



ELSEVIER

Contents lists available at ScienceDirect

Journal of Symbolic Computation

www.elsevier.com/locate/jsc



On the probability of generating a lattice

Felix Fontein^a, Pawel Wocjan^{b,1}^a Institute of Mathematics, University of Zurich, Winterthurerstrasse 190, 8057 Zurich, Switzerland^b Mathematics Department and Center for Theoretical Physics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

ARTICLE INFO

Article history:

Received 27 November 2012

Accepted 19 September 2013

Available online 6 December 2013

Keywords:

Random generation of lattice

Riemann zeta function

Infrastructure

Analysis of quantum algorithm

ABSTRACT

We study the problem of determining the probability that m vectors selected uniformly at random from the intersection of the full-rank lattice Λ in \mathbb{R}^n and the window $[0, B]^n$ generate Λ when B is chosen to be appropriately large. This problem plays an important role in the analysis of the success probability of quantum algorithms for solving the Discrete Logarithm Problem in infrastructures obtained from number fields and also for computing fundamental units of number fields.

We provide the first complete and rigorous proof that $2n + 1$ vectors suffice to generate Λ with constant probability (provided that B is chosen to be sufficiently large in terms of n and the covering radius of Λ and the last $n + 1$ vectors are sampled from a slightly larger window). Based on extensive computer simulations, we conjecture that only $n + 1$ vectors sampled from one window suffice to generate Λ with constant success probability. If this conjecture is true, then a significantly better success probability of the above quantum algorithms can be guaranteed.

© 2013 Elsevier B.V. All rights reserved.

1. Introduction

Let G be a finite group. Denote by $p_m(G)$ the probability that m elements drawn uniformly at random from G with replacement generate G . The problem of determining or bounding this probability is of fundamental interest in group theory and has been extensively studied for various families of groups (Acciario, 1996; Pomerance, 2001).

E-mail addresses: felix.fontein@math.uzh.ch (F. Fontein), wocjan@eecs.ucf.edu (P. Wocjan).

¹ On sabbatical leave from Department of Electrical Engineering and Computer Science, University of Central Florida, Orlando, FL 32816, USA.

The purpose of this paper is to study a very natural generalization of this problem from finite abelian groups to finitely generated abelian torsion-free groups. More precisely, we consider the case of lattices, i.e., discrete subgroups of \mathbb{R}^n . The problem is now to determine the probability that m vectors selected uniformly at random with replacement from the intersection of the full-rank lattice Λ in \mathbb{R}^n and a window $X \subset \mathbb{R}^n$ generate Λ . We denote this probability by $p_m(\Lambda, X)$.

Our study of this problem was initially motivated by its relevance to quantum algorithms and quantum cryptanalysis, which we explain in more detail at the end of the paper. But we also believe that this problem is interesting on its own due to its appeal as a very natural and fundamental problem in lattice theory. In fact, it can be viewed as a generalization of the following elementary problem in number theory. For $\Lambda = \mathbb{Z}$ and $X = [1, B]$, the probability $p_m(\Lambda, X)$ corresponds to the probability that m integers chosen uniformly at random from the set $\{1, \dots, B\}$ with replacement are coprime. It is known that $\lim_{B \rightarrow \infty} p_m(\mathbb{Z}, B) = 1/\zeta(m)$ where ζ denotes the Riemann zeta function. For $\Lambda = \mathbb{Z}^n$ and X , the probability $p_m(\Lambda, X)$ is equal to the probability that the $m \times n$ matrix whose column vectors are selected uniformly at random from $\Lambda \cap X$ is unimodular. This problem was studied for special forms of X asymptotically. For $X = [-B, B]^n$, $B \rightarrow \infty$, it was studied by Maze et al. (2011), and for $X = v + [-B, B]^n$, $B \rightarrow \infty$, where the entries of the vector v are bounded polynomially in terms of B , by Elizalde and Woods (2007). In both cases, it was shown that the limit of the probability for $B \rightarrow \infty$ is $\prod_{j=m-n+1}^m \zeta(j)^{-1}$. Both works did not study the problem of bounding the probability in the non-asymptotic case, i.e., in the case where B is fixed.

In this paper, we consider the case where Λ is an arbitrary full-rank lattice and $X = [0, B]^n$ for a sufficiently large but fixed B . Ideally, we want to minimize m , while at the same time ensure that the probability $p_m(\Lambda, [0, B]^n)$ is bounded from below by a nonzero constant. We use $\nu(\Lambda)$ to denote the covering radius of Λ , $\lambda_1(\Lambda)$ the length of a shortest (nonzero) vector of Λ , and $\det(\Lambda)$ the determinant of Λ .

Our two major contributions to the study of this problem are:

Theorem 1.1. *Let Λ be a lattice of full rank in \mathbb{R}^n , and assume that $B \geq 8n^{n/2} \cdot \nu(\Lambda)$ and $B_1 \geq 8n^2(n+1)B$. Assume that n vectors are selected uniformly at random from $\Lambda \cap [0, B]^n$ and $n+1$ vectors uniformly at random from $\Lambda \cap [0, B_1]^n$. If the vectors are sampled independently, then the probability that all these vectors generate Λ is at least*

$$\alpha_n := \left(\prod_{i=2}^{n+1} \zeta(i)^{-1} - \frac{1}{4} \right) \cdot \prod_{k=0}^{n-1} \left(1 - n^{k/2} \frac{(4n^{n/2} + 1)^k}{(4n^{n/2} - 1)^n} \right) \geq 0.092.$$

Unfortunately, our current approach requires $m = 2n + 1$ samples and two windows of different sizes to be able to prove that the probability of generating the lattice Λ is bounded from below by a nonzero constant. However, based on extensive numerical evidence, we formulate the following conjecture, which states that only $m = n + 1$ samples and only one window size suffice to attain a constant probability of generating the lattice.

Conjecture 1.2. *For every $n \in \mathbb{N}$, there exists a constant $0 < c_n < 1$ and a rational function $f_n \in \mathbb{R}(x, y)$ satisfying*

$$\forall x_0 > 0, \forall y_0 \in (0, x_0^{1/n}]: \sup\{f_n(x, y) \mid 0 < x \leq x_0, y_0 \leq y \leq x^{1/n}\} < \infty$$

such that the following holds:

Let Λ be a lattice in \mathbb{R}^n and let $B > f_n(\det \Lambda, \lambda_1(\Lambda))$. Then the probability that $n + 1$ vectors chosen uniformly at random from $\Lambda \cap [0, B]^n$ generate the lattice Λ is at least c_n . Moreover, the constant c_n can be chosen close to $\prod_{k=2}^{n+1} \zeta(k)^{-1}$.

2. Solving the lattice generation problem

We break down the lattice generation problem into two subproblems. First, we consider the probability that n vectors sampled uniformly at random from Λ generate a sublattice Λ_1 of full rank, i.e.

do not lie in a hyperplane. Then, we compute the probability that the residue classes of the next $n + 1$ vectors generate the finite abelian quotient group Λ/Λ_1 . Finally, we combine these two results.

In the following, we assume that $n > 1$. We discuss a result for the case $n = 1$ in Section 3.

The idea to prove a lower bound on the probability by considering the above two steps was proposed by Schmidt (2007). We present a correct proof of the problem arising in the first step, fixing a mistake in Schmidt’s proof. Our approach to analyzing the problem arising in the second step is entirely different from the approach undertaken by Schmidt. The differences will be discussed in Sections 2.1 and 2.3.

2.1. Generating a sublattice of full rank

Note that $\lambda_1, \dots, \lambda_n \in \Lambda \cap [0, B)^n$ generate a sublattice of full rank if and only if they are linearly independent over \mathbb{R} . This is the case if λ_i is not contained in the $(i - 1)$ -dimensional hyperplane spanned by $\lambda_1, \dots, \lambda_{i-1}$. Thus to bound the probability that n uniformly random vectors from $\Lambda \cap [0, B)^n$ generate a full-rank sublattice, we bound the number of lattice elements in the intersection as well as the number of lattice elements lying both in the intersection and a k -dimensional hyperplane, $1 \leq k < n$. We find such bounds using Voronoi cells; see also Section 1.2 of Chapter 8 in Micciancio and Goldwasser (2002). To state the results, we need to introduce some notation, most notably the covering radius of a lattice.

Let Λ be a lattice in \mathbb{R}^n of full rank. For $\lambda \in \Lambda$, let

$$V_\Lambda(\lambda) = \{x \in \mathbb{R}^n \mid \forall \lambda' \in \Lambda \setminus \{\lambda\}: \|x - \lambda\|_2 < \|x - \lambda'\|_2\}$$

be its (open) Voronoi cell. We know that $V_\Lambda(\lambda)$ is contained in an open ball of radius $\nu(\Lambda)$ centered around λ , where $\nu(\Lambda)$ is the covering radius of Λ , and that the volume of $V_\Lambda(\lambda)$ is $\det \Lambda$. Moreover, if $\lambda \neq \lambda'$, $V_\Lambda(\lambda) \cap V_\Lambda(\lambda') = \emptyset$, and $\bigcup_{\lambda \in \Lambda} V_\Lambda(\lambda) = \mathbb{R}^n$. Details can be found in Micciancio and Goldwasser (2002, Chapter 8).

Note that $\nu(\Lambda) \leq \frac{1}{2}n^{n/2+1} \frac{\det \Lambda}{\lambda_1(\Lambda)^{n-1}}$, where $\lambda_1(\Lambda)$ denotes the first successive minimum of Λ (Micciancio and Goldwasser, 2002), i.e. the length of a shortest nonzero vector in Λ .

Lemma 2.1. *If $B > 2\nu(\Lambda)$. Then*

$$\frac{(B - 2\nu(\Lambda))^n}{\det \Lambda} \leq |\Lambda \cap [0, B)^n| \leq \frac{(B + 2\nu(\Lambda))^n}{\det \Lambda}.$$

Lemma 2.2. *Let $B > 0$ and H be a k -dimensional hyperplane, $1 \leq k < n$. Then*

$$|\Lambda \cap H \cap [0, B)^n| \leq \frac{n^{k/2} (B + 2\nu(\Lambda))^k (2\nu(\Lambda))^{n-k}}{\det \Lambda}.$$

The proofs are similar to the one of Proposition 8.7 in Micciancio and Goldwasser (2002):

Proof of Lemma 2.1. If $\lambda \in \Lambda$ satisfies $V_\Lambda(\lambda) \cap [\nu(\Lambda), B - \nu(\Lambda))^n \neq \emptyset$, then we must have $\lambda \in [0, B)^n$. Therefore, $(B - 2\nu(\Lambda))^n / \det \Lambda \leq |\Lambda \cap [0, B)^n|$.

If $\lambda \in \Lambda \cap [0, B)^n$, then we must have $V_\Lambda(\lambda) \subseteq [-\nu(\Lambda), B + \nu(\Lambda))^n$. Therefore, $|\Lambda \cap [0, B)^n| \leq (B + 2\nu(\Lambda))^n / \det \Lambda$. \square

Proof of Lemma 2.2. Let $\lambda \in \Lambda \cap H \cap [0, B)^n$. Then $V_\Lambda(\lambda) \subseteq X := [-\nu(\Lambda), B + \nu(\Lambda))^n \cap (H + B_{\nu(\Lambda)}(0))$, where $B_{\nu(\Lambda)}(0)$ is a ball of radius $\nu(\Lambda)$ centered around 0. Therefore, $|\Lambda \cap H \cap [0, B)^n| \leq \text{vol}(X) / \det \Lambda$, and we have to estimate $\text{vol}(X)$.

Clearly, if $\text{vol}_k(Y)$ denotes the k -dimensional volume of $Y := H \cap [-\nu(\Lambda), B + \nu(\Lambda))^n$, we have that $\text{vol}(X) \leq \text{vol}_k(Y) \cdot (2\nu(\Lambda))^{n-k}$. (In fact, we can replace $(2\nu(\Lambda))^{n-k}$ by the volume of an $(n - k)$ -dimensional ball of radius $\nu(\Lambda)$.)

Let b_1, \dots, b_k be an orthonormal basis of H . Set $T := \{(x_1, \dots, x_k) \in \mathbb{R}^k \mid \sum_{i=1}^k x_i b_i \in [-\nu(\Lambda), B + \nu(\Lambda))^n\}$; then $\text{vol}(T) = \text{vol}_k(Y)$. A point $y \in Y$ corresponds to $(\langle y, b_1 \rangle, \dots, \langle y, b_k \rangle) \in T$. Write $b_i =$

(b_{i1}, \dots, b_{in}) and $y = (y_1, \dots, y_n) \in [-\nu(\Lambda), B + \nu(\Lambda)]^n$, set $A_{ij} := B + \nu(\Lambda)$ if $b_{ij} \geq 0$ and $A_{ij} := \nu(\Lambda)$ if $b_{ij} < 0$. Then

$$\sum_{j=1}^n |b_{ij}|(A_{ij} - (B + 2\nu(\Lambda))) \leq \langle y, b_i \rangle = \sum_{j=1}^n y_j b_{ij} \leq \sum_{j=1}^n |b_{ij}| A_{ij},$$

implying that $\langle y, b_i \rangle$ ranges over an interval of length $\|b_i\|_1(B + 2\nu(\Lambda)) \leq \sqrt{n}(B + 2\nu(\Lambda))$. Therefore,

$$\text{vol}(T) \leq n^{k/2}(B + 2\nu(\Lambda))^k. \quad \square$$

The lemmas allow us to find the following bound on the probability that n random vectors generate a sublattice of full rank:

Corollary 2.3. Assume that $B \geq 8n^{n/2} \cdot \nu(\Lambda)$. Let

$$X := (\Lambda \cap [0, B]^n)^n \quad \text{and} \\ Y := \{(y_1, \dots, y_n) \in X \mid \text{span}_{\mathbb{R}}(y_1, \dots, y_n) = \mathbb{R}^n\}.$$

Then $|Y| \geq \frac{1}{2}|X|$.

Proof. Assume that $y_1, \dots, y_k \in X$ are linearly independent, $0 \leq k < n$. We bound the probability from above that $y_{k+1} \in X$ is not contained in the hyperplane generated by y_1, \dots, y_k , which is of dimension k . Write $B = j \cdot \nu(\Lambda)$ with $j \geq 8n^{n/2}$. By Lemmas 2.1 and 2.2, the probability that y_{k+1} is in a k -dimensional hyperplane is bounded from above by

$$P_k := \frac{n^{k/2}(B + 2\nu(\Lambda))^k(2\nu(\Lambda))^{n-k}}{\det \Lambda} \cdot \frac{\det \Lambda}{(B - 2\nu(\Lambda))^n} = n^{k/2} \frac{(j + 2)^k 2^{n-k}}{(j - 2)^n}.$$

The success probability is bounded from below by $\prod_{k=0}^{n-1} (1 - P_k)$. Using induction on n , we can prove that

$$\prod_{k=0}^{n-1} (1 - P_k) \geq 1 - \sum_{k=0}^{n-1} P_k.$$

The sum $\sum_{k=0}^{n-1} P_k$ can be bounded from above as follows:

$$\begin{aligned} \sum_{k=0}^{n-1} P_k &= \frac{2^n}{(j - 2)^n} \sum_{k=0}^{n-1} \left(\frac{\sqrt{n}(j + 2)}{2} \right)^k \\ &= \frac{2^n}{(j - 2)^n} \left[\left(\frac{\sqrt{n}(j + 2)}{2} \right)^n - 1 \right] \left[\left(\frac{\sqrt{n}(j + 2)}{2} \right) - 1 \right]^{-1} \\ &< \frac{2^n}{(j - 2)^n} \left(\frac{\sqrt{n}(j + 2)}{2} \right)^n \left[\left(\frac{\sqrt{n}(j + 2)}{2} \right) - 1 \right]^{-1} \\ &= n^{n/2} \left(1 + \frac{4}{j - 2} \right)^n \left[\left(\frac{\sqrt{n}(j + 2)}{2} \right) - 1 \right]^{-1}. \end{aligned}$$

Now $(1 + \frac{4}{j-2})^n \leq \exp(\frac{4n}{j-2}) \leq \exp(\frac{4}{8n^{n/2-1}-2/n}) \leq 2$ for all $n \geq 1$ and $\sqrt{n}(j + 2)/2 - 1 \geq j/2$, whence

$$\sum_{k=0}^{n-1} P_k < 2n^{n/2} \cdot \frac{2}{j} \leq \frac{4n^{n/2}}{8n^{n/2}} = \frac{1}{2}. \quad \square$$

Note that our lower bound is far from optimal. If one considers the value $\sum_{k=0}^{n-1} P_k$ from the proof and substitutes j by $8n^{n/2}$, one obtains the lower bound

$$\prod_{k=0}^{n-1} \left(1 - n^{k/2} \frac{(4n^{n/2} + 1)^k}{(4n^{n/2} - 1)^n} \right).$$

For $n = 1$ this is $\frac{2}{3}$, and the product grows to 1 for $n \rightarrow \infty$. For small n , the values are:

Dimension n	1	2	3	4	5	6	7
Lower bound	0.666	0.725	0.812	0.859	0.883	0.896	0.905

Remark 2.4. The basic idea of the proof of this corollary is similar to the proof of the first part of Satz 2.4.23 in Schmidt (2007). Note that the proof in Schmidt (2007) is not correct: the ratio $|M_{i-1} \cap \mathcal{B}|/|M_i \cap \mathcal{B}|$ considered in the proof can be $> \frac{1}{2}$; for example, consider $r = 3$, $M = \mathbb{Z}^3$, $n > 0$ arbitrary (in Schmidt (2007), $n\nu(M)$ is what we denote by b , i.e., $\mathcal{B} = [0, n\nu(M))^n$), $x_1 = (1, n\nu(M) - 1, 0)$, $x_2 = (0, 1, n\nu(M) - 1)$, $x_3 = (0, 0, 1)$; then $M_1 \cap \mathcal{B}$ contains two elements, while $M_2 \cap \mathcal{B}$ contains three elements. Therefore, $|M_1 \cap \mathcal{B}|/|M_2 \cap \mathcal{B}| = \frac{2}{3} > \frac{1}{2}$. The problem is that $\det M_i$ cannot be bounded linearly in terms of $\nu(M)$ and $\det M_{i-1}$, as it was claimed in that proof; in this example, $\det M_1 = \sqrt{1 + (n-1)^2}$, $\det M_2 = \sqrt{1 + (n-1)^2 + (n-1)^4}$ and $\nu(M) = 1$. In our proof, we proceed differently by considering the ratio $|M_i \cap \mathcal{B}|/|M \cap \mathcal{B}|$ directly, and both our bound on the probability and our bound on the minimal size of \mathcal{B} is in fact better than the corresponding bounds given in Schmidt (2007).

2.2. Generating a finite abelian group

In case Λ_1 is a sublattice of full rank of Λ , the quotient group $G = \Lambda/\Lambda_1$ is a finite abelian group. Its order equals the index $[\Lambda : \Lambda_1]$, and by the Elementary Divisor Theorem, it can be generated by n elements.

Proposition 2.5. *Let G be a finite abelian group known to be generated by n elements. Then the probability that $n + 1$ elements drawn uniformly at random from G generate G is at least*

$$\hat{\zeta} := \prod_{i=2}^{\infty} \zeta(i)^{-1} \geq 0.434,$$

where ζ denotes the Riemann zeta function.

For the decimal expansion of $\hat{\zeta}$, see Integer sequence A021002 (no date). The probability that a finite group is generated by a certain number of random elements has been studied extensively. Formulas for the probability for p -groups and products of finite groups of coprime orders have been derived by Acciario (1996, Lemma 4 and Corollary 3) (see also Pomerance, 2001). Our result is essentially a corollary of these two results, which we have not found in this form in the literature.

Proof of Proposition 2.5. For a finite group H , let $\lambda_t(H)$ be the probability that t group elements chosen uniformly at random generate H . In Acciario (1996), it is shown that if H is a p -group with minimal number d of generators, then

$$\lambda_t(H) = \prod_{i=1}^d (1 - p^{-i}) \cdot \prod_{i=d+1}^t \frac{p^{i-d} - p^{-d}}{p^{i-d} - 1} = \prod_{i=t-d+1}^t (1 - p^{-i})$$

for $t \geq d$ (Lemma 4), and that if $H = H_1 \times H_2$ with $|H_1|, |H_2|$ being coprime, then $\lambda_t(H) = \lambda_t(H_1)\lambda_t(H_2)$ (Corollary 3).

Let p_1, \dots, p_k be the distinct prime divisors of $|G|$, and let G_i be the p_i -Sylow subgroup of G . Then $G = G_1 \oplus \dots \oplus G_k$. Now Acciario (1996, Corollary 3) yields $\lambda_t(G) = \prod_{i=1}^k \lambda_t(G_i)$ since $|G_i|$ is a p_i -group, and $p_i \neq p_j$ for $i \neq j$. Let d_i be the minimal number of generators for G_i ; since the minimal number for G is n , we must have $d_i \leq n$. Thus, by Acciario (1996, Lemma 4)

$$\lambda_{n+1}(G_i) = \prod_{i=n+1-d_i}^{n+1} (1 - p^{-i}) \geq \prod_{i=2}^{n+1} (1 - p^{-i}).$$

Therefore, the probability that n elements of an arbitrary finite abelian group G which can be generated by n elements generate the group is at least

$$\prod_p \prod_{i=2}^{n+1} (1 - p^{-i}) = \prod_{i=2}^{n+1} \prod_p (1 - p^{-i}) = \left(\prod_{i=2}^{n+1} \zeta(i) \right)^{-1}$$

using the Euler product representation of the Riemann zeta function. Now

$$\prod_{i=2}^{n+1} \zeta(i) \leq \prod_{i=2}^{\infty} \zeta(i) = \hat{\zeta}^{-1}. \quad \square$$

Observe that our approach only works if we have at least $n + 1$ elements. If we chose just n elements randomly, the final product would include $\zeta(1)^{-1} = 0$ and the probability would drop down to zero. However, a different approach can result in a nonzero probability for n elements. This probability will necessarily not be constant anymore, but has to depend on n or $|G|$. For example, if p_1, \dots, p_k are distinct primes and $G = \prod_{i=1}^k \mathbb{F}_{p_i}^n \cong (\mathbb{Z}/(p_1 \cdots p_k)\mathbb{Z})^n$, then G can be generated by n elements, but the probability that n random elements from G generates G is exactly $\prod_{i=1}^k \prod_{j=1}^n (1 - p_i^{-j})$, which goes to zero for $k \rightarrow \infty$ for exactly the above reasons. Hence, any non-trivial bound on the probability must take n or p_1, \dots, p_k into account.

This shows that our approach will not work with fewer than $2n + 1$ elements, if the desired bound on the probability should be independent of n .

2.3. The final result

Assume that the first n sampled vectors from $\Lambda \cap [0, B]^n$ generate a sublattice Λ_1 of full rank. Then $G = \Lambda/\Lambda_1$ is a finite abelian group which can be generated by n elements. Thus if we sample $n + 1$ elements $\lambda + \Lambda_1$ from G in a uniform random manner, we can bound the probability that they generate G . In case $G = \langle \lambda_{n+1} + \Lambda_1, \dots, \lambda_{2n+1} + \Lambda_1 \rangle$ and $\Lambda_1 = \langle \lambda_1, \dots, \lambda_n \rangle$, we have $\Lambda = \langle \lambda_1, \dots, \lambda_n, \lambda_{n+1}, \dots, \lambda_{2n+1} \rangle$.

The main problem is that we cannot directly sample uniformly at random from G : if we choose $\lambda \in \Lambda \cap [0, B]^n$ uniformly at random, then $\lambda + \Lambda_1$ will in general not be uniformly distributed in $G = \Lambda/\Lambda_1$. By enlarging the window $[0, B]^n$ to $[0, B_1]^n$ with $B_1 > B$ large enough, we ensure that the residue classes of the samples $\lambda \in \Lambda \cap [0, B_1]^n$ are essentially distributed uniformly at random in G . More precisely, we can show that the statistical distance between the distribution and the perfectly uniform distribution is small enough. This is established by the following result:

Lemma 2.6. *Let Λ_1 be an arbitrary full-rank sublattice of Λ . Assume that $B_1 > 2\nu(\Lambda_1)$ and we can sample uniformly at random from $\Lambda \cap [0, B_1]^n$. Denote the sample by λ . Then, the total variation distance between the uniform distribution over Λ/Λ_1 and the distribution of $\lambda + \Lambda_1$, where $\lambda \in \Lambda \cap [0, B_1]^n$ is uniformly distributed, is at most*

$$1 - \frac{(B_1 - 2\nu(\Lambda_1))^n}{(B_1 + 2\nu(\Lambda))^n}.$$

Proof. First note that $V_{\Lambda_1}(\lambda_1) = \lambda_1 + V_{\Lambda_1}(0)$ and $\overline{V_{\Lambda_1}(\lambda_1)} = \lambda_1 + \overline{V_{\Lambda_1}(0)}$. Now, as $\bigcup_{\lambda_1 \in \Lambda_1} (\lambda_1 + \overline{V_{\Lambda_1}(0)}) = \mathbb{R}^n$ and two translates of $V_{\Lambda_1}(0)$ by different elements of Λ_1 do not intersect, there exists a set V with $V_{\Lambda_1}(0) \subseteq V \subseteq \overline{V_{\Lambda_1}(0)}$ satisfying

$$\bigcup_{\lambda_1 \in \Lambda_1} (\lambda_1 + V) = \mathbb{R}^n \quad \text{and} \quad \forall \lambda_1 \in \Lambda_1 \setminus \{0\}: (\lambda_1 + V) \cap V = \emptyset.$$

Note that $\text{vol}(V) = \text{vol}(V_{\Lambda_1}(0)) = \det \Lambda_1$.

We first assume that the window has the form $[0, B_1]^n$ instead of the form $[0, B_1]^n$ and later argue that the bounds derived also apply to the actual window $[0, B_1]^n$. We need the following three facts:

- There are exactly $m = \det \Lambda_1 / \det \Lambda$ points of Λ in each translate of V , i.e.,

$$|(\lambda_1 + V) \cap \Lambda| = m \quad \text{for all } \lambda_1 \in \Lambda_1.$$

This can be shown by using asymptotic arguments similarly to those used in the proof that each translate of the elementary parallelepipeds of Λ_1 contains exactly m elements of Λ (see e.g. [Barvinok, no date](#)).

- There are at least

$$\ell = \frac{(B_1 - 2\nu(\Lambda_1))^n}{\det \Lambda_1}$$

translates of V that are entirely contained inside the window $[0, B_1]^n$ since $V \subseteq \overline{B_{\nu(\Lambda_1)}(0)}$.

- There are at most

$$u = \frac{(B_1 + 2\nu(\Lambda))^n}{\det \Lambda}$$

points of Λ inside $[0, B_1]^n$.

Let $\Omega = \Lambda \cap [0, B_1]^n$. We call $\lambda \in \Omega$ *good* if there exists $\lambda_1 \in \Lambda_1$ such that

$$\lambda \in \lambda_1 + V \subseteq [0, B_1]^n.$$

In words, $\lambda \in \Omega$ is good if it belongs to a translate of V that is entirely inside the window $[0, B_1]^n$. Let Ω_{good} denote the set of good points. Using the first two facts above, we deduce that $|\Omega_{\text{good}}| \geq m\ell$.

Let \mathcal{P} denote the uniform distribution on Ω and $\tilde{\mathcal{P}}$ the uniform distribution on the set of good points. We view $\tilde{\mathcal{P}}$ as a probability distribution on Ω by assigning the probability 0 to any point that is not good. Then, the total variation distance between \mathcal{P} and $\tilde{\mathcal{P}}$ is bounded from above by

$$\begin{aligned} \frac{1}{2} \sum_{\lambda \in \Omega} |\mathcal{P}(\lambda) - \tilde{\mathcal{P}}(\lambda)| &= \frac{1}{2} |\Omega_{\text{good}}| \left(\frac{1}{|\Omega_{\text{good}}|} - \frac{1}{|\Omega|} \right) + \frac{1}{2} (|\Omega| - |\Omega_{\text{good}}|) \frac{1}{|\Omega|} \\ &= 1 - \frac{|\Omega_{\text{good}}|}{|\Omega|} \leq 1 - \frac{m\ell}{u} \leq 1 - \frac{(B_1 - 2\nu(\Lambda_1))^n}{(B_1 + 2\nu(\Lambda))^n}. \end{aligned}$$

Let $\kappa : \Lambda \rightarrow \Lambda/\Lambda_1$ denote the canonical projection map. Let \mathcal{Q} and $\tilde{\mathcal{Q}}$ be the probability distributions on the cosets Λ/Λ_1 induced by the following two-step process: (1) sample λ according to \mathcal{P} and $\tilde{\mathcal{P}}$, respectively, and (2) apply κ to the obtained sample λ . Observe that $\tilde{\mathcal{Q}}$ is the uniform distribution on Λ/Λ_1 . Unfortunately, we cannot sample according to $\tilde{\mathcal{Q}}$ but only according to \mathcal{Q} . However, the total variation distance between \mathcal{Q} and $\tilde{\mathcal{Q}}$ must be less or equal to the one between \mathcal{P} and $\tilde{\mathcal{P}}$ since the total variation distance satisfies a so-called data processing inequality.

Note that so far, we have considered $[0, B_1]^n$ instead of $[0, B_1]^n$. As Λ is discrete, there exists some $2\nu(\Lambda_1) < B'_1 < B_1$ with $[0, B'_1]^n \cap \Lambda = [0, B_1]^n$. Applying the result above to $[0, B'_1]^n$ and then using that

$$x \mapsto 1 - \frac{(x - 2\nu(\Lambda_1))^n}{(x + 2\nu(\Lambda_1))^n}$$

is increasing yields the stated claim for $[0, B_1]^n$. \square

Combining the lemma and Proposition 2.5 and using the additivity of the total variation distance under composition provided that the components are independent, we obtain the following result:

Corollary 2.7. Assume that $B \geq 8n^{n/2} \cdot \nu(\Lambda)$ and $B_1 \geq 8n^2(n + 1)B$. Let Y be as in Corollary 2.3 and $(y_1, \dots, y_n) \in Y$. Let

$$X_1 := (\Lambda \cap [0, B_1]^n)^{n+1},$$

$$Z = \{(z_1, \dots, z_{n+1}) \in X_1^{n+1} \mid \text{span}_{\mathbb{Z}}\{y_1, \dots, y_n, z_1, \dots, z_{n+1}\} = \Lambda\}.$$

Then $|Z| \geq (\hat{\zeta} - \frac{1}{4})|X_1| \geq 0.184|X_1|$.

Proof. Let Λ_1 be the full-rank sublattice generated by y_1, \dots, y_n . We have the following simple bound on the covering radius

$$\nu(\Lambda_1) \leq \frac{\sqrt{n}}{2} \lambda_n(\Lambda_1) \leq \frac{\sqrt{n}}{2} \max_{i=1, \dots, n} \|y_i\|_2 \leq \frac{\sqrt{n}}{2} \sqrt{n}B = \frac{nB}{2}$$

since the y_i are linearly independent and every vector in $[0, B]^n$ is shorter than $\sqrt{n}B$. Moreover, $\nu(\Lambda_1) \geq \nu(\Lambda)$.

Let z_i be uniformly distributed in $\Lambda \cap [0, B_1]^n$. Then, Lemma 2.6 implies that $z_i + \Lambda_1$ (for $i = n + 1, \dots, 2n + 1$) are distributed almost uniformly at random from Λ/Λ_1 . The total variation distance from the uniform distribution is bounded from above as follows:

$$1 - \frac{(B_1 - 2\nu(\Lambda_1))^n}{(B_1 + 2\nu(\Lambda_1))^n} \leq 1 - \frac{(B_1 - 2\nu(\Lambda_1))^n}{(B_1 + 2\nu(\Lambda_1))^n} = 1 - \left(1 - \frac{4\nu(\Lambda_1)}{B_1 + 2\nu(\Lambda_1)}\right)^n$$

$$\leq 1 - \left(1 - n \frac{4\nu(\Lambda_1)}{B_1 + 2\nu(\Lambda_1)}\right) \leq \frac{4n\nu(\Lambda_1)}{B_1} \leq \frac{2n^2B}{B_1} \leq \frac{1}{4(n + 1)}.$$

Consider now the uniform probability distribution on the $(n + 1)$ -fold direct product of Λ/Λ_1 and the probability distribution that arises from sampling almost uniformly at random on each of the components as above. Then the total variation between these two distributions is bounded from above by $(n + 1) \cdot \frac{1}{4(n + 1)} = \frac{1}{4}$. This is because the total variation distance is subadditive under composition provided that the components are independent (see e.g. Micciancio and Goldwasser (2002, Subsection 1.3 “Statistical distance” in Chapter 8) for more information on the total variation distance).

Clearly, the abelian group Λ/Λ_1 can be generated with only n generators. Hence, Proposition 2.5 implies that $n + 1$ samples (provided that they are distributed uniformly at random over the group) form a generating set with probability greater or equal to $\hat{\zeta}$. Due to the deviation from the uniform distribution on the $(n + 1)$ -fold direct product of Λ/Λ_1 this probability may decrease. However it is at least $\hat{\zeta} - 1/4$ since the total variation distance is at most $1/4$. The claim follows now by translating the lower bound on the probability to a lower bound on the fraction of elements with the desired property. \square

Combining this corollary with Corollary 2.3, we obtain a proof of Theorem 1.1. This theorem is similar to Satz 2.4.23 in Schmidt (2007). We emphasize that our bound on the success probability is constant, whereas the bound presented in Satz 2.4.23 decreases exponentially fast with the dimension n . The first part of the proof of Satz 2.4.23 (concerning the generation of a full-rank sublattice) is unfortunately not correct, but can be corrected as we have shown in our proof of Corollary 2.3. The idea behind the second part is completely different from our proof and cannot be used to prove a constant success probability. Perhaps it could be used to prove that only $2n$ random elements (as opposed to $2n + 1$ elements) are needed to guarantee a nonzero success probability.

Note that for a fixed dimension n , one obtains bounds larger than 0.092. For $n = 2, 3, 4$ and 5 , α_n is larger than 0.238, 0.185, 0.176, 0.172 and 0.170, respectively.

3. Conjecture

Let b_1, \dots, b_n be any basis of the lattice Λ . Consider the natural isomorphism $\Phi : \mathbb{R}^n \rightarrow \mathbb{R}^n$ mapping the i -th standard unit vector e_i to b_i . Then $\Phi(\mathbb{Z}^n) = \Lambda$. Let

$$X := \Phi^{-1}([0, B]^n) = \left\{ (a_1, \dots, a_n) \in \mathbb{R}^n \mid \sum_{i=1}^n a_i b_i \in [0, B]^n \right\};$$

this is a parallelepiped in \mathbb{R}^n of volume $\frac{B^n}{\det \Lambda}$ having 0 as a vertex. If we assume that the basis b_1, \dots, b_n is reduced, then this parallelepiped is not too skewed.

Now let $v_1, \dots, v_m \in \Lambda$ be vectors, $m \geq n$, and consider $\hat{v}_i := \Phi^{-1}(v_i) \in \mathbb{Z}^n$ for $i = 1, \dots, m$. We have that $\langle v_1, \dots, v_m \rangle = \Lambda$ if and only if $\langle \hat{v}_1, \dots, \hat{v}_m \rangle = \mathbb{Z}^n$, and this is the case if and only if the matrix $(\hat{v}_1, \dots, \hat{v}_m) \in \mathbb{Z}^{n \times m}$ is unimodular.

Therefore, the probability that $m \geq n$ vectors selected uniformly at random in $\Lambda \cap [0, B]^n$ generate Λ equals the probability that an $n \times m$ integer matrix whose columns are chosen uniformly at random in X is unimodular.

As indicated in the introduction, this problem was studied for special forms of X asymptotically. For $X = [-B, B]^n$, $B \rightarrow \infty$, it was studied in Maze et al. (2011), and for $X = v + [-B, B]^n$, $B \rightarrow \infty$ while v is bounded polynomially in terms of B , in Elizalde and Woods (2007). In both cases, it was shown that the limit of the probability for $B \rightarrow \infty$ is $\prod_{j=m-n+1}^m \zeta(j)^{-1}$ – which for $m = n + 1$, not very surprisingly, equals the probability given in Section 2.2. This can be bounded from below by $\hat{\zeta} > 0.434$ as soon as $m > n$. This implies that for a certain $\hat{B} > 0$, we have that the probability is at least 0.434 for all $B > \hat{B}$.

While it seems probable that the proof of Elizalde and Woods (2007) can yield effective non-trivial bounds for any such \hat{B} . However, it is unclear whether this would help for the general case, as the proof only considers the special case $X = v + [-B, B]^n$, while we have to consider essentially arbitrary parallelepipeds with 0 as a vertex.

We have run computer experiments to study the probability for arbitrary parallelepipeds. We restricted to the case $m = n + 1$. For the experiments, we generated a random parallelepiped by choosing n vectors from $[-C, C]^n$ and considering the parallelepiped spanned by them. We generated 1000 such parallelepipeds, and for every parallelepiped we generated 10000 integer matrices with columns taken uniformly at random from the parallelepiped. Every matrix was tested whether it is unimodular. We used three different bounds for C , namely $C = 10^4$, $C = 10^9$ and $C = 10^{18}$. For every combination of $n \times m = n \times (n + 1)$ and C , we computed both the average probability that an $n \times m$ integer matrix taken from a parallelepiped is unimodular, and the minimal probability (over all parallelepipeds for given $n \times m$ and C). The results are shown in Tables 1 (average probabilities) and 2 (minimal probabilities) on p. 12. They also include the “ideal” probabilities $\prod_{j=2}^{n+1} \zeta(j)^{-1}$ predicted for the special parallelepiped with $B \rightarrow \infty$ in Maze et al. (2011).

As one can clearly see, the average values are very close to the ideal ones. But also the minimal probabilities observed in the experiments were always close to the ideal values. In fact, the difference between minimal and maximal probabilities never exceeded 3.66%. If one compares these probabilities to the ones given at the end of Section 2.3, one sees that the probabilities obtained there are far too low.

Our conjecture is based on the evidence sketched above. The conditions on f ensure that given a family of lattices where we have an upper bound on $\det \Lambda$ and a lower bound on $\lambda_1(\Lambda)$, we can find a lower bound on B such that the result holds for all lattices of this family. This is for example the case for unit lattices of number fields. There, one has a lower bound on $\lambda_1(\Lambda)$ depending only on the degree of the number field (Remak, 1932), and an upper bound on $\det \Lambda$ in terms of the degree and discriminant of the number field (Sands, 1991).

Table 1

Average empirical probability that a random $n \times (n + 1)$ integer matrix from a random parallelepiped inside $[-C, C]^n$ is unimodular.

n	$C = 10^4$	$C = 10^9$	$C = 10^{18}$	Ideal probability
1	60.7273%	60.8094%	60.8103%	60.7927%
2	50.5849%	50.5899%	50.5649%	50.5739%
3	46.7040%	46.7257%	46.7367%	46.7272%
4	45.0382%	45.0252%	45.0080%	45.0631%
5	44.2531%	44.2315%	44.2052%	44.2949%
6	43.8661%	43.8894%	43.8740%	43.9281%
7	43.6945%	43.6773%	43.7059%	43.7497%
8	43.6003%	43.6162%	43.6049%	43.6620%
9	43.5529%	43.5662%	43.5447%	43.6187%
10	43.5369%	43.5343%	43.5332%	43.5971%
11	43.5124%	43.5463%	43.5556%	43.5864%
12	43.5314%	43.5488%	43.5218%	43.5810%
13	43.5329%	43.5314%	43.5224%	43.5784%
14	43.5217%	43.5322%	43.5679%	43.5770%
15	43.5113%	43.5273%	43.4947%	43.5764%

Table 2

Minimal empirical probability that a random $n \times (n + 1)$ integer matrix from a random parallelepiped inside $[-C, C]^n$ is unimodular.

n	$C = 10^4$	$C = 10^9$	$C = 10^{18}$	Ideal probability
1	58.98%	59.17%	59.31%	60.7927%
2	49.03%	48.91%	49.17%	50.5739%
3	45.16%	44.96%	45.34%	46.7272%
4	43.09%	43.31%	43.60%	45.0631%
5	42.39%	42.61%	42.61%	44.2949%
6	42.27%	42.06%	42.06%	43.9281%
7	42.24%	42.37%	41.72%	43.7497%
8	41.99%	42.17%	41.83%	43.6620%
9	42.18%	42.14%	41.78%	43.6187%
10	42.14%	42.02%	42.14%	43.5971%
11	41.94%	41.97%	42.09%	43.5864%
12	41.86%	41.81%	42.09%	43.5810%
13	41.98%	42.12%	42.05%	43.5784%
14	41.65%	42.10%	42.06%	43.5770%
15	41.99%	42.00%	42.13%	43.5764%

The only case in which we know how to prove the conjecture is $n = 1$. In that case, we have $\Lambda = v\mathbb{Z}$ for some real number $v > 0$. Given two elements $av, bv \in \Lambda \cap [0, B)$, we have that $\langle av, bv \rangle = v\mathbb{Z}$ if and only if a and b are coprime. Therefore, we are interested in the probability that two random integers in $[0, \frac{B}{\det \Lambda})$ are coprime. For $\frac{B}{\det \Lambda} \rightarrow \infty$, it is well-known that this probability goes to $\zeta(2)^{-1} = \frac{6}{\pi^2} \approx 0.607927$. One can easily make this more precise, for example by using the computations from [Lehmer \(1900\)](#) and additional computer computations for $n \leq 1000$:

Proposition 3.1. *Let $n \geq 1$ be a natural number and*

$$p_n = \frac{|\{(x, y) \in \mathbb{N}^2 \mid 0 \leq x, y \leq n, \gcd(x, y) = 1\}|}{(n + 1)^2}.$$

Then

$$p_n \geq \frac{13}{22} > 0.5909$$

with equality in the first inequality if and only if $n = 10$.

Proof. For $n \geq 1$, let

$$A(n) := \left| \{ (x, y) \in \mathbb{N}^2 \mid 0 \leq x, y \leq n, \gcd(x, y) = 1 \} \right|.$$

Clearly, $p_n = \frac{A(n)}{(n+1)^2}$ and $A(n) = 2 \sum_{k=1}^n \phi(k) + 1$, where

$$\phi(k) = \left| \{ x \in \mathbb{N} \mid 0 \leq x < k, \gcd(x, k) = 1 \} \right|$$

is Euler's totient function. Now in [Lehmer \(1900, Theorem IV and proof\)](#), it is proven that

$$\sum_{k=1}^n \phi(k) = \frac{n^2}{2} \cdot \frac{1}{\zeta(2)} + \Delta(n), \quad \text{where } |\Delta(n)| \leq n \sum_{k=1}^n \frac{1}{k} + \frac{n^2}{2} \cdot \frac{1}{n}$$

and ζ is the Riemann ζ function. Now $\sum_{k=1}^n \frac{1}{k} \leq 1 + \int_1^n \frac{1}{x} dx = 1 + \log n$, whence

$$|\Delta(n)| \leq n(1 + \log n) + \frac{1}{2}n = \frac{3}{2}n + n \log n.$$

This together with $\zeta(2) = \frac{\pi^2}{6}$ shows that

$$\begin{aligned} p_n &= \frac{1 + 2 \sum_{k=1}^n \phi(k)}{(n+1)^2} \geq \frac{1 + 2(\frac{3}{\pi^2}n^2 - \frac{3}{2}n - n \log n)}{(n+1)^2} \\ &= \frac{6}{\pi^2} \cdot \frac{n^2}{(n+1)^2} - \frac{n \log n}{(n+1)^2} - \frac{3n}{2(n+1)^2} + \frac{1}{(n+1)^2}. \end{aligned}$$

Using a computer program, one quickly verifies that $p_n \geq \frac{13}{22}$ for all $n \in \mathbb{Z} \cap [1, 1000]$, with equality if and only if $n = 10$. For $n > n_0 := 1000$, the above inequality yields

$$p_n > \frac{6}{\pi^2} \cdot \frac{n_0^2}{(n_0+1)^2} - \frac{n_0 \log n_0}{(n_0+1)^2} - \frac{3n_0}{2(n_0+1)^2} > \frac{13}{22}. \quad \square$$

Therefore, the conjecture is true for $n = 1$ with $c_1 = \frac{13}{22}$ and $f_1(x, y) = x$.

Finally, note that in case $n = m$, the result in [Maze et al. \(2011\)](#) shows that one expects that the only lower bound one can give is 0. We have run a few experiments here as well, and already for $C = 10^4$, not a single unimodular matrix was found during the experiments.

4. Relevance of lattice generation to quantum algorithms and quantum cryptanalysis

The *Discrete Logarithm Problem* (DLP) is a mathematical primitive on which many public-key cryptosystems are based. Examples of groups, in which the DLP is considered to be computationally hard, include the multiplicative group of \mathbb{F}_q ([Menezes et al., 1997](#)), the group of \mathbb{F}_q -rational points of an elliptic curve ([Cohen et al., 2006](#)), the divisor class and ideal class groups of an algebraic curve, or the infrastructure of an algebraic number field ([Buchmann, 1991](#); [Scheidler et al., 1994](#)). For the cryptographically relevant instances, the best know classical algorithms have a subexponential running time. In contrast, there are efficient quantum algorithms that solve these DLPs in polynomial time ([Shor, 1997](#); [Cheung and Mosca, 2001](#); [Hallgren, 2002](#); [Sarvepalli and Wocjan, 2013](#); [Hallgren, 2005](#); [Schmidt and Vollmer, 2005](#); [Schmidt, 2007](#)).

The statement on the running time of the quantum algorithm for solving the DLP in the infrastructure of a number field needs to be made more precise. It scales polynomially in the logarithm of the discriminant $\Delta_{K/\mathbb{Q}}$ of the number field K , but exponentially in a polynomial expression $q([K : \mathbb{Q}])$ of its degree $[K : \mathbb{Q}]$. While the exponential scaling seems to be unavoidable for fundamental reasons – one has to compute shortest lattice basis vectors in dimension $[K : \mathbb{Q}]$ to be able to perform basic arithmetic operations in the infrastructure – it is important to reduce the magnitude of the polynomial q . We now explain why our theorem and conjecture on lattice generation can be used to achieve such reduction.

The quantum algorithm solves the DLP problem by reducing it to the problem of finding a basis of a certain full-rank lattice $L \subset \mathbb{R}^n$ where $n = d + 1$. We explain this reduction in more detail at the end of this section. To find a basis of L , the quantum algorithm has a mechanism which, with a certain probability $p_1 > 0$, outputs an essentially uniformly distributed vector $\lambda \in \Lambda \cap [0, B)^{n+1}$, where $\Lambda = L^*$ is the dual lattice of L and $B > 0$ is sufficiently large. With probability $1 - p_1$ it outputs a vector that is not an element of Λ . Unfortunately, this unfavorable case cannot be recognized efficiently. If one has $\lambda_1, \dots, \lambda_m$ with $\Lambda = \langle \lambda_1, \dots, \lambda_m \rangle_{\mathbb{Z}}$, one can compute a basis of Λ from these vectors and then use linear algebra (matrix inversion) to retrieve a basis of $L = \Lambda^*$ itself.

To compute the success probability, one has to consider the probability that the m sampled vectors are actually in Λ , and the probability that the m random vectors from $\Lambda \cap [0, B)^{n+1}$ generate Λ . If the latter probability is p_2 , then the overall success probability is $\approx p_1^m p_2$, and one expects that one has to run the algorithm $\approx (p_1^m p_2)^{-1}$ times before it outputs a basis of Λ and thus of L itself. The main problem is that for $n > 1$, the lower bound one can prove for p_1 is quite small. In fact, it seems unavoidable that p_1 is bounded away from 1 by a nonzero constant. Therefore, it becomes evident why it is so important to minimize m without decreasing p_2 too much.

Theorem 1.1 shows that the quantum algorithm can recover a basis of L with constant probability p_2 conditioned on the event that it has obtained $m = 2n + 1$ samples of Λ , which occurs with probability $p_1^m = p_1^{2n+1}$. For this, we use two different window sizes: the first n vectors are sampled from a smaller window $[0, B)^{n+1}$, and the latter $n + 1$ vectors from a larger window $[0, B_1)^{n+1}$ with $B_1 > B$. To the best of our knowledge, this is the first explicit result that leads to a rigorous bound on the running time of the quantum algorithm.

Conjecture implies that the quantum algorithm can recover a basis of L with constant probability p_2 conditioned on the event that it has obtained only $m_c = n + 1$ samples of Λ . This event occurs with probability $p_1^{m_c} = p_1^{n+1}$, which is greater by the exponential factor $1/p_1^n$. Moreover, the quantum algorithm becomes simpler because it suffices to sample vectors from one window.

For the sake of completeness, we now describe the reduction in more detail. The infrastructure of a number field is isomorphic to a torus $T = \mathbb{R}^d/M$, where M is a full-rank lattice in \mathbb{R}^d and the coefficients of all non-trivial vectors of M are transcendental numbers (Fontein, 2011). This forces one to work with approximations, which is ultimately responsible for the poor performance when the dimension d increases.

Assume that we are given two elements $x, y \in T$ and we want to find the “discrete logarithm” $\ell \in \mathbb{Z}$ with $\ell x = y$, assuming that such number ℓ exists. For this, it suffices to know $v, w \in \mathbb{R}^d$ with $v + M = x$ and $w + M = y$ together with a basis of M : then one can use linear algebra to recover ℓ (and decide whether it exists). In the infrastructure of a number field, one has a representation of T which allows us to compute the projection $\mathbb{R}^d \rightarrow T$ easily, but recovering a preimage of a random $x \in T$ is hard.

Finding a preimage can be reformulated as a lattice problem: consider the map

$$\phi : \mathbb{Z} \times \mathbb{R}^d \rightarrow T, \quad (z, v) \mapsto zx + v;$$

this is a group homomorphism, and the kernel $L := \ker \phi$ is a lattice in \mathbb{R}^{d+1} of full rank. The kernel contains elements of the form $(-1, v)$ with $v \in \mathbb{R}^d$; any such element satisfies $v + M = x$. If we have a basis of L , we can again use linear algebra to recover such an element. Thus, the task is to find a basis of a lattice $L \subseteq \mathbb{R}^{d+1}$ of full rank.

Finally, we want to note that using different distributions on the lattice vectors can lead to much better results which are simpler to obtain. For example, when using the discrete Gaussian distribution on the lattice points, a result similar to ours in Theorem 1.1 follows from works by Micciancio and Regev (2007), Gentry et al. (2008). Unfortunately, it is not known how to sample from this distribution on a quantum computer, even in case a basis of the lattice is given. For our problem, where we want to determine a basis, we are only given an indirect description of the lattice. Therefore, these results cannot be used for solving the DLP without new ideas.

Acknowledgements

P.W. gratefully acknowledges the support from the NSF grant CCF-0726771 and the NSF CAREER Award CCF-0746600. P.W. would also like to thank Joachim Rosenthal and his group members for their hospitality during his visit at the Institute of Mathematics, University of Zurich. F.F. gratefully acknowledges partial support from the SNF grant No. 132256. Both authors would like to thank A. Schmidt for pointing out an error in an earlier preprint, both referees for their valuable comments, and S. Elizalde and K. Woods for pointing us to their paper (Elizalde and Woods, 2007).

References

- Acciario, V., 1996. The probability of generating some common families of finite groups. *Util. Math.* 49, 243–254.
- Barvinok, A., no date. Math669: Combinatorics, geometry and complexity of integer points. <http://www.math.lsa.umich.edu/~barvinok/latticenotes669.pdf>.
- Buchmann, J.A., 1991. Number theoretic algorithms and cryptology. In: FCT'91: Proceedings of the 8th International Symposium on Fundamentals of Computation Theory. Springer-Verlag, London, UK, pp. 16–21.
- Cohen, H., Frey, G., Avanzi, R., Doche, C., Lange, T., Nguyen, K., Vercauteren, F. (Eds.), 2006. Handbook of Elliptic and Hyperelliptic Curve Cryptography. Discrete Math. Appl. (Boca Raton), Chapman & Hall/CRC, Boca Raton, FL.
- Cheung, K.K.H., Mosca, M., 2001. Decomposing finite abelian groups. *Quantum Inf. Comput.* 1 (3), 26–32.
- Elizalde, S., Woods, K., 2007. The probability of choosing primitive sets. *J. Number Theory* 125 (1), 39–49.
- Fontein, F., 2011. The infrastructure of a global field of arbitrary unit rank. *Math. Comput.* 80 (276), 2325–2357.
- Gentry, Craig, Peikert, Chris, Vaikuntanathan, Vinod, 2008. Trapdoors for hard lattices and new cryptographic constructions [extended abstract]. In: STOC'08. ACM, New York, pp. 197–206.
- Hallgren, S., 2002. Polynomial-time quantum algorithms for Pell's equation and the principal ideal problem. In: Proceedings of the Thirty-Fourth Annual ACM Symposium on Theory of Computing. ACM, New York, pp. 653–658 (electronic).
- Hallgren, S., 2005. Fast quantum algorithms for computing the unit group and class group of a number field. In: STOC'05: Proceedings of the 37th Annual ACM Symposium on Theory of Computing. ACM, New York, pp. 468–474.
- Integer sequence A021002, no date. The on-line encyclopedia of integer sequences. <http://oeis.org/A021002>.
- Lehmer, D.N., 1900. Asymptotic evaluation of certain totient sums. *Am. J. Math.* 22 (4), 293–335.
- Micciancio, D., Goldwasser, S., 2002. Complexity of Lattice Problems. Kluwer Int. Ser. Eng. Comput. Sci., vol. 671. Kluwer Academic Publishers, Boston, MA. A cryptographic perspective.
- Micciancio, Daniele, Regev, Oded, 2007. Worst-case to average-case reductions based on Gaussian measures. *SIAM J. Comput.* 37 (1), 267–302 (electronic).
- Maze, G., Rosenthal, J., Wagner, U., 2011. Natural density of rectangular unimodular integer matrices. *Linear Algebra Appl.* 434 (5), 1319–1324.
- Menezes, A.J., van Oorschot, P.C., Vanstone, S.A., 1997. Handbook of Applied Cryptography. Discrete Math. Appl. (Boca Raton). CRC Press, Boca Raton, FL. With a foreword by Ronald L. Rivest.
- Pomerance, C., 2001. The expected number of random elements to generate a finite abelian group. *Period. Math. Hung.* 43 (1–2), 191–198.
- Remak, R., 1932. Über die Abschätzung des absoluten Betrages des Regulators eines algebraischen Zahlkörpers nach unten. *J. Reine Angew. Math.* 167, 360–378.
- Sands, J.W., 1991. Generalization of a theorem of Siegel. *Acta Arith.* 58 (1), 47–57.
- Scheidler, R., Buchmann, J.A., Williams, H.C., 1994. A key-exchange protocol using real quadratic fields. *J. Cryptol.* 7 (3), 171–199.
- Schmidt, A., 2007. Zur Lösung von zahlentheoretischen Problemen mit klassischen und Quantencomputern. Ph.D. thesis. Technische Universität Darmstadt.
- Shor, Peter W., 1997. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.* 26 (5), 1484–1509.
- Schmidt, A., Vollmer, U., 2005. Polynomial time quantum algorithm for the computation of the unit group of a number field (extended abstract). In: STOC'05: Proceedings of the 37th Annual ACM Symposium on Theory of Computing. ACM, New York, pp. 475–480.
- Sarvepalli, P., Wocjan, P., 2013. Quantum algorithms for one-dimensional infrastructures. *Quantum Inf. Comput.* 14 (1 & 2), 56–90.