



HAL
open science

Identity-Based Encryption from Lattices Using Approximate Trapdoors

Malika Izabachène, Lucas Prabel, Adeline Roux-Langlois

► **To cite this version:**

Malika Izabachène, Lucas Prabel, Adeline Roux-Langlois. Identity-Based Encryption from Lattices Using Approximate Trapdoors. ACISP 2023 - 28th Australasian Conference on Information Security and Privacy, Jul 2023, Brisbane, Australia. pp.270-290, 10.1007/978-3-031-35486-1_13. hal-04163136

HAL Id: hal-04163136

<https://hal.science/hal-04163136v1>

Submitted on 17 Jul 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Identity-Based Encryption from Lattices Using Approximate Trapdoors

Malika Izabachène¹, Lucas Prabel²(✉), and Adeline Roux-Langlois³

¹ Independent Scholar

² Univ Rennes, CNRS, IRISA, Rennes, France

lucas.prabel@irisa.fr

³ Normandie Univ, UNICAEN, ENSICAEN, CNRS, GREYC, 14000 Caen, France

Abstract. Practical implementations of advanced lattice-based constructions have received much attention since the first practical scheme instantiated over NTRU lattices, proposed by Prest et al (Asiacrypt 2014). They are using powerful lattice-based building blocks which allow to build Gaussian preimage sampling and trapdoor generation efficiently. In this paper, we propose two different constructions and implementations of identity-based encryption schemes (IBE) using approximate variants of "gadget-based" trapdoors introduced by Chen et al. (Asiacrypt 2019). Both constructions are proven secure.

Our first IBE scheme is an adaptation of the Bert et al. scheme (PQCrypto 2021) to the approximate setting, relying on the Module-LWE hardness assumption and making use of the Micciancio-Peikert paradigm with approximate trapdoors. The second IBE relies on a variant of the NTRU hardness assumption.

We provide several timings and a comparison analysis to explain our results. The two different instantiations give interesting trade-offs in terms of security and efficiency and both benefit from the use of approximate trapdoors. Though our second IBE construction is less efficient than other NTRU-based IBEs, we believe our work provides useful insights into efficient advanced lattice-based constructions.

Keywords: Lattice-based cryptography · approximate trapdoors · Gaussian preimage sampling · module lattices · IBE

1 Introduction

Identity-based encryption (IBE) is an advanced public key encryption scheme in which an identity, such as a username, email address, or social security number, acts as the public key. In identity-based encryption, the sender encrypts a message with the unique identity of the recipient, and the recipient then decrypts

This preprint has not undergone peer review (when applicable) or any post-submission improvements or corrections. The Version of Record of this contribution is published in Information Security and Privacy - 28th Australasian Conference, ACISP 2023, Brisbane, QLD, Australia, July 5-7, 2023, Proceedings, and is available online at https://doi.org/10.1007/978-3-031-35486-1_13.

the ciphertext with their private key to obtain the original message. In this way, one party can send an encrypted message to any other party without requesting the recipient's public key beforehand. The pair "identity" and "associated secret key" acts as a classical public key and secret key pair in a classical public key encryption scheme.

The idea was to eliminate the need for a public certificate across email systems. These schemes allowed secure communication without exchanging user keys. In [31] Shamir presented a solution for an identity-based signature scheme but the first IBE constructions appeared only in 2001 in [6, 10] and were based respectively on bilinear maps and quadratic residue-based assumptions. However, these schemes were vulnerable to quantum attacks due to Shor's algorithm.

Lattice-based IBE constructions. In [18], Gentry, Peikert and Vaikuntanathan described the first lattice-based IBE, relying on the Dual-Regev encryption scheme. An important contribution of their work was a sampling algorithm (known as GPV sampling) which showed how to use a short basis as a trapdoor for generating short lattice vectors. This sampler was then used to construct a lattice-based IBE scheme, proven adaptively secure against chosen-plaintext attack in the random oracle model as defined in [6, 10]. However, the master public key and user secret keys had large sizes in $O(n^2)$ bits. Later on, a construction of a Hierarchical IBE (HIBE) scheme in the standard model was proposed in [8] based on a new mechanism for users' keys delegation. This IBE scheme was proven secure in the selective model where the adversary needs to target an identity beforehand. In 2010, Agrawal et al. [1] proposed a Learning With Errors (LWE)-based IBE scheme with a trapdoor structure and with performance comparable to the GPV scheme. Their construction viewed an identity as a sequence of bits and then assigned a matrix to each bit. It used a sampling algorithm to obtain a basis with low Gram-Schmidt norm for the master secret key and formed a lattice family with two associated trapdoors to generate short vectors; one for all lattices in the family and the other one for all but one.

The first Ring-LWE based IBE scheme has been proposed by Ducas, Lyubashevsky and Prest [11] (DLP-IBE), which is still considered the most efficient scheme to date due to smaller key sizes. The use of the ring variant increased efficiency by reducing the public key size and ciphertext size to $O(n)$. The security of their scheme holds in the random oracle model and is related to the NTRU hardness assumption. An efficient C implementation of the DLP-IBE scheme and a detailed performance analysis was provided in [27]. In 2017, Campbell and Grover introduced a HIBE scheme, called LATTE, which can be viewed as a combination of the DLP scheme with the delegation mechanism from [8]. An optimized implementation and refined analysis of LATTE, has recently been proposed in [33].

The work from [5] constructed an IBE using the notion of gadget-based trapdoors in the ring setting, introduced by [28]. Such trapdoors can be seen as linear transformations mapping hard instances of cryptographic problems on some lattices to easy instances on a lattice defined by a public "gadget matrix".

The IBE from [5] also made use of the efficient Gaussian preimage sampling algorithms from [17] to propose an implementation of their scheme. In [4], this IBE and its associated sampling algorithms were adapted to the module setting and instantiated. The use of module lattices of dimension nd , where d is the rank module, led to a more flexible choice of parameters. In [32], the authors proposed new efficient gadget sampling algorithms which didn't need floating-point arithmetic, and as fast as the original [17] sampler.

All those constructions generally make use of dedicated trapdoors, needed by the authority to generate the secret key of a user. In that case, building the trapdoor and sampling particular short vectors are quite costly, and represent the main bottleneck in the efficiency of such schemes. More recently, [9] introduced approximate trapdoors to improve the efficiency of schemes built from lattice trapdoors while keeping the same concrete security. [16] showed that those approximate trapdoors, relying on the [28] framework, exist on a family of NTRU lattices. Our work explores the application of those approximate trapdoors, and in particular of this family of NTRU lattices, to more advanced schemes where the generation of private keys in a multi-user setting is needed.

Our contributions. Our main contribution is to provide and implement two lattice-based IBE schemes (IBE1 and IBE2), which make use of families of gadget-based approximate trapdoors. Our two constructions rely on the LWE problem over modules (Module-LWE) and the inhomogeneous NTRU problem (iNTRU) respectively. We investigate how to instantiate and parametrize the approximate trapdoor preimage sampling over these two families of trapdoors in a way to obtain provable and efficient quantum-safe IBE schemes. The IBE1 construction is an adaptation of the identity-based encryption scheme from [1, 5] to the module setting, using approximate trapdoors. The IBE2 construction follows the same blueprint as the DLP scheme except that it makes use of iNTRU gadget-based approximate trapdoors. As in previous IBE constructions, encryption is based on the Dual-Regev encryption scheme. We provide a complete public and open-source C implementation with performance benchmarking. The implementation is modular and makes it easy to change the building blocks of our algorithms according to the desired properties that we want to get.

Our work explores the use of approximate trapdoors for the construction of IBE schemes and the potential practical insights we can gain from it. We first adapted the construction of [4] using approximate trapdoors in our IBE1. The error induced by the approximate setting requires changes at several levels, either for the choice of encoding or for the sampling algorithms. As expected, we obtain better timings for all four algorithms composing our IBE by using approximate trapdoors rather than exact ones.

The second scheme IBE2 makes use of approximate trapdoors relying on a variant of NTRU rather than Module-LWE. Our approach was motivated by the fact that other efficient IBEs, such as DLP and Latte, used the NTRU hardness assumption. However, unlike us, these last two schemes used the GPV paradigm

https://github.com/lucasprabel/approx_lattice

to generate trapdoors, which significantly changes the way their schemes are constructed compared to ours.

More details on our implementation choices. Our IBE1 proof relies on a statistical trapdoor instantiation. Although the size of the parameters increases consequently, the use of approximate trapdoors allowed us to mitigate this loss in efficiency induced by the use of a statistical instantiation. In order to ensure decryption correctness, we also need to use a large modulus (see Section 4.1). This leads us to perform calculations carefully on 64-bit integers, so as not to affect our scheme efficiency. The IBE1 also makes use of a small-norm encoding, instead of the low-degree encoding used in [4] to ensure that the noise is still not too large in order to decrypt. The encoding we use sets constraints on the structure of the ring \mathcal{R}_q which is not compatible with the NTT for polynomial multiplications. Instead, we use a "partial NTT" based on [26] results, which reduces multiplication in \mathcal{R}_q to multiplication in smaller rings. We also have optimized the underlying CRT representation algorithm compared to [4]. Finally, the sampling algorithms have been adapted to the approximate setting. Table 2 shows some applicable parameter sets together with their concrete bit security using the LWE estimator from [2] with BKZ as a reduction cost model. All the algorithms comprising the IBE1 over module lattices are more efficient than their exact counterpart at the same security level. This improvement concerns in particular Setup and Extract which are optimized by a factor ≈ 1.5 . We give more details in Section 4.2.

The IBE2 scheme is instantiated using gadget-based trapdoors on a family of NTRU lattices. Table 4 provides a set of applicable parameter sets together with their concrete bit security. The computational trapdoors instantiation lowers the bounds on parameters required for the correctness compared to the IBE1 scheme, which allows the use of smaller moduli. For the IBE2 construction, an identity is encoded as $H(\text{id})$ with H having special properties so that we are able to respond to private key queries except for the target identity chosen by the adversary.

Table 1: Timings comparison in ms of the different operations of the IBE2 scheme in this paper and in Latte [33] for different sets of parameters.

Scheme	$(n, \lceil \log_2(q) \rceil)$	Security level	Setup	Extract	Encrypt	Decrypt
[33]	(1024, 24)	128	102	0.82	0.05	0.06
[33]	(2048, 25)	256	292	2.62	0.10	0.13
This paper	(1024, 25)	159	3.32	5.92	1.10	0.07
This paper	(2048, 25)	293	10.21	12.79	2.96	0.16

In Table 1, we provide timings comparison with the [33] IBE scheme. A more complete and detailed comparison is given in Section 5.2 between our IBE2 scheme and the [11] and [33] IBE schemes which also rely on the NTRU assumption. The use of approximate trapdoors and of the iNTRU assumption allows to obtain better timings for the Setup algorithm and for the Extract

algorithm for some sets of parameters. Unfortunately, we obtain an overhead for encryption which essentially comes from the Gaussian sampling phase. We can save a factor 3.5 using binomial samples as in [11] and [33] but there is still an overhead due to the extra number of samples we need. We make $n(2+m)$ calls to the integer Gaussian sampler while encryption in [11] and [33] makes use of $3n$ binomial sampling calls, where n is the dimension of the underlying polynomial ring and m is the size of one of the vector used in encryption. As an example, we obtain a timing of 0.87ms for one encryption against 0.13 for [33] for a security level of 128 bits.

Organization of the paper. The paper is structured as follows: Section 2 reviews the necessary linear algebra and lattice backgrounds. The Section also presents the notions of approximate trapdoor. Then, Section 4 and Section 5 contain a detailed description of the two IBEs, based on the Module-LWE hardness assumption and based on a variant of the NTRU assumption respectively. Both Sections provide a security analysis and detailed performance results together with a comparison of their respective analogue in terms of lattice-based building blocks and assumptions.

2 Preliminaries

Notations. Throughout the paper, the vectors are written in lowercase bold letters (e.g. \mathbf{v}) and matrices in uppercase bold letters (e.g. \mathbf{A}). We refer to their entries with a subscript index $v_i, A_{i,j}$. We denote $\|\cdot\|$ and $\|\cdot\|_\infty$ the euclidean norm and the infinity norm respectively. The norm of a vector over \mathbb{Z}_q is the norm of the corresponding vector over \mathbb{Z} obtained by choosing its entries in the set $\{-\lfloor q/2 \rfloor, \dots, \lfloor q/2 \rfloor\}$. The norm of a polynomial $a = \sum_{i=0}^{n-1} a_i X^i$ is the norm of the coefficients vector (a_0, \dots, a_{n-1}) . Finally, the norm of a matrix is the maximum norm of its column vectors. If x is sampled from a distribution D , we write $x \leftarrow D$. We denote $\mathcal{U}(S)$ the uniform distribution over a finite set S .

We say that a function f is negligible when $f(n) = o(n^{-c})$ for all $c > 0$ as $n \rightarrow \infty$. An event is said to happen with overwhelming probability if its probability of not happening is negligible. Two distributions D_0 and D_1 over the same countable domain Ω are said to be statistically indistinguishable if their statistical distance $\Delta(D_0, D_1) = \frac{1}{2} \sum_{\omega \in \Omega} |D_0(\omega) - D_1(\omega)|$ is negligible. Two distributions are said to be computationally indistinguishable if no probabilistic polynomial time algorithm can distinguish them with non-negligible advantage.

2.1 Lattices background and discrete Gaussian distributions.

Lattices. A lattice A in \mathbb{R}^m is the set $\{\sum_{i=1}^k \lambda_i \mathbf{b}_i, \lambda_i \in \mathbb{Z}\}$ of all integer linear combinations of some linearly independent vectors $\mathbf{B} = \{\mathbf{b}_1, \dots, \mathbf{b}_k\} \subset \mathbb{R}^m$. We call k the rank of the lattice and m its dimension. When $k = m$, the lattice is said to be full-rank.

Given $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and $\mathbf{u} \in \mathbb{Z}_q^n$, we define the following m -dimensional q -ary lattice $\Lambda_q^\perp(\mathbf{A}) = \{\mathbf{x} \in \mathbb{Z}^m \mid \mathbf{A}\mathbf{x} = \mathbf{0} \pmod{q}\}$, and its coset $\Lambda_q^{\mathbf{u}}(\mathbf{A}) = \{\mathbf{x} \in \mathbb{Z}^m \mid \mathbf{A}\mathbf{x} = \mathbf{u} \pmod{q}\}$.

Module lattices are particular lattices that have a polynomial structure. We denote d the module rank of those lattices. When $d = 1$, module lattices are in fact ideal lattices. We consider the ones that are based on the rings $\mathcal{R} = \mathbb{Z}[X]/\langle X^n + 1 \rangle$ and $\mathcal{R}_q = \mathbb{Z}_q[X]/\langle X^n + 1 \rangle$, where n is a power of two and q is prime. They are sublattices of the full lattice \mathcal{R}^m , which is isomorphic to the integer lattice \mathbb{Z}^{nm} .

Gaussian distributions. We recall that a symmetric matrix $\mathbf{M} \in \mathbb{R}^{n \times n}$ is positive definite (resp. positive semidefinite) if $\mathbf{x}^T \mathbf{M} \mathbf{x} > 0$ (resp. $\mathbf{x}^T \mathbf{M} \mathbf{x} \geq 0$) for all nonzero $\mathbf{x} \in \mathbb{R}^n$. In this case we write $\mathbf{M} \succ 0$ (resp. $\mathbf{M} \succeq 0$). We say that $\mathbf{M} \succeq \mathbf{N}$ when $\mathbf{M} - \mathbf{N} \succeq 0$, and write $\mathbf{M} \succeq \eta$ instead of $\mathbf{M} \succeq \eta \mathbf{I}_n$ when $\eta \geq 0$ is a real. The spherical continuous Gaussian function of center $\mathbf{c} \in \mathbb{R}^n$ and parameter σ is defined on \mathbb{R}^n by $\rho_{\mathbf{c}, \sigma}(\mathbf{x}) = \exp(-\frac{\pi \|\mathbf{x} - \mathbf{c}\|^2}{\sigma^2})$. For a positive definite matrix $\mathbf{\Sigma} \in \mathbb{R}^{n \times n}$, we also define the (skewed) Gaussian function $\rho_{\mathbf{c}, \sqrt{\mathbf{\Sigma}}}(\mathbf{x}) = \exp(-\pi(\mathbf{x} - \mathbf{c})^T \mathbf{\Sigma}^{-1}(\mathbf{x} - \mathbf{c}))$. Then, for a full-rank lattice $\Lambda \subset \mathbb{Z}^n$, we denote $D_{\Lambda, \sigma, \mathbf{c}}$ (respectively $D_{\Lambda, \sqrt{\mathbf{\Sigma}}, \mathbf{c}}$) the spherical (resp. ellipsoid) discrete Gaussian distribution of center $\mathbf{c} \in \mathbb{R}^n$ and parameter $\sigma > 0$, associated with the density $\rho_{\mathbf{c}, \sigma}$ (resp. $\rho_{\mathbf{c}, \sqrt{\mathbf{\Sigma}}}$).

Smoothing parameter. The smoothing parameter of a lattice Λ , denoted $\eta_\varepsilon(\Lambda)$, was first introduced in [29]. Lemma 1 gives an upper bound on it.

Lemma 1 ([18, Lemma 3.1]). *Let $\Lambda \subset \mathbb{R}^n$ be a lattice with basis \mathbf{B} , and $\tilde{\mathbf{B}}$ the Gram-Schmidt orthogonalization of \mathbf{B} . Then, for any $\varepsilon > 0$, we have $\eta_\varepsilon(\Lambda) \leq \|\tilde{\mathbf{B}}\| \cdot \sqrt{\ln(2n(1 + 1/\varepsilon))}/\pi$.*

Gaussian tailcut. We will use the following result to bound the euclidean norm of vectors that follow a discrete Gaussian distribution.

Lemma 2 ([24, Lemma 4.4]). *For any integer $n \geq 1$ and reals $\sigma > 0$, $t > 1$, $\Pr[\|\mathbf{x}\| > t\sigma\sqrt{n} \mid \mathbf{x} \leftarrow D_{\mathbb{Z}^n, \sigma}] < t^n e^{\frac{n}{2}(1-t^2)}$.*

We call t a tailcut of the discrete Gaussian of parameter σ when a vector \mathbf{x} sampled from $D_{\mathbb{Z}^n, \sigma}$ verifies the inequality $\|\mathbf{x}\| \leq t\sigma\sqrt{n}$ on its euclidean norm with overwhelming probability.

2.2 Cryptographic problems on lattices

We work on the rings $\mathcal{R} = \mathbb{Z}[X]/\langle X^n + 1 \rangle$ and $\mathcal{R}_q = \mathbb{Z}_q[X]/\langle X^n + 1 \rangle$, where n is a power of two and q a prime modulus. Let us now define the hard problems on which our constructions rely. In most practical lattice-based schemes, security is based on structured variants of LWE rather than on LWE itself. In the Module-SIS and Module-LWE variants, the parameter d is the rank of the module, and nd

is the dimension of the corresponding module lattice. The hardness of those problems is proven by worst-case to average-case reductions from hard problems on module lattices (see [22]).

Definition 1 (Module-SIS $_{n,d,m,q,\beta}$). *Given a uniform $\mathbf{A} \in \mathcal{R}_q^{d \times m}$, find a vector $\mathbf{x} \in \mathcal{R}^m$ such that $\mathbf{A}\mathbf{x} = \mathbf{0} \pmod{q}$, and $0 < \|\mathbf{x}\| \leq \beta$.*

Definition 2 (Decision Module-LWE $_{n,d,q,\sigma}$). *Given a uniform $\mathbf{A} \in \mathcal{R}_q^{d \times m}$ and the vector $\mathbf{b}^T = \mathbf{s}^T \mathbf{A} + \mathbf{e}^T \pmod{q} \in \mathcal{R}_q^m$, where $\mathbf{s} \leftarrow \mathcal{U}(\mathcal{R}_q^d)$ and $\mathbf{e} \leftarrow D_{\mathcal{R}^m, \sigma}$, distinguish the distribution of (\mathbf{A}, \mathbf{b}) from the uniform distribution over $\mathcal{R}_q^{d \times m} \times \mathcal{R}_q^m$.*

As stated in [7], the Module-LWE problem remains hard when the short secret \mathbf{s} is sampled from the Gaussian distribution $D_{\mathcal{R}^d, \sigma}$.

Note that Module-LWE $_{n,d,q,\sigma}$ with rank $d = 1$ corresponds to the ring version of LWE, which will be denoted Ring-LWE $_{n,q,\sigma}$. We define $\mathbf{G} = \mathbf{I}_d \otimes \mathbf{g}^T$ and its associated lattice named the module \mathbf{G} -lattice, for which the Module-SIS problem is easy.

To build our second IBE, we use gadget-based trapdoors whose pseudorandomness is based on the inhomogeneous NTRU problem (iNTRU), a variant of NTRU introduced in [15].

Definition 3 (iNTRU $_{q,\chi}$). *Let k, q be integers and χ a distribution over \mathcal{R} . The input of the iNTRU $_{q,\chi}$ problem is a vector $\mathbf{a} \in \mathcal{R}_q^k$ which is either taken uniform in \mathcal{R}_q^k or either set as $\mathbf{a} = r^{-1}(\mathbf{g} + \mathbf{e})$ where (r, \mathbf{e}) is drawn from χ^{k+1} . The goal is to decide which is the case.*

In their paper, the authors from [15] showed a reduction of a matrix-variant of iNTRU called MiNTRU from the non-standard Ring-LWE with a trapdoor oracle access problem (more details are given in [15, Section 4.3]). The reduction can be adapted to the iNTRU problem as well.

2.3 Encoding identities with full-rank differences

We use the notion of full-rank differences encoding (FRD) in our first scheme, with different properties that the ones used in [1].

Definition 4 ([4]). *An encoding with full-rank differences from the set \mathcal{M} to a ring \mathcal{R} is a map $H : \mathcal{M} \rightarrow \mathcal{R}$ such that:*

- for any $m \in \mathcal{M}$, $H(m)$ is invertible,
- for any $m_1, m_2 \in \mathcal{M}$ such that $m_1 \neq m_2$, $H(m_1) - H(m_2)$ is invertible,
- H is computable in polynomial time.

As shown in [4], we construct an FRD encoding in the module setting (i.e. over $\mathcal{R}_q^{d \times d}$), by first constructing one in the ring setting (i.e. over \mathcal{R}_q). The encodings used in our IBE scheme impose a structure on the ring \mathcal{R}_q which is not compatible with the Number Theoretic Transform (NTT). To speed up polynomial multiplications and to mitigate the loss of performance due to the inability of using the NTT, we use ideas from [26]. This method can be thought of as a "partial NTT". It is based on the following result.

Lemma 3 ([26, Corollary 1.2]). *Let $n \geq r > 1$ be powers of 2, and q a prime such that $q \equiv 2r + 1 \pmod{4r}$. Then the cyclotomic polynomial $X^n + 1$ factors in $\mathbb{Z}_q[X]$ as $X^n + 1 = \prod_{i=1}^r (X^{n/r} - s_i)$, for some distinct $s_i \in \mathbb{Z}_q^*$ such that the $(X^{n/r} - s_i)$ are irreducible in $\mathbb{Z}_q[X]$. Moreover, if $y \in \mathbb{Z}_q[X]/(X^n + 1)$ satisfies $0 < \|y\|_\infty < \frac{q^{1/r}}{\sqrt{r}}$, then y has an inverse in $\mathbb{Z}_q[X]/(X^n + 1)$.*

This lemma was used in [4] to build a "low-degree" FRD (the identities were encoded as polynomials of low-degree). In our case, we construct a "small-norm" FRD, encoding polynomials with ℓ_∞ -norm smaller than $\frac{q^{1/r}}{2\sqrt{r}}$. We make use of the "small-norm" FRD described in Proposition 1 rather than "low-degree" FRD encoding so that the correctness of the decryption algorithm still holds when using approximate trapdoors, because they introduce an additional error term that must be bounded.

Proposition 1. *Let $n \geq r > 1$ be powers of 2, q a prime such that $q \equiv 2r + 1 \pmod{4r}$ and $1 \leq D \leq \frac{q^{1/r}}{2\sqrt{r}}$ an integer. We define $\mathcal{M} = \{-D, \dots, D\}^n \setminus \{(0, \dots, 0)\}$ the set of identities. Then the following map $H_M : \mathcal{M} \rightarrow \mathcal{R}_q$, such that $H_M(m_0, \dots, m_{n-1}) = \sum_{i=0}^{n-1} m_i X^i$ is an FRD encoding.*

Proof. We have $\|H_M(m)\|_\infty \leq D < \frac{q^{1/r}}{2\sqrt{r}}$ for all m because of the choice of \mathcal{M} . So according to Lemma 3, $H_M(m)$ is invertible. For all $m_1, m_2 \in \mathcal{M}$, we also have $\|H_M(m_1) - H_M(m_2)\|_\infty \leq 2D < \frac{q^{1/r}}{\sqrt{r}}$ so $H_M(m_1) - H_M(m_2)$ is invertible. Finally, H_M is an FRD encoding. \square

FRD on modules. As explained in [4], FRD encoding in the module setting can be built using an existing FRD encoding in the ring setting $H_M : \mathcal{M} \rightarrow \mathcal{R}_q$ by constructing $H_M(m) \cdot \mathbf{I}_d \in R_q^{d \times d}$ for $m \in \mathcal{M}$ and $\mathbf{I}_d \in \mathcal{R}_q^{d \times d}$ is the identity matrix.

3 Trapdoors on lattices

The two IBE constructions make use of efficient trapdoors, called gadget-trapdoors, introduced in [28]. Such trapdoors were generalized to ideal lattices in [21] and to module lattices in [4]. An efficient instantiation of the sampling algorithms was given in [17]. All those results allow us to efficiently instantiate trapdoor generation and sampling algorithms. Having introduced the \mathbf{G} -lattice in the previous Section, we can now define the associated \mathbf{G} -trapdoors.

Definition 5. *A \mathbf{G} -trapdoor for a matrix $\mathbf{A} \in \mathcal{R}_q^{d \times m}$ is a matrix $\mathbf{R} \in \mathcal{R}^{(m-dk) \times dk}$ such that*

$$\mathbf{A} \begin{bmatrix} \mathbf{R} \\ \mathbf{I}_{dk} \end{bmatrix} = \mathbf{H}\mathbf{G}$$

for some invertible $\mathbf{H} \in \mathcal{R}_q^{d \times d}$, called the tag of \mathbf{A} .

The knowledge of the trapdoor \mathbf{R} allows one to compute small vectors on any coset of a given lattice, by sampling on the \mathbf{G} -lattice first, and then by using \mathbf{R} to map the sample to a small vector in the given lattice.

3.1 Approximate trapdoors for module lattices.

As shown in [9], the gadget trapdoor proposed by Micciancio and Peikert can be modified to an approximate trapdoor, in a way that further reduces the sizes of the public matrix, the trapdoor and the preimage.

Definition 6 (AISIS $_{n,d,m,q,\alpha,\beta}$). *For any $n, d, m, q \in \mathbb{N}$ and $\alpha, \beta \in \mathbb{R}$, the approximate inhomogeneous short integer solution problem AISIS $_{n,d,m,q,\alpha,\beta}$ is defined as follows: given $\mathbf{A} \in \mathcal{R}_q^{d \times m}$, $\mathbf{y} \in \mathcal{R}_q^d$, find a vector $\mathbf{x} \in \mathcal{R}^m$ such that $\|\mathbf{x}\| \leq \beta$ and such that there is a vector $\mathbf{z} \in \mathcal{R}^d$ satisfying: $\|\mathbf{z}\| \leq \alpha$ and $\mathbf{Ax} = \mathbf{y} + \mathbf{z} \bmod q$.*

Definition 7 (Approximate Trapdoor). *A string τ is called an (α, β) -approximate trapdoor for a matrix $\mathbf{A} \in \mathcal{R}_q^{d \times m}$ if there is a probabilistic polynomial time algorithm that given τ, \mathbf{A} and any $\mathbf{y} \in \mathcal{R}_q^d$, outputs a non-zero vector $\mathbf{x} \in \mathcal{R}^m$ such that $\|\mathbf{x}\| \leq \beta$ and there is a vector $\mathbf{z} \in \mathcal{R}^d$ satisfying $\|\mathbf{z}\| \leq \alpha$ and $\mathbf{Ax} = \mathbf{y} + \mathbf{z} \bmod q$.*

In practice, we generate approximate trapdoors by dropping ℓ entries corresponding to the small powers of b from the gadget matrix \mathbf{G} to get the following matrix: $\mathbf{F} = \mathbf{I}_d \otimes \mathbf{f}^T = \mathbf{I}_d \otimes [b^\ell \ b^{\ell+1} \ b^{\ell+2} \ \dots \ b^{k-1}] \in \mathcal{R}^{d \times d(k-\ell)}$.

3.2 Module-LWE Approximate trapdoors.

In [9], the authors described the approximate trapdoor generation and the approximate preimage sampling algorithms, adapted from [28] by making use of the approximate gadget matrix \mathbf{F} instead of \mathbf{G} . Its trapdoor generation and trapdoor preimage sampling algorithms over modules are recalled in Figure 1 and the result of the approximate sampling is stated in Theorem 1.

Theorem 1 ([9, Theorem 4.1]). *There exists probabilistic, polynomial time algorithms $\text{ApproxTrapGen1}(\cdot)$ and $\text{ApproxSamplePre1}(\cdot)$ such that:*

1. $\text{ApproxTrapGen1}(\mathbf{H}, \sigma)$ takes as input public parameters, a tag matrix $\mathbf{H} \in \mathcal{R}^{d \times d}$ and parameter $\sigma > 0$ and returns a matrix-approximate trapdoor pair $(\mathbf{A}, \mathbf{R}) \in \mathcal{R}_q^{d \times m} \times \mathcal{R}^{\tilde{m} \times \omega}$ where \mathbf{R} coefficients are drawn from a Gaussian distribution of parameter σ over \mathcal{R} .
2. Let \mathbf{A} be generated with an approximate trapdoor as above. The following two distributions are statistically indistinguishable:

$$\{(\mathbf{A}, \mathbf{x}, \mathbf{u}, \mathbf{y}) \mid \mathbf{u} \leftarrow \mathcal{U}(\mathcal{R}_q^d), \mathbf{x} \leftarrow \text{ApproxSamplePre1}(\mathbf{A}, \mathbf{R}, \mathbf{0}, \mathbf{u}, \zeta), \mathbf{y} = \mathbf{u} - \mathbf{Ax} \bmod q\}$$

$$\text{and } \{(\mathbf{A}, \mathbf{x}, \mathbf{u}, \mathbf{y}) \mid \mathbf{x} \leftarrow D_{\mathcal{R}^m, \zeta}, \mathbf{y} \leftarrow D_{\mathcal{R}^d, \sigma \sqrt{(b^{2\ell}-1)/(b^2-1)}}, \mathbf{u} = \mathbf{Ax} + \mathbf{y} \bmod q\}$$

for any $\sigma \geq \sqrt{b^2 + 1} \cdot \omega(\sqrt{\log d})$ and $\zeta \gtrsim \sqrt{b^2 + 1} \frac{s_1^2(\mathbf{R})}{s_{2d}(\mathbf{R})} \eta_\epsilon(\mathbb{Z}^{dk})$.

3.3 iNTRU Approximate trapdoors.

In [16], Genise and Li introduced a family of Ring-SIS approximate trapdoors whose pseudorandomness is based on the iNTRU problem and then showed that the efficient gadget-based trapdoor framework of [28] exists on a family of NTRU lattices. Their trapdoor scheme enjoys small secret keys and is compatible with applications requiring tag matrices.

In their second trapdoor scheme, the matrix $\begin{bmatrix} -\mathbf{e}^T \\ r\mathbf{I} \end{bmatrix}$ is used as a \mathbf{f} -trapdoor for $[1 \ \mathbf{a}] = [1 \ r^{-1}(\mathbf{f}^T + \mathbf{e}^T)]$ where $(r, \mathbf{e}) \leftarrow \chi^{k-\ell+1}$ is drawn from a distribution with small entries and \mathbf{f} is the approximate gadget vector. They got the following results, similar to the ones from [9].

<p><u>ApproxTrapGen1(\mathbf{H}, σ)</u></p> <ol style="list-style-type: none"> 1. $\bar{\mathbf{A}} \leftarrow \mathcal{U}(\mathcal{R}_q^{d \times \bar{m}})$ 2. $\mathbf{R} \leftarrow D_{\mathcal{R}^{\bar{m} \times \omega}, \sigma}$ 3. set $\mathbf{A} = [\bar{\mathbf{A}} \mid \mathbf{H}\mathbf{F} - \bar{\mathbf{A}}\mathbf{R}] \in \mathcal{R}_q^{d \times m}$ 4. return (\mathbf{A}, \mathbf{R}) <hr/> <p><u>ApproxSamplePre1($\mathbf{A}, \mathbf{R}, \mathbf{H}, \mathbf{u}, \zeta$)</u></p> <ol style="list-style-type: none"> 1. sample perturbation $\mathbf{p} \leftarrow D_{\mathcal{R}^m, \sqrt{\Sigma_p}}$ 2. set coset $\mathbf{v} = \mathbf{H}^{-1}(\mathbf{u} - \mathbf{A}\mathbf{p})$ 3. sample $\mathbf{z} = (\mathbf{z}_1^T, \mathbf{z}_2^T)^T \leftarrow D_{\Lambda_q^y(\mathbf{G}), \sigma_g}$ 4. set $\mathbf{x} = \mathbf{p} + \begin{bmatrix} \mathbf{R} \\ \mathbf{I} \end{bmatrix} \mathbf{z}_2 \in \mathcal{R}^m$ 5. return \mathbf{x} 	<p><u>ApproxTrapGen2(σ)</u></p> <ol style="list-style-type: none"> 1. sample $r \leftarrow D_{\mathcal{R}, \sigma}$, $\mathbf{e} \leftarrow D_{\mathcal{R}^m, \sigma}$ 2. set $\mathbf{a}' = r^{-1}(\mathbf{f} + \mathbf{e}) \in \mathcal{R}^m$ 3. set $\mathbf{a} = (1, \mathbf{a}') \in \mathcal{R}_q^{m+1}$ 4. set $\mathbf{R} = \begin{bmatrix} -\mathbf{e}^T \\ r\mathbf{I}_m \end{bmatrix} \in \mathcal{R}^{(m+1) \times m}$ 5. return (\mathbf{a}, \mathbf{R}) <hr/> <p><u>ApproxSamplePre2($\mathbf{a}, \mathbf{R}, \mathbf{u}, \zeta$)</u></p> <ol style="list-style-type: none"> 1. sample perturbation $\mathbf{p} \leftarrow D_{\mathcal{R}^{m+1}, \sqrt{\Sigma_p}}$ 2. set coset $v = (u - \mathbf{a}^T \mathbf{p}) \in \mathcal{R}_q$ 3. sample $\mathbf{z} = (\mathbf{z}_1^T, \mathbf{z}_2^T)^T \leftarrow D_{\Lambda_q^y(\mathbf{g}^T), \sigma_g}$ 4. set $\mathbf{x} = \mathbf{p} + \mathbf{R}\mathbf{z}_2 \in \mathcal{R}^{m+1}$ 5. return \mathbf{x}
--	---

Fig. 1: Description of the approximate trapdoor generation and preimage sampling algorithms based on [9] (on the left-hand side) and [16] (on the right-hand side) instantiations, where the perturbation is sampled with parameter $\Sigma_p = \zeta^2 \mathbf{I}_m - \sigma_g^2 \begin{bmatrix} \mathbf{R} \\ \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{R}^T & \mathbf{I} \end{bmatrix}$ on the left side and $\Sigma_p = \zeta^2 \mathbf{I}_{m+1} - \sigma_g^2 \mathbf{R}\mathbf{R}^T$ on the right side. Both are independent of the target \mathbf{u} . We refer to [9, 16] for more information on the iNTRU based trapdoor generation.

Theorem 2. *Let $r \leftarrow \chi$ and $\mathbf{e}^T \leftarrow \chi^m$ and set the trapdoor function description as $\mathbf{a} = [1 \ \mathbf{a}'] = [1 \ r^{-1}(\mathbf{f} + \mathbf{e})] \in \mathcal{R}_q^{m+1}$. Let $\eta = \eta_\epsilon(\mathbb{Z}^{n \times m})$ and $\sigma_g = \eta_\epsilon(\Lambda_q^\perp(\mathbf{g}^T)) \geq \sqrt{b^2 + 1} \cdot \eta_\epsilon(\mathbb{Z}^{n \times m})$ for some $\epsilon \in (0, 1)$ and $\zeta \geq \sqrt{\sigma_g^2 \mathbf{R}\mathbf{R}^T + \eta^2 \mathbf{I}_{m+1}}$. Then, the following distributions are within a max-log distance $3 \log \frac{1+\epsilon}{1-\epsilon} \leq \frac{6\epsilon}{1-\epsilon}$:*

$\{(\mathbf{a}, \mathbf{x}, u, y) \mid u \leftarrow \mathcal{U}(\mathcal{R}_q), \mathbf{x} \leftarrow \text{ApproxSamplePre2}(\mathbf{a}, \mathbf{R}, u, \zeta), y = u - \mathbf{a}^T \mathbf{x} \in \mathcal{R}_q\}$
and $\{(\mathbf{a}, \mathbf{x}, u, y) \mid \mathbf{x} \leftarrow D_{\mathcal{R}^{m+1}, \zeta}, y \leftarrow D_{\mathcal{R}, \sigma_e} \pmod q, u = \mathbf{a}^T \mathbf{x} + y \in \mathcal{R}_q\}$

for $\sigma_e = \sigma_g \sqrt{(b^{2\ell} - 1)/(b^2 - 1)}$.

4 An Approximate-IBE scheme on modules

The first IBE scheme we present is the IBE scheme from [1, 5] adapted to the module setting, instantiated using the approximate trapdoors from [9]. At a high level, the scheme will make use of the following blocks:

- The master secret key is an \mathbf{F} -approximate trapdoor $\mathbf{R} \in \mathcal{R}^{\bar{m} \times \omega}$ associated with \mathbf{A} with tag $\mathbf{0}$, with subgaussian coefficients of parameter σ , with $\omega = d(k - \ell)$. The master public key is a tuple consisting of a uniformly random vector $\mathbf{u} \in \mathcal{R}_q^d$ and the matrix $\mathbf{A} \in \mathcal{R}_q^{d \times m}$, with $m = \bar{m} + \omega$ chosen as: $\mathbf{A} = [\bar{\mathbf{A}} \mid -\bar{\mathbf{A}}\mathbf{R}]$. Taking $\bar{m} = d \log q$, we get that $\bar{\mathbf{A}}$ is full rank with high probability according to [7, Lemma 2.6]. Moreover, by taking $\sigma > 4 \cdot 4^{\frac{1}{nd}} \sqrt{n}$, we obtain that \mathbf{A} is statistically close to uniform by Corollary 7.5 from [25].
- A "small-norm" FRD encoding H_M as described in Section 2.3. This allows anyone, with the knowledge of the master public key \mathbf{A} , to compute a public matrix $\mathbf{A}_{\text{id}} = \mathbf{A} + \begin{bmatrix} \mathbf{0}_{d, \bar{m}} & \mathbf{H}_{\text{id}}\mathbf{F} \end{bmatrix}$ associated with the identity id of a user. Then, the secret key for id is an approximate short vector $\mathbf{x}_{\text{id}} \in \mathcal{R}^m$ obtained by using the ApproxSamplePre1 algorithm. Such a vector satisfies the relation $\mathbf{A}_{\text{id}}\mathbf{x}_{\text{id}} \approx \mathbf{u} \pmod{q}$. We can bound the approximate error in this relation by using Theorem 1 and the fact that we use a small-norm FRD encoding.
- Finally, we use the Dual-Regev encryption scheme for the encryption and decryption algorithms, taking care of the additional error which appears when using approximate trapdoors.

4.1 Construction

We detail the first IBE where users' keys are defined based on the approximate preimage sampling from Theorem 1 and encryption is based on the Dual-Regev scheme.

- Setup(1^n) \longrightarrow (mpk, msk):
 - $(\mathbf{A}, \mathbf{R}) \leftarrow \text{ApproxTrapGen1}(\mathbf{0}, \sigma) \in \mathcal{R}_q^{d \times m} \times \mathcal{R}^{\bar{m} \times \omega}$, $\mathbf{u} \leftarrow \mathcal{U}(\mathcal{R}_q^d)$;
 - mpk = (\mathbf{A}, \mathbf{u}) , msk = \mathbf{R} .
- Extract(mpk, msk, id) \longrightarrow sk_{id} $\mathbf{x} \in \mathcal{R}^m$:
 - $\mathbf{H}_{\text{id}} \leftarrow H_M(\text{id})$; $\mathbf{A}_{\text{id}} \leftarrow \mathbf{A} + \begin{bmatrix} \mathbf{0}_{d, \bar{m}} & \mathbf{H}_{\text{id}}\mathbf{F} \end{bmatrix} \in \mathcal{R}_q^{d \times m}$;
 - $\mathbf{x} \leftarrow \text{ApproxSamplePre1}(\mathbf{A}_{\text{id}}, \mathbf{R}, \mathbf{H}_{\text{id}}, \mathbf{u}, \zeta)$;
- Encrypt(mpk, id, M) \longrightarrow $C = (\mathbf{b}, c) \in \mathcal{R}_q^{m+1}$:
 - $\mathbf{H}_{\text{id}} \leftarrow H_M(\text{id})$; $\mathbf{A}_{\text{id}} \leftarrow \mathbf{A} + \begin{bmatrix} \mathbf{0}_{d, \bar{m}} & \mathbf{H}_{\text{id}}\mathbf{F} \end{bmatrix} \in \mathcal{R}_q^{d \times m}$;
 - $\mathbf{s} \leftarrow D_{\mathcal{R}_q^d, \tau}$, $\mathbf{e}_0 \leftarrow D_{\mathcal{R}^{m-\omega}, \tau}$, $\mathbf{e}_1 \leftarrow D_{\mathcal{R}^{\omega}, \gamma}$, and $e' \leftarrow D_{\mathcal{R}, \tau}$;
 - $\mathbf{b} \leftarrow (\mathbf{s}^T \mathbf{A}_{\text{id}})^T + (\mathbf{e}_0^T \mid \mathbf{e}_1^T)^T$ and $c \leftarrow \mathbf{s}^T \mathbf{u} + e' + \lfloor q/2 \rfloor M$;
- Decrypt(sk_{id}, C) \rightarrow M :

- set $\mathbf{x} = \text{sk}_{\text{id}}$ and compute $\text{res} \leftarrow c - \mathbf{b}^T \mathbf{x}$ which has integer coefficients;
- for each i , if the coefficient $\text{res}_i \in \mathbb{Z}$ is closer to $\lfloor q/2 \rfloor$ than to 0, then $M_i = 1$, otherwise $M_i = 0$.

Correctness. To use approximate trapdoors with the Dual-Regev approach, we need to sample the LWE secret term with a small norm instead of sampling from the uniform distribution, in order to maintain the correctness of the schemes. Let's write $\mathbf{y} \in \mathcal{R}_q^d$ the additional error we get by using approximate trapdoors instead of exact ones. The correctness of the decryption holds if the error term $\|e' - (\mathbf{e}_0^T \mid \mathbf{e}_1^T) \mathbf{x} - \mathbf{s}^T \mathbf{y}\|$ is small enough, i.e. less than $\lfloor q/4 \rfloor$.

$$\begin{aligned}
\text{res} &= c - \mathbf{b}^T \mathbf{x} \\
&= \mathbf{u} \cdot \mathbf{s} + e' + \lfloor q/2 \rfloor M - [(\mathbf{s}^T \mathbf{A}_{\text{id}})^T + (\mathbf{e}_0^T \mid \mathbf{e}_1^T)^T]^T \mathbf{x} \\
&= \mathbf{u} \cdot \mathbf{s} + e' + \lfloor q/2 \rfloor M - \mathbf{s}^T \mathbf{A}_{\text{id}} \mathbf{x} - (\mathbf{e}_0^T \mid \mathbf{e}_1^T) \mathbf{x} \\
&= \mathbf{u} \cdot \mathbf{s} + e' + \lfloor q/2 \rfloor M - \mathbf{s}^T (\mathbf{u} + \mathbf{y}) - (\mathbf{e}_0^T \mid \mathbf{e}_1^T) \mathbf{x} \\
&= e' - \underbrace{(\mathbf{e}_0^T \mid \mathbf{e}_1^T) \mathbf{x} - \mathbf{s}^T \mathbf{y}}_{\text{error term}} + \lfloor q/2 \rfloor M \in \mathcal{R}.
\end{aligned}$$

So we need to choose our parameters properly for the correctness of the Dual-Regev encryption to hold. We can bound as follows the euclidean norms of the quantities that appear in the error term:

- $\|e'\| \leq t\tau\sqrt{n}$ from Lemma 2.
- $\|\mathbf{e}_0^T \mathbf{x}_0\| \leq 2t^2\tau\zeta nd$ from Lemma 2 and Theorem 1.
- $\|\mathbf{e}_1^T \mathbf{x}_1\| \leq t^2\gamma\zeta nd(k - \ell)$ from Lemma 2 and Theorem 1.
- $\|\mathbf{s}^T \mathbf{y}\| \leq t^2\tau n^{5/2} d \sigma_g \frac{q^{1/r}}{\sqrt{r}} \sqrt{(b^{2\ell} - 1)/(b^2 - 1)}$ from Lemma 2 and Theorem 1.

By substituting these bounds, we get the following constraints:

$$\begin{aligned}
\|e' - (\mathbf{e}_0^T \mid \mathbf{e}_1^T) \mathbf{x} - \mathbf{s}^T \mathbf{y}\| &\leq \|e'\| + \|(\mathbf{e}_0^T \mid \mathbf{e}_1^T)\| + \|\mathbf{s}^T \mathbf{y}\| \\
&\leq t\tau\sqrt{n} + t^2 nd \left[\zeta (2\tau + \gamma(k - \ell)) + \tau n^{3/2} \sigma_g \frac{q^{1/r}}{\sqrt{r}} \sqrt{(b^{2\ell} - 1)/(b^2 - 1)} \right] \\
&\leq \lfloor q/4 \rfloor.
\end{aligned}$$

Parameter constraints. The following constraints, combined with the norms bounds above, must be met to ensure correctness:

- The Gaussian parameter σ_g used for the \mathbf{G} -sampling in the ApproxSamplePre1 algorithm must verify $\sigma_g \geq \sqrt{2b} \cdot (2b + 1) \cdot \sqrt{\log(2nw(1 + 1/\epsilon))/\pi}$ (see [17], Corollary 3.1).
- The Gaussian width for preimage sampling ζ must follow the condition $\zeta > \sqrt{(\sigma_g^2 + 1)s_1^2(\mathbf{R}) + \eta_\varepsilon^2(\mathbb{Z}^{nm})}$, knowing that $s_1(\mathbf{R}) < 1.1\sigma(\sqrt{2nd} + \sqrt{nw} + 4.7)$ with high probability (see [4], Section A.3), where $s_1(\mathbf{R})$ is the spectral norm of the trapdoor \mathbf{R} .

- The Gaussian width for approximate trapdoor generation σ must verify $\sigma > 4 \cdot 4^{\frac{1}{\pi d}} \sqrt{n}$ to ensure the public matrix \mathbf{A} we use is statistically close to uniform (see [25], Corollary 7.5).
- We choose to set the Gaussian parameter γ of the Gaussian error $\mathbf{e}_1 \in \mathcal{R}^w$ as $\gamma^2 = \sigma^2 \|\mathbf{e}_0\|^2 + 2nt^2\sigma^2\tau^2$.

The proof of the following Theorem 3 is standard and can be found in the Supplementary materials, Appendix B.

Theorem 3. *The IBE construction with parameters $n, d, m, q, k, \ell, \sigma, \alpha, \zeta, \tau$ and γ , chosen as in the above description, is IND-sID-CPA secure in the standard model under the hardness of Module-LWE $_{n,d,q,\tau}$.*

4.2 Implementation and performance.

The scheme is implemented in C, using [4] libraries, inheriting its modularity. It relies on several basic blocks that can be swapped out: the arithmetic over \mathbb{Z}_q and \mathcal{R}_q , a pseudorandom number generator, and a (constant-time) sampler of discrete Gaussian distributions over \mathbb{Z} . To generate our specific discrete Gaussian distributions, we make use of the following building blocks: an AES-based pseudorandom number generator (implemented using AES-NI instructions for x86 architectures), and a sampler of discrete Gaussians over \mathbb{Z} similar to Karney’s sampler [19]. We chose this sampler as it can generate samples in constant time, independently of the center, Gaussian parameter, and output value. All the computations that deal with non-integers are carried out with floating-point operations that do not involve subnormal numbers. We rely on results from [26, Lemma 3] to reduce multiplications in \mathcal{R}_q to polynomial multiplications in rings of the form $\mathbb{Z}_q[X]/\langle X^n+1 \rangle$. The CRT reduction we used then allowed us to speed up polynomial arithmetic in \mathcal{R}_q .

To assess the security of each parameter set, we estimate the pseudorandomness of the public key (corresponding to the LWE security) and the hardness of breaking AISIS. The estimation of the LWE security is done with the LWE estimator of [2] with BKZ as the reduction model. We approximate our instances by an instance of an unstructured LWE problem in dimension nd . We follow a very pessimistic core-SVP hardness, where the cost of a BKZ algorithm with blocksize κ is taken to be the cost of only one call to an SVP oracle in dimension κ . For the AISIS problem, we follow the approach of [9] which consists in computing the smallest blocksize achieving the target root Hermite factor corresponding to forging a signature.

Timings. As expected, the 4 algorithms are more efficient for low value of d . Concerning the Decrypt algorithm, its execution time relies mostly on the value of n rather than d . Our timings have been obtained on an Intel i7-8650U CPU running at 1.9 GHz, and then scaled at 4.2GHz to compare ourselves with other schemes. We provide concrete parameters sets and the associated concrete results in Table 2.

$\lceil \log_2 q \rceil$	(n, d)	ℓ	ζ	Setup	Extract	Encrypt	Decrypt	M – $\text{LWE}_{n,d,q,\tau}$
58	(256, 4)	0	1137729	203.52	56.88	17.69	2.72	81
58	(256, 4)	8	1068989	174.94	49.82	15.00	2.40	80
58	(256, 4)	15	1004226	152.70	41.57	12.79	2.19	78
60	(512, 3)	0	1398812	210.85	67.28	19.20	4.06	110
60	(512, 3)	8	1315427	189.35	59.33	17.13	3.80	109
60	(512, 3)	15	1236981	160.33	53.39	14.00	3.22	107

Table 2: Proposed IBE parameters for our first construction of Section 4 with different pairs of polynomial ring dimension n and rank d , for different moduli sizes and taking $\sigma_g = 54.9$, $\sigma = 64.1$. σ_g is the Gaussian parameter for the \mathbf{G} -Sampling, σ is the Gaussian width of the trapdoor \mathbf{R} used in the Setup algorithm and ζ is the standard deviation for the Gaussian preimage sampling in the Extract algorithm. Timings in columns Setup-to-Encrypt are given in ms. M – $\text{LWE}_{n,d,q,\tau}$ denotes the concrete bit security of the scheme.

Scheme	(n, d)	Setup	Extract	Encrypt	Decrypt
[4] ($\ell = 0$)	(512, 2)	123.62	45.80	13.12	2.91
[4] ($\ell = 0$)	(256, 4)	234.45	67.90	17.88	2.69
This paper ($\ell = 15$)	(512, 2)	79.10	29.44	10.09	2.36
This paper ($\ell = 15$)	(256, 4)	152.70	41.57	12.79	2.19

Table 3: Comparison of timings in ms of the different operations of the IBE scheme between this paper and [4] for parameters giving an equivalent level of security.

Comparison with related work. We compare our IBE performance with the IBE from [4], which corresponds to the case $\ell = 0$, that is to say the use of exact trapdoors instead of approximate ones. We observe that for a fixed pair (n, d) , the larger ℓ is, the better timings are. Overall, the use of approximate trapdoors allows us to obtain better timings for all the algorithms comprising the IBE scheme.

5 An IBE scheme based on the hardness of iNTRU

In this Section, we introduce the IBE2 scheme, instantiated using gadget-based approximate trapdoors over iNTRU trapdoors combined with the Dual-Regev encryption scheme over modules.

We recall that $m = k - \ell$ where $k = \lceil \log_b q \rceil$ and that the approximate gadget vector \mathbf{f} is defined as $\mathbf{f}^T = [b^\ell \ b^{\ell+1} \ b^{\ell+2} \ \dots \ b^{k-1}] \in \mathcal{R}^{k-\ell}$. Here, the master public key is a vector $\mathbf{a} \in \mathcal{R}_q^{m+1}$ generated with the ApproxTrapGen2 algorithm and whose pseudorandomness is based on the iNTRU problem. The master secret key $r, \mathbf{e} \in \mathcal{R}^{m+1}$ defines an \mathbf{f} -approximate trapdoor associated with \mathbf{a} . An identity is mapped to an element in \mathcal{R}_q by the use of a hash function modeled as a random oracle in the security proof; the secret key associated with an identity id is an approximate short vector $\mathbf{x} \in \mathcal{R}^{m+1}$.

5.1 IBE construction details

We detail below the four algorithms of the IBE2:

- Setup(1^n) \longrightarrow (mpk, msk):
 - let $\mathbf{a} = [1 \ \mathbf{a}_0^T]^T \in \mathcal{R}_q^{m+1}$ and $\mathbf{R} = \begin{bmatrix} -\mathbf{e}^T \\ r\mathbf{I}_m \end{bmatrix} \in \mathcal{R}^{(m+1) \times m}$ output by ApproxTrapGen2(σ); and let $\mathcal{H} : \{0, 1\}^* \rightarrow \mathcal{R}_q$ a hash function;
 - output mpk = (\mathbf{a}, \mathcal{H}) and msk = \mathbf{R} .
- Extract(mpk, msk, id) $\rightarrow \mathbf{x}_{\text{id}} = \mathbf{x}_2 \in \mathcal{R}^m$:
 - define the tag $h_{\text{id}} = \mathcal{H}(\text{id}) \in \mathcal{R}_q$;
 - sample a short preimage $\mathbf{x} = (x_1, \mathbf{x}_2^T)^T \leftarrow \text{ApproxSamplePre2}(\mathbf{a}, \mathbf{R}, h_{\text{id}}, \zeta)$;
- Encrypt(mpk, id, M) $\rightarrow \mathbf{C} = (\mathbf{b}, c) \in \mathcal{R}_q^{m+1}$:
 - compute $h_{\text{id}} = \mathcal{H}(\text{id})$; sample $s \leftarrow D_{\mathcal{R}, \tau}$, $\mathbf{e}_1 \leftarrow D_{\mathcal{R}^m, \tau}$, $e_2 \leftarrow D_{\mathcal{R}, \tau}$;
 - compute $\mathbf{b} = s\mathbf{a}_0 + \mathbf{e}_1 \in \mathcal{R}_q^m$ and $c = h_{\text{id}} \cdot s + e_2 + \lfloor q/2 \rfloor M \in \mathcal{R}_q$, where a message is encoded as $M \in \mathcal{R}_2$;
- Decrypt($\mathbf{x}_{\text{id}}, \mathbf{C}$) $\rightarrow M$:
 - parse \mathbf{x}_{id} as (x_1, \mathbf{x}_2) ; and compute $\text{res} \leftarrow c - \mathbf{b}^T \mathbf{x}_2$;
 - for each i , if the coefficient $\text{res}_i \in \mathbb{Z}$ is closer to $\lfloor q/2 \rfloor$ than to 0, $M_i = 1$, otherwise $M_i = 0$.

Correctness. We have the following equality:

$$\begin{aligned}
 c - \mathbf{b}^T \mathbf{x}_2 &= h_{\text{id}} \cdot s + e_2 + \lfloor q/2 \rfloor M - (s\mathbf{a}_0 + \mathbf{e}_1)^T \mathbf{x}_2 \\
 &= s \cdot x_1 + s\mathbf{a}_0^T \mathbf{x}_2 - y \cdot s + e_2 + \lfloor q/2 \rfloor M - s\mathbf{a}_0^T \mathbf{x}_2 - \mathbf{e}_1^T \mathbf{x}_2 \\
 &= \lfloor q/2 \rfloor M + s \cdot (x_1 - y) + e_2 - \mathbf{e}_1^T \mathbf{x}_2.
 \end{aligned}$$

Furthermore, the following bounds apply:

- $\|s \cdot x_1\| \leq t^2 \tau \zeta n$ from Lemma 2 and Theorem 2.
- $\|y \cdot s\| \leq t^2 \tau \sigma_g \sqrt{(b^{2\ell} - 1)/(b^2 - 1)} n$ from Lemma 2 and Theorem 2.
- $\|e_2\| \leq t\tau \sqrt{n}$ from Lemma 2.
- $\|\mathbf{e}_1^T \mathbf{x}_2\| \leq t^2 \tau \zeta n m$ from Lemma 2 and Theorem 2.

By substituting these bounds, we obtain:

$$\begin{aligned}
 \|s \cdot (x_1 - y) + e_2 - \mathbf{e}_1^T \mathbf{x}_2\| &\leq \|s \cdot x_1\| + \|y \cdot s\| + \|e_2\| + \|\mathbf{e}_1^T \mathbf{x}_2\| \\
 &\leq t^2 \tau \left[\zeta (m+1) + \sigma_g \sqrt{(b^{2\ell} - 1)/(b^2 - 1)} \right] + t\tau \sqrt{n} \\
 &\leq \lfloor q/4 \rfloor.
 \end{aligned}$$

Parameters constraints. The following constraints combined with the errors norm constraints above should be satisfied:

- The Gaussian parameter σ_g used for the \mathbf{G} -sampling in the ApproxSamplePre2 algorithm must verify $\sigma_g \geq \sqrt{2b} \cdot (2b+1) \cdot \sqrt{\log(2nw(1+1/\epsilon))/\pi}$ (see [17], Corollary 3.1).
- The Gaussian width for preimage sampling ζ must follow the condition $\zeta > \sqrt{(\sigma_g^2 + 1)s_1^2(\mathbf{R}) + \eta_\epsilon^2(\mathbb{Z}^{nm})}$, where $s_1(\mathbf{R})$ is the spectral norm of the trapdoor \mathbf{R} (see [4], Section A.3).

The proof of the following is adapted from the GPV IBE proof [18, Section 7.2] and is given in the Supplementary materials, Appendix C.

Theorem 4. *Our IBE construction with parameters $n, m, q, k, \ell, \sigma, \sigma_g, \zeta, \tau$ and γ is IND-sID-CPA secure in the random oracle model under the hardness of $\text{iNTRU}_{q, D_{R, \sigma}}$ and Ring-LWE $_{n, q, \tau}$.*

5.2 Implementation and performance

In [15], the authors only propose a reduction from a non-standard version of LWE to iNTRU. Therefore, in the absence of a thorough study on the asymptotic and practical security of the iNTRU problem, which we leave for future work, we have chosen to estimate the security of our iNTRU instances by relying on the existing cryptanalysis on NTRU. As explained in [15], there is a syntactic link between NTRU and iNTRU. To the best of our knowledge, there is no known reduction between the two problems and the analysis of the iNTRU assumption might deserve additional study. Still, we additionally consider the practical security of NTRU for our targetted sets of parameters and we take into account the known cryptanalysis efforts on iNTRU.

For security estimation, we follow the same approach as in [16] to assess the concrete security of the IBE2 scheme. We determine the hardness of our underlying lattice problem by computing the root Hermite factor, introduced in [14]. Then, we use the following heuristic relation between the blocksize κ and the root Hermite factor δ to find the smallest blocksize which would break our underlying lattice problem:

$$\delta \approx \left(\frac{\kappa}{2\pi e} (\pi\kappa)^{1/\kappa}\right)^{1/2(\kappa-1)}.$$

Finally, our experiments estimate the running time of the BKZ algorithm to analyze the concrete security of the scheme. This algorithm makes use of an oracle to solve the Shortest Vector Problem (SVP) in smaller lattices. We chose the "Core-SVP" model introduced in [3] in the sieving regime as the SVP oracle for the BKZ algorithm with time complexity $2^{0.292\kappa+16.4}$ in the blocksize κ .

Overstretched parameters. [23] adapts the attack from [20] on iNTRU, which applies to the parameters originally proposed for the homomorphic encryption scheme from [15]. The attack can be performed when the modulus q is much larger than the dimension of the associated lattices. The target iNTRU instance has an error and secret that follow a uniform binary distribution. Following [20] attack, [23] uses the fact that a very dense sublattice can be found in this NTRU-like lattice, because of the overstretched regime. Relying on a lemma from Pataki and Tural [30, Lemma 1], they can bound the volume of this sublattice and run a BKZ-reduction which leads to a full recovery of the iNTRU secret. [12] improves the asymptotic bound given by Kirchner and Fouque [20] by conducting a refined analysis which lowers the overstretched regime for NTRU with ternary distribution to the value $q = n^{2.484+o(1)}$.

They also provide a concrete analysis, computing a bound on the modulus q for which the attacks exploiting the overstretched regime are more efficient than standard secret key recovery attacks. The authors of [12] run experiments which allow detecting the fatigue "point", which separates these two regimes.

The cryptanalysis carried out by [23] and [12] can be adapted for the iNTRU instances we consider, where the secret and the error both follow a Gaussian distribution. [12] also consider the case where the NTRU secret distributions are Gaussian. Indeed, in both cases, the cryptanalysis in the overstretched regime requires performing lattice reductions on sublattices of a NTRU-like lattice, in order to retrieve the very dense sublattice. We leave the detailed adaptation of this attack to iNTRU and its experimental study for future work. For our concrete parameters analysis, we took care to fall outside the range of parameters affected in the overstretched regime.

Timings. Our timings have been obtained on an Intel i7-8650U CPU running at 1.9 GHz, and then scaled at 4.2GHz to compare ourselves with other schemes. Results are provided in Table 4. The use of the efficient gadget-based approximate trapdoor framework together with the iNTRU hardness assumption allows us to obtain efficient algorithms. The slowest of the 4 algorithms of the IBE2 scheme are the Setup and Extract algorithms, which correspond respectively to the approximate trapdoor generation and preimage sampling phases of the scheme. However, the Setup algorithm is usually not performed often, and the subroutine algorithms used by Extract for sampling are really modular, leaving the way for possible future improvements. Moreover, our Setup and Extract algorithms are competitive with other NTRU-based IBE (see Table 5 and Table 6).

Comparison with related works. We compare the IBE2 timings with the ones of [11] (re-implemented in [27]) and with [33], two IBE schemes whose security is based on the NTRU hardness assumption. Our comparison experiments were carried out using equivalent parameters sets between the different schemes. In particular, we have been careful to use equivalent module sizes and equivalent noise rates when performing Gaussian Sampling.

(n, q)	$\lceil \log_2(q) \rceil$	Setup	Extract	Encrypt	Decrypt	Bit security
(256, 1073741441)	30	1.01	2.13	0.39	0.03	64
(512, 1073741441)	30	2.12	3.79	0.74	0.06	115
(1024, 1073741441)	30	4.08	7.65	1.49	0.11	223
(256, 16777601)	25	0.78	1.66	0.23	0.02	35
(512, 16777601)	25	1.48	3.00	0.49	0.04	82
(1024, 16777601)	25	3.32	5.92	1.10	0.07	159

Table 4: Timings of the different operations for different values of n , given in ms from Setup to Decrypt.

Scheme	$(n, \lceil \log_2(q) \rceil)$	Security level	Setup	Extract	Encrypt	Decrypt
[11]	(512, 26)	< 80	3.84s	1.77	0.10	0.05
[11]	(1024, 26)	< 192	23.93s	6.95	0.27	0.09
This paper	(512, 25)	82	1.48	3.00	0.49	0.04
This paper	(1024, 25)	159	3.32	5.92	1.10	0.07

Table 5: Timings comparison of the different operations of the IBE scheme between this paper and [11] (the timings are extracted from the [27] article, scaled up to account for CPU differences) for different parameter sets. The timings are given in ms, except for the Setup algorithm from [11], which is given in seconds.

Scheme	$(n, \lceil \log_2(q) \rceil)$	Security level	Setup	Extract	Encrypt	Decrypt
[33]	(1024, 24)	128	102	0.82	0.05	0.06
[33]	(2048, 25)	256	292	2.62	0.10	0.13
[33]	(1024, 36)	80	165	26.4	0.08	0.09
[33]	(2048, 38)	160	643	57.8	0.16	0.18
This paper	(1024, 25)	159	3.32	5.92	1.10	0.07
This paper	(2048, 25)	293	10.21	12.79	2.96	0.16
This paper	(1024, 30)	191	4.08	7.65	1.49	0.11
This paper	(2048, 30)	412	12.51	16.03	3.89	0.23

Table 6: Timings comparison of the different operations of the IBE scheme in this paper and in [33] for different sets of parameters in ms.

We observe that we obtain better timings for the Setup algorithm than [11] and [33] and for the Extract for some sets of parameters. Furthermore, our Decrypt algorithm is slightly faster than [11]. However, the Encrypt algorithm is less efficient than theirs. The use of binomial distribution improves the timings for encryption. An improvement can be obtained in our case by using the binomial distribution, but we need more samples in our case which still affects performance for encryption; we make $n(2 + m)$ calls to the integer Gaussian sampler while encryption in [11] and [33] can make use of only $3n$ binomial sampling calls. As in [33], our Extract algorithm is slower than the Sign algorithm from the Falcon signature scheme [13]. Note also that for a similar security level, Falcon can use smaller parameters than us.

We obtain an overhead in terms of parameters sizes compared to [11] and [33], but the trapdoor generations rely on a different paradigm. Results on public and secret sizes are provided in Appendix D. Nonetheless, as stated in [9], the use of approximate trapdoors instead of exact ones helps us to reduce the sizes of the public key and signatures by up to two times. Therefore, as expected, our obtained sizes for the master and the users' private keys and the ciphertexts are close to the ones of [16] whose signature scheme relies on the same paradigm as the IBE2 scheme.

Acknowledgements. This work is supported by the ANR ASTRID project AMIRAL (ANR-21-ASTR-0016).

Supplementary materials

A Definition of an IBE scheme

An IBE scheme is composed of 4 polynomial time algorithms defined as follows:

$\text{Setup}(1^n) \rightarrow (\text{mpk}, \text{msk})$: takes as input a security parameter n and outputs the master public key mpk and the master secret key msk . We assume public parameters which are part of mpk are taken as input to all 3 next algorithms.

$\text{Extract}(\text{mpk}, \text{msk}, \text{id}) \rightarrow \text{sk}_{\text{id}}$: takes as input the security parameter, the master keys mpk , msk and an identity $\text{id} \in \mathcal{ID}$ and outputs a private key sk_{id} associated with the identity id .

$\text{Encrypt}(\text{mpk}, \text{id}, M) \rightarrow C$: takes as input the security parameter, the master public key mpk , an identity id and a message M and outputs a ciphertext C which is the encryption of M for identity id .

$\text{Decrypt}(\text{mpk}, \text{sk}_{\text{id}}, C) \rightarrow M$ or *Error*: takes as input the master public key mpk , a private key sk_{id} associated with the identity id and a ciphertext C , it outputs either a message M or an *Error* message indicating the ciphertext is invalid.

Correctness. An IBE is correct if for any message M in the message space and identity id such that $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^n)$ and $\text{sk}_{\text{id}} \leftarrow \text{Extract}(\text{mpk}, \text{msk}, \text{id})$, $\text{Decrypt}(\text{mpk}, \text{sk}_{\text{id}}, \text{Encrypt}(\text{mpk}, \text{id}, M)) = M$ holds with overwhelming probability.

Security Game. We describe the notion of ciphertext indistinguishability under a selective-identity chosen-plaintext attack (IND-sID-CPA), meaning that the adversary indicates the identity on which he would like to be challenged. The notion is defined as a game between an adversary \mathcal{A} and a challenger, and works as follows:

- **Initialisation:** The adversary \mathcal{A} chooses the challenge identity id^* .
- **Setup :** The challenger runs $\text{Setup}(1^n)$ and gives the master public key mpk to \mathcal{A} .
- **Extract queries - Phase 1:** The adversary can make private-key extraction queries on identities $\text{id} \neq \text{id}^*$, and the challenger answers by running $\text{sk}_{\text{id}} \leftarrow \text{Extract}(\text{mpk}, \text{msk}, \text{id})$.
- **Challenge:** The adversary \mathcal{A} outputs two plaintexts M_0, M_1 . The challenger chooses a random bit $b^* \leftarrow \{0, 1\}$ and sets the challenge ciphertext to $C^* = \text{Encrypt}(\text{mpk}, \text{id}^*, M_{b^*})$. The challenge ciphertext C^* is sent to \mathcal{A} .
- **Extract queries - Phase 2:** The adversary can make additional queries (answered as in Phase 1).
- **Guess:** The adversary outputs a bit b and wins if $b^* = b$.

The advantage of the adversary \mathcal{A} playing the IND-sID-CPA security game above is $\text{Adv}(\mathcal{A}_{\text{IBE}}^{\text{ind-sID-cpa}}) = |Pr(b = b^*) - \frac{1}{2}|$. An IBE scheme is IND-sID-CPA secure if, for all PPT adversary \mathcal{A} , his advantage $\text{Adv}(\mathcal{A}_{\text{IBE}}^{\text{ind-sID-cpa}})$ is negligible.

B Proof of Theorem 3

Proof. The proof proceeds in a sequence of games G_i where the first game G_0 is identical to the original IND-sID-CPA game. In the last game in the sequence, the adversary has no information left about the initial message and hence has advantage zero. We show that a PPT adversary cannot distinguish between the intermediate games.

Game G_0 . This is the initial IND-sID-CPA game between an adversary \mathcal{A} and an IND-sID-CPA challenger.

Game G_1 . In this game, we change the way the challenger generates the public matrix \mathbf{A} by adding information about the identity id^* that \mathcal{A} targets for his attack. In Game 0, \mathbf{A} was generated thanks to $\text{ApproxTrapGen1}(\mathbf{0}, \sigma)$, together with the associated trapdoor \mathbf{R} . We had $\mathbf{A} = [\bar{\mathbf{A}} \mid -\bar{\mathbf{A}}\mathbf{R}]$. The public matrix \mathbf{A} is now generated in G_1 thanks to the algorithm $\text{ApproxTrapGen1}(-\mathbf{H}_{\text{id}^*}, \sigma)$, so that we have $\mathbf{A} = [\bar{\mathbf{A}} \mid -\mathbf{H}_{\text{id}^*}\mathbf{F} - \bar{\mathbf{A}}\mathbf{R}] \in \mathcal{R}_q^{d \times m}$.

This matrix \mathbf{A} is statistically indistinguishable from a uniformly random matrix.

The challenger answers \mathcal{A} 's private key queries on identities $\text{id} \neq \text{id}^*$ by calling $\text{Extract}((\mathbf{A}, \mathbf{u}), \mathbf{R}, \text{id})$ and then using $\text{ApproxSamplePre1}(\mathbf{A}_{\text{id}}, \mathbf{R}, \mathbf{H}_{\text{id}}, \mathbf{u}, \zeta)$. We have $\mathbf{A}_{\text{id}} = [\bar{\mathbf{A}} \mid (\mathbf{H}_{\text{id}} - \mathbf{H}_{\text{id}^*})\mathbf{F} - \bar{\mathbf{A}}\mathbf{R}]$ and $\mathbf{H}_{\text{id}} - \mathbf{H}_{\text{id}^*}$ is invertible because of the FRD construction, which allows the challenger to get a private key associated with the identity id .

Game G_2 . This game is identical to G_1 except that the challenge ciphertext \mathbf{C}^* is chosen as a uniformly random element in \mathcal{R}_q^{m+1} instead of being chosen by calling $\text{Encrypt}((\mathbf{A}, \mathbf{u}), \text{id}^*, M_{b^*})$.

Reduction from Module-LWE. We show that G_2 and G_1 are computationally indistinguishable for the adversary \mathcal{A} under the Module-LWE assumption.

Suppose \mathcal{A} has a non-negligible advantage in distinguishing G_2 and G_1 . We will use \mathcal{A} to construct a simulator \mathcal{B} who will be able to solve the Module-LWE problem with non-negligible advantage.

The simulator \mathcal{B} receives $\bar{m} + 1$ samples $(\mathbf{A}_i, b_i)_{0 \leq i \leq \bar{m}}$ with $\mathbf{A}_i \in \mathcal{R}_q^d$ and $b_i \in \mathcal{R}_q$ as an instance of the decisional Module-LWE problem. The simulator also receives the identity id^* targeted by the adversary \mathcal{A} . The simulator sets $\mathbf{A}' = (\mathbf{A}_1, \dots, \mathbf{A}_{\bar{m}}) \in \mathcal{R}_q^{d \times \bar{m}}$ and $\mathbf{b}' = (b_1, \dots, b_{\bar{m}}) \in \mathcal{R}_q^{\bar{m}}$ and then runs $\text{ApproxTrapGen1}(-\mathbf{H}_{\text{id}^*}, \sigma)$ to get $\mathbf{A} = [\mathbf{A}' \mid -\mathbf{H}_{\text{id}^*}\mathbf{F} - \mathbf{A}'\mathbf{R}]$ together with the trapdoor. Next, \mathcal{B} sets $\mathbf{u} = \mathbf{A}_0 \in \mathcal{R}_q^d$ and he sends (\mathbf{A}, \mathbf{u}) to \mathcal{A} as the master public key. Then, \mathcal{B} answers \mathcal{A} 's private key queries as in G_1 .

Afterwards, the adversary outputs two plaintexts M_0, M_1 and sends them to \mathcal{B} . The simulator generates a random bit b^* , and the challenge ciphertext $\mathbf{C}^* = (b^*, c^*)$ as follows:

$$\mathbf{b}^* = (\mathbf{b}'^T \mid -\mathbf{b}'^T \mathbf{R} + \hat{\mathbf{e}}^T)^T, \quad c^* = b_0 + \lfloor q/2 \rfloor M_{b^*},$$

where $\hat{e} \leftarrow D_{\mathcal{R}^w, \mu}$ for some real μ , that is explicited below.

Samples from the LWE distribution. If the Module-LWE samples are drawn from the Module-LWE distribution, we have:

$$\mathbf{b}^T = \mathbf{s}^T \mathbf{A}' + \mathbf{e}'^T \text{ and } b_0 = \mathbf{s}^T \mathbf{A}_0 + e_0,$$

for some $\mathbf{s} \in \mathcal{R}_q^d$, $\mathbf{e}' \leftarrow D_{\mathcal{R}^{\bar{m}}, \tau}$ and $e_0 \leftarrow D_{\mathcal{R}, \tau}$. Therefore, the first part of the ciphertext is $\mathbf{b}^* = \mathbf{A}_{\text{id}^*}^T \mathbf{s} + (\mathbf{e}'^T - \mathbf{e}'^T \mathbf{R} + \hat{\mathbf{e}}^T)^T$ and the second part is $c^* = \mathbf{s}^T \mathbf{A}_0 + e_0 + \lfloor q/2 \rfloor M_{b^*}$.

Then, for a fixed error \mathbf{e}' , the term $-\mathbf{e}'^T \mathbf{R}$ is distributed as $D_{\mathcal{R}^w, \sigma \|\mathbf{e}'\|}$ since $\mathbb{E}[\mathbf{e}'^T \mathbf{R}] = \mathbf{e}'^T \mathbb{E}[\mathbf{R}] = \mathbf{0}$ and $\text{cov}(\mathbf{e}'^T \mathbf{R}) = \sigma^2 \mathbf{e}'^T \mathbf{e}' \mathbf{I}_w = \sigma^2 \|\mathbf{e}'\|^2 \mathbf{I}_w$ by linearity of expectation and bilinearity of covariance. Moreover, the random variable $\hat{\mathbf{e}}$ follows the Gaussian distribution $D_{\mathcal{R}^w, \mu}$ and is independent from the random variable $\mathbf{e}'^T \mathbf{R}$. Therefore, $-\mathbf{e}'^T \mathbf{R} + \hat{\mathbf{e}}$ is indistinguishable from a sample drawn from the distribution $D_{\mathcal{R}^w, \gamma}$ for μ satisfying $\gamma^2 = (\sigma \|\mathbf{e}'\|)^2 + \mu^2$.

Then, the challenge ciphertext (\mathbf{b}^*, c^*) follows the same distribution as in G_1 .

Samples from the uniform distribution. If the Module-LWE samples are drawn from the uniform distribution, the ciphertext challenge also looks uniform as in G_2 .

Finally, the adversary \mathcal{A} outputs a guess b . Because we supposed \mathcal{A} has a non-negligible advantage in distinguishing G_1 and G_2 , if $b = b'$ with overwhelming probability, the simulator concludes that the challenged instance was drawn from the Module-LWE distribution. Otherwise, \mathcal{B} concludes that the Module-LWE instance was drawn from the uniform distribution.

C Proof of Theorem 4

For the sake of completeness, we give the IND-sID-CPA proof of the IBE2 scheme, which is directly adapted from [18, Theorem 7.2].

Proof. The idea of the proof is to simulate the view of the adversary given a Ring-LWE instance (a'_i, b'_i) for $0 \leq i \leq m$ for which we build an attacker \mathcal{S} who makes $Q_{\mathcal{H}}$ queries to \mathcal{H} .

We construct a new simulation game given $\mathbf{a}_0 = [a'_1 \dots a'_m]$. \mathcal{S} chooses an index i^* at random in $[1, Q_{\mathcal{H}}]$ and sets $h_{\text{id}^*} = a'_{i^*}$.

A hash query on input id is answered as follows: it first checks whether an entry of the form $(\text{id}, *, *, *)$ already exists in the hash table. If it is not the case, it responds with a value u_{id} such that an approximated preimage for $\mathcal{H}(\text{id}) = h_{\text{id}}$ is known i.e: it samples $\mathbf{x} \leftarrow D_{\mathcal{R}^{m+1}, \zeta}$ and $y_{\text{id}} \leftarrow D_{\mathcal{R}, \sigma_e}$ such that $y_{\text{id}} = (x_1^T, \mathbf{x}_2^T)^T [1 \ \mathbf{a}_0]$, sets $h_{\text{id}} = (x_1^T, \mathbf{x}_2^T)^T [1 \ \mathbf{a}_0] - y_{\text{id}}$ and adds the tuple $(\text{id}, h_{\text{id}}, \mathbf{x}_2, y_{\text{id}})$ in the hash table.

An extraction query for an identity id is responded as follows: assuming that an entry for id already exists in the table, the corresponding \mathbf{x}_2 is output by \mathcal{S} .

Note that the response to the extract queries is close to the response provided in the real game by Theorem 2 if \mathcal{H} is modeled as a random oracle and the way the public key is generated is indistinguishable from that in the real game under the iNTRU assumption.

When the attacker outputs two messages m_0, m_1 and id^* , if id^* has not been queried to \mathcal{H} , then \mathcal{S} aborts; otherwise the challenger sets the challenger ciphertext as:

$$(\mathbf{b}, c) = ([b'_1 \dots b'_m], b'_0 + \lfloor q/2 \rfloor m_b) \in \mathcal{R}_q^{m+1}.$$

The attacker \mathcal{S} outputs the same bit as \mathcal{A} . Assuming no abort has occurred, they both have the same advantage, which concludes the proof.

D Parameters' sizes comparison

The use of the [28] paradigm doesn't make us competitive compared to [11] and [33] in terms of parameters sizes. As explained in Section 5.2, the trapdoor generations rely on a different paradigm. However, for completeness, we provide comparisons with [11] in Table 7 and with [33] in Table 8 in terms of parameters' sizes.

Table 7: Parameters sizes comparison between IBE2 parameters sizes and [11] for different sets of parameters at a similar security level, in Bytes.

Scheme	$(n, \lceil \log_2(q) \rceil)$	Security	mpk	msk	sk_{id}	Ciphertext
[11]	(512, 23)	< 80	1536	6144	1375	1625
[11]	(1024, 27)	< 192	3584	14336	3375	3750
This paper	(512, 25)	\approx 80	14720	16192	16192	16192
This paper	(1024, 25)	\approx 192	32000	35200	35200	35200

Table 8: Parameters sizes comparison between IBE2 parameters sizes and [33] for different sets of parameters at a similar security level, in Bytes.

Scheme	$(n, \lceil \log_2(q) \rceil)$	Security	mpk	msk	sk_{id}	Ciphertext
[33]	(1024, 36)	80	4608	18432	4608	9248
[33]	(1024, 24)	128	3072	12288	3072	6176
[33]	(2048, 25)	256	6400	25600	6400	12832
This paper	(512, 25)	> 80	14720	16192	16192	16192
This paper	(1024, 25)	> 128	32000	35200	35200	35200
This paper	(2048, 25)	> 256	64000	70400	70400	70400

We observe that using approximate trapdoors instead of exact ones can help reduce the size of public keys and signatures by up to two times.

References

- [1] S. Agrawal, D. Boneh, and X. Boyen. “Efficient Lattice (H)IBE in the Standard Model”. In: *EUROCRYPT*. LNCS. Springer, 2010.
- [2] M. R. Albrecht, R. Player, and S. Scott. “On the concrete hardness of Learning with Errors”. In: *J. Mathematical Cryptology* (2015).
- [3] E. Alkim, L. Ducas, T. Pöppelmann, and P. Schwabe. “Post-quantum Key Exchange - A New Hope”. In: *USENIX Security Symposium*. 2016.
- [4] P. Bert, G. Eberhart, L. Prabel, A. Roux-Langlois, and M. Sabt. “Implementation of Lattice Trapdoors on Modules and Applications”. In: *PQCrypto*. Lecture Notes in Computer Science. Springer, 2021.
- [5] P. Bert, P. Fouque, A. Roux-Langlois, and M. Sabt. “Practical Implementation of Ring-SIS/LWE Based Signature and IBE”. In: 2018.
- [6] D. Boneh and M. K. Franklin. “Identity-Based Encryption from the Weil Pairing”. In: *CRYPTO*. LNCS. Springer, 2001.
- [7] K. Boudgoust, C. Jeudy, A. Roux-Langlois, and W. Wen. “On the Hardness of Module Learning With Errors with Short Distributions”. In: *J. Cryptology* 28.1 (2023).
- [8] D. Cash, D. Hofheinz, E. Kiltz, and C. Peikert. “Bonsai Trees, or How to Delegate a Lattice Basis”. In: *EUROCRYPT*. 2010.
- [9] Y. Chen, N. Genise, and P. Mukherjee. “Approximate Trapdoors for Lattices and Smaller Hash-Sign Signatures”. In: *ASIACRYPT*. 2019.
- [10] C. C. Cocks. “An Identity Based Encryption Scheme Based on Quadratic Residues”. In: *IMACC*. Springer, 2001.
- [11] L. Ducas, V. Lyubashevsky, and T. Prest. “Efficient Identity-Based Encryption over NTRU Lattices”. In: *ASIACRYPT*. 2014.
- [12] L. Ducas and W. P. J. van Woerden. “NTRU Fatigue: How Stretched is Overstretched?” In: *ASIACRYPT (4)*. 2021.
- [13] P.-A. Fouque, J. Hoffstein, P. Kirchner, V. Lyubashevsky, T. Pornin, T. Prest, T. Ricosset, G. Seiler, W. Whyte, and Z. Zhang. *FALCON: Fast-Fourier Lattice-based Compact Signatures over NTRU*. 2017.
- [14] N. Gama and P. Q. Nguyen. “Predicting Lattice Reduction”. In: *EUROCRYPT*. Lecture Notes in Computer Science. Springer, 2008.
- [15] N. Genise, C. Gentry, S. Halevi, B. Li, and D. Micciancio. “Homomorphic Encryption for Finite Automata”. In: *ASIACRYPT*. 2019.
- [16] N. Genise and B. Li. “Gadget-Based iNTRU Lattice Trapdoors”. In: *INDOCRYPT*. 2020.
- [17] N. Genise and D. Micciancio. “Faster Gaussian Sampling for Trapdoor Lattices with Arbitrary Modulus”. In: *EUROCRYPT*. 2018.
- [18] C. Gentry, C. Peikert, and V. Vaikuntanathan. “Trapdoors for hard lattices and new cryptographic constructions”. In: *STOC*. 2008.
- [19] C. F. F. Karney. “Sampling Exactly from the Normal Distribution”. In: *ACM Trans. Math. Softw.* 42.1 (2016), 3:1–3:14.
- [20] P. Kirchner and P. Fouque. “Revisiting Lattice Attacks on Overstretched NTRU Parameters”. In: *EUROCRYPT (1)*. Vol. 10210. Lecture Notes in Computer Science. 2017, pp. 3–26.

- [21] R. W. F. Lai, H. K. F. Cheung, and S. S. M. Chow. “Trapdoors for Ideal Lattices with Applications”. In: *Inscrypt*. 2014.
- [22] A. Langlois and D. Stehlé. “Worst-case to average-case reductions for module lattices”. In: *DCC 75.3* (2015), pp. 565–599.
- [23] C. Lee and A. Wallet. *Lattice analysis on MiNTRU problem*. Cryptology ePrint Archive, Paper 2020/230. 2020.
- [24] V. Lyubashevsky. “Lattice Signatures without Trapdoors”. In: *EUROCRYPT*. Vol. 7237. LNCS. Springer, 2012, pp. 738–755.
- [25] V. Lyubashevsky, C. Peikert, and O. Regev. “A Toolkit for Ring-LWE Cryptography”. In: *EUROCRYPT*. Springer, 2013.
- [26] V. Lyubashevsky and G. Seiler. “Short, Invertible Elements in Partially Splitting Cyclotomic Rings and Applications to Lattice-Based Zero-Knowledge Proofs”. In: *EUROCRYPT*. 2018.
- [27] S. McCarthy, N. Smyth, and E. O’Sullivan. “A Practical Implementation of Identity-Based Encryption Over NTRU Lattices”. In: 2017.
- [28] D. Micciancio and C. Peikert. “Trapdoors for Lattices: Simpler, Tighter, Faster, Smaller”. In: *EUROCRYPT*. 2012.
- [29] D. Micciancio and O. Regev. “Worst-Case to Average-Case Reductions Based on Gaussian Measures”. In: *SIAM J. Comput.* (2007).
- [30] G. Pataki and M. Tural. “On sublattice determinants in reduced bases”. In: *arXiv preprint arXiv:0804.4014* (2008).
- [31] A. Shamir. “Identity-Based Cryptosystems and Signature Schemes”. In: *CRYPTO*. 1984.
- [32] S. Zhang and Y. Yu. “Towards a Simpler Lattice Gadget Toolkit”. In: *Public Key Cryptography*. 2022.
- [33] R. K. Zhao, S. McCarthy, R. Steinfeld, A. Sakzad, and M. O’Neill. *Quantum-safe HIBE: does it cost a Latte?* ePrint Archive. 2021.