On the Quantum Complexity of the Continuous Hidden Subgroup Problem *

Koen de Boer[1], Léo Ducas[1], and Serge Fehr[1,2]

[1] Cryptology Group, Centrum Wiskunde & Informatica (CWI), Amsterdam, The Netherlands
[2] Mathematical Institute, Leiden University, The Netherlands

Abstract. The Hidden Subgroup Problem (HSP) aims at capturing all problems that are susceptible to be solvable in quantum polynomial time following the blueprints of Shor's celebrated algorithm. Successful solutions to this problems over various commutative groups allow to efficiently perform number-theoretic tasks such as factoring or finding discrete logarithms.
The latest successful generalization (Eisenträger et al. STOC 2014) considers the problem of finding a full-rank lattice as the hidden subgroup of the continuous vector space $\mathbb{R}^m$, even for large dimensions m. It unlocked new cryptanalytic algorithms (Biasse-Song SODA 2016, Cramer et al. EUROCRYPT 2016 and 2017), in particular to find mildly short vectors in ideal lattices.
The cryptanalytic relevance of such a problem raises the question of a more refined and quantitative complexity analysis. In the light of the increasing physical difficulty of maintaining a large entanglement of qubits, the degree of concern may be different whether the above algorithm requires only linearly many qubits or a much larger polynomial amount of qubits.
This is the question we start addressing with this work. We propose a detailed analysis of (a variation of) the aforementioned HSP algorithm, and conclude on its complexity as a function of all the relevant parameters. Our modular analysis is tailored to support the optimization of future specialization to cases of cryptanalytic interests. We suggest a few ideas in this direction.

Keywords: Quantum Algorithm, Hidden Subgroup, Period Finding, Fourier Transform, Cryptanalysis.

1 Introduction

The Hidden Subgroup Problem. Among all quantum algorithms, Shor's algorithm [32] for factoring and finding discrete logarithms stands out as demon-

* All three authors were supported by the European Union H2020 Research and Innovation Program Grant 780701 (PROMETHEUS). Additionally, K.d.B. was supported by the ERC Advanced Grant 740972 (ALGSTRONGCRYPTO) and L.D. was supported by the Veni Innovational Research Grant from NWO under project number 639.021.645.

strating the largest complexity gap between classical and quantum computing. It is also singular by its cryptanalytic implications, and, due to progress toward the realization of large quantum computers, this celebrated algorithm is now motivating the standardization of quantum-resistant schemes [23], in preparation of a global update of widely deployed encryption and authentication protocols.

The core idea of quantum period finding from [32] is not limited to factoring and discrete logarithm, and the Hidden Subgroup Problem formalized in [22] serves as a convenient interface between the quantum-algorithmic techniques for period finding, and applications to solve other computational problems, in particular problems arising from number theory. We will here discuss only the case of commutative groups. The cases of non-abelian groups such as dihedral groups are very interesting as well and have fascinating connections with lattice problems [29]; however, no polynomial time algorithm is known for those cases, and the best known algorithm has sub-exponential complexity [19], using very different techniques.

The simplest version of the Hidden Subgroup Problem consists of finding a hidden subgroup H in a *finite* abelian group G, when given access to a strictly H-periodic function $\boldsymbol{f} : G \to R$. Here, in the language of representation theory, the off-the-shelf period-finding quantum algorithm finds a uniformly random character $\chi \in \hat{G}$ that acts trivially on H. Shor's original algorithm [32] for integer factoring finds a hidden subgroup H in the ambient group $\mathbb{Z}$. The infiniteness of $\mathbb{Z}$ induces some "cut-off" error; nevertheless, the distribution of the algorithm's output is still concentrated around the multiples of the inverse period.

A generalization to the real line $H = \mathbb{R}$ was given by Hallgren [16] and allows to solve Pell's equation. The case of real vector space of constant dimension $H = \mathbb{R}^c$ has also been studied in [15,31], and permits the computation of unit groups of number fields of finite degree.

The *Continuous* Hidden Subgroup Problem in large dimension. The latest generalization of the HSP algorithm, given by Eisenträger, Hallgren, Kitaev and Song in an extended abstract [11], targets the ambient group $G = \mathbb{R}^m$ (for a non-constant dimension m) with a hidden discrete subgroup $H = \Lambda$, i.e. a *lattice*. Next to the ambient group $\mathbb{R}^m$ being *continuous*, an additional special feature is that the Λ-periodic function $\boldsymbol{f}$ is assumed to produce a "quantum output". More formally, $\boldsymbol{f} : \mathbb{R}^m \to \mathcal{S}$, $x \mapsto |\boldsymbol{f}(x)\rangle$, where $\mathcal{S}$ is the state space of a quantum system, and the HSP algorithm is given access to a unitary that maps $|x\rangle|0\rangle$ to $|x\rangle|\boldsymbol{f}(x)\rangle$. A crucial observation here is that $|\boldsymbol{f}(x)\rangle$ and $|\boldsymbol{f}(y)\rangle$ are *not* necessarily orthogonal (or even distinct) for distinct x and y modulo Λ. In other words, it is not assumed that $\boldsymbol{f}$ is *strictly* periodic, but merely that $|\boldsymbol{f}(x)\rangle$ and $|\boldsymbol{f}(y)\rangle$ are "somewhat orthogonal" for x and y that are "not too close" modulo Λ, and that $\boldsymbol{f}$ is Lipschitz continuous.

Mor specifically. they consider a variation of the standard HSP algorithm in order to tackle the Continuous Hidden Subroup Problem (CHSP). In order to deal with the continuous nature of the domain $\mathbb{R}^m$ of $\boldsymbol{f}$, the given HSP algorithm acts on a bounded "grid" of points within $\mathbb{R}^m$. Additionally, the algorithm

is modified in the following ways: (1) The initial state is not a uniform superposition (over the considered grid points in $\mathbb{R}^n$) but follows a trigonometric distribution, and (2) the quantum Fourier transform is done "remotely", i.e., rather than applying it to the actual register, the register is entangled with an ancilla and the quantum Fourier transform is then applied to the ancilla instead. According to [11], applying the quantum Fourier transform directly would make the resulting approximation errors difficult to analyze.

As an application, the work of [11] also gave a quantum polynomial time algorithm for computing the unit group of a number field in their article [11]. This was generalized by Biasse and Song [2] to the computation of S-unit groups, and therefore to the computation of class groups and to finding a generator of a principal ideals. This led to solving the short vector problem in certain ideal lattices for non-trivial approximation factors [7,8,27]. While the cryptanalytic consequences for ideal-lattice based cryptography seems limited so far [10], these results demonstrate a hardness gap between ideal lattices and general ones.

The algorithm of [11] has proved itself to be a key tool in quantum cryptanalysis, and, as such, the question of its precise range of application, and of its practical efficiency are therefore of cryptographic interest. Unfortunately, [11] offers only an informal treatment of the algorithm, both in terms of the analysis and in terms of the formulation of the result. Also, at the time of preparing this article, there was no full version publicly available.[3]

The extended abstract [11] explains convincingly that in the limit of choosing an unbounded and infinitely fine grid in $\mathbb{R}^m$ the algorithm does what it is supposed to do; however, the "rate of convergence" and thus the quantitative aspects of their result are not provided. Furthermore, it was not clear to us what "polynomial-time" formally meant when the input is an oracle, specified by various parameters. For example, in an application of the Continuous HSP algorithm it may be critical to know whether the running time grows polynomially in the Lipschitz constant of f (which is one of the 3 parameters of the Continuous HSP), or polynomially in its logarithm.

In an email from September 2018, Fang Song [33] partially answered early questions we had; technically his comments corresponds to a claim on the error term ϵ_{lip} in Part 2 Step 2 of our analysis of the Dual Lattice Sampling step (Section 5.2). We found that this claim could be related to Yudin-Jackson Theorem [37]. To make the analysis tighter, we found it preferable to generalize Yudin-Jackson Theorem to multi-dimensional ranges (see Appendix D of the full version [3]).

The urge to understand the security post-quantum cryptography motivates the elevation of the powerful result of [11] into an open and lively research topic.

Our work. The goal of this paper is to provide a complete, modular, and quantitative analysis of (a slightly modified version of) the Continuous HSP

[3] The STOC 2014 submitted version [11] has been made publicly available online on November 2019 (after submissison of this paper) `http://www.cse.psu.edu/~sjh26/units-stoc-submission.pdf` . A full version is announced to be in preparation.

quantum algorithm given by [11]. More concretely, we provide an explicit bound on the number of qubits needed by the algorithm, clarifying the dependency on the parameters of the Continuous HSP instance and on the required precision and success probability. This shows explicitly in what parameters the algorithm is polynomial time and with what exponent.

The algorithm that we consider and analyze differs from the one considered [11] in the following points:

- First, we specify the initial state of the algorithm to have Gaussian amplitudes, while [11, Sec. 6.2] suggests to use a cropped trigonometric function; as far as we can see, our choice makes the analysis simpler and tighter thanks to the well known tail-cut and smoothness bounds of Banaszczyk [1] and Micciancio and Regev [20].
- Secondly, we do not make use of a "remote" Fourier transform but instead follow the blueprint of Shor's original algorithm in that respect; the claimed advantage of the "remote" Fourier transform is unclear to us.

These modifications simplify the algorithm and its analysis. Due to the lack of details given in [11], we can not state a complexity comparison, but we think this variation is at least as efficient as the original algorithm.

Our analysis is divided into four parts, each summarized by a formal statement given in Sections 2.3 to 2.6, leading to the main theorem (Section 2.2). We insist on this modular presentation, so as to enable future work on optimization and specialization of this algorithm to instances of interests; specific suggestions follow.

In the first part (*Dual Lattice Sampling*), which is the technically more involved one, we show that the appropriately discretized and finitized, but otherwise (almost) standard HSP quantum algorithm produces sample points in $\mathbb{R}^m$ that lie close to the dual lattice Λ^* with high probability. More precisely, and more technically speaking, we show that the algorithm's output is a sample point close to $\ell^* \in \Lambda^*$ with probability close to $\langle c_{\ell^*} | c_{\ell^*} \rangle$, where the vectors $|c_{\ell^*}\rangle$ are the Fourier coefficients of the function $\boldsymbol{f}$. This is in line with the general HSP approach, where for instance Shor's algorithm for period finding over $\mathbb{Z}$ produces a point that is close to a random multiple of the inverse period, except with bounded probability.

In this first part (Section 4 and Section 5), we bound the complexity of the core algorithm in terms of the error that we allow in the above context of a sampling algorithm, and depending on the Lipschitz constant of $\boldsymbol{f}$. In particular, we show that the number of qubits grows as mQ, where Q, the "number of qubits per dimension", grows linearly in the logarithm of the Lipschitz constant of $\boldsymbol{f}$, the logarithm of the inverse of the error probability and the logarithm of the inverse of the (absolute) precision, and quasi-linearly in m. The running time of the algorithm is then bounded[4] by $O(m^2Q^2)$.

[4] This complexity estimate can be lowered to $O(mQ\log(kmQ))$ if we allow an error in the L_2-distance of $< 1/k^2$ [14], see Remark 1.

In the second part (*Full Dual Recovery*, Section 6), we then relate the parameters of the Continuous HSP instance to the number of sample points, and thus to how often the core algorithm needs to be repeated, necessary in order to have an approximation of the entire dual lattice Λ^*.

In the third part (*Primal Basis Reconstruction*, see Appendix B of the full version [3]), we study the numerical stability of reconstructing an approximate basis of the primal lattice Λ from a set of approximate generators of the dual lattice Λ^*. This is based on the Buchmann-Pohst algorithm [4] already mentioned in [11]. The claim of [11] involves intricate quantities related to sublattices of Λ, making the final complexity hard to derive; we provide a simpler statement with a detailed proof.

Finally, in the last part (see Appendix C of the full version [3]), we revisit the quantum poly-time algorithm for *Gaussian State Preparation* [13,18] used as a black-box in our first part, and provide its precise complexity.

These four parts leads to our formal and quantitative version of the informal CHSP Theorem of [11, Theorem 6.1], stated as Theorem 1 in Section 2.2.

Conclusion and Research Directions. Our conclusion is that, in its generic form, the Continuous Hidden Subgroup Problem is rather expensive to solve; not accounting for other parameters than the dimension m, it already requires $\tilde{O}(m^3)$ qubits and $\tilde{O}(m^7)$ quantum gates (or, $\tilde{O}(m^4)$ quantum gates if an approximate quantum Fourier transform is used). However, this inefficiency seems to be a consequence of its genericness. In particular, the core algorithm for Dual Lattice Sampling would only need $\tilde{O}(m^2)$ qubits, if it wasn't for accommodating for the terrible numerical stability of the Primal Basis Reconstruction step. Similarly, we expect the number of samples needed to generate the dual lattice to be significantly smaller for smoother oracle functions.

All in all, our modular analysis of the generic steps of the CHSP algorithm sets the stage for analyzing and optimizing its specializations, in particular to cryptanalytic applications [7,8]. We propose as few research directions towards this objective:

- Study the costs (qubits, quantum gates) and the parameters of the oracle functions from [11,2,34] for solving the Unit Group Problem, the Principal Ideal Problem (PIP), and for the computation of the class-group.
- Find stronger hypotheses satisfied by the above oracle functions (or by variant thereof) that improve this generic analysis of the CHSP algorithm; or resort to an ad-hoc analysis of the Full Dual Recovery step by directly studying the spectrum of these oracle functions.
- Explore the possibility of a trade-off between the (classical) Primal Basis Reconstruction step and the (quantum) Dual Lattice Sampling step, possibly up to small sub-exponential classical complexity. More specifically, does replacing LLL by BKZ with an medium block-size substantially improve the numerical stability of Buchmann-Pohst algorithm?
- Exploit prior knowledge of sublattices (potentially close to full-rank) of the hidden lattice to accelerate or skip the Full Dual Recovery and Primal Basis

Reconstruction steps. This is for example the case when solving PIP [2] while already knowing the unit group and the class group of a given number field. This would be applicable in the context of [7,8].
– Exploit known symmetries of the hidden sublattice to improve the Full Dual Recovery and Primal Basis Reconstruction steps. Such symmetries are for example induced by the Galois action on the log-unit lattice and the lattice of class relation, in particular in the case of the cyclotomic number fields. This would again be applicable in the context of [7,8].

Acknowledgments. We would like to thank Stacey Jeffery, Oded Regev and Ronald de Wolf for helpful discussions on the topic of this article.

2 Problem Statements and Results

2.1 Notation and Set-Up

Here and throughout the paper, $\mathcal{H}$ is a complex Hilbert space of dimension $N = 2^n$, and $\mathcal{S}$ is the unit sphere in $\mathcal{H}$; thus, a vector in $\mathcal{S}$ describes the state of a system of n qubits. For an arbitrary positive integer m, we consider a function

$$\boldsymbol{f} : \mathbb{R}^m \to \mathcal{S} \subset \mathcal{H}\,,\ x \mapsto |\boldsymbol{f}(x)\rangle$$

that is periodic with respect to a full rank lattice $\Lambda \subset \mathbb{R}^m$; hence, $\boldsymbol{f}$ may be understood as a function $\mathbb{R}^m/\Lambda \to \mathcal{S}$. The function $\boldsymbol{f}$ is assumed to be Lipschitz continuous with Lipschitz constant

$$\mathrm{Lip}(\boldsymbol{f}) = \inf\{L > 0 \mid \| |\boldsymbol{f}(x)\rangle - |\boldsymbol{f}(y)\rangle \|_{\mathcal{H}} \leq L \, \|x - y\|_{2,\mathbb{T}^m}\}.$$

Later, we will also require $\boldsymbol{f}$ to be "sufficiently non-constant". One should think of $\boldsymbol{f}$ as an oracle that maps a classical input x to a quantum state over n qubits, which is denoted $|\boldsymbol{f}(x)\rangle$.

We write Λ^* for the dual lattice of Λ. By $\lambda_1(\Lambda)$ we denote the length of a shortest non-zero vector of Λ, and correspondingly for $\lambda_1(\Lambda^*)$. Since Λ is typically clear from the context, we may just write λ_1 and λ_1^* instead of $\lambda_1(\Lambda)$ and $\lambda_1(\Lambda^*)$.

We denote by $\mathcal{B}_r(x) = \{y \in \mathbb{R}^m \mid \|y - x\| < r\}$ the open Euclidean ball with radius r around x, and by $B_r(x) = \mathcal{B}_r(x) \cap \mathbb{Z}^m$ its integer analogue. For the open ball around 0 we just denote $\mathcal{B}_r$, and for a set $X \subset \mathbb{R}^m$ we write $\mathcal{B}_r(X) = \bigcup_x \mathcal{B}_r(x)$ and $B_r(X) = \bigcup_x B_r(x)$ where the union is over all $x \in X$.

Definition 1 (Definition 1.1 from [11]). *A function $\boldsymbol{f} : \mathbb{R}^m \to \mathcal{S} \subset \mathcal{H}$ is said to be an (a, r, ϵ)-HSP oracle of the full-rank lattice $\Lambda \subset \mathbb{R}^m$ if*

– *$\boldsymbol{f}$ is Λ-periodic,*
– *$\boldsymbol{f}$ is a-Lipschitz:* $\mathrm{Lip}(f) \leq a$,
– *For all $x, y \in \mathbb{R}^m$ such that $d_{\mathbb{R}^m/\Lambda}(x, y) \geq r$, it holds that $|\langle \boldsymbol{f}(x) | \boldsymbol{f}(y) \rangle| \leq \epsilon$,*

where $d_{\mathbb{R}^m/\Lambda}(x, y) = \min_{v \in \Lambda} \|x - y - v\|$ denotes the distance induced by the Euclidean distance of $\mathbb{R}^n$ modulo Λ.

2.2 Main Theorem: Continuous Hidden Subgroup Problem

Theorem 1. *There exists a quantum algorithm that, given access to an (a, r, ϵ)-HSP oracle with period lattice Λ, $r < \lambda_1(\Lambda)/6$ and $\epsilon < 1/4$, computes, with constant success probability, an approximate basis $\tilde{B} = B + \Delta_B$ of this lattice Λ, satisfying $\|\Delta_B\| < \tau$.*

This algorithm makes k quantum oracle calls to the (a, r, ϵ)-HSP oracle, and uses $mQ+n$ qubits, $O(km^2Q^2)$ quantum gates and $poly(m, \log \frac{a}{\lambda_1^*}, \log \frac{a}{\tau})$ classical bit operations, where

$$Q = O(mk) + O\left(\log \frac{a}{\lambda_1^*}\right) + O\left(\log \frac{1}{\lambda_1^* \cdot \tau}\right), \tag{1}$$

$$k = O\left(m \cdot \log\left(\sqrt{m} \cdot a \cdot (\det \Lambda)^{1/m}\right)\right) \tag{2}$$

Remark 1. The quantum gate complexity in this theorem can be lowered to $O(kmQ \log(kmQ))$ if we *approximate* the quantum Fourier transform [14] over $\mathbb{Z}/q^m\mathbb{Z}$. For example, an approximation that is $1/k^2$-close in the induced matrix norm – which is sufficient for our purposes – can be computed using $O(mQ \log(kmQ))$ quantum gates (where $Q = \log q$). Repeating this approximate Fourier transform k times, one arrives at the complexity $O(kmQ \log(kmQ))$.

Remark 2. Note that the quantities inside logarithms are homogeneous. In particular, scaling the lattice Λ by a factor f, also scales τ, $1/a$ and $1/\lambda_1^*$ by the same factor f, leaving the complexity parameters Q and k unaffected.

Remark 3. The expert reader may expect the "distortion" parameter $\lambda_1 \cdot \lambda_1^*$ of the lattice Λ to have a bearing on the complexity of this algorithm. It is indeed implicitly the case: the assumption the HSP definition implies that $ar \geq 1 - \epsilon^2$, and therefore the theorem's hypothesis requires $a \geq \frac{45}{8\lambda_1}$.

Proof. This is obtained by instantiating Theorems 2 to 5. First, we obtain k samples close to the dual lattice by invoking k times Algorithm 1, whose correctness and complexity is given in Theorem 2. Samples whose Euclidean length exceed a certain threshold R are rejected. The approximate samples are collected into a matrix $\tilde{G}$.

The above step requires to prepare Gaussian states with parameter s over a grid of granularity q; this is obtained by k calls to Algorithm 1, whose cost and correctness is stated in Theorem 5. The cost of this subroutine is dominated by the cost of Algorithm 1.

According to Theorem 3, the approximated dual samples generate the dual lattice Λ^* with constant probability. Finally, one applies the Buchmann-Pohst algorithm [5,4] and matrix inversion to $\tilde{G}$, in order to recover an approximate basis of the primal lattice Λ. The loss of precision induced by this computation is given in Theorem 4. The parameters are instantiated as follows:

- the failure probability η of dual lattice sampling is set to $\eta = 1/k^2$,
- the parameter α driving the success of dual reconstruction is set to $\alpha = 1$,

- the relative error on dual lattice sample is set to

$$\delta = \frac{(\lambda_1^*)^2 \cdot \det(\Lambda^*)}{2^{O(mk)} \cdot \|\tilde{G}\|_\infty^{m+1}} \cdot \tau,$$

- the maximal entry size of the dual samples is $\|\tilde{G}\|_\infty \leq R$ where $R = \sqrt{m} \cdot a$,
- the discretization granularity is set to $q = 2^Q$,
- the Gaussian windowing parameter s is set to $s = O(\sqrt{m \log(\eta^{-1})})$.

We defer the detailed bookkeeping for deriving the parameters Q and k to Appendix A of the full version [3]. □

2.3 Dual Lattice Sampling Problem

Following our modular approach as outlined in the introduction, we first consider the following *Dual Lattice Sampling Problem* instead. Informally, the task is to sample points in $\mathbb{R}^m$ that are respectively close to points $\ell^* \in \Lambda^*$ that follow the distribution $\mathcal{D}_{ideal}(\ell^*) = \langle c_{\ell^*} | c_{\ell^*} \rangle$, where $|c_{\ell^*}\rangle$ are the vectorial Fourier coefficients of $\boldsymbol{f} : \mathbb{R}^m/\Lambda \to \mathcal{S}$ (see Section 3).

Problem 1 (Dual Lattice Sampling Problem). Given error parameter $\eta > 0$ and a relative distance parameter $\frac{1}{2} > \delta > 0$, and given oracle access to an HSP oracle $\boldsymbol{f}$ as above, sample according to a (finite) distribution $\mathcal{D}$ on $\mathbb{R}^m$ that satisfies, for any $S \subseteq \Lambda^*$,

$$p_S := \mathcal{D}\big(\mathcal{B}_{\delta\lambda_1^*}(S)\big) \geq \left(\sum_{\ell^* \in S} \langle c_{\ell^*} | c_{\ell^*} \rangle \right) - \eta\,. \tag{3}$$

In the problem statement above, $\mathcal{D}\big(\mathcal{B}_{\delta\lambda_1^*}(S)\big)$ denotes the cumulative weight of the set $\mathcal{B}_{\delta\lambda_1^*}(S)$ with respect to the distribution $\mathcal{D}$.

Theorem 2. *Algorithm 1 solves the Dual Lattice Sampling Problem with parameters η and δ; it uses m calls to the Gaussian superposition subroutine (see Theorem 5), one quantum oracle call to $\boldsymbol{f}$, $mQ + n$ qubits, and $O(m^2Q^2)$ quantum gates, where*

$$Q = O\left(m \log\left(m \log \frac{1}{\eta}\right)\right) + O\left(\log\left(\frac{\mathrm{Lip}(\boldsymbol{f})}{\eta \cdot \delta\lambda_1^*}\right)\right). \tag{4}$$

Remark 4. Note that this step only requires smoothness of the HSP oracle (via the Lipchitz constant), but does not rely on the "separability" assumption (third item of Definition 1). Indeed this third assumption will only play a role to ensure that those samples are actually non-trivial and usable.

2.4 Full Dual Lattice Recovery

Recovering the full lattice (or equivalently its dual) requires an extra assumption on the oracle function $\boldsymbol{f}$, as captured by the third condition in the following definition, reformatted from Definition 1.1 of [11].

According to Eisenträger et al. [11], for (some undetermined) adequate parameters, Definition 1 ensures that the distribution on the dual lattice Λ^* is not concentrated on any proper sublattice, hence sufficiently many samples will generate the lattice fully. We formalize and quantify this proof strategy, and obtain the following quantitative conclusion. We note that the constraints on r and ϵ are milder that one could think, for example ϵ does not need to tend to 0 as a function of n or m.

Theorem 3. *Let $\boldsymbol{f} : \mathbb{R}^m \to \mathcal{S}$ be an (a, r, ϵ)-HSP oracle with $r \leq \lambda_1(\Lambda)/6$ and $\epsilon \in [0, 1/3)$, and let $\mathcal{D}_{\text{ideal}}$ be the distribution described above, given by $\mathcal{D}_{\text{ideal}}(\ell^*) = \langle c_{\ell^*} | c_{\ell^*} \rangle$ for $\ell^* \in \Lambda^*$. Furthermore, denote by S the random variable defined by the number of samples that need to be drawn from $\mathcal{D}_{\text{ideal}}$ such that the samples together generate Λ^* as a lattice. Then, for any $\alpha > 0$,*

$$\Pr\left[S > (2+\alpha)\frac{t+m}{\frac{1}{2} - \frac{1}{4\pi^2} - \epsilon}\right] \leq \exp(-\alpha(t+m)/2)$$

where $t = m \log_2(\sqrt{m} \cdot a) + \log_2(\det(\Lambda))$.

The above Theorem is obtained by combining Lemmata 5 and 8 from Section 6, instantiating the parameter R to $R^2 = ma^2$. This choice is somewhat arbitrary and given for concreteness, however it does not have a critical quantitative impact.

2.5 Primal Basis Reconstruction

Theorem 4. *There exists a polynomial time algorithm, that, for any matrix $G \in \mathbb{R}^{k \times m}$ of k generators of a (dual) lattice Λ^*, and given an approximation $\tilde{G} = G + \Delta_G \in \mathbb{Q}^{k \times n}$, computes an approximation $\tilde{B} = B + \Delta_B$ of a basis B of the primal lattice Λ, such that*

$$\|\Delta_B\|_\infty \leq \frac{2^{O(mk)} \cdot \|\tilde{G}\|_\infty^{m+1}}{(\lambda_1^*)^3 \cdot \det(\Lambda^*)} \cdot \|\Delta_G\|_\infty,$$

under the assumption that $\|\Delta_G\|_\infty < \frac{\min(1, (\lambda_1^*)^2) \cdot \det(\Lambda^*)}{2^{O(km)} \cdot \|\tilde{G}\|_\infty^{m+1}}$.

Remark 5. More specifically, the algorithm from Theorem 4 essentially consists of the Buchmann-Pohst algorithm [5,4] and a matrix inversion. Its complexity is dominated by two calls to LLL on matrices of dimension $(m + k) \times k$ and entry bitsize $O(k^2 \log(\|\tilde{G}\|/\lambda_1^*))$ (see the discussion before [4, Cor. 4.1]). One can optimize the final running time by choosing the adequate variant of LLL [24,26] depending on the relative dimension and bitsizes of these inputs.

Our contribution on this step is merely a completed numerical analysis, with the help of a theorem from [6]. A claim with a similar purpose is given in [11], yet involves more intricate lattice quantities.

2.6 Gaussian State Preparation

The main algorithm of this paper requires the preparation of a multidimensional Gaussian initial state, which can be obtained by generating the one-dimensional Gaussian state on m parallel quantum registers. This task is known to be polynomial time [13,18], and we provide a quantitative analysis in Appendix C of the full version [3]. The precise running time of preparing this Gaussian state is summarized below.

Theorem 5. *For any positive integers q, p and for any $s > 1$, there exists a quantum algorithm that prepares the one-dimensional Gaussian state*

$$\frac{1}{\sqrt{\rho_{1/s}(\frac{1}{q}[q]_c)}} \sum_{x \in \frac{1}{q}[q]_c} \sqrt{\rho_{1/s}(x)}|x\rangle \tag{5}$$

up to trace distance $se^{-\pi s^2/8} + Q \cdot 2^{-p}$ using $O(Q+p)$ qubits and $O(Q \cdot p^{3/2} \cdot \mathrm{polylog}(p))$ quantum gates, where $Q = \log(q)$ and $\frac{1}{q}[q]_c = [-\frac{1}{2}, \frac{1}{2}) \cap \frac{1}{q}\mathbb{Z}$.

The above theorem is obtained by instantiating Theorem 12 in Appendix C of the full version [3] with parameters $\mu = q/2$, $k = p$ and $\sigma = \sqrt{2}q/s$ and relabeling the basis states. Whenever above theorem is used as a subroutine in Theorem 2, choosing $p = \log(mQ/\eta^2)$ is sufficient, causing merely an extra error of η^2.

Remark 6. In Theorem 1, we chose η to be $1/k^2$, yielding $p = \log(mk^4Q)$. Therefore, one call to the one-dimensional Gaussian state preperation with the parameters of Theorem 1 takes $O(Q)$ qubits and $O(Q\log(kmQ))$ quantum gates. As Theorem 1 requires k subsequent preparations of the m-dimensional Gaussian state, the total costs of the Gaussian state preparation steps are $O(mQ)$ qubits and $\tilde{O}(kmQ)$ quantum gates. As this is negligible to the overall complexity of Theorem 1, we can ignore these costs.

3 Preliminaries

We start with a brief introduction to Fourier analysis over arbitrary locally compact Abelian groups. Our general treatment allows us to then apply the general principles to the different groups that play a role in this work. For the reader that is unfamiliar with such a general treatment, it is useful — and almost sufficient — to think of $\mathbb{R}$, of $\mathbb{T} = \mathbb{R}/\mathbb{Z}$, and a finite group. For more details and for the proofs we refer to [9].

3.1 Groups

Here and below we consider a *locally compact Abelian* group G. Such a group admits a *Haar measure* μ that is unique up to a normalization factor. The crucial property of such a Haar measure is that it is invariant under the group action.

Simple examples are $G = \mathbb{R}$ with μ the Lebesgue measure λ, or a finite group G with μ the counting measure $\#$.

The *dual group* $\hat{G}$, consisting of the continuous group homomorphisms χ from G into the multiplicative group of complex numbers of absolute value 1, is again a locally compact Abelian group. As we shall see soon, for a fixed choice of the normalization factor of the Haar measure μ for G, there is a natural choice for the normalization factor of the Haar measure $\hat{\mu}$ for $\hat{G}$.

Examples of locally compact Abelian groups that play an important role in this work are: the m-dimensional real vector space $\mathbb{R}^m$; the m-fold torus $\mathbb{T}^m := \mathbb{R}^m/\mathbb{Z}^m$ and more generally $\mathbb{R}^m/\Lambda$ for an arbitrary lattice Λ in $\mathbb{R}^m$; and the finite group $\mathbb{D}^m := \frac{1}{q}\mathbb{Z}^m/\mathbb{Z}^m \subset \mathbb{T}^m$ (which is isomorphic to $\mathbb{Z}^m/q\mathbb{Z}^m$) for a positive integer q. Figure 1 below shows the corresponding dual groups as well as the respective (dual) Haar measures as used in this paper.

G	μ	$\hat{G}$	$\hat{\mu}$
$\mathbb{R}^m$	λ	$\hat{\mathbb{R}}^m \simeq \mathbb{R}^m$	λ
$\mathbb{T}^m := \mathbb{R}^m/\mathbb{Z}^m$	λ	$\hat{\mathbb{T}}^m \simeq \mathbb{Z}^m$	$\#$
$\mathbb{D}^m := \frac{1}{q}\mathbb{Z}^m/\mathbb{Z}^m$	$\frac{1}{q^m}\#$	$\hat{\mathbb{D}}^m \simeq \mathbb{Z}^m/q\mathbb{Z}^m$	$\#$
$\mathbb{R}^m/\Lambda$	$\frac{1}{\det(\Lambda)}\lambda$	$\widehat{(\mathbb{R}^m/\Lambda)} \simeq \Lambda^*$	$\#$

Fig. 1. Some groups G and their respective dual groups $\hat{G}$, plus the considered (dual) Haar measures μ and $\hat{\mu}$. Here, λ denotes the Lebesgue and $\#$ the counting measure.

In some cases it will be useful to identify the quotient groups $\mathbb{T}^m = \mathbb{R}^m/\mathbb{Z}^m$ and $\mathbb{D}^m = \frac{1}{q}\mathbb{Z}^m/\mathbb{Z}^m$ with the respective representing sets

$$\mathbb{T}^m_{\mathrm{rep}} := [-\tfrac{1}{2}, \tfrac{1}{2})^m \subset \mathbb{R}^m \qquad \text{and} \qquad \mathbb{D}^m_{\mathrm{rep}} := \tfrac{1}{q}\mathbb{Z}^m \cap \mathbb{T}^m_{\mathrm{rep}}\,,$$

and similarly $\hat{\mathbb{D}}^m \simeq \mathbb{Z}^m/q\mathbb{Z}^m$ with

$$\hat{\mathbb{D}}^m_{\mathrm{rep}} := [q]^m_c := \mathbb{Z}^m \cap [-\tfrac{q}{2}, \tfrac{q}{2})^m\,.$$

It will be useful to understand that if $H \subset G$ is a closed subgroup then G/H and H have dual groups that satisfy the following natural isomorphisms.

$$\widehat{G/H} \simeq H^\perp := \{\chi \in \hat{G} \mid \chi(h) = 1 \,\forall\, h \in H\} \subset \hat{G} \quad \text{and} \quad \hat{H} \simeq \hat{G}/H^\perp.$$

As we shall see soon, for any choice of the Haar measure μ_H for H there is a natural choice for the Haar measure $\mu_{G/H}$ for G/H, and vice versa.

3.2 Norms and Fourier Transforms

Let G be as above with a fixed choice for the Haar measure μ. For any $p \in [1, \infty]$, $L_p(G)$ denotes the vector space of measurable functions $f : G \to \mathbb{C}$ with finite

norm $\|f\|_p$ (modulo the functions with vanishing norm), where

$$\|f\|_p^p := \int_{g\in G} |f(g)|^p d\mu \qquad \text{for } p < \infty,$$

and

$$\|f\|_\infty := \operatorname*{ess\,sup}_{g\in G} |f(g)|,$$

the essential supremum of $|f|$. We write $\|f\|_{p,G}$ if we want to make G explicit. For any function $f \in L^1(G)$, the *Fourier transform* of f is the function

$$\mathcal{F}_G\{f\} : \hat{G} \to \mathbb{C},\ \chi \mapsto \int_{g\in G} f(g)\bar{\chi}(g) d\mu\,,$$

also denoted by $\hat{f}$ when G is clear from the context. The Fourier transform of $f \in L^1(G)$ is continuous, but not necessarily in $L^1(\hat{G})$.

For example, for the group $\mathbb{D}^m := \frac{1}{q}\mathbb{Z}^m/\mathbb{Z}^m$ with the Haar measure as fixed in Figure 1, the L_2-norm and the Fourier transform are respectively given by

$$\|f\|_2^2 = \frac{1}{q^m} \sum_{x\in\mathbb{D}^m} |f(x)|^2 \qquad \text{and} \qquad \mathcal{F}\{f\}(y) = \frac{1}{q^m} \sum_{x\in\mathbb{D}^m} f(x) e^{-2\pi i\langle x,y\rangle}\,.$$

We note that we use a different convention on the scaling than what is common in the context of the quantum Fourier transform.

Given the Haar measure μ for G, there exists a unique *dual* Haar measure $\hat{\mu}$ for $\hat{G}$ with the property that, for any $f \in L^1(G)$, if $\hat{f} = \mathcal{F}_G\{f\} \in L^1(\hat{G})$, then $f = \mathcal{F}_G^{-1}\{\hat{f}\}$, where

$$\mathcal{F}_G^{-1}\{\hat{f}\} : G \to \mathbb{C},\ g \mapsto \int_{\chi\in\hat{G}} \hat{f}(\chi)\chi(g) d\hat{\mu}$$

is the *inverse Fourier transform*. From now on it is always understood that the Haar measure of the dual group is chosen to be the dual of the Haar measure of the primal group. With this choice, we also have the following well known fact [9, Thm. 3.4.8].

Theorem 6 (Plancherel's Identity). *For all $f \in L^1(G) \cap L^2(G)$,*

$$\|f\|_{2,G} = \|\mathcal{F}_G\{f\}\|_{2,\hat{G}}\,.$$

Finally, we recall the *convolution theorem*, which states that $\widehat{fg} = \hat{f} \star \hat{g} = \int_{x\in G} \hat{f}(x)\hat{g}(\cdot\ - x)d\mu(x)$ for all functions $f, g \in L^1(G)$ that have Fourier transforms $\hat{f}, \hat{g} \in L^1(G)$. This extends to functions $f \in L^1(G/H)$ and $g \in L^1(G)$, with f understood as an H-periodic function on G. Tailored to $G = \mathbb{R}^m$ and $H = \Lambda$, where $\mathbb{R}^m/\Lambda$ has dual group Λ^*, it then states that

$$\mathcal{F}_{\mathbb{R}^m}\{fg\}(y) = \mathcal{F}_{\mathbb{R}^m/\Lambda}\{f\} \star \mathcal{F}_{\mathbb{R}^m}\{g\}(y) = \sum_{\ell^*\in\Lambda^*} \mathcal{F}_{\mathbb{R}^m/\Lambda}\{f\}(\ell^*)\, \mathcal{F}_{\mathbb{R}^m}\{g\}(y - \ell^*)$$

for any $y \in \mathbb{R}^m$.

3.3 The Poisson Summation Formula

Poisson summation formula is well-known for the group $G = \mathbb{R}$, where it states that $\sum_{k\in\mathbb{Z}} \hat{f}(k) = \sum_{x\in\mathbb{Z}} f(x)$. In the case $G = \mathbb{Z}/N\mathbb{Z}$, it states that

$$\sum_{i=0}^{N/s} \hat{f}(is) = \sum_{j=1}^{s} f(j\tfrac{N}{s})$$

for any integer s that divides N. In order to formulate the Poisson summation formula for an arbitrary locally compact Abelian group G, we need to introduce the notion of *restriction* and *periodization* of functions.

Definition 2 (Restriction). *Let $H \subseteq G$ be a subset or a subgroup. For any continuous function $f : G \to \mathbb{C}$ we define $f\big|_H : H \to \mathbb{C}, h \mapsto f(h)$.*

Definition 3 (Periodization). *Let H be a closed subgroup of G with Haar measure μ_H. For any function $f \in L^1(G)$, we define*

$$f|^{G/H} : G/H \to \mathbb{C},\ g + H \mapsto \int_{h\in H} f(g+h)d\mu_H\,.$$

For any closed subgroup of G with some fixed Haar measure μ and any choice of the Haar measure μ_H for H, there exists a Haar measure $\mu_{G/H}$ for G/H such that the *quotient integral formula*

$$\int_{G/H} \left(\int_H f(g+h)d\mu_H(h) \right) d\mu_{G/H}(g+H) = \int_G f(g)d\mu(g) \tag{6}$$

holds for any continuous function $f : G \to \mathbb{C}$ with compact support (see [9, Section 1.5]).

With this choice of Haar measure for G/H, and with the dual measures for the respective dual groups, we are ready to state the general form of the Poisson summation formula (obtained from [9, Section 3.6], see also Fig. 2).

Theorem 7 (Poisson Summation Formula). *For continuous $f \in L^1(G)$,*

$$\mathcal{F}_H\{f\big|_H\} = \mathcal{F}_G\{f\}|^{\hat{H}} \quad \textit{and} \quad \mathcal{F}_{G/H}\{f|^{G/H}\} = \mathcal{F}_G\{f\}\big|_{\widehat{G/H}}.$$

$$\begin{array}{ccccc}
L^1(H) & \xleftarrow{\big|_H} & L^1(G) & \xrightarrow{|^{G/H}} & L^1(G/H) \\
\downarrow{\scriptstyle \mathcal{F}_H} & & \downarrow{\scriptstyle \mathcal{F}_G} & & \downarrow{\scriptstyle \mathcal{F}_{G/H}} \\
L^1\big(\hat{G}/\widehat{G/H}\big) & \xleftarrow[|^{\hat{H}}]{} & L^1(\hat{G}) & \xrightarrow[\big|_{\widehat{G/H}}]{} & L^1\big(\widehat{G/H}\big)
\end{array}$$

Fig. 2. Informal illustration of Theorem 7 by means of a diagram that commutes whenever the maps are well defined.

Applied to $G = \mathbb{R}^m$ and $H = \mathbb{Z}^m$, so that $G/H = \mathbb{T}^m$ and $\widehat{G/H} \simeq \mathbb{Z}^m$; and applied to $G = \mathbb{T}^m$ and $H = \mathbb{D}^m$ below, we obtain the following.

Corollary 1. *For continuous $h \in L^1(\mathbb{R}^m)$, we have $\mathcal{F}_{\mathbb{T}^m}\{h|^{\mathbb{T}^m}\} = \mathcal{F}_{\mathbb{R}^m}\{h\}\big|_{\mathbb{Z}^m}$.*

Corollary 2. *For continuous $t \in L^1(\mathbb{T}^m)$, we have $\mathcal{F}_{\mathbb{D}^m}\left\{t\big|_{\mathbb{D}^m}\right\} = \mathcal{F}_{\mathbb{T}^m}\{t\}|^{\hat{\mathbb{D}}^m}$.*

3.4 The Fourier Transform of Vector-Valued Functions

The Fourier transform as discussed above generalizes to vector-valued functions $\boldsymbol{f} : G \to \mathbb{C}^N$ simply by applying $\mathcal{F}$ to the N coordinate functions, resulting in a function $\mathcal{F}\{\boldsymbol{f}\} : \hat{G} \to \mathbb{C}^N$. By fixing an orthonormal basis, this extends to functions $\boldsymbol{f} : G \to \mathcal{H}$ for an arbitrary finite-dimensional complex Hilbert space, where, by linearity of the Fourier transform, $\mathcal{F}\{\boldsymbol{f}\} : \hat{G} \to \mathcal{H}$ is independent of the choice of the basis.

The norm $\|\cdot\|_{2,G}$ on functions $G \to \mathbb{C}$ generalizes to vector-valued functions $\boldsymbol{f} : G \to \mathcal{H}$, as well, by defining $\|\boldsymbol{f}\|_{2,G}$ to be the norm of the scalar function $x \mapsto \|\boldsymbol{f}(x)\|_{\mathcal{H}} = \sqrt{\langle \boldsymbol{f}(x) | \boldsymbol{f}(x) \rangle}$. The vectorial Fourier transforms and norms are compatible with each other, in the sense that Plancherel's identity (see Theorem 6) still holds; that is,

$$\|\boldsymbol{f}\|_{2,G} = \|\mathcal{F}_G\{\boldsymbol{f}\}\|_{2,\hat{G}}\,.$$

Also the Poisson summation formula (see Theorem 7) is still valid, as well as the convolution theorem whenever one of the functions in the product is scalar:

$$\mathcal{F}_G\{\boldsymbol{f}g\} = \mathcal{F}_G\{\boldsymbol{f}\} \star \mathcal{F}_G\{g\}.$$

An important example is the case $\boldsymbol{f} : \mathbb{R}^m/\Lambda \to \mathcal{H}$. Spelling out the above, we get

$$\mathcal{F}_{\mathbb{R}^m/\Lambda}\{\boldsymbol{f}\} : \Lambda^* \to \mathcal{H},\ \ell^* \mapsto |c_{\ell^*}\rangle := \frac{1}{\det \Lambda} \int_{x \in \mathbb{R}^m/\Lambda} |\boldsymbol{f}(x)\rangle e^{-2\pi i \langle x, \ell^* \rangle} dx\,,$$

where the vectors $|c_{\ell^*}\rangle$ are also referred to as the *(vectorial) Fourier coefficients* of $\boldsymbol{f}$. The Parseval-Plancherel identity then becomes

$$\sum_{\ell^* \in \Lambda^*} \langle c_{\ell^*} | c_{\ell^*} \rangle = \|\boldsymbol{f}\|^2_{2,\mathbb{R}^m/\Lambda} := \frac{1}{\det \Lambda} \int_{x \in \mathbb{R}^m/\Lambda} \langle \boldsymbol{f}(x) | \boldsymbol{f}(x) \rangle dx\,.$$

3.5 Trigonometric Approximation

As another application of the Poisson summation formula, we derive a relation between the Lipschitz constant of a function on $\mathbb{T}^m$ and the 'error of discretization' in the Fourier transform when restricting the function to $\mathbb{D}^m$.

Theorem 8. *For any Lipschitz function $\boldsymbol{h} : \mathbb{T}^m \to \mathcal{H}$ with Lipschitz constant* $\mathrm{Lip}(\boldsymbol{h})$*, and any subset $C \subseteq \hat{\mathbb{D}}^m$, we have*

$$\left| \|1_C \cdot \mathcal{F}_{\mathbb{D}^m}\{\boldsymbol{h}\}\|_{2,\hat{\mathbb{D}}^m} - \|1_C \cdot \mathcal{F}_{\mathbb{T}^m}\{\boldsymbol{h}\}\|_{2,\mathbb{Z}^m} \right| \leq \frac{4\pi\sqrt{m}\,\mathrm{Lip}(\boldsymbol{h})}{q}$$

Here and below, we slightly abuse notation and use 1_C as indicator function acting on $\hat{\mathbb{D}}^m$ and on $\mathbb{Z}^m$, justified by identifying $\hat{\mathbb{D}}^m$ with $\hat{\mathbb{D}}^m_{\mathrm{rep}} = [q]^m_c \subset \mathbb{Z}^m$. Also, we write $\mathcal{F}_{\mathbb{D}^m}\{\boldsymbol{h}\}$ instead of $\mathcal{F}_{\mathbb{D}^m}\{\boldsymbol{h}|_{\mathbb{D}^m}\}$, taking it as understood that $\boldsymbol{h}$ is restricted to $\mathbb{D}^m$ when applying $\mathcal{F}_{\mathbb{D}^m}$.

Proof. Using a result of Yudin ([37, Example I after Theorem 2], see also[5] Appendix D of the full version [3]), there exists a trigonometric approximation $\boldsymbol{t}$ of $\boldsymbol{h}$, i.e. a function $\boldsymbol{t} : \mathbb{T}^m \to \mathbb{C}$ with $\hat{\boldsymbol{t}}(x) := \mathcal{F}_{\mathbb{T}^m}\{\boldsymbol{t}\}(x) = 0$ for all $x \notin [q]^m_c$ so that $\|\boldsymbol{h} - \boldsymbol{t}\|_\infty \leq \pi\sqrt{m}\,\mathrm{Lip}(\boldsymbol{h})/q$. Recalling that $\hat{\mathbb{D}}^m \simeq \mathbb{Z}^m/q\mathbb{Z}^m$, the fact that $\hat{\boldsymbol{t}} : \mathbb{Z}^m \to \mathbb{C}$ vanishes outside of $[q]^m_c$ implies for all $x \in [q]^m_c$ that

$$\hat{\boldsymbol{t}}(x) = \sum_{d \in q\mathbb{Z}^m} \hat{\boldsymbol{t}}(x+d) = \hat{\boldsymbol{t}}\big|^{\hat{\mathbb{D}}^m}(x + q\mathbb{Z}^m) = \mathcal{F}_{\mathbb{D}^m}\{\boldsymbol{t}\}(x + q\mathbb{Z}^m)\,,$$

where the last equality holds by Corollary 2 (and our convention of omitting the restriction to $\mathbb{D}^m$). In particular, we have $\|1_C \cdot \mathcal{F}_{\mathbb{D}^m}\{\boldsymbol{t}\}\|_{2,\hat{\mathbb{D}}^m} = \|1_C \cdot \mathcal{F}_{\mathbb{T}^m}\{\boldsymbol{t}\}\|_{2,\mathbb{Z}^m}$. Therefore, by the (reverse) triangle inequality and the linearity of the Fourier transform, one obtains

$$\begin{aligned}&\left| \|1_C \cdot \mathcal{F}_{\mathbb{D}^m}\{\boldsymbol{h}\}\|_{2,\hat{\mathbb{D}}^m} - \|1_C \cdot \mathcal{F}_{\mathbb{T}^m}\{\boldsymbol{h}\}\|_{2,\mathbb{Z}^m} \right| \\ &\qquad \leq \|1_C \cdot \mathcal{F}_{\mathbb{D}^m}\{\boldsymbol{h}-\boldsymbol{t}\}\|_{2,\hat{\mathbb{D}}^m} + \|1_C \cdot \mathcal{F}_{\mathbb{T}^m}\{\boldsymbol{h}-\boldsymbol{t}\}\|_{2,\mathbb{Z}^m}\,.\end{aligned}$$

We now observe that

$$\|1_C \cdot \mathcal{F}_G\{\boldsymbol{h}-\boldsymbol{t}\}\|_{2,\hat{G}} \leq \|\mathcal{F}_G\{\boldsymbol{h}-\boldsymbol{t}\}\|_{2,\hat{G}} = \|\boldsymbol{h}-\boldsymbol{t}\|_{2,G} \leq \sqrt{\mu(G)}\,\|\boldsymbol{h}-\boldsymbol{t}\|_\infty$$

where $\mu(G) = \int_G d\mu$ denotes the total measure of G. We conclude by noting that $\mu(G) = 1$ for both groups at hand $G = \mathbb{D}^m$ and $G = \mathbb{T}^m$. □

3.6 The Gaussian Function and Smoothing Errors

Let m be a fixed positive integer. For any parameter $\sigma > 0$, we consider the m-dimensional *Gaussian function*

$$\rho_\sigma : \mathbb{R}^m \to \mathbb{C}\,,\ x \mapsto e^{-\frac{\pi\|x\|^2}{\sigma^2}}\,,$$

which is well known to satisfy the following basic properties.

[5] In Appendix D of the full version [3], we provide a slight generalization of Yudin's paper [37] to functions with vectorial output. In principle the bound of Theorem 8 can also derived without this generalization, but at the cost of an undesirable extra factor $\dim \mathcal{H} = 2^n$.

Lemma 1. *For all $\sigma > 0$, $m \in \mathbb{N}$ and $x, y \in \mathbb{R}^m$, we have $\int_{z\in\mathbb{R}^m} \rho_\sigma(z)dz = \sigma^m$, $\mathcal{F}_{\mathbb{R}^m}\{\rho_\sigma\} = \sigma^m \rho_{1/\sigma}$, $\sqrt{\rho_\sigma(x)} = \rho_{\sqrt{2}\sigma}(x)$ and $\rho_\sigma(x)\rho_\sigma(y) = \rho_{\frac{\sigma}{\sqrt{2}}}(\frac{x+y}{2})\rho_{\frac{\sigma}{\sqrt{2}}}(\frac{x-y}{2})$.*

Remark 7. From these properties it follows that the integral of the L_2-norm of $x \mapsto \sigma^{m/2} \cdot \sqrt{\rho_{1/\sigma}(x)}$ equals 1, i.e., $\left\|\sigma^{m/2} \cdot \sqrt{\rho_{1/\sigma}(x)}\right\|^2_{2,\mathbb{R}^m} = 1$.

The following two results (and the variations we discuss below) will play an important role and will be used several times in this paper: *Banaszczyk's bound*, originating from [1], and the *smoothing error* [6] , as introduced by Micciancio and Regev [20]. They allow us to control

$$\rho_\sigma(X) := \sum_{x\in X} \rho_\sigma(x)\,,$$

for certain discrete subsets $X \subseteq \mathbb{R}^m$. For ease of notation, we let

$$\beta_z^{(m)} := \left(\frac{2\pi e z^2}{m}\right)^{m/2} e^{-\pi z^2}\,,$$

which decays super-exponentially in z (for fixed m). The following formulation of Banaszczyk's lemma is obtained from [21, Equation (1.1)].

Lemma 2 (Banaszczyk's Bound). *Whenever $r/\sigma \geq \sqrt{\frac{m}{2\pi}}$,*

$$\rho_\sigma\big((\Lambda + t) \setminus \mathcal{B}_r\big) \leq \beta_{r/\sigma}^{(m)} \cdot \rho_\sigma(\Lambda)\,,$$

where $\mathcal{B}_r = \mathcal{B}_r(0) = \{x \in \mathbb{R}^m \;\big|\; |x| < r\}$.

Imitating techniques from [20, Lemma 3.2], we have:

Lemma 3. *Let $\sigma \geq \frac{\sqrt{m}}{\lambda_1(\Lambda^*)}$. Then $\rho_{1/\sigma}(\Lambda^*\backslash 0) \leq 2 \cdot \beta^{(m)}_{\sigma\lambda_1(\Lambda^*)}$.*

As a direct corollary, we have the following result.

Corollary 3. *Let $\sigma \geq 2\sqrt{m}$, and let $x \in \mathbb{R}^m$ with $\|x\|_\infty \leq 1/2$. Then*

$$\rho_{1/\sigma}\big(\mathbb{Z}^m\backslash\{0\} + x\big) \leq 2\beta^{(m)}_{\sigma/2}.$$

Proof. We have $\rho_{1/\sigma}\big(\mathbb{Z}^m\backslash\{0\}+x\big) \leq \rho_{1/\sigma}\big((\mathbb{Z}^m+x)\backslash\mathcal{B}_{\frac{1}{2}}\big) \leq \beta^{(m)}_{\sigma/2}\rho_{1/\sigma}(\mathbb{Z}^m)$, where the second inequality follows from Lemma 2. Using Lemma 3 to argue that $\rho_{1/\sigma}(\mathbb{Z}^m) = 1 + \rho_{1/\sigma}(\mathbb{Z}^m\backslash 0) \leq 1 + 2\beta^{(m)}_\sigma \leq 2$ then proves the claim. □

The following lemma, which combines [20, Lemma 4.1] and [20, Lemma 3.2], controls the fluctuation of the sum $\rho_s(\Lambda + t)$ for varying $t \in \mathbb{R}^m$.

[6] Although most literature on lattices analyze smoothing errors in terms of the *smoothing parameter* η_ϵ, we chose not to do so. Instead, this paper addresses smoothing errors in a reversed and more direct way, making the errors occurring in the later analysis more easy to describe.

Lemma 4 (Smoothing Error). *Let $\Lambda \in \mathbb{R}^m$ be a full rank lattice, and let $\sigma \geq \sqrt{m}/\lambda_1(\Lambda^*)$. Then, for any $t \in \mathbb{R}^m$,*

$$(1 - 2\beta^{(m)}_{\sigma\lambda_1(\Lambda^*)})\frac{\sigma^m}{\det \Lambda} \leq \rho_\sigma(\Lambda + t) \leq (1 + 2\beta^{(m)}_{\sigma\lambda_1(\Lambda^*)})\frac{\sigma^m}{\det \Lambda}\,. \tag{7}$$

Corollary 4. *For $\sigma \geq \frac{\sqrt{m}}{\lambda_1(\Lambda^*)}$ and for any $t \in \mathbb{R}^m$, we have $\rho_\sigma(\Lambda + t) \leq 2\frac{\sigma^m}{\det \Lambda}$.*

Proof. Using Lemma 4 and noticing $2\beta^{(m)}_{\sigma\lambda_1(\Lambda^*)} \leq 2\beta^{(m)}_{\sqrt{m}} \leq 1$ yields the result. □

3.7 Lipschitz Condition

Theorem 9 (Rademacher's theorem). *A Lipschitz function $\boldsymbol{f} : \mathbb{R}^m/\Lambda \to \mathcal{H}$ has weak partial derivatives $\partial_{x_j}\boldsymbol{f} : \mathbb{R}^m/\Lambda \to \mathcal{H}$ lying in $L_2(\mathbb{R}^m/\Lambda)$. In particular, $\sum_{j=1}^m \left\|\partial_{x_j}\boldsymbol{f}\right\|^2_{2,\mathbb{R}^m/\Lambda} \leq \mathrm{Lip}(\boldsymbol{f})^2$.*

Proof. Combining the proof of [17, Theorem 4.1 and 4.9] and [35, Theorem 2] on measures of compact sets, we obtain this result. □

Corollary 5. *Let $\boldsymbol{f} : \mathbb{R}^m/\Lambda \to \mathcal{H}$ be a Lipschitz-continuous function, and denote by $|c_{\ell^*}\rangle$ the vectorial Fourier coefficients of $\boldsymbol{f}$. Then,*

$$\sum_{\substack{\ell^* \in \Lambda^* \\ \|\ell^*\| \geq B}} \langle c_{\ell^*}|c_{\ell^*}\rangle \leq \frac{\mathrm{Lip}(\boldsymbol{f})^2}{4\pi^2 B^2}\,.$$

Proof. Since $\boldsymbol{f}$ is Lipschitz, we can apply Theorem 9. Furthermore, the identity $|\boldsymbol{f}(x)\rangle = \sum_{\ell^* \in \Lambda^*} |c_{\ell^*}\rangle e^{2\pi i\langle x,\ell^*\rangle}$ implies $|\partial_{x_j}\boldsymbol{f}(x)\rangle = 2\pi i \sum_{\ell^* \in \Lambda^*} \ell^*_j |c_{\ell^*}\rangle e^{2\pi i\langle x,\ell^*\rangle}$ almost everywhere ([36, Lemma V.2.11] or [30, Lemma 2.16]). Finally, given that $\sum_{j=1}^m \left\|\partial_{x_j}\boldsymbol{f}\right\|^2_{2,\mathbb{R}^m/\Lambda} \leq \mathrm{Lip}(\boldsymbol{f})^2$, Plancherel's identity implies that

$$\begin{aligned}\mathrm{Lip}(\boldsymbol{f})^2 &\geq \sum_{j=1}^m \left\|\partial_{x_j}\boldsymbol{f}\right\|^2_{2,\mathbb{R}^m/\Lambda} = 4\pi^2 \sum_{\ell^* \in \Lambda^*} \|\ell^*\|_2^2 \langle c_{\ell^*}|c_{\ell^*}\rangle \\ &\geq 4\pi^2 \sum_{\substack{\ell^* \in \Lambda^* \\ \|\ell^*\|_2 \geq B}} \|\ell^*\|_2^2 \langle c_{\ell^*}|c_{\ell^*}\rangle \geq 4B^2\pi^2 \sum_{\substack{\ell^* \in \Lambda^* \\ \|\ell^*\|_2 \geq B}} \langle c_{\ell^*}|c_{\ell^*}\rangle\,,\end{aligned}$$

from which the claim follows. □

4 Algorithm

4.1 The Algorithm

Given a Λ-periodic function $\boldsymbol{f} : \mathbb{R}^m \to \mathcal{S}$ as discussed in Section 2, which maps a classical input x to a quantum state $|\boldsymbol{f}(x)\rangle$, we consider the following quantum algorithm (see Figure 3). The algorithm has oracle access to $\boldsymbol{f}$, meaning that

it has access to a unitary that maps $|x\rangle|0\rangle$ to $|x\rangle|\boldsymbol{f}(x)\rangle$. As a matter of fact, we may obviously assume the algorithm to have oracle access to a unitary that maps $|x\rangle|0\rangle$ to $|x\rangle|\boldsymbol{f}(Vx)\rangle$ for a parameter $V \in \mathbb{R}$ chosen by the algorithm. Per se, x may be arbitrary in $\mathbb{R}^m$; for any concrete algorithm it is of course necessary to restrict x to some finite subset of $\mathbb{R}^m$.

The algorithm we consider follows the blueprint of the standard hidden-subgroup algorithm. Notable differences are that we need to discretize (and finitize) the continuous domain $\mathbb{R}^m$ of the function, and the algorithm starts off with a superposition that is not uniform but follows a (discretized and finitized) Gaussian distribution. The reason for the latter choice is that Gaussian distributions decay very fast and behave nicely under the Fourier transform (as they are eigenfunctions of the Fourier transform).

The algorithm is given in Figure 3 below. It uses two quantum registers, each one consisting of a certain number of qubits. Associated to the first register are orthonormal bases $\{|x\rangle_{\mathbb{D}^m}\}_{x\in\mathbb{D}^m}$ and $\{|y\rangle_{\hat{\mathbb{D}}^m}\}_{y\in\hat{\mathbb{D}}^m}$ where the basis vectors are labeled by $x \in \mathbb{D}^m$ and $y \in \hat{\mathbb{D}}^m$, respectively, which we identify with elements $x \in \mathbb{D}^m_{\mathrm{rep}}$ and $y \in \hat{\mathbb{D}}^m_{\mathrm{rep}}$ (see Section 3.1). The second register has state space $\mathcal{H}$. The algorithm is parameterized by $q \in \mathbb{N}$ (which determines $\mathbb{D}^m$), $s > 0$ and $V > 0$. Intuitively, the fraction $\frac{s}{V}$ is tightly related to the absolute precision of the output, whereas q is connected with the number of qubits needed.

Algorithm 1: Quantum algorithm for the dual lattice sampling problem

1 **Prepare the Gaussian state** $|\psi_\circ\rangle := \sum_{x\in\mathbb{D}^m} \sqrt{\rho_{1/s}(x)} \cdot |x\rangle_{\mathbb{D}^m}|0\rangle$;

2 **Apply the $\boldsymbol{f}$-oracle**, yielding $\sum_{x\in\mathbb{D}^m} \sqrt{\rho_{1/s}(x)} \cdot |x\rangle_{\mathbb{D}^m}|\boldsymbol{f}(Vx)\rangle$;

3 **Apply the quantum Fourier transform on the first register**, yielding the unnormalized state $\sum_{x\in\mathbb{D}^m}\sum_{y\in\hat{\mathbb{D}}^m} \sqrt{\rho_{1/s}(x)} \cdot e^{-2\pi i\langle x,y\rangle} \cdot |y\rangle_{\hat{\mathbb{D}}^m}|\boldsymbol{f}(Vx)\rangle$;

4 **Measure the first register in the $\hat{\mathbb{D}}^m_{\mathrm{rep}}$-basis** yielding some $y \in \hat{\mathbb{D}}^m_{\mathrm{rep}}$, and output $\frac{y}{V}$;

Fig. 3. The continuous-hidden-subgroup quantum algorithm.

The description and Analysis of Step 1 is deferred to Appendix C of the full version [3]. It will be shown (as summarized in Theorem 5) that its cost is negligible compared to the main cost of Algorithm 1, while contributing an error of at most $o(\eta)$ in the trace distance.

4.2 The Figure of Merit

Recall that $N = \dim \mathcal{H} = 2^n$. Then the state after step (2) of Algorithm 1 equals, up to normalization,

$$|\psi\rangle := s^{m/2} \sum_{x \in \mathbb{D}^m} \sqrt{\rho_{1/s}(x)}\, |x\rangle_{\mathbb{D}^m} |\boldsymbol{f}(Vx)\rangle$$

which we can rewrite as

$$|\psi\rangle = \sum_{x \in \mathbb{D}^m} |x\rangle_{\mathbb{D}^m} |\boldsymbol{h}(x)\rangle$$

where

$$\boldsymbol{h}(x) := s^{m/2} \sqrt{\rho_{1/s}(x)} \cdot |\boldsymbol{f}(Vx)\rangle\,.$$

Applying the quantum Fourier transform in step (3) maps this to

$$|\hat{\psi}\rangle = q^{-m/2} \sum_{x \in \mathbb{D}^m} \sum_{y \in \hat{\mathbb{D}}^m} e^{-2\pi i \langle x, y\rangle} |y\rangle_{\hat{\mathbb{D}}^m} |\boldsymbol{h}(x)\rangle = q^{m/2} \sum_{y \in \hat{\mathbb{D}}^m} |y\rangle_{\hat{\mathbb{D}}^m} |\mathcal{F}_{\mathbb{D}^m}\{\boldsymbol{h}\}(y)\rangle\,,$$

where the factor $q^{m/2}$ comes from the fact that, by our convention, the Fourier transform $\mathcal{F}_{\mathbb{D}^m}$ is scaled with the factor q^{-m}, while the quantum Fourier transform comes with a scaling factor $q^{-m/2}$.

Up to normalization, the probability to observe outcome y in step (4) thus is

$$\langle \hat{\psi} | (|y\rangle\langle y| \otimes \mathbb{I}) | \hat{\psi} \rangle = q^m \left\| \mathcal{F}_{\mathbb{D}^m}\{\boldsymbol{h}\}(y) \right\|_{\mathcal{H}}^2\,,$$

and so, for any "target" subset $C \subset \hat{\mathbb{D}}^m$, the probability for the algorithm to produce an outcome $y \in C$ equals

$$\mathcal{D}(C) = \sum_{y \in C} \frac{\langle \hat{\psi} | (|y\rangle\langle y| \otimes \mathbb{I}) | \hat{\psi} \rangle}{\langle \psi_\circ | \psi_\circ \rangle} = \frac{\| 1_C \cdot \mathcal{F}_{\mathbb{D}^m}\{\boldsymbol{h}\} \|_{2,\hat{\mathbb{D}}^m}^2}{\frac{s^m}{q^m} \sum_{x \in \mathbb{D}^m} \rho_{1/s}(x)}\,. \tag{8}$$

Intuitively, in the limit $q \to \infty$, the grid $\frac{1}{q}\mathbb{Z}^m$ becomes $\mathbb{R}^m$; thus, neglecting constant factors, the function $\mathcal{F}_{\mathbb{D}^m}\{\boldsymbol{h}\}$ is expected to converge to

$$\mathcal{F}_{\mathbb{R}^m}\{\rho_{\sqrt{2}/s} \boldsymbol{f}(V\cdot)\} = \rho_{s/\sqrt{2}} \star \mathcal{F}_{\mathbb{R}^m}\{\boldsymbol{f}(V\cdot)\}\,.$$

Furthermore, when V is large enough compared to s then, relative to the dual lattice $V\Lambda^*$, the Gaussian function behaves as a Dirac delta function. Thus, the above function is then supported by $V\Lambda^*$ and takes on the values $|c_{\ell^*}\rangle$. Hence, by taking square norms, we get the claimed $\langle c_{\ell^*} | c_{\ell^*} \rangle$.

Below, we prove that this intuition is indeed correct, and we work out the actual "rate of convergence".

5 Analysis

5.1 Proof Overview

In the overview here and in the formal analysis in the next section, we consider the case $V = 1$. This is without loss of generality; in order to deal with an arbitrary V we simply apply our analysis to the function $\boldsymbol{f}_V := \boldsymbol{f}(V\cdot)$, with the effect that in the error term, Λ^* becomes $V\Lambda^*$ and $\mathrm{Lip}(\boldsymbol{f}_V)$ becomes $V\,\mathrm{Lip}(\boldsymbol{f})$.

The error analysis (for $V = 1$) is divided into three parts. The first part consists of showing that the denominator from Equation (8) satisfies

$$\frac{s^m}{q^m} \sum_{x \in \mathbb{D}^m} \rho_{1/s}(x) \approx 1\,.$$

In the second part, which is the most technical one, we show that for any $C \subset \hat{\mathbb{D}}^m$, also understood as a subset of $\hat{\mathbb{D}}^m_{\mathrm{rep}} = [q]^m_c \subset \mathbb{Z}^m$,

$$\|1_C \cdot \mathcal{F}_{\mathbb{D}^m}\{\boldsymbol{h}\}\|^2_{2,\hat{\mathbb{D}}^m} \gtrsim \sum_{\substack{\ell^* \in \Lambda^* \\ B_{\delta\lambda_1^*}(\ell^*) \subseteq C}} \langle c_{\ell^*} | c_{\ell^*} \rangle\,. \tag{9}$$

We recall that $|c_{\ell^*}\rangle$ are the vectorial Fourier coefficients of $\boldsymbol{f}$ and $B_{\delta\lambda_1^*}(\ell^*) = \mathcal{B}_{\delta\lambda_1^*}(\ell^*) \cap \mathbb{Z}^m$. This approximation (9) is divided into the following five steps:

$$\|1_C \mathcal{F}_{\mathbb{D}^m}\{\boldsymbol{h}\}\|^2_{2,\hat{\mathbb{D}}^m} \overset{(1)}{\approx} \left\|1_C \mathcal{F}_{\mathbb{D}^m}\left\{\boldsymbol{h}|^{\mathbb{T}^m}\right\}\right\|^2_{2,\hat{\mathbb{D}}^m} \overset{(2)}{\approx} \left\|1_C \mathcal{F}_{\mathbb{T}^m}\{\boldsymbol{h}|^{\mathbb{T}^m}\}\right\|^2_{2,\mathbb{Z}^m}$$
$$\overset{(3)}{=} \|1_C \mathcal{F}_{\mathbb{R}^m}\{\boldsymbol{h}\}\|^2_{2,\mathbb{Z}^m} \overset{(4)}{\approx} \sum_{\ell^* \in \Lambda^*} \langle c_{\ell^*} | c_{\ell^*} \rangle \cdot \iota_C(\ell^*) \overset{(5)}{\geq} \sum_{\substack{\ell^* \in \Lambda^* \\ B_{\delta\lambda_1^*}(\ell^*) \subseteq C}} \langle c_{\ell^*} | c_{\ell^*} \rangle\,.$$

It thus follows that

$$\mathcal{D}(C) \gtrsim \sum_{\substack{\ell^* \in \Lambda^* \\ B_{\delta\lambda_1^*}(\ell^*) \subseteq C}} \langle c_{\ell^*} | c_{\ell^*} \rangle\,,$$

and therefore, applied to $C := B_{\delta\lambda_1^*}(S)$, that for any $S \subset \Lambda^*$ for which $B_{\delta\lambda_1^*}(S) \subset [q]^m_c$, requirement (3) is satisfied.

The third part of the analysis is to show that (3) is satisfied also for $S \subset \Lambda^*$ for which $B_{\delta\lambda_1^*}(S)$ is not fully contained in $[q]^m_c$. For such S, it is then sufficient to show that $\sum_{\ell^* \in S \setminus S_0} \langle c_{\ell^*} | c_{\ell^*} \rangle \approx 0$ then, where $S_0 = \{\ell^* \in S \mid B_{\delta\lambda_1^*}(\ell^*) \subseteq [q]^m_c\}$. We prove this by means of Corollary 5.

We emphasize that in the formal proof below, we explicitly follow this 3-part structure of the proof, with part 2 being divided into 5 steps as indicated above.

5.2 Formal Analysis

Part 1 By Lemma 4, we have (whenever $q/s \geq \sqrt{m}$),

$$\frac{s^m}{q^m} \sum_{x \in \mathbb{D}^m} \rho_{1/s}(x) \leq \frac{s^m}{q^m} \cdot \rho_{1/s}\left(\frac{1}{q}\mathbb{Z}^m\right) \leq 1 + 2\beta^{(m)}_{q/s}. \tag{10}$$

Therefore,

$$\frac{\|1_C \cdot \mathcal{F}_{\mathbb{D}^m}\{\boldsymbol{h}\}\|^2_{2,\hat{\mathbb{D}}^m}}{\frac{s^m}{q^m}\sum_{x\in\mathbb{D}^m}\rho_{1/s}(x)} \geq \|1_C \cdot \mathcal{F}_{\mathbb{D}^m}\{\boldsymbol{h}\}\|^2_{2,\hat{\mathbb{D}}^m} - \varepsilon_{\text{denom}} \tag{11}$$

with $\varepsilon_{\text{denom}} = 2\beta^{(m)}_{q/s}$.

Part 2 Recall that $\boldsymbol{h} = s^{m/2}\cdot\boldsymbol{f}\cdot\rho_{\sqrt{2}/s}$ is a function $\boldsymbol{h}:\mathbb{R}^m\to\mathcal{H}$. In the following, by slightly abusing notation, we also understand $\boldsymbol{h}$ as a function $\boldsymbol{h}:\mathbb{T}^m\to\mathcal{H}$ by considering the restriction of $\boldsymbol{h}$ to $\mathbb{T}^m_{\text{rep}} = [-\frac{1}{2},\frac{1}{2})^m$. Similarly, we understand $\boldsymbol{h}$ as a function $\boldsymbol{h}:\mathbb{D}^m\to\mathcal{H}$ by considering its restriction to $\mathbb{D}^m_{\text{rep}} = \mathbb{T}^m_{\text{rep}}\cap\frac{1}{q}\mathbb{Z}^m$.

Step 1. Observe that

$$\left\|1_C \cdot \mathcal{F}_{\mathbb{D}^m}\{\boldsymbol{h}\} - 1_C\cdot\mathcal{F}_{\mathbb{D}^m}\left\{\boldsymbol{h}|^{\mathbb{T}^m}\right\}\right\|_{2,\hat{\mathbb{D}}^m} \leq \left\|\mathcal{F}_{\mathbb{D}^m}\left\{\boldsymbol{h}-\boldsymbol{h}|^{\mathbb{T}^m}\right\}\right\|_{2,\hat{\mathbb{D}}^m} = \left\|\boldsymbol{h}|^{\mathbb{T}^m}-\boldsymbol{h}\right\|_{2,\mathbb{D}^m}.$$

Writing out the definition of $\boldsymbol{h}|^{\mathbb{T}^m}$ and $\boldsymbol{h}$, we obtain (provided that $\frac{s}{2\sqrt{2}}\geq\sqrt{m}$)

$$\left\|\boldsymbol{h}|^{\mathbb{T}^m}-\boldsymbol{h}\right\|^2_{2,\mathbb{D}^m} = \frac{1}{q^m}\sum_{x\in\mathbb{D}^m}\left\|\sum_{z\in\mathbb{Z}^m\setminus 0}\boldsymbol{h}(x+z)\right\|^2_{\mathcal{H}}$$

$$\leq \frac{\|\boldsymbol{f}\|^2_\infty s^m}{q^m}\sum_{x\in\mathbb{D}^m}\left(\sum_{z\in\mathbb{Z}^m\setminus 0}\rho_{\sqrt{2}/s}(x+z)\right)^2 \leq 4s^m(\beta^{(m)}_{\frac{s}{2\sqrt{2}}})^2,$$

as $\rho_{\sqrt{2}/s}\big(\mathbb{Z}^m\setminus\{0\}+x\big)\leq 2\beta^{(m)}_{\frac{s}{2\sqrt{2}}}$, from Corollary 3, combining with the fact that $\|f\|_\infty = 1$. Taking square roots and using the reverse triangle inequality yields

$$\left|\|1_C\cdot\mathcal{F}_{\mathbb{D}^m}\{\boldsymbol{h}\}\|_{2,\hat{\mathbb{D}}^m} - \left\|1_C\cdot\mathcal{F}_{\mathbb{D}^m}\left\{\boldsymbol{h}|^{\mathbb{T}^m}\right\}\right\|_{2,\hat{\mathbb{D}}^m}\right| \leq 2s^{m/2}\beta^{(m)}_{\frac{s}{2\sqrt{2}}} =: \varepsilon_{\text{per}}$$

Step 2. Using Theorem 8 with $\boldsymbol{h}|^{\mathbb{T}^m}$, one obtains

$$\left|\left\|1_C\cdot\mathcal{F}_{\mathbb{D}^m}\left\{\boldsymbol{h}|^{\mathbb{T}^m}\right\}\right\|_{2,\hat{\mathbb{D}}^m} - \left\|1_C\cdot\mathcal{F}_{\mathbb{T}^m}\{\boldsymbol{h}|^{\mathbb{T}^m}\}\right\|_{2,\mathbb{Z}^m}\right| \leq \varepsilon_{\text{lip}},$$

where $\varepsilon_{\text{lip}} = \frac{4\pi\sqrt{m}\,\text{Lip}(\boldsymbol{h}|^{\mathbb{T}^m})}{q}$. Recall that we use 1_C as indicator function acting on $\mathbb{Z}^m$ and on $\hat{\mathbb{D}}^m\simeq\mathbb{Z}^m/q\mathbb{Z}^m$ in the obvious way.

The Lipschitz constant of $\boldsymbol{h}|^{\mathbb{T}^m}$ can be obtained by taking the maximum value of the absolute value of the derivative.

$$\frac{\partial}{\partial x_j}\left(\boldsymbol{h}|^{\mathbb{T}^m}\right) = s^{m/2}\sum_{z\in\mathbb{Z}^m}\left(\frac{\partial}{\partial x_j}\boldsymbol{f}(x+z)\cdot\rho_{\sqrt{2}/s}(x+z) + \boldsymbol{f}(x+z)\frac{\partial}{\partial x_j}\rho_{\sqrt{2}/s}(x+z)\right)$$

The norm of $\nabla\left(\boldsymbol{h}|^{\mathbb{T}^m}\right)$ is therefore bounded by

$$s^{m/2}\left(\mathrm{Lip}(\boldsymbol{f})\rho_{\sqrt{2}/s}(x+\mathbb{Z}^m)+\pi s^2\left\|\boldsymbol{f}\right\|_\infty\sum_{z\in\mathbb{Z}^m}\|x+z\|\,\rho_{\sqrt{2}/s}(x+z)\right)$$
$$\leq s^{m/2}\left(2\,\mathrm{Lip}(\boldsymbol{f})+2\pi s^2\right)$$

where we used $\|\nabla\boldsymbol{f}\|=\sqrt{\sum_{j=1}^m\left\|\frac{\partial}{\partial x_j}\boldsymbol{f}\right\|_{\mathcal{H}}^2}\leq\mathrm{Lip}(\boldsymbol{f})$, $\|\boldsymbol{f}\|_\infty\leq 1$, $\nabla\rho_{\sqrt{2}/s}(x)=\pi s^2 x\cdot\rho_{\sqrt{2}/s}(x)$, $\rho_{\sqrt{2}/s}(x+\mathbb{Z}^m)\leq 2$ and $\sum_{z\in\mathbb{Z}^m}\|x+z\|\,\rho_{\sqrt{2}/s}(x+z)\leq 2$. The second last inequality follows from $\rho_{\sqrt{2}/s}(x+\mathbb{Z}^m)\leq 1+\rho_{\sqrt{2}/s}(\mathbb{Z}^m\backslash\{0\}+x)\leq 1+2\beta^{(m)}_{\frac{s}{2\sqrt{2}}}\leq 2$, see Corollary 3. The last inequality can be obtained by the fact that $\|x+z\|\,\rho_{\sqrt{2}/s}(x+z)\leq\rho_{\sqrt{2}/(s-1)}(x+z)$, and repeating the former argument.

Step 3. Apply Corollary 1 to conclude that

$$\left\|1_C\cdot\mathcal{F}_{\mathbb{T}^m}\{\boldsymbol{h}|^{\mathbb{T}^m}\}\right\|_{2,\mathbb{Z}^m}=\|1_C\cdot\mathcal{F}_{\mathbb{R}^m}\{\boldsymbol{h}\}\|_{2,\mathbb{Z}^m}\ ,$$

where we continue to abuse notation here by identifying $\mathcal{F}_{\mathbb{R}^m}\{\boldsymbol{h}\}$ with its restriction to $\mathbb{Z}$.

Using $|a^2-b^2|=|a+b||a-b|\leq(|a-b|+2|a|)|a-b|$ and the fact that $\|1_C\cdot\mathcal{F}_{\mathbb{D}^m}\{\boldsymbol{h}\}\|_{2,\hat{\mathbb{D}}^m}\leq 2$ (which follows from Equation (8) and Equation (10)), we conclude that

$$\left|\|1_C\cdot\mathcal{F}_{\mathbb{D}^m}\{\boldsymbol{h}\}\|^2_{2,\hat{\mathbb{D}}^m}-\|1_C\cdot\mathcal{F}_{\mathbb{R}^m}\{\boldsymbol{h}\}\|^2_{2,\mathbb{Z}^m}\right|\leq 5(\varepsilon_{\mathrm{per}}+\varepsilon_{\mathrm{lip}}),$$

where we assume that $\varepsilon_{\mathrm{per}}+\varepsilon_{\mathrm{lip}}<1$.

Step 4. By applying the convolution theorem as outlined in Section 3.2, we see that

$$\mathcal{F}_{\mathbb{R}^m}\{\boldsymbol{h}\}[y]=\mathcal{F}_{\mathbb{R}^m/\Lambda}\{\boldsymbol{f}\}\star\mathcal{F}_{\mathbb{R}^m}\{s^{m/2}\rho_{s/\sqrt{2}}\}(y)=\left(\frac{2}{s}\right)^{m/2}\sum_{\ell^*\in\Lambda^*}|c_{\ell^*}\rangle\rho_{s/\sqrt{2}}(y-\ell^*)$$

where $|c_{\ell^*}\rangle$ are the vectorial Fourier coefficients of $\boldsymbol{f}$. Therefore,

$$\begin{aligned}\|\mathcal{F}_{\mathbb{R}^m}\{\boldsymbol{h}\}[y]\|^2_{\mathcal{H}}&=\left(\frac{2}{s}\right)^m\sum_{k^*\in\Lambda^*}\sum_{\ell^*\in\Lambda^*}\langle c_{\ell^*}|c_{k^*}\rangle\rho_{s/\sqrt{2}}(y-\ell^*)\rho_{s/\sqrt{2}}(y-k^*)\\&=\left(\frac{2}{s}\right)^m\sum_{u^*\in\frac{1}{2}\Lambda^*}\sum_{v^*\in u^*+\Lambda^*}\langle c_{v^*+u^*}|c_{v^*-u^*}\rangle\rho_{s/2}(u^*)\rho_{s/2}(y-v^*)\,,\end{aligned}$$

where the latter is obtained by the variable substitution $u^*=\frac{\ell^*-k^*}{2}$, $v^*=\frac{\ell^*+k^*}{2}$, and using Lemma 1. Summing over $y\in C$, setting

$$\iota_C(\ell^*):=\left(\frac{2}{s}\right)^m\sum_{y\in C}\rho_{s/2}(y-\ell^*)\,,$$

and splitting into $u^* = 0$ and $u^* \neq 0$, we obtain

$$\|1_C \mathcal{F}_{\mathbb{R}^m}\{\boldsymbol{h}\}\|^2_{2,\mathbb{Z}^m} = \sum_{v^* \in \Lambda^*} \langle c_{v^*} | c_{v^*} \rangle \cdot \iota_C(v^*) + \sum_{u^* \in \frac{1}{2}\Lambda^* \setminus 0} \rho_{s/2}(u^*) \sum_{v^* \in u^* + \Lambda^*} \langle c_{v^*+u^*} | c_{v^*-u^*} \rangle \cdot \iota_C(v^*)$$

We now bound the second term. Assuming $s \geq \sqrt{m}$, we have that $\iota_C(v^*) \leq \left(\frac{2}{s}\right)^m \rho_{s/2}(\mathbb{Z}^m + t) \leq 2$ (see Corollary 4). Furthermore, by the Cauchy-Schwartz inequality,

$$\left| \sum_{v^* \in u^* + \Lambda^*} \langle c_{v^*+u^*} | c_{v^*-u^*} \rangle \right| \leq \sum_{v^* \in \Lambda^*} \sqrt{\langle c_{v^*+2u^*} | c_{v^*+2u^*} \rangle \langle c_{v^*} | c_{v^*} \rangle}$$

$$\leq \sum_{v^* \in \Lambda^*} \left(\langle c_{v^*+2u^*} | c_{v^*+2u^*} \rangle + \langle c_{v^*} | c_{v^*} \rangle \right) = 2 \left\| \boldsymbol{f} \right\|^2_{2,\mathbb{R}^m/\Lambda} = 2$$

Finally, using Lemma 3, we have

$$\sum_{u^* \in \frac{1}{2}\Lambda^* \setminus 0} \rho_{s/2}(u^*) = \rho_s\left(\Lambda^* \setminus 0\right) \leq 2 \cdot \beta^{(m)}_{\frac{\lambda_1^*}{s}} .$$

Putting all together, we obtain that

$$\left| \|1_C \mathcal{F}_{\mathbb{R}^m}\{\boldsymbol{h}\}\|^2_{2,\mathbb{Z}^m} - \sum_{\ell^* \in \Lambda^*} \langle c_{\ell^*} | c_{\ell^*} \rangle \iota_C(\ell^*) \right| \leq \varepsilon_{\text{diag}} ,$$

where $\varepsilon_{\text{diag}} = 8 \cdot \beta^{(m)}_{\lambda_1^*/s}$.

Step 5. Recall the notation $B_{\delta\lambda_1^*}(\ell^*) = \{x \in \mathbb{Z}^m \mid |x - \ell^*| < \delta\lambda_1^*\}$. Whenever $\overline{B_{\delta\lambda_1^*}(\ell^*)} \subseteq C$, it obviously holds that

$$\iota_C(\ell^*) = \left(\frac{2}{s}\right)^m \sum_{y \in C} \rho_{s/2}(y - v^*) \geq \left(\frac{2}{s}\right)^m \sum_{y \in B_{\delta\lambda_1^*}(\ell^*)} \rho_{s/2}(y - \ell^*)$$

$$\geq \left(\frac{2}{s}\right)^m \rho_{s/2}(\mathbb{Z}^m) \left(1 - \beta^{(m)}_{2\delta\lambda_1^*/s}\right) \geq (1 - 2 \cdot \beta^{(m)}_{s/2})(1 - \beta^{(m)}_{2\delta\lambda_1^*/s}) ,$$

where the second inequality follows from Banaszczyk's bound (see Lemma 2) and the last from Lemma 4. It follows then that

$$\sum_{\ell^* \in \Lambda^*} \langle c_{\ell^*} | c_{\ell^*} \rangle \iota(\ell^*) \geq (1 - \varepsilon_{\text{smooth}}) \sum_{\substack{\ell^* \in \Lambda^* \\ B_{V\delta}(V\ell^*) \subseteq C}} \langle c_{\ell^*} | c_{\ell^*} \rangle .$$

where $\varepsilon_{\text{smooth}} = 2 \cdot \beta^{(m)}_{s/2} + \beta^{(m)}_{2\delta\lambda_1^*/s}$

Finalizing By collecting all the error terms, we obtain that

$$\|1_C \cdot \mathcal{F}_{\mathbb{D}^m}\{\boldsymbol{h}\}\|^2_{2,\hat{\mathbb{D}}^m}$$

$$\geq \sum_{\substack{\ell^* \in \Lambda^* \\ B_{\delta\lambda_1^*}(\ell^*) \subseteq C}} \langle c_{\ell^*} | c_{\ell^*} \rangle - \varepsilon_{\text{smooth}} - \varepsilon_{\text{diag}} - 5(\varepsilon_{\text{per}} + \varepsilon_{\text{lip}})$$

whenever s, δ and λ_1^* satisfy the following:

$$\frac{2\delta\lambda_1^*}{s} \geq \sqrt{m} \quad \text{and} \quad \frac{s}{2\sqrt{2}} \geq \sqrt{m}. \tag{12}$$

Part 3 Let $\mathcal{D}$ be the distribution defined by the output y of Algorithm 1 (recall that we assumed $V = 1$); note that $\mathcal{D}$ has support only on $[q]_c^m$. Throughout this part of the analysis, S denotes a subset of Λ^*.

By above analysis, we can conclude that whenever $B_{\delta\lambda_1^*}(S) \subseteq [q]_c^m$, we have (putting $C = B_{\delta\lambda_1^*}(S)$),

$$p_S := \mathcal{D}(B_{\delta\lambda_1^*}(S)) \geq \sum_{\ell^* \in S} \langle c_{\ell^*} | c_{\ell^*} \rangle - \eta',$$

where $\eta' = \varepsilon_{\text{smooth}} + \varepsilon_{\text{diag}} + \varepsilon_{\text{denom}} + 5(\varepsilon_{\text{per}} + \varepsilon_{\text{lip}})$.

For general $S \subseteq \Lambda^*$, write $S = S_0 \cup S_1$ as a disjoint union, where $S_0 = \{\ell^* \in S \mid B_{\delta\lambda_1^*}(\ell^*) \subseteq [q]_c^m\}$. Then it is evident that $S_1 \subseteq \Lambda^* \backslash [-\frac{q}{4}, \frac{q}{4}]^m$. Then, putting $\varepsilon_{\text{tail}} = \frac{4m\,\text{Lip}(\boldsymbol{f})^2}{\pi^2 q^2} \geq \sum_{\ell^* \in \Lambda^* \backslash [-\frac{q}{4}, \frac{q}{4}]^m} \langle c_{\ell^*} | c_{\ell^*} \rangle \geq \sum_{\ell^* \in S_1} \langle c_{\ell^*} | c_{\ell^*} \rangle$, (see Corollary 5), we have

$$\mathcal{D}(B_{\delta\lambda_1^*}(S)) \geq \mathcal{D}(B_{\delta\lambda_1^*}(S_0)) \geq \sum_{\ell^* \in S_0} \langle c_{\ell^*} | c_{\ell^*} \rangle - \eta' \geq \sum_{\ell^* \in S} \langle c_{\ell^*} | c_{\ell^*} \rangle - \varepsilon_{\text{tail}} - \eta',$$

$$= \sum_{\ell^* \in S} \langle c_{\ell^*} | c_{\ell^*} \rangle - \varepsilon_{\text{smooth}} - \varepsilon_{\text{diag}} - \varepsilon_{\text{denom}} - 5(\varepsilon_{\text{per}} + \varepsilon_{\text{lip}}) - \varepsilon_{\text{tail}} \tag{13}$$

5.3 Tuning Parameters

The left hand side of the table in Figure 4 collects the different error terms obtained above, considering $V = 1$. The general case is obtained simply by applying the above analysis to the function $\boldsymbol{f}_V := \boldsymbol{f}(V\cdot)$. The hidden lattice of $\boldsymbol{f}_V$ is $\frac{1}{V}\Lambda$, which has $V\Lambda^*$ as its dual, and the Lipschitz constant of $\boldsymbol{f}_V$ is $V\,\text{Lip}(\boldsymbol{f})$. Thus, the requirements on the parameters (see Equation (12)) change to

$$\frac{2\delta V\lambda_1^*}{s} \geq \sqrt{m} \quad \text{and} \quad \frac{s}{2\sqrt{2}} \geq \sqrt{m}, \tag{14}$$

and the different error terms become as listed in the table in Figure 4.

Error	$V = 1$	V arbitrary
$\varepsilon_{\text{denom}}$	$2\beta^{(m)}_{q/s}$	$2\beta^{(m)}_{q/s}$
$\varepsilon_{\text{smooth}}$	$2 \cdot \beta^{(m)}_{s/2} + \beta^{(m)}_{2\delta\lambda_1^*/s}$	$2 \cdot \beta^{(m)}_{s/2} + \beta^{(m)}_{2\delta V\lambda_1^*/s}$
$\varepsilon_{\text{diag}}$	$8\beta^{(m)}_{\lambda_1^*/s}$	$8\beta^{(m)}_{V\lambda_1^*/s}$
ε_{per}	$2s^{m/2}\beta^{(m)}_{\frac{s}{2\sqrt{2}}}$	$2s^{m/2}\beta^{(m)}_{\frac{s}{2\sqrt{2}}}$
ε_{lip}	$\frac{4\pi\sqrt{m}s^{m/2}\left(2\,\text{Lip}(\boldsymbol{f})+2\pi s^2\right)}{q}$	$\frac{4\pi\sqrt{m}s^{m/2}\left(2V\,\text{Lip}(\boldsymbol{f})+2\pi s^2\right)}{q}$
$\varepsilon_{\text{tail}}$	$\frac{m\,\text{Lip}(\boldsymbol{f})^2}{\pi^2 q^2}$	$\frac{mV^2\,\text{Lip}(\boldsymbol{f})^2}{\pi^2 q^2}$

Fig. 4. Change of the errors when applying the analysis to f_V

Recall that $\beta^{(m)}_z := \left(\frac{2\pi e z^2}{m}\right)^{m/2} e^{-\pi z^2}$ and $N = 2^n$. We can now choose the parameters s, V and q of the algorithm appropriately to enforce all the error terms to be small. In detail, we can select:

- $s \in O(\sqrt{m\log(\eta^{-1})})$ so that $5\varepsilon_{\text{per}} \leq \eta/6$, and $2\beta^{(m)}_{s/2} \leq \eta/12$ in $\varepsilon_{\text{smooth}}$.
- $V \in O(\frac{\sqrt{m\log(\eta^{-1})}s}{\delta\lambda_1^*}) = O(\frac{m\log(\eta^{-1})}{\delta\lambda_1^*})$ so that $\varepsilon_{\text{smooth}}, \varepsilon_{\text{diag}} \leq \eta/6$.
- $Q = \log(q) \in O(m\log(s)+\log(V)+\log(\text{Lip}(f))+\log(\eta^{-1}))$ so that $5\varepsilon_{\text{lip}} \leq \eta/6$ and $\varepsilon_{\text{tail}} \leq \eta/6$.

With the above choice of parameters, $\varepsilon_{\text{smooth}} + \varepsilon_{\text{diag}} + \varepsilon_{\text{denom}} + 5(\varepsilon_{\text{per}} + \varepsilon_{\text{lip}}) + \varepsilon_{\text{tail}} \leq \eta$ in Equation (13). Unrolling the expression of $Q = \log(q)$ and recalling that the quantum Fourier transform requires a quadratic number of gates [25, Ch. 5], we obtain the main theorem.

Theorem 2. *Algorithm 1 solves the Dual Lattice Sampling Problem with parameters η and δ; it uses m calls to the Gaussian superposition subroutine (see Theorem 5), one quantum oracle call to $\boldsymbol{f}$, $mQ + n$ qubits, and $O(m^2Q^2)$ quantum gates, where*

$$Q = O\left(m\log\left(m\log\frac{1}{\eta}\right)\right) + O\left(\log\left(\frac{\text{Lip}(\boldsymbol{f})}{\eta\cdot\delta\lambda_1^*}\right)\right). \tag{4}$$

6 From Sampling to Full Dual Lattice Recovery

We have so far focused on approximate sampling dual lattice points following weights $\langle c_{\ell^*}|c_{\ell^*}\rangle$ for $\ell^* \in \Lambda^*$, regardless of how useful this distribution may be. Indeed, until now, it could be that the function $\boldsymbol{f} : \mathbb{R}^m/\Lambda \to \mathcal{S}$ is constant,

and therefore that the weight is concentrated on $0 \in \Lambda^*$. We would like now make sure we can reconstruct (approximately) Λ^* from such samples, i.e., that a sufficient number of sampled vectors from Λ^* will generate it. Informally, an equivalent condition is that the weight $\langle c_{\ell^*} | c_{\ell^*} \rangle$ is not concentrated on any proper sublattice $M^* \subsetneq \Lambda^*$. More formally, we give the following sufficient conditions.

Definition 4. *Let $L \subseteq \mathbb{R}^m$ be a full-rank lattice. A distribution $\mathcal{D}$ on L is called* p-evenly distributed *whenever* $\Pr_{v \leftarrow \mathcal{D}}[v \in L'] \leq p$ *for any proper sublattice $L' \subsetneq L$.*

Definition 5. *Let $L \subseteq \mathbb{R}^m$ be a full-rank lattice. A distribution $\mathcal{D}$ on L is called* (R, q)-concentrated *whenever* $\Pr_{v \leftarrow \mathcal{D}}[\|v\| \geq R] \leq q$.

Lemma 5. *Let $L \subseteq \mathbb{R}^m$ be a full-rank lattice with a p-evenly distributed and (R, q)-concentrated distribution $\mathcal{D}$. Denote by S the random variable defined by the number of samples that needs to be drawn from $\mathcal{D}$ such that the samples together generate L as a lattice. Then, for all $\alpha > 0$,*

$$\Pr\left[S > (2+\alpha) \cdot \frac{(t+m)}{1-p-q}\right] \leq \exp(-\alpha(t+m)/2)$$

where $t = m \log_2(R) - \log_2(\det(L))$.

Proof. First, we define the following sublattices of L, for any $v_1, \ldots, v_{j-1} \in L$.

$$L_{v_1,\ldots,v_{j-1}} = \begin{cases} \mathrm{span}_{\mathbb{R}}(v_1, \ldots, v_{j-1}) \cap L & \text{if } \dim(\mathrm{span}_{\mathbb{R}}(v_1, \ldots, v_{j-1})) < m \\ \langle v_1, \ldots, v_{j-1} \rangle & \text{otherwise.} \end{cases}$$

Consider a sequence of samples $(v_i)_{i>0}$ (from $\mathcal{D}$). We call v_j 'good' whenever $\|v_j\| \leq R$ and $v_j \notin L_{v_1,\ldots,v_{j-1}}$. We argue that we need at most $m+t$ good vectors to generate L.

Denote L' for the lattice generated by the $m+t$ good vectors. Then the first m good vectors ensure that L' is of rank m, whereas the last t good vectors will reduce the index of the L' lattice in L. Calculating determinants, using the fact that all good vectors are bounded by R, we have $\det(L') \leq R^m/2^t \leq \det(L)$. This yields $L' = L$.

Denote by X the random variable having the negative binomial distribution with success probability $p+q$ and number of 'failures' $m+t$. That is, X is the number of independent samples from a $(p+q)$-Bernoulli distribution until $m+t$ 'failures' [7] are obtained. We argue that the random variable S is dominated by the random variable X, i.e., $\Pr[S > x] \leq \Pr[X > x]$ for every $x \in \mathbb{N}$.

Again, consider a sequence of samples $(v_i)_{i>0}$ (from $\mathcal{D}$). The probability of v_j being a 'good' vector is at least $1-p-q$, by the fact that $\mathcal{D}$ is (R, q)-concentrated and p-evenly distributed. Because at most $m+t$ 'good' vectors are needed to generate the whole lattice, S is indeed dominated by X. Therefore, for any $k \in \mathbb{N}$,

$$\Pr\left[S > \frac{t+m+k}{1-p-q}\right] \leq \Pr\left[X > \frac{t+m+k}{1-p-q}\right] \leq \Pr\left[B < m+t\right] \qquad (15)$$

[7] In our case, the failures are the 'good' vectors. We nonetheless chose the word 'failure' because it is standard nomenclature for the negative binomial distribution.

$$\leq \exp\left(-\frac{1}{2}\frac{k^2}{t+m+k}\right)$$

where B is binomially distributed with $\lfloor \frac{t+m+k}{1-p-q} \rfloor$ trials and success probability $1-p-q$. The first inequality follows from the fact that S is upper bounded by X. The second inequality comes from the close relationship between the negative binomial distribution and the binomial distribution [12, Ch. 8, Ex. 17]. The last inequality follows from Chernoff's bound. Putting $k = (1+\alpha)(t+m)$ into Equation (15) yields the claim. □

We conclude by relating the parameters (a, r, ϵ) of the HSP oracle (Definition 1) $\boldsymbol{f} : \mathbb{R}^m/\Lambda \to \mathcal{S}$ and the assumption used in the above Lemma 5.

Lemma 6. *Let Λ be a lattice, and let $M \supsetneq \Lambda$ a proper super-lattice of Λ. Then there exists a $v \in M$ such that $d(v, \Lambda) \geq \lambda_1(\Lambda)/3$.*

Proof. Let $w \in M$ be the shortest non-zero vector in M and write $\|w\| = \alpha\lambda_1(\Lambda)$ for $\alpha < 1$. We show that $v = \lceil \frac{1}{3\alpha} \rceil \cdot w \in M$ suffices. If $\alpha \geq 1/3$ this is certainly true. For $\alpha < 1/3$ it is clear that $\|v\| \geq \lambda_1(\Lambda)/3$ and $\|v\| \leq \lambda_1(\Lambda)/3 + \|w\| \leq \frac{2}{3}\lambda_1(\Lambda)$. In particular, for any $\ell \in \Lambda \setminus \{0\}$, $\|v - \ell\| \geq \lambda_1(\Lambda) - \|v\| \geq \lambda_1(\Lambda)/3$. Therefore, $d(v, \Lambda) \geq \lambda_1(\Lambda)/3$. □

Lemma 7. *Let Λ be a lattice and $M \supsetneq \Lambda$ a proper super-lattice of Λ. Then the number $N = \left|\left\{c \in M/\Lambda \mid d(c, \Lambda) < \frac{1}{6}\lambda_1(\Lambda)\right\}\right|$ of close cosets is at most $\frac{1}{2} \cdot |M/\Lambda|$.*

Proof. By Lemma 6 there exists a $v \in M$ such that $d(v, \Lambda) \geq \frac{1}{3}\lambda_1(\Lambda)$. Denoting $T = \left\{c \in M/\Lambda \mid d(c, \Lambda) < \frac{1}{6}\lambda_1(\Lambda)\right\}$, we can deduce that $T \cup (T + v)$ is a disjoint union in M/Λ. Indeed, elements $c \in T$ satisfy $d(c, \Lambda) \leq \frac{1}{6}\lambda_1(\Lambda)$, whereas $c' \in T + v$ satisfy $d(c', \Lambda) \geq d(v, \Lambda) - \frac{1}{6}\lambda_1(\Lambda) \geq \frac{1}{6}\lambda_1(\Lambda)$. Therefore $N = |T| \leq \frac{1}{2}|M/\Lambda|$. □

Lemma 8. *Let $\boldsymbol{f} : \mathbb{R}^m \to \mathcal{S}$ be an (a, r, ϵ)-HSP oracle of the full-rank lattice $\Lambda \subset \mathbb{R}^m$, with $r \leq \lambda_1(\Lambda)/6$. Let $\mathcal{D}_{\text{ideal}}$ be the distribution supported by Λ^*, with weight $\langle c_{\ell^*} | c_{\ell^*} \rangle$ at $\ell^* \in \Lambda^*$, where $|c_{\ell^*}\rangle$ are the vectorial Fourier coefficients of the function $\boldsymbol{f}$. Then $\mathcal{D}_{\text{ideal}}$ is both $(\frac{1}{2} + \epsilon)$-evenly distributed and $(R, \frac{ma^2}{4\pi^2 R^2})$-concentrated for any $R > 0$.*

Proof. The distribution $\mathcal{D}_{\text{ideal}}$ being $(R, \frac{ma^2}{4\pi^2 R^2})$-concentrated for any $R > 0$ is a direct consequence of Corollary 5. For the $(\frac{1}{2} + \epsilon)$-evenly distributed part, we argue as follows. Let M^* be any strict sublattice of Λ^*, and let M be its dual, which is then a superlattice of Λ. Put $\boldsymbol{f}|^{\mathbb{R}^m/M}(x) = \frac{1}{|M/\Lambda|}\sum_{v \in M/\Lambda} \boldsymbol{f}(x + v)$, the periodization of $\boldsymbol{f}$ with respect to $\mathbb{R}^m/M$ (c.f. Definition 3). We have the following sequence of equalities, of which the first follows from the Poisson

summation formula (see Theorem 7).

$$\begin{aligned}\sum_{v^*\in M^*}\langle c_{v^*}|c_{v^*}\rangle &= \left\|\boldsymbol{f}|^{\mathbb{R}^m/M}\right\|_{2,\mathbb{R}^m/M} = \frac{1}{\det M}\int_{x\in\mathbb{R}^m/M}\left\langle\boldsymbol{f}|^{\mathbb{R}^m/M}\middle|\boldsymbol{f}|^{\mathbb{R}^m/M}\right\rangle dx,\\ &= \frac{1}{|M/\Lambda|^2}\sum_{v,w\in M/\Lambda}\underbrace{\frac{1}{\det M}\int_{x\in\mathbb{R}^m/M}\langle\boldsymbol{f}(x+v)|\ \boldsymbol{f}(x+w)\rangle\,dx}_{I_{v,w}}\\ &= \frac{1}{|M/\Lambda|^2}\sum_{\substack{v,w\in M/\Lambda\\ d_{\mathbb{R}^m/\Lambda}(v,w)<r}} I_{v,w} + \frac{1}{|M/\Lambda|^2}\sum_{\substack{v,w\in M/\Lambda\\ d_{\mathbb{R}^m/\Lambda}(v,w)\geq r}} I_{v,w}\end{aligned}$$

By the definition of an (a, r, ϵ)-oracle, we have that $|I_{v,w}| \leq \epsilon$ whenever $d_{\mathbb{R}^m/\Lambda}(v, w) \geq r$. In the rest of the cases we have $|I_{v,w}| \leq 1$, because f maps to the unit sphere. Above expression is therefore bounded by $\frac{|M/\Lambda \cap \mathcal{B}_r|}{|M/\Lambda|} + \epsilon$, where $\mathcal{B}_r$ is the open unit ball with radius r. By Lemma 7, we have $\frac{|M/\Lambda \cap r\mathcal{B}|}{|M/\Lambda|} \leq \frac{1}{2}$ for $r \leq \lambda_1(\Lambda)/6$. Summarizing all results, we conclude that

$$\sum_{v^*\in M^*}\langle c_{v^*}|c_{v^*}\rangle \leq \frac{1}{2} + \epsilon.$$

Since M^* was chosen arbitrarily, we can conclude that $\mathcal{D}_{\text{ideal}}$ is $(\frac{1}{2} + \epsilon)$-evenly distributed. □

Remark 8. A similar reasoning happens in [28, Lecture 12], though it specifically targets the discrete Gaussian distribution on lattices. Despite being not general enough for our purposes, it may well be helpful for optimizing a future specialization.

References

1. Banaszczyk, W.: New bounds in some transference theorems in the geometry of numbers. Mathematische Annalen **296**(4), 625–636 (1993), `http://eudml.org/doc/165105`
2. Biasse, J.F., Song, F.: Efficient quantum algorithms for computing class groups and solving the principal ideal problem in arbitrary degree number fields. In: Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms. pp. 893–902. Society for Industrial and Applied Mathematics (2016)
3. de Boer, K., Ducas, L., Fehr, S.: On the quantum complexity of the continuous hidden subgroup problem. Cryptology ePrint Archive, Report 2019/716 (2019), `https://eprint.iacr.org/2019/716`
4. Buchmann, J., Kessler, V.: Computing a reduced lattice basis from a generating system (08 1996)
5. Buchmann, J., Pohst, M.: Computing a lattice basis from a system of generating vectors. In: Proceedings of the European Conference on Computer Algebra. pp. 54–63. EUROCAL '87, Springer-Verlag, London, UK, UK (1989), `http://dl.acm.org/citation.cfm?id=646658.700556`

6. Chang, X., Stehlé, D., Villard, G.: Perturbation analysis of the QR factor R in the context of LLL lattice basis reduction. Math. Comput. **81**(279), 1487–1511 (2012). https://doi.org/10.1090/S0025-5718-2012-02545-2
7. Cramer, R., Ducas, L., Peikert, C., Regev, O.: Recovering short generators of principal ideals in cyclotomic rings. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 559–585. Springer (2016)
8. Cramer, R., Ducas, L., Wesolowski, B.: Short Stickelberger class relations and application to ideal-SVP. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 324–348. Springer (2017)
9. Deitmar, A., Echterhoff, S.: Principles of Harmonic Analysis. Springer Publishing Company, Incorporated, 2nd edn. (2016)
10. Ducas, L., Plançon, M., Wesolowski, B.: On the shortness of vectors to be found by the ideal-SVP quantum algorithm. In: Boldyreva, A., Micciancio, D. (eds.) Advances in Cryptology – CRYPTO 2019. pp. 322–351. Springer, Cham (2019)
11. Eisenträger, K., Hallgren, S., Kitaev, A., Song, F.: A quantum algorithm for computing the unit group of an arbitrary degree number field. In: Proceedings of the forty-sixth annual ACM symposium on Theory of computing. pp. 293–302. ACM (2014)
12. Graham, R.L., Knuth, D.E., Patashnik, O.: Concrete Mathematics: A Foundation for Computer Science. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 2nd edn. (1994)
13. Grover, L., Rudolph, T.: Creating superpositions that correspond to efficiently integrable probability distributions. arXiv preprint quant-ph/0208112 (2002)
14. Hales, L., Hallgren, S.: An improved quantum fourier transform algorithm and applications. In: Proceedings 41st Annual Symposium on Foundations of Computer Science. pp. 515–525 (Nov 2000). https://doi.org/10.1109/SFCS.2000.892139
15. Hallgren, S.: Fast quantum algorithms for computing the unit group and class group of a number field. In: Proceedings of the thirty-seventh annual ACM symposium on Theory of computing. pp. 468–474. ACM (2005)
16. Hallgren, S.: Polynomial-time quantum algorithms for Pell's equation and the principal ideal problem. Journal of the ACM (JACM) **54**(1), 4 (2007)
17. Heinonen, J.: Lectures on Lipschitz analysis `http://www.math.jyu.fi/research/reports/rep100.pdf`
18. Kitaev, A., Webb, W.A.: Wavefunction preparation and resampling using a quantum computer. arXiv preprint arXiv:0801.0342 (2008)
19. Kuperberg, G.: A subexponential-time quantum algorithm for the dihedral hidden subgroup problem. SIAM Journal on Computing **35**(1), 170–188 (2005)
20. Micciancio, D., Regev, O.: Worst-case to average-case reductions based on gaussian measures. SIAM J. Comput. **37**(1), 267–302 (Apr 2007). https://doi.org/10.1137/S0097539705447360
21. Miller, S.D., Stephens-Davidowitz, N.: Generalizations of Banaszczyk's transference theorems and tail bound. arXiv preprint arXiv:1802.05708 (2018)
22. Mosca, M., Ekert, A.: The hidden subgroup problem and eigenvalue estimation on a quantum computer. In: NASA International Conference on Quantum Computing and Quantum Communications. pp. 174–188. Springer (1998)
23. National Institute of Standards and Technology: Post-quantum cryptography standardization (2017), `https://csrc.nist.gov/Projects/Post-Quantum-Cryptography/Post-Quantum-Cryptography-Standardization`
24. Nguyen, P.Q., Stehlé, D.: An lll algorithm with quadratic complexity. SIAM Journal on Computing **39**(3), 874–903 (2009)

25. Nielsen, M.A., Chuang, I.L.: Quantum Computation and Quantum Information: 10th Anniversary Edition. Cambridge University Press, New York, NY, USA, 10th edn. (2011)
26. Novocin, A., Stehlé, D., Villard, G.: An lll-reduction algorithm with quasi-linear time complexity. In: Proceedings of the forty-third annual ACM symposium on Theory of computing. pp. 403–412. ACM (2011)
27. Pellet-Mary, A., Hanrot, G., Stehlé, D.: Approx-SVP in ideal lattices with pre-processing. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 685–716. Springer (2019)
28. Regev, O.: Lecture notes in 'lattices in computer science' (November 2004)
29. Regev, O.: Quantum computation and lattice problems. SIAM Journal on Computing **33**(3), 738–760 (2004)
30. Reiter, M., Arthur, S.: Fourier transform & solobev spaces (lecture notes) (2008), `https://www.mat.univie.ac.at/~stein/teaching/SoSem08/sobolev_fourier.pdf`
31. Schmidt, A., Vollmer, U.: Polynomial time quantum algorithm for the computation of the unit group of a number field. In: Proceedings of the thirty-seventh annual ACM symposium on Theory of computing. pp. 475–480. ACM (2005)
32. Shor, P.W.: Algorithms for quantum computation: Discrete logarithms and factoring. In: Proceedings 35th annual symposium on foundations of computer science. pp. 124–134. IEEE (1994)
33. Song, F.: Email, from september 2018
34. Song, F.: Quantum Computing: A Cryptographic Perspective. Ph.D. thesis, The Pennsylvania State University (2013), `https://etda.libraries.psu.edu/files/final_submissions/8820`
35. Villani, A.: Another note on the inclusion $l^p(\mu) \subset l^q(\mu)$. The American Mathematical Monthly **92**(7), 485–487 (1985), `http://www.jstor.org/stable/2322503`
36. Werner, D.: Funktionalanalysis. Springer-Lehrbuch, Springer Berlin Heidelberg (2007)
37. Yudin, V.A.: The multidimensional Jackson theorem. Mathematical notes of the Academy of Sciences of the USSR **20**(3), 801–804 (Sep 1976). https://doi.org/10.1007/BF01097255