Léo Ducas

Cryptology Group
Centrum voor Wiskunde & Informatica
ducas@cwi.nl

Advances on quantum cryptanalysis of ideal lattices

Léo Ducas is a NWO Veni Laureate working in CWI's Cryptology Group on lattice-based cryptology. He co-authored the efficient quantum-safe key exchange algorithm 'New Hope' which was awarded the Facebook Internet Defense Award. Although using lattices with additional structure can lead to more efficient cryptographic algorithms, in this article Ducas explains how lattices that have too much additional structure may be insecure due to efficient quantum attacks.

The problem of finding a shortest vector of a Euclidean lattice (the shortest vector problem, or SVP) is a central hard problem in complexity theory. Approximated versions of this problem (e.g. α-SVP, the problem of finding a vector at most α times longer than the shortest one) have become the theoretical foundation for many cryptographic constructions. Indeed, lattice-based cryptography typically benefits from *worst-case hardness* [1, 14, 18]: it is sufficient that there exists *some* lattices in which finding short vectors is hard for those cryptosystems to be secure. Among several advantages, lattice-based cryptography is also praised for its apparent resistance to quantum algorithms, unlike the current public-key schemes based on factoring or discrete logarithm.

The main drawback of lattice-based cryptography is its large memory and bandwidth footprints: a lattice is represented by a basis, i.e. an $n \times n$ matrix for a dimension n of several hundreds. For efficiency reasons, it is tempting to rely on structured

lattices, such as lattices generated by a circulant matrix. The earliest example of such a cryptosystem is the NTRUᴇɴᴄʀʏᴘᴛ proposal from Hoffstein et al. [9] from 1998. Algebraically, those lattices can be viewed as ideals or modules over cyclotomic number fields.

Nevertheless, there is no guarantee that hard lattice problems remain hard on particular classes of structured lattices, and indeed, a series of results [4–8] have lead to new quantum algorithms solving certain ideal lattice problems. To the best of our

knowledge, the same problems remain hard over arbitrary lattices, even with a quantum computer. More precisely, for certain sub-exponential approximation factors α, α-SVP on ideal lattices admit a polynomial-time algorithm, as depicted in Figure 1. In this survey, we give an overview of the techniques that have lead to these results.

The first quantum attack on certain ideal lattices of cyclotomic fields was sketched by Campbell, Groves and Shefferd [5], and applies to a few schemes, in particular to one of the first Fully-Homomorphic Encryption schemes [17]. Yet those broken schemes were based on ad-hoc problems that do not benefit from worse-case hardness.

The first step of this attack does not actually solve a lattice problem: it does not provide guarantees about the shortness of


Figure 1 Best known quantum algorithm for general α-SVP (left), and for α-SVP in cyclotomic ideal lattices (right).

the solution. Namely, hinted by recent results of Eisenträger, Hallgren, Kitaev and Song [8], it is conjectured in [5] that the Principal Ideal Problem (finding a generator of a given principal ideal) could be solved in quantum polynomial time. This was soon confirmed by the work of Biasse and Song [4]. The second step was also only conjectured to be correct, but could easily be checked in practice. Precisely, taking logarithms, finding a short generator can be phrased as a lattice problem in a fixed lattice (Dirichlet's unit lattice), for which we know a seemingly good basis. A detailed geometric analysis of the cyclotomic units [6] confirmed that conjecture, using tools from analytic number theory.

While this initial attack concerned a particular distribution of principal ideal lattices, the work of [6] also considers what can be done in the worst-case: using similar algorithms, one can always recover a generator longer than the shortest vector by a factor at most $\alpha = \exp(\tilde{O}(\sqrt{n}))$. This constitutes a first worst-case hardness gap between generic lattices and structured ones. The gap was widened in a follow-up result of Cramer, Ducas and Wesolowski [7], showing how to extend these algorithms to non-principal ideals. Naturally, one would look for an ideal $\mathfrak{a}\mathfrak{b} \subset \mathfrak{a}$ which is a multiple of $\mathfrak{a}$, that is principal, and with a small relative index $\#(\mathfrak{a}/\mathfrak{a}\mathfrak{b})$. Again, this problem can be translated to a lattice problem in a fixed lattice, namely the lattice underlying Stickelberger's class group annihilation theorem [19].

Lattices and computational problems

We recall that a lattice is a discrete subgroup of the vector space $\mathbb{R}^n$, equipped with its canonical Euclidean norm denoted $\|\cdot\|$. The minimal distance of a lattice Λ is defined by $\lambda_1(\Lambda) = \min_{\mathbf{x} \in \Lambda\setminus\{0\}} \|x\|$.

Our main goal is to solve the following problem in a particular class of lattices.

Definition. The Short Vector Problem with approximation factor α (α-SVP) is defined as:

– Given a basis $\mathbf{B}$ of a lattice $\Lambda \subset \mathbb{R}^n$,
– Find $\mathbf{v} \in \Lambda\setminus\{0\}$ such that $\|v\| \le \alpha \cdot \lambda_1(\Lambda)$.

For our purpose, we will also consider two related problems, namely the approximate Close Vector Problem (δ-CVP), and the Bounded Distance Decoding problem (δ-BDD). For convenience, we will define


Figure 2   Rounding with a good basis.

Figure 3   Rounding with a bad basis.

them with respect to an absolute distance δ, rather than an approximation factor α.

Definition. The Close Vector Problem up to distance δ (δ-CVP) is defined as:

– Given a basis $\mathbf{B}$ of a lattice $\Lambda \subset \mathbb{R}^n$,
– Given a target $\mathbf{t} \in \mathbb{R}^n$,
– Find $\mathbf{v} \in \Lambda\setminus\{0\}$ such that $\|v - t\| \le \delta$,

where δ is large enough so that a solution exists for any target $\mathbf{t}$ (namely δ is larger than the *covering radius* of Λ).

Definition. The Bounded Distance Decoding Problem up to distance δ (δ-BDD) is defined as:

– Given a basis $\mathbf{B}$ of a lattice $\Lambda \subset \mathbb{R}^n$,
– Given a target $\mathbf{t} \in \mathbb{R}^n$ at distance at most δ from Λ,
– Find $\mathbf{v} \in \Lambda\setminus\{0\}$ such that $\|v - t\| \le \delta$,

where δ is small enough so that at most one solution exists for any target $\mathbf{t}$ (namely $\delta < \lambda_1(\Lambda)/2$).

Both problems are somehow dual, in particular δ-CVP gets easier as δ increases, while δ-BDD gets easier as δ decreases. A very simple and efficient algorithm for those problems is given by a simple coordinate-wise rounding:

$$\mathbf{v} = \mathbf{B} \cdot \lfloor \mathbf{B}^{-1} \cdot t \rceil.$$

This algorithm induces a parallelepipedic tiling of the space as depicted in Figure 2, where the shape of the tile P is given by the basis $\mathbf{B}$:

$$P = \mathbf{B} \cdot \left[-\tfrac{1}{2}, \tfrac{1}{2}\right]^m = \left\{ \sum \mathbf{x}_i b_i \mid x_i \in \left[-\tfrac{1}{2}, \tfrac{1}{2}\right] \right\}.$$

This fast algorithm solves δ_1-CVP (respectively δ_2-BDD) for radii defined by the small-est enclosing sphere of (respectively largest enclosed sphere) of P. More precisely:

$$\delta_1 = \tfrac{1}{2}\max\left\|\sum \pm \mathbf{b}_i\right\| \le \tfrac{1}{2}\sum \|\mathbf{b}_i\|,$$

$$\delta_2 = \tfrac{1}{2}\min\left\|\mathbf{b}_i^{\vee}\right\|,$$

where $\mathbf{b}_i^{\vee}$ denotes the i-th vector of the dual basis $(\mathbf{B}^T)^{-1}$. We see that the ability to solve these problems highly depends on the *quality* of the basis $\mathbf{B}$. For comparison, we picture what happens with a *bad* basis of the same lattice in Figure 3: the CVP and BDD radii get worse.

Lattice-based cryptography

The gap between what can be done with *good* and *bad* bases is what gives rise to public key cryptography: the bad basis will be used as a public key (allowing to generate noisy lattice points as ciphertexts), while the good basis is kept secret (allowing to solve BDD for decryption). The secret-key owner is able to construct a good basis only because he controls the construction of the lattice.

In this brief overview, we explained why CVP and BDD are useful problems in cryptography, but it turns out that the core problem is SVP. For example, solving SVP a few times allows to construct a basis with small vectors, allowing in turn to solve CVP. And the converse is also true for certain variants of CVP and BDD, as demonstrated by the worst-case to average-case reductions of Ajtai and others [1, 14, 18]. These converse results allow to *prove* that breaking certain cryptosystems is at least as hard as solving α-SVP for some approximation factor, typically polynomially large in the dimension $\alpha = n^{O(1)}$.

The hardness of α-SVP decrease with growing approximation factor α. For small $\alpha = O(1)$, this problem is known to be NP-hard [12], unfortunately it seems impossible to base cryptosystems on α-SVP with such a small approximation factor. The best known algorithms for α-SVP for polynomial approximation factors $\alpha = n^{O(1)}$ in unstructured lattices require time exponential in n. The conjecture that it cannot be done much faster implies that lattice-based cryptosystems are unbreakable in an asymptotic sense. More generally, the best algorithms to solve $\exp(n^c)$-SVP is BKZ [15], a generalization of the Lenstra–Lenstra–Lovàsz algorithm (LLL) [11], and runs in time $\exp(\tilde{\Theta}(n^{1-c}))$, as depicted in Figure 1.

Cyclotomic ideal lattices
Consider the m-th cyclotomic number field $K = \mathbb{Q}(\zeta)$, where ζ denotes a formal m-th primitive root of unity. Its ring of integer is known to be $O_K = \mathbb{Z}[\zeta]$. The number field K is equipped with $n = \phi(m)$ complex embeddings, sending ζ to each of the primitive m-th roots of unity in $\mathbb{C}$: $\psi_i : \zeta \mapsto \omega^i$ for each $i \in (\mathbb{Z}/m\mathbb{Z})^\times$, where $\omega = \exp(2\imath\pi/m)$.

An (integral) ideal $\Im \subset O_K$ is an additive subgroup of O_K also closed under multiplication by elements of O_K. An ideal may be viewed as a euclidean lattice via the Minkowski embedding:

$$\psi : x \in K \mapsto (\psi_1(x), \dots, \psi_{m-1}(x))$$
$$\in \mathbb{H} = \mathbb{C}^n \simeq \mathbb{R}^{2n}.$$

Each embedding ψ_i is a field morphism; in particular ψ is linear, and multiplication in K corresponds to component-wise multiplication in $\mathbb{H}$.

Quantum algorithms and HSP
In 1994, Shor [16] formulated a factorization algorithm that would run in polynomial time on a quantum computer. Shor's algorithm exploits the properties of quantum mechanics to efficiently find the period of the function:

$$f : x \in \mathbb{Z} \mapsto a^x \bmod N$$

which reveals the order r of $a \in (\mathbb{Z}/N\mathbb{Z})^\times$. Unless $a^{r/2} = -1 \bmod N$, the quantities $\gcd(a^{r/2}+1, N)$ and $\gcd(a^{r/2}-1, N)$ will provide non-trivial factors of N. Very similar ideas also allow to solve the discrete logarithm problem over any cyclic group G. In a world with large general-purpose quantum computers, all the public key cryptographic schemes deployed nowadays would become insecure, as they are all based on factoring and discrete logarithm problems.

This motivates the development of cryptosystems based on other mathematical problems, such as lattice-based schemes. This also calls for a better understanding of the power of quantum computers, especially with respect to Shor's idea of period finding. This is generalized as the following problem.

Definition. The Hidden Subgroup Problem (HSP) over the Abelian group G is defined as:
- Given an efficient quantum computable function $f : G \to S$, which is exactly H-periodic for a subgroup $H \subset G$,
- Find the hidden subgroup H,

where S denotes the set of quantum states.

This problem admits an efficient quantum algorithm in many cases. For example, Shor's algorithm is an instance of the HSP algorithm where $G = \mathbb{Z}$, $H = r\mathbb{Z}$, and where the function f is injective modulo the period r. In this particular case, f produces a classical result, that is trivially encoded as a quantum state. More generally, quantum algorithms for HSP are now known for larger Abelian groups such as $\mathbb{Z}^n$, and even $\mathbb{R}^n$ [8] with some technical restriction on f.

Quantum encodings of lattices
As we will see in the next section, many problems in number theory can be phrased as HSP using a function f producing lattices rather than quantum states: $f : G \to \mathcal{L}_V$ where $\mathcal{L} = \{L \mid L \subset V$ is a lattice$\}$. To apply the known HSP algorithm (using $R \circ f$) one therefore needs to be able to compute canonical representation for lattices $R : \mathcal{L}_V \to S$, a task that is not always so easy.

For integer lattices $L \subset \mathbb{Z}^d$, such a representation is provided by the Hermite Normal Form, which is computable in classical polynomial time: the representation R is classical. When $L \subset \mathbb{R}^d$ is a lattice of small dimension d, one can also compute a normal form, for example using an Hermite–Korkin–Zolotarev (HKZ) reduced basis. Again, this representation of L is purely classical. But as dimension grows, HKZ reduction becomes exponentially hard.

This issue was circumvented by Eisenträger et al. [8], this time by resorting to

Figure 4   Discrete Gaussian distribution over a lattice.

quantum representations. While we do not know of an efficiently computable normal form for non-integer large dimensional lattices, we do know how to sample according to a canonical distribution over a lattice, with an efficient classical probabilistic algorithm. First, one would start by finding a weakly reduced basis with the LLL algorithm [11], followed by Klein's sampling algorithm [10], which produces a wide discrete Gaussian distribution supported by the lattice L.

This probabibilistic algorithm producing a classical distribution can be adapted [8, 14] to a quantum algorithm producing the corresponding quantum state S, namely S is a weighted quantum superposition of all lattice points:

$$R : \mathcal{L} \to S,$$
$$R(L) = \sum_{\mathbf{x} \in L} \exp\left(\frac{-\|\mathbf{x}\|^2}{2\sigma^2}\right) \cdot | \mathbf{x} \rangle.$$

The above quantum superposition is infinite, but can be tail-cut considering the rapid decay of Gaussian distributions. Extra effort is also needed to discretize $\mathbb{R}^d$ and represent each point $\mathbf{x} \in \mathbb{R}^d$ using a finite amount of qubits (see 'straddle encoding' in [8]).

HSPs in number theory
In this section, we describe algorithms [4, 5, 8] that apply to any number field K, given its ring of integers O_K. Let us start by recalling the definitions of ideals of K:

Definition. An *integral ideal* $\Im$ of K is an additive subgroup $\Im \subset O_K$ closed by multiplication by elements of O_K:

$$a \in \Im, x \in O_K \to ax \in \Im.$$

A *fractional ideal* $\mathfrak{f}$ of K is a set of the form $\mathfrak{f} = \frac{1}{z}\Im$ for some non-zero integer $z \in \mathbb{Z}$, and some integral ideal $\Im \subset O_K$.

An ideal $\mathfrak{f} \subset K$ is said *principal* if it is generated by a single element, i.e. if $\mathfrak{f} = gO_K = \{gx \mid x \in O_K\}$ for some $g \in K$; such a g is called a *generator* of $\mathfrak{f}$.

Principal ideals have multiple generators, more precisely, g and $g' \in K$ generate the same ideal if and only if $u = g/g'$ is a unit of O_K. The multiplicative group of units is denoted $O_K^\times = \{u \in O_K \mid u^{-1} \in O_K\}$.

Recall that ideals can be multiplied:

$$\mathfrak{a} \cdot \mathfrak{b} = \Big\{\sum a_i b_i \mid a_i \in \mathfrak{a}, b_i \in \mathfrak{b}\Big\},$$

which makes the set $\mathscr{F}_K$ of fractional ideals an Abelian group. The set of $\mathscr{P}_K \subset \mathscr{F}_K$ of principal ideals form a subgroup of $\mathscr{F}_K$.

With those definitions, we can already consider two important computational problems in number theory.

Definition. The Unit Group Problem (UGP) consists in:

– Given a number field K and its ring of integers O_K,
– Find a finite set of units $u_1, \ldots, u_d \in O_K^\times$ that generate $O_K^\times$.

Definition. The Principal Ideal Problem (PIP) consists in:

– Given a number field K and its ring of integers O_K,
– Given a principal ideal $\mathfrak{I} \subset K$,
– Find a generator g of $\mathfrak{I}$.

Both problems can be viewed as particular cases of the more general problem of computing the group of S-units for a well chosen set S of prime ideals, as done in the paper of Biasse and Song [4].

We start by phrasing the Unit Group Problem as a multiplicative Hidden Subgroup Problem. Note that $u \in K$ is a unit of O_K if and only if $uO_K = O_K$, and more generally $gO_K = g'O_K$ if and only if $u = g/g'$ is a unit of O_K. This means that the function:

$$f_{\text{UGP}} : K^\times \to \mathscr{P}_K$$
$$x \mapsto x \cdot O_K$$

is (multiplicatively) $O_K^\times$-periodic, and one easily checks that it is injective modulo $O_K^\times$ as well. The images of such a function are ideals, and can therefore be viewed as lattices. Using the strategy described in the previous section, one can efficiently construct a canonical quantum representation of these lattices.

A lot of technicalities remain to implement this approach. In particular, the domain of the function f described above is the multiplicative group $K^\times$, while one

wishes to apply the known HSP algorithm over the vector space $\mathbb{R}^n$. This issue is essentially dealt with by resorting to the well known logarithmic embeddings:

$$\text{Log} : K^\times \to \mathbb{R}^n$$
$$x \mapsto (\log |\psi_1(x)|, \ldots, \log |\psi_{m-1}(x)|).$$

In more details, define $\text{Exp} : \mathbb{C}^n \to \mathbb{H}$ by coordinate-wise application of $\exp : \mathbb{C} \to \mathbb{C}^\times$. Up to an appropriate quotient on the domain, one can set

$$f_{\text{UGP}} : \mathbb{C}^n \to \mathcal{L}$$
$$\mathbf{y} \mapsto \text{Exp}(\mathbf{y}) \odot \psi(O_K)$$

and recovers $\text{Log}\, O_K^\times$ from the period of f_{UGP}. Geometrically, $f_{\text{UGP}}(\mathbf{y})$ is a deformation the of the lattice $\psi(O_K)$, where the i-th coordinate axis of $\mathbb{H} = \mathbb{C}^n$ has been stretched by the complex factor $\exp(y_i)$; this deformation leaves $\psi(O_K)$ invariant precisely when $\text{Exp}(\mathbf{y})$ equals to $\psi(u)$ for some unit $u \in O_K^\times$.

While this strategy seems simple, proving its correctness requires an in-depth analysis of the metric properties of $R \circ f_{\text{UGP}}$: Lipschitz continuity and some strong form of injectivity. We refer to the original article for more details [8].

Finally, we sketch a (over-)simplified strategy to generalize the above to the Principal Ideal Problem. Given a principal ideal $\mathfrak{I}$, one extends the function f to:

$$f_{\text{PIP}(\mathfrak{I})} : (\mathbf{y}, i) \mapsto \text{Exp}(\mathbf{y}) \odot \psi(\mathfrak{I}^i).$$

The periods of this function contains the extension of the previous, namely, it is $(\text{Log}\, O_K^\times, 0)$ periodic; but it is also $(\text{Log}\, g, -1)$ periodic for any generator g of $\mathfrak{I}$, as $f(\text{Log}\, g, -1) = g \cdot \mathfrak{I}^{-1} = O_K$. With this function $f_{\text{PIP}(\mathfrak{I})}$, the quantum algorithm for Hidden Subgroup Problem of [8] allows not only to recover the unit group, but also a generator of the principal ideal $\mathfrak{I}$. Again, much care is required to ensure that this strategy will indeed work, see [4].

Short generators of principal ideals

So far, we have been concerned with problems that were purely of number theoretic nature, in the sense that the solutions to UGP and PIP have no guarantees in term of size. In this section we explain how one can recover a short generator of a principal ideal from an arbitrary generator in the particular case of cyclotomic number fields.

Definition. The Short Generator Problem α-SGP consists in:

– Given an element $h \in K^\times$, generating an ideal $\mathfrak{I} = hO_K$
– Find a generator $g \in K^\times$ of $\mathfrak{I}$ of small Euclidean length: $\|g\| \leq \alpha \cdot \lambda_1(\mathfrak{I})$.

Remembering that h and g generate the same ideal if and only if $g = hu$ for some unit $u \in O_K$, the idea consists in rephrasing this problem as a close vector problem using the logarithmic embedding: indeed, we have that $\text{Log}\, g = \text{Log}\, h + \text{Log}\, u$ must belong to the lattice coset $\log h + \text{Log}\, O_K^\times$. Furthermore, up to appropriate rescaling, it can be proved that the length of g is related to the length of its logarithmic embedding $\text{Log}\, g$. Minimizing g can therefore be rephrased as finding a unit u such that $\text{Log}\, u$ is close to $-\text{Log}\, h$, in other words solving a Close Vector Problem over Dirichlet's logarithmic unit lattice $\text{Log}\, O_K^\times$.

Moreover, in certain cryptosystems, we have additional constraints on the ideal $\mathfrak{I}$, ensuring that a unusually short generator g exists (which is used as the secret key). This suggested that the Close Vector Problem may actually become a BDD problem [5]. And indeed, experiments confirmed that this BDD is easily solved in practice. Running such experiments requires knowing the group of units $O_K^\times$. Fortunately, in the case of cyclotomic number fields, there are some well known units — that very often generate the whole group $O_K^\times$ — namely, the cyclotomic units:

$$\{\zeta\} \cup \left\{ u_i = \frac{1 - \zeta^i}{1 - \zeta} \mid i \in (\mathbb{Z}/m\mathbb{Z})^\times \right\}.$$

Geometric analysis of the cyclotomic units
The fact that the attack works in practice suggests that the matrix $\mathbf{U} = (\text{Log}\, u_i)_{i \in (\mathbb{Z}/m\mathbb{Z})^\times}$ forms a good basis for BDD. Yet it is not so straightforward to prove it: recalling the first section, one needs to show that the *dual* vectors $\mathbf{u}_i^\vee$ are short. To proceed with the analysis, Cramer et al. [6] instead considered the related matrix $\mathbf{M} = (\text{Log}(1 - \zeta^{i^{-1}}))_i$, where

$$\mathbf{M}_{i,j} = \log |\psi_j(1 - \zeta^{i^{-1}})| = \log |1 - \omega^{ji^{-1}}|$$

for indices i, j running over the group $G = (\mathbb{Z}/m\mathbb{Z})^\times$. Since this matrix is G-circulant, it can therefore be explicitly diagonalized, and a lower-bound on the diagonal coefficients will provide an upper-bound

on the length of the dual vectors $\| m_i^\vee \|$. The eigenvalue λ_χ associated to the character $\chi : G \to \mathbb{C}$ of G is given by:

$$\lambda_\chi = \sum_{i \in G} \chi(i) \log | 1 - \omega^i |.$$

Using classical techniques of analytic number theory (in this case, the Taylor series of $\log$, and separation of Gauss sums), the above formula can be massaged to

$$\lambda_\chi = \sqrt{f_\chi} \cdot L(\chi, 1),$$

where f_χ is the conductor of χ, and L denotes Dirichlet's L-series. Lower bounds on L-series at 1 have a very long history and play a crucial role in the study of the distribution of prime numbers. For example, Landau proved that $L(\chi, 1) \geq 1/O(\log f_\chi)$ for non-quadratic characters. With more effort, Cramer et al. [6] conclude on an upper bound on $\| m_i^\vee \|$ and then on $\| u_i^\vee \|$.

Extension to the worst-case
In addition, the article [6] also covers the performance of this strategy in the worst-case, that is when there is no guarantee of existence of a particularly short generator g, by quantifying how good is the basis $\mathbf{U}$ to solve CVP. This analysis is somewhat easier as it concerns the length of the primal vectors, whose length can be bounded using the following finite integral:

$$\int_0^1 (\log | 1 - \exp(2\imath\pi x) |)^2 dx < \infty.$$

Further efforts lead to a classical probabilistic polynomial time algorithm that solves α-SGP for sub-exponential $\alpha = \exp(\tilde{O}(\sqrt{n}))$. Combined with the previous algorithm of Biasse and Song for PIP [4], this provides a solution to α-SVP in quantum polynomial time over principal ideal lattices in the worst case, outperforming the best known generic algorithms LLL and BKZ.

Finally, it is also shown in [6] that this result is roughly optimal: there exist many ideals for which the shortest generator g is much larger than the shortest vector, by a factor $\exp(\tilde{O}(\sqrt{n}))$. This is established by a lower bound on the covering radius of the lattice $\mathrm{Log}\, O_K^\times$. Lowering the SVP approximation factor reachable in polynomial time will necessarily require algorithms that are not limited to finding generators.

Short vectors in arbitrary ideals
Considering that the previous result only applies to principal ideals, which form a

very small fraction of all ideals, it is unclear whether this previous result has an impact on lattice-based cryptography beyond the few aforementioned atypical cryptosystems [5, 17]. Indeed, most schemes are instead based on worst-case problems, and are not affected by the presence of a small fraction of weak ideal lattices.

The obstacle ahead to attack non-principal ideals is the class group, the quotient of all ideals by the principal ones:

$$\mathrm{Cl}_K = \mathcal{F}_K / \mathcal{P}_K.$$

The class of an ideal $\mathfrak{a} \subset K$ in this quotient is denoted $[\mathfrak{a}] \in \mathrm{Cl}_K$, and the neutral element is $[O_K]$ (the class of principal ideals). This quotient is always finite, and in the case of cyclotomic number fields, it has size about $\#\mathrm{Cl}_K = 2^{\Theta(n \log n)}$: the fraction of principal ideals is super-exponentially small.

To generalize the previous result to non-principal ideals, the natural strategy consists in trying to find sub-ideals that are principal. More formally, Cramer, Ducas and Wesolowski [7] define the following problem:

Definition. The Close Principal Multiple problem with approximation factor F (F-CPM) is defined as:
- Given an ideal $\mathfrak{a} \subset K$,
- Find an *integral* ideal $\mathfrak{b}$ such that $\mathfrak{a}\mathfrak{b}$ is principal (i.e. $[\mathfrak{a}\mathfrak{b}] = [O_K]$) and such that $\mathfrak{b}$ is a dense ideal: $N\mathfrak{b} \leq F$,

where $N\mathfrak{b} := \#(O_K/\mathfrak{b})$ denotes the algebraic norm (i.e. the sparsity) of the ideal $\mathfrak{b}$.

Combining algorithm for F-CPM with the previous algorithms provides solution to α-SVP over non-principal ideals within approximation factor:

$$\alpha = F^{1/n} \cdot \exp(\tilde{O}(\sqrt{n})).$$

In this section, we will sketch how this CPM problem was solved for $F = \exp(\tilde{O}(n^{3/2}))$, which leads to a similar SVP approximation factor $\alpha = \exp(\tilde{O}(\sqrt{n}))$ as in the principal case. Consider a factor basis of prime ideals $\mathfrak{B} = \{\mathfrak{p}_1, \ldots, \mathfrak{p}_d\}$ that are dense ($N\mathfrak{p}_i \leq n^{O(1)}$) and the morphism:

$$\phi : \mathbb{Z}^d \to \mathrm{Cl}_K, \quad \phi(\mathbf{e}) = \left[\prod \mathfrak{p}_i^{e_i} \right].$$

Assuming that ϕ is surjective, we can rephrase CPM as a CVP problem in the lattice of class relations $\Lambda = \ker \phi$. Indeed, consider $\mathbf{v} \in \phi^{-1}([\mathfrak{a}])$, and $\mathbf{c} \in \Lambda$ such that $\mathbf{c} - \mathbf{v}$

is small, and all positive. Then $\mathfrak{b} = \prod \mathfrak{p}_i^{c_i - v_i}$ will be an appropriate solution to CPM, up to a factor $F = \prod (N\mathfrak{p}_i)^{v_i - c_i} = n^{O(\| v - c \|_1)}$ where $\| x \|_1 = \sum |x_i|$ denotes the ℓ_1-norm of $\mathbf{x}$.

In general, there is no reason why CVP should be easy in the lattice Λ, which is not even explicitly known. Yet, by choosing an appropriate factor basis, namely, a basis composed of all the Galois conjugates of a single ideal, one can on the contrary get a very explicit description of the lattice Λ thanks to the classical theorem of Stickelberger [19]. It turns out that one may easily explicitly construct a short basis of Λ. Again, this overview is highly simplified, and hides several technicalities, see [7].

Conclusion and open questions
There remain serious obstacles for this approach to attack ideal lattice-based cryptosystems. First the approximation factor $\alpha = \exp(\tilde{O}(\sqrt{n}))$ is too large to affect cryptographic schemes. Second, these algorithms are limited to ideal lattices (i.e. module lattices of rank 1), while most cryptosystems in fact use module lattices of rank 2 or more.

Nevertheless, these recent works questioned our understanding of the hardness of lattice problems when using special classes of lattices. We now know of a specialized algorithm for relevant classes of structured lattices that outperforms generic ones (see Figure 1). Alternatives to cyclotomics ideal lattices are already being studied [2, 3, 13] from various point of view: complexity theory, concrete cryptographic design and cryptanalysis.

There are many cases where the volume of the log-unit lattice (the regulator) and the lattice of class relation (the class number) is well understood. We have seen here that in the case of cyclotomic number field, much more can be said about those lattices (known good bases, covering radius,...). Generalizing this geometric analysis to other number fields seems to be an interesting mathematical problem, with potential cryptanalytic implications.   ⟵⋯

References

1. M. Ajtai, Generating hard instances of the short basis problem, *ICALP 1999*.

2. J. Bauch, D. J. Bernstein, H. de Valence, T. Lange and C. van Vredendaal, Short generators without quantum computers: The case of multiquadratics, *Eurocrypt 2017*.

3. D. J. Bernstein, C. Chuengsatiansup, T. Lange and C. van Vredendaal, NTRU Prime, Preprint 2016.

4. J.-F. Biasse and F. Song, A polynomial time quantum algorithm for computing class groups and solving the principal ideal problem in arbitrary degree number fields, *SODA 2016*.

5. Peter Campbell, Michael Groves and Dan Shepherd, Soliloquy: A cautionary tale, *ETSI 2nd Quantum-Safe Crypto Workshop*, 2014.

6. R. Cramer, L. Ducas, C. Peikert and O. Regev, Recovering short generators of principal ideals in cyclotomic rings, *Eurocrypt 2016*.

7. R. Cramer, L. Ducas and B. Wesolowski, Short Stickelberger class relations and application to ideal-SVP, *Eurocrypt 2017*.

8. K. Eisenträger, S. Hallgren, A. Kitaev and F. Song, A quantum algorithm for computing the unit group of an arbitrary degree number field, *STOC 2014*.

9. J. Hoffstein, J. Pipher and J. H. Silverman, NTRUSIGN: Digital signatures using the NTRU lattice, *CT-RSA 2003*.

10. P. Klein, Finding the closest lattice vector when it's unusually close, *SODA 2000*.

11. A. K. Lenstra, H. W. Lenstra and L. Lovàsz, Factoring polynomials with rational coefficients, *Mathematische Annalen* 261(4) (1982).

12. M. Daniele, The shortest vector in a lattice is hard to approximate to within some constant, *SIAM Journal on Computing* 30(6) (2001).

13. C. Peikert, O. Regev and N. Stephens-Davidowitz, Pseudorandomness of Ring-LWE for any ring and modulus, *STOC 2017*.

14. O. Regev, On lattices, learning with errors, random linear codes, and cryptography, *STOC 2005*.

15. C.-P. Schnorr, A hierarchy of polynomial time lattice basis reduction algorithms, *Theoretical Computer Science* 53(2) (1987).

16. P. W. Shor, Algorithms for quantum computation: Discrete logarithms and factoring, *FOCS 1994*.

17. N. P. Smart and F. Vercauteren, Fully homomorphic encryption with relatively small key and ciphertext sizes, *PKC 2010*.

18. D. Stehlé, R. Steinfeld, K. Tanaka and K. Xagawa, Efficient public key encryption based on ideal lattices, *Asiacrypt 2009*.

19. L. C. Washington, *Introduction to Cyclotomic Fields*, Springer, 1997.