

Improving Lattice Based Cryptosystems Using the Hermite Normal Form

Daniele Micciancio^{1*}

Department of Computer Science and Engineering
University of California, San Diego
9500 Gilman Drive, La Jolla, CA 92093, USA
daniele@cs.ucsd.edu

Abstract. We describe a simple technique that can be used to substantially reduce the key and ciphertext size of various lattice based cryptosystems and trapdoor functions of the kind proposed by Goldreich, Goldwasser and Halevi (GGH). The improvement is significant both from the theoretical and practical point of view, reducing the size of both key and ciphertext by a factor n equal to the dimension of the lattice (i.e., several hundreds for typical values of the security parameter.) The efficiency improvement is obtained without decreasing the security of the functions: we formally prove that the new functions are at least as secure as the original ones, and possibly even better as the adversary gets less information in a strong information theoretical sense. The increased efficiency of the new cryptosystems allows the use of bigger values for the security parameter, making the functions secure against the best cryptanalytic attacks, while keeping the size of the key even below the smallest key size for which lattice cryptosystems were ever conjectured to be hard to break.

Keywords: Lattices, trapdoor functions, public-key encryption

1 Introduction

Recent results on the complexity of lattices [1] have drawn considerable attention to lattice problems as potential candidates to design cryptographic primitives, and encryption schemes in particular. The two most notable proposals are the Ajtai-Dwork cryptosystem (AD, [2]) and the Goldreich-Goldwasser-Halevi cryptosystem (GGH, [14]). Other lattice based cryptosystems designed along the lines of [2, 14], were subsequently proposed by Fischlin and Seifert [10] and Cai and Cusick [6]. While the AD cryptosystem is mainly of theoretical interest, the GGH cryptosystem was suggested as a practical alternative to number theory based schemes currently in use (e.g., the RSA cryptosystem [28]). Two other related cryptosystems are McEliece's [21] and NTRU [17]. Neither of them is a lattice based cryptosystem in the strict meaning of the term, as they use ideas

* Research supported in part by NSF Career Award CCR-0093029

from other areas of mathematics (polynomial ring and finite field arithmetics respectively). However, the security of NTRU is related to the hardness of certain lattice problems (see Sect. 7 for further discussion), and it definitely deserves a mention as the most practical of the above proposals, yielding keys of size $O(n \log n)$ instead of the $\Omega(n^2)$ key size of McEliece, GGH and related schemes. (In fact, we'll see that GGH keys can be as big as $O(n^3 \log n)$.) McEliece's scheme is based on the hardness of coding theoretic problems, rather than lattices, but it bears much resemblance to the GGH cryptosystem, and we will further discuss this scheme in Sect. 7, after introducing our new lattice based trapdoor function. As it is the case for RSA, no proof of security for the GGH cryptosystem (or any of the afore mentioned schemes with the only exception of Ajtai and Dwork's theoretical proposal) is currently known, and its conjectured security is based on the empirical evidence that certain lattice problems are hard. In the attempt of stimulating further cryptanalytic efforts against their system and determining appropriate key size, the authors of GGH published 5 numerical challenges [13] corresponding to increasing values of the security parameter $n = 200, 250, 300, 350, 400$, resulting in public key sizes ranging from 330 KB to over 2 MB. Despite the big key size even for the smallest dimension (330 KB), this cryptosystem was still competitive with number theoretic cryptosystems because the encryption time is essentially linear in the size of the key, while modular exponentiation typically requires $O(n^3)$ operations.

At the time the GGH cryptosystem was presented at Crypto'97, challenges in dimension $n = 150$ were known to be breakable [31] using lattice reduction techniques. Still these techniques didn't seem to apply to lattices in higher dimension $n > 200$. At Crypto'99 Nguyen [24] showed how to exploit a weakness specific to the way GGH challenges were chosen and break the first four of the five GGH challenges. As the only unbroken challenge had key size over 2 MB, the practical value of the GGH cryptosystem (and variants) seemed quite questionable and [24] concluded that "unless major improvements are found, lattice-based cryptography cannot provide a serious alternative to existing public-key encryption algorithms". In this paper we present one such improvement.

Although quite simple, the techniques described in this paper can be used to significantly reduce the key size of GGH-like cryptosystems (e.g., those presented in [14, 10]). In particular, we show how to build lattice based trapdoor functions with public keys of size below 330 KB, and still secure against the best known lattice attacks. The improvement gets even better as the dimension of the lattice grows: our techniques can be used to asymptotically reduce the key size from $O(n^3 \log n)$ to $O(n^2 \log n)$. The improvement in the encryption time is analogous, being roughly proportional to the size of the public key, and the size of the ciphertext is also significantly reduced going from $O(n^2 \log n)$ to $O(n \log n)$. Surprisingly, we can achieve these efficiency improvements without decreasing the security of the scheme: any attack to the new scheme can be provably transformed into an (at least) equally effective attack against the original GGH (and variant) schemes. The increased efficiency allows one to use greater values of the security parameter than considered in [14] and [24], while

maintaining the scheme reasonably practical. In particular, our techniques give encryption schemes that resist the best known lattice attacks and still have public keys even smaller than the weakest of the GGH challenges, bringing the feasibility of lattice based cryptography back into discussion.

The rest of the paper is organized as follows. In Sect. 2 we recall some basic definitions and properties of lattices. In Sect. 3 we describe the original GGH scheme and discuss its weaknesses. In Sect. 4 we define a new cryptographic function that significantly improves the GGH scheme both in terms of security and efficiency. The new scheme is analyzed in Sect. 5. To be precise, the GGH scheme, as well as the new one we are proposing in this paper, should be considered as *trapdoor functions* instead of ready-to-use *encryption schemes*. The reason is explained in Sect. 6, where we also discuss how to transform these functions into encryption schemes. In Sect. 6 we also describe some cryptanalytic experiments that we performed to validate our theoretical results. The McEliece and NTRU cryptosystems are discussed in Sect. 7. Section 8 concludes with final remarks and open problems.

2 Preliminaries

Let $B = \{\mathbf{b}_1, \dots, \mathbf{b}_n\}$ be a set of n linearly independent vectors in \mathbb{R}^m . The *lattice* generated by B is the set $\mathcal{L}(B) = \{\sum_i x_i \mathbf{b}_i \mid x_i \in \mathbb{Z}\}$ of all *integer* linear combinations of the vectors in B . The set B is called *basis* and it is usually identified with the matrix $\mathbf{B} = [\mathbf{b}_1, \dots, \mathbf{b}_n] \in \mathbb{R}^{m \times n}$ having the vectors \mathbf{b}_i as columns. In matrix notation $\mathcal{L}(B) = \{\mathbf{B}\mathbf{x} \mid \mathbf{x} \in \mathbb{Z}^n\}$. The lattice $\mathcal{L}(B)$ is *full rank* if $n = m$, i.e. if \mathbf{B} spans the entire vector space \mathbb{R}^m over the reals. For simplicity, in the rest of this paper we will consider only full rank lattices. Moreover, for computational purposes, we will assume the basis vectors $\mathbf{b}_i \in \mathbb{Z}^n$ have integer entries. Then, a basis is just a non-singular integer matrix $\mathbf{B} \in \mathbb{Z}^{n \times n}$.

The basis of a lattice is not unique. A particularly convenient basis for some applications is the *Hermite Normal Form* (HNF). A basis \mathbf{B} is in HNF if it is upper triangular, all elements on the diagonal are strictly positive, and any other element $b_{i,j}$ satisfies $0 \leq b_{i,j} < b_{i,i}$. It is easy to see that every integer lattice $\mathcal{L}(B)$ has a unique basis in Hermite Normal Form, denoted $\text{HNF}(\mathbf{B})$. Moreover, given any basis \mathbf{B} for the lattice, $\text{HNF}(\mathbf{B})$ can be efficiently computed (see [23] for a recent algorithm and a survey of previous results). Notice that $\text{HNF}(\mathbf{B})$ does not depend on the particular basis \mathbf{B} with started from, and it is uniquely defined by the lattice $\mathcal{L}(B)$ generated by B .

For every basis \mathbf{B} , the *orthogonalized basis* $\mathbf{B}^* = [\mathbf{b}_1^*, \dots, \mathbf{b}_n^*]$ is defined by the usual Gram-Schmidt orthogonalization process:

$$\begin{aligned} \mathbf{b}_1^* &= \mathbf{b}_1 \\ \mathbf{b}_i^* &= \mathbf{b}_i - \sum_{j < i} \frac{\langle \mathbf{b}_i, \mathbf{b}_j^* \rangle}{\langle \mathbf{b}_j^*, \mathbf{b}_j^* \rangle} \mathbf{b}_j^* . \end{aligned}$$

Notice that \mathbf{B}^* is a basis for the vector space \mathbb{R}^n , but is not in general a lattice basis for $\mathcal{L}(\mathbf{B})$. If \mathbf{B} is in HNF, then \mathbf{B}^* is simply the diagonal matrix with $b_{1,1}, \dots, b_{n,n}$ on the diagonal. The determinant of a lattice $\mathcal{L}(\mathbf{B})$ is the absolute value of the determinant of the matrix \mathbf{B} . The determinant is a lattice invariant, i.e. it does not depend on the particular choice of the basis \mathbf{B} . If \mathbf{B} is in HNF, then the determinant $\det(\mathbf{B})$ is just the product of the elements on the diagonal $\prod_{i=1}^n b_{i,i}$. An important property of integer lattices is that if two vectors \mathbf{v} and \mathbf{w} are congruent modulo $\det(\mathbf{B})$ (i.e. $\det(\mathbf{B})$ divides $v_i - w_i$ for all i), then $\mathbf{v} \in \mathcal{L}(\mathbf{B})$ if and only if $\mathbf{w} \in \mathcal{L}(\mathbf{B})$. In other words, the lattice repeats identically if translated by multiples of $\det(\mathbf{B})$ along the direction of any of the main axes.

The distance between two vectors \mathbf{v} and \mathbf{w} is defined by

$$\text{dist}(\mathbf{v}, \mathbf{w}) = \|\mathbf{v} - \mathbf{w}\| = \sqrt{\sum_i (v_i - w_i)^2} .$$

The distance function is extended to sets of vectors as usual

$$\text{dist}(S_1, S_2) = \inf\{\|\mathbf{v} - \mathbf{w}\| : \mathbf{v} \in S_1, \mathbf{w} \in S_2\} .$$

In particular the distance of a vector from a lattice is given by $\text{dist}(\mathbf{v}, \mathcal{L}(\mathbf{B})) = \min\{\|\mathbf{v} - \mathbf{w}\| : \mathbf{w} \in \mathcal{L}(\mathbf{B})\}$. In the *closest vector problem* (CVP), one is given a basis \mathbf{B} and a target vector \mathbf{v} (usually not in the lattice) and must find the lattice vector in $\mathcal{L}(\mathbf{B})$ closest to \mathbf{v} . CVP was proved NP-hard in [36], and it remains hard even if the lattice basis can be arbitrarily preprocessed [22], or one allows for approximate solutions with approximation factor $2^{\lg^{1-\epsilon} n}$ [3, 9]. To date, the best polynomial time algorithm to approximate CVP achieves only a worst case approximation factor which is almost exponential in the dimension of the lattice [19, 4, 29].

A closely related problem is the shortest vector problem (SVP): given a lattice $L = \mathcal{L}(\mathbf{B})$, find the length $\lambda(L)$ of the shortest non-zero vector in $\mathcal{L}(\mathbf{B})$. By linearity, $\lambda(\mathbf{B})$ equals the minimum distance between any two lattice points $\min\{\|\mathbf{v} - \mathbf{w}\| : \mathbf{v}, \mathbf{w} \in \mathcal{L}(\mathbf{B}), \mathbf{v} \neq \mathbf{w}\}$. It is easy to see that for any vector \mathbf{v} (not necessarily in the lattice) there exists at most one lattice point within distance $\lambda/2$ from \mathbf{v} . One can easily show that the length of the shortest vector in a lattice $L = \mathcal{L}(\mathbf{B})$ satisfies $\lambda(L) \geq \min_i \|\mathbf{b}_i^*\|$. Moreover, given a vector \mathbf{v} within distance $\rho = \frac{1}{2} \min_i \|\mathbf{b}_i^*\|$ from the lattice, the (unique) lattice vector within distance ρ from \mathbf{v} can be efficiently computed from \mathbf{B} and \mathbf{v} using Babai's *nearest plane* algorithm [4]. (See also [18].)

3 The GGH Encryption Scheme

The GGH encryption scheme [14] works essentially as follows. The private and public keys of the scheme are two bases \mathbf{B}, \mathbf{R} of the same lattice $L = \mathcal{L}(\mathbf{B}) = \mathcal{L}(\mathbf{R})$. The private key \mathbf{R} is an exceptionally good basis. In particular, \mathbf{R} is chosen in such a way that the quantity $\rho = \frac{1}{2} \min \|\mathbf{r}_i^*\|$ is relatively large, so that all

errors of length less than ρ can be efficiently corrected using \mathbf{R} . However, given the public basis this same task should be computationally hard. In particular, $\frac{1}{2} \min \|b_i^*\|$ is much smaller than ρ .

The two bases are used to define a trapdoor function that takes as input an integer vector \mathbf{x} and an error vector \mathbf{r} of length at most ρ , and returns the vector $\mathbf{c} = \mathbf{B}\mathbf{x} + \mathbf{r}$, i.e. the lattice vector with public coefficients \mathbf{x} perturbed by a small additive error \mathbf{r} . Notice that $\mathbf{B}\mathbf{x}$ can be recovered from \mathbf{c} using the private basis \mathbf{R} . Once $\mathbf{B}\mathbf{x}$ is recovered, one can easily compute \mathbf{x} and \mathbf{r} using simple linear algebra, therefore inverting the function.

A message m is “encrypted” by first encoding m in the input (\mathbf{x}, \mathbf{r}) , and then applying the trapdoor function to (\mathbf{x}, \mathbf{r}) . Two encoding methods are considered in [14]:

1. In the first method, the message m is encoded in the error vector \mathbf{r} , and \mathbf{x} is chosen at random
2. In the second method, the message m is encoded in the coefficients \mathbf{x} , and \mathbf{r} is chosen at random.

For concreteness, in the rest of the paper we will assume the first encoding method, but most of the techniques we describe can be easily adapted to the second method as well.

In order to fully specify the trapdoor function the following questions must be answered:

1. How is the private basis \mathbf{R} chosen?
2. How is the public basis \mathbf{B} obtained from \mathbf{R} ?
3. How is the random vector \mathbf{x} chosen?
4. How is the error vector \mathbf{r} chosen?

The authors of [14] suggest¹ to take $\mathbf{R} = \sqrt{n} \cdot \mathbf{I} + \mathbf{Q}$, where \mathbf{I} is the identity matrix, and \mathbf{Q} is a random perturbation matrix with entries in $\{-4, \dots, +4\}$. Then obtain the public basis \mathbf{B} applying a sufficiently long sequence of random elementary column operations to \mathbf{R} . Then the message m is “encrypted” by encoding it in an error vector \mathbf{r} with entries $r_i = \pm 3$, and adding it to a lattice vector $\mathbf{B}\mathbf{x}$ chosen at random from a sufficiently large region of space.

Notice that in order to avoid attacks based on exhaustive search, the sequence of operations applied to \mathbf{R} to obtain the public basis, and the region of space from which the lattice vector $\mathbf{B}\mathbf{x}$ is chosen must be sufficiently large. Since the lattice repeats identically if translated by $\det(L)$ along any of the main axes, we can always assume that the entries of \mathbf{B} and \mathbf{x} are reduced modulo $\det(L)$ without decreasing the security of the scheme. We can use this observation to estimate the proper size of the public key \mathbf{B} and the ciphertext $\mathbf{c} = \mathbf{B}\mathbf{x} + \mathbf{r}$ as $O(n^2 \cdot \lg(\det(L)))$ and $O(n \cdot \lg(\det(L)))$. The determinant $\det(L)$ can also be estimated to be $2^{O(n \lg n)}$ applying Hadamard inequality to the private basis,

¹ These are the choices used to generate the challenges in [13] and considered in cryptanalytic attacks [24]. In fact a broader range of possible choices is considered in the paper [14].

resulting in $O(n^3 \lg n)$ and $O(n^2 \lg n)$ estimates for the key and ciphertext size. At this point one can point out how the particular way questions 1-4 were answered in [13] introduced some serious weakness in the system. In fact:

1. The only property of \mathbf{R} required for the decryption algorithm to work is that the orthogonalized vectors \mathbf{r}_i^* are sufficiently long. Choosing them to be close to the main axes $\sqrt{n} \cdot \mathbf{e}_i$ seems a quite peculiar restriction that might make the scheme much weaker. Other work in the design of lattice based encryption schemes [35, 10] suggests that rotations might play a fundamental role in making these schemes secure.
2. The size of the GGH challenges published at [13] are about an order of magnitude smaller than what we estimated to be the proper size. This suggests that the public key of the GGH challenges had not been “randomized” enough, making them particularly easy to break.
3. The same arguments apply to the choice of the lattice vector $\mathbf{B}\mathbf{x}$.
4. Choosing error vectors whose entries have all the same absolute value also introduces serious weaknesses, as shown in [24].

Now it shouldn't be a surprise that the numerical challenges in [13] have been broken, but it should also be clear that known attacks say very little about the general methodology used by the GGH cryptosystem. In particular, Nguyen attack [24] relies, in an essential way, on the property that all the entries in the error vector have the same absolute value and is based on the following observation: if all entries of \mathbf{r} have absolute value σ , then $\mathbf{c} + \mathbf{s} \equiv \mathbf{B}\mathbf{x} \pmod{2\sigma}$, where \mathbf{s} is the vector $[\sigma, \dots, \sigma]^T$. This allows to find \mathbf{x} modulo 2σ and reduce the problem of finding a lattice vector within distance $\|\mathbf{r}\|$ from \mathbf{c} to the problem of finding a lattice vector within smaller distance $\|\mathbf{r}\|/(2\sigma)$ from $(\mathbf{c} - \mathbf{B}(\mathbf{x} \bmod 2\sigma))/(2\sigma)$. (In the GGH challenges $\sigma = 3$, reducing the length of the error by a factor 6.) As already noted in [24], removing the restriction that all entries in the error vector have the same modulus, immediately fixes the problem, and the main question is whether the security parameter in the GGH scheme can be made sufficiently large to achieve security, and maintaining the scheme practical at the same time.

4 An Optimal GGH-like Trapdoor Function

We first describe a general scheme to define GGH-like trapdoor functions, of which GGH is a special case. Then we show how to instantiate the scheme in an essentially optimal way, defining a specific trapdoor function that is both more secure and efficient than any other function from the scheme. In particular, the new trapdoor function is considerably more efficient than the original GGH.

Fix a probability distribution on the set of private bases \mathbf{R} and let ρ be a correction radius such that using \mathbf{R} one can correct any error of length less than ρ . (e.g. $\rho = \frac{1}{2} \min_i \|\mathbf{r}_i^*\|$.) We define a family of functions $f_{\beta, \gamma}$ parametrized by two algorithms β and γ as follows.

1. Let β be a (possibly randomized) function that on input a matrix \mathbf{R} outputs another basis \mathbf{B} for the same lattice that will be used as a public key.

2. γ is a (possibly randomized) function that on input the public basis \mathbf{B} and an error vector \mathbf{r} , outputs the coefficients \mathbf{x} of a lattice point $\mathbf{B}\mathbf{x}$ to be added to the error vector.
3. Let $f_{\beta,\gamma}$ be the (possibly randomized) function with domain the set of vectors shorter than ρ defined as follows:

$$f_{\beta,\gamma}(\mathbf{r}) = \mathbf{B}\mathbf{x} + \mathbf{r}$$

where $\mathbf{B} = \beta(\mathbf{R})$ and $\mathbf{x} = \gamma(\mathbf{B}, \mathbf{r})$.

Notice that if the input \mathbf{r} has length less than ρ , then the basis \mathbf{R} can be used to find the lattice point $\mathbf{B}\mathbf{x}$ closest to $f_{\beta,\gamma}(\mathbf{r})$ and recover \mathbf{r} .

Therefore, for any fixed β and γ , the probability distribution on \mathbf{R} defines a family of trapdoor functions $f_{\beta,\gamma}$ with public key (\mathbf{B}, ρ) and trapdoor information \mathbf{R} . The GGH trapdoor function can be defined as a special case of this scheme when $\beta(\mathbf{R})$ applies a sequence of elementary random operations to \mathbf{R} , and $\gamma(\mathbf{B}, \mathbf{r})$ outputs an integer vector \mathbf{x} chosen at random from a sufficiently large region of space.

We are now ready to define different β and γ , that greatly increase both the security and the efficiency of the scheme. The idea is to replace the random choice of the public basis \mathbf{B} and vector \mathbf{x} with deterministic choices that can be formally proved to be optimal from the security point of view. In the following subsections we first give some definitions that will be useful in the sequel, then show how to compute the public basis, and finally define the new trapdoor function.

4.1 Reducing Vectors Modulo a Basis

Every lattice $L = \mathcal{L}(\mathbf{B})$ induces an equivalence relation over \mathbb{Z}^n defined as follows: $\mathbf{v} \equiv_L \mathbf{w}$ if and only if $\mathbf{v} - \mathbf{w} \in L$. It is easy to see that for every point $\mathbf{v} \in \mathbb{Z}^n$, there exists a unique point \mathbf{w} in the *orthogonalized parallelepiped* $\mathcal{P}(\mathbf{B}^*) = \{\sum_i x_i \mathbf{b}_i^* \mid 0 \leq x_i < 1\}$ such that $\mathbf{w} \equiv_L \mathbf{v}$. Vector \mathbf{w} can be easily computed from \mathbf{v} and \mathbf{B} as follows. For all $i = n, \dots, 1$ (in decreasing order), let $\alpha_i = \frac{\langle \mathbf{v}, \mathbf{b}_i^* \rangle}{\langle \mathbf{b}_i^*, \mathbf{b}_i^* \rangle}$ the component of \mathbf{v} along \mathbf{b}_i^* and subtract $[\alpha] \mathbf{b}_i$ from \mathbf{v} . Let \mathbf{w} the final result. Since \mathbf{w} and \mathbf{v} differ only by integer multiples of basis vectors, we have $\mathbf{w} \equiv_L \mathbf{v}$. Moreover, it is easy to check that $\frac{\langle \mathbf{w}, \mathbf{b}_i^* \rangle}{\langle \mathbf{b}_i^*, \mathbf{b}_i^* \rangle} \in [0, 1)$ for all i , and therefore $\mathbf{w} \in \mathcal{P}(\mathbf{B}^*)$. The unique element of $\mathcal{P}(\mathbf{B}^*)$ congruent to \mathbf{v} modulo L is denoted $\mathbf{v} \bmod \mathbf{B}$. Notice that although the equivalence relation $\mathbf{v} \equiv_L \mathbf{w}$ does not depend on the particular choice of the basis \mathbf{B} for the lattice $L = \mathcal{L}(\mathbf{B})$, the definition of the reduced vector $(\mathbf{v} \bmod \mathbf{B})$ is basis dependent.

Pictorially, we can think the vector space \mathbb{R}^n as partitioned into parallelepipeds $\{\mathcal{P}(\mathbf{B}^*) + \mathbf{z} \mid \mathbf{z} \in \mathcal{L}(\mathbf{B})\}$. Then, the reduced vector $\mathbf{v} \bmod \mathbf{B}$ is the relative position of \mathbf{v} in the parallelepiped $\mathcal{P}(\mathbf{B}^*) + \mathbf{z}$ it belongs to. Notice that if \mathbf{v} is an integer vector, then also $\mathbf{v} \bmod \mathbf{B}$ is an integer vector. Therefore the function $\mathbf{x} \mapsto (\mathbf{x} \bmod \mathbf{B})$ defines a function from \mathbb{Z}^n to $\mathcal{P}(\mathbf{B}^*) \cap \mathbb{Z}^n$.

If \mathbf{B} is in Hermite Normal Form, then $\mathbf{w} = \mathbf{v} \bmod \mathbf{B}$ is an integer vector satisfying $0 \leq w_i < b_{i,i}$. In particular \mathbf{w} can be represented using roughly $\sum_{i=1}^n \lg b_{i,i} = \lg(\det(L))$ bits. This representation is essentially optimal because \equiv_L induces exactly $\det(L)$ congruence classes on \mathbb{Z}^n .

4.2 Choosing the Public Basis

Let's consider the choice of the public basis. The private basis \mathbf{R} we start from is an exceptionally good basis that allows to solve the closest vector problem in the lattice $\mathcal{L}(\mathbf{R})$ and consequently decrypt messages. We would like to transform it into another basis for the same lattice L that gives the least possible amount of information about \mathbf{R} . Instead of computing \mathbf{B} by applying a complex random transformation to \mathbf{R} , we set the public key of the new cryptosystem to be the Hermite Normal Form $\mathbf{B} = \text{HNF}(\mathbf{R})$ of \mathbf{R} . Since the $\text{HNF}(\mathbf{R})$ depends solely on the lattice $\mathcal{L}(\mathbf{R})$ generated by \mathbf{R} (and not on the particular basis \mathbf{R} we used to compute it), the new public key gives *no* information about the private key \mathbf{R} , other than the lattice L it generates. More formally, one can prove that any information about \mathbf{R} that can be efficiently computed from $\mathbf{B} = \text{HNF}(\mathbf{R})$ can also be efficiently computed starting from any other (possibly random) basis \mathbf{B}' . This is because if \mathbf{B}' and \mathbf{R} generate the same lattice $L = \mathcal{L}(\mathbf{B}') = \mathcal{L}(\mathbf{R})$ then $\mathbf{B} = \text{HNF}(\mathbf{B}') = \text{HNF}(\mathbf{R})$ and \mathbf{B} can be efficiently computed from \mathbf{B}' .

4.3 Adding a “Random” Lattice Point

Let's now look at how to simulate the addition of a “random” lattice vector $\mathbf{B}\mathbf{x}$ to the error vector \mathbf{r} . Ideally, we would like $\mathbf{B}\mathbf{x}$ to be a uniformly chosen vector from L . Unfortunately this is neither a computationally feasible nor a mathematically well defined operation. However, we notice that we can achieve exactly the same result by mapping the error vector \mathbf{r} to its equivalence class $[\mathbf{r}]_L$ modulo \equiv_L . An efficient way to do this is to use the reduced vector $\mathbf{r} \bmod \mathbf{B}$ as a representative for this class. So, instead of adding to \mathbf{r} a random lattice point $\mathbf{B}\mathbf{x}$, we reduce \mathbf{r} modulo the public basis \mathbf{B} to obtain the ciphertext $\mathbf{c} \in \mathcal{P}(\mathbf{B}^*)$.

The new trapdoor function is then defined as follows:

$$f(\mathbf{r}) = \mathbf{r} \bmod \mathbf{B}$$

where $\mathbf{B} = \text{HNF}(\mathbf{R})$ is the Hermite Normal Form of the trapdoor information \mathbf{R} . The triangular form of \mathbf{B} also makes the trapdoor function (i.e., the reduction modulo \mathbf{B}) extremely simple. Given \mathbf{r} , the reduced vector $\mathbf{r} \bmod \mathbf{B}$ can be easily determined as follow. Compute the integer vector \mathbf{x} one coordinate at a time (starting from x_n) using the formula

$$x_i = \left\lfloor \frac{r_i - \sum_{j>i} b_{i,j}x_j}{b_{i,i}} \right\rfloor.$$

The output of the trapdoor function is $\mathbf{c} = \mathbf{r} - \mathbf{B}\mathbf{x} \equiv \mathbf{r} \bmod \mathbf{B}$. The reader can easily check that for every i , $0 \leq c_i < b_{i,i}$, i.e., the result is the unique point in the parallelepiped $\mathcal{P}(\mathbf{B}^*) = \{\mathbf{w} \mid 0 \leq w_i < b_{i,i}\}$ which is congruent to \mathbf{r} modulo $\mathcal{L}(\mathbf{B})$.

4.4 The New Trapdoor Function

We now put all pieces together and define the new trapdoor function. Let \mathbf{R} be a private basis chosen in such a way that $\rho = \frac{1}{2} \min_i \|\mathbf{r}_i^*\|$ is relatively big (see Fig. 1). The public basis \mathbf{B} is the Hermite Normal Form of \mathbf{R} (see Fig. 2). One can see that the public basis \mathbf{B} and the corresponding orthogonalized parallelepiped $\mathcal{P}(\mathbf{B}^*)$ are very skewed. The public basis \mathbf{B} defines a function with domain the set of vectors of length at most ρ (the shaded circle in Fig. 2). The result of applying the function to vector \mathbf{r} is the point $(\mathbf{r} \bmod \mathbf{B})$ in the parallelepiped $\mathcal{P}(\mathbf{B}^*)$ congruent to \mathbf{r} modulo the lattice. Notice that even if we always start from a vector \mathbf{r} close to the origin, the result of performing the reduction operation is a point of $\mathcal{P}(\mathbf{B}^*)$ possibly closest to some other lattice point (see black regions in Fig. 2 and the corresponding closest lattice points). Notice that recovering the input vector \mathbf{r} from $f(\mathbf{r})$ involves finding the lattice point closest to $f(\mathbf{r})$, which is conjectured to be infeasible using only the public basis \mathbf{B} . However the lattice vector closest to $\mathbf{r} \bmod \mathbf{B}$ can be computed using the private basis \mathbf{R} as discussed in Sect. 2 because $\text{dist}(f(\mathbf{r}), L) = \text{dist}(\mathbf{r}, L) \leq \rho$. Fig. 3 shows the orthogonalized parallelepipeds $\mathcal{P}(\mathbf{R}^*)$ centered at every lattice point. Notice that the lattice point closest to $f(\mathbf{r})$ (i.e., any point of the black regions in the picture) is just the center of the parallelepiped $\mathcal{P}(\mathbf{R}^*)$ containing $f(\mathbf{r})$, which can be found using the private basis \mathbf{R} .




Fig. 1. A good lattice basis and the corresponding correction radius

5 Analysis

In this section we discuss the security and performance of the new scheme.




Fig. 2. HNF basis and corresponding orthogonalized parallelepiped

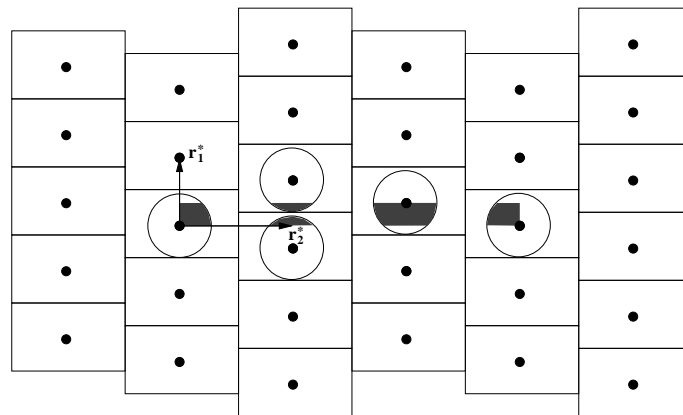


Fig. 3. Correcting small errors using the private basis

5.1 Security

We want to prove that the new trapdoor function $f(\mathbf{r})$ is at least as secure as the original GGH function. We actually prove that $f(\mathbf{r})$ is at least as secure as *any* GGH-like function $f_{\beta,\gamma}$ as defined in the previous section.

Theorem 1. *For any (efficiently computable) functions β, γ , and for any (efficient) algorithm that on input $f(\mathbf{r})$ finds some partial information about \mathbf{r} with non-trivial probability, there exists an efficient algorithm that on input $f_{\beta,\gamma}(\mathbf{r})$ finds the same partial information with the same success probability.*

Proof. The proof is a simple reduction argument. Assume A is an algorithm that breaks f . We show how to attack $f_{\beta,\gamma}$ using A as a subroutine. We are given public basis $\mathbf{B} = \beta(\mathbf{R})$ and a ciphertext $\mathbf{c} = \mathbf{B} \cdot \boldsymbol{\gamma} + \mathbf{r}$. The task is to find (some partial information about) \mathbf{r} . We first compute $\mathbf{B}' = \text{HNF}(\mathbf{B})$ and $\mathbf{c}' = \mathbf{c} \bmod \mathbf{B}'$. Notice that $\mathbf{B}' = \text{HNF}(\mathbf{R})$ and $\mathbf{c}' = (\mathbf{B}\boldsymbol{\gamma} + \mathbf{r}) \bmod \mathbf{B}' = \mathbf{r} \bmod \mathbf{B}'$. Therefore \mathbf{B}' and \mathbf{c}' have the right distribution for algorithm A . Running A on \mathbf{B}' and \mathbf{c}' we will recover (the partial information about) \mathbf{r} with the same success probability as A . \square

5.2 Space Efficiency

We now analyze the size of the keys and the ciphertext of the new encryption algorithm. We assume that the private key satisfies $|r_{i,j}| < \text{poly}(n)$. Therefore the size of the private key can be bounded by $O(n^2 \lg n)$. Using the Hadamard inequality we can also bound the size of the determinant by $O(n \lg n)$, and using the bounds proved in Sect. 3, we get that also the public basis is $O(n^2 \lg n)$ and the ciphertext has size $O(n \lg n)$.

Estimates of the key and ciphertext sizes for the GGH and the modified scheme are shown in Fig. 1. The estimates are based on the GGH challenges published at [13]. Notice that the size of the GGH challenges is much smaller than it should have been to assure adequate randomization. This discrepancy might be explained noticing that the authors of GGH applied the LLL reduction algorithm to the public basis to somehow reduce their size. Nevertheless, the new scheme results in keys and ciphertexts more than an order of magnitude smaller than GGH. We remark that the sizes relative to the modified scheme are only upper bounds obtained using the Hadamard inequality to estimate the determinant of the lattice, and the actual sizes of the keys and ciphertexts of the modified cryptosystem can be even smaller than shown in the table. On the other hand, public-key and ciphertext size might be bigger than the estimate shown in the table if the secret key is generated differently than the GGH challenges.

5.3 Running Time

The experiments described in the next section have been performed using a highly unoptimized prototype implementation, so we do not have meaningful

Table 1. Comparison of the key and ciphertext sizes in the GGH scheme and the modified scheme. All sizes are in kilobytes (KB).

dimension	Basis Size		Ciphertext	
	GGH	New scheme	GGH	New scheme
200	330	32	2	0.16
250	620	50	3	0.20
300	990	75	4	0.25
350	1630	100	5	0.30
400	2370	140	6	0.35

experimental data on the running time of the new lattice scheme. However a few remarks regarding the running time are due.

Encryption time for GGH-like schemes is roughly proportional to the size of the public key. So, if the original GGH cryptosystem was competitive with RSA (see [14]), we should expect the new scheme to outperform RSA in terms of encryption speed because of the reduced key size.

Key Generation was one of the major problems in GGH, requiring the application of LLL on a high dimensional matrix with very large entries. Here, key generation essentially consists of a Hermite Normal Form computation, a much simpler task than lattice reduction. Moreover, recent progress on the design of HNF algorithms [23] might lead to efficient key generation procedures.

The most critical part of lattice based encryption schemes at this point is probably *decryption*. In this paper, and in our experiments we have used Babai’s nearest plane algorithm [4], as we believe this is a quite natural choice. Although polynomial time, this algorithm is very slow compared to the linear time encryption procedure, and decrypting lattices in dimension 400 (using the private basis) can take several minutes with a straightforward implementation. In fact, [14] suggested to use the simpler (but less accurate) rounding off algorithm (also from [4]) instead of nearest plane. It should be noted that the nearest plane algorithm can be made considerably faster if the orthogonalized basis needed by the nearest plane algorithm is precomputed and stored as part of the secret key. Of course, this would increase the size of the secret key, but a relatively poor approximation of the orthogonalized basis should be enough to achieve reasonable correction radius. Decryption methods based on probabilistic rounding procedures as described in [18] are also an interesting alternative to be explored, and much work is still to be done. However, since the decryption procedure is strongly related to the choice of the secret basis, it is probably not worth focusing on the optimization of the decryption algorithm until we get more confidence on what’s a good way to generate the private basis for the lattice. (See Sect. 6 for further discussion about the choice of the private key.)

6 Discussion and Experiments

We defined a trapdoor function f that is at least as hard to break as the GGH “encryption” scheme. Notice that although the original GGH function was a randomized one, the new function is deterministic. Since any semantically secure encryption scheme must be probabilistic [16], the new function f cannot be a secure encryption scheme in the sense of [16]. It is now clear that also GGH (which is less secure than our new function) cannot be a semantically secure encryption scheme, although it is randomized. The situation is similar to other popular “encryption” functions, like RSA: also RSA is a deterministic function, and therefore cannot be a semantically secure encryption scheme if directly applied to the message. In fact, encrypting with RSA usually involves padding the message with some random bits, or applying some other randomized procedure.

In fact, standard techniques can be used to turn trapdoor functions into semantically secure encryption schemes. In the seminal paper [16] Goldwasser and Micali showed that if h is a hard-core predicate for a trapdoor function f , then $E(b, r) = (f(r), h(r) \oplus b)$ is a semantically secure encryption function for one bit messages $b \in \{0, 1\}$. Hard-core predicates from any one-way (or trapdoor) function can be easily constructed following [37, 15], giving a first theoretical construction of a semantically secure cryptosystem. The construction is easily generalized to hard-core functions. Namely, if h is a hard-core function for f , then $E(m, r) = (f(r), h(r) \oplus m)$ is a semantically secure encryption scheme for messages of length equal to the output size of h . Practical instantiations of this scheme can be obtained in the random oracle model [5], simply observing that random oracles are hard-core functions for any one-way function. See [5] also for simple constructions (still in the random oracle model) achieving security against chosen ciphertext attacks [27]. The constructions in [5] have also been recently extended in [26, 11, 12] to allow for probabilistic trapdoor functions, but these extensions are not required because our trapdoor function is deterministic.

A problem that needs further investigation is how to choose the private basis \mathbf{R} . In the GGH cryptosystem \mathbf{R} was chosen as the sum of a multiple of the identity matrix $\sqrt{n}\mathbf{I}$ plus a perturbation matrix \mathbf{Q} with small entries $q_{i,j} \in \{-4, \dots, +4\}$. It is not clear why one should prefer this probability distribution to other distributions, and in fact we think that disclosing the approximate direction of the vectors \mathbf{b}_i might actually weaken the system. At this point of our research, we believe it is better to leave the set of matrices from which \mathbf{R} is chosen as big as possible. A possible way to choose \mathbf{R} is² to choose each entry at random in the interval $\{-n, \dots, n\}$. It turns out that random matrices are pretty good on average [8] and running the LLL algorithm [19] on them rapidly yields matrices \mathbf{R} with relatively large correction radius $\rho = \min_i \|\mathbf{r}_i^*\| = O(n)$. On the other hand, if one applies the LLL basis reduction algorithm to the public basis $\mathbf{B} = \text{HNF}(\mathbf{R})$, although this time LLL takes a much longer time to terminate, the correction radius obtained is much smaller even in relatively low dimension.

² Similar distributions on \mathbf{R} were already considered in [14], thought not used in the construction of the challenges [13].

In Fig. 4 we show some preliminary experimental results obtained running the LLL algorithm on random matrices, and their Hermite normal forms. We observe that random matrices give a correction radius approximately equal to $n/2$, while running the LLL algorithm on the Hermite Normal Forms of the same matrices results in a correction radius that approaches zero as the dimension of the lattice grows.




Fig. 4. Correction radius obtained applying the LLL algorithm to random matrices and their Hermite Normal Forms.

These preliminary data clearly show that applying the HNF algorithm reduces the effectiveness of lattice reduction. Still, the plot in Fig. 4 is not clear evidence of the increased security of the scheme for various reasons: first of all better result can be achieved using more sophisticated basis reduction algorithms (e.g., [29, 30, 32]) than LLL. Moreover, the correction radius is only a worst case measure of the quality of a basis. It is still possible that the HNF basis allows to recover from most small errors.

To support our claim that the modified scheme is secure against the strongest lattice attacks we performed the following experiment. We generated a private basis \mathbf{R} in dimension 400×400 by choosing each entry at random in the interval $\{-400, \dots, +400\}$. After running the fast LLL reduction algorithm, the correc-

tion radius of the random matrix was $\rho = 340$. The public basis $\mathbf{B} = \text{HNF}(\mathbf{R})$ had size 260 KB. Notice that although the size of \mathbf{B} is even smaller than the first GGH challenge (which was about 330 KB), our theoretical analysis suggests that the new function should still be secure because the underlying lattice has much higher dimension.

We then generated random error messages \mathbf{r} choosing each entry at random in the interval $\{-28, \dots, +28\}$ and computed the ciphertexts $\mathbf{c} = \mathbf{r} \bmod \mathbf{B}$. The size of \mathbf{c} was around 800 B. The error vector had length 320 (i.e. less than the correction radius) and could be correctly recovered from \mathbf{c} using \mathbf{R} . We then tried to recover \mathbf{r} using the public basis \mathbf{B} .

Following [24], we first applied a strong basis reduction algorithm (Block Korkine-Zolotarev reduction [29] with block size 20, as implemented in Victor Shoup’s Number Theory Library [33]) to the public basis \mathbf{B} to obtain a reduced basis \mathbf{G} . The computation took over 10 days on a 700 MHz Pentium III workstation. Still the correction radius was only 6.77. Finally, we applied the “embedding technique” (see [24]) to recover \mathbf{r} from \mathbf{c} and \mathbf{G} using the Block KZ reduction algorithm with block size 60 and pruning factor 14 [32]. The clear-text could not be recovered, and even after several days of computation the best lattice vector found by the attack was 500,000 away from the target, over three orders of magnitude worse than the optimal solution $\|\mathbf{r}\| = 320$.

Previous cryptanalytic results [24] warn to be cautious about the security of the cryptosystem, and dimension $n = 400$ should be considered only borderline secure. In fact we suggest the dimension should be at least $n \geq 500$, but only after careful cryptanalysis the minimum dimension for which the system is secure in practice can be determined.

7 Other cryptosystems

In this section we discuss the McEliece and NTRU cryptosystems, and compare them to the general scheme presented in this paper.

7.1 Comparison with the McEliece Cryptosystem

In 1978, McEliece [21] suggested a “cryptosystem” based on the hardness of coding problems that in retrospect is very similar to the GGH cryptosystem. The proposal is to use as a secret key the generating matrix \mathbf{G} of a Goppa code, together with a random permutation matrix \mathbf{P} and a random invertible matrix \mathbf{S} (over $\text{GF}(2)$). The public key is given by the product $\mathbf{G}' = \mathbf{P}\mathbf{G}\mathbf{S}$. Then the trapdoor function is defined by $f(\mathbf{x}, \mathbf{r}) = \mathbf{G}'\mathbf{x} + \mathbf{r}$ (all arithmetic performed over $\text{GF}(2)$), where \mathbf{x} is a random binary vector, and \mathbf{r} is chosen at random among the binary vectors with small Hamming weight. Let $\mathbf{c} = \mathbf{G}'\mathbf{x} + \mathbf{r}$ be the output of the trapdoor function. Using the secret key, we can first compute the permuted target $\mathbf{P}^{-1}\mathbf{c} = \mathbf{G}\mathbf{S} + \mathbf{P}^{-1}\mathbf{r}$. Then we can correct from the small error $\mathbf{P}^{-1}\mathbf{r}$ using the decoding algorithm for Goppa codes, retrieving codeword $\mathbf{G}(\mathbf{S}\mathbf{x})$. Finally, we can compute \mathbf{x} from $\mathbf{S}\mathbf{x}$ solving a system of linear equations

over $\text{GF}(2)$. Notice that this is essentially the same of the GGH cryptosystem with the message m encoded in the coefficients of the lattice vector. A variant of McEliece cryptosystem roughly corresponding to encoding the message in the error vector has also been proposed [25] and the two are known to be equivalent [20]. Interestingly, essentially the same techniques presented in this paper for lattices, can also be used for codes (e.g., they can be applied to [21, 25] and many of their variants). Here instead of the Hermite normal form, we use systematic form for the public code \mathbf{G}' . If $[\mathbf{I}|\mathbf{H}]^T$ is the systematic generating matrix, then the trapdoor function is given by $f(\mathbf{x}, \mathbf{y}) = \mathbf{y} - \mathbf{H}^T \mathbf{x}$, where \mathbf{x} and \mathbf{y} have total weight less than the correction radius of the code. However, the advantage of this transformation for codes is not as good as in the lattice case, giving only a constant (typically a factor 2) improvement over the original McEliece scheme.

A detailed comparison of lattice and coding based cryptosystems is beyond the scope of this paper. However, an apparent advantage of coding based schemes is a potentially smaller public key: since the matrix defining the code has $\{0, 1\}$ entries, the size of the public key is only $O(n^2)$, as opposed to $O(n^2 \log n)$. Of course, this kind of comparisons is not necessarily meaningful if we first do not achieve a better understanding of the relation between the hardness of lattice and coding problems. For example, if decoding binary linear codes in dimension n is easier than breaking lattices in the same dimension, then for a fair comparison of the two schemes one should use different values of the security parameter. In fact, the original McEliece scheme [21] was already proposing codes with block-length $n = 1024$, and recent cryptanalytic work [7] shows that even this large dimension might be insufficient. Interestingly, the nearest codeword problem for binary (or ternary) codes can be efficiently reduced to the closest vector problem over the integers: in order to find the codeword $\mathbf{C}\mathbf{x}$ closest to \mathbf{y} , look for a lattice vector in $[\mathbf{C}|\mathbf{2I}]$ closest to \mathbf{y} . This suggests that lattice problems might be harder than coding problems, at least for the binary case. For larger alphabets, the relation between codes and lattices is less clear. However, if large alphabets are used, then the public key size for code based cryptosystems would increase. For example, if Reed-Solomon codes were used, then the alphabet size would be n , giving keys of total size $O(n^2 \log n)$, matching the asymptotic key size of lattice based cryptosystems. (Here Reed-Solomon codes are just an hypothetical example, as these codes are known not to be secure [34].) We hope that our work will stimulate further investigations on the relation between lattice and coding problems.

7.2 The NTRU cryptosystem

NTRU is a cryptosystem based on polynomial ring arithmetics proposed by Hoffstein, Pipher and Silverman in [17]. The system works essentially as follows. Let n, p, q be system parameters where n is a security parameter (say $n = 200$), p is a small prime (say, $p = 3$) and q is a relatively large prime (typically $q = \Omega(n)$). The secret key is a pair of degree $n - 1$ polynomials $a(X), b(X) \in \mathbb{Z}[X]$ with small coefficients, such that $a(X)$ is invertible modulo $(X^n - 1, pq)$. The public key is given by $c(X) = p \cdot g(X) / f(X) \bmod (X^n - 1, q)$. The encryption function

takes as input two polynomials $m(X), r(X) \in \mathbb{Z}[X]/(X^N - 1)$ with small coefficients (the first interpreted as a message, and the second as a randomizer), and outputs $c(X) \cdot r(X) + m(X) \bmod (X^n - 1, q)$. (For the decryption procedure, as well as a more detailed description of the system, see the original article [17].) As for the GGH and the McEliece cryptosystem, NTRU does not provide semantic security, so it is better described as a deterministic trapdoor function, instead of a full fledged probabilistic cryptosystem. Interestingly, this function can be formulated in terms of lattices as follows. Consider the lattice generated by the $2n \times 2n$ matrix

$$\left[\begin{array}{c|c} c \cdot \mathbf{I} & \mathbf{C} \\ \hline \mathbf{0} & \mathbf{I} \end{array} \right]$$

where \mathbf{C} is the $n \times n$ matrix whose rows are given by all cyclic permutations of the coefficients $[c_0, c_1, \dots, c_{n-1}]$ of polynomial $c(X) = \sum_{i=0}^{n-1} c_i X^i$, i.e., $\mathbf{C}_{i,j} = c_{(j-i \bmod n)}$. Notice that this public basis is in Hermite normal form. Moreover, it is easy to see that the output $c(X)r(X) + m(X)$ of the trapdoor function is exactly the result of reducing vector $[m_{n-1}, \dots, m_0, -r_{n-1}, \dots, -r_0]^T$ modulo the public lattice basis. So, when viewed as a lattice based trapdoor function, NTRU has the same high level structure of the functions described in this paper: the public key is an HNF lattice basis, and the function is computed reducing small “error” vector modulo the public basis. What sets NTRU apart from all other lattice based functions is the use of a class of lattices with special structure: convolutional modular lattices. While the security offered by this class of lattices is still largely to be investigated, the performance advantage is clear: as the HNF basis can be represented by only n ($\log n$)-bit numbers, the public key is much smaller than those obtained using general lattices which require $(n^2 \log n)$ bits.

8 Conclusion

We presented a new trapdoor function based on the hardness of lattice problems. The trapdoor function can be transformed into a full fledged encryption algorithm using standard techniques [16, 5]. The new function can be formally proved to be at least as secure as any other function from a general scheme to design lattice based trapdoor functions that includes the GGH trapdoor function [14] and the tensor based trapdoor function [10] as special cases. Moreover, the new function substantially improves previous proposals from the efficiency point of view: for the same level of security the new function reduces both the time and space requirements by a factor $O(n)$. The improved efficiency allows to use bigger values of the security parameter while maintaining the scheme reasonably practical. One last advantage of the new scheme is simplicity. While previous schemes computed the public key and function values using a substantial amount of randomness, in the new function these operations are substituted by simple deterministic procedures. This is important both from the theoretical and practical point of view, because it makes the algorithms easier to implement and also easier to analyze.

At this point the main question about lattice based cryptography is how to choose the private key, i.e., finding families of easily decodable lattices for which decoding becomes infeasible when the lattice is presented in Hermite normal form. We believe that in order to be really competitive with RSA, key sizes even smaller than those achieved in this paper would be desirable. However, the public key size cannot be further reduced unless one considers classes of lattices with special structure. (A simple counting argument shows that the number of lattices in a certain dimension is exponential in the bit-size representation of their Hermite normal forms, so the HNF representation is essentially optimal if one considers arbitrary lattices.) The search for easily decodable lattices for which decoding is hard when the lattice is presented in Hermite normal form becomes particularly interesting if one could find special classes of hard lattices that have small HNF representation. We observed that one such family of lattices is given by the NTRU trapdoor function. Again, the main question is security. Are the convolutional modular lattices used by NTRU really hard to decode? We hope that our work will stimulate further research on the computational complexity of decoding this and other classes of lattices.

9 Acknowledgements

The author wishes to thank Shai Halevi, Shafi Goldwasser and Salil Vadhan for many stimulating conversations on lattice based cryptography. Thanks also to the anonymous referees for their useful comments.

References

- [1] M. Ajtai. Generating hard instances of lattice problems (extended abstract). In *Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing*, pages 99–108, Philadelphia, Pennsylvania, 22–24 May 1996.
- [2] M. Ajtai and C. Dwork. A public-key cryptosystem with worst-case/average-case equivalence. In *Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing*, pages 284–293, El Paso, Texas, 4–6 May 1997.
- [3] S. Arora, L. Babai, J. Stern, and E. Z. Sweedyk. The hardness of approximate optima in lattices, codes, and systems of linear equations. *J. Comput. Syst. Sci.*, 54(2):317–331, Apr. 1997. Preliminary version in FOCS'93.
- [4] L. Babai. On Lovasz' lattice reduction and the nearest lattice point problem. *Combinatorica*, 6(1):1–13, 1986.
- [5] M. Bellare and P. Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *Proceedings of the first ACM Conference on Computer and Communications Security*. ACM, Nov. 1993.
- [6] J.-Y. Cai and T. W. Cusick. A lattice-based public-key cryptosystem. *Information and Computation*, 151(1–2):17–31, May–June 1999.
- [7] A. Canteaut and N. Sendrier. Cryptanalysis of the original McEliece cryptosystem. In K. Ohta and D. Pei, editors, *Advances in Cryptology — Proceedings of Asiacrypt 98*, volume 1514 of *Lecture Notes in Computer Science*, pages 187–199, Beijing, China, 1998.

- [8] H. Daude and B. Vallée. An upper bound on the average number of iterations of the LLL algorithm. *Theoretical Computer Science*, 123(1):95–115, Jan. 1994.
- [9] I. Dinur, G. Kindler, and S. Safra. Approximating CVP to within almost-polynomial factors is NP-hard. In *39th Annual Symposium on Foundations of Computer Science*, Palo Alto, California, 7–10 Nov. 1998. IEEE.
- [10] R. Fischlin and J.-P. Seifert. Tensor-based trapdoors for CVP and their application to public key cryptography. In *7th IMA International Conference "Cryptography and Coding"*, volume 1746 of *Lecture Notes in Computer Science*, pages 244–257. Springer-Verlag, 1999.
- [11] E. Fujisaki and T. Okamoto. Secure integration of asymmetric and symmetric encryption schemes. In M. Wiener, editor, *Advances in Cryptology—CRYPTO '99*, volume 1666 of *Lecture Notes in Computer Science*, pages 537–554, University of California, Santa Barbara, Aug. 1999. IACR, Springer-Verlag.
- [12] E. Fujisaki and T. Okamoto. How to enhance the security of public-key encryption at minimum cost. *IEICE Transaction of Fundamentals of electronic Communications and Computer Science*, E38-A(1):24–32, Jan. 2000.
- [13] O. Goldreich, S. Goldwasser, and S. Halevi. The GGH cryptosystem, challenge page. <http://theory.lcs.mit.edu/~cis/lattice/challenge.html>.
- [14] O. Goldreich, S. Goldwasser, and S. Halevi. Public-key cryptosystems from lattice reduction problems. In B. S. Kaliski Jr., editor, *Advances in Cryptology—CRYPTO '97*, volume 1294 of *Lecture Notes in Computer Science*, pages 112–131. Springer-Verlag, 17–21 Aug. 1997.
- [15] O. Goldreich and L. Levin. A hard predicate for all one-way functions. In *Proceedings of the 21st Annual Symposium on Theory of Computing (STOC)*. ACM, 1989.
- [16] S. Goldwasser and S. Micali. Probabilistic encryption. *Journal of Computer and System Science*, 28(2):270–299, 1984. Preliminary version in STOC'82.
- [17] J. Hoffstein, J. Pipher, and J. H. Silverman. NTRU: A ring based public key cryptosystem. In J. Buhler, editor, *Algorithmic Number Theory (ANTS III)*, volume 1423 of *Lecture Notes in Computer Science*, pages 267–288, Portland, OR, 1998. Springer.
- [18] P. Klein. Finding the closest lattice vector when it's unusually close. In *Proceedings of the 11th Symposium on Discrete Algorithms*, San Francisco, California, Jan. 2000. SIAM.
- [19] A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261:513–534, 1982.
- [20] Y. X. Li, R. H. Deng, and X. M. Wang. On the equivalence of McEliece's and Niederreiter's public-key cryptosystems. *IEEE Transactions on Information Theory*, 40(1):271–273, Jan. 1994.
- [21] R. J. McEliece. A public-key cryptosystem based on algebraic coding theory. DSN Progress Report 42-44, Jet Propulsion Laboratory, Pasadena, 1978.
- [22] D. Micciancio. The hardness of the closest vector problem with preprocessing. *IEEE Transactions on Information Theory*, 2001. To Appear.
- [23] D. Micciancio and B. Warinschi. A linear space algorithm for computing the Hermite Normal Form. In B. Mourrain, editor, *International Symposium on Symbolic and Algebraic Computation*. ACM, ACM, 2001. To Appear.
- [24] P. Nguyen. Cryptanalysis of the Goldreich-Goldwasser-Halevi cryptosystem from Crypto '97. In M. Wiener, editor, *Advances in Cryptology—CRYPTO '99*, volume 1666 of *Lecture Notes in Computer Science*. Springer-Verlag, Aug. 1999.
- [25] H. Niederreiter. Knapsack-type cryptosystems and algebraic coding theory. *Problems of Control and Information Theory*, 15(2):159–166, 1986.

- [26] T. Okamoto and D. Pointcheval. React: Rapid enhanced-security asymmetric cryptosystem transform. In D. Naccache, editor, *Proceedings of the Cryptographers' Track of the RSA Conference '2001 (RSA '2001)*, Lecture Notes in Computer Science, San Francisco, California, USA, 8–12Apr. 2001. Springer-Verlag.
- [27] C. Rackoff and D. R. Simon. Non-interactive zero-knowledge proof of knowledge and chosen ciphertext attack. In J. Feigenbaum, editor, *Advances in Cryptology: Proceedings of Crypto 91*, volume 576 of *Lecture Notes in Computer Science*, University of California, Santa Barbara, Aug. 1991. IACR, Springer-Verlag.
- [28] R. L. Rivest, A. Shamir, and L. Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Communications of the ACM*, 21:120–126, 1978.
- [29] C.-P. Schnorr. A hierarchy of polynomial time lattice basis reduction algorithms. *Theoretical Computer Science*, 53(2–3):201–224, 1987.
- [30] C.-P. Schnorr and M. Euchner. Lattice basis reduction: Improved practical algorithms and solving subset sum problems. In L. Budach, editor, *Proceedings of Fundamentals of Computation Theory*, volume 529 of *lncs*, pages 68–85. Springer-Verlag, 1991.
- [31] C.-P. Schnorr, M. Fischlin, H. Koy, and A. May. Lattice attacks on GGH cryptosystem. Rump session of Crypto'97, 1997.
- [32] C.-P. Schnorr and H. H. Hörner. Attacking the Chor–Rivest cryptosystem by improved lattice reduction. In L. C. Guillou and J.-J. Quisquater, editors, *Advances in Cryptology—EUROCRYPT 95*, volume 921 of *Lecture Notes in Computer Science*, pages 1–12. Springer-Verlag, 21–25 May 1995.
- [33] V. Shoup. NTL: A library for doing number theory. URL <http://www.shoup.net/ntl/index.html>.
- [34] V. Sidelnikov and S. Shestakov. On cryptosystems based on generalized Reed-Solomon codes. *Diskretnaya Math*, 4(3):57–63, 1992. In Russian.
- [35] N. J. A. Sloane. Encryption by random rotations. In *Workshop on Cryptography Burg Feuerstein 1982*, volume 149 of *Lecture notes in computer science*, pages 71–129, 1983.
- [36] P. van Emde Boas. Another NP-complete problem and the complexity of computing short vectors in a lattice. Technical Report 81-04, Mathematische Instituut, University of Amsterdam, 1981. Available on-line at URL <http://turing.wins.uva.nl/~peter/>.
- [37] A. Yao. Theory and applications of trapdoor functions. In *Proceedings of the 23rd IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 80–91, Chicago, IL, 1982. IEEE.