Explicit hard instances of the shortest vector problem

– Extended Version –

Johannes Buchmann, Richard Lindner, and Markus Rückert

Technische Universität Darmstadt, Department of Computer Science
Hochschulstraße 10, 64289 Darmstadt, Germany
buchmann,rlindner,rueckert@cdc.informatik.tu-darmstadt.de

Abstract. Building upon a famous result due to Ajtai, we propose a sequence of lattice bases with growing dimension, which can be expected to be hard instances of the shortest vector problem (SVP) and which can therefore be used to benchmark lattice reduction algorithms.
The SVP is the basis of security for potentially post-quantum cryptosystems. We use our sequence of lattice bases to create a challenge, which may be helpful in determining appropriate parameters for these schemes.

Keywords: Lattice reduction, lattice-based cryptography, challenge

1 Introduction

For the construction of post-quantum cryptosystems, it is necessary to identify computational problems, whose difficulty can be used as a basis of the security for such systems, and that remain difficult even in the presence of quantum computers. One candidate is the problem of approximating short vectors in a lattice (shortest vector problem — SVP). The quantum-hardness of this problem was analyzed by Ludwig [24] and Regev [32]. They both find that the computational advantage gained with quantum computers is marginal. There are several cryptographic schemes whose security is based on the intractability of the SVP in lattices of sufficiently large dimension (e.g. [15,16,3,33]). To determine appropriate parameters for these cryptosystems, it is necessary to assess the practical difficulty of this problem as precisely as possible.

In this paper, we present a sequence of lattice bases with increasing dimension, which we propose as a world wide challenge. The construction of these lattices is based both on theoretical and on practical considerations. On the theoretical side, we apply a result of Ajtai [2]. It states that being able to find a sufficiently short vector in a random lattice from a certain set, which also contains our challenge lattices, implies the ability to solve supposedly hard problems (cf. [34]) in all lattices with a slightly smaller dimension than that of the random lattice. Furthermore, we invoke a theorem of Dirichlet on Diophantine approximation (cf. [19]). It guarantees the existence of a short vector in each challenge lattice. On the practical side, using an analysis by Gama and Nguyen [12], we

argue that finding this vector is hard for the lattices in our challenge. We also present first experimental results that confirm the analysis.

Our challenge at `http://www.latticechallenge.org` can be considered as an analogue of similar challenges for the integer factoring problem [35] and the problems of computing discrete logarithms in the multiplicative group of a finite field [26], or in the group of points on an elliptic curve over a finite field [9].

Our aim is to evaluate the current state-of-the-art in practical lattice basis reduction by providing means for an immediate and well-founded comparison. As a first application of the proposed challenge, we compare the performance of LLL-type reduction methods — LLL [23], Stehlé's fpLLL [29], Koy and Schnorr's segment LLL (sLLL) [21] — and block-type algorithms — Schnorr's BKZ [38,37], Koy's primal-dual (PD) [20], Ludwig's practical random sampling [1] (PSR) [25]. To our knowledge, this is the first comparison of these algorithms.

Related work. Lattice reduction has been subject to intense studies over the last decades, where a couple of methods and reduction schemes, in particular the LLL algorithm by Lenstra, Lenstra, and Lovász [23], have been developed and successively improved. Especially, the block Korkine Zolorarev algorithm (BKZ), due to Schnorr [37,38], has become the standard method when strong lattice basis reduction is required.

There have been several approaches to measure the effectiveness of known lattice reduction algorithms, especially in the context of the NTRU cryptosystem [16]. Some of them, as in [17,18], base their analysis on cryptosystems while others, like [30,12], make a more general approach using random lattices.

To our knowledge, there has never been a unified challenge, one that is independent of a specific cryptosystem, for lattice reduction algorithms. In all previous challenges, the solution was always known to the creator.

Organization. In Section 2, we provide a brief introduction to lattices and state some fundamental definitions. In Section 3, we define a family of lattices and prove two properties, which are fundamental for our explicit construction presented in Section 4. Then, we give first experimental results comparing the performance of various lattice reduction algorithms in Section 5. Finally, Section 6 introduces the actual lattice challenge.

2 Preliminaries

Let $\mathbb{R}^n$ denote the n-dimensional real vectorspace. We write the vectors of this space in boldface to distinguish them from numbers. Any two vectors $\mathbf{v}, \mathbf{w} \in \mathbb{R}^n$ have an inner product $\langle \mathbf{v}, \mathbf{w} \rangle = \mathbf{v}^T \mathbf{w}$. Any $\mathbf{v} \in \mathbb{R}^n$ has a length given by the Euclidean norm $\|\mathbf{v}\|_2 = \sqrt{\langle \mathbf{v}, \mathbf{v} \rangle} = \sqrt{v_1^2 + \cdots + v_n^2}$. In addition to the Euclidean norm, we also use the maximum norm $\|\mathbf{v}\|_\infty = \max_{i=1,\ldots,n} \{|v_i|\}$.

A lattice in $\mathbb{R}^n$ is a set $L = \{\sum_{i=1}^m x_i \mathbf{b}_i \,|\, x_i \in \mathbb{Z}\}$, where $\mathbf{b}_1, \ldots, \mathbf{b}_m$ are linearly independent over $\mathbb{R}$. The matrix $B = [\mathbf{b}_1, \ldots, \mathbf{b}_m]$ is called a *basis* of

[1] A practical variant of Schnorr's random sampling reduction [39].

the lattice L and we write $L = L(B)$. The number of linearly independent vectors in the basis is the dimension of the lattice. If $\dim(L(B)) = n$ the lattice is full-dimensional.

An m-dimensional lattice $L = L(B)$ has many different bases, namely all the matrices in the orbit $B\,\mathrm{GL}_m(\mathbb{Z}) = \{BT \,|\, T \in \mathrm{GL}_m(\mathbb{Z})\}$. If the lattice is full-dimensional and integral, that is $L \subseteq \mathbb{Z}^n$, then there exists a unique basis $B = (b_{i,j})$ of L, which is in Hermite normal form (HNF), i.e.

i. $b_{i,j} = 0$ for all $1 \leq j < i \leq m$
ii. $b_{i,i} > b_{i,j} \geq 0$ for all $1 \leq i < j \leq m$

Furthermore, the volume $\mathrm{vol}(L)$ of a full-dimensional lattice is defined as $|\det(B)|$, for any basis B of L. For every m-dimensional lattice L there is a dual (or polar, reciprocal) lattice $L^* = \{\mathbf{x} \in \mathbb{R}^m \,|\, \forall \mathbf{y} \in L : \langle \mathbf{x}, \mathbf{y} \rangle \in \mathbb{Z}\}$. For any full-dimensional lattice $L = L(B)$, it holds that $L^* = L((B^{-1})^T)$. The length of the shortest lattice vector, denoted with $\lambda_1 = \lambda_1(L)$, is called first successive minimum.

3 Foundations of the challenge

In this section, we define a family of sets containing lattices, where each set will have two important properties:

1. All lattices in the set contain non-obvious short vectors;
2. Being able to find a short vector in a lattice chosen uniformly at random from the set, implies being able to solve difficult computational problems in all lattices of a certain smaller dimension.

The family of lattice sets. Let $n \in \mathbb{N}$, $c_1, c_2 \in \mathbb{R}_{>0}$, such that

$$\frac{1}{2\,\ln(2)} \leq c_2 \leq \frac{c_1}{4} \ln\left(\frac{n}{c_1 \ln(n)}\right) . \tag{1}$$

Furthermore, let

$$m = \lfloor c_1 n \ln(n) \rfloor , \tag{2}$$
$$q = \lfloor n^{c_2} \rfloor , \tag{3}$$

and $\mathbb{Z}_q = \{0, \ldots, q-1\}$. For a matrix $X \in \mathbb{Z}_q^{n \times m}$, with column vectors $\mathbf{x}_1, \ldots, \mathbf{x}_m$, let

$$L(c_1, c_2, n, X) = \left\{ (v_1, \ldots, v_m) \in \mathbb{Z}^m \;\middle|\; \sum_{i=1}^{m} v_i\, \mathbf{x}_i \equiv \mathbf{0} \pmod{q} \right\} .$$

All lattices in the set $L(c_1, c_2, n, \cdot) = \{L(c_1, c_2, n, X) | X \in \mathbb{Z}_q^{n \times m}\}$ are of dimension m and the family of lattices $\mathfrak{L}$ is the set of all $L(c_1, c_2, n, \cdot)$, such that c_1, c_2, n are chosen according to (1).

In the following theorems, we prove that all lattices in the sets of the family $\mathfrak{L}$ have the desired properties.

Existence of short vectors. We prove that all lattices in $L(c_1, c_2, n, \cdot)$ of the family $\mathfrak{L}$ contain a vector with Euclidean norm less than n.

Theorem 1. *Let $n \in \mathbb{N}, c_1, c_2 \in \mathbb{R}_{>0}$, and $q, m \in \mathbb{N}$ be as described above. Then, any lattice in $L(c_1, c_2, n, \cdot) \in \mathfrak{L}$ contains a vector with Euclidean norm less than n.*

Proof. Let $L(c_1, c_2, n, X) \in L(c_1, c_2, n, \cdot) \in \mathfrak{L}$. We first show that any solution of a certain Diophantine approximation problem corresponds to a vector in $L(c_1, c_2, n, X)$. Then, we use a theorem of Dirichlet to establish the existence of a non-zero lattice vector of length less than n.

Let $\mathbf{v} \in L(c_1, c_2, n, X)$, then there exists $\mathbf{w} \in \mathbb{Z}^n$, such that

$$\frac{1}{q} X \mathbf{v} - \mathbf{w} = 0 \,.$$

This is equivalent to

$$\left\| \frac{1}{q} X \mathbf{v} - \mathbf{w} \right\|_\infty < \frac{1}{q} \,. \tag{4}$$

Dirichlet's theorem (cf. [36,19]) states that for any $t > 1$, there is $\mathbf{v} \in \mathbb{Z}^m$ and $\mathbf{w} \in \mathbb{Z}^n$, such that

$$\left\| \frac{1}{q} X \mathbf{v} - \mathbf{w} \right\|_\infty < e^{-\frac{t}{n}} \qquad \text{and} \tag{5}$$

$$\|\mathbf{v}\|_\infty < e^{\frac{t}{m}} \,. \tag{6}$$

We set $t = n \ln(q)$. Then, (5) implies that (4) is satisfied. It remains to prove that $\|\mathbf{v}\|_\infty < n/\sqrt{m}$ because this implies $\|\mathbf{v}\|_2 < n$. Using (6), we have

$$\|\mathbf{v}\|_\infty < e^{\frac{t}{m}} \leq e^{\frac{n \ln(q)}{m}} \leq e^{\frac{n \ln(\lfloor n^{c_2} \rfloor)}{\lfloor c_1 n \ln(n) \rfloor}} \overset{*}{\leq} e^{\frac{2 n c_2 \ln(n)}{c_1 n \ln(n)}} \leq e^{\frac{2 c_2}{c_1}} \,.$$

For a rigorous proof of inequality $*$ see Appendix A. Together with (1), this evaluates to

$$e^{\frac{2 c_2}{c_1}} \leq e^{\frac{2 c_1}{4 c_1} \ln\left(\frac{n}{c_1 \ln(n)}\right)} \leq \sqrt{\frac{n}{c_1 \ln(n)}} \leq \frac{n}{\sqrt{m}} \,,$$

which completes the proof. □

Hardness of finding short vectors. In the following, we show that being able to find short vectors in an m-dimensional lattice chosen uniformly at random from $L(c_1, c_2, n, \cdot) \in \mathfrak{L}$, implies being able to solve (conjectured) hard lattice problems for *all* lattices of dimension n.

In his seminal work [2], Ajtai proved the following theorem that connects average-case instances of certain lattice problems to worst-case instances. The problems are defined as follows.

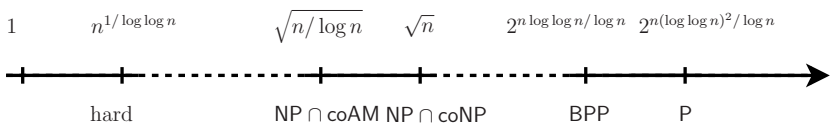

Fig. 1. The complexity of γ-SVP for increasing γ (some constants omitted).

Lattice problems. Let $L \subseteq \mathbb{Z}^n$ be an n-dimensional lattice and $\gamma \geq 1$. We define the

- Approximate shortest length problem (γ-SLP):
  Find $l \in \mathbb{R}$, such that $l \leq \lambda_1(L) \leq \gamma\, l$.
- Approximate shortest vector problem (γ-SVP):
  Find a vector $\mathbf{v} \in L \setminus \{\mathbf{0}\}$, such that for all $\mathbf{w} \in L : \|\mathbf{v}\|_2 \leq \gamma\, \|\mathbf{w}\|_2$.
- Approximate shortest basis problem (γ-SBP):
  Find a basis B of L, such that for all $C \in B\,\mathrm{GL}_m(\mathbb{Z})$:

$$\max_{i=1,2,\ldots,n} \|\mathbf{b}_i\|_2 \leq \gamma \max_{i=1,2,\ldots,n} \|\mathbf{c}_i\|_2 \ .$$

Theorem 2 ([2, Theorem 1]). *Let $c > 1$ be an absolute constant. If there exists a probabilistic polynomial time (in n) algorithm $\mathcal{A}$ that finds a vector of norm $< n$ in a random m-dimensional lattice from $L(c_1, c_2, n, \cdot) \in \mathfrak{L}$ with probability $\geq 1/2$ then there exists*

1. *an algorithm $\mathcal{B}_1$ that solves the γ-SLP;*
2. *an algorithm $\mathcal{B}_2$ that solves the SVP, provided that the shortest vector is γ-unique* [2]*;*
3. *an algorithm $\mathcal{B}_3$ that solves the γ-SBP.*

Algorithms $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3$ solve the respective problem (each with $\gamma = n^c$) with probability exponentially close to 1 in all *lattices of dimension n, i.e. especially in the worst-case. $\mathcal{B}_1, \mathcal{B}_2$, and $\mathcal{B}_3$ run in probabilistic polynomial time in n.*

As for the constant c in Theorem 2, there have been several improvements to Ajtai's reduction with $c \geq 8$ [8]. The first improvement ($c = 3.5 + \epsilon$) is due to Cai and Nerurkar [8], whereas the most recent works by Miccancio [27] and Micciancio and Regev [28], improve c to almost [3] 1.

Asymptotic and practical hardness of the above problems depends on the choice of γ. A recent survey [34] by Regev states the currently known "approximability" and "inapproximability" results. As for the complexity of lattice problems, it focuses on the works of Lagarias, Lenstra, and Schnorr [22], Banaszczyk [7], Goldreich and Goldwasser [14], Ajtai, Kumar, and Sivakumar [4], Aharonov and Regev [1], and Peikert [31]. Since it is very helpful and descriptive, we adopted Figure 1 from the survey.

[2] A shortest vector $\mathbf{v} \in L$ is γ-*unique* if for all $\mathbf{w} \in L$ with $\|\mathbf{w}\|_2 \leq \gamma\ \|\mathbf{v}\|_2 \Rightarrow \mathbf{w} = \pm\mathbf{v}$.
[3] Omitting poly-logarithmic terms in the resulting approximation factor.

On the left, there are provably NP-hard problems, followed by a gap for which the hardness is unknown. In the center, there are problems conjectured not to be NP-hard because their NP-hardness would contradict the general perception that coNP $\neq$ NP. Finally, on the right, there are problems that can be solved in probabilistic polynomial time.

We emphasize that the problems in Theorem 2 are *not* believed to be NP-hard because $\gamma > \sqrt{n}$. Nevertheless, there is no known algorithm that efficiently solves worst-case instances of lattice problems for sufficiently large dimensions n, with an approximation factor polynomial in n. So Theorem 2 strongly supports our claim that computing short vectors in the lattice family is hard. This is also supported by a heuristic argument of Gama and Nguyen [12], which we refer to in Section 4.

4 Construction of explicit bases

Ajtai's construction in [2] defines all lattices implicitly. In this section, we show how to generate explicit integral bases for these lattices.

For any $m \geq 500$, we now construct a lattice L_m of dimension m, which is our hard instance of the SVP. The lattice L_m is of the form $L(c_1, c_2, n, X)$, where the parameters c_1, c_2, n, X are chosen as a function of the dimension m as follows.

We start with a desired lattice dimension m, set $c_2 = 1$, and choose $c_1, n = n(m)$ such that (1) and (2) hold. This is done by setting

$$c_1 = \inf\{c \in \mathbb{R} \,|\, \exists n \in \mathbb{N} : m = \lfloor c\, n\, \ln(n) \rfloor \wedge c_2 \leq c\, \ln(n/(c\, \ln(n)))/4\}\,, \quad (7)$$

$$n(m) = \max\{n \in \mathbb{N} \,|\, m = \lfloor c_1\, n\, \ln(n) \rfloor \wedge c_2 \leq c_1\, \ln(n/(c_1\, \ln(n)))/4\}\,. \quad (8)$$

With $m = 500$, for example, we get $c_1 = 1.9453, c_2 = 1$, and $n = q = 63$.

Having selected the set $L(c_1, c_2, n, \cdot)$, we "randomly" pick a lattice from it. We use the digits of π as a source of "randomness" [4]. This approach is supported by the conjectured normalcy of π in [5,6]. We write

$$3.\pi_1\, \pi_2\, \pi_3\, \pi_4\, \ldots\,,$$

so π_i, for $i \geq 1$, is the ith decimal digit of π in the expansion after the decimal point. In order to compensate for potential statistical bias, we define

$$\pi_i^* = \pi_{2\,i} + \pi_{2\,i-1} \mod 2 \qquad \text{for } i \geq 1\,.$$

Now, we use the sequence $(\pi_1^*, \pi_2^*, \pi_3^*, \pi_4^*, \ldots)$ as a substitute for a sequence of uniformly distributed random bits.

The matrix $X = (x_{i,j}) \in \mathbb{Z}_q^{n \times m}$ is chosen via

$$x_{i,j} = \sum_{l=k}^{k+\lfloor \log_2(q) \rfloor} 2^{l-k}\, \pi_l^* \bmod q \qquad \text{for } 1 \leq i \leq n, 1 \leq j \leq m\,,$$

$$\text{with} \qquad k = k(i,j) = ((i-1)\, m + (j-1))\ \lfloor \log_2(q) \rfloor + 1\,.$$

[4] The digits of π can be optained from `ftp://pi.super-computing.org/`.

With that, we have selected a "random" element $L(c_1, c_2, n, X)$, for which we will now generate an integral basis.

Let I_m be the m-dimensional identity matrix. We start with the matrix

$$Y_1 = (X^T \mid q\, I_m) = \left(\begin{array}{ccc|ccc} x_{1,1} & \cdots & x_{n,1} & q & 0 & \cdots & 0 \\ x_{1,2} & \cdots & x_{n,2} & 0 & q & & \vdots \\ \vdots & \ddots & \vdots & \vdots & & \ddots & 0 \\ x_{1,m} & \cdots & x_{n,m} & 0 & \cdots & 0 & q \end{array}\right).$$

Let Y_2 be the Hermite normal form of Y_1, we compute the transformation matrix T_1, which satisfies

$$Y_2\, T_1 = Y_1 = (X^T \mid qI_m)\,.$$

We set T_2 to be equal to T_1, but without the n leading columns. This guarantees that

$$Y_2\, T_2 = q\, I_m\,. \tag{9}$$

Finally, we set the basis to $B = T_2^T$.

Now, we have to show that B is an integral basis of $L(c_1, c_2, n, X)$. Clearly, B is an integral matrix because the transformation T_1, given by the HNF computation, is in $\mathbb{Z}^{m\times(n+m)}$ and T_2 is the same matrix with the n leading columns removed.

By the uniqueness of inverses, (9) shows that $B = ((Y_2/q)^{-1})^T$. This implies that B is a basis for the dual lattice of $L(Y_2/q)$ (cf. Section 2). Since Y_2 is an integral transformation of Y_1, they span the same lattice. Thus, $L(Y_2/q) = L(Y_1/q)$.

By the defining property of the dual lattice, we have that for any $\mathbf{v} \in L(B)$ and $\mathbf{w} \in L(Y_1/q)$, it holds that $\langle \mathbf{v}, \mathbf{w}\rangle \in \mathbb{Z}$. So especially for all columns $\mathbf{x}$ of X^T, it holds that $\langle \mathbf{v}, \mathbf{x}/q\rangle \in \mathbb{Z}$, or equivalently $\langle \mathbf{v}, \mathbf{x}\rangle \in q\mathbb{Z}$. This implies $\langle \mathbf{v}, \mathbf{x}\rangle \bmod q = 0$, which in turn gives us $L(B) \subseteq L(c_1, c_2, n, X)$.

Now let $\mathbf{v} \in L(c_1, c_2, n, X)$, so for any column $\mathbf{x}$ of X^T we have that the inner product $\langle \mathbf{v}, \mathbf{x}\rangle \bmod q = 0$, or equivalently $\langle \mathbf{v}, \mathbf{x}/q\rangle \in \mathbb{Z}$. Since we know $L(c_1, c_2, n, X) \subseteq \mathbb{Z}^m$, it also holds that $\langle \mathbf{v}, \mathbf{e}\rangle \in \mathbb{Z}$ for any column $\mathbf{e}$ of the identity matrix I_m. Since $\mathbf{v}$ has an integral inner product with each column vector in Y_1/q, this means $\mathbf{v}$ is in the dual lattice of $L(Y_1/q)$, which we know to be $L(B)$. Finally, we have $L(B) = L(c_1, c_2, n, X)$.

For a small example of such a basis, refer to Appendix C.

The choice of parameters. We now argue that our choice of the paramters leads to m-dimensional lattices $L_m = L(c_1, c_2, n, X)$, in which vectors of norm less than $n(m)$ are hard to find.

We have chosen $c_2 = 1$. By Theorem 1, this guarantees the existence of lattice vectors with norm less than $n(m) = q$ in L_m.

A choice of $c_2 < 1$, and thus $q < n$, would imply that all q-vectors, namely vectors that are zero except for one entry q, in $\mathbb{Z}^m$ have Euclidean norm less than $n(m)$. This renders the lattice challenge preposterous because q-vectors are

m	n, q	γ
500	63	1.0072^m
825	127	1.0050^m
1000	160	1.0042^m
1250	208	1.0036^m
1500	256	1.0031^m
1750	304	1.0027^m
2000	348	1.0024^m

Table 1. Lattice parameters with the necessary Hermite factor γ.

easy to find. Moreover, Theorem 1 only guarantees the existence of one short vector, which in this case might be a q-vector.

On the other hand, choosing $c_2 > 1$ enlarges c_1, and because of (2) decreases $n(m)$. Then, the hardness of lattice problems in a large dimension m would be based on the worst-case hardness of lattice problems in a very small dimension n. As n decreases, our hardness argument becomes less meaningful because even worst-case lattice problems in small dimensions are believed to be easy.

Table 1 shows how m and n are related for the selected lattices L_m. For a graphical overview, up to $m = 2000$, refer to Appendix B. Thus, in order to apply Theorem 2 as a strong indication for hardness, we keep $n(m)$ close to m in the above construction. We choose a pseudo-random X to get a random element in $L(c_1, c_2, n, \cdot)$, as required by Theorem 2. Using the recent improvement of Ajtai's result due to Gentry, Peikert, and Vaikuntanathan [13], it is possible to choose c_2 arbitrarily close to 1. Their results can also be used to improve our construction, by providing an even stronger indication of hardness.

To give an even stronger argument for the hardness of the SVP in our lattices, we use a result by Gama and Nguyen [12]. They argue that finding vectors $\mathbf{v}$ in a lattice L is difficult if

$$\|\mathbf{v}\| < \gamma \mathrm{vol}(L)^{1/m}, \tag{10}$$

where $\gamma \leq 1.01^m$ and m is the dimension of L. In this inequality, γ is called Hermite factor. For $\gamma \leq 1.005^m$ Gama and Nguyen state that computing vectors $\mathbf{v}$ that satisfy (10) is "totally out of reach".

Finding a vector $\mathbf{v} \in L_m$ of length less than $n(m)$ means finding a vector $\mathbf{v}$ that satisfies (10) with Hermite factor

$$\gamma < \frac{n(m)}{\mathrm{vol}(L_m)^{1/m}}.$$

Such Hermite factors are tabulated in column 3 of Table 1.

In combination with the analysis of Gama and Nguyen, the table suggests that while finding a vector shorter than $n(m)$ in L_{500} is still possible, the respective problem in L_{825} will be very hard in practice. As the dimension increases, the necessary Hermite factor falls below 1.004^n and 1.003^n. We think that finding short vectors in the corresponding lattices will require entirely new algorithms.

5 Experiments with lattice reduction algorithms

As a first application of our explicit construction of lattices L_m, we show how various lattice reduction algorithms perform on them. Basically, there are two types of algorithms: the LLL-type and the block-type. Building upon LLL, block-type algorithms are typically stronger, in the sense that they are able to find significantly shorter vectors. Block-type algorithms, however, are impractical for large block sizes because their running time increases at least exponentially in this parameter.

Toy challenges. In Section 4, we have seen that the problem of finding a vector of length less than $n(m)$ in lattices L_m starts to become difficult for $m \geq 500$ and it should be infeasible for $m \geq 825$.

Thus, we define a relaxed variant of the family $\mathfrak{L}$. It is the family of all lattice sets $L(2, 1, n, \cdot)$, i.e. we set $c_2 = 1$ and $c_1 = 2$, so (1) does not necessarily hold. Although, in such lattices, there is no guarantee for the existence of lattice vectors of norm less than $n(m)$, such vectors indeed exist in practice. Moreover, our explicit construction in Section 4 still works and produces bases for lattices L_m, $m < 500$. In the following, the lattices L_m, $200 \leq m < 500$, will be referred to as toy challenges. Explicit parameters for this range can be found in Appendix D. There, we also compute the necessary Hermite factor as in Section 4. The factors suggest that current lattice reduction methods are supposed to find lattice vectors of norm less than $n(m)$. Our experiments with block-type methods confirm this.

All experiments were run on a single core AMD Opteron at 2.6 GHz, using Shoup's NTL [40] in version 5.4.2 and GCC 4.1.2. .

Implementations. For LLL and BKZ, we used the famous implementations integrated in the NTL. We thank Filipović and Koy for making available their implementations of sLLL and PD, which were part of the diploma thesis [11]. We also thank Ludwig for making available and updating his implementation of PSR that was part of his PhD thesis [25]. Finally, we thank Cadé and Stehlé for making available their implementation of fpLLL, which can be obtained from [41].

Figure 2 and Figure 3 depict the performance, i.e. the length of the shortest obtained vector and the logarithmic running time in seconds, for LLL-type and block-type methods, respectively. The boxed line in the left figures shows the norm bound $n(m)$ that has to be undercut. While block-type methods reliably find vectors of norm less than $n(m)$ up to a dimension around 500, the best LLL-type algorithms merely succeed in dimensions < 300.

While being arguably efficient with our choice of parameters, sLLL is unable to find sufficiently short vectors even in dimension 200. For larger dimensions, however, the approximation results of all LLL-type algorithms seem to converge, whereas the running time performance of fpLLL is significantly surpassed by that of the other two.

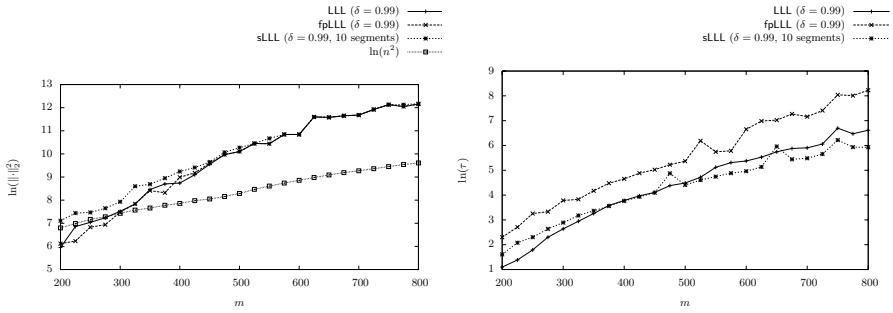

(a) Shortest vectors (b) Run time

Fig. 2. Performance of LLL-type lattice reduction with comparable parameters.

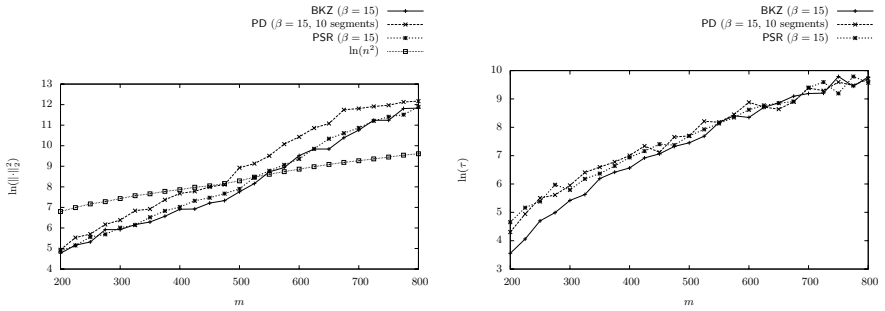

(a) Shortest vectors (b) Run time

Fig. 3. Performance of block-type lattice reduction with comparable parameters.

In Figure 3a, observe that BKZ and PSR perform slightly better than PD, which is mostly due to the internal sLLL step in PD. Accordingly, the graphs seem to converge at the right end, similarly to those in Figure 2a. While the approximation performance of block-type algorithms can be further improved using higher block sizes, this approach is limited by the resulting running time. Extrapolating to higher dimensions, it becomes obvious that finding sufficiently short vectors in L_m requires a significantly larger effort for dimensions that are somewhat higher than 600. This coincides with our observation on the Hermite factor in Section 4.

As for the running time performance of the block-type schemes, observe in Figure 3b that all three behave similarly. In lower dimensions, up to about $m = 450$, BKZ performs strictly better. In higher dimensions, the differences even out and the random character of PSR becomes obvious in its slightly erratic timing.

To conclude, we have reviewed the current state-of-the-art performance of lattice reduction algorithms, using reasonable parameters. We did not, however, explore the limits of the block-type methods. This assessment, we leave to the contestants of the actual lattice challenge that is defined in the next section.

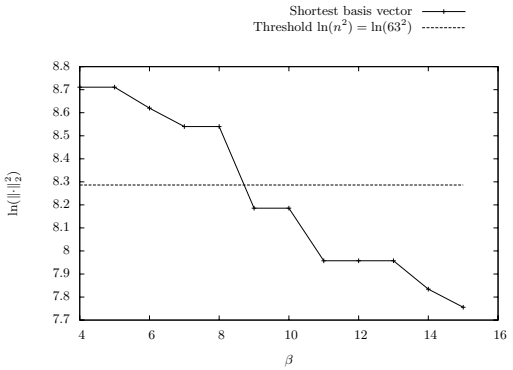

Fig. 4. Shortest vectors found by β-BKZ in dimension $m = 500$.

6 The challenge

In Section 4, we have constructed challenge lattices L_m of dimension m, for $m \geq 500$. The results in Section 3 together with the pseudo-random choice of L_m guarantee the existence of vectors $\mathbf{v} \in L_m$ with $\|\mathbf{v}\|_2 < n(m)$, which are hard to find. For a toy example, refer to Appendix C.

As stated before, we want the lattice challenge to be *open* in the sense that it does not terminate when the *first* short vector is found. Having proven the existence of just one solution might suggest that there are no more, but during practical experiments, we found that many successively shorter vectors exist. For example in Figure 4, we display that in dimension $m = 500$ BKZ with increasing block size subsequently finds smaller and smaller lattice vectors.

We propose the following challenge to all researchers and students.

Lattice Challenge

The contestants are given lattice bases of lattices L_m, together with a norm bound ν. Initially, we set $\nu = n(m)$.

The goal is to find a vector $\mathbf{v} \in L_m$, with $\|\mathbf{v}\|_2 < \nu$.
Each solution $\mathbf{v}$ to the challenge decreases ν to $\|\mathbf{v}\|_2$.
The challenge is hosted at `http://www.latticechallenge.org`.

Acknowledgements

We would like to thank Oded Regev for his helpful remarks and suggestions. Furthermore, we thank the program committee and the anonymous reviewers for their valuable comments.

References

1. D. Aharonov and O. Regev. Lattice problems in $\mathsf{NP} \cap \mathsf{coNP}$. *J. ACM*, 52(5):749–765, 2005.
2. M. Ajtai. Generating hard instances of lattice problems. In *Proceedings of the Annual Symposium on the Theory of Computing (STOC)*, pages 99–108. ACM Press, 1996.
3. M. Ajtai and C. Dwork. A public-key cryptosystem with worst-case/average-case equivalence. In *Proceedings of the Annual Symposium on the Theory of Computing (STOC)*, pages 284–293. ACM Press, 1997.
4. M. Ajtai, R. Kumar, and D. Sivakumar. A sieve algorithm for the shortest lattice vector problem. In *Proceedings of the Annual Symposium on the Theory of Computing (STOC)*, pages 601–610. ACM Press, 2001.
5. D. Bailey and R. Crandall. On the random character of fundamental constant expansions. *Experimental Mathematics*, 10(2):175–190, 2001.
6. D. Bailey and R. Crandall. Random generators and normal numbers. *Experimental Mathematics*, 11(4):527–546, 2002.
7. W. Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 296(4):625–635, 1993.
8. J. Cai and A. Nerurkar. An improved worst-case to average-case connection for lattice problems. In *Proceedings of the Annual Symposium on Foundations of Computer Science (FOCS)*, pages 468–477, 1997.
9. Certicom Corp. The Certicom ECC Challenge. `http://www.certicom.com/index.php/the-certicom-ecc-challenge`.
10. D. Coppersmith and A. Shamir. Lattice attacks on NTRU. In *Advances in Cryptology — Eurocrypt 1997*, pages 52–61, 1997.
11. B. Filipović. Implementierung der gitterbasenreduktion in segmenten. Master's thesis, Johann Wolfgang Göthe-Universität Frankfurt am Main, 2002.
12. N. Gama and P. Q. Nguyen. Predicting lattice reduction. In N. P. Smart, editor, *Advances in Cryptology — Eurocrypt 2008*, volume 4965 of *Lecture Notes in Computer Science*, pages 31–51. Springer-Verlag, 2008.
13. C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In R. E. Ladner and C. Dwork, editors, *STOC*, pages 197–206. ACM, 2008.
14. O. Goldreich and S. Goldwasser. On the limits of nonapproximability of lattice problems. *J. Comput. Syst. Sci.*, 60(3):540–563, 2000.
15. O. Goldreich, S. Goldwasser, and S. Halevi. Public-key cryptosystems from lattice reduction problems. In *Advances in Cryptology — Crypto 1997*, volume 1294 of *Lecture Notes in Computer Science*, pages 112–131. Springer-Verlag, 1997.
16. J. Hoffstein, J. Pipher, and J. H. Silverman. NTRU: A ring-based public key cryptosystem. In J. Buhler, editor, *Algorithmic Number Theory Symposium — ANTS*, volume 1423 of *Lecture Notes in Computer Science*, pages 267–288. Springer-Verlag, 1998.
17. J. Hoffstein, J. H. Silverman, and W. Whyte. Estimated breaking times for NTRU lattices. Technical Report 012, Version 2, NTRU Cryptosystems, 2003. `http://ntru.com/cryptolab/tech_notes.htm`.
18. N. Howgrave-Graham, H. J., J. Pipher, and W. Whyte. On estimating the lattice security of NTRU. Technical Report 104, Cryptology ePrint Archive, 2005. `http://eprint.iacr.org/2005/104/`.

19. D. Kleinbock and B. Weiss. Dirichlet's theorem on diophantine approximation and homogeneous flows. *J.MOD.DYN.*, 4:43, 2008.
20. H. Koy. Primale-duale Segment-Reduktion. `http://www.mi.informatik.uni-frankfurt.de/research/papers.html`, 2004.
21. H. Koy and C.-P. Schnorr. Segment LLL-reduction of lattice bases. In J. H. Silverman, editor, *CaLC*, volume 2146 of *Lecture Notes in Computer Science*, pages 67–80. Springer, 2001.
22. J. C. Lagarias, H. W. L. Jr., and C.-P. Schnorr. Korkin-Zolotarev bases and successive minima of a lattice and its reciprocal lattice. *Combinatorica*, 10(4):333–348, 1990.
23. A. Lenstra, H. Lenstra, and L. Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261(4):515–534, 1982.
24. C. Ludwig. A faster lattice reduction method using quantum search. In *Algorithms and Computation*, volume 2906 of *Lecture Notes in Computer Science*, pages 199–208. Springer-Verlag, 2003.
25. C. Ludwig. *Practical Lattice Basis Sampling Reduction*. PhD thesis, Technische Universität Darmstadt, 2005. `http://elib.tu-darmstadt.de/diss/000640/`.
26. K. S. McCurley. The discrete logarithm problem. In C. Pomerance, editor, *Cryptology and computational number theory*, pages 49–74, Providence, 1990. American Mathematical Society.
27. D. Micciancio. Almost perfect lattices, the covering radius problem, and applications to Ajtai's connection factor. *SIAM Journal on Computing*, 34(1):118–169, 2004.
28. D. Micciancio and O. Regev. Worst-case to average-case reductions based on gaussian measures. *SIAM Journal on Computing*, 37(1):267–302, 2007.
29. P. Q. Nguyen and D. Stehlé. Floating-point LLL revisited. In R. Cramer, editor, *Advances in Cryptology — Eurocrypt 2005*, volume 3494 of *Lecture Notes in Computer Science*, pages 215–233. Springer-Verlag, 2005.
30. P. Q. Nguyen and D. Stehlé. LLL on the average. In F. Hess, S. Pauli, and M. E. Pohst, editors, *ANTS*, volume 4076 of *Lecture Notes in Computer Science*, pages 238–256. Springer-Verlag, 2006.
31. C. Peikert. Limits on the hardness of lattice problems in ℓ_p norms. In *IEEE Conference on Computational Complexity*, pages 333–346. IEEE Computer Society, 2007.
32. O. Regev. Quantum computation and lattice problems. *SIAM J. Comput.*, 33(3):738–760, 2004.
33. O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In *Proceedings of the 37th annual ACM symposium on Theory of computing*, pages 84–93. ACM Press, 2005.
34. O. Regev. On the complexity of lattice problems with polynomial approximation factors, 2007. A survey for the LLL+25 conference.
35. RSA Security Inc. The RSA Challenge Numbers. `http://www.rsa.com/rsalabs/node.asp?id=2093`.
36. W. Schmidt. *Diophantine Approximation*, volume 785 of *Lecture Notes in Mathematics*. Springer-Verlag, 1980.
37. C. Schnorr. A hierarchy of polynomial time lattice basis reduction algorithms. *Theoretical Computer Science*, 53:201–224, 1987.
38. C. Schnorr. Block reduced lattice bases and successive minima. *Combinatorics, Probability and Computing*, 4:1–16, 1994.

39. C. Schnorr. Lattice reduction by random sampling and birthday methods. In *STACS 2003: 20th Annual Symposium on Theoretical Aspects of Computer Science*, volume 2607 of *Lecture Notes in Computer Science*, pages 146–156. Springer-Verlag, 2003.

40. V. Shoup. Number theory library (NTL) for C++. `http://www.shoup.net/ntl/`.

41. D. Stehlé. Damien Stehlé's homepage at école normale supérieure de Lyon. `http://perso.ens-lyon.fr/damien.stehle/english.html`.

A Completing the proof of Theorem 1

With parameters c_1, c_2, n as in the theorem, we want to show that

$$\lfloor c_1 n \ln(n) \rfloor \geq c_1 n \ln(n)/2 \tag{11}$$

holds. By (1), we have that $c_1 \geq 1/(2\ln(2))$. Evaluating both sides of (11) with $n = 1, 2, 3$, we find that the inequality holds for these n. For all $n \geq 4$, consider the following.

We have that $c_1 \geq 1/(2\ln(2)) \geq 2/4\ln(4)$, which implies

$$\begin{aligned}
\lfloor c_1 n \ln(n) \rfloor &\geq c_1 n \ln(n) - 1 \\
&\geq c_1 n \ln(n)/2 + c_1 n \ln(n)/2 - 1 \\
&\geq c_1 n \ln(n)/2 + c_1 4 \ln(4)/2 - 1 \\
&\geq c_1 n \ln(n)/2
\end{aligned}$$

This completes the proof.

B Ratio between m and n

In order to get an idea of the ratio m/n in our challenge lattices, refer to Figure 5. The bend at $m = 500$ reflects our choice of c_1 and c_2 in the toy challenges, where we "cap" the value of c_2 at 2.0.

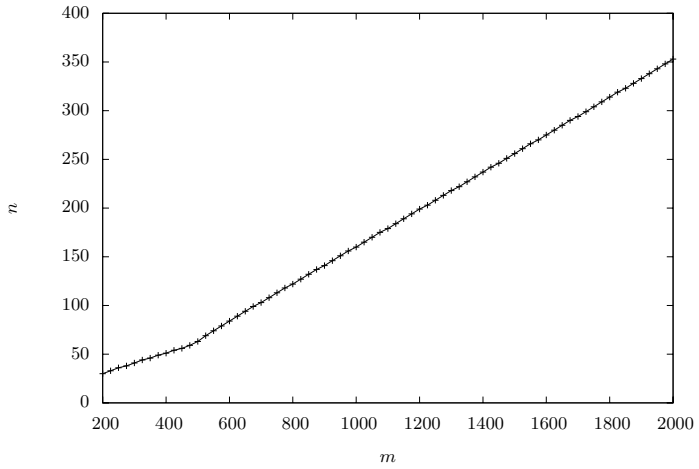

Fig. 5. Ratio between challenge dimension m and reference dimension n.

C Challenge example

The following low-dimensional example gives an idea of what the challenge lattices, and the short vectors in them, essentially look like. Its block structure is similar to the one found by Coppersmith and Shamir for NTRU lattices [10]. This is not surprising because both belong to the class of modular lattices.

Example 1. The transposed challenge basis for $m = 30, n = q = 8$ looks like:

```
[
[1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 -4 -7 -4 -7 -6 -2 -3 -7]
[0 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 -7 -4 -1  0 -6 -7 -1 -5]
[0 0 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 -2 -2 -6 -2 -6 -6 -4 -6]
[0 0 0 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 -6 -7 -1 -5 -5 -1 -4 -3]
[0 0 0 0 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 -7 -4 -2 -3 -1  0 -1 -3]
[0 0 0 0 0 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 -6 -3 -5 -7 -3 -7  0 -2]
[0 0 0 0 0 0 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 -5 -1 -6 -6 -6 -4 -3 -5]
[0 0 0 0 0 0 0 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0  0  0 -1  0 -2 -2 -2 -7]
[0 0 0 0 0 0 0 0 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 -4 -4 -3  0 -5 -7 -6 -4]
[0 0 0 0 0 0 0 0 0 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 -3 -2 -4 -6 -4 -3 -2 -3]
[0 0 0 0 0 0 0 0 0 0 1 0 0 0 0 0 0 0 0 0 0 0 0 0 -7 -6 -4  0  0 -2 -7 -4]
[0 0 0 0 0 0 0 0 0 0 0 1 0 0 0 0 0 0 0 0 0 0 0 0 -4 -1  0  0 -7 -3 -7  0]
[0 0 0 0 0 0 0 0 0 0 0 0 1 0 0 0 0 0 0 0 0 0 0 0 -1 -6 -3  0 -4 -1 -2 -3]
[0 0 0 0 0 0 0 0 0 0 0 0 0 1 0 0 0 0 0 0 0 0 0 0 -3 -1  0 -4 -3 -3 -2  0]
[0 0 0 0 0 0 0 0 0 0 0 0 0 0 1 0 0 0 0 0 0 0 0 0 -6 -6 -2 -2 -1 -3 -6 -6]
[0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 1 0 0 0 0 0 0 0 0  0 -7 -7 -4 -2 -1 -2 -5]
[0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 1 0 0 0 0 0 0 0 -6 -2 -1 -4 -4 -3 -2 -6]
[0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 1 0 0 0 0 0 0 -2 -6 -1 -1 -5 -4 -3 -3]
[0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 1 0 0 0 0 0 -4  0 -5 -4 -6 -7 -5 -2]
```

```
[0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 1 0 0 -4 -3 -3  0 -5 -3 -3 -7]
[0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 1 0 -4  0 -3 -2 -2 -6 -4 -4]
[0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 1 -5 -5 -3  0 -1 -3  0 -6]
[0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0  8  0  0  0  0  0  0  0]
[0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0  0  8  0  0  0  0  0  0]
[0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0  0  0  8  0  0  0  0  0]
[0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0  0  0  0  8  0  0  0  0]
[0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0  0  0  0  0  8  0  0  0]
[0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0  0  0  0  0  0  8  0  0]
[0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0  0  0  0  0  0  0  8  0]
[0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0  0  0  0  0  0  0  0  8]
]
```

The shortest vector in the respective lattice is

```
[0 0 0 0 -1 1 0 0 -1 0 0 0 1 0 0 0 0 0 1 0 0 0 0 -1 0 0 1 0 0 0]
```

and its Euclidean norm is $\sqrt{7} < n = 8$.

D Toy challenges

Table 2 depicts the parameters of our toy challenges.

m	n, q	γ
200	30	1.0146^m
225	33	1.0133^m
250	36	1.0123^m
275	38	1.0115^m
300	41	1.0107^m
325	44	1.0101^m
350	46	1.0095^m
375	49	1.0091^m
400	51	1.0086^m
425	54	1.0082^m
450	56	1.0079^m
475	59	1.0075^m

Table 2. Lattice parameters with the necessary Hermite factor γ.