameter β . The cost of the attack grows with β , however, modeling this cost precisely is at the moment rather delicate, as the state of the art seems to still be in motion. Numerous NIST candidates choose to underestimate this cost, keeping a margin to accommodate future improvements, and there seems to be no clear consensus on which model to use (see [1] for a summary of existing cost models).

While this problem is orthogonal to our work, we still wish to be able to formulate quantitative security losses. We therefore express all concrete security estimates using the blocksize β as our measure of the level of security, and treat the latter as a measurement of the security level in a unit called the *bikz*. We thereby leave the question of the exact bikz-to-bit conversion estimate outside the scope of this paper, and recall that those conversion formulae are not necessarily linear, and may have small dependency in other parameters. For the sake of concreteness, we note that certain choose, for example, to claim 128 bits of security for 380 bikz, and in this range, most models suggest a security increase of one bit every 2 to 4 bikz.

Remark 19. We also clarify that the estimates given in this paper only concern the pure lattice attack via the uSVP embedding discussed above. In particular, we note that some NIST candidates with ternary secrets [25] also consider the hybrid attack of [20], which we ignore in this work. We nevertheless think that the compatibility with our framework is plausible, with some effort.

Predicting β from a uSVP instance. The state-of-the-art predictions for solving uSVP instances using BKZ were given in [3, 4]. Namely, for Λ a lattice of dimension $\dim(\Lambda)$, it is predicted that BKZ- β can solve a uSVP $_{\Lambda}$ instance with secret \mathbf{s} when

$$\sqrt{\beta / \dim(\Lambda)} \cdot \|\mathbf{s}\| \leq \delta_{\beta}^{2\beta - \dim(\Lambda) - 1} \cdot \text{Vol}(\Lambda)^{1 / \dim(\Lambda)} \tag{7}$$

where δ_{β} is the so called root-Hermite-Factor of BKZ- β . For $\beta \geq 50$, the Root-Hermite-Factor is predictable using the Gaussian Heuristic [11]:

$$\delta_{\beta} = \left((\pi\beta)^{\frac{1}{\beta}} \cdot \frac{\beta}{2\pi e} \right)^{1 / (2\beta - 2)} \tag{8}$$

Note that the uSVP instances we generate are isotropic and centered so that the secret has covariance $\Sigma = \mathbf{I}$ (or $\Sigma = \Pi_{\Lambda}$ if Λ is not of full rank) and $\mu = \mathbf{0}$. Thus, on average, we have $\|\mathbf{s}\|^2 = \text{rank}(\Sigma) = \dim(\Lambda)$. Therefore, β can be estimated as the minimum integer that satisfies

$$\sqrt{\beta} \leq \delta_{\beta}^{2\beta - \dim(\Lambda) - 1} \cdot \text{Vol}(\Lambda)^{1 / \dim(\Lambda)} \tag{9}$$

While β must be an integer as a BKZ parameter, we nevertheless provide a continuous value, for a finer comparison of the difficulty of an instance. Below, we will call this method the ‘‘GSA-Intersect’’ method.

Remark 20. To predict security, one does not need the basis of Λ , but only its dimension and its volume. Similarly, it is not necessary to explicitly compute the isotropization matrix \mathbf{M} of Sect. 3.3, thanks to Fact 14: $\text{Vol}(\Lambda \cdot \mathbf{M}) = \text{rdet}(\mathbf{M}) \text{Vol}(\Lambda) = \text{rdet}(\Sigma')^{-1/2} \text{Vol}(\Lambda)$. These two shortcuts will allow us to efficiently make predictions for cryptographically large instances, in our *lightweight* implementation of Sect. 5.

Refined prediction for small block sizes. For experimental validation purposes of our work, we prefer to have accurate prediction even for small block sizes; a regime where those predictions are not accurate with the current state of the art. We therefore present a refined strategy using BKZ-simulation and a probabilistic model in the full version of our paper [13].

As depicted in Fig. 4, this methodology (coined Probabilistic-simulation) leads to much more satisfactory estimates compared to the model from the literature [3, 4]. In particular, for low block size the literature widely underestimates the required block size, which is due to only considering detectability at position $d - \beta$. For large block size, it somewhat overestimates it, which could be attributed to the fact that it does not account for luck. On the contrary, our new methodology seems quite precise in all regimes, making errors of at most 1 bikz. This new methodology certainly deserves further study and refinement, which we leave to future work.

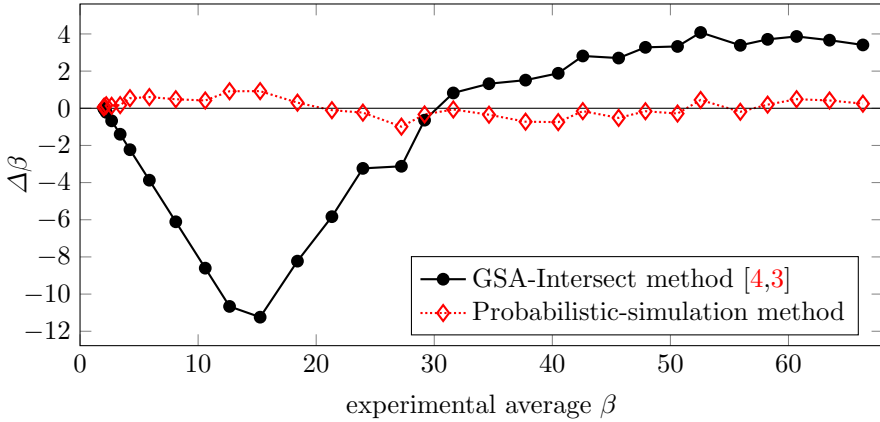


Fig. 4. The difference $\Delta\beta = \text{real} - \text{predicted}$, as a function of the average experimental beta β . The experiment consists in running a single tour of BKZ- β for $\beta = 2, 3, 4, \dots$ until the secret short vector is found. This was averaged over 256 many LWE instances per data-point, for parameters $q = 3301$, $\sigma = 20$ and $n = m \in \{30, 32, 34, \dots, 88\}$.

4 Hints and Their Integration

In this Section, we define several categories of hints—**perfect hints**, **modular hints**, **approximate hints** (**conditioning** and *a posteriori*), and **short vector hints**—and show that these types of hints can be integrated into a DBDD instance. Hints belonging to these categories typically have the form of a linear equation in \mathbf{s} (and possibly additional variables). As emphasized in Sect. 1, these hints have lattice-friendly forms and their usefulness in realistic applications may not be obvious. We refer to Sect. 6 for detailed applications of these hints.

The technical challenge, therefore, is to characterize the effect of such hints on the DBDD instance—i.e. determine the resulting $(\Lambda', \mu', \Sigma')$ of the new DBDD instance, after the hint is incorporated.

Henceforth, let $\mathcal{I} = \text{DBDD}_{\Lambda, \mu, \Sigma}$ be a fixed instance constructed from an LWE instance with secret $\mathbf{s} = (\mathbf{z}, \mathbf{e})$. Each hint will introduce new constraints on \mathbf{s} and will ultimately decrease the security level.

Non-Commutativity. It should be noted that many types of hints commute: Integrating them in any order will lead to the same DBDD instance. Potential exceptions are **non-smooth modular hints** (See later in Sect. 4.2) and **a posteriori approximate hints** (See later in Sect. 4.4): they do not always commute with the other types of hints, and do not always commute between themselves, unless the vectors \mathbf{v} 's of those hints are all orthogonal to each other. The reason is: in these cases, the distribution in the direction of \mathbf{v} is redefined which erases the prior information.

4.1 Perfect Hints

Definition 21 (Perfect hint). *A perfect hint on the secret \mathbf{s} is the knowledge of $\mathbf{v} \in \mathbb{Z}^{d-1}$ and $l \in \mathbb{Z}$, such that*

$$\langle \mathbf{s}, \mathbf{v} \rangle = l.$$

A perfect hint is quite strong in terms of additional knowledge. It allows decreasing the dimension of the lattice by one and increases its volume. One could expect such hints to arise from the following scenarios:

- The full leak without noise of an original coefficient, or even an unreduced intermediate register since most of the computations are linear. For the second case, one may note that optimized implementations of NTT typically attempt to delay the first reduction modulo q , so leaking a register on one of the first few levels of the NTT would indeed lead to such a hint.
- A noisy leakage of the same registers, but with still a rather high guessing confidence. In that case it may be worth making the guess while decreasing the success probability of the attack.¹ This could happen in a cold-boot attack scenario. This is also the case in the single trace attack on Frodo [9] that we will study as one of our examples in Sect. 6.1.
- More surprisingly, certain schemes, including some NIST candidates offer such a hint ‘by design’. Indeed, LAC, Round5 and NTRU-HPS all choose ternary secret vectors with a prescribed number of 1’s and -1 ’s, which directly induce one or two such perfect hints. This will be detailed in Sect. 6.3.

Integrating a perfect hint into a DBDD instance. Let $\mathbf{v} \in \mathbb{Z}^{d-1}$ and $l \in \mathbb{Z}$ be such that $\langle \mathbf{s}, \mathbf{v} \rangle = l$. Note that the hint can also be written as

$$\langle \bar{\mathbf{s}}, \bar{\mathbf{v}} \rangle = 0,$$

where $\bar{\mathbf{s}}$ is the extended LWE secret as defined in Eq. (5) and $\bar{\mathbf{v}} := (\mathbf{v}; -l)$.

Remark 22. Here we understand the interest of using Kannan’s embedding *before* integrating hints rather than after: it allows to also homogenize the hint, and therefore to make Λ' a proper lattice rather than a lattice coset (i.e. a shifted lattice).

Including this hint is done by modifying the $\text{DBDD}_{\Lambda, \mu, \Sigma}$ to $\text{DBDD}_{\Lambda', \mu', \Sigma'}$, where:

$$\Lambda' = \Lambda \cap \{ \mathbf{x} \in \mathbb{Z}^d \mid \langle \mathbf{x}, \bar{\mathbf{v}} \rangle = 0 \}$$

$$\Sigma' = \Sigma - \frac{(\bar{\mathbf{v}}\Sigma)^T \bar{\mathbf{v}}\Sigma}{\bar{\mathbf{v}}\Sigma\bar{\mathbf{v}}^T} \tag{10}$$

$$\mu' = \mu - \frac{\langle \bar{\mathbf{v}}, \mu \rangle}{\bar{\mathbf{v}}\Sigma\bar{\mathbf{v}}^T} \bar{\mathbf{v}}\Sigma \tag{11}$$

¹ One may then re-amplify the success probability by retrying the attack making guesses at different locations.

We now explain how to derive the new mean $\boldsymbol{\mu}'$ and the new covariance $\boldsymbol{\Sigma}'$. Let y be the random variable $\langle \bar{\mathbf{s}}, \bar{\mathbf{v}} \rangle$, where $\bar{\mathbf{s}}$ has mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$. Then $\boldsymbol{\mu}'$ is the mean of $\bar{\mathbf{s}}$ conditioned on $y = 0$, and $\boldsymbol{\Sigma}'$ is the covariance of $\bar{\mathbf{s}}$ conditioned on $y = 0$. Using Corollary 7, we obtain the corresponding conditional mean and covariance.

We note that lattice Λ' is an intersection of Λ and a hyperplane orthogonal to $\bar{\mathbf{v}}$. Given \mathbf{B} as basis of Λ , by Lemma 9 a basis of Λ' can be computed as follows:

1. Let \mathbf{D} be dual basis of \mathbf{B} . Compute $\mathbf{D}_\perp := \mathbf{D} \cdot \boldsymbol{\Pi}_{\bar{\mathbf{v}}}^\perp$.
2. Apply the LLL algorithm on \mathbf{D}_\perp to eliminate linear dependencies. Then delete the first row of \mathbf{D}_\perp (which is $\mathbf{0}$ because with the hyperplane intersection, the dimension of the lattice is decremented).
3. Output the dual of the resulting matrix.

While polynomial time, the above computation is quite heavy, especially as there is no convenient library offering a parallel version of LLL. Fortunately, for predicting attack costs, one only needs the dimension of the lattice Λ and its volume. These can easily be computed assuming $\bar{\mathbf{v}}$ is a primitive vector (see Definition 10) of the dual lattice: the dimension decreases by 1, and the volume increases by a factor $\|\bar{\mathbf{v}}\|$. This is stated and proved in Lemma 11. Intuitively, the primitivity condition is needed since then one can scale the leak to $\langle \mathbf{s}, f\mathbf{v} \rangle = fl$ for any non-zero factor $f \in \mathbb{R}$ and get an equivalent leak; however there is only one factor f that can ensure that $f\bar{\mathbf{v}} \in \Lambda^*$, and is primitive in it.

Remark 23. Note that if $\bar{\mathbf{v}}$ is not in the span of Λ —as typically occurs if other non-orthogonal perfect hints have already been integrated—Lemma 11 should be applied to the orthogonal projection $\bar{\mathbf{v}}' = \bar{\mathbf{v}} \cdot \boldsymbol{\Pi}_\Lambda$ of $\bar{\mathbf{v}}$ onto Λ . Indeed, the perfect hint $\langle \bar{\mathbf{s}}, \bar{\mathbf{v}}' \rangle = 0$ replacing $\bar{\mathbf{v}}$ by $\bar{\mathbf{v}}'$ is equally valid.

4.2 Modular Hints

Definition 24 (Modular hint). A modular hint on the secret \mathbf{s} is the knowledge of $\mathbf{v} \in \mathbb{Z}^{d-1}$, $k \in \mathbb{Z}$ and $l \in \mathbb{Z}$, such that

$$\langle \mathbf{s}, \mathbf{v} \rangle = l \pmod k.$$

We can expect such hints to arise from several scenarios:

- obtaining the value of an intermediate register during LWE decryption would likely correspond to giving such a modular equation modulo q . This is also the case if an NTT coefficient leaks in a Ring-LWE scheme. It can also occur “by design” if the LWE secret is chosen so that certain NTT coordinates are fixed to 0 modulo q , as is the case in some instances of Order LWE [6].
- obtaining the absolute value $a = |s|$ of a coefficient s implies $s = a \pmod{2a}$, and such a hint could be obtained by a timing attack on an unprotected implementation of a table-based sampler, in the spirit of [17].

- obtaining the Hamming weight of the string $b_1b_2 \dots b'_1b'_2 \dots$ used to sample a centered binomial coefficient $s = \sum b_i - \sum b'_i$ (as done in NewHope and Kyber [31,34]) reveals in particular $s \bmod 2$. Indeed, the latter string (or at least some parts of it) is more likely to be leaked than the Hamming weight of s .

Integrating a modular hint into a DBDD instance. Let $\mathbf{v} \in \mathbb{Z}^{d-1}$; $k \in \mathbb{Z}$ and $l \in \mathbb{Z}$ be such that $\langle \mathbf{s}, \mathbf{v} \rangle = l \pmod k$. Note that the hint can also be written as

$$\langle \bar{\mathbf{s}}, \bar{\mathbf{v}} \rangle = 0 \pmod k \tag{12}$$

where $\bar{\mathbf{s}}$ is the extended LWE secret as defined in Eq. 5 and $\bar{\mathbf{v}} := (\mathbf{v}; -l)$. We refer to Remark 22 for the legitimacy of such dimension increase.

Smooth case. Intuitively, such a hint should only sparsify the lattice, and leave the average and the variance unchanged. This is not entirely true, this is only (approximately) true when the variance is sufficiently large in the direction of \mathbf{v} to ensure smoothness, i.e. when $k^2 \ll \mathbf{v}\Sigma\mathbf{v}^T$; one can refer to [28, Lemma 3.3 and Lemma 4.2] for the quality of that approximation. In this smooth case, we therefore have:

$$\Lambda' = \Lambda \cap \{ \mathbf{x} \in \mathbb{Z}^d \mid \langle \mathbf{x}, \bar{\mathbf{v}} \rangle = 0 \pmod k \} \tag{13}$$

$$\boldsymbol{\mu}' = \boldsymbol{\mu} \tag{14}$$

$$\Sigma' = \Sigma \tag{15}$$

On the other hand, if $k^2 \gg \mathbf{v}\Sigma\mathbf{v}^T$, then the residual distribution will be highly concentrated on a single value, and one should therefore instead use a perfect $\langle \mathbf{s}, \mathbf{v} \rangle = l + ik$ for some i .

General case. In the general case, one can resort to a numerical computation of the average μ_c and the variance σ_c^2 of the one-dimensional centered discrete Gaussian of variance $\sigma^2 = \mathbf{v}\Sigma\mathbf{v}^T$ over the coset $l+k\mathbb{Z}$, and apply the corrections:

$$\boldsymbol{\mu}' = \boldsymbol{\mu} + \frac{\mu_c - \langle \bar{\mathbf{v}}, \boldsymbol{\mu} \rangle}{\bar{\mathbf{v}}\Sigma\bar{\mathbf{v}}^T} \bar{\mathbf{v}}\Sigma \tag{16}$$

$$\Sigma' = \Sigma + \left(\frac{\sigma_c^2}{(\bar{\mathbf{v}}\Sigma\bar{\mathbf{v}}^T)^2} - \frac{1}{\bar{\mathbf{v}}\Sigma\bar{\mathbf{v}}^T} \right) (\bar{\mathbf{v}}\Sigma)^T (\bar{\mathbf{v}}\Sigma) \tag{17}$$

Intuitively, these formulae completely erase prior information on $\langle \mathbf{s}, \bar{\mathbf{v}} \rangle$, before it is replaced by the new average and variance in the adequate direction. Both can be derived² using Corollary 7.

As for perfect hints, the computation of Λ' can be done by working on the dual lattice. More specifically:

² We are thankful to Thibault Feneuil for pointing out an incorrect equation in a previous version of this paper.

1. Let \mathbf{D} be dual basis of \mathbf{B} .
2. Redefine $\bar{\mathbf{v}} \leftarrow \bar{\mathbf{v}} \cdot \mathbf{\Pi}_A$, noting that this does not affect the validity of the hint.
3. Append $\bar{\mathbf{v}}/k$ to \mathbf{D} and obtain \mathbf{D}'
4. Apply the LLL algorithm on \mathbf{D}' to eliminate linear dependencies. Then delete the first row of \mathbf{D}' (which is $\mathbf{0}$ since we introduced a linear dependency).
5. Output the dual of the resulting matrix.

Also, as for perfect hints the parameters of the new lattice Λ' can be predicted: the dimension is unchanged, and the volume increases by a factor k under a primitivity condition, which is proved by Lemma 12.

4.3 Approximate Hints (conditioning)

Definition 25 (Approximate hint). *An approximate hint on the secret \mathbf{s} is the knowledge of $\mathbf{v} \in \mathbb{Z}^{d-1}$ and $l \in \mathbb{Z}$, such that*

$$\langle \mathbf{s}, \mathbf{v} \rangle + e = l,$$

where e models noise following a distribution $N_1(0, \sigma_e^2)$, independent of \mathbf{s} .

One can expect such hints from:

- any noisy side channel information about a secret coefficient. This is the case of our study in Sect. 6.1.
- decryption failures. In Sect. 6.2, we show how this type of hint can represent the information gained by a decryption failure.

To include this knowledge in the DBDD instance, we must combine this knowledge with the prior knowledge on the solution \mathbf{s} of the instance.

Integrating an approximate hint into a DBDD instance. Let $\mathbf{v} \in \mathbb{Z}^{d-1}$ and $l \in \mathbb{Z}$ be such that $\langle \mathbf{s}, \mathbf{v} \rangle \approx l$. Note that the hint can also be written as

$$\langle \bar{\mathbf{s}}, \bar{\mathbf{v}} \rangle + e = 0 \tag{18}$$

where $\bar{\mathbf{s}}$ is the extended LWE secret as defined in Eq. (5), $\bar{\mathbf{v}} := (\mathbf{v}; -l)$, and e has $N_1(0, \sigma_e^2)$ distribution. The unique shortest non-zero solution of $\text{DBDD}_{\Lambda, \mu, \Sigma}$, is also the unique solution of the instance $\text{DBDD}_{\Lambda', \mu', \Sigma'}$ where

$$\Lambda' = \Lambda \tag{19}$$

$$\Sigma' = \Sigma - \frac{(\bar{\mathbf{v}}\Sigma)^T \bar{\mathbf{v}}\Sigma}{\bar{\mathbf{v}}\Sigma\bar{\mathbf{v}}^T + \sigma_e^2} \tag{20}$$

$$\mu' = \mu - \frac{\langle \bar{\mathbf{v}}, \mu \rangle}{\bar{\mathbf{v}}\Sigma\bar{\mathbf{v}}^T + \sigma_e^2} \bar{\mathbf{v}}\Sigma \tag{21}$$

We note that Eq. (19) comes from

$$\Lambda' := \Lambda \cap \{ \mathbf{x} \in \mathbb{Z}^d \mid \langle \mathbf{x}, \bar{\mathbf{v}} \rangle + e = 0, \text{ for all possible } e \sim N_1(0, \sigma_e^2) \} = \Lambda.$$

The new covariance and mean follow from Corollary 7.

Consistency with Perfect Hint. Note that if $\sigma_e = 0$, we fall back to a perfect hint $\langle \mathbf{s}, \mathbf{v} \rangle = l$. The above computation of Σ' (20) (resp. $\boldsymbol{\mu}'$ (21)) is indeed equivalent to Eq. (10) (resp. Eq. (11)) from Sect. 4.1. Note however, in our implementation, that to avoid singularities, we require the span of $\text{Span}(\Sigma + \boldsymbol{\mu}^T \boldsymbol{\mu}) = \text{Span}(A)$ (See the requirement in Eq. (2)): If $\sigma_e = 0$, one *must* instead use a Perfect hint.

Multi-dimensional approximate hints. The formulae of [24] are even more general, and one could consider a multidimensional hint of the form $\mathbf{s}\mathbf{V} + \mathbf{e} = \mathbf{l}$, where $\mathbf{V} \in \mathbb{R}^{n \times k}$ and \mathbf{e} a gaussian noise of any covariance $\Sigma_{\mathbf{e}}$. However, those general formulae require explicit matrix inversion which becomes impractical in large dimension. We therefore only implemented full-dimensional ($k = n$) hint integration in the *super-lightweight* version of our tool, which assumes all covariance matrices to be diagonal. These will be used for hints obtained from decryption failures in Sect. 6.2.

4.4 Approximate Hint (*a posteriori*)

In certain scenarios, one may more naturally obtain directly the a posteriori distribution of $\langle \mathbf{s}, \mathbf{v} \rangle$, rather than a hint $\langle \mathbf{s}, \mathbf{v} \rangle + e = l$ for some error e independent of \mathbf{s} . Such a scenario is typical in template attacks, as we exemplify via the single trace attack on Frodo from [9], which we study in Sect. 6.1.

Given the a posteriori distribution of $\langle \bar{\mathbf{s}}, \bar{\mathbf{v}} \rangle$, one can derive its mean μ_{ap} and variance σ_{ap}^2 and apply the corrections to compute the new mean and covariance exactly as in Eqs. (16) and (17).

4.5 Short Vector Hints

Definition 26 (Short vector hint). *A short vector hint on the lattice Λ is the knowledge of a short vector $\bar{\mathbf{v}}$ such that*

$$\bar{\mathbf{v}} \in \Lambda.$$

Note that such hints are not related to the secret, and are not expected to be obtained by side-channel information, but rather by the very design of the scheme. In particular, the lattice Λ underlying LWE instance modulo q contains the so-called q -vectors, i.e. the vectors $(q, 0, 0, \dots, 0)$ and its permutations. These vectors are in fact implicitly exploited in the literature on the cryptanalysis of LWE since at least [23]. Indeed, in some regimes, the best attacks are obtained by ‘forgetting’ certain LWE equations, which can be geometrically interpreted as a projection orthogonally to a q -vector. Note that, among all hints, the short vector hints should be the last to be integrated. In our context, we need to generalize this idea beyond q -vector because the q -vectors may simply disappear after the integration of a perfect or modular hint. For example, after the integration of a perfect hint $\langle \mathbf{s}, (1, 1, \dots, 1) \rangle = 0$, all the q -vectors are no longer in the lattice, but $(q, -q, 0, \dots, 0)$ still is, and so are all its permutations.

Resolving the DBDD problem resulting from this projection will not directly lead to the original secret, as projection is not injective. However, as long as we keep $n + 1$ dimensions out of the $n + m + 1$ dimensions of the original LWE instance, we can still efficiently reconstruct the full LWE secret by solving a linear system over the rationals.

Integrating a short vector hint into a DBDD instance. It is the case when the secret vector is short enough to be a solution after applying projection $\Pi_{\bar{\mathbf{v}}}^\perp$ on $\text{DBDD}_{\Lambda, \Sigma, \mu}$.

$$\Lambda' = \Lambda \cdot \Pi_{\bar{\mathbf{v}}}^\perp \tag{22}$$

$$\Sigma' = (\Pi_{\bar{\mathbf{v}}}^\perp)^T \cdot \Sigma \cdot \Pi_{\bar{\mathbf{v}}}^\perp \tag{23}$$

$$\mu' = \mu \cdot \Pi_{\bar{\mathbf{v}}}^\perp \tag{24}$$

To compute a basis of Λ' one can simply apply the projection to all the vectors of its current basis, and then eliminate linear dependencies in the resulting basis using LLL.

Remark 27. Once a short vector hint $\bar{\mathbf{v}} \in \Lambda$ has been integrated, Λ has been transformed into Λ' . And, if one has to perform another short vector hint integration $\bar{\mathbf{v}}_1 \in \Lambda$, $\bar{\mathbf{v}}_1$ should be projected onto Λ' with $\bar{\mathbf{v}} \cdot \Pi_{\Lambda'} \in \Lambda'$. In our implementation however, this has been taken into account and one can simply apply the same transformation as above, replacing a single vector $\bar{\mathbf{v}}$ by a matrix \mathbf{V} .

The dimension of the lattice decreases by one (or by k , if one directly integrates a matrix of k vectors) and the volume of the lattice also decreases according to Fact 13. One can also predict the decrease of the determinant of Σ via the identity:

$$\text{rdet}(\Sigma') = \text{rdet}(\Sigma) \cdot \frac{\|\bar{\mathbf{v}}\|^2}{\bar{\mathbf{v}} \Sigma \bar{\mathbf{v}}^T}, \quad \text{or} \quad \text{rdet}(\Sigma') = \text{rdet}(\Sigma) \cdot \frac{\det(\mathbf{V}\mathbf{V}^T)}{\det(\mathbf{V}\Sigma\mathbf{V}^T)}. \tag{25}$$

Worthiness and choice of short vector hints. Integrating such a hint induces a trade-off between the dimension and the volume, and therefore it is not always advantageous to integrate.

This raises the following potentially hard problem: given a set \mathbf{W} of short vectors of Λ (viewed as a matrix), which subset $\mathbf{V} \subset \mathbf{W}$ of size k lead to the easiest DBDD instance? Because the hardness of the new problem grows with

$$\frac{\text{rdet}(\Sigma')}{\text{Vol}(\Lambda')^2} = \frac{\text{rdet}(\Sigma)}{\text{Vol}(\Lambda)^2} \cdot \frac{\det(\mathbf{V}\mathbf{V}^T)^2}{\det(\mathbf{V}\Sigma\mathbf{V}^T)} \tag{26}$$

In the case of an un-hinted DBDD instance directly obtained from the LWE problem, for \mathbf{V} being the set of (primitive) q -vectors, the problem is easier: all subsets of size k lead to instances with the same parameters.

But this is not true anymore as soon as Σ has been altered or if the set \mathbf{W} is arbitrary. For example, setting $\Sigma = \mathbf{I}$, one simply wishes to minimize $\det(\mathbf{V}\mathbf{V}^T)$; but for an arbitrary set \mathbf{W} the problem of finding the optimal subset $\mathbf{V} \subset \mathbf{W}$ is NP-hard [22], and remains NP-hard up to exponential approximation factors.

A natural approach to try to get an approximate solution in polynomial time consists in making sequential greedy choices. This involves computing $|\mathbf{V}| \cdot |\mathbf{W}|$ many matrix-vector products over increasingly large rationals, and appeared painfully slow in practice for making prediction on cryptographically large instances. Fortunately, in the typical cases where the vectors of \mathbf{W} are the q -vectors, this can be made somewhat practical (See Sect. 6.3 for example).

Remark 28. When the basis of an LWE-lattice is given in its systematic form, the q -vectors are already explicitly given to lattice reduction algorithms, and these algorithms will implicitly make use of them when they are worthy, as if we had integrated them. The reason is that lattice reduction algorithm naturally work with projected sublattices, and if a q -vector is shorter than what the algorithm can produce, those q -vectors will remain untouched at the beginning of the basis; the reduction algorithm will effectively work on the lattice projected orthogonally to them. In other words, integrating q -vectors is important to understand and predict how lattice reduction algorithm will work, but, in certain cases they may be automatically detected and exploited by lattice reduction algorithms themselves.

5 Implementation

5.1 Our Sage Implementation

We propose three implementations of our framework, all following the same python/sage 9.0 API.³ More specifically, the API and some common functions are defined in `DBDD_generic.sage`, as a class `DBDD_Generic`. Three derived classes are then given:

1. The class `DBDD` (provided in `DBDD.sage`) is the *full-fledged* implementation: i.e. it fully maintains all information about a `DBDD` instance as one integrates hints: the lattice Λ , the covariance matrix Σ and the average μ . While polynomial time, maintaining the lattice information can be quite slow, especially since consecutive intersections with hyperplanes can lead to manipulations on rationals with large denominators. It also allows to finalize the attack, running the homogenization, isotropization and lattice reduction, based on the `fpLLL` [15] library available through sage.

We note that if one were to repeatedly use perfect or modular hints, a lot of effort would be spent on uselessly alternating between the primal and the dual lattice. Instead, we implement a caching mechanism for the primal and dual basis, and only update them when necessary.

³ While we would have preferred a full python implementation, we are making a heavy use of linear algebra over the rationals for which we could find no convenient python library.

2. The class `DBDD_predict` (provided in `DBDD_predict.sage`) is the *lightweight* implementation: it only fully maintains the covariance information, and the parameters of the lattice (dimension, volume). It must therefore work under assumptions about the primitivity of the vector \mathbf{v} ; in particular, it cannot detect hints that are redundant. If one must resort to this faster variant on large instances, it is advised to consider potential (even partial) redundancy between the given hints, and to run a comparison with the previous on small instances with similarly generated hints.
3. The class `DBDD_predict_diag` (provided in `DBDD_predict_diag.sage`) is the *super-lightweight* implementation. It maintains the same information as the above, but requires the covariance matrix to remain diagonal at all times. In particular, one can only integrate hints for which the directional vector \mathbf{v} is colinear with a canonical vector.

5.2 Tests and Validation

In the full version of our paper, we present a demonstration of our tool with some extracts of Sage 9.0 code. We implement two tests to verify the correctness of our scripts, and more generally the validity of our predictions.

Consistency checks. Our first test (`check_consistency.sage`) simply verifies that all three classes always agree perfectly. More specifically we run all three versions on a given instances, integrating the same random hint in all of them, and compare their hardness prediction. We first test using the full-fledged version that the primitivity condition does hold, and discard the hint if not, as we know that predictions cannot be correct on such hints. This verification passes.

Prediction verifications. We now verify experimentally the prediction made by our tool for various types of hints, by comparing those predictions to actual attack experiments (see `compare_usvp_models.sage` for the prediction without hints and `prediction_verifications.sage` for the prediction with hints). This is done for a given set of LWE parameters, and increasing the number of hints. The details of the experiments and the results are given in Fig. 5.

While our predictions seem overall accurate, we still note a minor discrepancy of up to 2 or 3 bikz in the low blocksize regime. This exceeds the error made by prediction on the attack without any hint, which was below 1 bikz, even in the same low blocksize regime. We suspected that this discrepancy is due to residual q -vectors, or small combinations of them, that are hard to predict for randomly generated hints, but would still benefit by lattice reduction. We tested that hypothesis by running similar experiments, but leaving certain coordinates untouched by hints, so to still explicitly know some q -vectors for short-vector hint integration, if they are “worthy”. This didn’t to improve the accuracy of our prediction, which infirms our suspected explanation. We are at the moment unable to explain this inaccuracy. We nevertheless find our predictions satisfactory, considering that even without hints, previous predictions [3] were much less accurate (see Fig. 4).

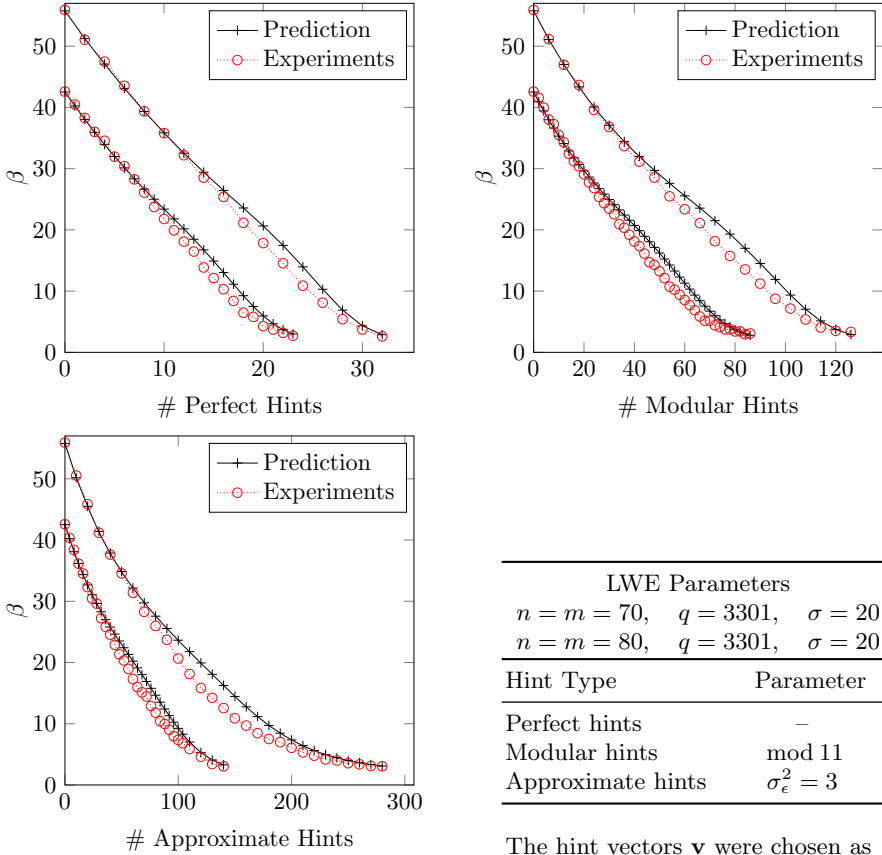


Fig. 5. Experimental verification of the security decay predictions for each type of hints. Each data point was averaged over 256 samples.

6 Applications Examples

6.1 Hints from Side Channels

In [9], W. Bos et al. study the feasibility of a single-trace power analysis of the Frodo Key Encapsulation Mechanism (FrodoKEM) [29]. Specifically, in the first approach, they analyze the possibility of a divide-and-conquer attack targeting a multiplication in the key generation. This attack was claimed unsuccessful in [9] because the bruteforce phase after recovering a candidate for the private key was too expensive. Along with this unsuccessful result, a successful powerful extend-and-prune attack is provided in [9].

We emphasize that the purpose of this section is to exemplify our tool on a standard side-channel attack, and this is why we choose the former unsuccessful

divide-and-conquer attack of [9]. The point of this section is to show that our framework can indeed lead to improvements in the algorithmic phase of a side-channel attack, once the leak has been fixed.

FrodoKEM. FrodoKEM is based on small-secret-LWE; we outline here some details necessary to understand the attack. Note that we use different letter notations from [29] for consistency. For parameters n and q , the private key is $(\mathbf{z} \in \mathbb{Z}_q^n, \mathbf{e} \in \mathbb{Z}_q^n)$ where the coefficients of \mathbf{z} and \mathbf{e} , denoted \mathbf{z}_i and \mathbf{e}_i , can take several values in a small set that we denote L . The public key is $(\mathbf{A} \in \mathbb{Z}_q^{n \times n}, \mathbf{b} = \mathbf{z}\mathbf{A} + \mathbf{e})$. The goal of the attack is to recover \mathbf{z} by making measurements during the multiplication between \mathbf{z} and \mathbf{A} when computing \mathbf{b} in the key generation. Note that there is no multiplication involving \mathbf{e} and thus it is not targeted in this attack. Six sets of parameters are considered: CCS1, CCS2, CCS3 and CCS4 introduced in [8] and NIST1 and NIST2 introduced in [29]. For example, with NIST1 parameters, $n = 640$, $q = 2^{15}$ and $L = \{-11, \dots, 11\}$.

Side-channel simulation. The divide-and-conquer attack provided by [9] simulates side-channel information using ELMO, a power simulator for a Cortex M0 [27]. This tool outputs simulated power traces using an elaborate leakage model with Gaussian noise. Thus, it is parametrized by the standard deviation of the side-channel noise. For proofs of concept, the authors of [27] suggest to choose the standard deviation of the simulated noise as $\sigma_{\text{SimNoise}} := 0.0045$ for realistic leakage modeling. This standard deviation was also the one chosen in [9, Fig. 2b] and W. Bos et al. implemented a Matlab script that calls ELMO to simulate the side-channel information applied on Frodo. This precise side-channel simulator was provided to us by the authors of [9] and we were able to re-generate all their data with Matlab, again using $\sigma_{\text{SimNoise}} = 0.0045$.

Template attack. The divide-and-conquer side-channel attack proposed by W. Bos et al. belongs in the template attack family. Template attacks were introduced in [10]. In a nutshell, these attacks include a profiling phase and an online phase. Let us detail the template attack for Frodo implemented in [9].

1. The profiling phase consists in using a copy of the device and recording a large number of traces using many different known secret values. From these measures, the attacker can derive the multidimensional distribution of several points of interest when the traces share the same secret coefficient. More precisely, in the case of FrodoKEM, for a given index $i \in [0, n - 1]$, the points of interest will be the instants in the trace when \mathbf{z}_i is multiplied by the coefficients of \mathbf{A} (n interest points in total). Let us define

$$\mathbf{c}_i := (T[t_{i,0}], \dots, T[t_{i,n-1}]) \quad \mathbf{c} \in \mathbb{R}^n, \quad (27)$$

where T denotes the trace measurement and $(t_{i,k})$ denotes the instants of the multiplication of \mathbf{z}_i with the coefficients $\mathbf{A}_{i,k}$ for $(i, k) \in [0, n - 1]$. The random variable vector associated to \mathbf{c}_i is denoted by C_i . For each $i \in [0, n - 1]$ and

$x \in L$, the goal of the profiling phase is to learn the center of the probability distribution

$$A_{i,x}(\mathbf{c}) := P[C_i = \mathbf{c} \mid \mathbf{z}_i = x].$$

By hypothesis, for template attacks (see [10, Section 2.1]), $A_{i,x}$ is assumed to follow a multidimensional normal distribution of standard deviation $\sigma_{\text{SimNoise}} \cdot \mathbf{I}_n$. Thus, the attacker recovers the center of $A_{i,x}$ for each $i \in [0, n-1]$ and $x \in L$ by averaging all the measured \mathbf{c}_i that validate $\mathbf{z}_i = x$. The center of $A_{i,x}$ is denoted $\mathbf{t}_{i,x}$ and we call it a *template*. W. Bos et al. [9] actually assume that $\mathbf{t}_{i,x}$ depends only on x and is independent from the index i . Thus, $\mathbf{t}_{i,x} = \mathbf{t}_x$. Essentially, this common assumption implies that the index $i \in [0, n-1]$ of the target coefficient does not influence the leakage. Consequently, the attacker only has to derive $\mathbf{t}_{0,x}$, for example.

2. In a second step, the attacker knows the templates \mathbf{t}_x for all $x \in L$. She also knows the points of interest $t_{i,k}$ as defined above in Eq. 27. She will construct a candidate $\tilde{\mathbf{z}}$ for the secret \mathbf{z} by recovering the coefficients one by one. For each unknown secret coefficient \mathbf{z}_i , she takes the measurement \mathbf{c}_i as defined in Eq. 27. Using this measurement, she can derive an a posteriori probability distribution: With her fixed $i \in [0, n-1]$ and measured $\mathbf{c}_i \in \mathbb{R}$, she computes for all $x \in L$,

$$P[\mathbf{z}_i = x \mid C_i = \mathbf{c}_i] = \frac{P[\mathbf{z}_i = x]}{P[C_i = \mathbf{c}_i]} \cdot P[C_i = \mathbf{c}_i \mid \mathbf{z}_i = x] \tag{28}$$

$$\propto P[\mathbf{z}_i = x] \cdot \exp\left(-\frac{\|\mathbf{c}_i - \mathbf{t}_x\|_2^2}{2\sigma_{\text{SimNoise}}^2}\right) \tag{29}$$

In [9], a score table, denoted $(S_i[x])_{x \in L}$ is derived from the a posteriori distribution as follows,

$$S_i[x] := \ln(P[\mathbf{z}_i = x \mid C_i = \mathbf{c}_i]) \tag{30}$$

$$= \ln(P[\mathbf{z}_i = x]) - \frac{\|\mathbf{c}_i - \mathbf{t}_x\|_2^2}{2\sigma_{\text{SimNoise}}^2}. \tag{31}$$

Finally, the output candidate for \mathbf{z}_i is $\tilde{\mathbf{z}}_i := \operatorname{argmax}_{x \in L}(S_i[x])$.

One can use the presented attack as a “black-box” to generate the score tables using the script from [9]. As an example, using the NIST1 parameters, we show several measured scores $(S[-11], \dots, S[11])$ corresponding to several secret coefficients in Table 1. The first line corresponds to a secret equal to 0, the second line to 1 and the third and fourth line to -1 . The last line is an example of failed guessing because we see that the outputted candidate is not -1 . We remark that the values having the opposite sign are assigned a very low score, we conjecture that it is because the sign is filling the register and then the Hamming weight of the register will be very far from the correct one.

Table 1. Examples of scores associated to the secret values $\mathbf{s}_i \in \{0, \pm 1\}$, after the side-channel analysis of [9] for NIST1 parameters. The best score in each score table is highlighted. This best guess is correct for the first 3 score table, but incorrect for the last one.

z_i	S											
	-11	-10	-9	-8	-7	-6	-5	-4	-3	-2	-1	0
0	-4098	-3918	-4344	-2580	-3212	-3108	-3758	-3155	-3583	-3498	-3900	-340
1	-3273	-3114	-3491	-1951	-2495	-2405	-2972	-2445	-2819	-2744	-3098	-365
-1	-341	-335	-352	-465	-358	-369	-329	-362	-331	-334	-328	-3712
-1	-306	-298	-319	-414	-314	-323	-290	-317	-291	-293	-291	-3608

	...	1	2	3	4	5	6	7	8	9	10	11
0	...	-380	-367	-452	-818	-975	-933	-1084	-368	-459	-453	-592
1	...	-325	-328	-338	-546	-657	-627	-737	-333	-344	-342	-407
-1	...	-3079	-3195	-2656	-1696	-1461	-1521	-1329	-3231	-2648	-2685	-2201
-1	...	-2982	-3097	-2564	-1617	-1385	-1444	-1256	-3132	-2556	-2593	-2115

With this template attack, one can recover $\tilde{\mathbf{z}} \approx \mathbf{z}$. However, W. Bos et al. [9] could not conclude the attack with a key recovery even though much information leaked about the secret. Frustratingly, a bruteforce phase to derive \mathbf{z} from $\tilde{\mathbf{z}}$ did not lead to any security threat as stated in [9, Section 3]. They actually pointed out an interesting open question of whether “novel lattice reduction algorithms [can] take into account side-channel information”. Our work solves this open question by combining the knowledge obtained in the divide-and-conquer template attack of [9] with our framework.

From scores to hints. We first instantiate a DBDD instance with a chosen set of parameters. Then we assume that, for each secret coefficient z_i , we are given the associated score table S_i , thanks to the template attack that has already been carried out. We go back to the a posteriori distribution in Eq. 29 by applying the $\exp()$ function and renormalizing the score table. As an example, we show the probability distributions derived from Table 1, along with their variances and centers, in Table 2.

Finally, we use our framework to introduce n a posteriori *approximate hints* to our DBDD instance with the derived centers and variances for each score table. When the variance is exactly 0, we integrate perfect hints instead.

Results. One can reproduce this attack using the Sage 9.0 script `exploiting_SCA_from_Bos_et_al.sage`. The experimentally derived data containing the score tables is in the folder `Scores_tables_SCA` for which, as mentioned earlier, was generated with a simulated noise variance of 0.0045. One can note that the obtained security fluctuates a bit from instance to instance, as it depends on the strength of the hints, which themselves depend on the randomness of the scheme. In the first two lines of Table 3, we show the new security with the inclusion of the approximate hints averaged on 50 tests per set of parameters.

Table 2. Probability distributions derived from Table 1, along with variances and centers.

z_i	A posteriori distribution												center	variance	
	-11	-10	-9	-8	-7	-6	-5	-4	-3	-2	-1	0			
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
-1	0	0	0	0	0	0	0.26	0	0.04	0.00	0.70	0	0	0	0
-1	0	0	0	0	0	0	0.56	0	0.21	0.03	0.21	0	0	0	0

	...	1	2	3	4	5	6	7	8	9	10	11	center	variance
0	...	0	0	0	0	0	0	0	0	0	0	0	0	0
1	...	0.95	0.04	0	0	0	0	0	0.01	0	0	0	1.05	0.06
-1	...	0	0	0	0	0	0	0	0	0	0	0	-2.11	3.11
-1	...	0	0	0	0	0	0	0	0	0	0	0	-3.68	2.63

Table 3. Cost of the attacks without/with hints without/with guesses.

		NIST1	NIST2	CCS1	CCS2	CCS3	CCS4
Attack without hints	(bikz)	487	708	239	448	492	584
Attack with hints	(bikz)	330	423	128	123	219	230
Attack with hints & guesses	(bikz)	292	298	70	29	124	129
Number of guesses g		100	250	200	300	250	250
Success probability		0.86	0.64	0.87	0.77	0.81	0.84

Guessing. To improve the attack further, one can note from Table 2 that certain key values have a very high probability of being correct, and assuming each of these values are correct, one can replace an approximate hint with a perfect one. For example, considering the second line of Table 2, the secret has a probability of 0.95 to be 1 and thus guessing it trades a perfect hint for a decrease of the success probability of the attack by 5%. This hybrid attack exploiting hints, guesses and lattice reduction, works as follows. Let g be a parameter.

1. Include all the approximate and perfect hints given by the score tables,
2. Order the coefficients of the secret z_i according to the maximum value of their a posteriori distribution table,
3. Include perfect hints for the g first coefficients and then solve and check the solution.

Increasing the number of guesses g leads to a trade-off between the cost of the attack and its success probability. We have chosen here a success probability larger than 0.6, while reducing the attack cost by 38 to 145 bikz depending on the parameter set. Given that 1 bit of security corresponds roughly to 3 or 4 bikz, this is undoubtedly advantageous.

Remark 29. The refinement presented above are very recent (lastly improved on June 2020). We are grateful to the authors of [9] of for helping us reconstructing distributions from the score table.

We remark that, with these results, the attacks with guesses on the parameters CCS1 and CCS2 seem doable in practice while it was not the case with our original results. However, some improvements of the implementation remain to be done in order to actually mount the attack. The full-fledged implementation cannot handle in reasonable time the large matrices of the original DBDD instance. We require another class of implementation which fully maintains all information about the instance, like the DBDD class, and assumes that the covariance matrix Σ is diagonal to simplify the computations, like the `DBDD_predict_diag` class. We hope to report on such an implementation in a future update of this report.

Remark 30. It should be noted that, given a single trace, one cannot naively retry the attack to boost its success probability. Indeed, the “second-best” guess may already have a much lower success probability than the first. Setting up such an hybrid attack mixing lattice reduction within our framework and key-ranking appears to be an interesting problem.

6.2 Hints from Decryption Failures

Another kind of hint our framework can model are hints provided by decryption failures. Using our framework, we produce prediction on a decryption failure attack on FrodoKEM-976 that match very closely the ad-hoc analysis of [14]. Our analysis is deferred to the full version of this paper [13].

6.3 Structural Hints from Design

Interestingly, we can also incorporate structural information on the secret or error that is present in certain schemes. We present (slightly) improved attacks on several Round 2 NIST submissions (such as LAC, Round5, and NTRU) which use ternary distribution for secrets, with a prescribed numbers of 1’s and -1 ’s in the full version of our paper [13].

Acknowledgments. The authors would like to thank Marco Martinoli and his co-authors [9] for sharing their source code. We express our gratitude to Jan-Pieter D’anvers for sharing precious insights and intuitions, guiding toward the proper formalization of decryption failures as approximate hints. We also thank John Schanck for valuable references and discussions that lead to refinements of the section on NTRU. We are also grateful to Martin Albrecht, Henri Gilbert, Ange Martinelli, Thomas Prest and Thibault Feneuil and to the anonymous CRYPTO’2020 reviewers for valuable feedback on a preliminary version of this work.

References

1. Albrecht, M.R., et al.: Estimate all the {LWE, NTRU} schemes! In: Catalano, D., De Prisco, R. (eds.) SCN 2018. LNCS, vol. 11035, pp. 351–367. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-98113-0_19
2. Albrecht, M.R., Deo, A., Paterson, K.G.: Cold boot attacks on ring and module LWE keys under the NTT. IACR TCHES **2018**, 173–213 (2018)
3. 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
4. Alkim, E., Ducas, L., Pöppelmann, T., Schwabe, P.: Post-quantum key exchange—a new hope. In: 25th USENIX Security Symposium (USENIX Security 2016), pp. 327–343 (2016)
5. Bai, S., Miller, S., Wen, W.: A refined analysis of the cost for solving LWE via uSVP. Cryptology ePrint Archive, Report 2019/502 (2019)
6. Bolboceanu, M., Brakerski, Z., Perlman, R., Sharma, D.: Order-LWE and the hardness of ring-LWE with entropic secrets. In: Galbraith, S.D., Moriai, S. (eds.) ASIACRYPT 2019. LNCS, vol. 11922, pp. 91–120. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-34621-8_4
7. Bootle, J., Delaplace, C., Espitau, T., Fouque, P.-A., Tibouchi, M.: LWE without modular reduction and improved side-channel attacks against BLISS. In: Peyrin, T., Galbraith, S. (eds.) ASIACRYPT 2018. LNCS, vol. 11272, pp. 494–524. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-03326-2_17
8. Bos, J.W., et al.: Frodo: take off the ring! Practical, quantum-secure key exchange from LWE, pp. 1006–1018 (2016)
9. Bos, J.W., Friedberger, S., Martinoli, M., Oswald, E., Stam, M.: Assessing the feasibility of single trace power analysis of Frodo. In: Cid, C., Jacobson Jr., M. (eds.) SAC 2018. LNCS, vol. 11349, pp. 216–234. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-10970-7_10
10. Chari, S., Rao, J.R., Rohatgi, P.: Template attacks. In: Kaliski, B.S., Koç, K., Paar, C. (eds.) CHES 2002. LNCS, vol. 2523, pp. 13–28. Springer, Heidelberg (2003). https://doi.org/10.1007/3-540-36400-5_3
11. 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
12. Cheon, J.H., Kim, D., Lee, J., Song, Y.: Lizard: cut off the tail! A practical post-quantum public-key encryption from LWE and LWR. In: Catalano, D., De Prisco, R. (eds.) SCN 2018. LNCS, vol. 11035, pp. 160–177. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-98113-0_9
13. Dachman-Soled, D., Ducas, L., Gong, H., Rossi, M.: LWE with side information: attacks and concrete security estimation. Cryptology ePrint Archive, Report 2020/292 (2020). <https://eprint.iacr.org/2020/292>
14. D’Anvers, J.-P., Vercauteren, F., Verbaauwhede, I.: On the impact of decryption failures on the security of LWE/LWR based schemes. IACR Cryptology ePrint Archive, 2018:1089 (2018)
15. The FPLLL Development Team: FPLLL, a lattice reduction library (2016). <https://github.com/fplll/fplll>
16. Garcia-Morchon, O., et al.: Round5. Technical report, NIST (2019)

17. Groot Bruinderink, L., Hülsing, A., Lange, T., Yarom, Y.: Flush, gauss, and reload—a cache attack on the BLISS lattice-based signature scheme. In: Gierlichs, B., Poschmann, A.Y. (eds.) CHES 2016. LNCS, vol. 9813, pp. 323–345. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53140-2_16
18. Groot Bruinderink, L., Pessl, P.: Differential fault attacks on deterministic lattice signatures. *IACR TCHES* **2018**, 21–43 (2018)
19. Hoffstein, J., Howgrave-Graham, N., Pipher, J., Whyte, W.: Practical lattice-based cryptography: NTRUEncrypt and NTRUSign. In: Nguyen, P., Vallée, B. (eds.) The LLL Algorithm. ISC, pp. 349–390. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-02295-1_11
20. Howgrave-Graham, N.: A hybrid lattice-reduction and meet-in-the-middle attack against NTRU. In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 150–169. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-74143-5_9
21. Kannan, R.: Minkowski’s convex body theorem and integer programming. *Math. Oper. Res.* **12**, 415–440 (1987)
22. Khachiyan, L.: On the complexity of approximating extremal determinants in matrices. *J. Complex.* **11**, 138–153 (1995)
23. Lindner, R., Peikert, C.: Better key sizes (and attacks) for LWE-based encryption. In: Kiayias, A. (ed.) CT-RSA 2011. LNCS, vol. 6558, pp. 319–339. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-19074-2_21
24. Liu, L.-P.: Linear transformation of multivariate normal distribution: marginal, joint and posterior. http://www.cs.columbia.edu/~liulp/pdf/linear_normal_dist.pdf. Accessed Sept 2019
25. Lu, X., et al.: PQC round-2 candidate: LAC. Technical report, NIST (2019). <https://csrc.nist.gov/projects/post-quantum-cryptography/round-2-submissions>
26. Martinet, J.: Perfect Lattices in Euclidean Spaces, vol. 327. Springer, Heidelberg (2013). <https://doi.org/10.1007/978-3-662-05167-2>
27. McCann, D., Oswald, E., Whitnall, C.: Towards practical tools for side channel aware software engineering: ‘grey box’ modelling for instruction leakages. In: 26th USENIX Security Symposium (USENIX Security 2017), Vancouver, BC, pp. 199–216. USENIX Association, August 2017
28. Micciancio, D., Regev, O.: Worst-case to average-case reductions based on Gaussian measures. *SIAM J. Comput.* **37**, 267–302 (2007)
29. Naehrig, M., et al.: FrodoKEM. Technical report, National Institute of Standards and Technology (2017). <https://csrc.nist.gov/projects/post-quantum-cryptography/round-1-submissions>
30. Nguyen, P.: Giophanthus and *LWR-based submissions (2019). Comment on the NIST PQC forum. <https://groups.google.com/a/list.nist.gov/d/msg/pqc-forum/nZBIBvYmmUI/J0pug16CBgAJ>
31. Pöppelmann, T., et al.: NewHope. Technical report, NIST (2019)
32. Ravi, P., Jhanwar, M.P., Howe, J., Chattopadhyay, A., Bhasin, S.: Side-channel assisted existential forgery attack on Dilithium - a NIST PQC candidate. *Cryptology ePrint Archive, Report 2018/821* (2018)
33. Ravi, P., Jhanwar, M.P., Howe, J., Chattopadhyay, A., Bhasin, S.: Exploiting determinism in lattice-based signatures: practical fault attacks on pqm4 implementations of NIST candidates. In: Asia CCS 2019, pp. 427–440. Association for Computing Machinery (2019)
34. Schwabe, P., et al.: CRYSTALS-KYBER. Technical report, NIST (2019)
35. Zhang, Z., et al.: PQC round-2 candidate: NTRU. Technical report, NIST (2019). <https://csrc.nist.gov/projects/post-quantum-cryptography/round-2-submissions>