This is a corrected version of Proc. of Cryptography and Lattices Conference
(March 29–30, 2001, Providence, Rhode Island, USA)
J. Silverman (Ed.), vol. 2146 of Lecture Notes in Computer Science, pages ???–???

The Two Faces of Lattices in Cryptology

Phong Q. Nguyen and Jacques Stern

École Normale Supérieure, Département d'Informatique,
45 rue d'Ulm, 75005 Paris, France
pnguyen@ens.fr and http://www.di.ens.fr/~pnguyen/
stern@di.ens.fr and http://www.di.ens.fr/~stern/

Abstract. Lattices are regular arrangements of points in n-dimensional space, whose study appeared in the 19th century in both number theory and crystallography. Since the appearance of the celebrated Lenstra-Lenstra-Lovász lattice basis reduction algorithm twenty years ago, lattices have had surprising applications in cryptology. Until recently, the applications of lattices to cryptology were only negative, as lattices were used to break various cryptographic schemes. Paradoxically, several positive cryptographic applications of lattices have emerged in the past five years: there now exist public-key cryptosystems based on the hardness of lattice problems, and lattices play a crucial rôle in a few security proofs. We survey the main examples of the two faces of lattices in cryptology.

1 Introduction

Lattices are discrete subgroups of $\mathbb{R}^n$. A lattice has infinitely many $\mathbb{Z}$-bases, but some are more useful than others. The goal of *lattice reduction* is to find interesting lattice bases, such as bases consisting of reasonably short and almost orthogonal vectors. From the mathematical point of view, the history of lattice reduction goes back to the reduction theory of quadratic forms developed by Lagrange [86], Gauss [55], Hermite [68], Korkine and Zolotarev [82, 83], among others, and to Minkowski's geometry of numbers [103]. With the advent of algorithmic number theory, the subject had a revival in 1981 with Lenstra's celebrated work on integer programming (see [89, 90]), which was, among others, based on a novel lattice reduction technique (which can be found in the preliminary version [89] of [90]). Lenstra's reduction technique was only polynomial-time for fixed dimension, which was however enough in [89]. That inspired Lovász to develop a polynomial-time variant of the algorithm, which computes a so-called *reduced* basis of a lattice. The algorithm reached a final form in the seminal paper [88] where Lenstra, Lenstra and Lovász applied it to factor rational polynomials in polynomial time (back then, a famous problem), from which the name LLL comes. Further refinements of the LLL algorithm were later proposed, notably by Schnorr [121, 122].

Those algorithms have proved invaluable in many areas of mathematics and computer science (see [91, 78, 132, 64, 36, 84]). In particular, their relevance to cryptology was immediately understood, and they were used to break schemes based on the knapsack problem (see [119, 29]), which were early alternatives to the RSA cryptosystem [120]. The success of reduction algorithms at breaking various cryptographic schemes over the past twenty years (see [75]) have arguably established lattice reduction techniques as the most popular tool in public-key cryptanalysis. As a matter of fact, applications of lattices to cryptology have been mainly negative. Interestingly, it was noticed in many cryptanalytic experiments that LLL, as well as other lattice reduction algorithms, behave much more nicely than what was expected from the worst-case proved bounds. This led to a common belief among cryptographers, that lattice reduction is an easy problem, at least in practice.

That belief has recently been challenged by some exciting progress on the complexity of lattice problems, which originated in large part in two seminal papers written by Ajtai in 1996 and in 1997 respectively. Prior to 1996, little was known on the complexity of lattice problems. In his 1996 paper [3], Ajtai discovered a fascinating connection between the worst-case complexity and the average-case complexity of some well-known lattice problems. Such a connection is not known to hold for any other problem in NP believed to be outside P. In his 1997 paper [4], building on previous work by Adleman [2], Ajtai further proved the NP-hardness (under randomized reductions) of the most famous lattice problem, the shortest vector problem (SVP). The NP-hardness of SVP has been a long standing open problem. Ajtai's breakthroughs initiated a series of new results on the complexity of lattice problems, which are nicely surveyed by Cai [30, 31].

Those complexity results opened the door to positive applications in cryptology. Indeed, several cryptographic schemes based on the hardness of lattice problems were proposed shortly after Ajtai's discoveries (see [5, 61, 69, 32, 99, 50]). Some have been broken, while others seem to resist state-of-the-art attacks, for now. Those schemes attracted interest for at least two reasons: on the one hand, there are very few public-key cryptosystems based on problems different from integer factorization or the discrete logarithm problem, and on the other hand, some of those schemes offered encryption/decryption rates asymptotically higher than classical schemes. Besides, one of those schemes, by Ajtai and Dwork [5], enjoyed a surprising security proof based on worst-case (instead of average-case) hardness assumptions.

Independently of those developments, there has been renewed cryptographic interest in lattice reduction, following a beautiful work by Coppersmith [38] in 1996. Coppersmith showed, by means of lattice reduction, how to solve rigorously certain problems, apparently non-linear, related to the question of finding small roots of low-degree polynomial equations. In particular, this has led to surprising attacks on the RSA [120] cryptosystem in special settings such as low public or private exponent, but curiously, also to new security proofs [128, 18]. Coppersmith's results differ from "traditional" applications of lattice reduction in cryptanalysis, where the underlying problem is already linear, and the attack often heuristic by requiring (at least) that current lattice reduction algorithms behave ideally, as opposed to what is theoretically guaranteed. The use of lattice reduction techniques to solve polynomial equations goes back to the eighties [66, 133]. The first result of that kind, the broadcast attack on low-exponent RSA due to Håstad [66], can be viewed as a weaker version of Coppersmith's theorem on univariate modular polynomial equations.

A shorter version of this survey previously appeared in [118]. The rest of the paper is organized as follows. In Section 2, we give basic definitions and results on lattices and their algorithmic problems. In Section 3, we survey an old application of lattice reduction in cryptology: finding small solutions of multivariate linear equations, which includes the well-known subset sum or knapsack problem as a special case. In Section 4, we review a related problem: the hidden number problem. In Section 5, we discuss lattice-based cryptography, somehow a revival for knapsack-based cryptography. In Section 6, we discuss developments on the problem of finding small roots of polynomial equations, inspired by Coppersmith's discoveries in 1996. In Section 7, we survey the surprising links between lattice reduction, the RSA cryptosystem, and integer factorization.

2 Lattice problems

2.1 Definitions

Recall that a *lattice* is a discrete (additive) subgroup of $\mathbb{R}^n$. In particular, any subgroup of $\mathbb{Z}^n$ is a lattice, and such lattices are called *integer lattices.* An equivalent definition is that a lattice consists of all integral linear combinations of a set of linearly independent vectors, that is,

$$L = \left\{ \sum_{i=1}^{d} n_i \mathbf{b}_i \mid n_i \in \mathbb{Z} \right\},$$

where the $\mathbf{b}_i$'s are linearly independent over $\mathbb{R}$. Such a set of vectors $\mathbf{b}_i$'s is called a lattice *basis.* All the bases have the same number $\dim(L)$ of elements, called the *dimension* (or *rank*) of the lattice since it matches the dimension of the vector subspace $\mathrm{span}(L)$ spanned by L.

There are infinitely many lattice bases when $\dim(L) \geq 2$. Any two bases are related to each other by some unimodular matrix (integral matrix of determinant ± 1), and therefore all the bases share the same Gramian determinant $\det_{1 \leq i,j \leq d} \langle \mathbf{b}_i, \mathbf{b}_j \rangle$. The *volume* $\mathrm{vol}(L)$ (or *determinant*) of the lattice is by definition the square root of that Gramian determinant, thus corresponding to the d-dimensional volume of the parallelepiped spanned by the $\mathbf{b}_i$'s. In the important case of full-dimensional lattices where $\dim(L) = n$, the volume is equal to the absolute value of the determinant of any lattice basis (hence the name determinant). If the lattice is further an integer lattice, then the volume is also equal to the index $[\mathbb{Z}^n : L]$ of L in $\mathbb{Z}^n$.

Since a lattice is discrete, it has a shortest non-zero vector: the Euclidean norm of such a vector is called the lattice *first minimum*, denoted by $\lambda_1(L)$ or $\|L\|$. Of course, one can use other norms as well: we will use $\|L\|_\infty$ to denote the first minimum for the infinity norm. More generally, for all $1 \leq i \leq \dim(L)$, Minkowski's i-th *minimum* $\lambda_i(L)$ is defined as the minimum of $\max_{1 \leq j \leq i} \|\mathbf{v}_j\|$ over all i linearly independent lattice vectors $\mathbf{v}_1, \ldots, \mathbf{v}_i \in L$. There always exist linearly independent lattice vectors $\mathbf{v}_1, \ldots, \mathbf{v}_d$ reaching the minima, that is $\|\mathbf{v}_i\| = \lambda_i(L)$. However, surprisingly, as soon as $\dim(L) \geq 4$, such vectors do not necessarily form a lattice basis, and when $\dim(L) \geq 5$, there may not even exist a lattice basis reaching the minima. This is one of the reasons why there exist several notions of basis reduction in high dimension, without any "optimal" one. It will be convenient to define the *lattice gap* as the ratio $\lambda_2(L)/\lambda_1(L)$ between the first two minima.

Minkowski's Convex Body Theorem guarantees the existence of short vectors in lattices: a careful application shows that any d-dimensional lattice L satisfies $\|L\|_\infty \leq \mathrm{vol}(L)^{1/d}$, which is obviously the best possible bound. It follows that $\|L\| \leq \sqrt{d}\mathrm{vol}(L)^{1/d}$, which is not optimal, but shows that the value $\lambda_1(L)/\mathrm{vol}(L)^{1/d}$ is bounded when L runs over all d-dimensional lattices. The supremum of $\lambda_1(L)^2/\mathrm{vol}(L)^{2/d}$ is denoted by γ_d, and called Hermite's constant[1] of dimension d, because Hermite was the first to establish its existence in the language of quadratic forms. The exact value of Hermite's constant is only known for $d \leq 8$. The best asymptotic bounds known for Hermite's constant are the following ones (see [102, Chapter II] for the lower bound, and [37, Chapter 9] for the upper bound):

$$\frac{d}{2\pi e} + \frac{\log(\pi d)}{2\pi e} + o(1) \leq \gamma_d \leq \frac{1.744d}{2\pi e}(1 + o(1)).$$

Minkowski proved more generally:

Theorem 1 (Minkowski). *For all d-dimensional lattices L and all $r \leq d$:*

$$\prod_{i=1}^{r} \lambda_i(L) \leq \sqrt{\gamma_d^r}\mathrm{vol}(L)^{r/d}.$$

A general principle, dating back to Gauss, estimates the number of lattice points (in a full-rank lattice) in nice sets of $\mathbb{R}^n$ by the volume of the set divided by the volume of the lattice, with a small error term. This approach can be proved to be rigorous in certain settings, such as when the lattice dimension is fixed and the set is the ball centered at the origin with radius growing to infinity. Thus, one often heuristically approximates the successive minima of a d-dimensional lattice L by $\sqrt{\frac{d}{2\pi e}}\mathrm{vol}(L)^{1/d}$. This is of course only an intuitive estimate, which may be far away from the truth.

For any lattice L of $\mathbb{R}^n$, one defines the *dual lattice* (also called *polar lattice*) of L as:

$$L^* = \{\mathbf{x} \in \mathrm{span}(L) : \forall \mathbf{y} \in L, \langle \mathbf{x}, \mathbf{y} \rangle \in \mathbb{Z}\}.$$

[1] For historical reasons, Hermite's constant refers to $\max \lambda_1(L)^2/\mathrm{vol}(L)^{2/d}$ and not to $\max \lambda_1(L)/\mathrm{vol}(L)^{1/d}$.

If $(\mathbf{b}_1, \ldots, \mathbf{b}_d)$ is a basis of L, then the dual family $(\mathbf{b}_1^*, \ldots, \mathbf{b}_d^*)$ is a basis of L^* (the dual family is the unique linearly independent family of $\mathrm{span}(L)$ such that $\langle \mathbf{b}_i^*, \mathbf{b}_j \rangle$ is equal to 1 if $i = j$, and to 0 otherwise). Thus, $(L^*)^* = L$ and $\mathrm{vol}(L)\mathrm{vol}(L^*) = 1$. The so-called transference theorems relate the successive minima of a lattice and its dual lattice. The first transference theorem follows from the definition of Hermite's constant:

$$\lambda_1(L)\lambda_1(L^*) \leq \gamma_d.$$

A more difficult transference theorem (see [9]) ensures that for all $1 \leq r \leq d$:

$$\lambda_r(L)\lambda_{d-r+1}(L^*) \leq d.$$

Both these transference bounds are optimal up to a constant. More information on lattice theory can be found in numerous textbooks, such as [65, 131, 92].

2.2 Algorithmic problems

In the rest of this section, we assume implicitly that lattices are rational lattices (lattices in $\mathbb{Q}^n$), and d will denote the lattice dimension.

The most famous lattice problem is the *shortest vector problem* (SVP): given a basis of a lattice L, find $\mathbf{u} \in L$ such that $\|\mathbf{u}\| = \|L\|$ (recall that $\|L\| = \lambda_1(L)$). SVP_∞ will denote the analogue for the infinity norm. One defines approximate short vector problems by asking a non-zero $\mathbf{v} \in L$ with norm bounded by some approximation factor: $\|\mathbf{v}\| \leq f(d)\|L\|$.

The *closest vector problem* (CVP), also called the *nearest lattice point problem*, is a non-homogeneous version of the shortest vector problem: given a basis of a lattice L and a vector $\mathbf{v} \in \mathbb{R}^n$ (it does not matter whether or not $\mathbf{v} \in \mathrm{span}(L)$), find a lattice vector minimizing the distance to $\mathbf{v}$. Again, one defines approximate closest vector problems by asking $\mathbf{u} \in L$ such that for all $\mathbf{w} \in L$, $\|\mathbf{u} - \mathbf{v}\| \leq f(d)\|\mathbf{w} - \mathbf{v}\|$.

Another problem is the *smallest basis problem* (SBP), which has many variants depending on the exact meaning of "smallest". The variant currently in vogue (see [3, 14]) is the following: find a lattice basis minimizing the maximum of the lengths of its elements. A more geometric variant asks instead to minimize the product of the lengths (see [64]), since the product is always larger than the lattice volume, with equality if and only if the basis is orthogonal.

2.3 Complexity results

We refer to Cai [30, 31] for an up-to-date survey of complexity results. Ajtai [4] recently proved that SVP is NP-hard under randomized reductions. Micciancio [98, 97] simplified and improved the result by showing that approximating SVP to within a factor $< \sqrt{2}$ is also NP-hard under randomized reductions. The NP-hardness of SVP under deterministic (Karp) reductions remains an open problem.

CVP seems to be a more difficult problem. Goldreich *et al.* [62] recently noticed that CVP cannot be easier than SVP: given an oracle that approximates CVP to within a factor $f(d)$, one can approximate SVP in polynomial time to within the same factor $f(d)$. Reciprocally, Kannan proved in [78, Section 7] that any algorithm approximating SVP to within a non-decreasing function $f(d)$ can be used to approximate CVP to within $d^{3/2}f(d)^2$. CVP was shown to be NP-hard as early as in 1981 [49] (for a much simpler "one-line" proof using the knapsack problem, see [100]). Approximating CVP to within a quasi-polynomial factor $2^{\log^{1-\varepsilon} d}$ is NP-hard [7, 45].

However, NP-hardness results for SVP and CVP have limits. Goldreich and Goldwasser [58] showed that approximating SVP or CVP to within $\sqrt{d/\log d}$ is not NP-hard, unless the polynomial-time hierarchy collapses.

Interestingly, SVP and CVP problems seem to be more difficult with the infinity norm. It was shown that SVP_∞ and CVP_∞ are NP-hard in 1981 [49]. In fact, approximating $\text{SVP}_\infty/\text{CVP}_\infty$ to within an almost-polynomial factor $d^{1/\log\log d}$ is NP-hard [44]. On the other hand, Goldreich and Goldwasser [58] showed that approximating $\text{SVP}_\infty/\text{CVP}_\infty$ to within $d/\log d$ is not NP-hard, unless the polynomial-time hierarchy collapses.

We will not discuss Ajtai's worst-case/average-case equivalence [3, 33], which refers to special versions of SVP and SBP (see [30, 31, 14]) such as SVP when the lattice gap λ_2/λ_1 is at least polynomial in the dimension.

2.4 Algorithmic results

The main algorithmic results are surveyed in [91, 78, 132, 64, 36, 84, 30, 109]. No polynomial-time algorithm is known for approximating either SVP, CVP or SBP to within a polynomial factor in the dimension d. In fact, the existence of such algorithms is an important open problem. The best polynomial time algorithms achieve only slightly subexponential factors, and are based on the LLL algorithm [88], which can approximate SVP and SBP. However, it should be emphasized that these algorithms typically perform much better than is theoretically guaranteed, on instances of practical interest. Given as input any basis of a lattice L, LLL provably outputs in polynomial time a basis $(\mathbf{b}_1, \ldots, \mathbf{b}_d)$ satisfying:

$$\|\mathbf{b}_1\| \le 2^{(d-1)/4}\text{vol}(L)^{1/d}, \|\mathbf{b}_i\| \le 2^{(d-1)/2}\lambda_i(L) \text{ and } \prod_{i=1}^{d} \|\mathbf{b}_i\| \le 2^{\binom{d}{2}/2}\text{vol}(L).$$

Thus, LLL can approximate SVP to within $2^{(d-1)/2}$. Schnorr[2] [121] improved the bound to $2^{O(d(\log\log d)^2/\log d)}$, and Ajtai *et al.* [6] recently further improved it to $2^{O(d\log\log d/\log d)}$ in randomized polynomial time thanks to a new randomized algorithm to find the shortest vector. In fact, Schnorr defined an LLL-based

[2] Schnorr's result is usually cited in the literature as an approximation algorithm to within $(1+\varepsilon)^n$ for any constant $\varepsilon > 0$. However, Goldreich and Håstad noticed about two years ago that one can choose some $\varepsilon = o(1)$ and still have polynomial running time, for instance using the blocksize $k = \log d/\log\log d$ in [121].

family of algorithms [121] (named BKZ for blockwise Korkine-Zolotarev) whose performances depend on a parameter called the blocksize. These algorithms use some kind of exhaustive search super-exponential in the blocksize. So far, the best reduction algorithms in practice are variants [124, 125] of those BKZ-algorithms, which apply a heuristic to reduce exhaustive search. But little is known on the average-case (and even worst-case) complexity of reduction algorithms.

Babai's nearest plane algorithm [8] uses LLL to approximate CVP to within $2^{d/2}$, in polynomial time (see also [80]). Using Schnorr's algorithm [121], this can be improved to $2^{O(d(\log\log d)^2/\log d)}$ in polynomial time, and even further to $2^{O(d\log\log d/\log d)}$ in randomized polynomial time using [6], due to Kannan's link between CVP and SVP (see previous section). In practice however, the best strategy seems to be the *embedding method* (see [61, 108]), which uses the previous algorithms for SVP and a simple heuristic reduction from CVP to SVP. Namely, given a lattice basis $(\mathbf{b}_1, \ldots, \mathbf{b}_d)$ and a vector $\mathbf{v} \in \mathbb{R}^n$, the embedding method builds the $(d+1)$-dimensional lattice (in $\mathbb{R}^{n+1}$) spanned by the row vectors $(\mathbf{b}_i, 0)$ and $(\mathbf{v}, 1)$. Depending on the lattice, one should choose a coefficient different than 1 in $(\mathbf{v}, 1)$. It is hoped that a shortest vector of that lattice is of the form $(\mathbf{v} - \mathbf{u}, 1)$ where $\mathbf{u}$ is a closest vector (in the original lattice) to $\mathbf{v}$, whenever the distance to the lattice is smaller than the lattice first minimum. This heuristic may fail (see for instance [97] for some simple counterexamples), but it can also sometimes be proved, notably in the case of lattices arising from low-density knapsacks.

For exact SVP, the best algorithm known (in theory) is the recent randomized $2^{O(d)}$-time algorithm by Ajtai *et al.* [6], which improved Kannan's super-exponential algorithm [77, 79] (see also [67]). For exact CVP, the best algorithm remains Kannan's super-exponential algorithm [77, 79], with running time $2^{O(d\log d)}$ (see also [67] for an improved constant).

3 Finding small roots of multivariate linear equations

One of the early and most natural applications of lattice reduction in cryptology was to find small roots of multivariate linear equations, where the equations are either integer equations or modular equations.

3.1 Knapsacks

Cryptology and lattices share a long history with the *knapsack* (also called *subset sum*) problem, a well-known NP-hard problem considered by Karp, and a particular case of multivariate linear equation: given a set $\{a_1, a_2, \ldots, a_n\}$ of positive integers and a sum $s = \sum_{i=1}^{n} x_i a_i$, where $x_i \in \{0, 1\}$, recover the x_i's.

In 1978, Merkle and Hellman[96] invented one of the first public-key cryptosystems, by converting some easy knapsacks into what they believed were hard knapsacks. It was basically the unique alternative to RSA until 1982, when Shamir [126] proposed a (heuristic) attack against the simplest version of the Merkle-Hellman scheme. Shamir used Lenstra's integer programming algorithm [89, 90] but, the same year, Adleman [1] showed how to use LLL instead, making experiments much easier. Brickell [27, 28] later extended the attacks to the more general "iterated" Merkle-Hellman scheme, and showed that Merkle-Hellman was insecure for all realistic parameters. The cryptanalysis of Merkle-Hellman schemes was the first application of lattice reduction in cryptology.

Despite the failure of Merkle-Hellman cryptosystems, researchers continued to search for knapsack cryptosystems because such systems are very easy to implement and can attain very high encryption/decryption rates. But basically, all knapsack cryptosystems have been broken (for a survey, see [119]), either by specific (often lattice-based) attacks or by the low-density attacks. The last significant candidate to survive was the Chor-Rivest cryptosystem [35], broken by Vaudenay [135] in 1997 with algebraic (not lattice) methods.

3.2 Low-density attacks on knapsacks

We only describe the basic link between lattices and knapsacks. Note that Ajtai's original proof [4] for the NP-hardness (under randomized reductions) of SVP used a connection between the subset sum problem and SVP.

Solving the knapsack problem amounts to find a $0,1$-solution of an inhomogeneous linear equation, which can be viewed as a closest vector problem in a natural way, by considering the corresponding homogeneous linear equation, together with an arbitrary solution of the inhomogeneous equation. Indeed, let $s = \sum_{i=1}^n x_i a_i$ be a knapsack instance. One can compute in polynomial time integers $y_1, \ldots, y_n$ such that $s = \sum_{i=1}^n y_i a_i$, using for instance an extended gcd algorithm. Then the vector $(y_1 - x_1, \ldots, y_n - x_n)$ belongs to the $(n-1)$-dimensional lattice L formed by all the solutions of the homogeneous equation, that is the vectors $(z_1, \ldots, z_n) \in \mathbb{Z}^n$ such that:

$$z_1 a_1 + \cdots + z_n a_n = 0.$$

And this lattice vector is fairly close to the vector $(y_1, \ldots, y_n)$, since the distance is roughly $\sqrt{n/2}$. But because $x_i \in \{0, 1\}$, the lattice vector is even closer to the vector $\mathbf{y} = (y_1 - 1/2, \ldots, y_n - 1/2)$ for which the distance is exactly $\sqrt{n/4}$. In fact, it is easy to see that $\mathbf{x} = (y_1 - x_1, \ldots, y_n - x_n)$ is a closest vector to $\mathbf{y}$ in the lattice L, and that any lattice vector whose distance to $\mathbf{y}$ is exactly $\sqrt{n/4}$ is necessarily of the form $(y_1 - x'_1, \ldots, y_n - x'_n)$ where $s = \sum_{i=1}^n x'_i a_i$ and $x'_i \in \{0, 1\}$. This gives a deterministic polynomial-time reduction from the knapsack problem to CVP (this reduction appeared in [100] with a slightly different lattice).

One can derive from this reduction a provable method to solve the knapsack problem in polynomial time with high probability when the knapsack *density* defined as $d = n / \max_{1 \le i \le n} \log_2 a_i$ is low (see [85, 51, 54]). Indeed, if $\|\mathbf{x} - \mathbf{y}\| = \sqrt{n/4}$ is strictly less than $2^{-(n-1)/2-1}\|L\|$, then by applying Babai's nearest plane CVP approximation algorithm to L and $\mathbf{y}$, one obtains $\mathbf{z} \in L$ such that $\|\mathbf{z} - \mathbf{y}\| < 2^{n/2}\|\mathbf{x} - \mathbf{y}\| < \|L\|/2$, and thus $\|\mathbf{z} - \mathbf{x}\| < \|L\|$ where $\mathbf{z} - \mathbf{x} \in L$, which implies that $\mathbf{z} = \mathbf{x}$, disclosing the x_i's. It remains to estimate the first minimum $\|L\|$. With high probability, the a_i's are coprime, and then:

$$\mathrm{vol}(L) = \left(\sum_{i=1}^{n} a_i^2\right)^{1/2} \approx 2^{n/d}\sqrt{n}.$$

Thus, one expects $\|L\| \approx 2^{1/d}\sqrt{\frac{n}{2\pi e}}$. It follows that the method should work whenever

$$\sqrt{\frac{n}{4}} < 2^{-(n-1)/2-1} 2^{1/d}\sqrt{\frac{n}{2\pi e}},$$

that is, roughly $d \leq 2/n$. This volume argument can be made rigorous because the probability that a fixed non-zero vector belongs to L is less than $1/A$ when the a_i's are chosen uniformly at random from $[0, A]$. One deduces that most knapsacks of density roughly less than $2/n$ are solvable in polynomial time (see [85, 51, 54]).

One does not know how to provably solve the knapsack problem in polynomial time when the density lies between $2/n$ and 1, which is typically the case for cryptographic knapsacks (where the density should be less than 1, otherwise heuristically, there would be several solutions, causing decryption troubles). However, one can hope that the embedding method that heuristically reduces CVP to SVP works, as while as the distance to the lattice (which is $\sqrt{n/4}$) is smaller than the first minimum $\|L\|$. By the previous reasoning, this should happen when

$$\sqrt{\frac{n}{4}} \leq 2^{1/d}\sqrt{\frac{n}{2\pi e}},$$

that is,

$$d \leq \frac{1}{\log_2 \sqrt{\pi e/2}} \approx 0.955\ldots$$

This heuristic bound turns out to be not too far away from the truth. Indeed, one can show that the target vector $(x_1 - 1/2, \ldots, x_n - 1/2, 1)$ is with high probability (over the choice of the a_i's) the shortest vector in the embedding lattice, when the density $d \leq 0.9408\ldots$ (see [41] who used a slightly different lattice, but the proof carries through). This is done by enumerating all possible short vectors, and using bounds on the number of integral points in high-dimensional spheres [93]. The result improved the earlier bound of $0.6463\ldots$ from Lagarias and Odlyzko [85], which was essentially obtained by approximating the vector $(y_1, \ldots, y_n)$ in the lattice L, instead of $(y_1 - 1/2, \ldots, y_n - 1/2)$. This rigorous bound of $0.6463\ldots$ matches the heuristic bound obtained by a volume argument on the corresponding embedding lattice:

$$\sqrt{\frac{n}{2}} \leq 2^{1/d}\sqrt{\frac{n}{2\pi e}}.$$

To summarize, the subset sum problem can always be efficiently reduced to CVP, and this reduction leads to an efficient probabilistic reduction to SVP in low density, and to a polynomial-time solution in extremely low density. In the light of recent results on the complexity of SVP, those reductions from knapsack to SVP may seem useless. Indeed, the NP-hardness of SVP under randomized reductions suggests that there is no polynomial-time algorithm that solves SVP. However, it turns out that in practice, one can hope that standard lattice reduction algorithms behave like SVP-oracles, up to reasonably high dimensions. Experiments carried out in [85, 124, 125] show the effectiveness of such an approach for solving low-density subset sums, up to n about the range of 100–200. It does not prove nor disprove that one can solve, in theory or in practice, low-density knapsacks with n over several hundreds. But it was sufficient to show that knapsack cryptography was impractical: indeed, the keysize of knapsack schemes grows in general at least quadratically with n, so that high values of n (as required by lattice attacks) are not practical.

Thus, lattice methods to solve the subset sum problem are mainly heuristic. And lattice attacks against knapsack cryptosystems are somehow even more heuristic, because the reductions from knapsack to SVP assume a uniform distribution of the weights a_i's, which is in general not necessarily satisfied by knapsacks arising from cryptosystems.

3.3 The orthogonal lattice

Recently, Nguyen and Stern proposed in [113] a natural generalization of the lattices arising from the knapsack problem. More precisely, they defined for any integer lattice L in $\mathbb{Z}^n$, the *orthogonal lattice* $L^\perp$ as the set of integral vectors orthogonal to L, that is, the set of $\mathbf{x} \in \mathbb{Z}^n$ such that the dot product $\langle \mathbf{x}, \mathbf{y} \rangle = 0$ for all $\mathbf{y} \in L$. Note that the lattice $L^\perp$ has dimension $n - \dim(L)$, and can be computed in polynomial time from L (see [36]). Interestingly, the links between duality and orthogonality (see Martinet's book [92, pages 34–35]) enable to prove that the volume of $L^\perp$ is equal to the volume of the lattice $\mathrm{span}(L) \cap \mathbb{Z}^n$ which we denote by $\bar{L}$. Thus, if a lattice in $\mathbb{Z}^n$ is low-dimensional, its orthogonal lattice is high-dimensional with a volume at most equal: the successive minima of the orthogonal lattice are likely to be much shorter than the ones of the original lattice. That property of orthogonal lattices has led to effective (though heuristic) lattice-based attacks on various cryptographic schemes [113, 115, 116, 114, 117]. We refer to [109] for more information. In particular, it was used in [117] to solve the *hidden subset sum problem* (used in [26]) in low density. The hidden subset sum problem was apparently a non-linear version of the subset sum problem: given M and n in $\mathbb{N}$, and $b_1, \ldots, b_m \in \mathbb{Z}_M$, find $\alpha_1, \ldots, \alpha_n \in \mathbb{Z}_M$ such that each b_i is some subset sum modulo M of $\alpha_1, \ldots, \alpha_n$.

We sketch the solution of [117] to give a flavour of cryptanalyses based on orthogonal lattices. We first restate the hidden subset sum problem in terms of vectors. We are given an integer M, and a vector $\mathbf{b} = (b_1, \ldots, b_m) \in \mathbb{Z}^m$ with entries in $[0..M-1]$ such that there exist integers $\alpha_1, \ldots, \alpha_n \in [0..M-1]$, and vectors $\mathbf{x}_1, \ldots, \mathbf{x}_n \in \mathbb{Z}^m$ with entries in $\{0, 1\}$ satisfying:

$$\mathbf{b} = \alpha_1 \mathbf{x}_1 + \alpha_2 \mathbf{x}_2 + \cdots + \alpha_n \mathbf{x}_n \pmod{M}.$$

We want to determine the α_i's. There exists a vector $\mathbf{k} \in \mathbb{Z}^m$ such that:

$$\mathbf{b} = \alpha_1 \mathbf{x}_1 + \alpha_2 \mathbf{x}_2 + \cdots + \alpha_n \mathbf{x}_n + M\mathbf{k}.$$

Notice that if $\mathbf{u}$ in $\mathbb{Z}^n$ is orthogonal to $\mathbf{b}$, then $\mathbf{p_u} = (\langle \mathbf{u}, \mathbf{x}_1 \rangle, \ldots, \langle \mathbf{u}, \mathbf{x}_n \rangle, \langle \mathbf{u}, \mathbf{k} \rangle)$ is orthogonal to the vector $\mathbf{v}_\alpha = (\alpha_1, \ldots, \alpha_n, M)$. But $\mathbf{v}_\alpha$ is independent of m, and so is the n-dimensional lattice $\mathbf{v}_\alpha^\perp$. On the other hand, as m grows for a fixed M, most of the vectors of any reduced basis of the $(m-1)$-dimensional lattice $\mathbf{b}^\perp$ should get shorter and shorter, because they should have norm close to $\mathrm{vol}(\mathbf{b}^\perp)^{1/(m-1)} \leq \mathrm{vol}(\mathbf{b})^{1/(m-1)} = \|\mathbf{b}\|^{1/(m-1)} \approx (M\sqrt{m})^{1/(m-1)}$. For such vectors $\mathbf{u}$, the corresponding vectors $\mathbf{p}_u$ also get shorter and shorter. But if $\mathbf{p}_u$ gets smaller than $\lambda_1(\mathbf{v}_\alpha^\perp)$ (which is independent of m), then it is actually zero, that is, $\mathbf{u}$ is orthogonal to all the $\mathbf{x}_j$'s and $\mathbf{k}$. Note that one expects $\lambda_1(\mathbf{v}_\alpha^\perp)$ to be of the order of $\|\mathbf{v}_\alpha\|^{1/n} \approx (M\sqrt{n})^{1/n}$.

This suggests that if $(\mathbf{u}_1, \ldots, \mathbf{u}_{m-1})$ is a sufficiently reduced basis of $\mathbf{b}^\perp$, then the first $m-(n+1)$ vectors $\mathbf{u}_1, \ldots, \mathbf{u}_{m-(n+1)}$ should heuristically be orthogonal to all the $\mathbf{x}_j$'s and $\mathbf{k}$. One cannot expect that more than $m-(n+1)$ vectors are orthogonal because the lattice L_x spanned by the $\mathbf{x}_j$'s and $\mathbf{k}$ is likely to have dimension $(n+1)$. From the previous discussion, one can hope that the heuristic condition is satisfied when the density $n/\log(M)$ is very small (so that $\lambda_1(\mathbf{v}_\alpha^\perp)$ is not too small), and m is sufficiently large. And if the heuristic condition is satisfied, the lattice $\bar{L}_x$ is disclosed, because it is then equal to the orthogonal lattice $(\mathbf{u}_1, \ldots, \mathbf{u}_{m-(n+1)})^\perp$. Once $\bar{L}_x$ is known, it is not difficult to recover (heuristically) the vectors $\mathbf{x}_j$'s by lattice reduction, because they are very short vectors. One eventually determines the coefficients α_j's from a linear modular system. The method is quite heuristic, but it works in practice for small parameters in low density (see [117] for more details).

3.4 Multivariate modular linear equations

The technique described in Section 3.2 to solve the knapsack problem can easily be extended to find small solutions of a system of multivariate linear equations over the integers: one views the problem as a closest vector problem in the lattice corresponding to the homogenized equations, which is an orthogonal lattice. Naturally, a similar method can be applied to a system of multivariate linear modular equations, except that in this case, the corresponding lattice is not an orthogonal lattice.

Let $A = (a_{i,j})$ be an $\ell \times k$ integral matrix, $\mathbf{c} \in \mathbb{Z}^\ell$ be a column vector and q be a prime number. The problem is to find a short column vector $\mathbf{x} \in \mathbb{Z}^k$ such that:

$$A\mathbf{x} \equiv \mathbf{c} \pmod{q}.$$

The interesting case is when the number of unknowns k is larger than the number of equations ℓ. Following Section 3, one computes an arbitrary solution $\mathbf{y} \in \mathbb{Z}^k$ such that $A\mathbf{y} \equiv \mathbf{c} \pmod{q}$, for instance by finding a solution of a solvable system of linear equations over the integers (if the system is not solvable, then the original problem has no solution). And one computes a basis of the full-dimensional lattice L of all column vectors $\mathbf{z} \in \mathbb{Z}^k$ such that

$$A\mathbf{z} \equiv 0 \pmod{q}.$$

Then any short solution $\mathbf{x}$ to $A\mathbf{x} \equiv \mathbf{c} \pmod{q}$ corresponds to a lattice vector $\mathbf{y}-\mathbf{x} \in L$ close to $\mathbf{y}$. Thus, there is at most one $\mathbf{x} \in \mathbb{Z}^k$ such that $A\mathbf{x} \equiv \mathbf{c} \pmod{q}$ and $\|\mathbf{x}\| < \|L\|/2$. And if ever there is an unusually short vector $\mathbf{x} \in \mathbb{Z}^k$ such that $A\mathbf{x} \equiv \mathbf{c} \pmod{q}$ and $\|\mathbf{x}\| < \|L\| 2^{-k/2-1}$, then Babai's CVP approximation algorithm will disclose it, as in Section 3. It remains to lower bound the first minimum of the lattice.

One can see that the volume of L is an integer dividing q^ℓ, because it is the index of L in $\mathbb{Z}^k$. In fact, for most matrices A, one expects the volume to be exactly q^ℓ, so that:

$$\|L\| \approx \sqrt{\frac{k}{2\pi e}} q^{\ell/k}.$$

This estimate is not far from the truth, since for any fixed non-zero vector $\mathbf{z} \in \mathbb{Z}^k$ such that $\|\mathbf{x}\|_\infty < q$, the probability that $\mathbf{z} \in L$ (when A is uniformly distributed) is exactly $q^{-\ell}$. It follows that for most matrices, if ever there exists $\mathbf{x} \in \mathbb{Z}^k$ such that $A\mathbf{x} \equiv \mathbf{c} \pmod{q}$ and $\|\mathbf{x}\|$ roughly less than $q^{\ell/k} 2^{-k/2-1}$, then one can find such an $\mathbf{x}$ in polynomial time. For a precise statement, we refer to [52] who actually used a dual approach requiring transference theorems (which we do not need here). An interesting application is that if we know a few bits of each entry of an arbitrary solution $\mathbf{x}$ of a system of linear modular equations, then we can recover all of $\mathbf{x}$, because if the number of bits is sufficiently large, the problem is reduced to finding an unusually short solution of a system of linear modular equations. This was used to show the insecurity of truncated linear congruential pseudo-random number generators in [52].

The result can in fact be extended to a wider class of parameters, when the modulus q is not necessarily prime (see [52]), and when the equations may have different modulus (see [10]). We note that the exponent $-k/2$ can be suppressed when a CVP-oracle is available, which is the case when k is fixed. Furthermore, the previous reasoning not only shows how to find unusually short solutions, it also shows how to find reasonably short solutions when the matrix A is random. Indeed, a tighter analysis then shows that all the minima of the lattice L are in fact not too far away from $\sqrt{k/(2\pi e)} q^{\ell/k}$, so that all points are reasonably close to the lattice. In this case, one can find in polynomial time a vector $\mathbf{x} \in \mathbb{Z}^k$ such that $A\mathbf{x} \equiv \mathbf{c} \pmod{q}$ and $\|\mathbf{x}\|$ is very roughly less than $\sqrt{k/(2\pi e)} q^{\ell/k} 2^{k/2}$. This was used to attack certain RSA padding signature schemes in which the messages have a fixed pattern (see [104, 57]), and it was also used to complete the proof of security of the RSA–OAEP encryption scheme (see [53]).

However, the previous results are weak in a certain sense. First, the results depend strongly on the distribution of the coefficients of the linear equations. More precisely, the first minimum of the lattice can be arbitrary small, leading to possibly much weaker bounds: hence, one must perform a new analysis of the lattice for any system of equations which is not uniformly distributed. This was the case in [52] where linear congruential generators gave rise to special systems of equations. Furthermore, the exponential or slightly subexponential factors in the polynomial-time approximation of CVP imply that the bounds obtained are rather weak as the number k of unknowns increases. The situation is somewhat similar to that of knapsacks for which only knapsacks of very low density can provably be solved. This is one of the reasons why the attack of [104] was only heuristic. On the other hand, k was as small as 2 in [53], making provable results useful. We will see in the next section a particular case of a system of linear modular equations for which the generic method can be replaced by another lattice-based method.

4 The hidden number problem

4.1 Hardness of Diffie-Hellman bits

In [24], Boneh and Venkatesan used the LLL algorithm to solve the *hidden number problem*, which enables to prove the hardness of the most significant bits of secret keys in Diffie-Hellman and related schemes in prime fields. This was the first positive application of LLL in cryptology. Recall the Diffie-Hellman key exchange protocol [43]: Alice and Bob fix a finite cyclic G and a generator g. They respectively pick random $a, b \in [1, |G|]$ and exchange g^a and g^b. The secret key is g^{ab}. Proving the security of the protocol under "reasonable" assumptions has been a challenging problem in cryptography (see [15]). Computing the most significant bits of g^{ab} is as hard as computing g^{ab} itself, in the case of prime fields:

Theorem 2 (Boneh-Venkatesan). *Let q be an n-bit prime and g be a generator of $\mathbb{Z}_q^*$. Let $\varepsilon > 0$ be fixed, and set $\ell = \ell(n) = \lceil \varepsilon\sqrt{n} \rceil$. Suppose there exists an expected polynomial time (in n) algorithm $\mathcal{A}$, that on input q, g, g^a and g^b, outputs the ℓ most significant bits of g^{ab}. Then there is also an expected polynomial time algorithm that on input q, g, g^a, g^b and the factorization of $q-1$, computes all of g^{ab}.*

The above result is slightly different from [24], due to a small gap in the proof of [24] spotted by [63]. The same result holds for the least significant bits. For a more general statement when g is not necessarily a generator, and the factorization of $q-1$ is unknown, see [63]. For analogous results in other groups, we refer to [136] for finite fields and to [23] for the elliptic curve case.

The proof goes as follows. We are given some g^a and g^b, and want to compute g^{ab}. We repeatedly pick a random r until g^{a+r} is a generator of $\mathbb{Z}_q^*$ (testing is easy thanks to the factorization of $q-1$). For each r, the probability of success is $\phi(q-1)/(q-1) \geq C/\log\log q$. Next, we apply $\mathcal{A}$ to the points g^{a+r} and g^{b+t} for many random values of t, so that we learn the most significant bits of $g^{(a+r)b}g^{(a+r)t}$, where $g^{(a+r)t}$ is a random element of $\mathbb{Z}_q^*$ since g^{a+r} is a generator. Note that one can easily recover g^{ab} from $\alpha = g^{(a+r)b}$. The problem becomes the *hidden number problem* (HNP): given $t_1, \ldots, t_d$ chosen uniformly and independently at random in $\mathbb{Z}_q^*$, and $\mathrm{MSB}_\ell(\alpha t_i \bmod q)$ for all i, recover $\alpha \in \mathbb{Z}_q$. Here, $\mathrm{MSB}_\ell(x)$ for $x \in \mathbb{Z}_q$ denotes any integer z satisfying $|x - z| < q/2^{\ell+1}$.

To achieve the proof, Boneh and Venkatesan presented a simple solution to HNP when ℓ is not too small, by reducing HNP to a lattice closest vector problem. We sketch this solution in the next section.

4.2 Solving the hidden number problem by lattice reduction

Consider an HNP-instance: let $t_1, \ldots, t_d$ be chosen uniformly and independently at random in $\mathbb{Z}_q^*$, and $a_i = \mathrm{MSB}_\ell(\alpha t_i \bmod q)$ where $\alpha \in \mathbb{Z}_q$ is hidden. Clearly, the vector $\mathbf{t} = (t_1\alpha \bmod q, \ldots t_d\alpha \bmod q, \alpha/2^{\ell+1})$ belongs to the $(d+1)$-dimensional lattice $L = L(q, \ell, t_1, \ldots, t_d)$ spanned by the rows of the following matrix:

$$\begin{pmatrix} q & 0 & \cdots & 0 & 0 \\ 0 & q & \ddots & \vdots & \vdots \\ \vdots & \ddots & \ddots & 0 & \vdots \\ 0 & \ldots & 0 & q & 0 \\ t_1 & \ldots & \ldots & t_d & 1/2^{\ell+1} \end{pmatrix}$$

The vector $\mathbf{a} = (a_1, \ldots, a_d, 0)$ is very close to L, because it is very close to $\mathbf{t}$. Indeed, $\|\mathbf{t} - \mathbf{a}\| \leq q\sqrt{d+1}/2^{\ell+1}$. It is not difficult to show that any lattice point sufficiently close to $\mathbf{a}$ discloses the hidden number α, because sufficiently short lattice vectors must have their first d coordinates equal to zero (see [24, Theorem 5] or [110, 112]):

Lemma 3 (Uniqueness). *Set $d = 2\lceil\sqrt{\log q}\rceil$ and $\mu = \frac{1}{2}\sqrt{\log q} + 3$. Let α be in $\mathbb{Z}_q^*$. Choose integers $t_1, \ldots, t_d$ uniformly and independently at random in $\mathbb{Z}_q^*$. Let $\mathbf{a} = (a_1, \ldots, a_d, 0)$ be such that $|(\alpha t_i \bmod q) - a_i| < q/2^\mu$. Then with probability at least $\frac{1}{2}$, all $\mathbf{u} \in L$ with $\|\mathbf{u} - \mathbf{a}\| < \frac{q}{2^\mu}$ are of the form:*

$$\mathbf{u} = (t_1\beta \bmod q, \ldots t_d\beta \bmod q, \beta/2^{\ell+1}) \text{ where } \alpha \equiv \beta \pmod q.$$

Since $\mathbf{a}$ is close enough to L, Babai's nearest plane CVP approximation algorithm [8] yields a lattice point sufficiently close to $\mathbf{a}$, which leads to:

Theorem 4 (Boneh-Venkatesan). *Let α be in $\mathbb{Z}_q^*$. Let $\mathcal{O}$ be a function defined by $\mathcal{O}(t) = \mathrm{MSB}_\ell(\alpha t \bmod q)$ with $\ell = \lceil\sqrt{\log q}\rceil + \lceil\log\log q\rceil$. There exists a deterministic polynomial time algorithm $\mathcal{A}$ which, on input $t_1, \ldots, t_d, \mathcal{O}(t_1), \ldots, \mathcal{O}(t_d)$, outputs α with probability at least $1/2$ over $t_1, \ldots, t_d$ chosen uniformly and independently at random from $\mathbb{Z}_q^*$, where $d = 2\lceil\sqrt{\log q}\rceil$.*

Thus, the hidden number problem can be solved using $\ell = \sqrt{\log q} + \log\log q$ bits. Using the best polynomial-time CVP approximation algorithm known, this can be asymptotically improved to $O(\sqrt{\log q}\log\log\log q/\log\log q)$. Theorem 2 is a simple consequence.

We note that Theorem 4 could have alternatively be obtained from the generic method described in Section 3.4. Indeed, the hidden number problem can be viewed as a system of d modular linear equations in the $d+1$ unknowns α and $(\alpha t_i \bmod q) - \mathrm{MSB}_\ell(\alpha t_i \bmod q)$ where $1 \leq i \leq d$. Among these $d+1$ unknowns, only α may be large. One may transform the system to eliminate the possibly large unknown α. One then obtains a new system of $d-1$ modular linear equations in the d unknowns $(\alpha t_i \bmod q) - \mathrm{MSB}_\ell(\alpha t_i \bmod q)$ all smaller than $q/2^{\ell+1}$ in absolute value. Although this system does not correspond to a uniformly distributed matrix, one can easily obtain the same lower bound on the first minimum of the lattice as in the random case (see Section 3.4). It follows that one can find the (unique) small solution of the system in polynomial time (and thus, α) provided that roughly:

$$\frac{q}{2^{\ell+1}} \leq q^{(d-1)/d} 2^{-d/2-1},$$

that is $\ell \geq d/2 + 1 + \log(q)/d$, where the right-hand term is minimized for $d \approx \sqrt{2 \log q}$, leading to ℓ larger than roughly $\sqrt{2 \log q}$. Thus, one can obtain essentially the same bounds.

4.3 Variants of the hidden number problem

It was recently realized that the condition that the t_i's are uniformly distributed is often too restrictive for applications. The previous solution to the hidden number problem can in fact be extended to cases where the distribution of the t_i's is not necessarily perfectly uniform (see [63, 111]). A precise definition of this relaxed uniformity property can be made with the classical notion of discrepancy (see [111] for more details). To apply the solution to this generalized hidden number problem, it suffices to show that the distribution of the t_i's is sufficiently uniform, which is usually obtained by exponential sum techniques (see [63, 111, 112, 48, 130, 129] for some examples).

One may also extend the solution to the hidden number problem to the case when an oracle for CVP (in the Euclidean norm or the infinity norm) is available, which significantly decreases the number of necessary bits (see [110, 111]). This is useful to estimate what can be achieved in practice, especially when the lattice dimension is small. It turns out that the required number of bits becomes $O(\log \log q)$ and 2 respectively, with a CVP-oracle and a CVP_∞-oracle.

One may also study the hidden number problem with arbitrary bits instead of most significant bits. It is easy to see that the HNP with ℓ least significant bits can be reduced to the original HNP with ℓ most significant bits, but the situation worsens with arbitrary bits. By multiplying the t_i's with an appropriate number independent of the t_i's (see [111]), one obtains a deterministic polynomial-time reduction from the HNP with ℓ consecutive bits at a known position to the original HNP with $\ell/2$ most significant bits (the prime field $\mathbb{Z}_q$ and the number of random multipliers remain the same). This appropriate number can be found either by continued fractions or lattice reduction in dimension 2. More generally, by using high-dimensional lattice reduction, it is not difficult to show that there is a deterministic polynomial-time reduction from the HNP with ℓ arbitrary bits at known positions such that the number of blocks of consecutive unknown bits is m, to the original HNP with $\ell/m + 1 - \log m$ most significant bits. Thus, the HNP with arbitrary bits seems to be harder, especially when there are many blocks of consecutive unknown bits.

Finally, variants of the hidden number problem in settings other than prime fields have been studied in [130, 129, 23].

4.4 Lattice attacks on DSA

Interestingly, the previous solution of the hidden number problem also has a dark side: it leads to a simple attack against the Digital Signature Algorithm [106, 95] (DSA) in special settings (see [73, 110]). Recall that the DSA uses a public element $g \in \mathbb{Z}_p$ of order q, a 160-bit prime dividing $p-1$ where p is a large prime (at least 512 bits). The signer has a secret key $\alpha \in \mathbb{Z}_q^*$ and a public key $\beta = g^\alpha \bmod p$. The DSA signature of a message m is $(r, s) \in \mathbb{Z}_q^2$ where $r = (g^k \bmod p) \bmod q$, $s = k^{-1}(h(m) + \alpha r) \bmod q$, h is SHA-1 hash function and k is a random element in $\mathbb{Z}_q^*$ chosen at each signature.

It is well-known that the secret key α can easily be recovered if the random nonce k is disclosed, or if k is produced by a cryptographically weak pseudorandom generator such as a linear congruential generator with known parameters [10] and a few signatures are available. Recently, Howgrave-Graham and Smart [73] noticed that Babai's nearest plane CVP algorithm could heuristically recover α, provided that sufficiently many signatures and sufficiently many bits of the corresponding nonces k are known. This is not surprising, because the underlying problem is in fact a generalized hidden number problem.

Indeed, assume that for d signatures (r_i, s_i) of messages m_i, the ℓ least significant bits of the random nonce k_i are known to the attacker: one knows $a_i < 2^\ell$ such that $k_i - a_i$ is of the form $2^\ell b_i$. Then $\alpha r_i \equiv s_i(a_i + b_i 2^\ell) - h(m_i) \pmod q$, which can be rewritten as: $\alpha r_i 2^{-\ell} s_i^{-1} \equiv (a_i - s_i^{-1} h(m_i)) \cdot 2^{-\ell} + b_i \pmod q$. Letting $t_i = r_i 2^{-\ell} s_i^{-1} \bmod q$, one sees that $\mathrm{MSB}_\ell(\alpha t_i \bmod q)$ is known. Recovering the secret key α is therefore a generalized hidden number problem in which the t_i's are not assumed to be independent and uniformly distributed over $\mathbb{Z}_q$, but are of the form $r_i 2^{-\ell} s_i^{-1}$ where the underlying k_i's are independent and uniformly distributed over $\mathbb{Z}_q^*$. Nguyen and Shparlinski [111] proved that under a reasonable assumption on the hash function, the t_i's are sufficiently uniform to make the corresponding hidden number problem provably tractable with the same number of bits as in Theorem 4, that is, essentially $\sqrt{\log q}$. Since lattice reduction algorithms can behave much better than theoretically expected, one may even hope to solve CVP exactly, yielding better bounds to Theorem 4. For the case of a 160-bit prime q as in DSA, one obtains that the DSA–HNP can be solved using respectively $\ell = 2$ bits and $d = 160$, or $\ell = 6$ bits and $d = 100$ respectively, when an oracle for CVP_∞ or CVP is available (see [110, 112]). In fact, the bounds are better in practice. It turns out that using standard lattice reduction algorithms implemented in Shoup's NTL library [127], one can often solve HNP for a 160-bit prime q using $\ell = 3$ bits and $d = 100$ (see [110, 112]).

Naturally, this attack can also be applied to similar signature algorithms (see [111]), such as the elliptic curve variant of DSA (see [112]), or the Nyberg-Rueppel signature scheme and related schemes (see [48]). The only difference is that one needs to establish the uniformity of different types of multipliers. This usually requires different kinds of exponential sums.

5 Lattice-based cryptography

We review state-of-the-art results on the main lattice-based cryptosystems. To keep the presentation simple, descriptions of the schemes are intuitive, referring to the original papers for more details. Only one of these schemes (the GGH cryptosystem [61]) explicitly works with lattices.

5.1 The Ajtai-Dwork cryptosystem

Description. The Ajtai-Dwork cryptosystem [5] (AD) works in $\mathbb{R}^n$, with some finite precision depending on n. Its security is based on a variant of SVP.

The private key is a uniformly chosen vector u in the n-dimensional unit ball. One then defines a distribution $\mathcal{H}_u$ of points $\mathbf{a}$ in a large n-dimensional cube such that the dot product $\langle \mathbf{a}, \mathbf{u} \rangle$ is very close to $\mathbb{Z}$.

The public key is obtained by picking $\mathbf{w}_1, \ldots, \mathbf{w}_n, \mathbf{v}_1, \ldots, \mathbf{v}_m$ (where $m = n^3$) independently at random from the distribution $\mathcal{H}_u$, subject to the constraint that the parallelepiped w spanned by the $\mathbf{w}_i$'s is not flat. Thus, the public key consists of a polynomial number of points close to a collection of parallel affine hyperplanes, which is kept secret.

The scheme is mainly of theoretical purpose, as encryption is bit-by-bit. To encrypt a '0', one randomly selects $b_1, \ldots, b_m$ in $\{0, 1\}$, and reduces $\sum_{i=1}^{m} b_i \mathbf{v}_i$ modulo the parallelepiped w. The vector obtained is the ciphertext. The ciphertext of '1' is just a randomly chosen vector in the parallelepiped w. To decrypt a ciphertext $\mathbf{x}$ with the private key $\mathbf{u}$, one computes $\tau = \langle \mathbf{x}, \mathbf{u} \rangle$. If τ is sufficiently close to $\mathbb{Z}$, then $\mathbf{x}$ is decrypted as '0', and otherwise as '1'. Thus, an encryption of '0' will always be decrypted as '0', and an encryption of '1' has a small probability to be decrypted as '0'. These decryption errors can be removed (see [60]).

Security. The Ajtai-Dwork [5] cryptosystem received wide attention due to a surprising security proof based on worst-case assumptions. Indeed, it was shown that any probabilistic algorithm distinguishing encryptions of a '0' from encryptions of a '1' with some polynomial advantage can be used to solve SVP in any n-dimensional lattice with gap λ_2/λ_1 larger than n^8. There is a converse, due to Nguyen and Stern [115]: one can decrypt in polynomial time with high probability, provided an oracle that approximates SVP to within $n^{0.5-\varepsilon}$, or one that approximates CVP to within $n^{1.33}$. It follows that the problem of decrypting ciphertexts is unlikely to be NP-hard, due to the result of Goldreich-Goldwasser [58].

Nguyen and Stern [115] further presented a heuristic attack to recover the secret key. Experiments suggest that the attack is likely to succeed up to at least $n = 32$. For such parameters, the system is already impractical, as the public key requires 20 megabytes and the ciphertext for each bit has bit-length 6144. This shows that unless major improvements[3] are found, the Ajtai-Dwork cryptosystem is only of theoretical importance.

Cryptanalysis overview. At this point, the reader might wonder how lattices come into play, since the description of AD does not involve lattices. Any ciphertext of '0' is a sum of $\mathbf{v}_i$'s minus some integer linear combination of the $\mathbf{w}_i$'s. Since the parallelepiped spanned by the $\mathbf{w}_i$'s is not too flat, the coefficients of the linear combination are relatively small. On the other hand, any linear combination of the $\mathbf{v}_i$'s and the $\mathbf{w}_i$'s with small coefficients is close to the hidden hyperplanes. This enables to build a particular lattice of dimension $n+m$ such that any ciphertext of '0' is in some sense close to the lattice, and reciprocally, any point sufficiently close to the lattice gives rise to a ciphertext of '0'. Thus, one can decrypt ciphertexts provided an oracle that approximates CVP sufficiently well. The analogous version for SVP uses related ideas, but is technically more complicated. For more details, see [115].

The attack to recover the secret key can be described quite easily. One knows that each $\langle \mathbf{v}_i, \mathbf{u} \rangle$ is close to some unknown integer V_i. It can be shown that any sufficiently short linear combination of the $\mathbf{v}_i$'s give information on the V_i's. More precisely, if $\sum_i \lambda_i \mathbf{v}_i$ is sufficiently short and the λ_i's are sufficiently small, then $\sum_i \lambda_i V_i = 0$ (because it is a too small integer). Note that the V_i's are disclosed if enough such equations are found. And each V_i gives an approximate linear equation satisfied by the coefficients of the secret key $\mathbf{u}$. Thus, one can compute a sufficiently good approximation of $\mathbf{u}$ from the V_i's. To find the V_i's, we produce many short combinations $\sum_i \lambda_i \mathbf{v}_i$ with small λ_i's, using lattice reduction. Heuristic arguments can justify that there exist enough such combinations. Experiments showed that the assumption was reasonable in practice.

5.2 The Goldreich-Goldwasser-Halevi cryptosystem

The Goldreich-Goldwasser-Halevi cryptosystem [61] (GGH) can be viewed as a lattice-analog to the McEliece [94] cryptosystem based on algebraic coding theory. In both schemes, a ciphertext is the addition of a random noise vector to a vector corresponding to the plaintext. The public key and the private key are two representations of the same object (a lattice for GGH, a linear code for McEliece). The private key has a particular structure allowing to cancel noise vectors up to a certain bound. However, the domains in which all these operations take place are quite different.

[3] A variant of AD with less message expansion was proposed in [32], however without any security proof. It mixes AD with a knapsack.

Description. The GGH scheme works in $\mathbb{Z}^n$. The private key is a non-singular $n \times n$ integral matrix R, with very short row vectors[4] (entries polynomial in n). The lattice L is the full-dimensional lattice in $\mathbb{Z}^n$ spanned by the rows of R. The basis R is then transformed to a non-reduced basis B, which will be public. In the original scheme, B is the multiplication of R by sufficiently many small unimodular matrices. Computing a basis as "good" as the private basis R, given only the non-reduced basis B, means approximating SBP.

The message space is a "large enough" cube in $\mathbb{Z}^n$. A message $\mathbf{m} \in \mathbb{Z}^n$ is encrypted into $\mathbf{c} = \mathbf{m}B + \mathbf{e}$ where $\mathbf{e}$ is an error vector uniformly chosen from $\{-\sigma, \sigma\}^n$, where σ is a security parameter. A ciphertext $\mathbf{c}$ is decrypted as $\lfloor \mathbf{c}R^{-1} \rceil RB^{-1}$ (note: this is Babai's round method [8] to solve CVP). But an eavesdropper is left with the CVP-instance defined by $\mathbf{c}$ and B. The private basis R is generated in such a way that the decryption process succeeds with high probability. The larger σ is, the harder the CVP-instances are expected to be. But σ must be small for the decryption process to succeed.

Improvements. In the original scheme, the public matrix B is the multiplication of the secret matrix by sufficiently many unimodular matrices. This means that without appropriate precaution, the public matrix may be as large as $O(n^3 \log n)$ bits. Micciancio [99, 101] therefore suggested to define instead B as the Hermite normal form (HNF) of R. Recall that the HNF of an integer square matrix R in row notation is the unique lower triangular matrix with coefficients in $\mathbb{N}$ such that: the rows span the same lattice as R, and any entry below the diagonal is strictly less than the diagonal entry in its column. Here, one can see that the HNF of R is $O(n^2 \log n)$ bits, which is much better but still big. When using the HNF, one should encode messages into the error vector $\mathbf{e}$ instead of a lattice point, because the HNF is unbalanced. The ciphertext is defined as the reduction of $\mathbf{e}$ modulo the HNF, and hence uses less than $O(n \log n)$ bits. One can easily prove that the new scheme (which is now deterministic) cannot be less secure than the original GGH scheme (see [99, 101]).

Security. GGH has no proven worst-case/average-case property, but it is much more efficient than AD. Specifically, for security parameter n, key-size and encryption time can be $O(n^2 \log n)$ for GGH (note that McEliece is slightly better though), *vs.* at least $O(n^4)$ for AD. For RSA and El-Gamal systems, key size is $O(n)$ and computation time is $O(n^3)$. The authors of GGH argued that the increase in size of the keys was more than compensated by the decrease in computation time. To bring confidence in their scheme, they published on the Internet a series of five numerical challenges [59], in dimensions 200, 250, 300, 350 and 400. In each of these challenges, a public key and a ciphertext were given, and the challenge was to recover the plaintext.

The GGH scheme is now considered broken, at least in its original form, due to an attack recently developed by Nguyen [108]. As an application, using

[4] A different construction for R based on tensor product was proposed in [50].

small computing power and Shoup's NTL library [127], Nguyen was able to solve all the GGH challenges, except the last one in dimension 400. But already in dimension 400, GGH is not very practical: in the 400-challenge, the public key takes 1.8 Mbytes without HNF or 124 Kbytes using the HNF.[5]

Nguyen's attack used two "qualitatively different" weaknesses of GGH. The first one is inherent to the GGH construction: the error vectors used in the encryption process are always much shorter[6] than lattice vectors. This makes CVP-instances arising from GGH easier than general CVP-instances. The second weakness is the particular form of the error vectors in the encryption process. Recall that $\mathbf{c} = \mathbf{m}B + \mathbf{e}$ where $\mathbf{e} \in \{\pm\sigma\}^n$. The form of $\mathbf{e}$ was apparently chosen to maximize the Euclidean norm under requirements on the infinity norm. However, if we let $\mathbf{s} = (\sigma, \ldots, \sigma)$ then $\mathbf{c} + \mathbf{s} \equiv \mathbf{m}B \pmod{2\sigma}$, which allows to guess $\mathbf{m}$ mod 2σ. Then the original closest vector problem can be reduced to finding a lattice vector within (smaller) distance $\mathbf{e}/(2\sigma)$ from $(\mathbf{c} - (\mathbf{m} \mod 2\sigma)B)/(2\sigma)$. The simplified closest vector problem happens to be within reach (in practice) of current lattice reduction algorithms, thanks to the embedding strategy that heuristically reduces CVP to SVP. We refer to [108] for more information.

It is easy to fix the second weakness by selecting the entries of the error vector $\mathbf{e}$ at random in $\{-\sigma, \ldots, +\sigma\}$ instead of $\{\pm\sigma\}$. However, one can argue that the resulting GGH system would still not be much practical, even using [99, 101]. Indeed, Nguyen's experiments [108] showed that SVP could be solved in practice up to dimensions as high as 350, for (certain) lattices with gap as small as 10. To be competitive, the new GGH system would require the hardness (in lower dimensions due to the size of the public key, even using [99]) of SVP for certain lattices of only slightly smaller gap, which means a rather smaller improvement in terms of reduction. Note also that those experiments do not support the practical hardness of Ajtai's variant of SVP in which the gap is polynomial in the lattice dimension. Besides, it is not clear how to make decryption efficient without a huge secret key (Babai's rounding requires the storage of R^{-1} or a good approximation, which could be in [61] over 1 Mbytes in dimension 400).

5.3 The NTRU cryptosystem

Description. The NTRU cryptosystem [69], proposed by Hoffstein, Pipher and Silverman, works in the ring $R = \mathbb{Z}[X]/(X^N - 1)$. An element $F \in R$ is seen as a polynomial or a row vector: $F = \sum_{i=0}^{N-1} F_i x^i = [F_0, F_1, \ldots, F_{N-1}]$. To select keys, one uses the set $\mathcal{L}(d_1, d_2)$ of polynomials $F \in R$ such that d_1 coefficients are equal to 1, d_2 coeffients are equal to -1, and the rest are zero. There are two small coprime moduli $p < q$: a possible choice is $q = 128$ and $p = 3$. There are also three integer parameters d_f, d_g and d_ϕ quite a bit smaller than the prime number N (which is around a few hundreds).

[5] The challenges do not use the HNF, as they were proposed before [99]. Note that 124 Kbytes is about twice as large as McEliece for the recommended parameters.

[6] In all GGH-like constructions known, the error vector is always at least twice as short.

The private keys are $f \in \mathcal{L}(d_f, d_f - 1)$ and $g \in \mathcal{L}(d_g, d_g)$. With high probability, f is invertible mod q. The public key $h \in R$ is defined as $h = g/f \mod q$. A message $m \in \{-(p-1)/2 \cdots + (p-1)/2\}^N$ is encrypted into: $e = (p\phi * h + m) \mod q$, where ϕ is randomly chosen in $\mathcal{L}(d_\phi, d_\phi)$. The user can decrypt thanks to the congruence $e * f \equiv p\phi * g + m * f \pmod{q}$, where the reduction is centered (one takes the smallest residue in absolute value). Since ϕ, f, g and m all have small coefficients and many zeroes (except possibly m), that congruence is likely to be a polynomial equality over $\mathbb{Z}$. By further reducing $e * f$ modulo p, one thus recovers $m * f \mod q$, hence m.

Security. The best attack known against NTRU is based on lattice reduction, but this does not mean that lattice reduction is necessary to break NTRU. The simplest lattice-based attack can be described as follows. Coppersmith and Shamir [40] noticed that the target vector $f \| g \in \mathbb{Z}^{2N}$ (the symbol $\|$ denotes vector concatenation) belongs to the following natural lattice:

$$L_{CS} = \{F \| G \in \mathbb{Z}^{2N} \mid F \equiv h * G \mod q \text{ where } F, G \in R\}.$$

It is not difficult to see that L_{CS} is a full-dimensional lattice in $\mathbb{Z}^{2N}$, with volume q^N. The volume suggests that the target vector is a shortest vector of L_{CS} (but with small gap), so that a SVP-oracle should heuristically output the private keys f and g. However, based on numerous experiments with Shoup's NTL library [127], the authors of NTRU claimed in [69] that all such attacks are exponential in N, so that even reasonable choices of N ensure sufficient security. The parameter N must be prime, otherwise the lattice attacks can be improved due to the factorization of $X^N - 1$ (see [56]). Note that the keysize of NTRU is only $O(N \log q)$, which makes NTRU the leading candidate among knapsack-based and lattice-based cryptosystems, and allows high lattice dimensions. It seems that better attacks or better lattice reduction algorithms are required in order to break NTRU. To date, none of the numerical challenges proposed in [69] has been solved. However, it is probably too early to tell whether or not NTRU is secure. Note that NTRU, like RSA, should only be used with appropriate preprocessing. Indeed, NTRU without padding cannot be semantically secure since $e(1) \equiv m(1) \pmod{q}$ as polynomials, and it is easily malleable using multiplications by X of polynomials (circular shifts). And there exist simple chosen ciphertext attacks [74] that can recover the secret key.

6 Finding small roots of low-degree polynomial equations

We survey an important application of lattice reduction found in 1996 by Coppersmith [38, 39], and its developments. These results illustrate the power of linearization combined with lattice reduction.

6.1 Univariate modular equations

The general problem of solving univariate polynomial equations modulo some integer N of unknown factorization seems to be hard. Indeed, notice that for some polynomials, it is equivalent to the knowledge of the factorization of N. And the particular case of extracting e-th roots modulo N is the problem of decrypting ciphertexts in the RSA cryptosystem, for an eavesdropper. Curiously, Coppersmith [38] showed using LLL that the special problem of finding small roots is easy:

Theorem 5 (Coppersmith). *Let P be a monic polynomial of degree δ in one variable modulo an integer N of unknown factorization. Then one can find in time polynomial in $(\log N, \delta)$ all integers x_0 such that $P(x_0) \equiv 0 \ (mod N)$ and $|x_0| \leq N^{1/\delta}$.*

Related (but weaker) results appeared in the eighties [66, 133].[7] Incidentally, the result implies that the number of roots less than $N^{1/\delta}$ is polynomial, which was also proved in [81] (using elementary techniques).

We sketch a proof of Theorem 5, in the spirit of Howgrave-Graham [70], who simplified Coppersmith's original proof (see also [76]) by working in the dual lattice of the lattice originally considered by Coppersmith. More details can be found in [39]. Coppersmith's method reduces the problem of finding small modular roots to the (easy) problem of solving polynomial equations over $\mathbb{Z}$. More precisely, it applies lattice reduction to find an integral polynomial equation satisfied by all small modular roots of P. The intuition is to linearize all the equations of the form $x^i P(x)^j \equiv 0 \pmod{N^j}$ for appropriate integral values of i and j. Such equations are satisfied by any solution of $P(x) \equiv 0 \pmod{N}$. Small solutions x_0 give rise to unusually short solutions to the resulting linear system. To transform modular equations into integer equations, the following elementary lemma[8] is used, with the notation $\|r(x)\| = \sqrt{\sum a_i^2}$ for any polynomial $r(x) = \sum a_i x^i \in \mathbb{Q}[x]$:

Lemma 6. *Let $r(x) \in \mathbb{Q}[x]$ be a polynomial of degree $< n$ and let X be a positive integer. Suppose $\|r(xX)\| < 1/\sqrt{n}$. If $r(x_0) \in \mathbb{Z}$ with $|x_0| < X$, then $r(x_0) = 0$ holds over the integers.*

This is just because any sufficiently small integer must be zero. Now the trick is to, given a parameter h, consider the $n = (h+1)\delta$ polynomials $q_{u,v}(x) = x^u (P(x)/N)^v$, where $0 \leq u \leq \delta - 1$ and $0 \leq v \leq h$. Notice that the degree of $q_{u,v}(x)$ is strictly less than n, and that the evaluation of $q_{u,v}(x)$ at any root x_0 of $P(x)$ modulo N is necessarily an integer. The same is true for any integral linear combination $r(x)$ of the $q_{u,v}(x)$'s. If such a combination $r(x)$ further satisfies $\|r(xX)\| < 1/\sqrt{n}$, then by Lemma 6, solving the equation $r(x) = 0$ over $\mathbb{Z}$ yields all roots of $P(x)$ modulo N less than X in absolute value. This suggests to look for a short vector in the lattice corresponding to the $q_{u,v}(xX)$'s. More precisely, define the $n \times n$ matrix M whose i-th row consists of the coefficients of $q_{u,v}(xX)$, starting by the low-degree terms, where $v = \lfloor (i-1)/\delta \rfloor$ and $u = (i-1) - \delta v$.

[7] Håstad [66] presented his result in terms of system of low-degree modular equations, but he actually studies the same problem, and his approach achieves the weaker bound $N^{2/(\delta(\delta+1))}$.

[8] A similar lemma is used in [66]. Note also the resemblance with [88, Prop. 2.7].

Notice that M is lower triangular, and a simple calculation leads to $\det(M) = X^{n(n-1)/2}N^{-nh/2}$. We apply an LLL-reduction to the full-dimensional lattice spanned by the rows of M. The first vector of the reduced basis corresponds to a polynomial of the form $r(xX)$, and has Euclidean norm $\|r(xX)\|$. The theoretical bounds of the LLL algorithm ensure that:

$$\|r(xX)\| \leq 2^{(n-1)/4} \det(M)^{1/n} = 2^{(n-1)/4} X^{(n-1)/2} N^{-h/2}.$$

Recall that we need $\|r(xX)\| \leq 1/\sqrt{n}$ to apply the lemma. Hence, for a given h, the method is guaranteed to find modular roots up to X if:

$$X \leq \frac{1}{\sqrt{2}} N^{h/(n-1)} n^{-1/(n-1)}.$$

The limit of the upper bound, when h grows to ∞, is $\frac{1}{\sqrt{2}}N^{1/\delta}$. Theorem 5 follows from an appropriate choice of h. This result is practical (see [42, 71] for experimental results) and has many applications. It can be used to attack RSA encryption when a very low public exponent is used (see [16] for a survey). Boneh *et al.* [21] applied it to factor efficiently numbers of the form $N = p^r q$ for large r. Boneh [17] used a variant to find smooth numbers in short interval. See also [13] for an application to Chinese remaindering in the presence of noise, and [72] to find approximate integer common divisors. Curiously, Coppersmith's theorem was also recently used in security proofs of factoring-based schemes (see [128, 18]).

Remarks. Theorem 5 is trivial if $P(x) = x^\delta + c$. Note also that one cannot hope to improve the (natural) bound $N^{1/\delta}$ for all polynomials and all moduli N. Indeed, for the polynomial $P(x) = x^\delta$ and $N = p^\delta$ where p is prime, the roots of P mod N are the multiples of p. Thus, one cannot hope to find all the small roots (slightly) beyond $N^{1/\delta} = p$, because there are too many of them. This suggests that even an SVP-oracle (instead of LLL) should not help Theorem 5 in general, as evidenced by the value of the lattice volume (the fudge factor $2^{(n-1)/4}$ yielded by LLL is negligible compared to $\det(M)^{1/n}$). It was recently noticed in [13] that if one only looks for the smallest root mod N, an SVP-oracle can improve the bound $N^{1/\delta}$ for very particular moduli (namely, squarefree N of known factorization, without too small factors). Note that in such cases, finding modular roots can still be difficult, because the number of modular roots can be exponential in the number of prime factors of N. Coppersmith discusses potential improvements in [39].

6.2 Multivariate modular equations

Interestingly, Theorem 5 can heuristically extend to multivariate polynomial modular equations. Assume for instance that one would like to find all small roots of $P(x, y) \equiv 0 \pmod N$, where $P(x, y)$ has total degree δ and has at least one monic monomial $x^\alpha y^{\delta-\alpha}$ of maximal total degree. If one could obtain two algebraically independent integral polynomial equations satisfied by all sufficiently small modular roots (x, y), then one could compute (by resultant) a univariate integral polynomial equation satisfied by x, and hence find efficiently all small (x, y). To find such equations, one can use an analogue of lemma 6 to bivariate polynomials, with the (natural) notation $\|r(x,y)\| = \sqrt{\sum_{i,j} a_{i,j}^2}$ for $r(x,y) = \sum_{i,j} a_{i,j} x^i y^j$:

Lemma 7. *Let $r(x,y) \in \mathbb{Q}[x,y]$ be a sum of at most w monomials. Assume $\|r(xX, yY)\| < 1/\sqrt{w}$ for some $X, Y \geq 0$. If $r(x_0, y_0) \in \mathbb{Z}$ with $|x_0| < X$ and $|y_0| < Y$, then $r(x_0, y_0) = 0$ holds over the integers.*

By analogy, one chooses a parameter h and select $r(x,y)$ as a linear combination of the polynomials $q_{u_1,u_2,v}(x,y) = x^{u_1} y^{u_2} (P(x,y)/N)^v$, where $u_1 + u_2 + \delta v \leq h\delta$ and $u_1, u_2, v \geq 0$ with $u_1 < \alpha$ or $u_2 < \delta - \alpha$. Such polynomials have total degree less than $h\delta$, and therefore are linear combinations of the $n = (h\delta+1)(h\delta+2)/2$ monic monomials of total degree $\leq \delta h$. Due to the condition $u_1 < \alpha$ or $u_2 < \delta - \alpha$, such polynomials are in bijective correspondence with the n monic monomials (associate to $q_{u_1,u_2,v}(x,y)$ the monomial $x^{u_1+v\alpha} y^{u_2+v(\delta-\alpha)}$). One can represent the polynomials as n-dimensional vectors in such a way that the $n \times n$ matrix consisting of the $q_{u_1,u_2,v}(xX, yY)$'s (for some ordering) is lower triangular with coefficients $N^{-v} X^{u_1+v\delta} y^{u_2+v(\delta-\alpha)}$ on the diagonal.

Now consider the first two vectors $r_1(xX, yY)$ and $r_2(xX, yY)$ of an LLL-reduced basis of the lattice spanned by the rows of that matrix. Since the rational $q_{u_1,u_2,v}(x_0, y_0)$ is actually an integer for any root (x_0, y_0) of $P(x,y)$ modulo N, we need $\|r_1(xX, yY)\|$ and $\|r_2(xX, yY)\|$ to be less than $1/\sqrt{n}$ to apply Lemma 7. A (tedious) computation of the triangular matrix determinant enables to prove that $r_1(x,y)$ and $r_2(x,y)$ satisfy that bound when $XY < N^{1/\delta - \varepsilon}$ and h is sufficiently large (see [76]). Thus, one obtains two integer polynomial bivariate equations satisfied by all small modular roots of $P(x,y)$.

The problem is that, although such polynomial equations are linearly independent as vectors, they might be algebraically dependent, making the method heuristic. This heuristic assumption is unusual: many lattice-based attacks are heuristic in the sense that they require traditional lattice reduction algorithms to behave like SVP-oracles. An important open problem is to find sufficient conditions to make Coppersmith's method provable for bivariate (or multivariate) equations. Note that the method cannot work all the time. For instance, the polynomial $x - y$ has clearly too many roots over $\mathbb{Z}^2$ and hence too many roots mod any N (see [38] for more general counterexamples).

Such a result may enable to prove several attacks which are for now, only heuristic. Indeed, there are applications to the security of the RSA encryption scheme when a very low public exponent or a low private exponent is used (see [16] for a survey), and related schemes such as the KMOV cryptosystem (see [12]). In particular, the experimental evidence of [19, 12, 46] shows that the method is very effective in practice for certain polynomials.

Remarks. In the case of univariate polynomials, there was basically no choice over the polynomials $q_{u,v}(x) = x^u(P(x)/N)^v$ used to generate the appropriate univariate integer polynomial equation satisfied by all small modular roots. There is much more freedom with bivariate modular equations. Indeed, in the description above, we selected the indices of the polynomials $q_{u_1,u_2,v}(x,y)$ in such a way that they corresponded to all the monomials of total degree $\leq h\delta$, which form a triangle in $\mathbb{Z}^2$ when a monomial x^iy^j is represented by the point (i,j). This corresponds to the general case where a polynomial may have several monomials of maximal total degree. However, depending on the shape of the polynomial $P(x,y)$ and the bounds X and Y, other regions of (u_1,u_2,v) might lead to better bounds.

Assume for instance $P(x,y)$ is of the form $x^{\delta_x}y^{\delta_y}$ plus a linear combination of x^iy^j's where $i \leq \delta_x$, $j \leq \delta_y$ and $i+j < \delta_x+\delta_y$. Intuitively, it is better to select the (u_1,u_2,v)'s to cover the rectangle of sides $h\delta_x$ and $h\delta_y$ instead of the previous triangle, by picking all $q_{u_1,u_2,v}(x,y)$ such that $u_1 + v\delta_x \leq h\delta_x$ and $u_2 + v\delta_y \leq h\delta_y$, with $u_1 < \delta_x$ or $u_2 < \delta_y$. One can show that the polynomials $r_1(x,y)$ and $r_2(x,y)$ obtained from the first two vectors of an LLL-reduced basis of the appropriate lattice satisfy Lemma 7, provided that h is sufficiently large, and the bounds satisfy $X^{\delta_x}Y^{\delta_y} \leq N^{2/3-\varepsilon}$. Boneh and Durfee [19] applied similar and other tricks to a polynomial of the form $P(x,y) = xy + ax + b$. This allowed better bounds than the generic bound, leading to improved attacks on RSA with low secret exponent (see also [46] for an extension to the trivariate case, useful when the RSA primes are unbalanced).

6.3 Multivariate integer equations

The general problem of solving multivariate polynomial equations over $\mathbb{Z}$ is also hard, as integer factorization is a special case. Coppersmith [38] showed that a similar[9] lattice-based approach can be used to find small roots of bivariate polynomial equations over $\mathbb{Z}$:

Theorem 8 (Coppersmith). *Let $P(x,y)$ be a polynomial in two variables over $\mathbb{Z}$, of maximum degree δ in each variable separately, and assume the coefficients of f are relatively prime as a set. Let X,Y be bounds on the desired solutions x_0,y_0. Define $\hat{P}(x,y) = P(Xx,Yy)$ and let D be the absolute value of the largest coefficient of $\hat{P}$. If $XY < D^{2/(3\delta)}$, then in time polynomial in $(\log D, \delta)$, we can find all integer pairs (x_0,y_0) such that $P(x_0,y_0) = 0$, $|x_0| < X$ and $|y_0| < Y$.*

Again, the method extends heuristically to more than two variables, and there can be improved bounds depending on the shape[10] of the polynomial (see [38]). Theorem 8 was introduced to factor in polynomial time an RSA–modulus[11]

[9] However current proofs are somehow more technical than for Theorem 5. A simplification analogue to what has been obtained for Theorem 5 would be useful.

[10] The coefficient 2/3 is natural from the remarks at the end of the previous section for the bivariate modular case. If we had assumed P to have total degree δ, the bound would be $XY < D^{1\delta}$.

[11] p and q are assumed to have similar size.

$N = pq$ provided that half of the (either least or most significant) bits of either p or q are known (see [38, 17, 20]). This was sufficient to break an ID-based RSA encryption scheme proposed by Vanstone and Zuccherato [134]. Boneh *et al.* [20] provide another application, for recovering the RSA secret key when a large fraction of the bits of the secret exponent is known. Curiously, none of the applications cited above happen to be "true" applications of Theorem 8. It was later realized in [71, 21] that those results could alternatively be obtained from a (simple) variant of the univariate modular case (Theorem 5).

7 Lattices and RSA

Section 3 and 6 suggest to clarify the links existing between lattice reduction and RSA [120], the most famous public-key cryptosystem. We refer to [95] for an exposition of RSA, and to [16] for a survey of attacks on RSA encryption. Recall that in RSA, one selects two prime numbers p and q of approximately the same size. The number $N = pq$ is public. One selects an integer d coprime with $\phi(N) = (p-1)(q-1)$. The integer d is the private key, and is called the RSA *secret exponent*. The *public exponent* is the inverse e of d modulo $\phi(N)$.

7.1 Lattice attacks on RSA encryption

Small public exponent. When the public exponent e is very small, such as 3, one can apply Coppersmith's method (seen in the previous section) for univariate polynomials in various settings (see [16, 38, 42] for exact statements):

- An attacker can recover the plaintext of a given ciphertext, provided a large part of the plaintext is known.
- If a message is randomized before encryption, by simply padding random bits at a known place, an attacker can recover the message provided the amount of randomness is small.
- Håstad [66] attacks can be improved. An attacker can recover a message broadcasted (by RSA encryption and known affine transformation) to sufficiently many participants, each holding a different modulus N. This precisely happens if one sends a similar message with different known headers or time-stamps which are part of the encryption block.

None of the attacks recover the secret exponent d: they can only recover the plaintext. The attacks do not work if appropriate padding is used (see current standards and [95]), or if the public exponent is not too small. For instance, the popular choice $e = 65537$ is not threatened by these attacks.

Small private exponent. When $d \leq N^{0.25}$, an old result of Wiener [137] shows that one can easily recover the secret exponent d (and thus the factorization of N) from the continued fractions algorithm. Boneh and Durfee [19] recently improved the bound to $d \leq N^{0.292}$, by applying Coppersmith's technique to bivariate modular polynomials and improving the generic bound. Note that the attack is heuristic (see Section 6), but experiments showed that it works well in practice (no counterexample has ever been found). This bound holds when the RSA primes are balanced: Durfee and Nguyen [46] improved the bound when the primes are unbalanced, using an extension to trivariate modular polynomials. All those attacks on RSA with small private exponent also hold against the RSA signature scheme, since they only use the public key. A related result (using Coppersmith's technique for either bivariate integer or univariate modular polynomials) is an attack [20] to recover d when a large portion of the bits of d is known (see [16]).

7.2 Lattice attacks on RSA signature

The RSA cryptosystem is often used as a digital signature scheme. To prevent various attacks, one must apply a preprocessing scheme to the message, prior to signature. The recommended solution is to use hash functions and appropriate padding (see current standards and [95]). However, several alternative simple solutions not involving hashing have been proposed, and sometimes accepted as standards. Today, all such solutions have been broken (see [57]), some of them by lattice reduction techniques (see [104, 57]). Those lattice attacks are heuristic but work well in practice. They apply lattice reduction algorithms to find small solutions to modular linear systems, which leads to signature forgeries for certain proposed RSA signature schemes. Finding such small solutions is viewed as a closest vector problem for some norm, as seen in Section 3.4.

7.3 Security of RSA–OAEP

Although no efficient method is known to invert the RSA encryption function in general, it is widely accepted that the RSA encryption scheme should not be directly used as such, because it does not satisfy strong security notions (see for instance [22, 95] for a simple explanation): a preprocessing function should be applied to the message prior to encryption. The most famous preprocessing scheme for RSA is OAEP proposed by Bellare and Rogaway [11], which is standardized in PKCS. The RSA–OAEP scheme was only recently proved to be strongly secure (semantic security against adaptive chosen-ciphertext attacks), under the assumption that the RSA function is hard to invert and the random oracle model. This was first proved by Shoup [128] for the particular case of public exponent 3 using Coppersmith's theorem on univariate polynomial equations, and later extended to any exponent by Fujisaki *et al.* [53]. Interestingly, the last part of the proof of [53] relied on lattices (in dimension 2) to find a small solution to a linear modular equation (see Section 3.4). Note however that the result could also have been obtained with continued fractions.

Boneh [18] recently proposed a simpler version of OAEP for the RSA and Rabin encryption functions. The proof for Rabin is based on Coppersmith's lattice-based theorem on univariate polynomial equations, while the proof for RSA uses lattices again to find small solutions of linear modular equations. It is somewhat surprising that lattices are used both to attack RSA in certain settings, and to prove the security of industrial uses of RSA.

7.4 Factoring and lattice reduction

In the general case, the best attack against RSA encryption or signature is integer factorization. Note that to prove (or disprove) the equivalence between integer factorization and breaking RSA encryption remains an important open problem in cryptology (latest results [25] suggest that breaking RSA encryption may actually be easier). We already pointed out that in some special cases, lattice reduction leads to efficient factorization: when the factors are partially known [38], or when the number to factor has the form $p^r q$ with large r [21].

Schnorr [123] was the first to establish a link between integer factorization and lattice reduction, which was later extended by Adleman [2]. Schnorr [123] proposed a heuristic method to factor general numbers, using lattice reduction to approximate the closest vector problem in the infinity or the L_1 norm. Adleman [2] showed how to use the Euclidean norm instead, which is more suited to current lattice reduction algorithms. Those methods use the same underlying ideas as sieving algorithms (see [36]): to factor a number n, they try to find many congruences of smooth numbers to produce random square congruences of the form $x^2 \equiv y^2 \pmod{n}$, after a linear algebra step. Heuristic assumptions are needed to ensure the existence of appropriate congruences. The problem of finding such congruences is seen as a closest vector problem. Still, it should be noted that those methods are theoretical, since they are not adapted to currently known lattice reduction algorithms. To be useful, they would require very good lattice reduction for lattices of dimension over at least several thousands.

We close this review by mentioning that current versions of the Number Field Sieve (NFS) (see [87, 36]), the best algorithm known for factoring large integers, use lattice reduction. Indeed, LLL plays a crucial role in the last stage of NFS where one has to compute an algebraic square root of a huge algebraic number given as a product of hundreds of thousands of small ones. The best algorithm known to solve this problem is due to Montgomery (see [105, 107]). It has been used in all recent large factorizations, notably the record factorization [34] of a 512-bit RSA-number of 155 decimal digits proposed in the RSA challenges. There, LLL is applied many times in low dimension (less than 10) to find nice algebraic integers in integral ideals. But the overall running time of NFS is dominated by other stages, such as sieving and linear algebra.

8 Conclusions

The LLL algorithm and other lattice basis reduction algorithms have proved invaluable in cryptology. They have become the most popular tool in public-key cryptanalysis. In particular, they play a crucial rôle in several attacks against the RSA cryptosystem. The past few years have seen new, sometimes provable, lattice-based methods for solving problems which were *a priori* not linear, and this definitely opens new fields of applications. Interestingly, several provable lattice-based results introduced in cryptanalysis have also recently been used in the area of security proofs. Paradoxically, at the same time, a series of complexity results on lattice reduction has emerged, giving rise to another family of cryptographic schemes based on the hardness of lattice problems. The resulting cryptosystems have enjoyed different fates, but it is probably too early to tell whether or not secure and practical cryptography can be built using hardness of lattice problems. Indeed, several questions on lattices remain open. In particular, we still do not know whether or not it is easy to approximate the shortest vector problem up to some polynomial factor, or to find the shortest vector when the lattice gap is larger than some polynomial in the dimension. Besides, only very few lattice basis reduction algorithms are known, and their behaviour (both complexity and output quality) is still not well understood. And so far, there has not been any massive computer experiment in lattice reduction comparable to what has been done for integer factorization or the elliptic curve discrete logarithm problem. Twenty years of lattice reduction yielded surprising applications in cryptology. We hope the next twenty years will prove as exciting.

Acknowledgements. We thank Dan Boneh, Don Coppersmith, Glenn Durfee, Arjen and Hendrik Lenstra, László Lovász, Daniele Micciancio, Igor Shparlinski and Joe Silverman for helpful discussions and comments.

References

1. L. M. Adleman. On breaking generalized knapsack publick key cryptosystems. In *Proc. of 15th STOC*, pages 402–412. ACM, 1983.
2. L. M. Adleman. Factoring and lattice reduction. Unpublished manuscript, 1995.
3. M. Ajtai. Generating hard instances of lattice problems. In *Proc. of 28th STOC*, pages 99–108. ACM, 1996. Available at [47] as TR96-007.
4. M. Ajtai. The shortest vector problem in L_2 is NP-hard for randomized reductions. In *Proc. of 30th STOC*. ACM, 1998. Available at [47] as TR97-047.
5. M. Ajtai and C. Dwork. A public-key cryptosystem with worst-case/average-case equivalence. In *Proc. of 29th STOC*, pages 284–293. ACM, 1997. Available at [47] as TR96-065.
6. M. Ajtai, R. Kumar, and D. Sivakumar. A sieve algorithm for the shortest lattice vector problem. In *Proc. 33rd STOC*, pages 601–610. ACM, 2001.
7. S. Arora, L. Babai, J. Stern, and Z. Sweedyk. The hardness of approximate optima in lattices, codes, and systems of linear equations. *Journal of Computer and System Sciences*, 54(2):317–331, 1997.
8. L. Babai. On Lovász lattice reduction and the nearest lattice point problem. *Combinatorica*, 6:1–13, 1986.
9. W. Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 296:625–635, 1993.
10. M. Bellare, S. Goldwasser, and D. Micciancio. "Pseudo-random" number generation within cryptographic algorithms: The DSS case. In *Proc. of Crypto '97*, volume 1294 of *LNCS*. IACR, Springer-Verlag, 1997.
11. M. Bellare and P. Rogaway. Optimal asymmetric encryption. In *Proc. of Eurocrypt '94*, volume 950 of *LNCS*, pages 92–111. IACR, Springer-Verlag, 1995.

12. D. Bleichenbacher. On the security of the KMOV public key cryptosystem. In *Proc. of Crypto '97*, volume 1294 of *LNCS*, pages 235–248. IACR, Springer-Verlag, 1997.
13. D. Bleichenbacher and P. Q. Nguyen. Noisy polynomial interpolation and noisy Chinese remaindering. In *Proc. of Eurocrypt '00*, volume 1807 of *LNCS*. IACR, Springer-Verlag, 2000.
14. J. Blömer and J.-P. Seifert. On the complexity of computing short linearly independent vectors and short bases in a lattice. In *Proc. of 31st STOC*. ACM, 1999.
15. D. Boneh. The decision Diffie-Hellman problem. In *Algorithmic Number Theory – Proc. of ANTS-III*, volume 1423 of *LNCS*. Springer-Verlag, 1998.
16. D. Boneh. Twenty years of attacks on the RSA cryptosystem. *Notices of the AMS*, 46(2):203–213, 1999.
17. D. Boneh. Finding smooth integers in short intervals using CRT decoding. In *Proc. of 32nd STOC*. ACM, 2000.
18. D. Boneh. Simplified OAEP for the RSA and Rabin functions. In *Proc. of Crypto '01*, LNCS. IACR, Springer-Verlag, 2001.
19. D. Boneh and G. Durfee. Cryptanalysis of RSA with private key d less than $N^{0.292}$. In *Proc. of Eurocrypt '99*, volume 1592 of *LNCS*, pages 1–11. IACR, Springer-Verlag, 1999.
20. D. Boneh, G. Durfee, and Y. Frankel. An attack on RSA given a small fraction of the private key bits. In *Proc. of Asiacrypt '98*, volume 1514 of *LNCS*, pages 25–34. Springer-Verlag, 1998.
21. D. Boneh, G. Durfee, and N. A. Howgrave-Graham. Factoring $n = p^r q$ for large r. In *Proc. of Crypto '99*, volume 1666 of *LNCS*. IACR, Springer-Verlag, 1999.
22. D. Boneh, A. Joux, and P. Q. Nguyen. Why textbook ElGamal and RSA encryption are insecure. In *Proc. of Asiacrypt '00*, volume 1976 of *LNCS*. IACR, Springer-Verlag, 2000.
23. D. Boneh and I. E. Shparlinski. Hard core bits for the elliptic curve Diffie-Hellman secret. In *Proc. of Crypto '01*, LNCS. IACR, Springer-Verlag, 2001.
24. D. Boneh and R. Venkatesan. Hardness of computing the most significant bits of secret keys in Diffie-Hellman and related schemes. In *Proc. of Crypto '96*, LNCS. IACR, Springer-Verlag, 1996.
25. D. Boneh and R. Venkatesan. Breaking RSA may not be equivalent to factoring. In *Proc. of Eurocrypt '98*, volume 1233 of *LNCS*, pages 59–71. Springer-Verlag, 1998.
26. V. Boyko, M. Peinado, and R. Venkatesan. Speeding up discrete log and factoring based schemes via precomputations. In *Proc. of Eurocrypt '98*, volume 1403 of *LNCS*, pages 221–235. IACR, Springer-Verlag, 1998.
27. E. F. Brickell. Solving low density knapsacks. In *Proc. of Crypto '83*. Plenum Press, 1984.
28. E. F. Brickell. Breaking iterated knapsacks. In *Proc. of Crypto '84*, volume 196 of *LNCS*. Springer-Verlag, 1985.
29. E. F. Brickell and A. M. Odlyzko. Cryptanalysis: A survey of recent results. In G. J. Simmons, editor, *Contemporary Cryptology*, pages 501–540. IEEE Press, 1991.
30. J.-Y. Cai. Some recent progress on the complexity of lattice problems. In *Proc. of FCRC*, 1999. Available at [47] as TR99-006.
31. J.-Y. Cai. The complexity of some lattice problems. In *Proc. of ANTS-IV*, volume 1838 of *LNCS*. Springer-Verlag, 2000.

32. J.-Y. Cai and T. W. Cusick. A lattice-based public-key cryptosystem. *Information and Computation*, 151:17–31, 1999.
33. J.-Y. Cai and A. P. Nerurkar. An improved worst-case to average-case connection for lattice problems. In *Proc. of 38th FOCS*, pages 468–477. IEEE, 1997.
34. S. Cavallar, B. Dodson, A. K. Lenstra, W. Lioen, P. L. Montgomery, B. Murphy, H. te Riele, K. Aardal, J. Gilchrist, G. Guillerm, P. Leyland, J. Marchand, F. Morain, A. Muffett, C. Putnam, and P. Zimmermann. Factorization of 512-bit RSA key using the number field sieve. In *Proc. of Eurocrypt '00*, volume 1807 of *LNCS*. IACR, Springer-Verlag, 2000.
35. B. Chor and R.L. Rivest. A knapsack-type public key cryptosystem based on arithmetic in finite fields. *IEEE Trans. Inform. Theory*, 34, 1988.
36. H. Cohen. *A Course in Computational Algebraic Number Theory*. Springer-Verlag, 1995. Second edition.
37. J.H. Conway and N.J.A. Sloane. *Sphere Packings, Lattices and Groups*. Springer-Verlag, 1998. Third edition.
38. D. Coppersmith. Small solutions to polynomial equations, and low exponent RSA vulnerabilities. *J. of Cryptology*, 10(4):233–260, 1997. Revised version of two articles from Eurocrypt '96.
39. D. Coppersmith. Finding small solutions to small degree polynomials. In *Proc. of CALC '01*, LNCS. Springer-Verlag, 2001.
40. D. Coppersmith and A. Shamir. Lattice attacks on NTRU. In *Proc. of Eurocrypt '97*, LNCS. IACR, Springer-Verlag, 1997.
41. M.J. Coster, A. Joux, B.A. LaMacchia, A.M. Odlyzko, C.-P. Schnorr, and J. Stern. Improved low-density subset sum algorithms. *Comput. Complexity*, 2:111–128, 1992.
42. C. Coupé, P. Q. Nguyen, and J. Stern. The effectiveness of lattice attacks against low-exponent RSA. In *Proc. of PKC'98*, volume 1431 of *LNCS*. Springer-Verlag, 1999.
43. W. Diffie and M. E. Hellman. New directions in cryptography. *IEEE Trans. Inform. Theory*, IT-22:644–654, Nov 1976.
44. I. Dinur. Approximating SVP_∞ to within almost-polynomial factors is NP-hard. Available at [47] as TR99-016.
45. I. Dinur, G. Kindler, and S. Safra. Approximating CVP to within almost-polynomial factors is NP-hard. In *Proc. of 39th FOCS*, pages 99–109. IEEE, 1998. Available at [47] as TR98-048.
46. G. Durfee and P. Q. Nguyen. Cryptanalysis of the RSA schemes with short secret exponent from asiacrypt '99. In *Proc. of Asiacrypt '00*, volume 1976 of *LNCS*. IACR, Springer-Verlag, 2000.
47. ECCC. `http://www.eccc.uni-trier.de/eccc/`. The Electronic Colloquium on Computational Complexity.
48. E. El Mahassni, P. Q. Nguyen, and I. E. Shparlinski. The insecurity of Nyberg–Rueppel and other DSA-like signature schemes with partially known nonces. In *Proc. of CALC '01*, LNCS. Springer-Verlag, 2001.
49. P. van Emde Boas. Another NP-complete problem and the complexity of computing short vectors in a lattice. Technical report, Mathematische Instituut, University of Amsterdam, 1981. Report 81-04. Available at `http://turing.wins.uva.nl/~peter/`.
50. R. Fischlin and J.-P. Seifert. Tensor-based trapdoors for CVP and their application to public key cryptography. In *IMA Conference on Cryptography and Coding*, LNCS. Springer-Verlag, 1999.

51. A. M. Frieze. On the Lagarias-Odlyzko algorithm for the subset sum problem. *SIAM J. Comput*, 15(2):536–539, 1986.
52. A. M. Frieze, J. Håstad, R. Kannan, J. C. Lagarias, and A. Shamir. Reconstructing truncated integer variables satisfying linear congruences. *SIAM J. Comput.*, 17(2):262–280, 1988. Special issue on cryptography.
53. E. Fujisaki, T. Okamoto, D. Pointcheval, and J. Stern. RSA–OAEP is secure under the RSA assumption. In *Proc. of Crypto '01*, LNCS. IACR, Springer-Verlag, 2001.
54. M. L. Furst and R. Kannan. Succinct certificates for almost all subset sum problems. *SIAM J. Comput*, 18(3):550–558, 1989.
55. C.F. Gauss. *Disquisitiones Arithmeticæ*. Leipzig, 1801.
56. C. Gentry. Key recovery and message attacks on NTRU-composite. In *Proc. of Eurocrypt '01*, volume 2045 of *LNCS*. IACR, Springer-Verlag, 2001.
57. M. Girault and J.-F. Misarsky. Cryptanalysis of countermeasures proposed for repairing ISO 9796–1. In *Proc. of Eurocrypt '00*, volume 1807 of *LNCS*. IACR, Springer-Verlag, 2000.
58. O. Goldreich and S. Goldwasser. On the limits of non-approximability of lattice problems. In *Proc. of 30th STOC*. ACM, 1998. Available at [47] as TR97-031.
59. O. Goldreich, S. Goldwasser, and S. Halevi. Challenges for the GGH cryptosystem. Available at `http://theory.lcs.mit.edu/~shaih/challenge.html`.
60. O. Goldreich, S. Goldwasser, and S. Halevi. Eliminating decryption errors in the Ajtai-Dwork cryptosystem. In *Proc. of Crypto '97*, volume 1294 of *LNCS*, pages 105–111. IACR, Springer-Verlag, 1997. Available at [47] as TR97-018.
61. O. Goldreich, S. Goldwasser, and S. Halevi. Public-key cryptosystems from lattice reduction problems. In *Proc. of Crypto '97*, volume 1294 of *LNCS*, pages 112–131. IACR, Springer-Verlag, 1997. Available at [47] as TR96-056.
62. O. Goldreich, D. Micciancio, S. Safra, and J.-P. Seifert. Approximating shortest lattice vectors is not harder than approximating closest lattice vectors, 1999. Available at [47] as TR99-002.
63. M. I. González Vasco and I. E. Shparlinski. On the security of Diffie-Hellman bits. In K.-Y. Lam, I. E. Shparlinski, H. Wang, and C. Xing, editors, *Proc. Workshop on Cryptography and Comp. Number Theory (CCNT'99)*. Birkhauser, 2000.
64. M. Grötschel, L. Lovász, and A. Schrijver. *Geometric Algorithms and Combinatorial Optimization*. Springer-Verlag, 1993.
65. M. Gruber and C. G. Lekkerkerker. *Geometry of Numbers*. North-Holland, 1987.
66. J. Håstad. Solving simultaneous modular equations of low degree. *SIAM J. Comput.*, 17(2):336–341, April 1988. Preliminary version in Proc. of Crypto '85.
67. B. Helfrich. Algorithms to construct Minkowski reduced and Hermite reduced bases. *Theoretical Computer Science*, 41:125–139, 1985.
68. C. Hermite. Extraits de lettres de M. Hermite à M. Jacobi sur différents objets de la théorie des nombres, deuxième lettre. *J. Reine Angew. Math.*, 40:279–290, 1850. Also available in the first volume of Hermite's complete works, published by Gauthier-Villars.
69. J. Hoffstein, J. Pipher, and J.H. Silverman. NTRU: a ring based public key cryptosystem. In *Proc. of ANTS III*, volume 1423 of *LNCS*, pages 267–288. Springer-Verlag, 1998. Additional information at `http://www.ntru.com`.
70. N. A. Howgrave-Graham. Finding small roots of univariate modular equations revisited. In *Cryptography and Coding*, volume 1355 of *LNCS*, pages 131–142. Springer-Verlag, 1997.
71. N. A. Howgrave-Graham. *Computational Mathematics Inspired by RSA*. PhD thesis, University of Bath, 1998.

72. N. A. Howgrave-Graham. Approximate integer common divisors. In *Proc. of CALC '01*, LNCS. Springer-Verlag, 2001.
73. N. A. Howgrave-Graham and N. P. Smart. Lattice attacks on digital signature schemes. Technical report, HP Labs, 1999. HPL-1999-90. To appear in *Designs, Codes and Cryptography*.
74. E. Jaulmes and A. Joux. A chosen ciphertext attack on NTRU. In *Proc. of Crypto '00*, volume 1880 of *LNCS*. IACR, Springer-Verlag, 2000.
75. A. Joux and J. Stern. Lattice reduction: A toolbox for the cryptanalyst. *J. of Cryptology*, 11:161–185, 1998.
76. C. S. Jutla. On finding small solutions of modular multivariate polynomial equations. In *Proc. of Eurocrypt '98*, volume 1403 of *LNCS*, pages 158–170. IACR, Springer-Verlag, 1998.
77. R. Kannan. Improved algorithms for integer programming and related lattice problems. In *Proc. of 15th STOC*, pages 193–206. ACM, 1983.
78. R. Kannan. Algorithmic geometry of numbers. *Annual review of computer science*, 2:231–267, 1987.
79. R. Kannan. Minkowski's convex body theorem and integer programming. *Math. Oper. Res.*, 12(3):415–440, 1987.
80. P. Klein. Finding the closest lattice vector when it's unusually close. In *Proc. of SODA '00*. ACM–SIAM, 2000.
81. S. V. Konyagin and T. Seger. On polynomial congruences. *Mathematical Notes*, 55(6):596–600, 1994.
82. A. Korkine and G. Zolotareff. Sur les formes quadratiques positives ternaires. *Math. Ann.*, 5:581–583, 1872.
83. A. Korkine and G. Zolotareff. Sur les formes quadratiques. *Math. Ann.*, 6:336–389, 1873.
84. J. C. Lagarias. Point lattices. In R. Graham, M. Grötschel, and L. Lovász, editors, *Handbook of Combinatorics*, volume 1, chapter 19. Elsevier, 1995.
85. J. C. Lagarias and A. M. Odlyzko. Solving low-density subset sum problems. *Journal of the Association for Computing Machinery*, January 1985.
86. L. Lagrange. Recherches d'arithmétique. *Nouv. Mém. Acad.*, 1773.
87. A. K. Lenstra and H. W. Lenstra, Jr. *The Development of the Number Field Sieve*, volume 1554 of *Lecture Notes in Mathematics*. Springer-Verlag, 1993.
88. A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász. Factoring polynomials with rational coefficients. *Mathematische Ann.*, 261:513–534, 1982.
89. H. W. Lenstra, Jr. Integer programming with a fixed number of variables. Technical report, Mathematisch Instituut, Universiteit van Amsterdam, April 1981. Report 81-03.
90. H. W. Lenstra, Jr. Integer programming with a fixed number of variables. *Math. Oper. Res.*, 8(4):538–548, 1983.
91. L. Lovász. *An Algorithmic Theory of Numbers, Graphs and Convexity*, volume 50. SIAM Publications, 1986. CBMS-NSF Regional Conference Series in Applied Mathematics.
92. J. Martinet. *Les Réseaux Parfaits des Espaces Euclidiens.* Éditions Masson, 1996. English translation to appear at Springer-Verlag.
93. J. E. Mazo and A. M. Odlyzko. Lattice points in high-dimensional spheres. *Monatsh. Math.*, 110:47–61, 1990.
94. R.J. McEliece. A public-key cryptosystem based on algebraic number theory. Technical report, Jet Propulsion Laboratory, 1978. DSN Progress Report 42-44.
95. A. Menezes, P. Van Oorschot, and S. Vanstone. *Handbook of Applied Cryptography.* CRC Press, 1997.

96. R. Merkle and M. Hellman. Hiding information and signatures in trapdoor knapsacks. *IEEE Trans. Inform. Theory*, IT-24:525–530, September 1978.
97. D. Micciancio. *On the Hardness of the Shortest Vector Problem.* PhD thesis, Massachusetts Institute of Technology, 1998.
98. D. Micciancio. The shortest vector problem is NP-hard to approximate within some constant. In *Proc. of 39th FOCS*. IEEE, 1998. Available at [47] as TR98-016.
99. D. Micciancio. Lattice based cryptography: A global improvement. Technical report, Theory of Cryptography Library, 1999. Report 99-05.
100. D. Micciancio. The hardness of the closest vector problem with preprocessing. *IEEE Trans. Inform. Theory*, 47(3):1212–1215, 2001.
101. D. Micciancio. Improving lattice-based cryptosystems using the Hermite normal form. In *Proc. of CALC '01*, LNCS. Springer-Verlag, 2001.
102. J. Milnor and D. Husemoller. *Symmetric Bilinear Forms*. Springer-Verlag, 1973.
103. H. Minkowski. *Geometrie der Zahlen*. Teubner-Verlag, Leipzig, 1896.
104. J.-F. Misarsky. A multiplicative attack using LLL algorithm on RSA signatures with redundancy. In *Proc. of Crypto '97*, volume 1294 of *LNCS*, pages 221–234. IACR, Springer-Verlag, 1997.
105. P. L. Montgomery. Square roots of products of algebraic numbers. In Walter Gautschi, editor, *Mathematics of Computation 1943-1993: a Half-Century of Computational Mathematics*, Proc. of Symposia in Applied Mathematics, pages 567–571. American Mathematical Society, 1994.
106. National Institute of Standards and Technology (NIST). *FIPS Publication 186: Digital Signature Standard*, May 1994.
107. P. Q. Nguyen. A Montgomery-like square root for the number field sieve. In *Algorithmic Number Theory – Proc. of ANTS-III*, volume 1423 of *LNCS*. Springer-Verlag, 1998.
108. P. Q. Nguyen. Cryptanalysis of the Goldreich-Goldwasser-Halevi cryptosystem from Crypto '97. In *Proc. of Crypto '99*, volume 1666 of *LNCS*, pages 288–304. IACR, Springer-Verlag, 1999.
109. P. Q. Nguyen. *La Géométrie des Nombres en Cryptologie*. PhD thesis, Université Paris 7, November 1999. Available at `http://www.di.ens.fr/~pnguyen/`.
110. P. Q. Nguyen. The dark side of the hidden number problem: Lattice attacks on DSA. In K.-Y. Lam, I. E. Shparlinski, H. Wang, and C. Xing, editors, *Proc. Workshop on Cryptography and Comp. Number Theory (CCNT'99)*. Birkhauser, 2000.
111. P. Q. Nguyen and I. E. Shparlinski. The insecurity of the Digital Signature Algorithm with partially known nonces. *J. of Cryptology*, 2001. To appear.
112. P. Q. Nguyen and I. E. Shparlinski. The insecurity of the elliptic curve Digital Signature Algorithm with partially known nonces. *Preprint*, 2001.
113. P. Q. Nguyen and J. Stern. Merkle-Hellman revisited: a cryptanalysis of the Qu-Vanstone cryptosystem based on group factorizations. In *Proc. of Crypto '97*, volume 1294 of *LNCS*, pages 198–212. IACR, Springer-Verlag, 1997.
114. P. Q. Nguyen and J. Stern. Cryptanalysis of a fast public key cryptosystem presented at SAC '97. In *Selected Areas in Cryptography – Proc. of SAC '98*, volume 1556 of *LNCS*. Springer-Verlag, 1998.
115. P. Q. Nguyen and J. Stern. Cryptanalysis of the Ajtai-Dwork cryptosystem. In *Proc. of Crypto '98*, volume 1462 of *LNCS*, pages 223–242. IACR, Springer-Verlag, 1998.
116. P. Q. Nguyen and J. Stern. The Béguin-Quisquater server-aided RSA protocol from Crypto '95 is not secure. In *Proc. of Asiacrypt '98*, volume 1514 of *LNCS*, pages 372–379. Springer-Verlag, 1998.

117. P. Q. Nguyen and J. Stern. The hardness of the hidden subset sum problem and its cryptographic implications. In *Proc. of Crypto '99*, volume 1666 of *LNCS*, pages 31–46. IACR, Springer-Verlag, 1999.
118. P. Q. Nguyen and J. Stern. Lattice reduction in cryptology: An update. In *Proc. of ANTS-IV*, volume 1838 of *LNCS*. Springer-Verlag, 2000.
119. A. M. Odlyzko. The rise and fall of knapsack cryptosystems. In *Cryptology and Computational Number Theory*, volume 42 of *Proc. of Symposia in Applied Mathematics*, pages 75–88. A.M.S., 1990.
120. R. L. Rivest, A. Shamir, and L. M. Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Communications of the ACM*, 21(2):120–126, 1978.
121. C. P. Schnorr. A hierarchy of polynomial lattice basis reduction algorithms. *Theoretical Computer Science*, 53:201–224, 1987.
122. C. P. Schnorr. A more efficient algorithm for lattice basis reduction. *J. of algorithms*, 9(1):47–62, 1988.
123. C. P. Schnorr. Factoring integers and computing discrete logarithms via diophantine approximation. In *Proc. of Eurocrypt '91*, volume 547 of *LNCS*, pages 171–181. IACR, Springer-Verlag, 1991.
124. C. P. Schnorr and M. Euchner. Lattice basis reduction: improved practical algorithms and solving subset sum problems. *Math. Programming*, 66:181–199, 1994.
125. C. P. Schnorr and H. H. Hörner. Attacking the Chor-Rivest cryptosystem by improved lattice reduction. In *Proc. of Eurocrypt '95*, volume 921 of *LNCS*, pages 1–12. IACR, Springer-Verlag, 1995.
126. A. Shamir. A polynomial time algorithm for breaking the basic Merkle-Hellman cryptosystem. In *Proc. of 23rd FOCS*, pages 145–152. IEEE, 1982.
127. V. Shoup. Number Theory C++ Library (NTL) version 3.6. Available at `http://www.shoup.net/ntl/`.
128. V. Shoup. OAEP reconsidered. In *Proc. of Crypto '01*, LNCS. IACR, Springer-Verlag, 2001.
129. I. E. Shparlinski. On the generalized hidden number problem and bit security of XTR. In *Proc. of 14th Symp. on Appl. Algebra, Algebraic Algorithms, and Error-Correcting Codes*, LNCS. Springer-Verlag, 2001.
130. I. E. Shparlinski. Sparse polynomial approximation in finite fields. In *Proc. 33rd STOC*. ACM, 2001.
131. C. L. Siegel. *Lectures on the Geometry of Numbers*. Springer-Verlag, 1989.
132. B. Vallée. La réduction des réseaux. autour de l'algorithme de Lenstra, Lenstra, Lovász. *RAIRO Inform. Théor. Appl*, 23(3):345–376, 1989.
133. B. Vallée, M. Girault, and P. Toffin. How to guess ℓ-th roots modulo n by reducing lattice bases. In *Proc. of AAEEC-6*, volume 357 of *LNCS*, pages 427–442. Springer-Verlag, 1988.
134. S. A. Vanstone and R. J. Zuccherato. Short RSA keys and their generation. *J. of Cryptology*, 8(2):101–114, 1995.
135. S. Vaudenay. Cryptanalysis of the Chor-Rivest cryptosystem. In *Proc. of Crypto '98*, volume 1462 of *LNCS*. IACR, Springer-Verlag, 1998.
136. E. R. Verheul. Certificates of recoverability with scalable recovery agent security. In *Proc. of PKC '00*, LNCS. Springer-Verlag, 2000.
137. M. Wiener. Cryptanalysis of short RSA secret exponents. *IEEE Trans. Inform. Theory*, 36(3):553–558, 1990.