

Quantum Algorithm for Computing the Period Lattice of an Infrastructure

Felix Fontein* Pawel Wocjan†

June 13, 2012

We present a quantum algorithm for computing the period lattice of infrastructures of fixed dimension. The algorithm applies to infrastructures that satisfy certain conditions. The latter are always fulfilled for infrastructures obtained from global fields, i.e., algebraic number fields and function fields with finite constant fields, as described in [Fon11].

The first of our main contributions is an exponentially better method for sampling approximations of vectors of the dual lattice of the period lattice than the methods outlined in the works of HALLGREN and SCHMIDT AND VOLLMER. This new method improves the success probability by a factor of at least 2^{n^2-1} where n is the dimension. The second main contribution is a rigorous and complete proof that the running time of the algorithm is polynomial in the logarithm of the determinant of the period lattice and exponential in n . The third contribution is the determination of an explicit lower bound on the success probability of our algorithm which greatly improves on the bounds given in the above works.

The exponential scaling seems inevitable because the best currently known methods for carrying out fundamental arithmetic operations in infrastructures obtained from algebraic number fields take exponential time. In contrast, the problem of computing the period lattice of infrastructures arising from function fields can be solved without the exponential dependence on the dimension n since this problem reduces efficiently to the abelian hidden subgroup problem. This is also true for other important computational problems in algebraic geometry. The running time of the best classical algorithms for infrastructures arising from global fields increases subexponentially with the determinant of the period lattice.

*Institute of Mathematics, University of Zurich, Winterthurerstrasse 190, 8057 Zurich, Switzerland; felix.fontein@math.uzh.ch

†Department of Electrical Engineering and Computer Science, University of Central Florida, Orlando, FL 32816-2362; wocjan@eecs.ucf.edu

Contents

1	Introduction	3
1.1	Informal definition of an infrastructure and the problem of computing the period lattice	3
1.2	Intuition behind the quantum algorithm and brief summary of new contributions	3
1.3	Efficient quantum algorithms for problems in arithmetic geometry	5
2	Formal definition of an infrastructure	6
3	Detailed outline of the quantum algorithm and new contributions	9
4	Computing the function f that hides the period lattice Λ	12
5	Preparing periodic states	17
6	Sampling approximations of vectors of the dual lattice Λ^*	21
6.1	Sampling in dimension greater than one	21
6.2	Sampling in dimension one	24
7	Lattice theoretic tools – Part 1	24
7.1	Lattices of dimension greater than one	24
7.1.1	Probability of generating a full-rank sublattice L_0 of L . .	24
7.1.2	Probability of generating finite abelian groups	27
7.1.3	Probability of generating the entire lattice L	28
7.2	Lattices of dimension one	31
8	Obtaining an approximate generating set of the dual lattice Λ^*	32
8.1	Lattices of dimension greater than one	32
8.1.1	Generating a full-rank sublattice of the dual lattice	32
8.1.2	Generating the entire dual lattice	34
8.1.3	Bounding the probability	35
8.2	Dimension one	36
9	Lattice theoretic tools – Part 2	37
9.1	Computing an approximate basis of L from an approximate generating set of L	37
9.2	Computing an approximate basis of the dual lattice L^* from an approximate basis of L	44
10	Final analysis of the quantum algorithm	46
	List of assumptions	50

1 Introduction

1.1 Informal definition of an infrastructure and the problem of computing the period lattice

An n -dimensional infrastructure \mathcal{I} is a finite set of distinguished points on an n -dimensional torus \mathbb{R}^n/Λ , where Λ is a lattice of full rank in \mathbb{R}^n . To every of these finitely many distinguished points, we assign a region on the torus, so that every point on the torus lies in exactly one such region. If x is such a distinguished point, every point y in this region can be represented by the difference $t := y - x$ together with x , i.e., as the pair (x, t) . These tuples (x, t) are essentially the f -representations of the infrastructure. Infrastructures can be obtained, for example, from global fields, i.e., from algebraic number fields as well as function fields with finite constant fields; in this case, the lattice corresponds to the free part of the unit group. We explain later that such infrastructures satisfy all assumptions we make on infrastructures in this paper.

We present a quantum algorithm for computing the period lattice Λ of infrastructures of fixed dimension n and provide a rigorous and detailed proof of its performance. We focus our attention on non-discrete infrastructures. An infrastructure is called discrete if its period lattice is integral and the coordinates of the distinguished points are integral (or more generally, if everything can be made integral by a suitable rescaling). Discrete and non-discrete infrastructures arise from function fields and number fields, respectively. The problem of computing the period lattice of discrete infrastructures is easy since this problem can be solved by using the same approach as for the abelian hidden subgroup problem. The reason is that the quantum algorithm for solving the abelian HSP can also be viewed as computing a hidden lattice in \mathbb{Z}^n .

1.2 Intuition behind the quantum algorithm and brief summary of new contributions

The idea behind the quantum algorithm for computing the period lattice of a (non-discrete) infrastructure is as follows. It is possible to define a function from the window $\mathcal{V} = \{0, \dots, qN - 1\}^n \subset \mathbb{Z}^n$ into a certain finite set, whose elements are related to f -representations, so that

$$f(v) = f(v') \Leftrightarrow \frac{v - v'}{N} \approx \lambda \text{ for some } \lambda \in \Lambda.$$

In words, there is a collision iff the two values v and v' differ approximately by an integer multiple of a lattice vector of the period lattice. This implies that the elements of the preimage $f^{-1}(v)$ have the special form

$$v' = v + N\lambda + \xi_\lambda,$$

where $\lambda \in \Lambda$ and ξ_λ is a certain error vector from $(-1, 1)^n$ such that $v' \in \mathcal{V}$. Moreover, for a constant fraction of v the cardinality of the corresponding

preimage is $f^{-1}(v)$ is close to $\frac{q^n}{\det(\Lambda)}$, which corresponds to the natural density of the lattice Λ in \mathbb{R}^n .

We prove that such function f exists and can always be evaluated correctly at all points of \mathcal{V} with constant probability. Our analysis takes into account the special nature of the shapes of the regions of the distinguished points and the way how these regions interlock with each other. This analysis closes a gap in the work [Hal05]. The works [SV05, Sch07] chose a different approach. They showed that it is not necessary that the function f can always be evaluated correctly. However, their resulting analysis leads to a significantly worse overall running time.

Efficiency means here that we can evaluate this function in time that is polynomial in the logarithm of the determinant of the period lattice Λ and exponential in the dimension n . This exponential scaling seems inevitable because the best methods for carrying out fundamental arithmetic operations in such infrastructures take exponential time.

Following the quantum algorithm for the abelian HSP, we start by evaluating the function f in superposition over the window \mathcal{V} and measuring the output register. The resulting post-measurement state is a “pseudo-periodic” state, i.e., a uniform superposition of the above v' . It is important that this superposition contains sufficiently many values of the form v' . The pseudo-periodic state corresponds to a uniform superposition of a randomly translated rectangular portion of the rescaled lattice $N\Lambda$ such that only few of its points are missing and the remaining points are only slightly perturbed. We present a new method for precisely analyzing the probability of obtaining a pseudo-periodic state with sufficiently values of the v' . This analysis also closes a gap in the work [Hal05].

Similarly to the situation in the abelian HSP, we effectively remove the undesired random offset v by applying a multidimensional quantum Fourier transform. This allows us to sample approximations of lattice vectors of the dual lattice Λ^* . To mitigate the perturbations effects caused by the error vectors ξ_λ , we have to perform the quantum Fourier transform over a larger window \mathcal{W} . But this comes at the price of an exponential decay of the success probability with increasing dimension n . The idea to use a larger window goes back to [Hal05] and [SV05, Sch07]. We obtain here a new better method for sampling approximations improving the success probability by the exponential factor 2^{n-1} compared to the less efficient methods in [Hal05] and [SV05, Sch07]. This is not just an improvement in the analysis, but an improvement of the algorithm.

We present lattice and group theoretic results, making it possible to prove that $2n + 1$ approximations obtained by the above sampling process form an approximate generating set of Λ^* with constant probability for fixed dimension. No such bound on the number of required samples was proved in the previous works. Once we have such approximate generating set, we recover an approximate basis of Λ^* . We describe an improved method for this purpose. We then determine an approximate basis of Λ from such approximate basis of Λ^* .

Finally, we discuss in detail how to choose all parameters to obtain an approximate basis of the period lattice Λ that has the desired approximation

quality. We obtain an explicit lower bound on the success probability of our algorithm, which reveals precisely how the complexity depends on the various parameters. We compare this probability to the ones presented in the works of SCHMIDT AND VOLLMER and SCHMIDT and conclude that our probability is exponentially better by at least 2^{n^2-1} . The work of HALLGREN gives no explicit probability.

Note that in the one-dimensional case more specialized algorithms lead to a much better probability of success; see, for example, [Hal02, SW11].

1.3 Efficient quantum algorithms for problems in arithmetic geometry

We conclude the introduction with some comments on the existence of efficient quantum algorithms for certain computationally hard problems in algebraic geometry. Readers not familiar with algebraic geometry may not be aware that many interesting problems can be reduced to the abelian HSP efficiently. The understanding of these reductions does require some specialized knowledge in algebraic geometry, but the necessary results are fairly standard. As noted previously, infrastructures obtained from function fields are easier to handle than general infrastructures. As shown in Theorem 7.1 of [Fon11], such infrastructures embed in a natural way into the divisor class group of degree zero, which is a finite abelian group in the case of function fields with finite constant field. There are polynomial time classical algorithms to do arithmetic in this group, for instance, the “algebraic” algorithm by F. Heß [Hes02, Die08]. Therefore, one can directly apply the standard algorithm for the abelian HSP [CM01] to compute the period lattice. Other important problems, such as determining discrete logarithms in the infrastructure, computing the whole divisor class group and the ideal class group, solving the principal ideal problem, as well as computing the Zeta function, can all be treated in the same way. The latter problem was solved in [Ked06] using this approach, while relying on a less efficient “geometric” method based on the Brill-Noether algorithm to do arithmetic.

Arithmetic geometry provides a unifying understanding and treatment of problems related to global fields. On the one hand, the discussion above shows that the algebro-geometric problems for function fields with finite constant fields (i.e., function fields of curves over finite fields) can be reduced to the abelian HSP. This presents an elegant and efficient quantum solution. On the other hand, the analysis of the quantum algorithms for the corresponding number-theoretic problems is significantly more challenging. We believe that our rigorous and improved treatment of the problem of computing the period lattice of non-discrete infrastructures can serve as a valuable starting point for addressing other number-theoretic problems and also finding more efficient quantum algorithms for them. A first stepping stone is our new method for sampling approximations of vectors of the dual lattice, which improves the success probability by an exponential factor.

2 Formal definition of an infrastructure

An n -dimensional infrastructure \mathcal{I} consists of

- a lattice Λ of full rank, called the *period lattice*,
- a finite non-empty set X , an injective map $d : X \rightarrow \mathbb{R}^n/\Lambda$, and
- a set of f -representations $\text{Rep}^f(\mathcal{I})$, i.e., a subset $\text{Rep}^f(\mathcal{I}) \subseteq X \times \mathbb{R}^n$ with $X \times \{0\} \subseteq \text{Rep}^f(\mathcal{I})$ such that the function

$$\Phi_{\mathcal{I}} : \text{Rep}^f(\mathcal{I}) \rightarrow \mathbb{R}^n/\Lambda, \quad (x, t) \mapsto d(x) + t$$

is a bijection.

Such a set of f -representations yields a *reduction map* $\text{red} : \mathbb{R}^n/\Lambda \rightarrow X$ satisfying $\text{red}(\Phi_{\mathcal{I}}(x, t)) = x$ for all $(x, t) \in \text{Rep}^f(\mathcal{I})$, as well as a *giant step* operation $\text{gs} : X \times X \rightarrow X$ by $\text{gs}(x, y) = \text{red}(d(x) + d(y))$. Note that the set of f -representations has a natural group structure using the pull-back of the group operation of \mathbb{R}^n/Λ via $\Phi_{\mathcal{I}} : (x, t) + (x', t') := \Phi_{\mathcal{I}}^{-1}(\Phi_{\mathcal{I}}(x, t) + \Phi_{\mathcal{I}}(x', t'))$.

Given such a set of f -representations, we can *unroll* the infrastructure. Let $\pi : \mathbb{R}^n \rightarrow \mathbb{R}^n/\Lambda$ be the canonical projection, and set

$$\hat{X} := \pi^{-1}(d(X)).$$

This is a discrete non-empty subset of \mathbb{R}^n satisfying $\hat{X} + \Lambda = \hat{X}$. Define $\hat{d}(\hat{x}) = \hat{x}$ for all $\hat{x} \in \hat{X}$ and

$$\hat{V}_{\hat{x}} := \{\hat{d}(\hat{x}) + t \mid (d^{-1}(\pi(\hat{x})), t) \in \text{Rep}^f(\mathcal{I})\}$$

for every $\hat{x} \in \hat{X}$. Then \mathbb{R}^n is the disjoint union of all $\hat{V}_{\hat{x}}$, $\hat{x} \in \hat{X}$, and one can define $\widehat{\text{red}} : \mathbb{R}^n \rightarrow \hat{X}$ by $\widehat{\text{red}}(v) = \hat{x}$ if $v \in \hat{V}_{\hat{x}}$.

The unrolled infrastructure is periodic with period lattice Λ in the sense that for $\hat{x} \in \hat{X}$, $v \in \mathbb{R}^n$ and $\lambda \in \Lambda$, we have $\hat{x} + \lambda \in \hat{X}$, $\hat{V}_{\hat{x} + \lambda} = \hat{V}_{\hat{x}} + \lambda$, $\widehat{\text{red}}(v + \lambda) = \widehat{\text{red}}(v) + \lambda$ and $\hat{d}(\hat{x} + \lambda) = \hat{d}(\hat{x}) + \lambda$. Moreover, $\pi(\hat{x}) = \pi(\hat{y})$ for $\hat{x}, \hat{y} \in \hat{X}$ if, and only if, $\hat{y} - \hat{x} \in \Lambda$.

For $s, t \in \mathbb{R}^n$, we write $[s, t]$ for $\{r \in \mathbb{R}^n \mid s \leq r \leq t\}$, where “ \leq ” denotes the component-wise inequality on \mathbb{R}^n . We say that a subset $U \subseteq \mathbb{R}^n$ is *cornered* with *corner* $s \in \mathbb{R}^n$ if $s \in U$ and for every $t \in U$, $t \in [s, t] \subseteq U$. In other words, $U = \bigcup_{t \in U} [s, t]$. Note that every cornered subset of \mathbb{R}^n has exactly one corner, which is its minimal element with respect to \leq . We say that \mathcal{I} is *cornered* if for all $\hat{x} \in \hat{X}$, $\hat{V}_{\hat{x}}$ is cornered with corner \hat{x} .

We make the following assumptions:

- A1) There exists a constant $A > 0$ such that for every $\hat{x} \in \hat{X}$,

$$\hat{V}_{\hat{x}} \subseteq \hat{x} + [0, A]^n.$$

A2) There exist constants $C, D > 0$ such that for every $r \in \mathbb{R}^n$, the set

$$(r + [0, C]^n) \cap \hat{X}$$

contains at most D elements.

A3) There exists a polynomial-time algorithm such that for given $k \in \mathbb{N}$ and $u \in \mathbb{Z}^n$, one can compute $(x, t) \in X \times 2^{-k}\mathbb{Z}^n$ such that

- (a) $\|\hat{x} + t - 2^{-k}u\|_\infty \leq 2^{-k}$ for some $\hat{x} \in \hat{X}$ with $d^{-1}(\pi(\hat{x})) = x$;
- (b) $(2^{-k}u + (-2^{-k}, 2^{-k})^n) \cap \hat{V}_{\hat{x}} \neq \emptyset$.

The running time is polynomial in k and $\log \|u\|_\infty$ when the dimension n is held constant.

Proposition 2.1. *Let K be a global field. Then any infrastructure obtained from K in the sense of [Fon11, Section 6] has f -representations in a natural way and is cornered. Moreover, it satisfies A1) to A3) with explicit constants A, C, D :*

If K is a number field of discriminant Δ and degree $d = [K : \mathbb{Q}]$, then one can choose $A = \frac{1}{2} \log |\Delta|$, $C = \log 2$ and $D = 4^d$. If K is a function field of genus g and degree $d = [K : k(x)]$, then one can choose $A = g + d - 1$, $C = 1 - \varepsilon$ for any $\varepsilon \in (0, 1)$, and $D = 1$.

Sketch of Proof. Assume that the infrastructure is $\mathcal{I} = (X^{\mathfrak{a}}, d^{\mathfrak{a}}, \text{red}^{\mathfrak{a}})$ in the notation of [Fon11]. Here, \mathfrak{a} is an ideal of the ring of integers \mathcal{O} (or the ring of holomorphic functions in case K is a function field), and $X^{\mathfrak{a}}$ is essentially the set of reduced ideals equivalent to \mathfrak{a} . If $|\bullet|_1, \dots, |\bullet|_{n+1}$ are the pairwise different absolute values of K , we define $\Lambda := \{(\log |\varepsilon|_1, \dots, \log |\varepsilon|_n) \mid \varepsilon \in \mathcal{O}^*\}$, which is isomorphic to the free part of the finitely generated abelian group \mathcal{O}^* of units of \mathcal{O} . The definition of f -representations is rather technical, whence we do not repeat it here, but just refer to Definition 6.3 of [Fon11]. For every $\hat{x} \in \hat{X}$,

$$\hat{V}_{\hat{x}} = \hat{x} + W(d^{-1}(\pi(\hat{x}))), \quad \text{where } W(x) = \{t \in \mathbb{R}^n \mid (x, t) \in \text{Rep}^f(\mathfrak{a})\} \text{ for } x \in X.$$

It is clear from Definition 6.3 in [Fon11] that $W(x)$ is cornered with corner 0. Hence, \mathcal{I} is a cornered infrastructure. Our assumption A1) follows from Proposition 8.1 of [Fon11]. The second assumption A2) holds trivially for function fields; for number fields, it follows from Lemma 3.2 in [Buc87b].

Assumption A3) will be discussed in an upcoming paper of the first author and M. J. Jacobson, Jr. In the case of function fields, the algorithms are of polynomial running time with respect to the genus of the function field as well as the size of its representation. In the case of number fields, the algorithms are polynomial with respect to the logarithm of the discriminant of the number field, but exponential in its degree $d = [K : \mathbb{Q}]$, as one has to find shortest vectors in lattices of dimension d . \square

Note that the algorithm we plan to use for **A3**) is exponential in n , but significantly more efficient than the algorithms that were proposed in [Hal05] and [SV05]. These are based on [Thi95a, Chapter 5 and 6], which essentially uses Buchmann’s baby step algorithm [Buc87a, Buc87c]. The latter is known for being practically unusable [BJP94]. Even on modern computers, computing all minima of one reduced ideal can take a long time for not too large number field degrees, say $[K : \mathbb{Q}] = 8$ (which yields $n = 7$); the first author verified this in 2010 when implementing that algorithm.

Note that Schoof’s Algorithm 10.7 in [Sch08] is also mentioned in [Hal05] as a more efficient alternative to Buchmann’s algorithm. Unfortunately, Schoof’s algorithm uses a different distance function from the one used by Hallgren and by us. Therefore, Schoof’s algorithm cannot be applied without non-trivial modifications if one wants to obtain a provably polynomial-time quantum algorithm for computing the period lattice.

Observe that **A3**) follows from the existence of two simpler algorithms. Before we list these, we need to define what an “approximate f -representation of error at most ε ” of a point $r \in \mathbb{R}^n$ is. This is a pair $(x, t) \in X \times \mathbb{R}^n$ satisfying

(a) $\|\hat{x} + t - r\|_\infty \leq \varepsilon$ for some $\hat{x} \in \hat{X}$ with $d^{-1}(\pi(\hat{x})) = x$;

(b) $(r + (-\varepsilon, \varepsilon))^n \cap \hat{V}_{\hat{x}} \neq \emptyset$,

Now we can describe the characteristics of the two simpler algorithms, which can be combined to obtain such an algorithm as described in **A3**):

(a) one algorithm which, given $\ell \in \mathbb{N}$ and $r \in 2^{-\ell}\{-2^\ell, -2^\ell + 1, \dots, 2^\ell\}^n \subset [-1, 1]^n$, computes an approximate f -representation (x, t) of error at most $2^{-\ell}$ such that $\|d(x) + t - r\|_\infty \leq 2^{-\ell}$ in time polynomial in ℓ ;

(b) a second algorithm which, given two approximate f -representations of error at most $2^{\ell'}$, computes an approximate f -representation of their sum of error at most $2^{\ell'+1}$ in time polynomial in ℓ' .

One can compute an approximate f -representation of any $r \in \mathbb{R}^n$ of error at most 2^{-k} in time polynomial in $\log \|r\|_\infty$ and k . This is done by using a double and add technique and by calling these algorithms to obtain approximate f -representations of error at most $2^{-(k+k')}$, where $k' = O(\log \|r\|_\infty)$.

The formal definition of the problem of computing the period lattice is as follows.

Definition 2.2. *Given $\gamma \in (0, 1)$, the task is to find $\tilde{\lambda}_1, \dots, \tilde{\lambda}_n \in \mathbb{R}^n$ such that there exists a basis $\lambda_1, \dots, \lambda_n$ of Λ with*

$$\|\tilde{\lambda}_j - \lambda_j\|_2 \leq \gamma$$

for $j = 1, \dots, n$. We call such $\tilde{\lambda}_1, \dots, \tilde{\lambda}_n$ a γ -approximate basis of Λ .

We present a quantum algorithm with running time polynomial in $\log \det(\Lambda)$ and $\log(1/\gamma)$ when $A, 1/C, D$ and $1/\lambda_1(\Lambda)$ can be bounded polynomially in

terms of $\log \det(\Lambda)$. Here, $\lambda_1(\Lambda)$ denotes the first consecutive minimum of Λ , i.e., the length of a shortest non-zero vector in Λ . Note that for number fields, $\lambda_1(\Lambda)$ can be bounded from below by a bound depending only on n ; see Satz 5.6 in [Buc87c].

In the case of computing units of a global field, computing a γ -approximate basis of Λ yields approximations of the logarithms of the absolute values of the units. These approximations can be refined to arbitrary precision in polynomial time. Note that one can also relatively efficiently recover the corresponding units themselves; since their representation is not of size polynomial in the genus respectively logarithm of the discriminant, explicitly computing them cannot be done in polynomial time. What can be done is computing a so-called *compact representation* of a unit, which was presented for number fields in [Thi95a, Thi95b] and for function fields in [EH12]; one can modify the quantum algorithm to output such compact representations of the units and still run in polynomial time.

Finally, we want to mention that our algorithm can be interpreted as an algorithm for solving certain instances of a *Hidden Subgroup Problem* for the group $G = \mathbb{R}^n$ provided that the group operation in $\text{Rep}^f(\mathcal{I})$ is effective. In case the infrastructure is obtained from a global field as in the above proposition, the group operation is effective and is described explicitly in Theorem 7.3 of [Fon11].

Now one can consider the group homomorphism $f : \mathbb{R}^n \rightarrow \text{Rep}^f(\mathcal{I})$ as the composition of the canonical projection $\pi : \mathbb{R}^n \rightarrow \mathbb{R}^n/\Lambda$ with $\Phi_{\mathcal{I}}^{-1}$. This map can be effectively computed – ignoring rounding and approximation issues – and it hides the lattice Λ by $\ker f = \Lambda$.

3 Detailed outline of the quantum algorithm and new contributions

Let $N \in \mathbb{N}$ and $s \in \mathbb{R}^n$ be fixed. Consider the function

$$f : \mathbb{R}^n \rightarrow X \times \mathbb{Z}^n, \quad v \mapsto (x, \lfloor Nt \rfloor) \text{ if } \Phi_{\mathcal{I}}^{-1}(\pi(s + \frac{1}{N}v)) = (x, t).$$

If $f(v) = f(v')$ for $v, v' \in \mathbb{Z}^n$, then $v - v'$ lies close to an element of $N\Lambda$. We want to use the quantum computer to find such collisions.

Let $\mathcal{V} = \{0, \dots, qN - 1\}^n$ and $\mathcal{W} = \{0, \dots, 2nqN - 1\}^n$ where q and N are positive integers that will be fixed later. Set $V = |\mathcal{V}|$ and $W = |\mathcal{W}|$. The input register is $\mathbb{C}^W = (\mathbb{C}^{2nqN})^{\otimes n}$. The output register is \mathbb{C}^d with d sufficiently large so it can store any element of the image $f(\mathcal{V})$. In the following we use f to denote the restriction of f to \mathcal{V} . We assume that we have a reversible version U_f of f that acts on the above input and output registers.

Algorithm

1. We start by preparing the state

$$\frac{1}{\sqrt{V}} \sum_{v \in \mathcal{V}} |v\rangle |f(v)\rangle.$$

Note that we evaluate f only on the subset \mathcal{V} of \mathcal{W} .

2. We measure the output register and denote the outcome by $f(v)$ for some $v \in \mathcal{V}$. The post-measurement state is then

$$\frac{1}{\sqrt{M}} \sum_{v' \in \mathcal{M}} |v'\rangle |f(v)\rangle$$

where $\mathcal{M} := \{v' \in \mathcal{V} \mid f(v') = f(v)\}$ and $M = |\mathcal{M}|$.

3. We apply the n -fold tensor product of the quantum Fourier transform of size $2nqN$ on the input register and obtain the state

$$\frac{1}{\sqrt{MW}} \sum_{w \in \mathcal{W}} \sum_{v \in \mathcal{M}} \exp\left(2\pi i v' \cdot \frac{w}{2nqN}\right) |w\rangle |(f(v))\rangle$$

where \cdot denotes the inner product on \mathbb{R}^n .

4. Finally, we measure the input register and denote the outcome by w .

This quantum procedure is repeated $2n + 1$ many times to obtain the samples w_1, \dots, w_{2n+1} . A subsequent classical post-processing step makes it possible to extract an approximate basis of Λ from these samples with a probability that can be bounded from below by a positive constant.

Organization of the paper and outline of technical results

In Section 4, we prove that with constant probability all evaluation points $v/N + s$ ($v \in \mathcal{V}$) are sufficiently far away from the boundary of $\hat{V}_{\hat{x}}$ for all $\hat{x} \in \hat{X}$. This is achieved by choosing the shift s uniformly at random from a certain finite set. This ensures that we can compute $f(v)$ correctly for all $v \in \mathcal{V}$ even though we may only determine approximate f -representations.

In Section 5, we show that the probability for post-measurement states being periodic states can be bounded from below by a constant. Roughly speaking, a periodic state corresponds to a (randomly) translated and perturbed finite portion of the lattice $N\Lambda$ that may be missing some points. In particular, we establish a lower bound on M showing that not too many points are missing in the superposition.

To derive the results in Sections 4 and 5, it is absolutely indispensable to take into account that the infrastructure is cornered. Relying only a lower bound on the minimal distance between two elements of \hat{X} is not sufficient because the union of ε -neighborhoods of the boundaries of $\hat{V}_{\hat{x}}$ of all $\hat{x} \in \hat{X}$ could still fill out

too much of \mathbb{R}^n . In the one-dimensional case, the regions $\hat{V}_{\hat{x}}$ are intervals. In contrast to that, in the n -dimensional case, their shapes can take on much more complicated forms. This makes the analysis more difficult. This problem was mentioned, but not resolved in [Hal05], while in [SV05, Sch07], this problem was solved differently by relaxing the conditions of the quantum algorithm on the function f .

In Section 6, we show that the last step of the above quantum procedure yields an approximation of an element of the dual lattice $\Lambda^* = \{\lambda^* \in \mathbb{R}^n \mid \forall \lambda \in \Lambda : \langle \lambda^*, \lambda \rangle \in \mathbb{Z}\}$ with a certain probability. It becomes essential here that the Fourier transform is taken over the larger window \mathcal{W} , while f is only evaluated inside \mathcal{V} . This makes it possible to mitigate the perturbation effects.

More precisely, we determine a lower bound on the probability the outcome w obtain in the final step is contained in the set \mathcal{R}_{λ^*} , where

$$\mathcal{R}_{\lambda^*} := \left\{ (w_1, \dots, w_n) \mid w_k \in \{ \lfloor 2nq\lambda_k^* \rfloor, \lfloor 2nq\lambda_k^* \rfloor + 1 \} \text{ for } k = 1, \dots, n \right\}$$

and $\lambda^* = (\lambda_1^*, \dots, \lambda_n^*) \in \Lambda^*$. Such elements yield good approximations of λ^* since

$$\left\| \frac{w}{2nq} - \lambda^* \right\|_2 \leq \frac{1}{2\sqrt{nq}}$$

for all $w \in \mathcal{R}_{\lambda^*}$.

The works [Hal05] nor [SV05, Sch07] consider only elements of the more restrictive form $\lfloor 2nq\lambda^* \rfloor$, where $\lfloor u \rfloor$ means that we round each coefficient of $u \in \mathbb{R}^n$ to the closest integer. This is why our method improves the success probability of obtaining a single good approximation by the exponential factor 2^{n-1} . It can be shown that at least $n+1$ samples are needed so our method provably leads to an overall improvement of the success probability by the factor 2^{n^2-1} .

In Section 7, we present lattice and group theoretic results, yielding a lower bound on the probability that n lattice vectors drawn uniformly at random from $L \cap [0, b)^n$ and $n+1$ lattice vectors drawn uniformly at random from $L \cap [0, b_0)^n$ generate together the entire lattice L , where L is a full-rank lattice in \mathbb{R}^n and $b < b_0$ are sufficiently large. Neither [Hal05] nor [SV05, Sch07] provide an explicit and proven upper bound on the complexity of generating a lattice by drawing samples. But this is a crucial result, directly affecting the success probability of the algorithm.

In Section 8, we specialize these lattice-theoretic results to $L := \Lambda^*$ and present an explicit lower bound on the probability that the $2n+1$ samples w_1, \dots, w_{2n+1} output by our quantum algorithm yield an approximate generating set for the dual lattice Λ^* .

In Section 9, we first present technical results based on [BK93] showing how to construct an approximate basis of L from an approximate generating set of L . Then, we show how to recover an approximate basis of the dual lattice L^* from the previously determined approximate basis of L .

Finally, in Section 10, we combine all results from the previous sections and show to find an approximate basis for the period lattice Λ . We explain in detail how to choose all parameters. We also bound the success probability of our

algorithm from below. There is a classical method for checking whether the computed basis vectors are indeed close to elements of Λ . If that is the case, we have computed Λ with a high probability.

Unfortunately, the success probability of this algorithm decreases exponentially in the dimension n of the infrastructure. This is a common problem of such algorithms which also applies to the algorithms described in [Hal05] and [SV05] (see also [Sch07, p. 122]). However the success probability of our algorithm decreases less rapidly than that of the previous works. It is better by the exponential factor 2^{n^2-1} .

4 Computing the function f that hides the period lattice Λ

We consider a computable version \tilde{f} of f and show under which conditions $f(v) = \tilde{f}(v)$ holds for all $v \in \mathcal{V}$ with high probability. Recall that v corresponds to the point $s + \frac{v}{N}$, where s is a random offset. We show that if s is chosen uniformly random at random from a certain finite set, then with high probability none of these evaluation points $u := s + \frac{1}{N}v$ (for $v \in \mathcal{V}$) falls into regions in which the method **A3**) may return a result that leads to a wrong evaluation of $f(v)$.

Let $v \in \mathcal{V}$ yield $f(v) = (x, \lfloor Nt \rfloor)$ with $(x, t) = \Phi_{\mathcal{I}}^{-1}(\pi(u))$. Let $\hat{x} \in \hat{X}$ with $u \in \hat{V}_{\hat{x}}$; then $\pi(\hat{x}) = d(x)$ and $u - \hat{x} = t$. If u is sufficiently far away from $\partial\hat{V}_{\hat{x}}$, then the oracle in **A3**) returns the correct $x \in X$. Moreover, if $t = (t_1, \dots, t_n) \in \mathbb{R}^n$ has no coordinate which comes close to an integer multiple of $\frac{1}{N}$, then the coordinates of Nt are bounded away from integers and $\lfloor Nt \rfloor = \lfloor Nt' \rfloor$ for all t' which are close enough to t . This ensures that the oracle in **A3**) outputs an approximation (x, t') of $\Phi_{\mathcal{I}}^{-1}(u) = (x, t)$ such that $(x, \lfloor Nt \rfloor) = (x, \lfloor Nt' \rfloor)$.

A *boundary point* of $\hat{\mathcal{I}}$ is a point $u \in \mathbb{R}^n$ such that every neighborhood of u contains points from at least two different $\hat{V}_{\hat{x}}$. Denote the set of all boundary points by H ; then

$$H = \bigcup_{\hat{x} \in \hat{X}} \partial\hat{V}_{\hat{x}}.$$

For a given $\varepsilon > 0$, define the enhanced boundary

$$H(\varepsilon) := H + [-\varepsilon, \varepsilon]^n.$$

Observe that $\hat{X} \subset H \subset H(\varepsilon)$ since by assumption all $\hat{V}_{\hat{x}}$ are cornered sets with corner \hat{x} . An example of how cornered sets could tile the plane \mathbb{R}^2 is shown in Figure 1, in which the enhanced boundary $H(\varepsilon)$ is highlighted.

If $u = s + \frac{v}{N} \notin H(\varepsilon)$, then the oracle in **A3**) can be used to correctly compute the x part of $f(v) = (x, \lfloor Nt \rfloor)$. To ensure that the $\lfloor Nt \rfloor$ part of $f(v) = (x, \lfloor Nt \rfloor)$ is also correctly computed, we need v to avoid a larger set. Formally, we define

$$H^{\text{grid}}(\varepsilon) := \bigcup_{\hat{x} \in \hat{X}} \left(\left(\frac{1}{N}\mathbb{N}^n + \partial\hat{V}_{\hat{x}} \right) \cap \overline{\hat{V}_{\hat{x}}} \right) + [-\varepsilon, \varepsilon]^n.$$




Figure 1: Demonstrating the tiling of \mathbb{R}^2 by cornered sets $\hat{V}_{\hat{x}}$, $\hat{x} \in \hat{X}$. The enhanced boundary region $H(\varepsilon)$ is highlighted.

Clearly, we have $H(\varepsilon) \subseteq H^{\text{grid}}(\varepsilon)$ for all $N \geq 1$. An example of what $H^{\text{grid}}(\varepsilon)$ may look like is shown in Figure 2.

Lemma 4.1. *Let $L, N, q \in \mathbb{N}$ with $L, N, q \geq 1$ and ε with $0 < \varepsilon \leq \frac{1}{2NL}$ be given. Let*

$$S := \frac{1}{NL} \{0, \dots, L-1\}^n.$$

For $s \in S$, consider the shifted grid

$$G(s) := \{s + \frac{1}{N}v \mid v \in \mathcal{V}\}.$$

Assume that for some $s \in S$ we have $G(s) \cap H^{\text{grid}}(\varepsilon) = \emptyset$. Then, this implies the following two conditions:

1. For every $v \in \mathcal{V}$, there is exactly one $\hat{x} \in \hat{X}$ with $\hat{V}_{\hat{x}} \cap (s + \frac{1}{N}v + (-\varepsilon, \varepsilon)^n) \neq \emptyset$.
2. Let $T := \{t \in \mathbb{R}^n \mid \exists v \in \mathcal{V} : \Phi_{\mathcal{I}}^{-1}(s + \frac{1}{N}v) = (x, t)\}$. Then, we have $T \cap ([-\varepsilon, \varepsilon]^n + \frac{1}{N}\mathbb{N}^n) = \emptyset$.

These conditions show that we can compute f correctly using **A3**) if the precision 2^{-k} used there is at most $\frac{\varepsilon}{2}$: the first condition ensures that the x part of $f(v) = (x, \lfloor Nt \rfloor)$ can be computed exactly, and the second condition ensures that $\lfloor Nt \rfloor$ is exact.

Proof. Observe that $G(s) \cap H^{\text{grid}}(\varepsilon) = \emptyset$ implies $G(s) \cap H(\varepsilon) = \emptyset$. We show that the latter implies the first condition of the lemma. The more general $G(s) \cap H^{\text{grid}}(\varepsilon) = \emptyset$ is needed to eliminate some sporadic cases in the second condition.




Figure 2: The set $H(\epsilon)$ from Figure 1 is depicted in black. The corresponding set $H^{\text{grid}}(\epsilon)$ is highlighted in gray.

1. Since $\bigcup_{\hat{x} \in \hat{X}} \hat{V}_{\hat{x}} = \mathbb{R}^n$ there must be at least one such \hat{x} . Let $\hat{x} \in \hat{X}$ be one such element. In the case that $s + \frac{1}{N}v + (-\epsilon, \epsilon)^n$ is not completely contained in $\hat{V}_{\hat{x}}$, the translated open disc $s + \frac{1}{N}v + (-\epsilon, \epsilon)^n$ must contain some $y \in \partial \hat{V}_{\hat{x}}$. But this implies that $G(s) \ni s + \frac{1}{N}v \in y + [-\epsilon, \epsilon]^n \subseteq H(\epsilon)$, contradicting $G(s) \cap H(\epsilon) = \emptyset$. Thus $s + \frac{1}{N}v + (-\epsilon, \epsilon)^n \subseteq \hat{V}_{\hat{x}}$ and we are done.
2. Assume that $t \in T$ can be written as $t = \frac{1}{N}w + e$ with $w \in \mathbb{N}^n$ and $e \in [-\epsilon, \epsilon]^n$, i.e., $t \in T \cap ([-\epsilon, \epsilon]^n + \frac{1}{N}\mathbb{N}^n)$. As $t \in T$ there exists some $u \in G(s)$ with $\Phi_T^{-1}(\pi(u)) = (x, t)$ for $x \in X$. Let $\hat{x} \in \hat{X}$ with $u \in \hat{V}_{\hat{x}}$; then $u = \hat{x} + \frac{1}{N}w + e$. But this yields $u \in (\frac{1}{N}\mathbb{N}^n + \partial \hat{V}_{\hat{x}}) \cap \overline{\hat{V}_{\hat{x}}} + [-\epsilon, \epsilon]^n$ and thus $u \in H^{\text{grid}}(\epsilon)$. Hence, $u \in G(s) \cap H^{\text{grid}}(\epsilon)$. \square

We now determine a lower bound on the probability that the desired condition $G(s) \cap H^{\text{grid}}(\epsilon) = \emptyset$ holds when s is chosen uniformly at random in S and L is sufficiently large.

Proposition 4.2. *Let $q, N \in \mathbb{N}$ with $q, N \geq 1$ and $p \in (0, 1)$ be given. Choose L and ϵ such that*

$$L \geq \frac{2nD(q + A + C + 2)^n}{(1 - p)C^n} \quad \text{and} \quad \epsilon \leq \frac{1}{2NL}.$$

If we select $s \in S$ uniformly at random, then

$$\Pr(G(s) \cap H^{\text{grid}}(\epsilon) = \emptyset) \geq p.$$

The main idea behind the proof of this proposition is as follows: while $\partial\hat{V}_{\hat{x}}$ for a single \hat{x} can be difficult to describe, the union of all $\partial\hat{V}_{\hat{x}}$, where \hat{x} ranges over all $\hat{x} \in \hat{X}$, has a much simpler structure. For instance, in the case of $n = 2$, i.e., in the plane, it suffices to consider only two faces of $\partial\hat{V}_{\hat{x}}$, namely, the ones incident with \hat{x} . Let us call these two faces the *principal boundaries* of $\hat{V}_{\hat{x}}$. Every boundary point is an element of a principal boundary of at least one $\hat{V}_{\hat{x}}$. The principal boundaries of some $\hat{V}_{\hat{x}}$ from Figure 1 are shown in Figure 3; note that we capped off the ends of the principal boundaries to make it possible to distinguish between different principal boundaries. The corners of the sets are marked by large dots, and the principal boundaries by thick lines. The proof works by covering the principal boundaries by larger sets which are known to cover them – for this, we need assumption **A1**).




Figure 3: Example showing the principal boundaries of some of the cornered sets of H from Figure 1. The principal boundaries are depicted in black. We capped them in the figure to make clear to which corners they belong.

Proof of Proposition 4.2. Define

$$\begin{aligned}\hat{X}' &:= \hat{X} \cap [-A - \varepsilon, q + 1 + \varepsilon]^n \\ F(s) &:= \left(s + \frac{1}{N}\mathbb{Z}^n\right) \cap \left(\bigcup_{\hat{x} \in \hat{X}'} \partial\hat{V}_{\hat{x}} + [-\varepsilon, \varepsilon]^n\right).\end{aligned}$$

We first show that $F(s) = \emptyset$ implies the desired condition $G(s) \cap H^{\text{grid}}(\varepsilon) = \emptyset$ and then bound the number of s in S for which it may be the case that $F(s) \neq \emptyset$.

We prove the first part by considering the contraposition of the implication. Assume that some $u \in G(s) \cap H^{\text{grid}}(\varepsilon)$ exists. Then there exists some $\hat{x} \in \hat{X}$ and some $e \in [-\varepsilon, \varepsilon]^n$ such that $u - e \in (\frac{1}{N}\mathbb{N}^n + \partial\hat{V}_{\hat{x}}) \cap \overline{\hat{V}_{\hat{x}}}$, and $u = s + \frac{1}{N}v$ with $v \in \mathcal{V} = \mathbb{Z}^n \cap [0, qN - 1]^n$. In particular, $u \in [0, q - \frac{1}{NL}]^n$ and hence $u - e \in [-\varepsilon, q + \varepsilon]^n$.

We have $\hat{V}_{\hat{x}} \subseteq \hat{x} + [0, A]^n$ since $\hat{V}_{\hat{x}}$ is cornered with corner \hat{x} and **A1**) holds. This implies $\hat{x} \in u - e - [0, A]^n \subseteq [-\varepsilon - A, q + \varepsilon]^n$ and hence $\hat{x} \in \hat{X}'$. As $u - e \in \frac{1}{N}\mathbb{N}^n + \partial\hat{V}_{\hat{x}}$, we can write $u - e = b + \frac{1}{N}w$ with $b \in \partial\hat{V}_{\hat{x}}$ and $w \in \mathbb{N}^n$. But this yields that $b + e \in F(s)$ and hence $F(s) \neq \emptyset$.

We now bound the number of $s \in S$ with $F(s) \neq \emptyset$. For $\hat{x} \in \hat{X}$, define

$$H'(\hat{x}, \varepsilon, i) := \{\hat{x} + (t_1, \dots, t_n) \mid -\varepsilon \leq t_j \leq A + \varepsilon, -\varepsilon \leq t_i \leq \varepsilon\}$$

and $H'(\hat{x}, \varepsilon) := \bigcup_{i=1}^n H'(\hat{x}, \varepsilon, i)$. This set $H'(\hat{x}, \varepsilon)$ covers the enhanced principal boundaries of $\hat{V}_{\hat{x}}$. The observation on which this proof is based (see Figure 3) can now be expressed by

$$H(\varepsilon) \subseteq \bigcup_{\hat{x} \in \hat{X}} H'(\hat{x}, \varepsilon),$$

which implies

$$F(s) \subseteq \bigcup_{\hat{x} \in \hat{X}'} H'(\hat{x}, \varepsilon) \cap (s + \frac{1}{N}\mathbb{Z}^n).$$

We count the number of s for which it may be the case that $H'(\hat{x}, \varepsilon) \cap (s + \frac{1}{N}\mathbb{Z}^n) \neq \emptyset$ for a fixed \hat{x} , and then multiply this by an upper bound on the number of elements in \hat{X}' . This allows us to obtain the formula from the theorem statement.

To obtain an upper bound on the cardinality $|\hat{X}'|$, we use **A2**). Since \hat{X}' is contained in at most $(\frac{q+A+1+2\varepsilon}{C}+1)^n$ blocks of size C , $|\hat{X}'| \leq D \cdot (\frac{q+A+1+2\varepsilon}{C}+1)^n$ by **A2**).

Now let $\hat{x} \in \mathbb{R}^n$ be arbitrary. As $\varepsilon \leq \frac{1}{2NL}$, there are at most $2L^{n-1}$ choices for $s \in S$ with $H'(\hat{x}, \varepsilon, i) \cap (s + \frac{1}{N}\mathbb{Z}^n) \neq \emptyset$. This shows that there are at most

$$2L^{n-1} \cdot n \cdot D \cdot \left(\frac{q + A + 1 + 2\varepsilon}{C} + 1 \right)^n$$

bad choices for $s \in S$, while $|S| = L^n$. This, together with $\varepsilon \leq \frac{1}{2NL}$, yields that the probability for $G(s) \cap H(\varepsilon) = \emptyset$ is at least

$$1 - \frac{1}{L} \cdot 2nD \left(\frac{q + A + 2 + C}{C} \right)^n. \quad \square$$

Corollary 4.3. *Let $N, q \in \mathbb{N}$ with $q, N \geq 1$ be given. Choose L and ε such that*

$$L \geq \frac{4nD(q + A + C + 2)^n}{C^n} \text{ and } \varepsilon \leq \frac{1}{2NL}. \quad (\text{I})$$

If we select $s \in S$ uniformly random, then

$$\Pr(G(s) \cap H^{\text{grid}}(\varepsilon) = \emptyset) \geq \frac{1}{2}.$$

This implies that we can compute $f(v)$ correctly for all $v \in \mathcal{V}$ and thus prepare the state

$$\frac{1}{\sqrt{V}} \sum_{v \in \mathcal{V}} |v\rangle |f(v)\rangle$$

in step 1 with probability greater or equal to $\frac{1}{2}$.

5 Preparing periodic states

The original function $\Phi_{\mathcal{T}}^{-1} \circ \pi : \mathbb{R}^n \rightarrow \mathbb{R}^n/\Lambda \rightarrow \text{Rep}^f(\mathcal{I})$ is perfectly periodic with period lattice Λ : if $u \in \mathbb{R}^n$ maps to $(x, t) \in \text{Rep}^f(\mathcal{I})$, then $u + \lambda$ will also map to (x, t) for all $\lambda \in \Lambda$.

Due to precision issues we have to work with the function $f : \mathbb{Z}^n \rightarrow \text{Rep}^f(\mathcal{I})$ defined by $v \mapsto (x, \lfloor Nt \rfloor)$ if $(\Phi_{\mathcal{T}}^{-1} \circ \pi)(s + \frac{1}{N}v) = (x, t)$. As $N\lambda$ will most certainly not have integral coordinates for $\lambda \in \Lambda$, we cannot directly obtain the collision $f(v) = f(v + N\lambda)$. And, if we round the coordinates of $N\lambda$ down or up to the nearest integers, it might happen that $f(v + \lfloor N\lambda \rfloor)$ yields an f -representation $(x', \lfloor Nt' \rfloor)$ with $x \neq x'$ – no matter to which integers we round the coordinates of $N\lambda$.

The first proposition of this section establishes a lower bound on the fraction of grid points for which this problem does not occur. For these grid points, the corresponding f -representation (x, t) is sufficiently far away from the boundaries, meaning that we remain in the same (translated) region $\hat{V}_{\hat{x}}$ when adding a suitably rounded version of $N\lambda$.

Similarly to the argument used in the proof of Proposition 4.2 in the previous section, we estimate the number of grid points lying in regions that are too close to the principal boundaries. The union of all such regions is denoted by H^{bound} . An example for $n = 2$ is shown in Figure 4 with H^{bound} highlighted.

Proposition 5.1. *Assume that $s \in S$ with $H(\varepsilon) \cap G(s) = \emptyset$ in the notation of the previous section. Consider $H^{\text{bound}} := (-\frac{1}{N}, 0]^n + H$. Then*

$$\frac{|G(s) \setminus H^{\text{bound}}|}{|G(s)|} \geq 1 - \frac{1}{N} \cdot \frac{nD(q+1+A+C)^n(A+2/N)^{n-1}}{(Cq)^n}.$$




Figure 4: The region H^{bound} is highlighted. For some of the cornered sets, the principal boundaries are shown.

Proof. Clearly $|G(s)| = (Nq)^n$. The cardinality of $G(s) \cap H^{\text{bound}}$ can be estimated by counting the number of cubes of the form $(-\frac{1}{N}, 0]^n$ needed to cover

$H^{\text{bound}} \cap [0, q]^n$. For $\hat{x} \in \hat{X}$, define

$$H''(\hat{x}, i) := \{\hat{x} + (t_1, \dots, t_n) \mid -\frac{1}{N} < t_j \leq A, -\frac{1}{N} < t_i \leq 0\}$$

and $\hat{X}'' := \hat{X} \cap [-A, q+1]^n$. Then

$$H^{\text{bound}} \subseteq \bigcup_{\hat{x} \in \hat{X}} \bigcup_{i=1}^n H''(\hat{x}, i) \quad \text{and} \quad G(s) \cap H^{\text{bound}} \subseteq \bigcup_{\hat{x} \in \hat{X}''} \bigcup_{i=1}^n H''(\hat{x}, i).$$

Now $H''(\hat{x}, i)$ can be covered by $\lceil NA+1 \rceil^{n-1}$ such cubes, whence the total number of cubes needed is less or equal to

$$|\hat{X}''| \cdot n \cdot \lceil NA+1 \rceil^{n-1}.$$

As above, \hat{X}'' is contained in at most $(\frac{q+1+A}{C} + 1)^n$ blocks of size C , whence

$$|\hat{X}''| \leq \frac{D}{C^n} (q+1+A+C)^n.$$

Therefore,

$$\frac{|G(s) \cap H^{\text{bound}}|}{|G(s)|} \leq \frac{\frac{D}{C^n} (q+1+A+C)^n \cdot n \cdot \lceil NA+1 \rceil^{n-1}}{(Nq)^n}. \quad \square$$

We now give an explicit lower bound on N guaranteeing that the fraction of grid points not contained in H^{bound} is sufficiently large.

Corollary 5.2. *We may assume w.l.o.g. that $C \leq 1$. Choose q and N such that*

$$q \geq 9 \max\{1, A\}, \quad \text{and} \quad N \geq \max \left\{ \frac{4}{A}, \frac{8(n+1)n \cdot 2^n D A^{n-1}}{3C^n} \right\}. \quad (\text{II})$$

If $s \in S$ is such that $H^{\text{grid}}(\varepsilon) \cap G(s) = \emptyset$, then

$$\frac{1}{N} \leq \frac{A}{4} \quad \text{and} \quad \frac{|G(s) \setminus H^{\text{bound}}|}{|G(s)|} \geq 1 - \frac{1}{4(n+1)}.$$

Note that we can always decrease C without invalidating assumption **A2**). Moreover, the recommended choice of C in Proposition 2.1 for infrastructures obtained from function fields or number fields satisfies $C \leq 1$.

Proof. Using $N \geq \frac{4}{A}$, the complement probability can be bounded by

$$\begin{aligned} & \frac{1}{N} \cdot \frac{nD(q+1+A+C)^n (A+2/N)^{n-1}}{(Cq)^n} \\ & \leq \frac{1}{N} \cdot \frac{nD(1+1/q+A/q+C/q)^n (A+A/2)^{n-1}}{C^n}. \end{aligned}$$

As $q \geq 9 \max\{1, A\}$, we have $1/q + A/q \leq \frac{2}{9}$, and as $C \leq 1$, we have $C/q \leq \frac{1}{9}$. Therefore, $1 + 1/q + A/q + C/q \leq 1 + \frac{1}{3} = \frac{4}{3}$, whence $(1 + 1/q + A/q + C/q)^n (1 + \frac{1}{2})^{n-1} \leq 2^n \cdot \frac{2}{3}$. Finally, the choice of N ensures that the complement probability is bounded by $\frac{1}{4(n+1)}$ from above. \square

Let $\mathcal{F} \subset X \times \mathbb{N}$ be the set of rounded f -representations. More precisely, it is defined by

$$(x, k) \in \mathcal{F} \quad \text{iff} \quad \text{there exists } v \in \mathcal{V} \text{ with } f(v) = (x, k) \text{ and } s + \frac{1}{N}v \notin H^{\text{bound}}.$$

Lemma 5.3. *Choose q and N according to (II). Assume $s \in S$ is such that $G(s) \cap H(\varepsilon) = \emptyset$. Let (x, k) be the measurement outcome obtained in step 2 of the quantum algorithm. Then,*

$$\Pr((x, k) \in \mathcal{F}) \geq 1 - \frac{1}{4(n+1)}.$$

Proof. For a fixed pair (x', k') , the probability that this pair is sampled is $\frac{1}{|\mathcal{V}|} |f^{-1}(x', k')|$. Let \mathcal{A} be the set of elements $v \in \mathcal{V}$ with $s + \frac{1}{N}v \notin H^{\text{bound}}$; then

$$\mathcal{A} \subseteq \bigcup_{(x', t') \in \mathcal{F}} f^{-1}(x', t'),$$

whence the probability we want to estimate can be bounded from below by $\frac{1}{|\mathcal{V}|} |\mathcal{A}|$. But this quantity equals $\frac{|G(s) \setminus H^{\text{bound}}|}{|G(s)|}$, and by Corollary 5.2 it can be bounded from below by $1 - \frac{1}{4(n+1)}$. \square

The next proposition makes precise statements on the periodicity of grid elements outside H^{bound} . First, we show that if $f(v) = f(v')$, then $\frac{1}{N}(v' - v)$ yields an approximation of some element $\lambda \in \Lambda$. Second, we show that for every $\lambda \in \Lambda$ such that $v + N\lambda$ stays within the boundaries of the grid there exists a unique v' with $f(v) = f(v')$ and $\frac{1}{N}(v' - v) \approx \lambda$. Finally, we estimate the number of collisions for one specific v , i.e., the numbers of v' in the grid such that $f(v) = f(v')$.

Proposition 5.4. *Choose q and N such that*

$$N \geq \frac{2\sqrt{n}}{\lambda_1(\Lambda)} \tag{III}$$

$$q > 2n\nu(\Lambda) + \frac{3n}{N}. \tag{IV}$$

Assume that $s \in S$ is such that $G(s) \cap H^{\text{grid}}(\frac{1}{2NL}) = \emptyset$. Let $v \in \mathcal{V}$ be such that $f(v)$ is equal to the measurement outcome. Assume that $s + \frac{1}{N}v \notin H^{\text{bound}}$. Let $\mathcal{M} = \{v' \in \mathcal{V} \mid f(v') = f(v)\}$ and $M = |\mathcal{M}|$.

- (i) *Let $v' \in \mathcal{M}$. We have $\|(v - v') - N\lambda\|_\infty < 1 - \frac{1}{L}$ for a unique $\lambda \in \Lambda$.*
- (ii) *Let $\lambda \in \Lambda$ such that $v + N\lambda \in [1, qN - 2]^n$. Then, there exists a unique $v' \in \mathcal{M}$ satisfying $\|(v - v') - N\lambda\|_\infty < 1 - \frac{1}{L}$.*
- (iii) *We have $M \geq M_\ell$, where*

$$M_\ell = \frac{q^n}{\det(\Lambda)} \left(1 - \frac{3n}{qN} - \frac{2n\nu(\Lambda)}{q} \right).$$

Proof.

- (i) Let $(\Phi_{\mathcal{I}}^{-1} \circ \pi)(s + \frac{1}{N}v) = (x, t)$ and $(\Phi_{\mathcal{I}}^{-1} \circ \pi)(s + \frac{1}{N}v') = (x', t')$; then $f(v) = (x, \lfloor Nt \rfloor)$ and $f(v') = (x', \lfloor Nt' \rfloor)$. Note that in \mathbb{R}^n/Λ , we have $d(x) + t = s + \frac{1}{N}v$ and $d(x') + t' = s + \frac{1}{N}v'$, whence $d(x) + t - (d(x') + t') = \frac{1}{N}(v - v')$.

We have $f(v) = f(v')$. Therefore, $x = x'$ and $\lfloor Nt \rfloor = \lfloor Nt' \rfloor$, which yields $\|t - t'\|_{\infty} < \frac{1}{N}$. By the assumption that $G(s) \cap H^{\text{grid}}(\frac{1}{2NL}) = \emptyset$, we have that the coefficients and Nt and Nt' are bounded away from an integer by at least $\frac{1}{2L}$ (compare Corollary 4.3 2), whence we actually have $\|t - t'\|_{\infty} < \frac{1}{N} - \frac{1}{NL}$.

Now $t - t' = d(x) + t - (d(x') + t') = \frac{1}{N}(v - v')$ in \mathbb{R}^n/Λ , whence there exists some $\lambda \in \Lambda$ such that $v - v' = N(t - t') + N\lambda$.

- (ii) Let $(\Phi_{\mathcal{I}}^{-1} \circ \pi)(s + \frac{1}{N}v) = (x, t)$; then $f(v) = (x, \lfloor Nt \rfloor)$ and $d(x) + t = s + \frac{1}{N}v$. Set $u := v + N\lambda$; then $d(x) + t = s + \frac{1}{N}u$ as an element of \mathbb{R}^n/Λ , whence $(x, t) = (\Phi_{\mathcal{I}}^{-1} \circ \pi)(s + \frac{1}{N}u)$.

There are at most two choices for each coordinate of the vector $e \in (-1, 1)^n$ such that $u + e$ has only integral coefficients. For each coordinate, there is exactly one choice if only 0 can be chosen; otherwise, there exists one choice $a \in (-1, 0)$ and the other is $1 + a$. Hence, there exists a unique $e \in (-1, 1)^n$ such that $\lfloor Nt \rfloor = \lfloor Nt + e \rfloor$ and $v' := u + e \in \mathbb{Z}^n$.

Clearly, $t + \frac{1}{N}e \geq 0$. First, $(x, t + \frac{1}{N}e) \in \text{Rep}^f(\mathcal{I})$ since $s + \frac{1}{N}v \notin H^{\text{bound}}$. Second, $d(x) + (t + \frac{1}{N}e) = s + \frac{1}{N}v'$ implies $f(v') = (x, \lfloor Nt + e \rfloor) = (x, \lfloor Nt \rfloor) = f(v)$. Third, $v' \in \mathcal{V}$ since $u \in [1, qN - 2]^n$.

It remains to show that v' is unique. Assume that $v', v'' \in \mathcal{V}$ satisfy $f(v') = f(v'')$, $\|(v - v') - N\lambda\|_{\infty} < 1 - \frac{1}{L}$, and $\|(v - v'') - N\lambda\|_{\infty} < 1 - \frac{1}{L}$.

By (i) of this proposition, the condition $f(v) = f(v')$ implies that there exists some $\lambda' \in \Lambda$ with $\|(v' - v'') - N\lambda'\|_{\infty} < 1$. By the triangle inequality, the two above conditions on the norms imply that $\|v' - v''\|_{\infty} < 2$. Since $v' - v'' \in \mathbb{Z}^n$, this yields $\|v' - v''\|_{\infty} \leq 1$.

By applying the triangle inequality again and dividing by N , we conclude that $\|\lambda'\|_{\infty} < \frac{2}{N}$. Now, if $v' \neq v''$, then $\|(v' - v'') - N\lambda'\|_{\infty} < 1$ would imply that $\lambda' \neq 0$. Then, $0 < \|\lambda'\|_2 < \sqrt{n} \cdot \frac{2}{N}$ would hold. But, this would violate $\lambda_1(\Lambda) \geq \frac{2\sqrt{n}}{N}$, which follows from (III). Therefore, we must have $v' = v''$ and, thus, v' is unique.

- (iii) Using (ii), we see that a lower bound M_{ℓ} on M is given by the cardinality of $N\Lambda \cap (-v + [1, qN - 2]^n)$. Let $\nu(N\Lambda)$ be the covering radius of $N\Lambda$. Let $\lambda \in N\Lambda$. If $\lambda \in (-v + [1 + \nu(N\Lambda), qN - 2 - \nu(N\Lambda)]^n)$, then the Voronoi cell $\hat{V}_{N\Lambda}(\lambda)$ of λ is entirely contained in $(-v + [1, qN - 2]^n)$. As the volume of $\hat{V}_{N\Lambda}(\lambda)$ is $\det(N\Lambda)$, this yields the lower bound

$$\frac{(qN - 3 - 2\nu(N\Lambda))^n}{\det(N\Lambda)} \geq \frac{q^n}{\det(\Lambda)} \left(1 - \frac{3n}{qN} - \frac{2n\nu(\Lambda)}{q}\right),$$

which is greater than 0 provided that assumption (III) holds. \square

6 Sampling approximations of vectors of the dual lattice Λ^*

6.1 Sampling in dimension greater than one

We present here our new method of sampling approximation of the vectors of the dual lattice Λ^* , which improves the success probability of the overall algorithm by at least the exponential factor 2^{n^2-1} .

We determine the probability that the quantum algorithm outputs a $w \in \mathcal{W}$ such that $\frac{1}{2nq}w$ is sufficiently close to some $\lambda^* \in \Lambda^*$. We have to impose certain conditions on w to be able to show that the probability of observing a good approximation is bounded away from 0. For $\lambda^* \in \Lambda^*$, let

$$\mathcal{R}_{\lambda^*} = \left\{ (w_1, \dots, w_n) \mid w_k \in \{ \lfloor 2nq\lambda_k^* \rfloor, \lfloor 2nq\lambda_k^* \rfloor + 1 \} \text{ for } k = 1, \dots, n \right\}.$$

Observe that for all $w \in \mathcal{R}_{\lambda^*}$, we have

$$\left\| \frac{w}{2nq} - \lambda^* \right\|_2 \leq \frac{1}{2\sqrt{nq}}.$$

The following proposition gives a lower bound on the probability of observing elements of \mathcal{R}_{λ^*} provided that $\mathcal{R}_{\lambda^*} \subset [0, 2nq\kappa N]^n$, where $\kappa \in (0, 1)$.

In the remainder of this section, we make the two following assumptions:

- (i) the random shift $s \in S$ is such that $G(s) \cap H^{\text{grid}}(\frac{1}{2NL}) = \emptyset$ and
- (ii) all measurement outcomes $f(v)$ are such that $s + \frac{v}{N} \notin H^{\text{bound}}$.

The relevant results can be stated in a more direct way if we do not have to include these two assumptions in the formulation of the propositions. Note that we can estimate the probabilities that they are satisfied with the help of Corollary 4.3 and Lemma 5.3. These will be included in the final analysis of the algorithm.

Proposition 6.1. *Choose q and N according to (III) and (IV). Choose κ such that*

$$\kappa < \frac{1}{8n} - \frac{1}{4nqN}. \quad (\text{V})$$

Then, for all $\lambda^ \in \Lambda^*$ with $\mathcal{R}_{\lambda^*} \subset [0, 2qn\kappa N]^n$, we have the lower bound*

$$\begin{aligned} \Pr(\mathcal{R}_{\lambda^*}) &= \sum_{w \in \mathcal{R}_{\lambda^*}} \Pr(w) \\ &= \sum_{w \in \mathcal{R}_{\lambda^*}} \left| \frac{1}{\sqrt{MW}} \sum_{v' \in \mathcal{M}} \exp\left(2\pi i v' \cdot \frac{w}{2nqN}\right) \right|^2 \\ &\geq \frac{2^{n-1}M_\ell}{W} \cos^2\left(\pi\left(\frac{1}{4} + \frac{1}{2qN} + 2\kappa n\right)\right). \end{aligned}$$

Proof. Let v' be an arbitrary but fixed element of \mathcal{M} . Proposition 5.4 (i) shows that $\|v' - v - N\lambda\|_\infty < 1 - \frac{1}{L}$ for some $\lambda \in \Lambda$ since condition (III) is satisfied. Define the error terms $e_1(v') = v' - v - N\lambda$ and $e_2(w) = \frac{w}{2nqN} - \frac{\lambda^*}{N}$ for $w \in \mathcal{R}_{\lambda^*}$. Both error types arise because both the rescaled lattice $N\Lambda$ and the dual lattice Λ^* are not necessarily integral.

To be able to show that the probability of observing a $w \in \mathcal{R}_{\lambda^*}$ is bounded away from zero by a constant, we have (i) to carry out the Fourier transform over a larger window and (ii) to disregard w whose infinity-norm is too large. These two measures makes it possible to mitigate the effects of the first and second errors, respectively. Unfortunately, both measures are also responsible for the exponentially decreasing success probability with increasing dimension n .

To understand the effects of these error terms, we expand the inner product $v' \cdot \frac{w}{2nqN}$ as follows

$$\begin{aligned} v' \cdot \frac{w}{2nqN} &= (v + N\lambda + e_1(v')) \cdot \frac{w}{2nqN} \\ &= (v + N\lambda) \cdot \frac{w}{2nqN} + e_1(v') \cdot \frac{w}{2nqN} \\ &= (v + N\lambda) \cdot \frac{\lambda^*}{N} + (v + N\lambda) \cdot e_2(w) + e_1(v') \cdot \frac{w}{2nqN} \\ &= v \cdot \frac{\lambda^*}{N} + \lambda \cdot \lambda^* + (v + N\lambda) \cdot e_2(w) + e_1(v') \cdot \frac{w}{2nqN}. \end{aligned}$$

Since $v \cdot \frac{\lambda^*}{N}$ is constant and $\lambda \cdot \lambda^* \in \mathbb{Z}$, we only have to consider the inner products $e_1(v') \cdot \frac{w}{2nqN}$ and $(v + N\lambda) \cdot e_2(w)$.

Using the upper bound $\|e_1(v')\|_\infty \leq 1 - \frac{1}{L}$, the absolute value of the first error term is seen to be bounded from above by

$$\left| e_1(v') \cdot \frac{w}{2nqN} \right| \leq \frac{n}{2nqN} \|e_1(v')\|_\infty \|w\|_\infty \leq \frac{1}{2qN} \left(1 - \frac{1}{L}\right) 2nq\kappa N < \kappa n.$$

To bound the norm of the second error term, we set

$$p_k = 2nq\lambda_k^* - \lfloor 2nq\lambda_k^* \rfloor$$

for $k = 1, \dots, n$. In words, the values p_k correspond to the errors caused by rounding down the coefficients of $2nq\lambda^*$ to the nearest integer. Set

$$A = \{k : w_k = \lfloor 2nq\lambda_k^* \rfloor\} \quad \text{and} \quad \bar{A} = \{\ell : w_\ell = \lfloor 2nq\lambda_\ell^* \rfloor + 1\}.$$

Observe that for $k \in A$ the k th coefficient of the error vector $e_2(w) := \frac{w}{2nqN} - \frac{\lambda^*}{N}$ is equal to $\frac{-p_k}{2nqN}$ and for $\ell \in \bar{A}$ the ℓ th coefficient is equal to $\frac{1-p_\ell}{2nqN}$. Set

$$L_A = \frac{1}{2n} \left(\sum_{k \in A} p_k + \sum_{\ell \in \bar{A}} (1 - p_\ell) \right),$$

which is equal to $qN\|e_2(w)\|_1$.

Since $v + N\lambda \in [-1 + \frac{1}{L}, qN - \frac{1}{L}]^n \subset (-1, qN)^n$, we have

$$\begin{aligned} -\frac{1}{2n} \sum_{k \in A} p_k - \frac{1}{2nqN} \sum_{\ell \in \bar{A}} (1 - p_\ell) &\leq (v + N\lambda) \cdot e_2(w) \\ \frac{1}{2n} \sum_{\ell \in \bar{A}} (1 - p_\ell) + \frac{1}{2nqN} \sum_{k \in A} p_k &\geq (v + N\lambda) \cdot e_2(w) \end{aligned}$$

Therefore, the sum $(v + N\lambda) \cdot e_2(w) + e_1(v') \cdot \frac{w}{2nqN}$ of both error terms ranges over an interval of length at most

$$L_A + \frac{1}{2qN} + 2\kappa n.$$

Clearly, the identity

$$L_{\bar{A}} = \frac{1}{2} - L_A$$

holds for all $A \subseteq \{1, \dots, n\}$. This simple fact implies the crucial inequality

$$\min\{L_A, L_{\bar{A}}\} \leq \frac{1}{4}.$$

The latter holds because otherwise we would have $L_A > \frac{1}{4}$ and $L_{\bar{A}} = \frac{1}{2} - L_A > \frac{1}{4}$, which would lead to the contradiction $\frac{1}{2} > \frac{1}{2}$.

In the remainder of the proof, without loss of generality A always denotes a subset of $\{1, \dots, n\}$ with $L_A \leq \frac{1}{4}$.

Let A be such subset and w the corresponding approximation of $2nqN\lambda^*$. This means that the sum we want to estimate can be written as

$$\sum_{v' \in \mathcal{M}} \exp(2\pi i(\alpha + \beta_{v'})) = \exp(2\pi i\alpha) \sum_{v' \in \mathcal{M}} \exp(2\pi i\beta_{v'})$$

with $\alpha, \beta_{v'} \in \mathbb{R}$ and $-\frac{1}{2}L_{\text{phase}} \leq \beta_{v'} \leq \frac{1}{2}L_{\text{phase}}$, where $L_{\text{phase}} = L_A + \frac{1}{2qN} + 2\kappa n$. Hence, the real part of every term $\exp(2\pi i\beta_{v'})$ is $\cos(2\pi\beta_{v'}) \geq \cos(\pi L_{\text{phase}})$ since $L_{\text{phase}} < \frac{1}{2}$ due to $L_A \leq \frac{1}{4}$ and the special choice of κ in (IV).

This implies that the absolute value of the sum is bounded from below by $M \cos(\pi L_{\text{phase}})$ for this particular w . Finally, we obtain the desired claim

$$\begin{aligned} \Pr(\mathcal{R}_{\lambda^*}) &\geq \frac{M}{W} \sum_{A: L_A \leq \frac{1}{4}} \cos^2\left(\pi\left(L_A + \frac{1}{2qN} + 2\kappa n\right)\right) \\ &\geq \frac{2^{n-1}M}{W} \cos^2\left(\pi\left(\frac{1}{4} + \frac{1}{2qN} + 2\kappa n\right)\right) \end{aligned}$$

by noting that there are at least 2^{n-1} subsets A with $L_A \leq \frac{1}{4}$. □

6.2 Sampling in dimension one

Remark 6.2 (One-dimensional infrastructures). In the special case of one-dimensional infrastructures, it is better to work with the sets

$$\mathcal{R}_{\lambda^*} = \{w \mid w = [2q\lambda^*]\}$$

for $\lambda^* \in \Lambda^*$. This is because we may then choose a slightly larger κ . The upper bound can be increased to

$$\kappa < \frac{1}{8} - \frac{1}{8qN},$$

which leads to the higher lower bound on the success probability

$$\Pr(\mathcal{R}_{\lambda^*}) \geq \frac{M}{W} \cos^2\left(\pi\left(\frac{1}{4} + \frac{1}{4qN} + 2\kappa n\right)\right).$$

This bound is established by using the same arguments as in the proof of the above proposition and by observing that the upper bound on $|e_2(w)|$ is reduced by a factor of 2. The latter statement is due to the fact that for all $\lambda^* \in \Lambda^*$, we have the better approximation

$$\left|\frac{w}{2q} - \lambda^*\right| \leq \frac{1}{4q},$$

where $w \in \mathcal{R}_{\lambda^*}$.

7 Lattice theoretic tools – Part 1

7.1 Lattices of dimension greater than one

We now show how to obtain a generating set of a full-rank lattice L in \mathbb{R}^n by first sampling n lattice vectors that are contained in the window $[0, b)^n$ and then $n + 1$ lattice vectors that are contained in the larger window $[0, b_0)^n$. If we chose b to be a sufficiently larger than the covering radius of L , then the first n lattice vectors generate a full-rank sublattice L_0 of L with probability greater or equal to $\frac{1}{4}$ (Subsection 7.1.1). Once we have such sublattice L_0 , the next $n + 1$ lattice vectors that we sample from the larger window $[0, b_0)$ generate together with the first n vectors the entire lattice L with probability greater or equal to $\hat{\zeta} - \frac{1}{4} \geq 0.184$, where $\hat{\zeta}$ is a certain constant (Subsection 7.1.3).

Our current proof requires that we use two windows. We think that it is possible to prove a similar result, while relying only on one window.

Note that these results will be used with $L = \Lambda^*$ throughout the rest of the paper.

7.1.1 Probability of generating a full-rank sublattice L_0 of L

Let L be a lattice in \mathbb{R}^n of full rank. For $\lambda \in L$, let $V_L(\lambda)$ be its (open) Voronoi cell. We know that $V_L(\lambda)$ is contained in an open sphere of radius $\nu(L)$ centered around λ , where $\nu(L)$ is the covering radius of L , and that the volume of $V_L(\lambda)$ is $\det(L)$. Moreover, if $\lambda \neq \lambda'$, $V_L(\lambda) \cap V_L(\lambda') = \emptyset$, and $\bigcup_{\lambda \in L} \overline{V_L(\lambda)} = \mathbb{R}^n$.

Lemma 7.1. *If $b > 2\nu(L)$. Then*

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

Proof. If $\lambda \in L$ satisfies $V_L(\lambda) \cap [\nu, b - \nu]^n \neq \emptyset$, then we must have $\lambda \in [0, b]^n$. Therefore, $(b - 2\nu)^n / \det(L) \leq |L \cap [0, b]^n|$.

If $\lambda \in L \cap [0, b]^n$, then we must have $V_L(\lambda) \subseteq [-\nu, b + \nu]^n$. Therefore, $|L \cap [0, b]^n| \leq (b + 2\nu)^n$. \square

Lemma 7.2. *Let $b > 0$ and H be a k -dimensional hyperplane, $1 \leq k < n$. Then*

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

Proof. Let $\lambda \in L \cap H \cap [0, b]^n$. Then $V_L(\lambda) \subseteq X := [-\nu, b + \nu]^n \cap (H + B_\nu(0))$, where $B_\nu(0)$ is a sphere of radius ν centered around 0. Therefore, $|L \cap H \cap [0, b]^n| \leq \text{vol}(X) / \det(L)$, and we have to estimate $\text{vol}(X)$.

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

Let b_1, \dots, b_k be an orthonormal basis of H . Set $T := \{(x_1, \dots, x_k) \in \mathbb{R}^k \mid \sum_{i=1}^k x_i b_i \in [-\nu, b + \nu]^n\}$; 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 $y = (y_1, \dots, y_n) \in [-\nu, b + \nu]^n$, set $A_{ij} := b + \nu$ if $b_{ij} \geq 0$ and $A_{ij} := \nu$ if $b_{ij} < 0$. Then

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

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

$$\text{vol}(T) \leq n^{k/2} (b + 2\nu)^k.$$

\square

Corollary 7.3. *Assume that $b \geq \max\{8n - 2, n^{(n-1)/2} 2^{n+1} - 2\} \cdot \nu(L)$. Let*

$$X := (L \cap [0, b]^n)^n$$

and $Y := \{(y_1, \dots, y_n) \in X \mid \text{span}_{\mathbb{R}}(y_1, \dots, y_n) = \mathbb{R}^n\}$.

Then

$$|Y| > 0.289 |X| > \frac{1}{4} |X|.$$

Note that $\max\{8n - 2, n^{(n-1)/2} 2^{n+1} - 2\} = n^{(n-1)/2} 2^{n+1} - 2$ unless $n \leq 2$, in which case the maximum is $8n - 2$.

The proof of this corollary is similar to the proof of the first part of Satz 2.4.23 in [Sch07]. Note that the proof in [Sch07] is not correct: the quantity $\frac{|M_i \cap \mathcal{B}|}{|M_{i-1} \cap \mathcal{B}|}$ in

the proof can be $> \frac{1}{2}$; for example, consider $r = 3$, $M = \mathbb{Z}^3$, $n > 0$ arbitrary (in [Sch07], $n\nu(M)$ is what we denote by b , i.e., $\mathcal{B} = [0, n\nu(M))^n$, $x_1 = (1, n\nu(M) - 1, -1)$, $x_2 = (0, 1, n\nu(M) - 1)$, $x_3 = (0, 0, 1)$; then $M_1 \cap \mathcal{B}$ contains three elements, while $M_2 \cap \mathcal{B}$ contains five elements. The problem is that $\det(M_i)$ cannot be bounded in terms of $\nu(M)$ and $\det(M_{i-1})$, as it was claimed in that proof. We proceed differently by considering the quantity $\frac{|M_i \cap \mathcal{B}|}{|M \cap \mathcal{B}|}$ directly, and our bound on the minimal size of \mathcal{B} is in fact better than the bound given in [Sch07].

Also, note that for specific small n , one can obtain better bounds of $|Y|$ in term of $|X|$. As the proof will show, a lower bound on $|Y|$ is given by $|X| \cdot \prod_{i=1}^{n-1} (1 - 2^{-i})$. The following table gives explicit values for this factor for small values of n , rounded down to a precision of 10^{-3} :

n	2	3	4	5	6
$\prod_{i=1}^{n-1} (1 - 2^{-i})$	0.500	0.375	0.328	0.307	0.298

Proof. Assume that $y_1, \dots, y_k \in X$ are linearly independent. We have to compute the probability 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 $b = j\nu(L)$ with $j \geq n^{(n-1)/2}2^{n+1} - 2$. By the above lemmata, the probability that y_{k+1} is in a k -dimensional hyperplane is bounded from above by

$$\begin{aligned} P_k &:= \frac{n^{k/2}(b + 2\nu)^k(2\nu)^{n-k}}{\det(L)} \cdot \frac{\det(L)}{(b - 2\nu)^n} \\ &= \frac{n^{k/2}(b + 2\nu)^k(2\nu)^{n-k}}{(b - 2\nu)^n} = n^{k/2} \frac{(j + 2)^k 2^{n-k}}{(j - 2)^n}. \end{aligned}$$

We now prove that $P_k \leq 2^{-k}$ holds, which is equivalent to

$$n^{k/2}(j + 2)^k 2^n \leq (j - 2)^n.$$

Clearly, the left-hand side is maximal for $k = n - 1$, giving the strictest condition

$$n^{(n-1)/2} 2^n \leq (j + 2) \left(\frac{j - 2}{j + 2} \right)^n.$$

The right-hand side is bounded from below by $(j + 2)/2$ provided that $j \geq 8n - 2$ (this follows from Bernoulli's inequality). Hence, the above condition is satisfied for $j \geq n^{(n-1)/2} 2^{n+1} - 2$.

To conclude the proof, note that the probability we look for is therefore bounded from below by

$$\prod_{i=1}^{n-1} (1 - 2^{-i}) \geq \prod_{i=1}^{\infty} (1 - 2^{-i}) > 0.289 > \frac{1}{4},$$

where the last two inequalities follows by Euler's Pentagon Number Theorem. \square

7.1.2 Probability of generating finite abelian groups

Proposition 7.4. *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.

Note that for small n , better lower bounds on the probability can be obtained. If G can be created by n elements, then a better lower bound is $\prod_{i=2}^{n+1} \zeta(i)^{-1}$; this is always larger than $\hat{\zeta}$. The following table gives explicit values for this product for small values of n , rounded down to a precision of 10^{-3} :

n	2	3	4	5	6
$\prod_{i=2}^{n+1} \zeta(i)^{-1}$	0.505	0.467	0.450	0.442	0.439

Proof. Let p_1, \dots, p_k be the prime divisors of $|G|$, and let G_i be the p_i -Sylow subgroup of G . Then $G = G_1 \oplus \dots \oplus G_k$. Let $(g_1, \dots, g_{n+1}) \in G^{n+1}$ be $n + 1$ elements of G ; then we can write $g_i = (g_{i1}, \dots, g_{ik}) \in G_1 \times \dots \times G_k$. Now

$$G = \langle g_1, \dots, g_{n+1} \rangle \iff \forall j : G_j = \langle g_{1j}, \dots, g_{n+1,j} \rangle.$$

Hence, it suffices to bound the probability for abelian p -groups.

In the proof of the theorem in [Pom01], it is shown that the probability that $n + 1$ elements in an abelian p -group of p -rank r generate the group is

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

We know that $r \leq n$, since G is generated by n elements.

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

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

The product $\prod_{i=2}^{\infty} \zeta(i)$ is well-known in group theory [Seq]. □

Note that it is essential for our proof to work that we use $n + 1$ elements instead of n , since if we choose 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 non-zero probability for n elements, but this probability will not be constant anymore, but 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 if $k \rightarrow \infty$ for exactly the above reasons. Hence, any non-trivial bound of the probability must take n or p_1, \dots, p_k into account.

7.1.3 Probability of generating the entire lattice L

Lemma 7.5 (Sampling almost uniformly at random from L/L_0). *Let L_0 be an arbitrary full-rank sublattice of L . Assume that $b_0 > 2\nu(L_0)$ and we can sample uniformly at random from*

$$L \cap [0, b_0)^n.$$

Denote the sample by λ . Then, $\lambda + L_0$ is distributed almost uniformly at random over the quotient group L/L_0 . More precisely, the total variation distance between the uniform distribution is at most

$$1 - \frac{(b_0 - 2\nu(L_0))^n}{(b_0 + 2\nu(L))^n}.$$

Proof. Let again $V_{L_0}(\lambda_0)$ denote the open Voronoi cell of the lattice L_0 centered around λ_0 . First note that $V_{L_0}(\lambda_0) = \lambda_0 + V_{L_0}(0)$ and $\overline{V_{L_0}(\lambda_0)} = \lambda_0 + \overline{V_{L_0}(0)}$. Now, as $\bigcup_{\lambda_0 \in L_0} (\lambda_0 + \overline{V_{L_0}(0)}) = \mathbb{R}^n$ and two translates of $V_{L_0}(0)$ by different elements of L_0 do not intersect, there exists a set V with $V_{L_0}(0) \subseteq V \subseteq \overline{V_{L_0}(0)}$ satisfying

$$\bigcup_{\lambda_0 \in L_0} (\lambda_0 + V) = \mathbb{R}^n \quad \text{and} \quad \forall \lambda_0 \in L_0 \setminus \{0\} : (\lambda_0 + V) \cap V = \emptyset.$$

Note that $\text{vol}(V) = \text{vol}(V_{L_0}(0)) = \det(L_0)$.

Every translate of V contains the same number of elements from L , and $|V \cap L|$ equals

$$m = \det(L_0) / \det(L);$$

this can be shown using asymptotic arguments similarly to the proof that any elementary parallelepiped of L_0 contains exactly m elements of L (see e.g. [Bar]).

For all $\lambda \in L \cap V$, the vectors $\lambda - \lambda_0$ form a transversal for L/L_0 .

As $V \subseteq \overline{B_{\nu(L_0)}(0)}$, there are at least

$$\ell_V = \frac{(b_0 - 2\nu(L_0))^n}{\det(L_0)}$$

translates of V that are contained inside the window $[0, b_0]^n$.

There are at most

$$u_P = \frac{(b_0 + 2\nu(L))^n}{\det(L)}$$

points of L inside $[0, b_0]^n$.

Let $d_{\max} = \lfloor u_P - m\ell_V \rfloor$ be the maximal possible deviation in the number of points of L inside $[0, b_0]^n$ from the lower bound $m\ell_V$. Let $d \in \{0, \dots, d_{\max}\}$ be the actual deviation.

Ideally, we would have the uniform distribution $p_j = 1/m$ on L/L_0 . But we only have the almost uniform distribution which necessarily has the form

$$\tilde{p}_j = \frac{\ell_V + d_j}{m\ell_V + d}$$

for $j = 1, \dots, m$, where d_1, \dots, d_m are integers with $0 \leq d_j \leq d$ and $\sum_{j=1}^m d_j = d$. The total variation distance can be bounded as follows

$$\begin{aligned} \frac{1}{2} \sum_{j=1}^m |p_j - \tilde{p}_j| &= \frac{1}{2} \sum_{j=1}^m \left| \frac{1}{m} - \frac{\ell_V + d_j}{m\ell_V + d} \right| \\ &= \frac{1}{2m} \sum_{j=1}^m \left| \frac{d - md_j}{m\ell_V + d} \right| \\ &\leq \frac{1}{2m} \sum_{j=1}^m \frac{d + md_j}{m\ell_V + d} \\ &= \frac{d}{m\ell_V + d} \\ &\leq \frac{d_{\max}}{m\ell_V + d_{\max}} \leq \frac{u_P - m\ell_V}{m\ell_V + u_P - m\ell_V} = 1 - \frac{m\ell_V}{u_P}. \end{aligned}$$

We have

$$1 - \frac{m\ell_V}{u_P} = 1 - \frac{(b_0 - 2\nu(L_0))^n}{(b_0 + 2\nu(L))^n}.$$

Note that so far, we have considered $[0, b_0]^n$ instead of $[0, b_0)^n$. As L is discrete, there exists some $2\nu(L_0) < b'_0 < b_0$ with $[0, b'_0]^n \cap L = [0, b_0)^n$. Applying the result above to $[0, b'_0]^n$ and then using that $x \mapsto 1 - \frac{(x - 2\nu(L_0))^n}{(x + 2\nu(L))^n}$ is increasing yields the stated claim for $[0, b_0)^n$. \square

Proposition 7.6. *Assume that $b \geq \max\{8n - 2, n^{(n-1)/2}2^{n+1} - 2\} \cdot \nu(L)$ and $b_0 \geq 8n^2(n+1)b$. Let Y be as in Corollary 7.3 and $(y_1, \dots, y_n) \in Y$. Let*

$$\begin{aligned} X_0 &:= (L \cap [0, b_0)^n)^{n+1} \\ Z &= \{(z_1, \dots, z_{n+1}) \in X_0^{n+1} \mid \text{span}_{\mathbb{Z}}\{y_1, \dots, y_n, z_1, \dots, z_{n+1}\} = L\}. \end{aligned}$$

Then

$$|Z| \geq \left(\hat{\zeta} - \frac{1}{4} \right) |X_0| \geq 0.184 |X_0|.$$

Proof. Let L_0 be the full-rank sublattice generated by y_1, \dots, y_n . We have the following simple bound on the covering radius

$$\nu(L_0) \leq \frac{\sqrt{n}}{2} \lambda_n(L_0) \leq \frac{\sqrt{n}}{2} \max_{i=1, \dots, n} \|y_i\|_\infty \leq \frac{\sqrt{n}}{2} \sqrt{nb} = \frac{nb}{2}$$

since the y_i are linearly independent and the longest vector in $[0, b]^n$ is shorter than \sqrt{nb} .

Let z_i be uniformly distributed in $L \cap [0, b_0]^n$. Then, Lemma 7.5 implies that $z_i + L_0$ (for $i = n+1, \dots, 2n+1$) are distributed almost uniformly at random from L/L_0 . The total variation distance from the uniform distribution is bounded from above as follows

$$\begin{aligned} 1 - \frac{(b_0 - 2\nu(L_0))^n}{(b_0 + 2\nu(L_0))^n} &\leq 1 - \frac{(b_0 - 2\nu(L_0))^n}{(b_0 + 2\nu(L_0))^n} \\ &= 1 - \left(1 - \frac{4\nu(L_0)}{b_0 + 2\nu(L_0)} \right)^n \\ &\leq 1 - \left(1 - n \frac{4\nu(L_0)}{b_0 + 2\nu(L_0)} \right) \\ &\leq \frac{4n\nu(L_0)}{b_0} \leq \frac{2n^2b}{b_0} \leq \frac{1}{4(n+1)}. \end{aligned}$$

Consider now the uniform probability distribution on the $(n+1)$ -fold direct product of L/L_0 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 bound from above by $(n+1) \cdot \frac{1}{4(n+1)} = \frac{1}{4}$. This is because total variation distance is additive under composition provided that the components are independent (see e.g. [MG02, Subsection 1.3 “Statistical distance” in Chapter 7] for more information total variation distance).

Clearly, the abelian group L/L_0 can be generated with only n generators. Hence, Proposition 7.4 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 L/L_0 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

Remark 7.7. The purpose of this proposition is similar to that of Satz 2.4.23 in [Sch07]. 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 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 7.3. 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 non-zero success probability.

Note that in [Hal05], neither a bound is given on how many lattice elements have to be sampled nor the probability is estimated with which the lattice is generated.

Lemma 7.8. *Assume*

$$b \geq \max\{8n - 2, n^{(n-1)/2}2^{n+1} - 2\} \cdot \frac{n}{2\lambda_1(\Lambda)} \quad \text{and}$$

$$b_0 \geq 8n^2(n+1)b.$$

Define

$$X := (\Lambda^* \cap [0, b]^n)^n$$

$$Y := \{(\lambda_1^*, \dots, \lambda_n^*) \in X \mid \text{span}_{\mathbb{R}}(\lambda_1^*, \dots, \lambda_n^*) = \mathbb{R}^n\}.$$

For each $(\lambda_1^, \dots, \lambda_n^*) \in Y$, define*

$$X_0 := (\Lambda^* \cap [0, b_0]^n)^{n+1}$$

$$Z := \{(\lambda_{n+1}^*, \dots, \lambda_{2n+1}^*) \in X_0 \mid \text{span}_{\mathbb{Z}}(\lambda_1^*, \dots, \lambda_n^*, \lambda_{n+1}^*, \dots, \lambda_{2n+1}^*) = L\}.$$

Then

$$|Y| \geq 0.289 |X| > \frac{1}{4} |X| \quad \text{and} \quad |Z| \geq (\hat{\zeta} - \frac{1}{4}) |X_0| \geq 0.184 |X_0|$$

Proof. The first lower bound follows from Corollary 7.3 and the inequality $\nu(\Lambda^*) \leq \frac{n}{2\lambda_1(\Lambda)}$ and the second from Proposition 7.6. \square

By combining the more precise bounds listed below Corollary 7.3 and Proposition 7.6, respectively, one obtains the following more precise bounds which depend on n :

$$|Y| \geq |X| \cdot \prod_{i=1}^{n-1} (1 - 2^{-i}) \quad \text{and} \quad |Z| \geq \left(\prod_{i=2}^{n+1} \zeta(i)^{-1} - \frac{1}{4} \right) \cdot |X_0|. \quad (*)$$

7.2 Lattices of dimension one

We now discuss the special case $n = 1$. For this case, $2n$ instead of $2n + 1$ vectors from one window suffice to generate the lattice with a significantly higher probability.

Lemma 7.9. *Let $L = \mathbb{Z}v$ be a one-dimensional lattice, where $v \in \mathbb{R}_{>0}$. Assume that $b \geq 3v + 1$. Then, two samples chosen uniformly at random in $L \cap [0, b)$ generate L with probability greater than $\frac{3^3}{\pi^2 2^3} > \frac{1}{3}$. Note that $\det(L) = v = \lambda_1(L)$, $\nu(L) = \frac{1}{2} \det(L)$ and that $L^* = \frac{1}{v}\mathbb{Z}$.*

Proof. Clearly, the number of lattice elements in $[0, b - 1]$ is $1 + \lfloor \frac{b-1}{v} \rfloor$, where 1 accounts for the zero vector. Hence, the probability that a random element of $L \cap [0, b - 1]$ is non-zero is

$$\frac{\lfloor \frac{b-1}{v} \rfloor}{1 + \lfloor \frac{b-1}{v} \rfloor} = 1 - \frac{1}{1 + \lfloor \frac{b-1}{v} \rfloor},$$

which is greater or equal to $\frac{3}{4}$ for $b \geq 3v + 1$. Further, note that this condition ensures that there are at least 3 non-zero elements. Assume that we obtained two non-zero elements; these have the form kv and ℓv , where k, ℓ are chosen uniformly at random in $\{1, \dots, m\}$ with $m \geq 3$. It is well-known that $\gcd(k, \ell) = 1$ with probability greater than $\frac{6}{\pi^2}$. This proves the bound $\frac{6}{\pi^2}(\frac{3}{4})^2 > \frac{1}{3}$. \square

8 Obtaining an approximate generating set of the dual lattice Λ^*

8.1 Lattices of dimension greater than one

The current result in Proposition 7.6 forces us to sample lattice vectors from windows of two different sizes. Recall that the parameter N directly determines the size of the portion of the dual lattice Λ^* from which we can sample. We refer to this parameter as N in Subsection 8.1.1 and as N_0 in Subsection 8.1.2. The other parameters q and κ can be chosen to be the same.

8.1.1 Generating a full-rank sublattice of the dual lattice

Lemma 8.1. *Choose q , N , and κ according to (III)–(V) and*

$$N \geq \frac{1}{\kappa} \left(\max\{8n - 2, n^{(n-1)/2} \cdot 2^{n+1} - 2\} \cdot \frac{n}{2\lambda_1(\Lambda)} + \frac{1}{2nq} \right), \quad (\text{VI})$$

$$N > \frac{1}{\kappa} \left(\frac{1}{2q} + \frac{n^2}{\lambda_1(\Lambda)} \right). \quad (\text{VII})$$

Run the quantum algorithm n times and denote the samples by w_1, \dots, w_n . Then, the probability that there exists $\lambda_1^, \dots, \lambda_n^* \in \Lambda^* \cap [0, \kappa N - \frac{1}{2nq})^n$ with*

- (i) *the lattice vectors $\lambda_1^*, \dots, \lambda_n^*$ span a full-rank sublattice of Λ^* and*
- (ii) *the samples w_i approximate these lattice vectors λ_i^* so that*

$$\left\| \frac{w_i}{2nq} - \lambda_i^* \right\|_2 \leq \frac{1}{2\sqrt{nq}} \text{ for } i = 1, \dots, n$$

is greater or equal to

$$\begin{aligned}
& \frac{1}{4} \left(\frac{2^{n-1} M_\ell L_\ell c}{W} \right)^n \\
& \geq \frac{1}{4} \left(\frac{c}{2} \right)^n \left(\frac{\kappa}{n} \right)^{n^2} \cdot \left[1 - \left(\frac{1}{2q} + \frac{n^2}{\lambda_1(\Lambda)} \right) \frac{1}{\kappa N} \right]^n \cdot \left[1 - \frac{3n}{qN} - \frac{2n\nu(\Lambda)}{q} \right]^n \\
& \approx \frac{1}{4} \left(\frac{c}{2} \right)^n \left(\frac{\kappa}{n} \right)^{n^2}.
\end{aligned}$$

Here $c := \cos^2\left(\pi\left(\frac{1}{4} + \frac{1}{2qN} + 2\kappa n\right)\right) > 0$ and L_ℓ is a lower bound on the cardinality of $\Lambda^* \cap [0, \kappa N - \frac{1}{2nq}]^n$. The approximation \approx indicates that L_ℓ and M_ℓ are close to 1 provided that b , N and q are sufficiently large.

Here, the factor $\frac{1}{4}$ can be replaced with 0.289 or $\prod_{i=1}^{n-1} (1 - 2^{-i})$ (compare Equation (*) on page 31).

Proof. Observe that $\mathcal{R}_{\lambda^*} \subset [0, 2nq\kappa N]^n$ for all $\lambda^* \in \Lambda^* \cap [0, \kappa N - \frac{1}{2nq}]^n$. Set $b := \kappa N - \frac{1}{2nq}$. For all $\lambda^* \in \Lambda^* \cap [0, b]$, Proposition 6.1 yields the lower bound

$$\Pr(w_i \in \mathcal{R}_{\lambda^*}) \geq \frac{2^{n-1} M_\ell c}{W}.$$

Clearly, if $w_i \in \mathcal{R}_{\lambda^*}$ then

$$\left\| \frac{w_i}{2nq} - \lambda^* \right\|_2 \leq \frac{1}{2\sqrt{nq}}.$$

We obtain the lower bound

$$\sum_{(\lambda_1^*, \dots, \lambda_n^*) \in (\Lambda_b^*)^n} \Pr(w_1 \in \mathcal{R}_{\lambda_1^*}, \dots, w_n \in \mathcal{R}_{\lambda_n^*}) \geq \left(\frac{2^{n-1} M_\ell L_\ell c}{W} \right)^n$$

where

$$L_\ell = (\kappa N)^n \det(\Lambda) \left[1 - \left(\frac{1}{2q} + \frac{n^2}{\lambda_1(\Lambda)} \right) \frac{1}{\kappa N} \right]$$

is a lower bound on the cardinality of $\Lambda^* \cap [0, b]^n$. We derive this particular lower bound by applying the argument based on Voronoi cells and

$$\begin{aligned}
\frac{(\kappa N - \frac{1}{2nq} - 2\nu(\Lambda^*))^n}{\det(\Lambda^*)} &= (\kappa N)^n \det(\Lambda) \left[1 - \left(\frac{1}{2nq} + 2\nu(\Lambda^*) \right) \frac{1}{\kappa N} \right]^n \\
&\geq (\kappa N)^n \det(\Lambda) \left[1 - \left(\frac{1}{2q} + 2n\nu(\Lambda^*) \right) \frac{1}{\kappa N} \right]^n \\
&\geq (\kappa N)^n \det(\Lambda) \left[1 - \left(\frac{1}{2q} + \frac{n^2}{\lambda_1(\Lambda)} \right) \frac{1}{\kappa N} \right]^n.
\end{aligned}$$

We used the Bernoulli inequality and the inequality $\lambda_1(\Lambda)\nu(\Lambda^*) \leq \frac{1}{2}n$. Observe that (VII) implies that L_ℓ is nontrivial.

Finally, (VI) implies that b is greater than the lower bound in Lemma 7.8. This shows that at least a fourth of the tuples $(\lambda_1^*, \dots, \lambda_n^*)$ with $\lambda_i \in \Lambda^* \cap [0, b]^n$ for $i = 1, \dots, n$ are such that the lattice vectors generate a full-rank sublattice. \square

8.1.2 Generating the entire dual lattice

Now we combine Proposition 8.1 and Proposition 7.6. We use the same parameters q and κ as in the previous section. We only have to use a larger value for N , which guarantees that we sample from a larger portion of the dual lattice Λ^* to satisfy the premises of Proposition 7.6. We denote this larger value by N_0 . Note that with this choice the conditions (III) and (IV) on q , N_0 , and κ are automatically satisfied. This is because it becomes easier to satisfy these conditions when N is made larger.

Lemma 8.2. *Let q , N , and κ be as in Lemma 8.1. Choose N_0 according to*

$$N_0 \geq 8n^2(n+1)N. \quad (\text{VIII})$$

Use the parameters q , N_0 , and κ for the quantum algorithm. Run it $n+1$ times and denote the samples by w_{n+1}, \dots, w_{2n+1} . Assume that $\lambda_1^, \dots, \lambda_n^*$ from Lemma 8.1 generate a full-rank sublattice of Λ^* . Then, the probability that there exist $\lambda_{n+1}^*, \dots, \lambda_{2n+1}^* \in \Lambda^* \cap [0, \kappa N_0 - \frac{1}{2nq}]^n$ with*

- (i) *the lattice vectors $\lambda_{n+1}^*, \dots, \lambda_{2n+1}^*$ together with the lattice vectors $\lambda_1^*, \dots, \lambda_n^*$ generate the entire dual lattice Λ^* and*
- (ii) *the samples w_{n+i} approximate these lattice vectors λ_{n+i}^* so that*

$$\left\| \frac{w_{n+i}}{2nq} - \lambda_{n+i}^* \right\|_2 \leq \frac{1}{2\sqrt{nq}} \text{ for } i = 1, \dots, n+1$$

is greater or equal to

$$\begin{aligned} & \left(\hat{\zeta} - \frac{1}{4} \right) \left(\frac{2^{n-1} M_\ell L_\ell c_0}{W} \right)^{n+1} \\ & \geq \left(\hat{\zeta} - \frac{1}{4} \right) \left(\frac{c}{2} \right)^{n+1} \left(\frac{\kappa}{n} \right)^{n(n+1)} \cdot \left[1 - \left(\frac{1}{2q} + \frac{n^2}{\lambda_1(\Lambda)} \right) \frac{1}{\kappa N_0} \right]^{n+1} \\ & \quad \left[1 - \frac{3n}{qN_0} - \frac{2n\nu(\Lambda)}{q} \right]^{n+1} \\ & \approx \left(\hat{\zeta} - \frac{1}{4} \right) \left(\frac{c}{2} \right)^{n+1} \left(\frac{\kappa}{n} \right)^{n(n+1)}. \end{aligned}$$

Here, the factor $\hat{\zeta} - \frac{1}{4}$ can be replaced with $\prod_{i=2}^{n+1} \zeta(i)^{-1} - \frac{1}{4}$ (compare Equation (*) on page 31).

The proof of this lemma is basically the same as that of Lemma 8.1. Here L_ℓ is the lower bound on $\Lambda^* \cap [0, b_0]^n$ where $b_0 := \kappa N_0 - \frac{1}{2nq}$, M_ℓ the lower bound

on M in Proposition 5.4 (iii), and $W = (2nqN_0)^n$, and $c_0 = \cos^2(\pi(\frac{1}{4} + \frac{1}{2qN_0} + 2\kappa n))$. The cosine factor c_0 is bounded from below by $c = \cos^2(\pi(\frac{1}{4} + \frac{1}{2qN} + 2\kappa n))$ since $N_0 > N$. The approximation \approx indicates that L_ℓ and M_ℓ are close to 1 provided that q and N_0 are sufficiently large.

There is one point that should be explained in more detail. It remains to verify that $b_0 \geq 8n(n^2 + 1)b$ so that we can apply Lemma 7.8. The condition on the relation of the window sizes is equivalent to

$$\kappa N_0 - \frac{1}{2nq} \geq 8n(n^2 + 1) \left(\kappa N - \frac{1}{2nq} \right).$$

This inequality is clearly satisfied due to (VIII).

8.1.3 Bounding the probability

We replace condition (VII) by the stricter condition

$$N \geq \frac{1}{\kappa} \left(\frac{n}{q} + \frac{2n^3}{\lambda_1(\Lambda)} \right). \quad (\text{VII}_1)$$

This, together with (VIII), implies

$$\left[1 - \left(\frac{1}{2q} + \frac{n^2}{\lambda_1(\Lambda)} \right) \frac{1}{\kappa N} \right]^n \cdot \left[1 - \left(\frac{1}{2q} + \frac{n^2}{\lambda_1(\Lambda)} \right) \frac{1}{\kappa N_0} \right]^{n+1} \geq \frac{1}{2^2}.$$

Moreover, we replace condition (IV) by the stricter condition

$$q \geq \frac{6n^2}{N} + 4n(n+1)\nu(\Lambda). \quad (\text{IV}_1)$$

This implies together with (VIII)

$$\left[1 - \frac{3n}{qN} - \frac{2n\nu(\Lambda)}{q} \right]^n \cdot \left[1 - \frac{3n}{qN_0} - \frac{2n\nu(\Lambda)}{q} \right]^{n+1} \geq \frac{1}{2^2}.$$

From the previous two subsections, under the assumption that (I)–(VIII) hold, we get that the probability that $2n+1$ samples from the algorithm generate the whole lattice Λ^* is at least

$$\begin{aligned} & \frac{1}{4} \left(\hat{\zeta} - \frac{1}{4} \right) \left(\frac{c}{2} \right)^{2n+1} \left(\frac{\kappa}{n} \right)^{2n^2+n} \cdot \left[1 - \left(\frac{1}{2q} + \frac{n^2}{\lambda_1(\Lambda)} \right) \frac{1}{\kappa N} \right]^n \\ & \cdot \left[1 - \frac{3n}{qN} - \frac{2n\nu(\Lambda)}{q} \right]^n \cdot \left[1 - \left(\frac{1}{2q} + \frac{n^2}{\lambda_1(\Lambda)} \right) \frac{1}{\kappa N_0} \right]^{n+1} \\ & \cdot \left[1 - \frac{3n}{qN_0} - \frac{2n\nu(\Lambda)}{q} \right]^{n+1}, \end{aligned}$$

where $c = \cos^2(\pi(\frac{1}{4} + \frac{1}{4qN} + 2\kappa n))$. Using the stricter conditions (VII₁) and (IV₁) from above, this can be bounded from below by

$$\frac{1}{2^6} \left(\hat{\zeta} - \frac{1}{4} \right) \left(\frac{c}{2} \right)^{2n+1} \left(\frac{\kappa}{n} \right)^{2n^2+n} \geq \frac{1}{2^9} \left(\frac{c}{2} \right)^{2n+1} \left(\frac{\kappa}{n} \right)^{2n^2+n}.$$

Here, the factor $\frac{1}{4}(\hat{\zeta} - \frac{1}{4})$ can be increased to 0.053176 or $(\prod_{i=2}^{n+1} \zeta(i)^{-1} - \frac{1}{4}) \cdot \prod_{i=1}^{n-1} (1 - 2^{-i})$ (compare Equation (*) on page 31). The latter would improve the lower bound on the probability that $2n + 1$ samples from the algorithm generate the whole lattice Λ^* to

$$\frac{1}{2^4} \left(\prod_{i=2}^{n+1} \zeta(i)^{-1} - \frac{1}{4} \right) \left(\prod_{i=1}^{n-1} (1 - 2^{-i}) \right) \left(\frac{c}{2} \right)^{2n+1} \left(\frac{\kappa}{n} \right)^{2n^2+n}.$$

8.2 Dimension one

Finally, we want to investigate the case $n = 1$ more closely. In this case, we have only one window and we sample only two vectors. If $b \geq 3 \det(L) + 1$, Lemma 7.9 yields that two randomly sampled vectors from $\Lambda^* \cap [0, b)$ generate Λ^* is larger than $\frac{1}{3}$. We proceed similarly to the proof of Proposition 8.1. For $b = \kappa N - \frac{1}{2q}$ to hold in conjunction with $b \geq 3 \det(L) + 1 = \frac{3}{\det(\Lambda)} + 1$, we must satisfy the new condition

$$N \geq \frac{1}{\kappa} \left(\frac{3}{\det(\Lambda)} + 1 + \frac{1}{2q} \right). \quad (\text{VI}_2)$$

Assume that the assumptions (I)–(V) and (VI₂) are satisfied. Let w_1, w_2 be the two samples output by our quantum algorithm. Then, the probability that all sampled w_i correspond to lattice vectors λ_i^* in $L_{[0,b)}$ for $i = 1, 2$ and that they generate L is at least

$$\begin{aligned} & \frac{1}{3} \left(\frac{M_\ell L_\ell c}{W} \right)^2 \\ & \geq \frac{1}{12} \kappa^2 c^2 \cdot \left[1 - \left(\frac{1}{2q} + \frac{1}{\det(\Lambda)} \right) \frac{1}{\kappa N} \right]^2 \cdot \left[1 - \frac{3}{qN} - \frac{\det(\Lambda)}{q} \right]^2, \end{aligned}$$

where L_ℓ is the lower bound on $L_{[0,b)}$ in Proposition 8.1, c the cosine-factor in Proposition 8.1, M_ℓ the lower bound on M in Proposition 5.4 (iii), and $W = 2qN$.

Let us introduce the two new assumptions

$$q \geq \frac{12}{N} + 4 \det(\Lambda) \quad (\text{IV}_2)$$

$$\text{and } N \geq \frac{1}{\kappa} \left(\frac{2}{q} + \frac{4}{\det(\Lambda)} \right); \quad (\text{VII}_2)$$

these imply (IV), and allow us to bound

$$\left(1 - \frac{1}{2q\kappa N} - \frac{1}{\kappa N \det(\Lambda)} \right)^2 \geq \frac{1}{2} \quad \text{and} \quad \left(1 - \frac{3}{qN} - \frac{\det(\Lambda)}{q} \right)^2 \geq \frac{1}{2}.$$

This yields the lower bound $\frac{1}{48} \kappa^2 c^2$ on the success probability.

9 Lattice theoretic tools – Part 2

First, we consider the problem to obtain an approximate basis of a lattice L from an approximate generating set of L . Second, we consider the problem to obtain an approximate basis of the dual lattice L^* from an approximate basis of L .

9.1 Computing an approximate basis of L from an approximate generating set of L

We address the problem of computing an approximate basis from an approximate generating set. In this subsection, we present BUCHMANN’S AND KESSLER’S approach in [BK93]. Our exposition simplifies and improves their results. Our more general analysis makes it possible to quantify the approximation quality when different lattice approximation algorithms can be used. The analysis in [BK93] is written only for the LLL algorithm. In the context of our quantum algorithm it is more advantageous to use algorithms to compute Korkine-Zolotarev reduced bases. In our analysis, the approximation quality is entirely expressed in terms of the lattice L . In contrast, in [BK93] the approximation quality depends on the characteristics of some sublattice of L .

Remark 9.1. An approach based on [BK93] was already suggested in [Sch07]. However, our requirements on the precision of the approximation can be stated in much simpler terms than those made in [Sch07]. For instance, an important simplification is that we do not have to consider any sublattice (compare to [Sch07, Satz 2.4.24]).

Note that [Hal05] suggested to use the precursor [BP89] for computing an approximate basis. The problem is that this earlier work does not make any statements on the size of the entries of a certain unimodular transformation matrix. Therefore, the results of this work cannot be directly applied because it not possible to quantify the quality of the resulting approximate basis. The major motivation for the follow-up work [BK93] to [BP89] was to bound the entries of the relevant transformation matrix (see [BK93, Introduction]).

Observe that both [BK93] and [BP89] rely on the LLL basis reduction algorithm to compute the transformation matrix. However, for the quantum algorithm it is significantly better to compute Korkine-Zolotarev-reduced bases in the classical post-processing step. This makes it possible to obtain a transformation matrix with exponentially smaller entries, which in turn yields an exponentially better approximation of the basis of the period lattice of the infrastructure. If the LLL algorithm is used, then it is necessary to evaluate the function f over an exponentially wider window to achieve the same quality of approximation of the period lattice. Note that the cost of computing Korkine-Zolotarev bases in the classical post-processing step is negligible compared to the time complexity of the quantum part.

Let L be a lattice in \mathbb{R}^n of rank $r \leq n$.

Definition 9.2 (Approximate basis). *We call $\mathbf{b}'_1, \dots, \mathbf{b}'_r$ a δ -approximate basis of L if there exists a basis $\mathbf{b}_1, \dots, \mathbf{b}_r$ of L with*

$$\|\mathbf{b}'_i - \mathbf{b}_i\|_2 \leq \delta$$

for $i = 1, \dots, r$.

Definition 9.3 (Approximate generating set). *We call $\mathbf{a}'_1, \dots, \mathbf{a}'_k$ an ε -approximate generating set of L if there exists a generating set $\mathbf{a}_1, \dots, \mathbf{a}_k$ of L with*

$$\|\mathbf{a}'_j - \mathbf{a}_j\|_2 \leq \varepsilon \tag{1}$$

for $j = 1, \dots, k$.

We assume

$$\begin{aligned} \mu &\leq \lambda_1(L) \\ \alpha &\geq \max_{j=1, \dots, k} \{\|\mathbf{a}_j\|_2\}. \end{aligned}$$

We need these bounds to derive the method for computing an approximate basis from an ε -approximate generating set and to bound its corresponding δ in terms of ε , μ , α , n , and k .

Remark 9.4. The approximate generating set arises in the following way in our quantum algorithm. We are given an algorithm that returns rational vectors of the special form $[t\mathbf{a}_j]$ where the vectors $\mathbf{a}_1, \dots, \mathbf{a}_k$ generate the lattice. The parameter t specifies the quality of the approximation and is under our control. The problem is to find a unimodular matrix $T \in \mathbb{Z}^{k \times r}$ that transforms the $\frac{\sqrt{n}}{2t}$ approximate generating set $\frac{1}{t}[t\mathbf{a}_j]$ into an approximate basis of L and to determine its corresponding δ .

We call a vector $\mathbf{z} = (z_1, \dots, z_k) \in \mathbb{Z}^k$ a (nontrivial) relation for the generating set if $\mathbf{z} \neq \mathbf{0}$ and

$$\sum_{j=1}^k z_j \mathbf{a}_j = \mathbf{0}, \tag{2}$$

where $\mathbf{0}$ denotes the (column) zero vector in either \mathbb{Z}^k or \mathbb{Z}^n .

Lemma 9.5 (Sufficient and necessary condition for relations). *Let $\mathbf{z} \in \mathbb{Z}^k$ and assume that*

$$2\varepsilon\|\mathbf{z}\|_1 < \mu. \tag{3}$$

Then \mathbf{z} is a relation for the generating set if and only if

$$\left\| \sum_{j=1}^k z_j \mathbf{a}'_j \right\|_2 \leq \varepsilon\|\mathbf{z}\|_1. \tag{4}$$

Proof. Let \mathbf{z} be an arbitrary relation. The condition in (4) follows then from (1) and (2)

$$\left\| \sum_{j=1}^k z_j \mathbf{a}'_j \right\|_2 \leq \sum_{j=1}^k |z_j| \|\mathbf{a}'_j - \mathbf{a}_j\|_2 + \left\| \sum_{j=1}^k z_j \mathbf{a}_j \right\|_2 \leq \varepsilon \|\mathbf{z}\|_1. \quad (5)$$

Now assume that (4) holds for some (nonzero) vector $\mathbf{z} \in \mathbb{Z}^k$. Using (1) and (3) we obtain

$$\begin{aligned} \left\| \sum_{j=1}^k z_j \mathbf{a}_j \right\|_2 &\leq \left\| \sum_{j=1}^k z_j (\mathbf{a}_j - \mathbf{a}'_j) \right\|_2 + \left\| \sum_{j=1}^k z_j \mathbf{a}'_j \right\|_2 \\ &\leq 2\varepsilon \|\mathbf{z}\|_1 < \mu. \end{aligned}$$

Since $\mu \leq \lambda_1(L)$ we must have that $\sum_{j=1}^k z_j \mathbf{a}_j = \mathbf{0}$. □

It is convenient to define the scaled approximation vectors

$$\hat{\mathbf{a}}_j = s \mathbf{a}'_j,$$

where s is a positive parameter that we fix later. Clearly, $\|\hat{\mathbf{a}}_j - s \mathbf{a}_j\|_2 \leq s\varepsilon$.

Definition 9.6 (Approximation lattice). *For $j = 1, \dots, k$, define the vectors $\tilde{\mathbf{a}}_j \in \mathbb{Z}^k \oplus \mathbb{R}^n$ by*

$$\tilde{\mathbf{a}}_j = \mathbf{e}_j \oplus \hat{\mathbf{a}}_j,$$

where \mathbf{e}_j is the j th standard basis vector of \mathbb{Z}^k . The vectors $\tilde{\mathbf{a}}_1, \dots, \tilde{\mathbf{a}}_k$ are linearly independent and form a basis of the approximation lattice

$$\tilde{L} = \bigoplus_{j=1}^k \mathbb{Z} \tilde{\mathbf{a}}_j.$$

The following lemma establishes that short lattice vectors of \tilde{L} give rise to relations for the generating set of L . For the sake of generality we introduce the parameter f that characterizes the approximation quality of basis reduction algorithms. We have $f = 2^{(k-1)/2}$ and $f = \frac{\sqrt{k+3}}{2}$ for the algorithms that compute LLL-reduced and Korkine-Zolotarev reduced bases.

Lemma 9.7 (Sufficient condition for relations). *Let $\lambda \geq 1$. Assume that the approximation error ε is bounded from above by*

$$\varepsilon \leq \frac{\mu}{2f\lambda\sqrt{k}}$$

and the scaling factor s is chosen so that

$$s > \frac{2f\lambda}{\mu}. \quad (6)$$

Let $\mathbf{z} = (z_1, \dots, z_k) \in \mathbb{Z}^k$ be an arbitrary vector and

$$\tilde{\mathbf{x}} = \sum_{j=1}^k z_j \tilde{\mathbf{a}}_j .$$

the corresponding lattice vector of \tilde{L} . If

$$\|\tilde{\mathbf{x}}\|_2 \leq f\lambda$$

then \mathbf{z} is a relation for the generating set $\mathbf{a}_1, \dots, \mathbf{a}_k$.

Proof. We prove the lemma by showing that the contraposition of the statement holds. Assume that \mathbf{z} is not a relation. We have to show that corresponding vector $\tilde{\mathbf{x}}$ is strictly longer than $f\lambda$.

We write $\tilde{\mathbf{x}} = \mathbf{z} \oplus \hat{\mathbf{x}}$ with $\hat{\mathbf{x}} = \sum_{j=1}^k z_j \hat{\mathbf{a}}_j$. Then we have

$$\|\tilde{\mathbf{x}}\|_2^2 = \|\mathbf{z}\|_2^2 + \|\hat{\mathbf{x}}\|_2^2 .$$

If $\|\mathbf{z}\|_2 > f\lambda$ holds then we are done. Otherwise we have

$$\begin{aligned} \|\tilde{\mathbf{x}}\|_2 &\geq \|\hat{\mathbf{x}}\|_2 = \left\| \sum_{j=1}^k z_j \hat{\mathbf{a}}_j \right\|_2 \\ &\geq s \left\| \sum_{j=1}^k z_j \mathbf{a}_j \right\|_2 - \left\| \sum_{j=1}^k z_j (s \mathbf{a}_j - \hat{\mathbf{a}}_j) \right\|_2 \\ &\geq s\mu - s\|\mathbf{z}\|_1 \varepsilon \geq s\mu - s\|\mathbf{z}\|_2 \sqrt{k} \varepsilon \\ &\geq s\mu - s f \lambda \sqrt{k} \varepsilon = s \left(\mu - f \lambda \sqrt{k} \varepsilon \right) \\ &\geq s \left(\mu - \frac{\mu}{2} \right) \geq s \frac{\mu}{2} > f\lambda . \end{aligned}$$

□

Lemma 9.8 (Linearly independent relations of bounded norm). *There exist $k - r$ linearly independent relations $\mathbf{m}_1, \dots, \mathbf{m}_{k-r}$ of the generating set with*

$$\|\mathbf{m}_j\|_\infty \leq \frac{\alpha^r}{\det(L)} .$$

Proof. We construct an isometric embedding of L into \mathbb{R}^r . Let $\mathbf{b}_1, \dots, \mathbf{b}_r$ be a basis of L and $\mathbf{b}_1^*, \dots, \mathbf{b}_r^*$ the corresponding orthonormal vectors obtained by the Gram-Schmidt process. Let $\mathbf{w}_1, \dots, \mathbf{w}_r$ be an arbitrary orthonormal basis of \mathbb{R}^r . The mapping Φ defined by

$$\Phi(\mathbf{b}_i^*) = \mathbf{w}_i$$

for $i = 1, \dots, r$ is an isometry between L and $L^\Phi := \Phi(L)$ and we have $\det(L) = \det(L^\Phi)$. We set $\mathbf{a}_i^\Phi := \Phi(\mathbf{a}_i)$. We assume w.l.o.g. that the first r vectors of the generating set $\mathbf{a}_1, \dots, \mathbf{a}_k$ are linearly independent. Define the matrices

$$\begin{aligned} A &= (\mathbf{a}_1^\Phi | \dots | \mathbf{a}_r^\Phi | \dots | \mathbf{a}_k^\Phi) \\ C &= (\mathbf{a}_1^\Phi | \dots | \mathbf{a}_r^\Phi) \end{aligned}$$

The submatrix $C \in \mathbb{R}^{r \times r}$ of $A \in \mathbb{R}^{r \times k}$ is nonsingular, which follows from the assumption that the first r generators of L are linearly independent. Let $\mathbf{v}_j \in \mathbb{R}^r$ be the solutions of the linear system

$$C\mathbf{v}_j = \mathbf{a}_{r+j}^\Phi$$

for $j = 1, \dots, k - r$. Define the (column) vectors

$$\mathbf{m}_j = \frac{\det(C)}{\det(L^\Phi)} (\mathbf{v}_j \oplus (-1)\mathbf{e}_j),$$

where \mathbf{e}_j are the standard basis vectors of \mathbb{R}^{k-r} for $j = 1, \dots, k - r$. Due to construction they are linearly independent and form a basis of the kernel of A (which has dimension $k - r$) since

$$A\mathbf{m}_j = \frac{\det(C)}{\det(L^\Phi)} (C\mathbf{v}_j - \mathbf{a}_{r+j}) = \frac{\det(C)}{\det(L^\Phi)} (\mathbf{a}_{r+j} - \mathbf{a}_{r+j}) = \mathbf{0}$$

for $j = 1, \dots, k - r$. Using Cramer's rule, we can express the coefficients v_{ij} of the vector \mathbf{v}_j as

$$v_{ij} = \frac{\det(\mathbf{a}_1^\Phi | \dots | \mathbf{a}_{i-1}^\Phi | \mathbf{a}_{r+j}^\Phi | \mathbf{a}_{i+1}^\Phi | \dots | \mathbf{a}_r^\Phi)}{\det(C)}.$$

Note that the values

$$\frac{\det(C)}{\det(L^\Phi)} v_{ij} = \frac{\det(\mathbf{a}_1^\Phi | \dots | \mathbf{a}_{i-1}^\Phi | \mathbf{a}_{r+j}^\Phi | \mathbf{a}_{i+1}^\Phi | \dots | \mathbf{a}_r^\Phi)}{\det(L^\Phi)}$$

are always either 0 or the indices of full-rank sublattices of L^Φ . The two mutually exclusive cases are: (i) \mathbf{a}_{r+j}^Φ is contained in the span of $\mathbf{a}_1^\Phi, \dots, \mathbf{a}_{i-1}^\Phi, \mathbf{a}_{i+1}^\Phi, \dots, \mathbf{a}_r^\Phi$, implying that the determinant is 0 and (ii) $\mathbf{a}_1^\Phi, \dots, \mathbf{a}_{i-1}^\Phi, \mathbf{a}_{r+j}^\Phi, \mathbf{a}_{i+1}^\Phi, \dots, \mathbf{a}_r^\Phi$, implying that they generate a full-rank sublattice. Therefore, all components of \mathbf{m}_j are integers. This concludes the proof that $\mathbf{m}_1, \dots, \mathbf{m}_{k-r}$ are relations for the generating set.

The upper bound on the $\|\cdot\|_\infty$ -norm of these relations follows directly from Minkowski's inequality. We can bound the absolute value of the determinants by the product of the norms of the column vectors, which can be at most α^r . \square

Lemma 9.9 (Upper bounds on minima of the approximate lattice). *Assume we set*

$$\lambda = 3\sqrt{k} \frac{\alpha^r}{\det(L)}$$

and choose the scaling factor s so that

$$s \leq \frac{4f\lambda}{\mu}. \quad (7)$$

The first $(k-r)$ minima are bounded from above by

$$\lambda_j(\tilde{L}) \leq \lambda$$

for $j = 1, \dots, k-r$.

The last r minima are bounded from above

$$\lambda_j(\tilde{L}) \leq \sqrt{s^2(\alpha + \varepsilon)^2 + 1} \leq \frac{6.5f\lambda\alpha}{\mu}$$

for $j = k-r+1, \dots, k$.

Proof. Let \mathbf{m}_j be the $(k-r)$ linearly independent relations constructed in the proof of Lemma 9.8. We define the vectors

$$\tilde{\mathbf{x}}_j = \sum_{i=1}^k m_{ij} \tilde{\mathbf{a}}_i = \mathbf{m}_j \oplus \sum_{i=1}^k m_{ij} \hat{\mathbf{a}}_i.$$

Obviously, the vectors $\tilde{\mathbf{x}}_j$ are linearly independent. Since \mathbf{m}_j is a relation we may apply the inequality in (5) from the first part of the proof of Lemma 9.5. We obtain

$$\begin{aligned} \|\tilde{\mathbf{x}}_j\|_2 &\leq \|\mathbf{m}_j\|_2 + \left\| \sum_{i=1}^k m_{ij} \hat{\mathbf{a}}_i \right\|_2 = \|\mathbf{m}_j\|_2 + s \left\| \sum_{i=1}^k m_{ij} \mathbf{a}'_i \right\|_2 \\ &\leq \|\mathbf{m}_j\|_2 + s\varepsilon \|\mathbf{m}_j\|_1 \leq \sqrt{k} \|\mathbf{m}_j\|_\infty + s\varepsilon k \|\mathbf{m}_j\|_\infty \\ &= \left(1 + s\varepsilon\sqrt{k}\right) \sqrt{k} \|\mathbf{m}_j\|_\infty \leq \left(1 + \frac{2}{\sqrt{k}}\sqrt{k}\right) \sqrt{k} \|\mathbf{m}_j\|_\infty \\ &\leq 3\sqrt{k} \frac{\alpha^r}{\det(L)} \leq \lambda. \end{aligned}$$

The upper bound on the last minima follows from

$$\lambda_j(\tilde{L}) \leq \max_{i=1, \dots, k} \|\tilde{\mathbf{a}}_i\|_2 \leq \sqrt{s^2(\alpha + \varepsilon)^2 + 1}.$$

The upper bound on the square root expression holds since the tangent to the square root at $s^2(\alpha + \varepsilon)^2 > 1$ has slope greater or equal to $1/2$ so a displacement by 1 can increase the value by at most $1/2$ and $s\varepsilon \leq 2/\sqrt{k} \leq 1$. This yields observations yield the upper bound $4f\lambda\alpha/\mu + 2.5$, which bounded from above by $6.5f\lambda\alpha/\mu$. \square

To simplify notation in the following we set

$$\tilde{\alpha} = \sqrt{s^2(\alpha + \varepsilon)^2 + 1}.$$

We apply the basis reduction algorithm to the lattice basis $\tilde{\mathbf{a}}_1, \dots, \tilde{\mathbf{a}}_k$ and obtain the reduced basis $\tilde{\mathbf{b}}_1, \dots, \tilde{\mathbf{b}}_k$. Denote by $M = (m_{ij}) \in \mathbb{Z}^{k \times k}$ the corresponding (unimodular) transformation matrix. We write the reduced basis vectors as

$$\tilde{\mathbf{b}}_j = (\mathbf{m}_j, \hat{\mathbf{b}}_j)$$

where $\mathbf{m}_j = (m_{1j}, \dots, m_{kj}) \in \mathbb{Z}^k$ are the column vectors of M and

$$\hat{\mathbf{b}}_j = \sum_{i=1}^k m_{ij} \hat{\mathbf{a}}_i.$$

The following lemma shows we can directly obtain a basis of L with the help of the transformation matrix M .

Lemma 9.10 (Basis and approximate bases for L). *Set*

$$\lambda = 3\sqrt{k} \frac{\alpha^r}{\det(L)}.$$

Assume that the approximation error is bounded from above by

$$\varepsilon \leq \frac{\mu}{2f\lambda\sqrt{k}}$$

and the scaling factor is bounded from below and above by

$$\frac{2f\lambda}{\mu} < s \leq \frac{4f\lambda}{\mu}.$$

Let M be the transformation matrix returned by the basis reduction algorithm when applied to the basis $\tilde{\mathbf{a}}_1, \dots, \tilde{\mathbf{a}}_k$ of the approximation lattice \tilde{L} .

Define the vectors

$$\begin{aligned} \mathbf{b}_j &= \sum_{i=1}^k m_{i,k-r+j} \mathbf{a}_i \\ \mathbf{b}'_j &= \sum_{i=1}^k m_{i,k-r+j} \mathbf{a}'_i \end{aligned}$$

for $j = 1, \dots, r$. Then we have

- *The vectors $\mathbf{b}_1, \dots, \mathbf{b}_r$ form a basis of L and their norms are bounded from above by*

$$\|\mathbf{b}_j\|_2 \leq f\sqrt{k} \tilde{\alpha} \alpha.$$

- *The vectors $\mathbf{b}'_1, \dots, \mathbf{b}'_r$ form a δ -approximate basis of L with*

$$\delta \leq f\sqrt{k} \tilde{\alpha} \varepsilon.$$

Proof. We know that the reduced basis vectors $\tilde{\mathbf{b}}_j$ satisfy

$$\|\tilde{\mathbf{b}}_\ell\|_2 \leq f\lambda_\ell(\tilde{L}).$$

Using the upper bounds on the first $(k-r)$ minima in Lemma 9.9 we obtain

$$\|\tilde{\mathbf{b}}_\ell\|_2 \leq f\lambda$$

for $\ell = 1, \dots, k-r$. These vectors are sufficiently short so that Lemma 9.7 applies. We conclude that $\mathbf{m}_1, \dots, \mathbf{m}_{k-r}$ are relations for the generating set $\mathbf{a}_1, \dots, \mathbf{a}_k$ of L .

Let $A = (\mathbf{a}_1 | \dots | \mathbf{a}_k) \in \mathbb{Z}^{n \times k}$. Then we have

$$AM = \left(\underbrace{\mathbf{0} | \dots | \mathbf{0}}_{k-r} | \mathbf{b}_1 | \dots | \mathbf{b}_r \right) \in \mathbb{Z}^{n \times k}$$

since the first $k-r$ columns of M are relations of the generating set. Since M is unimodular the lattice generated by $\mathbf{b}_1, \dots, \mathbf{b}_r$ is equal to L and, thus, $\mathbf{b}_1, \dots, \mathbf{b}_r$ form a basis.

We first determine an upper bound on the norm of the last r column vectors of M . For $j = 1, \dots, r$, we have

$$\|\mathbf{m}_{k-r+j}\|_2 \leq \|\tilde{\mathbf{b}}_j\|_2 \leq f\lambda_{k-r+j}(\tilde{L}) \leq f\tilde{\alpha}.$$

We have

$$\begin{aligned} \|\mathbf{b}_j\|_2 &\leq \|\mathbf{m}_{k-r+j}\|_1 \alpha \leq \sqrt{k} f \tilde{\alpha} \alpha \\ \|\mathbf{b}'_j - \mathbf{b}_j\|_2 &\leq \|\mathbf{m}_{k-r+j}\|_1 \varepsilon \leq \sqrt{k} f \tilde{\alpha} \varepsilon \end{aligned}$$

for $j = 1, \dots, r$. □

We assume in the following the lattice L has full rank, i.e., $r = n$. This situation occurs precisely in our quantum algorithm. To further simplify notation, we also set

$$g := f\sqrt{k}\tilde{\alpha}.$$

9.2 Computing an approximate basis of the dual lattice L^* from an approximate basis of L

Lemma 9.11. *Let $\mathbf{b}'_1, \dots, \mathbf{b}'_n$ be a δ -approximate basis of L with $\delta \leq g\varepsilon$ as in the lemma above. Then we can obtain a γ -approximate basis of the dual lattice L^* with*

$$\gamma \leq \frac{2n^{5/2}g^{2n-1}\alpha^{2(n-1)}}{\det(L)^2} \varepsilon.$$

provided that

$$\varepsilon \leq \frac{\det(L)}{2n^{3/2}g^n\alpha^{n-1}}.$$

Proof. Let $B = (\mathbf{b}_1 | \dots | \mathbf{b}_n)$ and $B' = (\mathbf{b}'_1 | \dots | \mathbf{b}'_n)$ be the matrices whose columns form the basis of L and the approximate basis of L , respectively. We compute the inverses of these matrices to obtain the basis and the approximate basis of the dual lattice L^* .

Denote the perturbation by $E = B' - B$. We use [SS90, Theorem 2.5] to estimate the sensitivity of the inverse under perturbation. If $\|B^{-1}\|_1 \|E\|_1 < 1$, then $B + E$ is nonsingular and

$$\|B'^{-1} - B^{-1}\|_1 = \|(B + E)^{-1} - B^{-1}\|_1 \leq \frac{\|B^{-1}\|_1^2 \|E\|_1}{1 - \|B^{-1}\|_1 \|E\|_1}.$$

We may apply the bound from this theorem because the matrix norm on $\mathbb{R}^{n \times n}$ defined by $\|X\|_1 = \max_{1 \leq j \leq n} \sum_{i=1}^n x_{ij}$ is multiplicative.

Let c_{ij} denote the entries of B^{-1} . Using Cramer's rule and Hadamard's inequality, we have

$$|c_{ij}| = \left| \frac{\det(\mathbf{b}_1, \dots, \mathbf{b}_{i-1}, \mathbf{e}_j, \mathbf{b}_{i+1}, \dots, \mathbf{b}_n)}{\det(B)} \right| \leq \frac{\prod_{i \neq j} \|\mathbf{b}_i\|_2}{\det(L)} \leq \frac{(g\alpha)^{n-1}}{\det(L)}.$$

This implies

$$\|B^{-1}\|_1 \leq \frac{n(g\alpha)^{n-1}}{\det(L)}.$$

Note that the Euclidean norm of the column vectors of E is bounded by δ from above since these vectors are equal to $\mathbf{b}'_i - \mathbf{b}_i$. This implies $\|E\|_1 \leq \sqrt{n}\delta \leq \sqrt{n}g\varepsilon$.

Assume that

$$\varepsilon \leq \frac{\det(L)}{2n^{3/2}g^n\alpha^{n-1}}, \quad (8)$$

which ensures that $\|B^{-1}\|_1 \|E\|_1 \leq 1/2$. Then we have

$$\|B'^{-1} - B^{-1}\|_1 \leq \frac{2n^{5/2}g^{2n-1}\alpha^{2(n-1)}}{\det(L)^2} \varepsilon. \quad (9)$$

This implies that the column vectors of B'^{-1} form a γ -approximate basis of L^* with

$$\gamma \leq \frac{2n^{5/2}g^{2n-1}\alpha^{2(n-1)}}{\det(L)^2} \varepsilon. \quad (10)$$

□

Corollary 9.12. *Recall that the quantum algorithm returns a generating set with $\varepsilon \leq 1/(4\sqrt{n}q)$. This and the above lemma imply that if*

$$q \geq \max \left\{ \frac{ng^n\alpha^{n-1}}{2\det(L)}, \frac{n^2g^{2n-1}\alpha^{2(n-1)}}{2\det(L)^2} \cdot \frac{1}{\gamma} \right\} \quad (11)$$

then we obtain a γ -approximate basis of the dual lattice L^* , where

$$g \leq \frac{19.5kf^2\alpha^{n+1}}{\det(L)\lambda_1(L)} \quad \text{and} \quad \alpha \geq \max_{j=1, \dots, k} \{\|\mathbf{a}_j\|_2\}. \quad (12)$$

Proof. This follows with $g = \sqrt{k}f\tilde{\alpha} \leq \sqrt{k}6.5f^2\lambda\alpha/\mu$ and $\lambda = 3\sqrt{k}\alpha^n/\det(L)$. \square

10 Final analysis of the quantum algorithm

By combining all material from the previous sections, we obtain the following result:

Theorem 10.1. *Assume that **A1)**–**A3)** hold with $C \leq 1$ and $A \geq 1$. Further, assume that $N, q, N_0, L \in \mathbb{N}$ are chosen such that*

$$N \geq \max \left\{ 32, \frac{8(n+1)n2^n D A^{n-1}}{3C^n}, \frac{9n^2}{32} + \frac{18n^4}{\lambda_1(\Lambda)}, \right. \\ \left. \max\{8n-2, n^{(n-1)/2}2^{n+1}-2\} \cdot \frac{9n^2}{2\lambda_1(\Lambda)} + \frac{9}{64} \right\},$$

$$N_0 \geq 8n^2(n+1)N,$$

$$q \geq \max \left\{ 32, 9A, \frac{6n^2}{N} + 2n^{(n+1)/2+1}(n+1) \frac{\det(\Lambda)}{\lambda_1(\Lambda)^{n-1}}, \right. \\ \frac{19.5^n n^{n+3/2} (1 + \frac{5}{2n} + \frac{1}{n^2})^n N_0^{n^2+2n-1} \det(\Lambda)^{2n+1}}{2 \cdot 9^{n^2+2n-1} \lambda_1(\Lambda)^{n^2-n}}, \\ \left. \frac{19.5^{2n} n^{2n+3/2} (1 + \frac{5}{2n} + \frac{1}{n^2})^{2n-1} N_0^{2n^2+3n-3} \det(\Lambda)^{4n}}{\gamma \cdot 39 \cdot 9^{2n^2+3n-3} \lambda_1(\Lambda)^{2n^2-3n-1}} \right\}$$

$$\text{and } L \geq \frac{4nD(q+A+C+2)^n}{C^n}.$$

Set $\kappa := \frac{1}{9n}$ and assume that $s \in S$ is chosen uniformly at random. Then the probability that the algorithm described in Section 3, applied n times with the parameters N, q, κ and $n+1$ times with the parameters N_0, q, κ , returns an $\frac{1}{4\sqrt{nq}}$ -approximate generating set of Λ^* is at least

$$\frac{\cos\left(\pi \frac{17417}{36864}\right)^{4n+2}}{2^{2n+6} 3^{4n^2+2n} n^{4n^2+2n}} \left(\prod_{i=2}^{n+1} \zeta(i)^{-1} - \frac{1}{4} \right) \prod_{i=1}^{n-1} (1 - 2^{-i}) \\ \geq \frac{6.198327 \cdot 1.54587777^n}{10^{6n+6} 81^{n^2} n^{4n^2+2n}}.$$

If such an approximate generating set of Λ^* is obtained, the algorithm described in Section 9 computes a γ -approximate basis of Λ .

We will prove this theorem further down (on page 48). In case $n = 1$, we can improve the bound from Theorem 10.1 significantly:

Proposition 10.2. *Assume that $\Lambda \subseteq \mathbb{R}$, i.e., that $n = 1$. Further, assume that **A1)**–**A3)** hold with $C \leq 1$ and $A \geq 1$, and assume that $N, q, L \in \mathbb{N}$ are chosen*

such that

$$\begin{aligned}
N &\geq \max\left\{32, \frac{4}{A}, \frac{32D}{3C}, \frac{36}{\det(\Lambda)} + \frac{9}{16}, \frac{27}{\det(\Lambda)} + 9 + \frac{9}{16}\right\}, \\
q &\geq \max\left\{32, 9A, \frac{12}{N} + 4\det(\Lambda), \frac{19.5}{9^2}N^2\det(\Lambda)^3 \cdot \max\left\{1, \frac{\det(\Lambda)}{\gamma}\right\}\right\} \\
\text{and } L &\geq \frac{4D(q + A + C + 2)}{C}.
\end{aligned}$$

Set $\kappa := \frac{1}{9}$ and assume that $s \in S$ is chosen uniformly at random. Then the probability that the algorithm described in Section 3, applied two times with the parameters N, q, κ , returns an $\frac{1}{4q}$ -approximate generating set of Λ^* is at least

$$\frac{\cos^4\left(\pi\frac{17417}{36864}\right)}{7776} \geq 7.163 \cdot 10^{-9}.$$

If such an approximate generating set of Λ^* is obtained, the algorithm described in Section 9 computes a γ -approximate basis of Λ .

We will also prove this proposition further down (on page 10.2).

One important remark is that it is not possible to determine whether our algorithm actually returns the lattice Λ or a proper sublattice of Λ . This is a problem of all such quantum algorithms, in particular the ones by HALLGREN and SCHMIDT AND VOLLMER. In case the infrastructure is obtained from a global field, checking whether the lattice computed by our algorithm is a sublattice of Λ can be done efficiently: one simply has to check whether the computed basis consists of units of the global field. However, even when one assumes that the Generalized Riemann Hypothesis holds, there is no efficient polynomial-time algorithm known which certifies that a given sublattice of Λ equals Λ . But we assume that the case that a basis returned by our algorithm (and any of the other algorithms, for that it matters) is a proper sublattice of Λ is somewhat pathological.

Note that the lower bound on the success probability is very small even for moderate n . More precisely, for $n = 1, \dots, 10$, the inverses of the probabilities, i.e., the expected number of iterations which have to be run, are bounded from above by

$$\begin{array}{ccccc}
1.40 \cdot 10^8, & 1.27 \cdot 10^{30}, & 4.67 \cdot 10^{59}, & 1.74 \cdot 10^{102}, & 6.47 \cdot 10^{158}, \\
1.39 \cdot 10^{230}, & 7.12 \cdot 10^{316}, & 2.92 \cdot 10^{419}, & 2.72 \cdot 10^{538}, & 1.43 \cdot 10^{674}.
\end{array}$$

(Note that for $n = 1$, we used the algorithm described in Proposition 10.2; the bound given by the formula in Theorem 10.1 is $1.26 \cdot 10^{12}$.) The success probability for the algorithm in [Sch07] is bounded from below by $2^{-20n^2 - 12n - 2n^{-4n^2}}$, as stated there in Satz 6.2.6. Hence, the expected number of iterations for $n = 1, \dots, 10$ for this algorithm are bounded by

$$\begin{array}{ccccc}
1.72 \cdot 10^{10}, & 5.32 \cdot 10^{36}, & 6.32 \cdot 10^{82}, & 8.18 \cdot 10^{149}, & 1.19 \cdot 10^{239}, \\
1.18 \cdot 10^{351}, & 3.45 \cdot 10^{486}, & 1.02 \cdot 10^{646}, & 9.05 \cdot 10^{829}, & 6.10 \cdot 10^{1038}.
\end{array}$$

Note that in [SV05] the success probability is given as $2^{-kn^{2+\varepsilon}}$ for some $k \in \mathbb{N}$ and $\varepsilon > 0$ without making these explicitly; the behavior for $n \rightarrow \infty$ will be similar to the analysis in [Sch07]. Finally, in [Hal05], no success probability is given at all. The current analyses can only prove expected running times which are impractical. Our analysis improves on the previous ones, though not substantially. We believe that it can be further optimized.

Assuming that n is constant, we obtain the following complexity theoretic result, which extends the results by HALLGREN and SCHMIDT AND VOLLMER to a larger class of infrastructures:

Corollary 10.3. *Assume that $n = O(1)$ and that \mathcal{I} is an infrastructure satisfying the assumptions **A1)–A3**). We obtain a quantum algorithm to compute Λ with a success probability bounded away from 0 by a constant which runs in time polynomial in $\log \det(\Lambda)$, $\log \frac{1}{\lambda_1(\Lambda)}$, $\log \frac{1}{\gamma}$, $\log A$, $\log \frac{1}{C}$ and $\log D$. \square*

Note that $\log L$, $\log N$, $\log N_0$ and $\log q$ can all be chosen to be *linear* in $\log \det(\Lambda)$, $\log \frac{1}{\lambda_1(\Lambda)}$, $\log \frac{1}{\gamma}$, $\log A$, $\log \frac{1}{C}$ and $\log D$.

Finally, we want to conclude with the proofs of Theorem 10.1 and Proposition 10.2.

Proof of Theorem 10.1. We have $C \leq 1$, $A \geq 1$, $N, N_0 \geq 32$, $q \geq \max\{32, 9A\}$ and $\kappa = \frac{1}{9n}$. Clearly, with $qN_0 \geq qN \geq 32^2 > 18$ we get assumption (V). Since $\frac{4}{A} \leq 4 \leq 32$ and since $N_0 \geq N \geq \frac{8(n+1)n2^n DA^{n-1}}{3C^n}$ we have assumption (II) for N and N_0 . The requirement $N_0 \geq 8n^2(n+1)N$ on N_0 is assumption (VIII).

Since $N_0 \geq N \geq \max\{8n - 2, n^{(n-1)/2}2^{n+1} - 2\} \cdot \frac{9n^2}{2\lambda_1(\Lambda)} + \frac{9}{64} \geq \max\{8n - 2, n^{(n-1)/2}2^{n+1} - 2\} \cdot \frac{9n^2}{2\lambda_1(\Lambda)} + \frac{9}{2q} \geq \frac{2\sqrt{n}}{\lambda_1(\Lambda)}$ we have assumptions (III) for N and N_0 as well as assumption (VI), and as $N \geq \frac{9n^2}{32} + \frac{18n^4}{\lambda_1(\Lambda)} \geq 9n(\frac{n}{q} + \frac{2n^3}{\lambda_1(\Lambda)})$ we get assumption (VII₁). Next, $q \geq 9A \geq 9$ yields assumption (II) for q . The third condition on q yields assumption (IV₁) using the bound $\nu(\Lambda) \leq \frac{1}{2}n^{(n+1)/2} \frac{\det(\Lambda)}{\lambda_1(\Lambda)^{n-1}}$. That bound follows by Theorem 7.9 in [MG02], stating that $\nu(\Lambda) \leq \frac{\sqrt{n}}{2}\lambda_n(\Lambda)$, and from

$$\lambda_n(\Lambda) \leq n^{n/2} \frac{\det(\Lambda)}{\prod_{i=1}^{n-1} \lambda_i(\Lambda)} \leq n^{n/2} \frac{\det(\Lambda)}{\lambda_1(\Lambda)^{n-1}}$$

by Minkowski's second theorem [MG02, Theorem 1.5].

The condition on L ensures that assumption (I), i.e. the hypotheses of Corollary 4.3, are satisfied. Hence, if $s \in S$ is uniformly picked, with probability at least $1/2$ we have $H^{\text{grid}}(1/(2NL)) \cap G(s) = \emptyset$, which guarantees that we can compute the function f for all $v \in \mathcal{V}$ exactly using **A3**).

Note that $\kappa = \frac{1}{9n}$ yields $c = \cos^2(\pi(\frac{1}{4} + \frac{1}{4qN} + 2\kappa n)) \geq \cos^2(\pi \frac{17417}{36864}) \geq 0.00746$ as $qN \geq 32^2$. Combining this with the bounds in Section 8.1.3 yields the lower bound

$$\frac{\cos(\pi \frac{17417}{36864})^{4n+2}}{22n+534n^2+2n_n4n^2+2n} p^* \geq \frac{1.239665 \cdot 1.54587777^n}{106n+581n^2n4n^2+2n}$$

for the probability that $2n + 1$ runs of the quantum algorithm (with fixed “good” s) yield a generating set of Λ^* ; here, $p^* \geq (\prod_{i=2}^{n+1} \zeta(i)^{-1} - \frac{1}{4}) \cdot \prod_{i=1}^{n-1} (1 - 2^{-i}) \geq (\hat{\zeta} - \frac{1}{4}) \cdot 0.289 \geq 0.184 \cdot \frac{1}{4}$ (compare Equation (*) on page 31). This has to be multiplied by $1/2$ for the above mentioned probability that a uniformly chosen $s \in S$ yields $H^{\text{grid}}(1/(2NL)) \cap G(s) = \emptyset$.

In the context of Corollary 9.12, we can bound α by $\sqrt{nb_0} = \sqrt{n}\kappa N_0 - \frac{\sqrt{n}}{2nq} \leq \frac{1}{9\sqrt{n}}N_0$, and $k = 2n + 1$ is the number of generating elements. When using Korkine-Zolotarev reduction, we can use $f = \frac{1}{2}\sqrt{2n+4}$. Since $L = \Lambda^*$, we see that $\det(L) = (\det(\Lambda))^{-1}$ and $\frac{1}{\lambda_1(L)} \leq \lambda_n(\Lambda) \leq \frac{n^{n/2} \det(\Lambda)}{\lambda_1(\Lambda)^{n-1}}$. This yields

$$g \leq \frac{13n^{3/2}(1 + \frac{5}{2n} + \frac{1}{n^2}) \det(\Lambda)^2 N_0^{n+1}}{6 \cdot 9^n \lambda_1(\Lambda)^{n-1}}.$$

Therefore, the algorithm in Section 9 computes a γ -approximate basis of Λ from a $\frac{1}{4\sqrt{nq}}$ -approximate generating set of $2n + 1$ vectors in Λ^* if

$$q \geq \max \left\{ \frac{19.5^n n^{n+3/2} (1 + \frac{5}{2n} + \frac{1}{n^2})^n N_0^{n^2+2n-1} \det(\Lambda)^{2n+1}}{2 \cdot 9^{n^2+2n-1} \lambda_1(\Lambda)^{n^2-n}}, \frac{19.5^{2n} n^{2n+3/2} (1 + \frac{5}{2n} + \frac{1}{n^2})^{2n-1} N_0^{2n^2+3n-3} \det(\Lambda)^{4n}}{\gamma \cdot 39 \cdot 9^{2n^2+3n-3} \lambda_1(\Lambda)^{2n^2-3n-1}} \right\}.$$

But this is satisfied by the fourth and fifth condition on q . \square

Proof of Proposition 10.2. We have $C \leq 1$, $A \geq 1$, $N \geq 32$, $q \geq \max\{32, 9A\}$ and $\kappa = \frac{1}{9}$. Clearly, with $qN \geq 32^2 > 18$ we get assumption (V). The second and third assumption on N yield the N -part of assumption (II), the fourth yields assumption (III) and (VII₂) and the fifth yields assumption (VI₂). The second assumption on q yields the q -part of assumption (II), and the third part yields assumption (IV₂). Note that $\lambda_1(\Lambda) = \det(\Lambda)$ and $\nu(\Lambda) = \frac{1}{2} \det(\Lambda)$.

Note that $\kappa = \frac{1}{9}$ yields $c = \cos^2(\pi(\frac{1}{4} + \frac{1}{4qN} + 2\kappa n)) \geq \cos^2(\pi \frac{17417}{36864}) \geq 0.00746$ as $qN \geq 32^2$. Combining this with the bounds in Section 8.2 yields the lower bound

$$\frac{1}{48} \kappa^2 c^2 \kappa^2 c^2 \geq \frac{\cos^4(\pi \frac{17417}{36864})}{48 \cdot 9^2}$$

for the probability that two runs of the quantum algorithm (with fixed s) yield a generating set of Λ^* . This has to be multiplied by $1/2$ for the above mentioned probability that a uniformly chosen $s \in S$ yields $H^{\text{grid}}(1/(2NL)) \cap G(s) = \emptyset$.

In the context of Corollary 9.12, we can bound α by $\frac{1}{9}N$, and $k = 2$ is the number of generating elements. Since in dimension one, one can reduce perfectly, we can use $f = 1$. Since $L = \Lambda^*$, we have $\det(L) = (\det(\Lambda))^{-1}$ and $\lambda_1(L) = \det(L) = (\det(\Lambda))^{-1}$. Using this, the algorithm in Section 9 computes a γ -approximate basis of Λ from a $\frac{1}{4q}$ -approximate generating set of two vectors in Λ^* if

$$q \geq \frac{19.5}{9^2} N^2 \det(\Lambda)^3 \cdot \max \left\{ 1, \frac{\det(\Lambda)}{\gamma} \right\}$$

But this is satisfied by the last condition on q . □

List of assumptions

Assumption	Page	Can be found in
(I)	16	Corollary 4.3
$L \geq \frac{4nD(q+A+C+2)^n}{C^n}$ and $\varepsilon \leq \frac{1}{2NL}$		
(II)	18	Corollary 5.2
$q \geq 9 \max\{1, A\}$ and $N \geq \max\left\{\frac{4}{A}, \frac{8(n+1)n \cdot 2^n DA^{n-1}}{3C^n}\right\}$		
(III)	19	Proposition 5.4
$N \geq \frac{2\sqrt{n}}{\lambda_1(\Lambda)}$		
(IV)	19	Proposition 5.4
$q > 2n\nu(\Lambda) + \frac{3n}{N}$		
(IV ₁)	35	Section 8.1.3
$q \geq \frac{6n^2}{N} + 4n(n+1)\nu(\Lambda)$		
(IV ₂)	36	Section 8.2
$q \geq \frac{12}{N} + 4 \det(\Lambda)$		
(V)	21	Proposition 6.1
$\kappa < \frac{1}{8n} - \frac{1}{4nqN}$		
(VI)	32	Lemma 8.1
$N \geq \frac{1}{\kappa} \left(\max\{8n-2, n^{(n-1)/2} \cdot 2^{n+1} - 2\} \cdot \frac{n}{2\lambda_1(\Lambda)} + \frac{1}{2nq} \right)$		
(VI ₂)	36	Section 8.2
$N \geq \frac{1}{\kappa} \left(\frac{3}{\det(\Lambda)} + 1 + \frac{1}{2q} \right)$		
(VII)	32	Lemma 8.1
$N > \frac{1}{\kappa} \left(\frac{1}{2q} + \frac{n^2}{\lambda_1(\Lambda)} \right)$		
(VII ₁)	35	Section 8.1.3
$N \geq \frac{1}{\kappa} \left(\frac{n}{q} + \frac{2n^3}{\lambda_1(\Lambda)} \right)$		
(VII ₂)	36	Section 8.2
$N \geq \frac{1}{\kappa} \left(\frac{2}{q} + \frac{4}{\det(\Lambda)} \right)$		
(VIII)	34	Lemma 8.2
$N_0 \geq 8n^2(n+1)N$		

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 the support from Armasuisse and the SNF grant No. 132256.

References

- [Bar] A. Barvinok. Math669: Combinatorics, geometry and complexity of integer points. <http://www.math.lsa.umich.edu/~barvinok/latticenotes669.pdf>.
- [BJP94] J. Buchmann, M. Jüntgen, and M. Pohst. A practical version of the generalized Lagrange algorithm. *Experiment. Math.*, 3(3):199–207, 1994.
- [BK93] J. Buchmann and V. Kessler. Computing a reduced lattice basis from a generating set. http://www.cdc.informatik.tu-darmstadt.de/reports/reports/reduced_basis.ps.gz, 1993.
- [BP89] J. Buchmann and M. Pohst. Computing a lattice basis from a system of generating vectors. *Proceedings of EUROCAL 1987, Lecture Notes in Computer Science*, 378, 1989.
- [Buc87a] J. A. Buchmann. On the computation of units and class numbers by a generalization of Lagrange’s algorithm. *J. Number Theory*, 26(1):8–30, 1987.
- [Buc87b] J. A. Buchmann. On the period length of the generalized Lagrange algorithm. *J. Number Theory*, 26(1):31–37, 1987.
- [Buc87c] J. A. Buchmann. Zur Komplexität der Berechnung von Einheiten und Klassenzahl algebraischer Zahlkörper. Habilitationsschrift, October 1987.
- [CM01] K. K. H. Cheung and M. Mosca. Decomposing finite abelian groups. *Quantum Information & Computation*, 1(3):26–32, 2001.
- [Die08] C. Diem. On arithmetic and the discrete logarithm problem in class groups of curves. Habilitationsschrift. Available at <http://www.math.uni-leipzig.de/~diem/preprints/english.html>, May 2008.
- [EH12] K. Eisenträger and S. Hallgren. Computing the unit group, class group and compact representations in algebraic function fields. To be presented at ANTS X., 2012.
- [Fon11] F. Fontein. The infrastructure of a global field of arbitrary unit rank. *Math. Comp.*, 80(276):2325–2357, 2011.

- [Hal02] S. Hallgren. 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*, pages 653–658 (electronic), New York, 2002. ACM.
- [Hal05] S. Hallgren. 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*, pages 468–474. ACM, New York, 2005.
- [Hes02] F. Hess. Computing Riemann-Roch spaces in algebraic function fields and related topics. *J. Symbolic Comput.*, 33(4):425–445, 2002.
- [Ked06] K. S. Kedlaya. Quantum computation of zeta functions of curves. *Computational Complexity*, 15(1):1–10, 2006.
- [MG02] D. Micciancio and S. Goldwasser. *Complexity of lattice problems*. The Kluwer International Series in Engineering and Computer Science, 671. Kluwer Academic Publishers, Boston, MA, 2002. A cryptographic perspective.
- [Pom01] C. Pomerance. The expected number of random elements to generate a finite abelian group. *Periodica Mathematica Hungarica*, 43(1–2):191–198, 2001.
- [Sch07] A. Schmidt. *Zur Lösung von zahlentheoretischen Problemen mit klassischen und Quantencomputern*. Ph.D. thesis, Technische Universität Darmstadt, 2007.
- [Sch08] R. J. Schoof. *Computing Arakelov class groups*, volume 44 of *MSRI Publications*, pages 447–495. Cambridge University Press, Cambridge, 2008.
- [Seq] Integer sequence A021002. The on-line encyclopedia of integer sequence <http://oeis.org/A021002>.
- [SS90] G. W. Stewart and J. G. Sun. *Matrix perturbation theory*. Academic Press, Inc., 1990.
- [SV05] A. Schmidt and U. Vollmer. 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*, pages 475–480. ACM, New York, 2005.
- [SW11] P. Sarvepalli and P. Wocjan. Quantum algorithms for one-dimensional infrastructures. <http://arxiv.org/abs/1106.6347>, 2011.
- [Thi95a] C. Thiel. *On the complexity of some problems in algorithmic algebraic number theory*. Ph.D. thesis, Universität des Saarlands, 1995.

- [Thi95b] C. Thiel. Short proofs using compact representations of algebraic integers. *J. Complexity*, 11(3):310–329, 1995.