Efficient Public Key Encryption Based on Ideal Lattices (Extended Abstract)

Damien Stehlé[1,2], Ron Steinfeld[2], Keisuke Tanaka[3], and Keita Xagawa[3]

[1] CNRS/Department of Mathematics and Statistics, University of Sydney NSW 2006, Australia.
[2] Centre for Advanced Computing - Algorithms and Cryptography, Department of Computing, Macquarie University, NSW 2109, Australia
[3] Department of Mathematical and Computing Sciences, Tokyo Institute of Technology, Japan

Abstract. We describe public key encryption schemes with security provably based on the worst case hardness of the approximate Shortest Vector Problem in some structured lattices, called ideal lattices. Under the assumption that the latter is exponentially hard to solve even with a quantum computer, we achieve CPA-security against subexponential attacks, with (quasi-)optimal asymptotic performance: if n is the security parameter, both keys are of bit-length $\widetilde{O}(n)$ and the amortized costs of both encryption and decryption are $\widetilde{O}(1)$ per message bit. Our construction adapts the trapdoor one-way function of Gentry *et al.* (STOC'08), based on the Learning With Errors problem, to structured lattices. Our main technical tools are an adaptation of Ajtai's trapdoor key generation algorithm (ICALP'99) and a re-interpretation of Regev's quantum reduction between the Bounded Distance Decoding problem and sampling short lattice vectors.

1 Introduction

Lattice-based cryptography has been rapidly developing in the last few years, inspired by the breakthrough result of Ajtai in 1996 [1], who constructed a one-way function with average-case security provably related to the worst-case complexity of hard lattice problems. The attractiveness of lattice-based cryptography stems from its provable security guarantees, well studied theoretical underpinnings, simplicity and potential efficiency (Ajtai's one-way function is a matrix-vector multiplication over a small finite field), and also the apparent security against quantum attacks. The main complexity assumption is the hardness of approximate versions of the Shortest Vector Problem (SVP). The $\mathrm{GapSVP}_{\gamma(n)}$ problem consists in, given a lattice of dimension n and a scalar d, replying YES if there exists a non-zero lattice vector of norm $\leq d$ and NO if all non-zero lattice vectors have norm $\geq \gamma(n)d$. The complexity of $\mathrm{GapSVP}_{\gamma(n)}$ increases with n, but decreases with $\gamma(n)$. Although the latter is believed to be exponential in n for any polynomial $\gamma(n)$, minimizing the degree of $\gamma(n)$ is very important in practice, to allow the use of a practical dimension n for a given security level.

LATTICE-BASED PUBLIC KEY ENCRYPTION. The first provably secure lattice-based cryptosystem was proposed by Ajtai and Dwork [3], and relied on a variant of GapSVP in arbitrary lattices (it is now known to also rely on GapSVP [19]). Subsequent works proposed more efficient alternatives [33, 30, 9, 28]. The current state of the art [9, 28] is a scheme with public/private key length $\widetilde{O}(n^2)$ and encryption/decryption throughput of $\widetilde{O}(n)$ bit operations per message bit. Its security relies on the quantum worst-case hardness of $\text{GapSVP}_{\widetilde{O}(n^{1.5})}$ in arbitrary lattices. The security can be de-quantumized at the expense of both increasing $\gamma(n)$ and decreasing the efficiency, or relying on a new and less studied problem [28]. In parallel to the provably secure schemes, there have also been heuristic proposals [11, 12]. In particular, unlike the above schemes which use unstructured random lattices, the NTRU encryption scheme [12] exploits the properties of *structured* lattices to achieve high efficiency with respect to key length ($\widetilde{O}(n)$ bits) and encryption/decryption cost ($\widetilde{O}(1)$ bit operation per message bit). Unfortunately, its security remains heuristic and it was an important open challenge to provide a provably secure scheme with comparable efficiency.

PROVABLY SECURE SCHEMES FROM IDEAL LATTICES. Micciancio [20] introduced the class of structured *cyclic* lattices, which correspond to ideals in polynomial rings $\mathbb{Z}[x]/(x^n - 1)$, and presented the first provably secure one-way function based on the worst-case hardness of the restriction of $\mathcal{P}oly(n)$-SVP to cyclic lattices. (The problem γ-SVP consists in computing a non-zero vector of a given lattice, whose norm is no more than γ times larger than the norm of a shortest non-zero lattice vector.) At the same time, thanks to its algebraic structure, this one-way function enjoys high efficiency comparable to the NTRU scheme ($\widetilde{O}(n)$ evaluation time and storage cost). Subsequently, Lyubashevsky and Micciancio [17] and independently Peikert and Rosen [29] showed how to modify Micciancio's function to construct an efficient and provably secure collision resistant hash function. For this, they introduced the more general class of *ideal* lattices, which correspond to ideals in polynomial rings $\mathbb{Z}[x]/f(x)$. The collision resistance relies on the hardness of the restriction of $\mathcal{P}oly(n)$-SVP to ideal lattices (called $\mathcal{P}oly(n)$-Ideal-SVP). The average-case collision-finding problem is a natural computational problem called Ideal-SIS, which has been shown to be as hard as the worst-case instances of Ideal-SVP. Provably secure efficient signature schemes from ideal lattices have also been proposed [18, 15, 16, 14], but constructing efficient provably secure public key encryption from ideal lattices was an interesting open problem.

OUR RESULTS. We describe the first provably CPA-secure public key encryption scheme whose security relies on the hardness of the worst-case instances of $\widetilde{O}(n^2)$-Ideal-SVP against subexponential quantum attacks. It achieves asymptotically optimal efficiency: the public/private key length is $\widetilde{O}(n)$ bits and the amortized encryption/decryption cost is $\widetilde{O}(1)$ bit operations per message bit (encrypting $\widetilde{\Omega}(n)$ bits at once, at a $\widetilde{O}(n)$ cost). Our security assumption is that $\widetilde{O}(n^2)$-Ideal-SVP cannot be solved by any subexponential time quantum algorithm, which is reasonable given the state-of-the art lattice algorithms [36]. Note that this is stronger than standard public key cryptography security as-

sumptions. On the other hand, contrary to most of public key cryptography, lattice-based cryptography allows security against subexponential quantum attacks. Our main technical tool is a re-interpretation of Regev's quantum reduction [33] between the Bounded Distance Decoding problem (BDD) and sampling short lattice vectors. Also, by adapting Ajtai's trapdoor generation algorithm [2] (or more precisely its recent improvement by Alwen and Peikert [5]) to structured ideal lattices, we are able to construct efficient provably secure trapdoor signatures, ID-based identification schemes, CCA-secure encryption and ID-based encryption. We think these techniques are very likely to find further applications.

Most of the cryptosystems based on general lattices [33, 30, 31, 9, 28] rely on the average-case hardness of the *Learning With Errors* (LWE) problem introduced in [33]. Our scheme is based on a structured variant of LWE, that we call Ideal-LWE. We introduce novel techniques to circumvent two main difficulties that arise from the restriction to ideal lattices. Firstly, the previous cryptosystems based on unstructured lattices all make use of Regev's worst-case to average-case classical reduction [33] from BDD to LWE (this is the *classical step* in the quantum reduction of [33] from SVP to LWE). This reduction exploits the unstructured-ness of the considered lattices, and does not seem to carry over to the structured lattices involved in Ideal-LWE. In particular, the probabilistic independence of the rows of the LWE matrices allows to consider a single row in [33, Cor. 3.10]. Secondly, the other ingredient used in previous cryptosystems, namely Regev's reduction [33] from the computational variant of LWE to its decisional variant, also seems to fail for Ideal-LWE: it relies on the probabilistic independence of the columns of the LWE matrices.

Our solution to the above difficulties avoids the *classical step* of the reduction from [33] altogether. Instead, we use the *quantum step* to construct a new quantum average-case reduction from SIS (the unstructured variant of Ideal-SIS) to LWE. It also works from Ideal-SIS to Ideal-LWE. Combined with the known reduction from worst-case Ideal-SVP to average-case Ideal-SIS [17], we obtain a quantum reduction from Ideal-SVP to Ideal-LWE. This shows the hardness of the computational variant of Ideal-LWE. Because we do not obtain the hardness of the decisional variant, we use a generic hardcore function to derive pseudorandom bits for encryption. This is why we need to assume the exponential hardness of SVP. The encryption scheme follows as an adaptation of [9, Sec. 7.1].

The main idea of our new quantum reduction from Ideal-SIS to Ideal-LWE is a re-interpretation of Regev's quantum step in [33]. The latter was presented as a worst-case quantum reduction from sampling short lattice vectors in a lattice L to solving BDD in the dual lattice $\hat{L}$. We observe that this reduction is actually stronger: it is an average-case reduction which works given an oracle for BDD in $\hat{L}$ with a normally distributed error vector. Also, as pointed out in [9], LWE can be seen as a BDD with a normally distributed error in a certain lattice whose dual is essentially the SIS lattice. This leads to our SIS to LWE reduction. Finally we show how to apply it to reduce Ideal-SIS to Ideal-LWE – this involves a probabilistic lower bound for the minimum of the Ideal-LWE lattice. We believe our new SIS to LWE reduction is of independent interest. Along with [22], it

provides an alternative to Regev's quantum reduction from GapSVP to LWE. Ours is weaker because the derived GapSVP factor increases with the number of LWE samples, but it has the advantage of carrying over to the ideal case. Also, when choosing practical parameters for lattice-based encryption (see, e.g., [23]), it is impractical to rely on the worst-case hardness of SVP. Instead, the practical average-case hardness of LWE is evaluated based on the best known attack which consists in solving SIS. Our reduction justifies this heuristic by showing that it is indeed necessary to (quantumly) break SIS in order to solve LWE.

ROAD-MAP. We provide some background in Section 2. Section 3 shows how to hide a trapdoor in the adaptation of SIS to ideal lattices. Section 4 contains the new reduction between SIS and LWE. Finally, in Section 5, we present our CPA-secure encryption scheme and briefly describe other cryptographic constructions.

NOTATION. Vectors will be denoted in bold. We denote by $\langle \cdot, \cdot \rangle$ and $\| \cdot \|$ the inner product and the Euclidean norm. We denote by $\rho_s(\boldsymbol{x})$ (resp. ν_s) the standard n-dimensional Gaussian function (resp. distribution) with center $\mathbf{0}$ and variance s, i.e., $\rho_s(\boldsymbol{x}) = \exp(-\pi\|\boldsymbol{x}\|^2/s^2)$ (resp. $\nu_s(\boldsymbol{x}) = \rho_s(\boldsymbol{x})/s^n$). We use the notations $\widetilde{O}(\cdot)$ and $\widetilde{\Omega}(\cdot)$ to hide poly-logarithmic factors. If D_1 and D_2 are two probability distributions over a discrete domain E, their statistical distance is $\Delta(D_1, D_2) = \frac{1}{2}\sum_{x\in E}|D_1(x) - D_2(x)|$. If a function f over a countable domain E takes non-negative real values, its sum over an arbitrary $F \subseteq E$ will be denoted by $f(F)$. If q is a prime number, we denote by $\mathbb{Z}_q$ the field of integers modulo q. We denote by Ψ_s the reduction modulo q of ν_s.

2 Reminders and Background Results on Lattices

We refer to [21] for a detailed introduction to the computational aspects of lattices. In the present section, we remind the reader very quickly some fundamental properties of lattices that we will need. We then introduce the so-called ideal lattices, and finally formally define some computational problems.

Euclidean lattices. An n-dimensional lattice L is the set of all integer linear combinations of some linearly independent vectors $\boldsymbol{b}_1, \ldots, \boldsymbol{b}_n \in \mathbb{R}^n$, i.e., $L = \sum \mathbb{Z}\boldsymbol{b}_i$. The $\boldsymbol{b}_i$'s are called a basis of L. The ith minimum $\lambda_i(L)$ is the smallest r such that L contains i linearly independent vectors of norms $\leq r$. We let $\lambda_1^\infty(L)$ denote the first minimum of L with respect to the infinity norm. If $B = (\boldsymbol{b}_1, \ldots, \boldsymbol{b}_n)$ is a basis, we define its norm by $\|B\| = \max\|\boldsymbol{b}_i\|$ and its fundamental parallelepiped by $P(B) = \{\sum_i c_i\boldsymbol{b}_i \mid \boldsymbol{c} \in [0,1)^n\}$. Given a basis B for lattice L and a vector $\boldsymbol{c} \in \mathbb{R}^n$, we define $\boldsymbol{c} \bmod L$ as the unique vector in $P(B)$ such that $\boldsymbol{c} - (\boldsymbol{c} \bmod L) \in L$ (the basis being implicit). For any lattice L and any $s > 0$, the sum $\rho_s(L)$ is finite. We define the lattice Gaussian distribution by $D_{L,s}(\boldsymbol{b}) = \frac{\rho_s(\boldsymbol{b})}{\rho_s(L)}$, for any $\boldsymbol{b} \in L$. If L is a lattice, its dual $\widehat{L}$ is the lattice $\{\widehat{\boldsymbol{b}} \in \mathbb{R}^n \mid \forall \boldsymbol{b} \in L, \langle\widehat{\boldsymbol{b}}, \boldsymbol{b}\rangle \in \mathbb{Z}\}$. We will use the following results.

Lemma 1 ([29, Lemma 2.11] and [27, Lemma 3.5]). *For any $\boldsymbol{x}$ in an n-dimensional lattice L and $s \geq 2\sqrt{\ln(10n)/\pi}/\lambda_1^\infty(\widehat{L})$, we have $D_{L,s}(\boldsymbol{x}) \leq 2^{-n+1}$.*

Lemma 2 ([22, Lemma 2.10]). *Given an n-dimensional lattice L, we have* $\Pr_{\boldsymbol{x}\sim D_{L,s}}[\|\boldsymbol{x}\| > s\sqrt{n}] \leq 2^{-n+1}$.

Ideal lattices. Ideal lattices are a subset of lattices with the computationally interesting property of being related to polynomials via structured matrices. The n-dimensional vector-matrix product costs $\widetilde{O}(n)$ arithmetic operations instead of $O(n^2)$. Let $f \in \mathbb{Z}[x]$ a monic degree n polynomial. For any $g \in \mathbb{Q}[x]$, there is a unique pair (q, r) with $\deg(r) < n$ and $g = qf + r$. We denote r by $g \bmod f$ and identify r with the vector $\boldsymbol{r} \in \mathbb{Q}^n$ of its coefficients. We define $\mathrm{rot}_f(r) \in \mathbb{Q}^{n\times n}$ as the matrix whose rows are the $x^i r(x) \bmod f(x)$'s, for $0 \leq i < n$. We extend that notation to the matrices A over $\mathbb{Q}[x]/f$, by applying rot_f component-wise. Note that $\mathrm{rot}_f(g_1)\mathrm{rot}_f(g_2) = \mathrm{rot}_f(g_1 g_2)$ for any $g_1, g_2 \in \mathbb{Q}[x]/f$. The strengths of our cryptographic constructions depend on the choice of f. Its quality is quantified by its expansion factor (we adapt the definition of [17] to the Euclidean norm):

$$\mathrm{EF}(f,k) = \max\left\{ \frac{\|g \bmod f\|}{\|g\|} \mid g \in \mathbb{Z}[x] \setminus \{0\} \text{ and } \deg(g) \leq k\,(\deg(f) - 1) \right\},$$

where we identified the polynomial $g \bmod f$ (resp. g) with the coefficients vector. Note that if $\deg(g) < n$, then $\|\mathrm{rot}_f(g)\| \leq \mathrm{EF}(f,2) \cdot \|g\|$. We will concentrate on the polynomials $x^{2^k} + 1$, although most of our results are more general. We recall some basic properties of $x^{2^k} + 1$ (see [7] for the last one).

Lemma 3. *Let $k \geq 0$ and $n = 2^k$. Then $f(x) = x^n + 1$ is irreducible in $\mathbb{Q}[x]$. Its expansion factor is $\leq \sqrt{2}$. Also, for any $g = \sum_{i<n} g_i x^i \in \mathbb{Q}[x]/f$, we have $\mathrm{rot}_f(g)^T = \mathrm{rot}_f(\bar{g})$ where $\bar{g} = g_0 - \sum_{1\leq i<n} g_{n-i} x^i$. Furthermore, if q is a prime such that $2n|(q-1)$, then f has n linear factors in $\mathbb{Z}_q[x]$. Finally, if $k \geq 2$ and q is a prime with $q \equiv 3 \bmod 8$, then $f = f_1 f_2 \bmod q$ where each f_i is irreducible in $\mathbb{Z}_q[x]$ and can be written $f_i = x^{n/2} + t_i x^{n/4} - 1$ with $t_i \in \mathbb{Z}_q$.*

Let I be an ideal of $\mathbb{Z}[x]/f$, i.e., a subset of $\mathbb{Z}[x]/f$ closed under addition and multiplication by any element of $\mathbb{Z}[x]/f$. It corresponds to a sublattice of $\mathbb{Z}^n$. An *f-ideal lattice* is a sublattice of $\mathbb{Z}^n$ that corresponds to an ideal $I \subseteq \mathbb{Z}[x]/f$.

Hard lattice problems. The most famous lattice problem is SVP. Given a basis of a lattice L, it aims at finding a shortest vector in $L \setminus \{\mathbf{0}\}$. It can be relaxed by asking for a non-zero vector that is no longer than $\gamma(n)$ times a solution to SVP, for a prescribed function $\gamma(\cdot)$. The best polynomial time algorithm [4, 35] solves γ-SVP only for a slightly subexponential γ. When γ is polynomial in n, then the most efficient algorithm [4] has an exponential worst-case complexity both in time and space. If we restrict the set of input lattices to ideal lattices, we obtain the problem Ideal-SVP (resp. γ-Ideal-SVP), which is implicitly parameterized by a sequence of polynomials f of growing degrees. No algorithm is known to perform non-negligibly better for Ideal-SVP than for SVP. It is believed that no subexponential quantum algorithm solves the computational variants of SVP or Ideal-SVP in the worst case. These worst-case problems can be reduced to the following average-case problems, introduced in [1] and [9].

Definition 1. *The Small Integer Solution problem with parameters $q(\cdot)$, $m(\cdot)$, $\beta(\cdot)$ ($\mathrm{SIS}_{q,m,\beta}$) is as follows: Given n and a matrix G sampled uniformly in $\mathbb{Z}_{q(n)}^{m(n)\times n}$, find $\boldsymbol{e} \in \mathbb{Z}^{m(n)} \setminus \{\mathbf{0}\}$ such that $\boldsymbol{e}^T G = \mathbf{0} \bmod q(n)$ (the modulus being taken component-wise) and $\|\boldsymbol{e}\| \leq \beta(n)$. The Ideal Small Integer Solution problem with parameters q, m, β and f ($\text{Ideal-SIS}^f_{q,m,\beta}$) is as follows: Given n and m polynomials $g_1, \ldots, g_m$ chosen uniformly and independently in $\mathbb{Z}_q[x]/f$, find $e_1, \ldots, e_m \in \mathbb{Z}[x]$ not all zero such that $\sum_{i \leq m} e_i g_i = 0$ in $\mathbb{Z}_q[x]/f$ and $\|\boldsymbol{e}\| \leq \beta$, where $\boldsymbol{e}$ is the vector obtained by concatenating the coefficients of the e_i's.*

The above problems can be interpreted as lattice problems. If $G \in \mathbb{Z}_q^{m\times n}$, then the set $G^\perp = \{\boldsymbol{b} \in \mathbb{Z}^m \mid \boldsymbol{b}^T G = \mathbf{0} \bmod q\}$ is an m-dimensional lattice and solving SIS corresponds to finding a short non-zero vector in it. Similarly, Ideal-SIS consists in finding a small non-zero element in the $\mathbb{Z}[x]/f$-module $M^\perp(\boldsymbol{g}) = \{\boldsymbol{b} \in (\mathbb{Z}[x]/f)^m \mid \langle \boldsymbol{b}, \boldsymbol{g}\rangle = 0 \bmod q\}$, where $\boldsymbol{g} = (g_1, \ldots, g_m)$. It can be seen as a lattice problem by applying the rot_f operator. Note that the m of SIS is n times larger than the m of Ideal-SIS. Lyubashevsky and Micciancio [17] reduced Ideal-SVP to Ideal-SIS. The approximation factors in [17] are given in terms of the infinity norm. For our purposes, it is more natural to use the Euclidean norm. To avoid losing a $\sqrt{n}$ factor by simply applying the norm equivalence formula, we modify the proof of [17]. We also adapt it to handle the case where the Ideal-SIS solver has a subexponentially small success probability, at the cost of an additional factor of $\widetilde{O}(\sqrt{n})$ in the SVP approximation factor.

Theorem 1. *Suppose that f is irreducible over $\mathbb{Q}$. Let $m = \mathcal{P}oly(n)$ and $q = \widetilde{\Omega}(\mathrm{EF}(f,3)\beta m^2 n)$ be integers. A polynomial-time (resp. subexponential-time) algorithm solving $\text{Ideal-SIS}^f_{q,m,\beta}$ with probability $1/\mathcal{P}oly(n)$ (resp. $2^{-o(n)}$) can be used to solve γ-Ideal-SVP in polynomial-time (resp. subexponential-time) with $\gamma = \widetilde{O}(\mathrm{EF}^2(f,2)\beta m n^{1/2})$ (resp. $\gamma = \widetilde{O}(\mathrm{EF}^2(f,2)\beta m n)$).*

The problem LWE is dual to SIS in the sense that if $G \in \mathbb{Z}_q^{m\times n}$ is the SIS-matrix, then LWE involves the dual of the lattice $G^\perp$. We have $\widehat{G^\perp} = \frac{1}{q}L(G)$ where $L(G) = \{\boldsymbol{b} \in \mathbb{Z}^m \mid \exists \boldsymbol{s} \in \mathbb{Z}_q^n, G\boldsymbol{s} = \boldsymbol{b} \bmod q\}$.

Definition 2. *The Learning With Errors problem with parameters q, m and a distribution χ on $\mathbb{R}/[0,q)$ ($\mathrm{LWE}_{q,m;\chi}$) is as follows: Given n, a matrix $G \in \mathbb{Z}_q^{m\times n}$ sampled uniformly at random and $G\boldsymbol{s} + \boldsymbol{e} \in (\mathbb{R}/[0,q))^n$, where $\boldsymbol{s} \in \mathbb{Z}_q^n$ is chosen uniformly at random and the coordinates of $\boldsymbol{e} \in (\mathbb{R}/[0,q))^m$ are independently sampled from χ, find $\boldsymbol{s}$. The Ideal Learning With Errors problem with parameters q, m, a distribution χ on $\mathbb{R}/[0,q)$ and f ($\text{Ideal-LWE}^f_{m,q;\chi}$) is the same as above, except that $G = \mathrm{rot}_f(\boldsymbol{g})$ with $\boldsymbol{g}$ chosen uniformly in $(\mathbb{Z}_q[x]/f)^m$.*

We will use the following results on the LWE and Ideal-LWE lattices.

Lemma 4. *Let n, m and q be integers with q prime, $m \geq 5n\log q$ and $n \geq 10$. Then for all but a fraction $\leq q^{-n}$ of the G's in $\mathbb{Z}_q^{m\times n}$, we have $\lambda_1^\infty(L(G)) \geq q/4$ and $\lambda_1(L(G)) \geq 0.07\sqrt{mq}$.*

Lemma 5. *Let n, m and q be integers with $q = 3 \bmod 4$ prime and $m \geq 41 \log q$ and $n = 2^k \geq 32$. Then for all but a fraction $\leq q^{-n}$ of the $\boldsymbol{g}$'s in $(\mathbb{Z}_q[x]/f)^m$, we have $\lambda_1^\infty(L(\mathrm{rot}_f(\boldsymbol{g}))) \geq q/4$ and $\lambda_1(L(\mathrm{rot}_f(\boldsymbol{g}))) \geq 0.017\sqrt{mn}q$.*

3 Hiding a Trapdoor in Ideal-SIS

In this section we show how to hide a trapdoor in the problem Ideal-SIS. Ajtai [2] showed how to simultaneously generate a (SIS) matrix $A \in \mathbb{Z}_q^{m \times n}$ and a (trapdoor) basis $S = (\boldsymbol{s}_1, \ldots, \boldsymbol{s}_m) \in \mathbb{Z}^{m \times m}$ of the lattice $A^\perp = \{\boldsymbol{b} \in \mathbb{Z}^m : \boldsymbol{b}^T A = \boldsymbol{0} \bmod q\}$, with the following properties:

1. The distribution of A is close to the uniform distribution over $\mathbb{Z}_q^{m \times n}$.
2. The basis vectors $\boldsymbol{s}_1, \ldots, \boldsymbol{s}_m$ are short.

Recently, Alwen and Peikert [5] improved Ajtai's construction in the sense that the created basis has shorter vectors: $\|S\| = \widetilde{O}(n \log q)$ with $m = \Omega(n \log q)$ and overwhelming probability and $\|S\| = O(\sqrt{n \log q})$ with $m = \Omega(n \log^2 q)$. We modify both constructions to obtain a trapdoor generation algorithm for the problem Ideal-SIS, with a resulting basis whose norm is as small as the one of [5].

Before describing the construction, we notice that the construction of [5] relies on the Hermite Normal Form (HNF), but that here there is no Hermite Normal Form for the rings under scope. We circumvent this issue by showing that except in negligibly rare cases we may use a matrix which is HNF-like.

Theorem 2. *There exists a probabilistic polynomial time algorithm with the following properties. It takes as inputs n, σ, r, an odd prime q, and integers m_1, m_2. It also takes as input a degree n polynomial $f \in \mathbb{Z}[x]$ and random polynomials $\boldsymbol{a}_1 \in (\mathbb{Z}_q[x]/f)^{m_1}$. We let $f = \prod_{i \leq t} f_i$ be the factorization of f over $\mathbb{Z}_q$. We let $\kappa = \lceil 1 + \log q \rceil$, $\Delta = \left(\prod_{i \leq t} \left(1 + \left(\frac{q}{3^r}\right)^{\deg f_i}\right) - 1\right)^{1/2}$ and $m = m_1 + m_2$. The algorithm succeeds with probability $\geq 1 - p$ over $\boldsymbol{a}_1$, where $p = (1 - \prod_{i \leq t}(1 - q^{-\deg f_i}))^\sigma$. When it does, it returns $\boldsymbol{a} = \begin{pmatrix} \boldsymbol{a}_1 \\ \boldsymbol{a}_2 \end{pmatrix} \in (\mathbb{Z}_q[x]/f)^m$ and a basis S of the lattice $\mathrm{rot}_f(\boldsymbol{a})^\perp$, such that:*

1. *The distance to uniformity of $\boldsymbol{a}$ is at most $p + m_2 \Delta$.*
2. *The quality of S is as follows:*
   - *If $m_1 \geq \max\{\sigma, \kappa, r\}$ and $m_2 \geq \kappa$, then $\|S\| \leq \mathrm{EF}(f, 2) \cdot \sqrt{2}\kappa r^{1/2} n^{3/2}$. Additionally, $\|S\| \leq \mathrm{EF}(f, 2)\sqrt{3a\kappa r \cdot n}$ with probability $1 - 2^{-a + O(\log n m_1 r)}$ for a super-logarithmic function $a = a(n) = \omega(\log n)$.*
   - *If $m_1 \geq \max\{\sigma, \kappa, r\}$ and $m_2 \geq \kappa m_1$, then $\|S\| \leq \mathrm{EF}(f, 2)(4\sqrt{nr} + 3)$.*
3. *In particular, for $f = x^{2^k} + 1$ with $k \geq 2$ and a prime q with $q \equiv 3 \bmod 8$, the following holds:*
   - *We can set $\sigma = 1$ and $r = \lceil 1 + \log_3 q \rceil$. Then, the error probability is $p = q^{-\Omega(n)}$ and the parameter Δ is $2^{-\Omega(n)}$.*
   - *If $m_1, m_2 \geq \kappa$, then $\|S\| \leq \sqrt{6a\kappa r} \cdot n = O(\sqrt{a} n \log q)$ with probability $1 - 2^{-a + O(\log n m_1 r)}$ for a super-logarithmic function $a = a(n) = \omega(\log n)$.*

– *If* $m_1 \geq \kappa$ *and* $m_2 \geq \kappa m_1$, *then* $\|S\| \leq \sqrt{2}(4\sqrt{nr}+3) = O(\sqrt{n \log q})$.

In the rest of this section, we only describe the analog of the second construction of Alwen and Peikert, i.e., the case $m_2 \geq \kappa m_1$, due to lack of space.

3.1 A trapdoor for Ideal-SIS

We now construct the trapdoor for Ideal-SIS. More precisely, we want to simultaneously construct a uniform $\boldsymbol{a} \in \mathcal{R}^m$ with $\mathcal{R} = \mathbb{Z}_q[x]/f$, and a small basis S of the lattice $A^\perp$ where $A = \mathrm{rot}_f(\boldsymbol{a})$. For this, it suffices to find a basis of the module $M^\perp(\boldsymbol{a}) = \{\boldsymbol{y} \in \mathcal{R}_0^m \mid \langle \boldsymbol{y}, \boldsymbol{a} \rangle \equiv 0 \bmod q\}$, with $\mathcal{R}_0 = \mathbb{Z}[x]/f$.

The principle of the design. In the following, for two matrices X and Y, $[X|Y]$ denotes the concatenation of the columns of X followed by Y and $[X;Y]$ denotes the concatenation of the rows of X and the rows of Y.

We mainly follow the Alwen-Peikert construction. Let $m_1 \geq \sigma, r$. Let us assume that we generate random polynomials $A_1 = [a_1, \ldots, a_{m_1}]^T \in \mathcal{R}^{m_1 \times 1}$. We will construct a random matrix $A_2 \in \mathcal{R}^{m_2 \times 1}$ with a structured matrix $S \in \mathcal{R}_0^{m \times m}$ such that $SA = 0$ and S is a basis of the module $M^\perp(\boldsymbol{a})$, where $A = [A_1; A_2]$. We first construct an HNF-like basis F of the module $M^\perp(\boldsymbol{a})$ with A. Next, we construct a unimodular matrix Q such that $S = QF$ is a short basis of the module. More precisely, S has the following form:

$$S = \begin{bmatrix} V & P \\ D & B \end{bmatrix} = \underbrace{\begin{bmatrix} -I_{m_1} & P \\ 0 & B \end{bmatrix}}_{Q} \cdot \underbrace{\begin{bmatrix} H & 0 \\ U & I_{m_2} \end{bmatrix}}_{F}.$$

Note that, by setting B lower triangular with diagonal coefficients equal to 1, the matrix Q is unimodular.

In this design principle, we want $FA = 0$. Hence, we should set

$$HA_1 = 0 \text{ and } A_2 = -UA_1.$$

Notice that, in order to prove that F is a basis of $A^\perp$, it suffices to show that H is a basis of $A_1^\perp$. The first equation is satisfied by setting H be an HNF-like matrix (see below). By setting $U = G + R$, with G to be defined later on and R a random matrix, we have that A_2 is almost uniformly random in $\mathcal{R}$ by Micciancio's regularity lemma (Lemma 6). More precisely, the i-th row of R is chosen from $(\{-1, 0, 1\}^n)^r \times (\{0\}^n)^{m_1 - r}$.

Lemma 6 (Adapted from [20, Th. 4.2]). *Let* $\mathbb{F}$ *be a finite field and* $f \in \mathbb{F}[x]$ *be monic and of degree* $n > 0$. *Let* R *be the ring* $\mathbb{F}[x]/f$. *Let* $D \subseteq \mathbb{F}$ *and* $r > 0$. *For* $a_1, \ldots, a_r \in R$, *we denote by* $H(a_1, \ldots, a_r)$ *the random variable* $\sum_{i \leq r} b_i a_i \in R$ *where the* b_i*'s are degree* $< n$ *polynomials with coefficients chosen independently and uniformly in* D. *If* $U_1, \ldots, U_r$ *are independent uniform random variables in* R, *then the statistical distance to uniformity*

of $(U_1, \ldots, U_r, H(U_1, \ldots, U_r))$ *is below:*

$$\frac{1}{2}\sqrt{\prod_{i \le t}\left(1 + \left(\frac{|\mathbb{F}|}{|D|^r}\right)^{\deg f_i}\right) - 1},$$

where $f = \prod_{i \le t} f_i$ *is the factorization of* f *over* $\mathbb{F}$.

We show below how to choose P and G such that $PG = H - I_{m_1}$. With this relation, the design principle form of S therefore implies that $V = -H + P(G + R) = PR - I_{m_1}$, and $D = B(G + R)$. Our constructions for P, G, B also ensure that P, B and BG have 'small' entries so that S has 'small' entries.

A construction of H without HNF. We start with how to construct H for $A_1 = [a_1, \ldots, a_{m_1}]^T \in \mathcal{R}^{m_1 \times 1}$. Since $m_1 \ge \max\{\sigma, \kappa, r\}$, we have $a_{i^*} \in \mathcal{R}^*$ for some index i^* with probability at least $1 - p$, where $\mathcal{R}^*$ denotes the set of invertible elements of $\mathcal{R}$. For now, we set $i^* = 1$ for simplicity. Using this a_{i^*}, we can construct an HNF-like matrix H: the first row is $q\boldsymbol{e}_1$ and the i-th row is $h_i\boldsymbol{e}_1 + \boldsymbol{e}_i$ for $i = 2, \ldots, m_1$, where $\boldsymbol{e}_i$ is a row vector in $\mathcal{R}_0^{m_1}$ such that the i-the element is 1 and others are 0, and $h_i = -a_i \cdot a_1^{-1} \bmod q$ such that $h_i \in [0, q)^n$. Let $\boldsymbol{h}_i$ denote the i-th row of H. By the definition of H, $H \cdot A_1 \equiv 0 \bmod q$. Thus, each row vector $\boldsymbol{h}_i$ is in $M^{\perp}(\boldsymbol{a}_1)$, where $\boldsymbol{a}_1 = A_1$. It is obvious that $\boldsymbol{h}_1, \ldots, \boldsymbol{h}_{m_1}$ are linearly independent over $\mathcal{R}_0$. Hence, we need to only show that H is indeed the basis of $M^{\perp}(\boldsymbol{a}_1)$, but this is a routine work.

Next, we consider the case where $i^* \neq 1$. In this case, we swap rows 1 and i^* of A_1 so that $a_1 \in \mathcal{R}^*$, and call it A_1'. Applying the method above, we get a basis H' of $\Lambda^{\perp}(A_1')$. By swapping columns 1 and i^* and rows 1 and i^* of H', we get a basis H of $\Lambda^{\perp}(A_1)$. In the following, we denote by i^* the index i such that $a_i \in \mathcal{R}^*$ and $h_{i,i} = q$. Note that our strategy fails if there is no index i such that $a_i \in \mathcal{R}^*$: this is not an issue, as this occurs only with small probability.

Preliminaries of the construction. Hereafter, we set $W = BG$. We often use the matrix $T_\kappa = (t_{i,j}) \in \mathcal{R}_0^{\kappa \times \kappa}$, where $t_{i,i} = 1$, $t_{i+1,i} = -2$, and all other $t_{i,j}$'s are 0. Notice that the i-th row of T_κ^{-1} is $(2^{i-1}, 2^{i-2}, \ldots, 1, 0, \ldots, 0) \in \mathcal{R}_0^{\kappa}$.

3.2 An analogue to the second Alwen-Peikert construction

The idea of the second construction in [5] is to have G contain the rows of $H - I_{m_1}$. This helps decrease the norms of the rows of P and V. To do so, we define $B = \mathrm{diag}(T_\kappa, \ldots, T_\kappa, I_{m_2 - m_1\kappa})$. Note that $B^{-1} = \mathrm{diag}(T_\kappa^{-1}, \ldots, T_\kappa^{-1}, I_{m_2 - m_1\kappa})$.

Let $\boldsymbol{h}_j'$ denote the j-th row of $H - I_{m_1}$. Let $W = [W_1; W_2; \ldots; W_{m_1}; 0]$, where $W_j = [\boldsymbol{w}_{j,\kappa}; \ldots; \boldsymbol{w}_{j,1}] \in \mathcal{R}_0^{\kappa \times m_1}$. We compute the $\boldsymbol{w}_{j,k}$'s such that $\boldsymbol{h}_j' = \sum_k 2^{k-1} \cdot \boldsymbol{w}_{j,k}$ and the components of all $\boldsymbol{w}_{j,k}$'s are polynomials with coefficients in $\{0, 1\}$. By this construction, $T_\kappa^{-1} \cdot W_j$ contains $\boldsymbol{h}_j'$ in the last row. Then, $G = B^{-1} \cdot W$ contains rows $\boldsymbol{h}_j'$ for $j = 1, \ldots, m_1$. The matrix $P = [\boldsymbol{p}_1; \ldots; \boldsymbol{p}_{m_1}]$ picks all rows $\boldsymbol{h}_1', \ldots, \boldsymbol{h}_{m_1}'$ in G by setting $\boldsymbol{p}_j = \boldsymbol{e}_{\kappa j} \in \mathcal{R}_0^{m_2}$.

The norm of S is $\max\{\|S_1\|, \|S_2\|\}$, where $S_1 = [V|P]$ and $S_2 = [D|B]$. For simplicity, we only consider the case where $f = x^n + 1$. In the general case, the bound on $\|S\|$ involves an extra $\mathrm{EF}(f, 2)$ factor.

We have that $\|BG\|^2 = \|W\|^2 \leq n$, since the entries of $\boldsymbol{h}'_j$ are all 0 except one which is either $h_{i^*,j}$ or $q - 1$. Hence, we obtain that

$$\|S_2\|^2 \leq \|D\|^2 + \|B\|^2 \leq (3\sqrt{nr} + \sqrt{n})^2 + 5 \leq (4\sqrt{nr} + 3)^2.$$

It is obvious that $\|P\| \leq 1$. Additionally, we have that $\|PR\|^2 \leq nr$. Therefore:

$$\|S_1\|^2 \leq \|V\|^2 + \|P\|^2 \leq (\sqrt{nr} + 1)^2 + 1 \leq (\sqrt{nr} + 2)^2,$$

which completes the proof of Theorem 2. □

4 From LWE to SIS

We show that any efficient algorithm solving LWE with some non-negligible probability may be used by a quantum machine to efficiently solve SIS with non-negligible probability. A crucial property of the reduction is that the matrix underlying the SIS and LWE instances is preserved. This allows the reduction to remain valid while working on Ideal-SIS and Ideal-LWE.

Theorem 3. *Let q, m, n be integers, and $\alpha \in (0, 1)$ with $n \geq 32$, $\mathcal{P}oly(n) \geq m \geq 5n \log q$ and $\alpha < \min\left(\frac{1}{10\sqrt{\ln(10m)}}, 0.006\right)$. Suppose that there exists an algorithm that solves $\mathrm{LWE}_{m,q;\Psi_{\alpha q}}$ in time T and with probability $\varepsilon \geq 4m \exp\left(-\frac{\pi}{4\alpha^2}\right)$. Then there exists a quantum algorithm that solves $\mathrm{SIS}_{m,q;\frac{\sqrt{m}}{2\alpha}}$ in time $\mathcal{P}oly(T, n)$ and with probability $\frac{\varepsilon^3}{64} - O(\varepsilon^5) - 2^{-\Omega(n)}$. The result still holds when replacing LWE by Ideal-LWE^f and SIS by Ideal-SIS^f, for $f = x^n + 1$ with $n = 2^k \geq 32$, $m \geq 41 \log q$ and $q \equiv 3 \bmod 8$.*

When $\alpha = O(1/\sqrt{n})$, the reduction applies even to a subexponential algorithm for LWE (with success probability $\varepsilon = 2^{-o(n)}$), transforming it into a subexponential quantum algorithm for SIS (with success probability $\varepsilon = 2^{-o(n)}$). The reduction works also for larger $\alpha = O(1/\sqrt{\log n})$, but in this case only applies to polynomial algorithms for LWE (with success probability $\varepsilon = \Omega(1/\mathcal{P}oly(n))$).

The reduction is made of two components. First, we argue that an algorithm solving LWE provides an algorithm that solves a certain bounded distance decoding problem, where the error vector is normally distributed. In a second step, we show that Regev's quantum algorithm [32, Lemma 3.14] can use such an algorithm to construct small solutions to SIS.

4.1 From LWE to BDD

An algorithm solving LWE allows us to solve, for certain lattices, a variation of the Bounded Distance Decoding problem. In that variation of BDD, the error vector is sampled according to a specified distribution.

Definition 3. *The problem* BDD_χ *with parameter distribution* $\chi(\cdot)$ *is as follows: Given an* n*-dimensional lattice* L *and a vector* $\boldsymbol{t} = \boldsymbol{b} + \boldsymbol{e}$ *where* $\boldsymbol{b} \in L$ *and* $\boldsymbol{e}$ *is distributed according to* $\chi(n)$*, the goal is to find* $\boldsymbol{b}$*. We say that a randomized algorithm* $\mathcal{A}$ *solves* BDD_χ *for a lattice* L *with success probability* $\geq \varepsilon$ *if, for* every $\boldsymbol{b} \in L$*, on input* $\boldsymbol{t} = \boldsymbol{b} + \boldsymbol{e}$*, algorithm* $\mathcal{A}$ *returns* $\boldsymbol{b}$ *with probability* $\geq \varepsilon$ *over the choice of* $\boldsymbol{e}$ *and the randomness of* $\mathcal{A}$*.*

For technical reasons, our reduction will require a randomized BDD_χ algorithm whose behaviour is independent of the solution vector $\boldsymbol{b}$, even when the error vector is fixed. This is made precise below.

Definition 4. *A randomized algorithm* $\mathcal{A}$ *solving* BDD_χ *for lattice* L *is said to be* strongly solution-independent *(SSI) if, for every fixed error vector* $\boldsymbol{e}$*, the probability (over the randomness of* $\mathcal{A}$*) that, given input* $\boldsymbol{t} = \boldsymbol{b} + \boldsymbol{e}$ *with* $\boldsymbol{b} \in L$*, algorithm* $\mathcal{A}$ *returns* $\boldsymbol{b}$ *is independent of* $\boldsymbol{b}$*.*

We show that if we have an algorithm that solves $\mathrm{LWE}_{m,q;\Psi_{\alpha q}}$, then we can construct an algorithm solving $\mathrm{BDD}_{\nu_{\alpha q}}$ for some lattices. Moreover, the constructed BDD algorithm is SSI.

Lemma 7. *Let* q, m, n *be integers and* $\alpha \in (0,1)$*, with* $m, \log q = \mathcal{P}oly(n)$*. Suppose that there exists an algorithm* $\mathcal{A}$ *that solves* $\mathrm{LWE}_{m,q;\Psi_{\alpha q}}$ *in time* T *and with probability* $\varepsilon \geq 4m \exp\left(-\frac{\pi}{4\alpha^2}\right)$*. Then there exists* $\mathcal{S} \subseteq \mathbb{Z}_q^{m \times n}$ *of proportion* $\geq \varepsilon/2$ *and an SSI algorithm* $\mathcal{A}'$ *such that if* $G \in \mathcal{S}$*, algorithm* $\mathcal{A}'$ *solves* $\mathrm{BDD}_{\nu_{\alpha q}}$ *for* $L(G)$ *in time* $T + \mathcal{P}oly(n)$ *and with probability* $\geq \varepsilon/4$*.*

Proof. If $G \in \mathbb{Z}_q^{m \times n}$ and $\boldsymbol{s} \in \mathbb{Z}_q^n$ are sampled uniformly and if the coordinates of $\boldsymbol{e}$ are sampled according to $\Psi_{\alpha q}$, then $\mathcal{A}$ finds $\boldsymbol{s}$ with probability $\geq \varepsilon$ over the choices of $G, \boldsymbol{s}$ and $\boldsymbol{e}$ and a string w of internal random bits. This implies that there exists a subset $\mathcal{S}$ of the G's of proportion $\geq \varepsilon/2$ such that for any $G \in \mathcal{S}$, algorithm $\mathcal{A}$ succeeds with probability $\geq \varepsilon/2$ over the choices of $\boldsymbol{s}$, $\boldsymbol{e}$ and w. For any $G \in \mathcal{S}$, we have $\Pr_{\boldsymbol{s},\boldsymbol{e},w}[\mathcal{A}(G\boldsymbol{s} + \boldsymbol{e}, w) = \boldsymbol{s}] \geq \varepsilon/2$.

On input $\boldsymbol{t} = \boldsymbol{b} + \boldsymbol{e}$, algorithm $\mathcal{A}'$ works as follows: it samples $\boldsymbol{s}$ uniformly in $\mathbb{Z}_q^n$; it computes $\boldsymbol{t}' = \boldsymbol{t} + A\boldsymbol{s}$, which is of the form $\boldsymbol{t}' = G\boldsymbol{s}' + q\boldsymbol{k} + \boldsymbol{e}$, where $\boldsymbol{k} \in \mathbb{Z}^m$; it calls $\mathcal{A}$ on $\boldsymbol{t}' \bmod q$ and finds $\boldsymbol{s}'$ (with probability $\geq \varepsilon/2$); it then computes $\boldsymbol{e}' = \boldsymbol{t}' - G\boldsymbol{s}' \bmod q$ and returns $\boldsymbol{t} - \boldsymbol{e}'$. Suppose that $\mathcal{A}$ succeeds, i.e., we have $\boldsymbol{s} = \boldsymbol{s}'$. Then $\boldsymbol{e}' = \boldsymbol{e} \bmod q$. Using the standard tail bound on the continuous Gaussian and the lower bound on ε we obtain that $\boldsymbol{e}$ has a component of magnitude $\geq q/2$ with probability $\leq m \exp(-\pi/(2\alpha)^2) \leq \varepsilon/4$. The algorithm thus succeeds with probability $\geq \varepsilon/2 - \varepsilon/4 = \varepsilon/4$. □

We now show that an algorithm solving $\mathrm{BDD}_{\nu_{\alpha q}}$ can be used to solve a quantized version of it. This quantization is required for the quantum part of our reduction. The intuition behind the proof is that the discretization grid is so fine (the parameter R can be chosen extremely large) that at the level of the grid the distribution ν_s looks constant.

Lemma 8. *Let $s > 0$ and L be an n-dimensional. Suppose that there exists an SSI algorithm $\mathcal{A}$ that solves BDD_{ν_s} for L in time T and with probability ε. Then there exists an R, whose bit-length is polynomial in $T, n, |\log s|$ and the bit-size of the given basis of L, and an SSI algorithm $\mathcal{A}'$ that solves $\mathrm{BDD}_{D_{L/R,s}}$ within a time polynomial in $\log R$ and with probability $\geq \varepsilon - 2^{-\Omega(n)}$.*

At this point, we have an R of bit-length polynomial in $T, n, |\log \alpha|$ and an SSI algorithm $\mathcal{B}$ with run-time polynomial in $\log R$ that solves $\mathrm{BDD}_{D_{L(G)/R,\alpha q}}$, for any G in a subset $\mathcal{S} \subseteq \mathbb{Z}_q^{m \times n}$ of proportion $\geq \varepsilon/2$, with probability $\geq \varepsilon/4 - 2^{-\Omega(n)}$ over the random choices of $\boldsymbol{e}$ and the internal randomness w. In the following we assume that on input $\boldsymbol{t} = \boldsymbol{b} + \boldsymbol{e}$, algorithm $\mathcal{B}$ outputs $\boldsymbol{e}$ when it succeeds, rather than $\boldsymbol{b}$. We implement $\mathcal{B}$ quantumly as follows: the quantum algorithm $\mathcal{B}_Q$ maps the state $|\boldsymbol{e}\rangle |\boldsymbol{b} + \boldsymbol{e}\rangle |w\rangle$ to the state $|\boldsymbol{e} - \mathcal{B}(\boldsymbol{b} + \boldsymbol{e}, w)\rangle |\boldsymbol{b} + \boldsymbol{e}\rangle |w\rangle$.

4.2 A new interpretation of Regev's quantum reduction

We first recall Regev's quantum reduction [32, Lemma 3.14]. It uses a randomized BDD oracle $\mathcal{B}^{wc}$ that finds the closest vector in a given lattice L to a given target vector, as long as the target is within a prescribed distance $d < \frac{\lambda_1(L)}{2}$ of L (as above, we assume that $\mathcal{B}^{wc}$ returns the error vector). It returns a sample from the distribution $D_{\widehat{L}, \frac{\sqrt{n}}{\sqrt{2}d}}$. We implement oracle $\mathcal{B}^{wc}$ as a quantum oracle $\mathcal{B}_Q^{wc}$ as above. We assume $\mathcal{B}_Q^{wc}$ accepts random inputs of length ℓ.

1. Set R to be a large constant and build a quantum state which is within ℓ_2 distance $2^{-\Omega(n)}$ of the normalized state corresponding to $\sum_{w \in \{0,1\}^\ell} \sum_{\boldsymbol{x} \in \frac{L}{R}, \|\boldsymbol{x}\| < d} \rho_{\frac{d}{\sqrt{n}}}(\boldsymbol{x}) |\boldsymbol{x}\rangle |\boldsymbol{x} \bmod L\rangle |w\rangle$.
2. Apply the BDD oracle $\mathcal{B}_Q^{wc}$ to the above state to remove the entanglement and obtain a state which is within ℓ_2 distance $2^{-\Omega(n)}$ of the normalized state corresponding to $\sum_{\boldsymbol{x} \in \frac{L}{R}, \|\boldsymbol{x}\| < d} \rho_{\frac{d}{\sqrt{n}}}(\boldsymbol{x}) |\boldsymbol{0}\rangle |\boldsymbol{x} \bmod L\rangle |w\rangle$.
3. Apply the quantum Fourier transform over $\mathbb{Z}_R^n$ to the second register to obtain a state that is within ℓ_2 distance $2^{-\Omega(n)}$ of the normalized state corresponding to $\sum_{\boldsymbol{x} \in \widehat{L}, \|\boldsymbol{x}\| < \frac{n}{d}} \rho_{\frac{\sqrt{n}}{d}}(\boldsymbol{x}) \left| \boldsymbol{x} \bmod (R \cdot \widehat{L}) \right\rangle$.
4. Measure the latter to obtain a vector $\widehat{\boldsymbol{b}} \bmod R \cdot \widehat{L}$. Using Babai's algorithm [6], recover $\widehat{\boldsymbol{b}}$ and output it. Its distribution is within statistical distance $2^{-\Omega(n)}$ of $D_{\widehat{L}, \frac{\sqrt{n}}{\sqrt{2}d}}$.

We now replace the perfect oracle $\mathcal{B}_Q^{wc}$ by an imperfect one.

Lemma 9. *Suppose we are given an n-dimensional lattice L, parameters $R > 2^{2n}\lambda_n(L)$ and $s < \frac{\lambda_1(L)}{2\sqrt{2n}}$, and an SSI algorithm $\mathcal{B}$ that solves $\mathrm{BDD}_{D_{\frac{L}{R}}, s}$ for L with run-time T and success probability ε. Then there exists a quantum algorithm $\mathcal{R}$ which outputs a vector $\widehat{\boldsymbol{b}} \in \widehat{L}$ whose distribution is within distance $1 - \varepsilon^2/2 + O(\varepsilon^4) + 2^{-\Omega(n)}$ of $D_{\widehat{L}, \frac{1}{2s}}$. It finishes in time polynomial in $T + \log R$.*

Proof. The quantum algorithm $\mathcal{R}$ is Regev's algorithm above with parameter $d = \sqrt{2n}s < \frac{\lambda_1(L)}{2}$, where $\mathcal{B}_Q^{wc}$ is replaced by the quantum implementation $\mathcal{B}_Q$ of $\mathcal{B}$. We just saw that if the $\mathrm{BDD}_{D_{L/R,s}}$ oracle was succeeding with probability $1-2^{-\Omega(n)}$, then the output vector $\widehat{\boldsymbol{b}}$ would follow a distribution whose statistical distance to $D_{\widehat{L},\frac{1}{2s}}$ would be $2^{-\Omega(n)}$. To work around the requirement that the oracle succeeds with overwhelming probability, we use the notion of trace distance between two quantum states, which is an adaptation of the statistical distance (see [25, Ch. 9]). The trace distance between two (pure) quantum states $|t_1\rangle$ and $|t_2\rangle$ is $\delta(|t_1\rangle, |t_2\rangle) = \sqrt{1 - |\langle t_1|t_2\rangle|^2}$. Its most important property is that for any generalized measurement (POVM), if D_1 (resp. D_2) is the resulting probability distribution when starting from $|t_1\rangle$ (resp. $|t_2\rangle$) then $\Delta(D_1, D_2) \leq \delta(|t_1\rangle, |t_2\rangle)$. Let $|t_1\rangle$ denote the state at the end of Step 2 of Regev's algorithm when we use $\mathcal{B}^{wc}$, and let $|t_2\rangle$ denote the state that we obtain at the end of Step 2 when we use $\mathcal{B}$. We upper bound $\delta(|t_1\rangle, |t_2\rangle)$ as follows.

Since $\mathcal{B}^{wc}(\boldsymbol{x} \bmod L, w) = \boldsymbol{x}$ for $\|\boldsymbol{x}\| < d$, we have that $|t_1\rangle$ is within ℓ_2 distance (and hence trace distance) $2^{-\Omega(n)}$ of the normalized state

$$|t_1'\rangle = 2^{-\ell/2} \sum_{w\in\{0,1\}^\ell} \sum_{\boldsymbol{x}\in\frac{L}{R}} \sqrt{D^d_{L/R,s}(\boldsymbol{x})}\, |\mathbf{0}\rangle\, |\boldsymbol{x} \bmod L\rangle\, |w\rangle\,,$$

where $D^d_{L/R,s}$ denotes the normalized distribution obtained by truncating $D_{L/R,s}$ to vectors of norm $< d$. On the other hand, for the imperfect oracle $\mathcal{B}$, we have that $|t_2\rangle$ is within trace distance $2^{-\Omega(n)}$ of the normalized state

$$|t_2'\rangle = 2^{-\ell/2} \sum_{w\in\{0,1\}^\ell} \sum_{\boldsymbol{x}\in\frac{L}{R}} \sqrt{D^d_{L/R,s}(\boldsymbol{x})}\, |\boldsymbol{x} - \mathcal{B}(\boldsymbol{x} \bmod L, w)\rangle\, |\boldsymbol{x} \bmod L\rangle\, |w\rangle\,.$$

Let $S_{\mathcal{B}} = \{(\boldsymbol{x}, w) \in \frac{L}{R} \times \{0,1\}^\ell \mid \|\boldsymbol{x}\| < d \text{ and } \mathcal{B}(\boldsymbol{x} \bmod L, w) = \boldsymbol{x}\}$. Notice that, if $(\boldsymbol{x}, w) \notin S_{\mathcal{B}}$, the states $|\boldsymbol{x} - \mathcal{B}(\boldsymbol{x} \bmod L, w)\rangle\, |\boldsymbol{x} \bmod L\rangle\, |w\rangle$ and $|\mathbf{0}\rangle\, |\boldsymbol{x'} \bmod L\rangle\, |w'\rangle$ are orthogonal for all $(\boldsymbol{x'}, w')$. Furthermore, if $(\boldsymbol{x}, w) \in S_{\mathcal{B}}$, the states $|\mathbf{0}\rangle\, |\boldsymbol{x} \bmod L\rangle\, |w\rangle$ and $|\mathbf{0}\rangle\, |\boldsymbol{x'} \bmod L\rangle\, |w'\rangle$ are orthogonal for all $(\boldsymbol{x'}, w') \neq (\boldsymbol{x}, w)$ with $\|\boldsymbol{x'}\| < d$, because the mapping $\boldsymbol{x} \mapsto \boldsymbol{x} \bmod L$ is 1-1 over $\boldsymbol{x}$ of norm $< d < \lambda_1(L)/2$. It follows that $|\langle t_1'|t_2'\rangle| = \sum_{(\boldsymbol{x},w)\in S_{\mathcal{B}}} 2^{-\ell} D^d_{L/R,s}(\boldsymbol{x})$. Hence, $|\langle t_1'|t_2'\rangle|$ is equal to the probability p that $\mathcal{B}(\boldsymbol{x} \bmod L, w) = \boldsymbol{x}$, over the choices of $\boldsymbol{x}$ from the distribution $D^d_{L/R,s}$ and $\boldsymbol{w}$ uniformly random in $\{0,1\}^\ell$. By Lemma 2, using the fact that $d > \sqrt{n}s$, we have $p \geq \hat{p} - 2^{-\Omega(n)}$, where $\widehat{p}$ is the corresponding probability when $\boldsymbol{x}$ is sampled from $D_{L/R,s}$. Finally, we have $\widehat{p} = \sum_{\boldsymbol{x}} D_{L/R,s}(\boldsymbol{x}) \Pr_w[\mathcal{B}(\boldsymbol{x} \bmod L, w) = \boldsymbol{x}]$. By the strong solution-independence of $\mathcal{B}$, we have $\Pr_w[\mathcal{B}(\boldsymbol{x} \bmod L, w) = \boldsymbol{x}] = \Pr_w[\mathcal{B}(\boldsymbol{b} + \boldsymbol{x}, w) = \boldsymbol{x}]$ for any fixed $\boldsymbol{b} \in L$. Therefore, $\widehat{p}$ is the success probability of $\mathcal{B}$ in solving $\mathrm{BDD}_{D_{L/R,s}}$, so $\widehat{p} \geq \varepsilon$ by assumption. Overall, we conclude that $\delta(|t_1\rangle, |t_2\rangle) \leq \sqrt{1 - \varepsilon^2 + 2^{-\Omega(n)}}$, and hence the output of $\mathcal{R}$ is within statistical distance $1 - \varepsilon^2/2 + O(\varepsilon^4) + 2^{-\Omega(n)}$ of $D_{\widehat{L},\frac{1}{2s}}$, as claimed. □

To prove Theorem 3, we apply Lemma 9 to the lattices $L(G)$ for $G \in \mathcal{S}$, with algorithm $\mathcal{B}$. For that, we need to ensure that the hypothesis $\alpha q < \frac{\lambda_1(L(G))}{2\sqrt{2m}}$ is

satisfied. From Lemma 4 (resp. Lemma 5 in the case of Ideal-LWE), we know that with probability $1-2^{-\Omega(n)}$ over the choice of G in $\mathbb{Z}_q^{m\times n}$, we have $\lambda_1^\infty(L(G)) \geq \frac{q}{4}$ and $\lambda_1(L(G)) \geq 0.07\sqrt{m}q$. For such 'good' G's, the hypothesis $\alpha q < \frac{\lambda_1(L(G))}{2\sqrt{2m}}$ is satisfied, since $\alpha < 0.006$. The set $\mathcal{S}'$ of the G's in $\mathcal{S}$ for which that condition is satisfied represents a proportion $\geq \varepsilon/2 - 2^{-\Omega(n)}$ of $\mathbb{Z}_q^{m\times n}$. Suppose now that $G \in \mathcal{S}'$. Lemma 9 shows that we can find a vector $\boldsymbol{s} \in G^\perp = q\widehat{L(G)}$ that follows a distribution whose distance to $D_{G^\perp,\frac{1}{2\alpha}}$ is $\Delta = 1 - \frac{\varepsilon^2}{32} + O(\varepsilon^4) + 2^{-\Omega(n)}$. Thanks to Lemmas 1 and 2 (since $G \in \mathcal{S}$ and $\alpha \leq 1/(10\sqrt{\ln(10m)})$, the hypothesis of Lemma 1 is satisfied), we have that with probability $\geq 1 - 2^{-\Omega(n)} - \Delta = \frac{\varepsilon^2}{32} - O(\varepsilon^4) - 2^{-\Omega(n)}$, the returned $\boldsymbol{s}$ is a non-zero vector of $G^\perp$ whose norm is $\leq \frac{\sqrt{m}}{2\alpha}$. Multiplying by the probability $\geq \varepsilon/2 - 2^{-\Omega(n)}$ that $G \in \mathcal{S}'$ gives the claimed success probability and completes the proof of Theorem 4. □

5 Cryptographic Applications

We now use the results of Sections 3 and 4 to construct efficient cryptographic primitives based on ideal lattices. This includes the first provably secure lattice-based public-key encryption scheme with asymptotically optimal encryption and decryption computation costs of $\widetilde{O}(1)$ bit operations per message bit.

5.1 Efficient public-key encryption scheme

Our scheme is constructed in two steps. Firstly, we use the LWE mapping $(\boldsymbol{s}, \boldsymbol{e}) \mapsto G \cdot \boldsymbol{s} + \boldsymbol{e} \bmod q$ as an injective trapdoor one-way function, with the trapdoor being the full-dimensional set of vectors in $G^\perp$ from Section 3, and the one-wayness being as hard as Ideal-SIS (and hence Ideal-SVP) by Theorem 3. This is an efficient ideal lattice analogue of some trapdoor functions presented in [9, 28] for arbitrary lattices. Secondly, we apply the Goldreich-Levin hardcore function based on Toeplitz matrices [10, Sec. 2.5] to our trapdoor function, and XOR the message with the hardcore bits to obtain a semantically secure encryption. To obtain the $\widetilde{O}(1)$ amortized bit complexity per message bit, we use $\widetilde{\Omega}(n)$ hardcore bits, which induces a subexponential loss in the security reduction.

Our trapdoor function family Id-Trap is defined in Figure 1. For security parameter $n = 2^k$, we fix $f(x) = x^n + 1$ and $q = \mathcal{P}oly(n)$ a prime satisfying $q \equiv 3 \bmod 8$. From Lemma 3, it follows that f splits modulo q into two irreducible factors of degree $n/2$. We set $\sigma = 1$, $r = 1 + \log_3 q = \widetilde{O}(1)$ and $m = (\lceil \log q \rceil + 1)\sigma + r = \widetilde{O}(1)$. We define $\mathcal{R} = \mathbb{Z}_q[x]/f$. The following lemma ensures the correctness of the scheme (this is essentially identical to [28, Sec. 4.1]) and asserts that the evaluation and inversion functions can be implemented efficiently.

Lemma 10. *Let $q > 2\sqrt{mn}L$ and $\alpha = o(1/(L\sqrt{\log n}))$. Then for any $s \in \mathcal{R}$ and for $\boldsymbol{e}$ sampled from $\overline{\Psi}_{\alpha q}$, the inversion algorithm recovers $(\boldsymbol{s}, \boldsymbol{e})$ with probability $1 - n^{-\omega(1)}$ over the choice of $\boldsymbol{e}$. Furthermore, the evaluation and inversion algorithms for $h_{\boldsymbol{g}}$ can be implemented with run-time $\widetilde{O}(n)$.*

- **Generating a function with trapdoor.** Run the algorithm from Theorem 2, using $f = x^n + 1, n, q, r, \sigma, m$ as inputs. Suppose it succeeds. It returns $\boldsymbol{g} \in (\mathbb{Z}_q[x]/f)^m$ (function index) and a trapdoor full-rank set S of linearly independent vectors in $\mathrm{rot}_f(\boldsymbol{g})^\perp \subseteq \mathbb{Z}_q^{mn \times mn}$ with $\|S\| \leq \sqrt{2}(4\sqrt{nr} + 3) =: L$ (we have $L = \widetilde{O}(\sqrt{n})$).
- **Function evaluation.** Given function index $\boldsymbol{g}$, we define the trapdoor function $h_{\boldsymbol{g}} : \mathbb{Z}_q^n \times \mathbb{Z}_q^{mn} \to \mathbb{Z}_q^{mn}$ as follows. On input $\boldsymbol{s}$ uniformly random in $\mathbb{Z}_q^n$ and $\boldsymbol{e} \in \mathbb{Z}_q^{mn}$ sampled from $\overline{\Psi}_{\alpha q}$ (defined as the rounding of $\Psi_{\alpha q}$ to the closest integer vector), we compute and return: $\boldsymbol{c} = h_{\boldsymbol{g}}(\boldsymbol{s}, \boldsymbol{e}) := \mathrm{rot}_f(\boldsymbol{g}) \cdot \boldsymbol{s} + \boldsymbol{e} \bmod q$.
- **Function inversion.** Given $\boldsymbol{c} = h_{\boldsymbol{g}}(\boldsymbol{s}, \boldsymbol{e})$ and trapdoor S, compute $\boldsymbol{d} = S^T \cdot \boldsymbol{c} \bmod q$ and $\boldsymbol{e}' = S^{-T} \cdot \boldsymbol{d}$ (in $\mathbb{Q}$). Compute $\boldsymbol{u} = \boldsymbol{c} - \boldsymbol{e}' \bmod q$ and $\boldsymbol{s}' = (\mathrm{rot}_f(g_1))^{-1} \cdot \boldsymbol{u}_1 \bmod q$, where $\boldsymbol{u}_1$ consists of the first n coordinates of $\boldsymbol{u}$. Return $(\boldsymbol{s}', \boldsymbol{e}')$.

Fig. 1. The trapdoor function family Id-Trap.

The one-wayness of Id-Trap is equivalent to the hardness of $\mathrm{LWE}_{m,q;\overline{\Psi}_{\alpha q}}$. Furthermore, an instance of $\mathrm{LWE}_{m,q;\Psi_{\alpha q}}$ can be efficiently converted by rounding to an instance of $\mathrm{LWE}_{m,q;\overline{\Psi}_{\alpha q}}$. This proves Lemma 11.

Lemma 11. *Any attacker against the one-wayness of* Id-Trap *(with parameters m, α, q) with run-time T and success probability ε provides an algorithm for $\mathrm{LWE}_{m,q;\Psi_{\alpha q}}$ with run-time T and success probability ε.*

By combining our trapdoor function with the GL hardcore function [10, Sec. 2.5] we get the encryption scheme of Figure 2.

- **Key generation.** For security parameter n, run the generation algorithm of Id-Trap to get an $h_{\boldsymbol{g}}$ and a trapdoor S. We can view the first component of the domain of $h_{\boldsymbol{g}}$ as a subset of $\mathbb{Z}_2^{\ell_I}$ for $\ell_I = O(n \log q) = \widetilde{O}(n)$. Generate $\boldsymbol{r} \in \mathbb{Z}_2^{\ell_I + \ell_M}$ uniformly and define the Toeplitz matrix $M_{GL} \in \mathbb{Z}_2^{\ell_M \times \ell_i}$ (allowing fast multiplication [26]) whose ith row is $[r_i, \ldots, r_{\ell_I + i - 1}]$. The public key is $(\boldsymbol{g}, \boldsymbol{r})$ and the secret key is S.
- **Encryption.** Given ℓ_M-bit message M with $\ell_M = n/\log n = \widetilde{\Omega}(n)$ and public key $(\boldsymbol{g}, \boldsymbol{r})$, sample $(\boldsymbol{s}, \boldsymbol{e})$ with $\boldsymbol{s} \in \mathbb{Z}_q^n$ uniform and $\boldsymbol{e}$ sampled from $\overline{\Psi}_{\alpha q}$, and evaluate $C_1 = h_{\boldsymbol{g}}(\boldsymbol{s}, \boldsymbol{e})$. Compute $C_2 = M \oplus (M_{GL} \cdot \boldsymbol{s})$, where the product $M_{GL} \cdot \boldsymbol{s}$ is computed over $\mathbb{Z}_2$, and $\boldsymbol{s}$ is viewed as a string over $\mathbb{Z}_2^{\ell_I}$. Return the ciphertext (C_1, C_2).
- **Decryption.** Given ciphertext (C_1, C_2) and secret key $(S, \boldsymbol{r})$, invert C_1 to compute $(\boldsymbol{s}, \boldsymbol{e})$ such that $h_{\boldsymbol{g}}(\boldsymbol{s}, \boldsymbol{e}) = C_1$, and return $M = C_2 \oplus (M_{GL} \cdot \boldsymbol{s})$.

Fig. 2. The semantically secure encryption scheme Id-Enc.

Theorem 4. *Any IND-CPA attacker against* Id-Enc *with run-time T and success probability $1/2 + \varepsilon$ provides an algorithm for* Ideal-$\mathrm{LWE}^f_{m,q;\Psi_{\alpha q}}$ *with run-time $O(2^{3\ell_M} n^3 \varepsilon^{-3} \cdot T)$ and success probability $\Omega(2^{-\ell_M} n^{-1} \cdot \varepsilon)$.*

Proof. The attacker can be converted to a GL hardcore function distinguisher that, given $C_1 = h_{\boldsymbol{g}}(\boldsymbol{s}, \boldsymbol{e})$, M_{GL}, and ℓ_M bit string z, for $\boldsymbol{s}$ sampled uniformly in $\mathbb{Z}_q^n$, $\boldsymbol{e}$ sampled from $\overline{\Psi}_{\alpha q}$, and M_{GL} constructed as in the key generation procedure, distinguishes whether z is uniformly random (independent of $\boldsymbol{s}$ and $\boldsymbol{e}$) or $z = M_{GL} \cdot \boldsymbol{s}$. It has run-time T and advantage ε. The result follows by applying Lemma 2.5.8, Proposition 2.5.7 and Proposition 2.5.3 in [10]. Note that we do not need to give the vector $\boldsymbol{e}$ additionally to $\boldsymbol{s}$ as input to the GL function, as $\boldsymbol{e}$ is uniquely determined once $\boldsymbol{s}$ is given (with overwhelming probability). □

By using Lemma 10 and Theorems 1, 3 and 4, we get our main result.

Corollary 1. *Any IND-CPA attacker against encryption scheme* Id-Enc *with run-time* $2^{o(n)}$ *and success probability* $1/2+2^{-o(n)}$ *provides a quantum algorithm for* $\widetilde{O}(n^2)$-Ideal-SVP *with* $f(x) = x^n + 1$ *and* $n = 2^k$*, with run-time* $2^{o(n)}$ *and overwhelming success probability. Furthermore, the scheme* Id-Enc *encrypts and decrypts* $\widetilde{\Omega}(n)$ *bits within* $\widetilde{O}(n)$ *bit operations, and its keys have* $\widetilde{O}(n)$ *bits.*

5.2 Further applications

Our results have several other applications, adapting various known constructions for unstructured lattices to ideal lattices, as summarised below.

CCA2-secure encryption. Peikert [28] derived a CCA2-secure encryption scheme from the non-structured variant of the trapdoor function family Id-Trap from Figure 1, using the framework of [31, 34] for building a CCA2-secure scheme from a collection of injective trapdoor functions that is secure under correlated product (i.e., one-wayness is preserved if several functions are evaluated on the same input). The approach of [28] can be applied to Id-Trap, using the equality between Ideal-LWE$_{km}$ and the product of k instances of Ideal-LWE$_m$, multiple hardcore bits as in Id-Enc, and instantiating the required strongly unforgeable signature with the Ideal-SVP-based scheme of [18]. By choosing $k = \widetilde{O}(n)$ (the bit-length of the verification key in [18]) and $\alpha = \widetilde{O}(n^{-3/2})$, we obtain a CCA2-secure scheme that encrypts $\widetilde{\Omega}(n)$ bits within $\widetilde{O}(n^2)$ bit operations and whose security relies on the exponential quantum hardness of $\widetilde{O}(n^4)$-Ideal-SVP.

Trapdoor signatures. Gentry *et al.* [9] give a construction of a trapdoor signature (in the random oracle model) from any family of collision-resistant preimage sampleable functions (PSFs). They show how to sample preimages of $f_G(\boldsymbol{x}) = \boldsymbol{x}^T G$, where $G \in \mathbb{Z}_q^{m \times n}$, using a full-dimensional set of short vectors in $G^\perp$. By applying this to $G = \mathrm{rot}_f(\boldsymbol{g})$ and using the trapdoor generation algorithm from Section 3, we obtain a PSF whose collision resistance relies on Ideal-SIS, and hence Ideal-SVP, and thus a structured variant of the trapdoor signature scheme of [9], with $\widetilde{O}(n)$ verification time and signature length.

ID-based identification. From lattice-based signatures, we derive ID-based identification (IBI) and ID-based signature (IBS). Applying the standard strategy, we construct lattice-based IBI schemes as follows: The master generates a key pair of a lattice-based signature scheme, say (G, S); Each user obtains from the master a short vector $\boldsymbol{e}$ such that $\boldsymbol{e}^T G = H(id)$, where H is a random oracle; The prover proves to the verifier that he/she has a short vector $\boldsymbol{e}$ through the Micciancio-Vadhan protocol [24]. This combination yields concurrently secure IBI schemes based on $\widetilde{O}(n^2)$-SVP and $\widetilde{O}(n^2)$-Ideal-SVP in the random oracle model. As the MV protocol is witness indistinguishable, we can use the Fiat-Shamir heuristic [8] and derive lattice-based IBS schemes.

ID-based encryption (IBE). It is shown in [9] that the unstructured variant of the above trapdoor signature can be used as the identity key extraction for an IBE scheme. This requires a 'dual' version of Id-Enc, in which the public key

is of the form $(\boldsymbol{g}, u)$, where $u = H(id)$ is the hashed identity, and the secret key is the signature of id, i.e., a short preimage of u under $f_{\boldsymbol{g}}(\boldsymbol{x}) = \boldsymbol{x}^T \mathrm{rot}_f(\boldsymbol{g})$. We construct the 'dual' encryption as (C_1, C_2) where $C_1 = h_{\boldsymbol{g}}(\boldsymbol{s}, \boldsymbol{e})$ and $C_2 = T_\ell(\mathrm{rot}_f(u) \cdot \boldsymbol{s}) + M$, where $M \in \mathbb{Z}_q^\ell$ contains the message and $T_\ell(\mathrm{rot}_f(u) \cdot \boldsymbol{s})$ denotes the first ℓ coordinates of $\mathrm{rot}_f(u) \cdot \boldsymbol{s} \bmod q$. By adapting the results of [13], we show that $T_\ell(\mathrm{rot}_f(u) \cdot \boldsymbol{s})$ is an exponentially-secure generic hardcore function for uniform $u \in \mathbb{Z}_q^n$, when $\ell = o(n)$. This allows us to prove the IND-CPA security of the resulting IBE scheme based on the hardness of Ideal-SVP.

Acknowledgements. We thank Chris Peikert and Oded Regev for helpful discussions. The first author was partly supported by the LaRedA ANR grant, the second author by a Macquarie University Research Fellowship (MQRF) and ARC Discovery Grant DP0987734, and the fourth author by KAKENHI 19-55201.

References

1. M. Ajtai. Generating hard instances of lattice problems (extended abstract). In *Proceedings of STOC 1996*, pages 99–108. ACM, 1996.
2. M. Ajtai. Generating hard instances of the short basis problem. In *Proceedings of ICALP 1999*, volume 1644 of *LNCS*, pages 1–9. Springer, 1999.
3. M. Ajtai and C. Dwork. A public-key cryptosystem with worst-case/average-case equivalence. In *Proceedings of STOC 1997*, pages 284–293. ACM, 1997.
4. M. Ajtai, R. Kumar, and D. Sivakumar. A sieve algorithm for the shortest lattice vector problem. In *Proceedings of STOC 2001*, pages 601–610. ACM, 2001.
5. J. Alwen and C. Peikert. Generating shorter bases for hard random lattices. In *Proceedings of STACS 2009*, LNCS. Springer, 2009.
6. L. Babai. On Lovász' lattice reduction and the nearest lattice point problem. *Combinatorica*, 6:1–13, 1986.
7. I. F. Blake, S. Gao, and R. C. Mullin. Explicit factorization of $x^{2^k} + 1$ over F_p with prime $p \equiv 3 \bmod 4$. *App. Alg. in Eng., Comm. and Comp*, 4:89–94, 1992.
8. A. Fiat and A. Shamir. How to prove yourself – practical solutions to identification and signature problems. In *Proceedings of Crypto 1986*, volume 263 of *LNCS*, pages 186–194. Springer, 1986.
9. C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *Proceedings of STOC 2008*, pages 197–206. ACM, 2008.
10. O. Goldreich. *Foundations of Cryptography*, volume II – Basic Applications. Cambridge University Press, 2001.
11. O. Goldreich, S. Goldwasser, and S. Halevi. Public-key cryptosystems from lattice reduction problems. In *Proceedings of Crypto 1997*, volume 1294 of *LNCS*, pages 112–131. Springer, 1997.
12. J. Hoffstein, J. Pipher, and J. H. Silverman. NTRU: a ring based public key cryptosystem. In *Proceedings of ANTS III*, volume 1423 of *LNCS*, pages 267–288. Springer, 1998.
13. T. Holenstein, U. Maurer, and J. Sjödin. Complete classification of bilinear hardcore functions. In *Proceedings of Crypto 2004*, volume 3152 of *LNCS*, pages 73–91. Springer, 2004.
14. A. Kawachi, K. Tanaka, and K. Xagawa. Concurrently secure identification schemes based on the worst-case hardness of lattice problems. In *Proceedings of Asiacrypt 2008*, volume 5350 of *LNCS*, pages 372–389. Springer, 2008.

15. V. Lyubashevsky. Lattice-based identification schemes secure under active attacks. In *Proceedings of PKC 2008*, volume 4939 of *LNCS*, pages 162–179. Springer, 2008.
16. V. Lyubashevsky. *Towards Practical Lattice-Based Cryptography.* PhD thesis, University of California, San Diego, 2008.
17. V. Lyubashevsky and D. Micciancio. Generalized compact knapsacks are collision resistant. In *Proceedings of ICALP 2006)*, volume 4052 of *LNCS*, pages 144–155. Springer, 2006.
18. V. Lyubashevsky and D. Micciancio. Asymptotically efficient lattice-based digital signatures. In *Proceedings of TCC 2008*, volume 4948 of *LNCS*, pages 37–54. Springer, 2008.
19. V. Lyubashevsky and D. Micciancio. On bounded distance decoding, unique shortest vectors, and the minimum distance problem. In *Proceedings of Crypto 2009*, volume 5677 of *LNCS*, pages 450–461. Springer, 2009.
20. D. Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions. *Computational Complexity*, 16(4):365–411, 2007.
21. D. Micciancio and S. Goldwasser. *Complexity of lattice problems: a cryptographic perspective.* Kluwer Academic Press, 2002.
22. D. Micciancio and O. Regev. Worst-case to average-case reductions based on gaussian measures. *SIAM J. Comput.*, 37(1):267–302, 2007.
23. D. Micciancio and O. Regev. *Post-Quantum Cryptography*, chapter Lattice-based Cryptography. Springer, 2008.
24. D. Micciancio and S. Vadhan. Statistical zero-knowledge proofs with efficient provers: Lattice problems and more. In *Proceedings of Crypto 2003*, volume 2729 of *LNCS*, pages 282–298. Springer, 2003.
25. M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information.* Cambridge University Press, 2000.
26. V. Y. Pan. *Structured matrices and polynomials, unified superfast algorithms.* Springer and Birkhäuser, 2001.
27. C. Peikert. Limits on the hardness of lattice problems in ℓ_p norms. *Computational Complexity*, 2(17):300–351, 2008.
28. C. Peikert. Public-key cryptosystems from the worst-case shortest vector problem. In *Proceedings of STOC 2009*, pages 333–342. ACM, 2009.
29. C. Peikert and A. Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In *Proceedings of TCC 2006*, pages 145–166, 2006.
30. C. Peikert, V. Vaikuntanathan, and B. Waters. A framework for efficient and composable oblivious transfer. In *Proceedings of Crypto 2008*, volume 5157 of *LNCS*, pages 554–571. Springer, 2008.
31. C. Peikert and B. Waters. Lossy trapdoor functions and their applications. In *Proceedings of STOC 2008*, pages 187–196. ACM, 2008.
32. O. Regev. On lattices, learning with errors, random linear codes, and cryptography. Extended version of [33] dated May 2, 2009. Available at the URL `http://www.cs.tau.ac.il/~odedr/`.
33. O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In *Proceedings of STOC 2005*, pages 84–93. ACM, 2005.
34. A. Rosen and G. Segev. Chosen-ciphertext security via correlated products. In *Proceedings of TCC 2009*, volume 5444 of *LNCS*, pages 419–436. Springer, 2009.
35. C. P. Schnorr. A hierarchy of polynomial lattice basis reduction algorithms. *Theor. Comput. Sci*, 53:201–224, 1987.
36. C. P. Schnorr. Hot topics of LLL and lattice reduction. To appear in the proceedings of the LLL+25 conference, 2009.