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 (Eisentrager 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. Incidentally, our work clarifies certain claims from the extended abstract of Eisentrager et al.

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

1 Introduction

The Hidden Subgroup Problem. Among all quantum algorithms, Shor's algorithm [25] for factoring and finding discrete logarithms stands out as demonstrating 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 [19], in preparation of a global update of widely deployed encryption and authentication protocols.

The core idea of quantum period finding from [25] is not limited to factoring and discrete logarithm, and the Hidden Subgroup Problem formalized in [18] 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 [22]; however, no polynomial time algorithm is known for those cases, and the best known algorithm has sub-exponential complexity [15], 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 $f : G \to R$. 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 [25] 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 proposed by Hallgren [12] 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 [11,24], 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, proposed by Eisentrager, Hallgren, Kitaev and Song in an extended abstract [8], 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 f is assumed to produce a "quantum output". More formally, $f : \mathbb{R}^m \to \mathcal{S}$, $x \mapsto |f(x)\rangle$, 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 $|x\rangle|f(x)\rangle$. A crucial observation here is that $|f(x)\rangle$ and $|f(y)\rangle$ are *not* necessarily orthogonal (or even distinct) for distinct x and y modulo Λ. In other words, it is not assumed that f is *strictly* periodic, but merely that $|f(x)\rangle$ and $|f(y)\rangle$ are "somewhat orthogonal" for x and y that are "not too close" modulo Λ, and that f is Lipschitz continuous.

In their extended abstract [8] Eisentrager et al. consider a variation of the standard HSP algorithm in order to tackle the Continuous HSP problem. In order to deal with the continuous nature of the domain $\mathbb{R}^m$ of f, the proposed 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 [8], applying the quantum Fourier transform directly would make the resulting approximation errors difficult to analyze.

As an application, Eisentrager et al. also propose a quantum polynomial time algorithm for computing the unit group of a number field in their article [8]. 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 [4,5,21]. While the cryptanalytic consequences for ideal-lattice based cryptography seems limited so far [7], these results demonstrate a hardness gap between ideal lattices and general ones.

The Analysis by Eisentrager et al. While demonstrating that a class of problems admits a quantum polynomial time algorithm is typically sufficiently satisfactory from a theoretical perspective, the potential cryptanalytic implication of efficiently solving the Continuous HSP invites us to refine the complexity analysis. This is the main purpose of this paper.

One difficulty for this study is that the extended abstract of Eisentrager et al. has to this date not been followed by a public full version. Certainly, the extended abstract does give a credible approach to the problem at hand, by illustrating 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, due to the absence of a full treatment of certain claims in the analysis and due the formulation of the main claim about the solvability of the continuous HSP, the quantitative aspects of their result remain unclear.

In more detail, neither the statement of Theorem 6.1 nor its analysis in [8] addresses the *dependency* of the (claimed to be polynomial time) running time on the parameters of the Continuous HSP. For example, a *constant* function satisfies Definition 1.1 in [8] of being a Continuous HSP instance for *any* lattice L with parameter $\epsilon = 1$, or, similarly, with $\epsilon < 1$ but r being greater than the covering radius of L; yet, such a function makes the Continuous HSP problem vacuously hard. Similarly, even when the parameters of the Continuous HSP are constant (and in a meaningful region so as to avoid the above kind of counter examples), it is unclear how the expected "quality" of the output (in terms of precision and success probability) affects the running time. The proposed algorithm is clearly polynomial time in the number of qubits it acts on; however, while [8] argues that in the theoretical limit of infinitely many qubits the algorithm works perfectly, the "rate of convergence" remains unclear.

An additional complication is that it may not be clear what polynomial-time formally means when the input is an oracle. For example, in an application of the Continuous HSP algorithm it may be critical to know whether the running time grows polynomially in the the Lipschitz constant of f (which is one of the

parameters of the Continuous HSP), or polynomially in its logarithm. We will show that it is actually the latter.

Our work. The goal of this paper is to provide a rigorous refined analysis of (a slightly modified version of) the Continuous HSP quantum algorithm proposed by Eisentrager et al. [8]. 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 proposed by Eisentrager et al. [8] in the following two points: First, we specify the initial state of the algorithm to have Gaussian weight, while [8, Sec. 6.2] suggests to use a cropped trigonometric function; 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 [16]. Secondly, we do not make use of a "remote" Fourier transform, as its advantages were unclear to us.

Our analysis is divided into two parts, which are summarized by formal statements in Section 2.2 and Section 2.3. In the first part, 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 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 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 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 by $O(m^2Q^2)$.

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

Remark 1. Recovering the *exact* hidden lattice is outside the scope of this work, since this task is application-dependent. For instance, when applying this algorithm to compute the unit group $\mathcal{O}_K^\times$ of a number field K, the hidden lattice will be the so-called Log-unit lattice which is irrational. Yet, a sufficiently good approximation of the logarithm of a unit yields the exact underlying unit, simply

by taking the exponential and rounding it to the closest element in the ring of integers $\mathcal{O}_K$.

Remark 2. An auxiliary task that we rely upon for the first step is the preparation of (an approximation of) a quantum superposition according to Gaussian weight. This task is known to take quantum polynomial time [10,14]. In Appendix A we propose a refined analysis for the cost of the algorithm of [14], demonstrating a complexity of $O(Q+k)$ qubits and $O(Qk^{3/2})$ for approximating it to a precision 2^{-k} over an interval of integers of length 2^Q. This is summarized and formalized as Theorem 3 in Section 2.4.

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

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

that is periodic with respect to a full rank lattice $\Lambda \subset \mathbb{R}^m$; hence, f may be understood as a function $\mathbb{R}^m/\Lambda \to \mathcal{S}$. The function f is assumed to be Lipschitz continuous with Lipschitz constant $\mathrm{Lip}(f)$. Later, we will also require f to be "sufficiently non-constant". One should think of f as an oracle that maps a classical input x to a quantum state $|f(x)\rangle$ over n qubits.

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 $x \in X$.

2.2 Dual Lattice Sampling Problem

Recalling from the introduction, the Continuous HSP is the problem of recovering the hidden lattice Λ when given oracle access to a function f as discussed above. For the purpose of a more modular analysis, we first consider the following 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 $f : \mathbb{R}^m/\Lambda \to \mathcal{S}$ (see Section 3).

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

$$p_S := \mathcal{D}\big(\mathcal{B}_{\delta\lambda_1^*}(S)\big) \geq \left(\sum_{\ell^* \in S} \langle c_{\ell^*} | c_{\ell^*} \rangle\right) - \eta. \tag{1}$$

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 1. *Algorithm 1 solves the Dual Lattice Sampling Problem with oracle function f, error parameter η and relative distance parameter δ, using m calls to the Gaussian superposition subroutine (see Theorem 3), one quantum oracle call to f, $O(mQ)$ qubits, and $O(m^2Q^2)$ quantum gates, where*

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

2.3 Full Dual Lattice Recovery

Recovering the full lattice (or equivalently its dual) requires an extra assumption on the oracle function f, as captured by the third condition in the following definition, reformatted from Definition 1.1 of [8].

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

- *f is Λ-periodic,*
- *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 $|\langle f(x)|f(y)\rangle| \leq \epsilon$,*

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.

According to Eisentrager et al. [8], for (some undetermined) adequate parameters, the above definition 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 2. *Let $f : \mathbb{R}^m \to \mathcal{S}$ be an (a, r, ϵ)-HSP oracle of the full-rank lattice $\Lambda \subset \mathbb{R}^m$, with $r \leq \lambda_1(\Lambda)/6$ and $\epsilon \in [0, 1/3]$. Let $\mathcal{D}_{\mathrm{ideal}}$ 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 f.*

Denote by S the random variable defined by the number of samples that needs 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 $t = m\log_2(ma^2) - \log_2(\det(L))$.

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

Remark 3. To optimize certain applications, it may be suboptimal to apply the above theorem in a black-box manner. For example, when solving multiple instances of the Principal Ideal Problem [2] in a fixed field K, one should note that all instances share a common hidden sublattice, namely the Logarithmic unit lattice $\text{Log}\,\mathcal{O}_K^\times$.

2.4 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 [10,14], and we provide a quantitative analysis in Appendix A. The precise running time of preparing this Gaussian state is summarized below.

Theorem 3. *For any positive integers q, k and for any $s > 1$, there exists a quantum algorithm that prepares the one-dimensional Gaussian state*

$$\frac{1}{\sqrt{\rho_{1/s}(\frac{1}{q}[q]_c)}} \sum_{x \in \frac{1}{q}[q]_c} \sqrt{\rho_{1/s}(x)}|x\rangle \tag{2}$$

up to trace distance $se^{-\pi s^2/8} + Q \cdot 2^{-k}$ using $O(Q + k)$ qubits and $O(Q \cdot k^{3/2} \cdot \text{polylog}(k))$ 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 8 with parameters $\mu = q/2$ and $\sigma = \sqrt{2}q/s$ and relabeling the basis states. Whenever above theorem is used as a subroutine in Theorem 1, choosing $k = \log(mQ/\eta^2)$ is sufficient, causing merely an extra error of η^2.

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 [6].

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, 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

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

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

$$\hat{\mathbb{D}}_{\text{rep}}^m := [q]_c^m := \mathbb{Z}^m \cap [-\tfrac{q}{2}, \tfrac{q}{2})^m \,.$$

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 $p \geq 1$, $L_p(G)$ denotes the vector space of measurable functions $f : G \to \mathbb{C}$ with finite norm $\|f\|_p$ (modulo the functions with vanishing norm), where

$$\|f\|_p^p := \int_{g \in G} |f(g)|^p d\mu \,.$$

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.

Theorem 4 (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 $\widehat{fg} = \hat{f} \star \hat{g}$ 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 and any choice of the Haar measure μ_H, there exists a Haar measure $\mu_{G/H}$ for G/H such that the *quotient integral formula*

$$\int_{G/H} \left(\int_H f(g + h) d\mu_H(h) \right) d\mu_{G/H}(g + H) = \int_G f(g) d\mu(g) \qquad (3)$$

holds for any continuous function $f : G \to \mathbb{C}$ with compact support (see [6, 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 [6, Section 3.6], see also Fig. 2).

Theorem 5 (Poisson Summation Formula). *For continuous $f \in L^1(G)$,*

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

Fig. 2. Informal illustration of Theorem 5 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* $\mathcal{F}_{\mathbb{D}^m}\{t|_{\mathbb{D}^m}\} = \mathcal{F}_{\mathbb{T}^m}\{t\}|^{\hat{\mathbb{D}}^m}$.

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

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

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 $\mathcal{F}_{\mathbb{D}^m}\{h\}$ instead of $\mathcal{F}_{\mathbb{D}^m}\{h|_{\mathbb{D}^m}\}$, taking it as understood that h is restricted to $\mathbb{D}^m$ when applying $\mathcal{F}_{\mathbb{D}^m}$.

Proof. Using a result of Yudin [28, Example I after Theorem 2], there exists a trigonometric approximation t of h, i.e. a function $t : \mathbb{T}^m \to \mathbb{C}$ with $\hat{t}(x) := \mathcal{F}_{\mathbb{T}^m}\{t\}(x) = 0$ for all $x \notin [q]^m_c$ so that $\|h - t\|_\infty \leq \pi\sqrt{m}\,\mathrm{Lip}(h)/q$. Recalling that $\hat{\mathbb{D}}^m \simeq \mathbb{Z}^m/q\mathbb{Z}^m$, the fact that $\hat{t} : \mathbb{Z}^m \to \mathbb{C}$ vanishes outside of $[q]^m_c$ implies for all $x \in [q]^m_c$ that

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

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

$$\big| \|1_C \cdot \mathcal{F}_{\mathbb{D}^m}\{h\}\|_{2,\hat{\mathbb{D}}^m} - \|1_C \cdot \mathcal{F}_{\mathbb{T}^m}\{h\}\|_{2,\mathbb{Z}^m} \big|$$
$$\leq \|1_C \cdot \mathcal{F}_{\mathbb{D}^m}\{h - t\}\|_{2,\hat{\mathbb{D}}^m} + \|1_C \cdot \mathcal{F}_{\mathbb{T}^m}\{h - t\}\|_{2,\mathbb{Z}^m}\,.$$

Observing that $\|1_C \cdot \mathcal{F}\{h - t\}\|_2 \leq \|\mathcal{F}\{h - t\}\|_2 = \|h - t\|_2 \leq \|h - t\|_\infty$ for both $\mathbb{D}^m$ and for $\mathbb{Z}^m$, this proves the claim.

11

3.5 Fourier Transform on Functions with Multidimensional Codomain

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

Important for us is the case $f : \mathbb{R}^m/\Lambda \to \mathcal{H}$. Spelling out the above, we get

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

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

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

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.

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 4. 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., $\left\|\sigma^{m/2} \cdot \sqrt{\rho_{1/\sigma}(x)}\right\|_{2,\mathbb{R}^m}^2 = 1$.

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* [3] , as introduced by Micciancio and Regev [16]. They allow us to control

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

---

[3] 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.

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 [17, 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 $\mathcal{B}_r = \mathcal{B}_r(0) = \{x \in \mathbb{R}^m \mid |x| < r\}$.

Imitating techniques from [16, Lemma 3.2], we have:

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

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*

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

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

The following lemma, which combines [16, Lemma 4.1] and [16, Lemma 3.2], controls the fluctuation of the sum $\rho_s(\Lambda + t)$ for varying $t \in \mathbb{R}^m$.

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* $t \in \mathbb{R}^m$,

$$(1 - 2\beta_{\sigma\lambda_1(\Lambda^*)}^{(m)})\frac{\sigma^m}{\det \Lambda} \leq \rho_\sigma(\Lambda + t) \leq (1 + 2\beta_{\sigma\lambda_1(\Lambda^*)}^{(m)})\frac{\sigma^m}{\det \Lambda}. \qquad (4)$$

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 $2\beta_{\sigma\lambda_1(\Lambda^*)}^{(m)} \leq 2\beta_{\sqrt{m}}^{(m)} \leq 1$ yields the result.

3.7 Lipschitz Condition

Theorem 7 (Rademacher's theorem). *A Lipschitz function* $f : \mathbb{R}^m/\Lambda \to \mathcal{H}$ *has weak partial derivatives* $\partial_{x_j} f : \mathbb{R}^m/\Lambda \to \mathcal{H}$ *lying in* $L_2(\mathbb{R}^m/\Lambda)$. *In particular,* $\left\|\partial_{x_j} f\right\|_{2,\mathbb{R}^m/\Lambda}^2 \leq \mathrm{Lip}(f)^2$.

Proof. Combining the proof of [13, Theorem 4.1 and 4.9] and [26, Theorem 2] on measures of compact sets, we obtain this result.

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

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

Proof. Since f is Lipschitz, we can apply Theorem 7. Furthermore, the identity $|f(x)\rangle = \sum_{\ell^* \in \Lambda^*} |c_{\ell^*}\rangle e^{2\pi i \langle x, \ell^* \rangle}$ implies $|\partial_{x_j} f(x)\rangle = 2\pi i \sum_{\ell^* \in \Lambda^*} \ell_j^* |c_{\ell^*}\rangle e^{2\pi i \langle x, \ell^* \rangle}$ almost everywhere ([27, Lemma V.2.11] or [23, Lemma 2.16]). Finally, given that $\left\| \partial_{x_j} f \right\|_{2, \mathbb{R}^m/\Lambda}^2 \leq \operatorname{Lip}(f)^2$, Plancherel's identity implies that

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

$$\geq 4\pi^2 \sum_{\substack{\ell^* \in \Lambda^* \\ \|\ell^*\|_2 \geq B}} \|\ell^*\|_2^2 \langle c_{\ell^*} | c_{\ell^*} \rangle \geq 4B^2 \pi^2 \sum_{\substack{\ell^* \in \Lambda^* \\ \|\ell^*\|_2 \geq B}} \langle c_{\ell^*} | c_{\ell^*} \rangle ,$$

from which the claim follows.

4 Algorithm

4.1 The Algorithm

Given a Λ-periodic function $f : \mathbb{R}^m \to \mathcal{S}$ as discussed in Section 2, which maps a classical input x to a quantum state $|f(x)\rangle$, we consider the following quantum algorithm (see Figure 3). The algorithm has oracle access to f, meaning that it has access to a unitary that maps $|x\rangle|0\rangle$ to $|x\rangle|f(x)\rangle$. 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 $|x\rangle|f(Vx)\rangle$ 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 $x \in \mathbb{D}_{\text{rep}}^m$ and $y \in \hat{\mathbb{D}}_{\text{rep}}^m$ (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

14

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

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

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.

$V > 0$. 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.

The description and Analysis of Step 1 is deferred to Appendix A. It will be shown (as summarized in Theorem 3) 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,

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

which we can rewrite as

$$|\psi\rangle = s^{m/2} \sum_{k=1}^{N} \sum_{x\in\mathbb{D}^m} \sqrt{\rho_{1/s}(x)}\, |x\rangle_{\mathbb{D}^m}|e_k\rangle\langle e_k|f(Vx)\rangle = \sum_{k=1}^{N} \sum_{x\in\mathbb{D}^m} h_k(x)\, |x\rangle_{\mathbb{D}^m}|e_k\rangle$$

by applying the identity operator $\sum_{k=1}^{N} |e_k\rangle\langle e_k|$ and putting

$$h_k(x) := s^{m/2} \sqrt{\rho_{1/s}(x)}\langle e_k|f(Vx)\rangle\,.$$

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

$$|\hat{\psi}\rangle = q^{m/2} \sum_{k=1}^{N} \sum_{y\in\hat{\mathbb{D}}^m} \mathcal{F}_{\mathbb{D}^m}\{h_k\}(y)\, |y\rangle_{\hat{\mathbb{D}}^m}|e_k\rangle\,,$$

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

$$\langle \hat{\psi} | (|y\rangle\langle y| \otimes \mathbb{I}) | \hat{\psi} \rangle = q^m \sum_{k=1}^{N} |\mathcal{F}_{\mathbb{D}^m} \{h_k\}(y)|^2 \,,$$

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

$$\mathcal{D}(C) = \sum_{y \in C} \frac{\langle \hat{\psi} | (|y\rangle\langle y| \otimes \mathbb{I}) | \hat{\psi} \rangle}{\langle \psi_\circ | \psi_\circ \rangle} = \frac{\sum_{k=1}^{N} \|1_C \cdot \mathcal{F}_{\mathbb{D}^m} \{h_k\}\|_{2,\hat{\mathbb{D}}^m}^2}{\frac{s^m}{q^m} \sum_{x \in \mathbb{D}^m} \rho_{1/s}(x)} \,. \tag{5}$$

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

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

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 funcion is then supported by $V\Lambda^*$ and takes on the values $\langle e_k | c_{\ell^*} \rangle$. Hence, by summing the squares over all k, 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 $f_V := f(V\cdot)$, with the effect that in the error term, Λ^* becomes $V\Lambda^*$ and $\mathrm{Lip}(f)$ becomes $V\,\mathrm{Lip}(f)$.

The error analysis (for $V = 1$) is divided into three parts. The first part consists of showing that the denominator from Equation (5) 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 $\hat{\mathbb{D}}_{\mathrm{rep}}^m = [q]_c^m \subset \mathbb{Z}^m$,

$$\|1_C \cdot \mathcal{F}_{\mathbb{D}^m} \{h_k\}\|_{2,\hat{\mathbb{D}}^m}^2 \gtrsim \sum_{\substack{\ell^* \in \Lambda^* \\ B_{\delta\lambda_1^*}(\ell^*) \subseteq C}} |\langle e_k | c_{\ell^*} \rangle|^2 \,. \tag{6}$$

We recall that $|c_{\ell^*}\rangle$ are the vectorial Fourier coefficients of f and $\mathcal{B}_{\delta\lambda_1^*}(\ell^*) = \mathcal{B}_{\delta\lambda_1^*}(\ell^*) \cap \mathbb{Z}^m$. This approximation (6) is divided into the following five steps:

$$\left\|1_C \mathcal{F}_{\mathbb{D}^m}\{h_k\}\right\|_{2,\hat{\mathbb{D}}^m}^2 \overset{(1)}{\approx} \left\|1_C \mathcal{F}_{\mathbb{D}^m}\left\{h_k\big|^{\mathbb{T}^m}\right\}\right\|_{2,\hat{\mathbb{D}}^m}^2 \overset{(2)}{\approx} \left\|1_C \mathcal{F}_{\mathbb{T}^m}\{h_k\big|^{\mathbb{T}^m}\}\right\|_{2,\mathbb{Z}^m}^2$$

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

It thus follows that

$$\mathcal{D}(C) \gtrsim \sum_{k=1}^{N} \sum_{\substack{\ell^* \in \Lambda^* \\ \mathcal{B}_{\delta\lambda_1^*}(\ell^*) \subseteq C}} |\langle e_k|c_{\ell^*}\rangle|^2 = \sum_{\substack{\ell^* \in \Lambda^* \\ \mathcal{B}_{\delta\lambda_1^*}(\ell^*) \subseteq C}} \langle c_{\ell^*}|c_{\ell^*}\rangle ,$$

and therefore, applied to $C := \mathcal{B}_{\delta\lambda_1^*}(S)$, that for any $S \subset \Lambda^*$ for which $\mathcal{B}_{\delta\lambda_1^*}(S) \subset [q]_c^m$, requirement (1) is satisfied.

The third part of the analysis is to show that (1) is satisfied also for $S \subset \Lambda^*$ for which $\mathcal{B}_{\delta\lambda_1^*}(S)$ is not fully contained in $[q]_c^m$. 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 $S_0 = \{\ell^* \in S \mid \mathcal{B}_{\delta\lambda_1^*}(\ell^*) \subseteq [q]_c^m\}$. 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 $q/s \geq \sqrt{m}$),

$$\frac{s^m}{q^m} \sum_{x \in \mathbb{D}^m} \rho_{1/s}(x) \leq \frac{s^m}{q^m} \cdot \rho_{1/s}\left(\frac{1}{q}\mathbb{Z}^m\right) \leq 1 + 2\beta_{q/s}^{(m)}. \tag{7}$$

Therefore,

$$\frac{\sum_{k=1}^{N} \left\|1_C \cdot \mathcal{F}_{\mathbb{D}^m}\{h_k\}\right\|_{2,\hat{\mathbb{D}}^m}^2}{\frac{s^m}{q^m} \sum_{x \in \mathbb{D}^m} \rho_{1/s}(x)} \geq \sum_{k=1}^{N} \left\|1_C \cdot \mathcal{F}_{\mathbb{D}^m}\{h_k\}\right\|_{2,\hat{\mathbb{D}}^m}^2 - \varepsilon_{\text{denom}} \tag{8}$$

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

Part 2 Recall that $h_k = s^{m/2} \cdot f_k \cdot \rho_{\sqrt{2}/s}$ is a function $h_k : \mathbb{R}^m \to \mathbb{C}$, where $f_k(x) = \langle e_k|f(x)\rangle$. In the following, by slightly abusing notation, we also understand h_k as a function $h_k : \mathbb{T}^m \to \mathbb{C}$ by considering the restriction of h_k to $\mathbb{T}_{\text{rep}}^m = [-\frac{1}{2}, \frac{1}{2})^m$. Similarly, we understand h_k as a function $h_k : \mathbb{D}^m \to \mathbb{C}$ by considering its restriction to $\mathbb{D}_{\text{rep}}^m = \mathbb{T}_{\text{rep}}^m \cap \frac{1}{q}\mathbb{Z}^m$.

Step 1. Observe that

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

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

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

$$\leq \frac{\|f_k\|_\infty^2 \, s^m}{q^m} \sum_{x \in \mathbb{D}^m} \left(\sum_{z \in \mathbb{Z}^m \backslash 0} \rho_{\sqrt{2}/s}(x+z) \right)^2 \leq 4 \|f_k\|_\infty^2 \, s^m (\beta_{s/\sqrt{8}}^{(m)})^2,$$

as $\rho_{\sqrt{2}/s}(\mathbb{Z}^m \backslash \{0\} + x) \leq 2\beta_{\frac{s}{2\sqrt{2}}}^{(m)}$, from Corollary 3. Therefore, noting that $\|f_k\|_\infty \leq 1$,

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

Step 2. Using Theorem 6 with $h_k|^{\mathbb{T}^m}$, one obtains

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

where $\varepsilon_{\text{lip}} = \frac{4\pi\sqrt{m} \operatorname{Lip}(h_k|^{\mathbb{T}^m})}{q}$. 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 $h_k|^{\mathbb{T}^m}$ can be obtained by taking the maximum value of the absolute value of the derivative.

$$\nabla \left(h_k|^{\mathbb{T}^m} \right) = s^{m/2} \sum_{z \in \mathbb{Z}^m} \left(\nabla f_k(x+z)\rho_{\sqrt{2}/s}(x+z) + f_k(x+z)\nabla\rho_{\sqrt{2}/s}(x+z) \right)$$

The norm of this expression is bounded by

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

$$\leq s^{m/2} \left(2\operatorname{Lip}(f) + 2\pi s^2 \right)$$

where we used $\|\nabla f_k\| \leq \operatorname{Lip}(f_k) \leq \operatorname{Lip}(f)$, $\|f_k\|_\infty \leq 1$, $\nabla\rho_{\sqrt{2}/s}(x) = \pi s^2 x \cdot \rho_{\sqrt{2}/s}(x)$, $\rho_{\sqrt{2}/s}(x+\mathbb{Z}^m) \leq 2$ and $\sum_{z \in \mathbb{Z}^m} \|x+z\| \, \rho_{\sqrt{2}/s}(x+z) \leq 2$. The second last inequality follows from $\rho_{\sqrt{2}/s}(x+\mathbb{Z}^m) \leq 1 + \rho_{\sqrt{2}/s}(\mathbb{Z}^m \backslash \{0\} + x) \leq 1 + 2\beta_{\frac{s}{2\sqrt{2}}}^{(m)} \leq 2$, 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

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

where we continue to abuse notation here by identifying $\mathcal{F}_{\mathbb{R}^m} \{ h_k \}$ 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 $\| 1_C \cdot \mathcal{F}_{\mathbb{D}^m} \{ h_k \} \|_{2,\hat{\mathbb{D}}^m} \leq 2$ (which follows from Equation (5) and Equation (7)), we conclude that

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

where we tacitly 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

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

where c_{ℓ^*} are the Fourier coefficients of f_k, i.e., $c_{\ell^*} = \mathcal{F}_{\mathbb{R}^m/\Lambda}[f_k](\ell^*)$. Therefore,

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

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

$$\begin{aligned}
\| 1_C \mathcal{F}_{\mathbb{R}^m} \{ h_k \} \|_{2,\mathbb{Z}^m}^2 &= \sum_{v^* \in \Lambda^*} |c_{v^*}|^2 \cdot \iota_C(v^*) \\
&\quad + \sum_{u^* \in \frac{1}{2}\Lambda^* \backslash 0} \rho_{s/2}(u^*) \sum_{v^* \in u^* + \Lambda^*} c_{v^*+u^*} \bar{c}_{v^*-u^*} \cdot \iota_C(v^*)
\end{aligned}$$

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,

$$\left| \sum_{v^* \in u^* + \Lambda^*} c_{v^*+u^*} \bar{c}_{v^*-u^*} \right| \leq \sum_{v^* \in \Lambda^*} |c_{v^*+2u^*}||c_{v^*}| \leq \sum_{v^* \in \Lambda^*} \left(|c_{v^*+2u^*}|^2 + |c_{v^*}|^2 \right) = 2 \| f_k \|_{2,\mathbb{R}^m/\Lambda}^2$$

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

$$\left| \|1_C \mathcal{F}_{\mathbb{R}^m}\{h_k\}\|^2_{2,\mathbb{Z}^m} - \sum_{\ell^* \in \Lambda^*} |c_{\ell^*}|^2 \iota_C(\ell^*) \right| \leq \varepsilon_{\mathrm{diag}} \cdot \|f_k\|^2_{2,\mathbb{R}^m/\Lambda} \ ,$$

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

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

$$\sum_{\ell^* \in \Lambda^*} |c_{\ell^*}|^2 \iota(\ell^*) \geq (1 - \varepsilon_{\mathrm{smooth}}) \sum_{\substack{\ell^* \in \Lambda^* \\ B_{V\delta}(V\ell^*) \subseteq C}} |c_{\ell^*}|^2 \ .$$

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

Finalizing By collecting all the error terms, we obtain that

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

$$\geq \sum_{\substack{\ell^* \in \Lambda^* \\ B_{\delta\lambda_1^*}(\ell^*) \subseteq C}} |\langle e_k | c_{\ell^*}\rangle|^2 - (\varepsilon_{\mathrm{smooth}} + \varepsilon_{\mathrm{diag}}) \|f_k\|^2_{2,\mathbb{R}^m/\Lambda} - 5(\varepsilon_{\mathrm{per}} + \varepsilon_{\mathrm{lip}})$$

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

$$\frac{2\delta\lambda_1^*}{s} \geq \sqrt{m} \quad \text{and} \quad \frac{s}{2\sqrt{2}} \geq \sqrt{m} \tag{9}$$

Summing over $k \in \{1, \ldots, N\}$ yields

$$\sum_{k=1}^{N} \|1_C \cdot \mathcal{F}_{\mathbb{D}^m}\{h_k\}\|^2_{2,\hat{\mathbb{D}}^m}$$

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

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

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 $\varepsilon_{\text{tail}} = \frac{4m \, \text{Lip}(f)^2}{\pi^2 q^2} \geq \sum_{\ell^* \in \Lambda^* \backslash [-\frac{q}{4}, \frac{q}{4}]^m} \langle c_{\ell^*} | c_{\ell^*} \rangle \geq \sum_{\ell^* \in S_1} \langle c_{\ell^*} | c_{\ell^*} \rangle$, (see Corollary 5), we have

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

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 $f_V := f(V \cdot)$. The hidden lattice of f_V is $\frac{1}{V}\Lambda$, which has $V\Lambda^*$ as its dual, and the Lipschitz constant of f_V is $V \, \text{Lip}(f)$. Thus, the requirements on the parameters (see Equation (9)) change to

$$\frac{2\delta V \lambda_1^*}{s} \geq \sqrt{m} \quad \text{and} \quad \frac{s}{2\sqrt{2}} \geq \sqrt{m}, \tag{10}$$

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

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

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

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

- $s \in O(\sqrt{nm\log(\eta^{-1})})$ so that $\varepsilon_{\text{per}} \cdot N \le \eta/6$, and $2\beta_{s/2}^{(m)} \le \eta/12$ in $\varepsilon_{\text{smooth}}$.
- $V \in O(\frac{\sqrt{m\log(\eta^{-1})s}}{\delta\lambda_1^*}) = O(\frac{m\sqrt{n}\log(\eta^{-1})}{\delta\lambda_1^*})$ so that $\varepsilon_{\text{smooth}}, \varepsilon_{\text{diag}} \le \eta/6$.
- $Q = \log(q) \in O(n + m\log(s) + \log(V) + \log(\mathrm{Lip}(f)) + \log(\eta^{-1}))$ so that $\varepsilon_{\text{lip}} \cdot N \le \eta/6$ and $\varepsilon_{\text{tail}} \le \eta/6$.

Unrolling the expression of $Q = \log(q)$ and recalling that the quantum Fourier transform requires a quadratic number of gates [20, Ch. 5], we obtain the main theorem.

Theorem 1. *Algorithm 1 solves the Dual Lattice Sampling Problem with oracle function f, error parameter η and relative distance parameter δ, using m calls to the Gaussian superposition subroutine (see Theorem 3), one quantum oracle call to f, $O(mQ)$ qubits, and $O(m^2Q^2)$ quantum gates, where*

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

6 From Sampling to Full Dual Lattice Recovery

We have so far focused on approximate sampling dual lattice points following weights $\|c_{\ell^*}\|^2$ for $\ell^* \in \Lambda^*$, regardless of how useful this distribution may be. Indeed, until now, it could be that the function $f : \mathbb{R}^m/\Lambda \to \mathcal{S}$ 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 $\|c_{\ell*}\|^2$ 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 $v_1, \ldots, v_{j-1} \in L$.

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

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' [4] 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 $k \in \mathbb{N}$,

$$\Pr\left[S > \frac{t + m + k}{1 - p - q}\right] \leq \Pr\left[X > \frac{t + m + k}{1 - p - q}\right] \leq \Pr\left[B < m + t\right] \qquad (11)$$

---

[4] 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 [9, Ch. 8, Ex. 17]. The last inequality follows from Chernoff's bound. Putting $k=(1+\alpha)(t+m)$ into Equation (11) yields the claim.

We conclude by relating the parameters (a,r,ϵ) of the HSP oracle (Definition 1) $f:\mathbb{R}^m/\Lambda\to\mathcal{S}$ 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 $f:\mathbb{R}^m\to\mathcal{S}$ 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 f. Then $\mathcal{D}_{\text{ideal}}$ is both $(\frac{1}{2}+\epsilon)$-evenly distributed and $(R,\frac{ma^2}{4\pi^2R^2})$-concentrated for any $R>0$.*

Proof. The distribution $\mathcal{D}_{\text{ideal}}$ being $(R,\frac{ma^2}{4\pi^2R^2})$-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 $f|^{\mathbb{R}^m/M}(x)=\frac{1}{|M/\Lambda|}\sum_{v\in M/\Lambda}f(x+v)$, the periodization of 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 5).

$$\sum_{v^* \in M^*} \langle c_{v^*} | c_{v^*} \rangle = \left\| f |^{\mathbb{R}^m/M} \right\|_{2,\mathbb{R}^m/M} = \frac{1}{\det M} \int_{x \in \mathbb{R}^m/M} \langle f |^{\mathbb{R}^m/M} | f |^{\mathbb{R}^m/M} \rangle dx,$$

$$= \frac{1}{|M/\Lambda|^2} \sum_{v,w \in M/\Lambda} \underbrace{\frac{1}{\det M} \int_{x \in \mathbb{R}^m/M} \langle f(x+v) | f(x+w) \rangle \, dx}_{I_{v,w}}$$

$$= \frac{1}{|M/\Lambda|^2} \sum_{\substack{v,w \in M/\Lambda \\ d_{\mathbb{R}^m/\Lambda}(v,w) < r}} I_{v,w} + \frac{1}{|M/\Lambda|^2} \sum_{\substack{v,w \in M/\Lambda \\ d_{\mathbb{R}^m/\Lambda}(v,w) \geq r}} I_{v,w}$$

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.

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. Brent, R.P.: Multiple-precision zero-finding methods and the complexity of elementary function evaluation. CoRR **abs/1004.3412** (2010), `http://arxiv.org/abs/1004.3412`
4. 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)
5. 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)
6. Deitmar, A., Echterhoff, S.: Principles of Harmonic Analysis. Springer Publishing Company, Incorporated, 2nd edn. (2016)
7. Ducas, L., Plançon, M., Wesolowski, B.: On the shortness of vectors to be found by the ideal-SVP quantum algorithm. Cryptology ePrint Archive, Report 2019/234 (2019), `https://eprint.iacr.org/2019/234`.

8. 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)

9. 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)

10. Grover, L., Rudolph, T.: Creating superpositions that correspond to efficiently integrable probability distributions. arXiv preprint quant-ph/0208112 (2002)

11. 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)

12. Hallgren, S.: Polynomial-time quantum algorithms for Pell's equation and the principal ideal problem. Journal of the ACM (JACM) **54**(1), 4 (2007)

13. Heinonen, J.: Lectures on Lipschitz analysis `http://www.math.jyu.fi/research/reports/rep100.pdf`

14. Kitaev, A., Webb, W.A.: Wavefunction preparation and resampling using a quantum computer. arXiv preprint arXiv:0801.0342 (2008)

15. Kuperberg, G.: A subexponential-time quantum algorithm for the dihedral hidden subgroup problem. SIAM Journal on Computing **35**(1), 170–188 (2005)

16. 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, `http://dx.doi.org/10.1137/S0097539705447360`

17. Miller, S.D., Stephens-Davidowitz, N.: Generalizations of Banaszczyk's transference theorems and tail bound. arXiv preprint arXiv:1802.05708 (2018)

18. 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)

19. National Institute of Standards and Technology: Post-quantum cryptography standardization (2017), `https://csrc.nist.gov/Projects/Post-Quantum-Cryptography/Post-Quantum-Cryptography-Standardization`

20. Nielsen, M.A., Chuang, I.L.: Quantum Computation and Quantum Information: 10th Anniversary Edition. Cambridge University Press, New York, NY, USA, 10th edn. (2011)

21. Pellet-Mary, A., Hanrot, G., Stehlé, D.: Approx-SVP in ideal lattices with preprocessing. Cryptology ePrint Archive, Report 2019/215 (2019), `https://eprint.iacr.org/2019/215`. To appear at EUROCRYPT 2019.

22. Regev, O.: Quantum computation and lattice problems. SIAM Journal on Computing **33**(3), 738–760 (2004)

23. Reiter, M., Arthur, S.: Fourier transform & solobev spaces (lecture notes) (2008), `https://www.mat.univie.ac.at/~stein/teaching/SoSem08/sobolev_fourier.pdf`

24. 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)

25. 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)

26. 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`

27. Werner, D.: Funktionalanalysis. Springer-Lehrbuch, Springer Berlin Heidelberg (2007), `https://books.google.nl/books?id=r1lIQ1ekLCcC`

28. 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, `https://doi.org/10.1007/BF01097255`

A Gaussian state

A.1 Reducing to the One-Dimensional Case

In this appendix, we are going to estimate the exact quantum complexity of obtaining an approximation (in the trace distance) of the state

$$\frac{1}{\sqrt{\rho_{1/s}(\mathbb{D}^m_{\mathrm{rep}})}} \sum_{x \in \mathbb{D}^m_{\mathrm{rep}}} \sqrt{\rho_{1/s}(x)}|x\rangle, \tag{12}$$

where $\mathbb{D}^m_{\mathrm{rep}} = \frac{1}{q}\mathbb{Z}^m \cap [0,1)^m$, and where $\rho_{1/s}$ is the Gaussian function (see Section 3.6).

An element $|x\rangle$ with $x = (x_1, \ldots, x_m) \in \mathbb{D}^m_{\mathrm{rep}}$ is represented as a tensor product $|x_1\rangle \otimes \ldots \otimes |x_m\rangle$. As the function $\sqrt{\rho_{1/s}(x)} = \rho_{\sqrt{2}/s}(x)$ can be written as a product of functions with separated variables as well, we obtain that Equation (12) equals

$$\bigotimes_{j=1}^{m} \frac{1}{\sqrt{\rho_{1/s}(\frac{1}{q}[q]_c)}} \sum_{a \in \frac{1}{q}[q]_c} \sqrt{\rho_{1/s}(a)}|a\rangle \tag{13}$$

where $\frac{1}{q}[q]_c = [-1/2, 1/2) \cap \frac{1}{q}\mathbb{Z}$. Therefore, the problem of approximating the state as in Equation (12) reduces to the one-dimensional case, i.e., as in Equation (13).

A.2 The Periodic and Non-Periodic Discrete Gaussian

To obtain a Gaussian superposition in one dimension, we follow Kitaev and Webb [14]. Their algorithm is an improvement of that of Grover and Rudolph [10].

Definition 6 (Gaussian function). *We denote by $\rho_{\mu,\sigma} : \mathbb{R} \to \mathbb{R}$ the function*

$$\rho_{\mu,\sigma}(x) = e^{-\pi \frac{(x-\mu)^2}{\sigma^2}}$$

The discrete Gaussian $\ddot{\rho}_{\mu,\sigma} : \mathbb{Z} \to \mathbb{R}$ is a rescaling of $\rho_{\mu,\sigma}$ such that $\sum_{j \in \mathbb{Z}} \ddot{\rho}_{\mu,\sigma}(j)^2 = 1$. Explicitly, $\ddot{\rho}_{\mu,\sigma}(j) = \frac{1}{\sqrt{\rho^2_{\mu,\sigma}(\mathbb{Z})}}\rho_{\mu,\sigma}(j)$.

Kitaev and Webb's algorithm actually doesn't compute a discrete Gaussian quantum state, but something very close; a *periodized* discrete Gaussian quantum state. This state has the advantage of having a more natural normalization and having a sum decomposition. These advantages lead to a slightly more efficient algorithm [14] computing the discrete Gaussian superposition, compared to the algorithm of Grover and Rudolph.

Definition 7 (Discrete Periodized Gaussian function). *We denote by $\xi_{\mu,\sigma,Q} : \mathbb{Z}/2^Q\mathbb{Z} \to \mathbb{R}_{>0}$ the function defined by the following rule*

$$\xi_{\mu,\sigma,Q}(x)^2 = \ddot{\rho}^2_{\mu,\sigma}(x + 2^Q\mathbb{Z}) = \sum_{t \in \mathbb{Z}} \ddot{\rho}_{\mu,\sigma}(x + 2^Q t)^2$$

The discrete periodized Gaussian state is then denoted by $|\xi_{\mu,\sigma,Q}\rangle := \sum_{j=0}^{2^Q-1} \xi_{\mu,\sigma,Q}(j)|j\rangle$. Note that the state is already normalized. As we already mentioned, the discrete periodized Gaussian state is very close to the discrete (non-periodized) Gaussian state. This is formalized in the next lemma.

Lemma 9. *Denote* $|\ddot{\rho}_{\mu,\sigma}\rangle = \frac{1}{\sqrt{c_{\mu,\sigma}}} \sum_{j=0}^{2^Q-1} \ddot{\rho}_{\mu,\sigma}(j) \cdot |j\rangle$, *with* $c_{\mu,\sigma} = \sum_{j=0}^{2^Q-1} \ddot{\rho}_{\mu,\sigma}(j)^2$. *Then, for* $\mu \in [0, 2^Q - 1]$ *and* $\sigma < 2^{Q-1}$,

$$D\left(|\ddot{\rho}_{\mu,\sigma}\rangle, |\xi_{\mu,\sigma,Q}\rangle\right) \leq \beta^{(1)}_{\frac{d_\mu}{\sigma}}$$

where $d_\mu := \min(\mu, 2^Q - \mu)$, D *is the trace distance, and* $\beta^{(1)}_{\frac{d_\mu}{\sigma}}$ *is Banaszczyk's function (see Section 3.6).*

Proof. Since $\xi_{\mu,\sigma,Q}(j) \geq \ddot{\rho}_{\mu,\sigma}(j)$, we have

$$\langle \xi_{\mu,\sigma,Q}|\ddot{\rho}_{\mu,\sigma}\rangle = \frac{1}{\sqrt{c_{\mu,\sigma}}} \sum_{j=0}^{2^Q-1} \ddot{\rho}_{\mu,\sigma}(j)\xi_{\mu,\sigma,Q}(j) \geq \frac{1}{\sqrt{c_{\mu,\sigma}}} \sum_{j=0}^{2^Q-1} \ddot{\rho}_{\mu,\sigma}(j)^2 = \sqrt{c_{\mu,\sigma}}$$

Since the trace distance between the pure states $|\ddot{\rho}_{\mu,\sigma}\rangle$ and $|\xi_{\mu,\sigma,Q}\rangle$ is equal to $\sqrt{1 - |\langle\xi_{\mu,\sigma,Q}|\ddot{\rho}_{\mu,\sigma}\rangle|^2}$ [20, §9.2], we obtain $D\left(|\ddot{\rho}_{\mu,\sigma}\rangle, |\xi_{\mu,\sigma,Q}\rangle\right) \leq \sqrt{1 - c_{\mu,\sigma}}$. As $1 - c_{\mu,\sigma} = \frac{\rho_{\mu,\sigma/\sqrt{2}}(\mathbb{Z}\backslash\{0,\ldots,2^Q-1\})}{\rho_{\mu,\sigma/\sqrt{2}}(\mathbb{Z})} \leq \beta^{(1)}_{\frac{\sqrt{2}d_\mu}{\sigma}}$ (see Lemma 2), and $\sqrt{\beta^{(1)}_z} \leq \beta^{(1)}_{z/\sqrt{2}}$, we obtain the claim.

Above lemma essentially states that whenever Q is relatively large, μ is relatively far away from the borders and σ is not too large, then the periodic discrete Gaussian and the (non-periodic) discrete Gaussian are very close in total variation distance.

A.3 Computing the Periodic Gaussian State

According to the last subsection, we can compute the state $|\xi_{\mu,\sigma,Q}\rangle$ instead of $|\ddot{\rho}_{\mu,\sigma}\rangle$, as they are close to each other for a suitable choice of parameters. As already mentioned, the quantum state $|\xi_{\mu,\sigma,Q}\rangle$ can be decomposed into a superposition that can be exploited algorithmically. The following lemma shows this decomposition.

Lemma 10 (Eq. 11 in [14]).

$$|\xi_{\mu,\sigma,Q}\rangle = |\xi_{\frac{\mu}{2},\frac{\sigma}{2},Q-1}\rangle \otimes \cos\alpha|0\rangle + |\xi_{\frac{\mu-1}{2},\frac{\sigma}{2},Q-1}\rangle \otimes \sin\alpha|1\rangle,$$

with $\alpha = \arccos\left(\sqrt{\rho_{\frac{\mu}{2},\frac{\sigma}{2\sqrt{2}}}(\mathbb{Z})\big/\rho_{\mu,\frac{\sigma}{\sqrt{2}}}(\mathbb{Z})}\right)$.

Proof. We have

$$|\xi_{\mu,\sigma,Q}\rangle = \sum_{j=0}^{2^Q-1} \xi_{\mu,\sigma,Q}(j)|j\rangle = \sum_{j=0}^{2^{Q-1}-1} \xi_{\mu,\sigma,Q}(2j)|j\rangle|0\rangle + \sum_{j=0}^{2^{Q-1}-1} \xi_{\mu,\sigma,Q}(2j+1)|j\rangle|1\rangle.$$

It remains to show that $\xi_{\mu,\sigma,Q}(2j) = \cos(\alpha) \cdot \xi_{\frac{\mu}{2},\frac{\sigma}{2},Q-1}(j)$ and $\xi_{\mu,\sigma,Q}(2j+1) = \sin(\alpha) \cdot \xi_{\frac{\mu-1}{2},\frac{\sigma}{2},Q-1}(j)$. We show the equality of the latter part, as the former part can be shown similarly.

$$\xi_{\mu,\sigma,Q}(2j+1)^2 = \ddot{\rho}_{\mu,\sigma}^2(2j+1+2^Q \cdot \mathbb{Z}) = \frac{\rho_{\mu-1,\sigma}^2(2j+2^Q \cdot \mathbb{Z})}{\rho_{\mu,\sigma}^2(\mathbb{Z})}$$

$$= \frac{\rho_{\frac{\mu-1}{2},\frac{\sigma}{2}}^2(j+2^{Q-1} \cdot \mathbb{Z})}{\rho_{\mu,\sigma}^2(\mathbb{Z})} = \frac{\rho_{\frac{\mu-1}{2},\frac{\sigma}{2}}^2(\mathbb{Z})}{\rho_{\mu,\sigma}^2(\mathbb{Z})} \cdot \xi_{\frac{\mu}{2},\frac{\sigma}{2},Q-1}(j)^2$$

Taking square roots, noting that $\rho_{\mu,\sigma}^2 = \rho_{\mu,\frac{\sigma}{\sqrt{2}}}$ and $\rho_{\frac{\mu-1}{2},\frac{\sigma}{2\sqrt{2}}}(\mathbb{Z}) + \rho_{\frac{\mu}{2},\frac{\sigma}{2\sqrt{2}}}(\mathbb{Z}) = \rho_{\mu,\frac{\sigma}{\sqrt{2}}}(\mathbb{Z})$ yields the result.

This lemma directly leads to an algorithm for computing (an approximation of) the state $|\xi_{\mu,\sigma,Q}\rangle$, which is spelled out in Algorithm 2.

Algorithm 2: Recursive algorithm preparing the periodic Gaussian state

Input : The parameters $\mu, \sigma \in \mathbb{R}_{>0}$, $k \in \mathbb{N}$ and $Q \in \mathbb{N}$.
Output: An approximation of the state $|\xi_{\mu,\sigma,Q}\rangle$

---

1 **Initial state**: $|0^k\rangle|\mu,\sigma,Q\rangle|0^Q\rangle$;

2 **Compute** α in the first register, yielding $|\alpha\rangle|\mu,\sigma,n\rangle|0^Q\rangle$, where
$\alpha = \arccos\left(\sqrt{\rho_{\frac{\mu}{2},\frac{\sigma}{2\sqrt{2}}}(\mathbb{Z})/\rho_{\mu,\frac{\sigma}{\sqrt{2}}}(\mathbb{Z})}\right)$;

3 **Apply the α-rotation on the last qubit**, yielding
$|\alpha\rangle|\mu,\sigma,Q\rangle|0^{Q-1}\rangle(\cos\alpha|0\rangle + \sin\alpha|1\rangle)$;

4 **Uncompute** α, yielding $|0^k\rangle|\mu,\sigma,Q\rangle|0^{Q-1}\rangle(\cos\alpha|0\rangle + \sin\alpha|1\rangle)$;

5 **Apply a parameter change, controlled by the last qubit** yielding
$\cos\alpha|0^k\rangle|\frac{\mu}{2},\frac{\sigma}{2},Q-1\rangle|0^{Q-1}\rangle|0\rangle + \sin\alpha|0^k\rangle|\frac{\mu-1}{2},\frac{\sigma}{2},Q-1\rangle|0^{Q-1}\rangle|1\rangle$;

6 **Apply quantum recursion (step 2 - 5) on all qubits except the last,**
whenever $Q > 1$, yielding
$\cos\alpha|0^k\rangle|\frac{\mu}{2},\frac{\sigma}{2},Q-1\rangle|\xi_{\frac{\mu}{2},\frac{\sigma}{2},Q-1}\rangle|0\rangle + \sin\alpha|0^k\rangle|\frac{\mu-1}{2},\frac{\sigma}{2},Q-1\rangle|\xi_{\frac{\mu-1}{2},\frac{\sigma}{2},Q-1}\rangle|1\rangle$;

7 **Un-apply the parameter change,** yielding
$|0^k\rangle|\mu,\sigma,Q\rangle\left(\cos\alpha|\xi_{\frac{\mu}{2},\frac{\sigma}{2},Q-1}\rangle|0\rangle + \sin\alpha|\xi_{\frac{\mu-1}{2},\frac{\sigma}{2},Q-1}\rangle|1\rangle\right) = |0^k\rangle|\mu,\sigma,Q\rangle|\xi_{\mu,\sigma,Q}\rangle$
;

A.4 Estimating the Complexity and Fidelity of Algorithm 2

We will discuss now how well Algorithm 2 approximates the state $|\xi_{\mu,\sigma,Q}\rangle$. For ease of analysis, we will assume (without loss of generality) that the operations on the parameters μ and σ (in step 5 and 7 of Algorithm 2) are exact. Then it turns out that the approximation error is primarily caused by the fact that the angle α in the algorithm is computed up to precision 2^{-k}. This is made precise in the following lemma.

Lemma 11. *Algorithm 2 with input parameters μ, σ, k, Q computes the periodic gaussian $|\xi_{\mu,\sigma,Q}\rangle$ within trace distance $Q2^{-k}$.*

Proof. The proof proceeds by induction on Q. We use the the identity $D(|\psi\rangle, |\phi\rangle)^2 + |\langle\psi|\phi\rangle|^2 = 1$ (see [20, §9.2]) multiple times throughout the proof. Let $\tilde{\alpha}$ be the k-bit approximation of α, and denote

$$|\tilde{\xi}_{\mu,\sigma,Q}\rangle = \cos\tilde{\alpha}|\tilde{\xi}_{\frac{\mu}{2},\frac{\sigma}{2},Q-1}\rangle|0\rangle + \sin\tilde{\alpha}|\tilde{\xi}_{\frac{\mu-1}{2},\frac{\sigma}{2},Q-1}\rangle|1\rangle$$

for the output of Algorithm 2 with input parameters μ, σ, k and Q. Then the inner product $\langle\tilde{\xi}_{\mu,\sigma,Q}|\xi_{\mu,\sigma,Q}\rangle$ equals

$$\cos(\alpha)\cos(\tilde{\alpha})\langle\tilde{\xi}_{\frac{\mu}{2},\frac{\sigma}{2},Q-1}|\xi_{\frac{\mu}{2},\frac{\sigma}{2},Q-1}\rangle + \sin(\alpha)\sin(\tilde{\alpha})\langle\tilde{\xi}_{\frac{\mu-1}{2},\frac{\sigma}{2},Q-1}|\xi_{\frac{\mu-1}{2},\frac{\sigma}{2},Q-1}\rangle.$$

By the induction hypothesis and the fact that the periodic Gaussian state has only positive amplitudes, above expression is at least

$$(\cos(\alpha)\cos(\tilde{\alpha})+\sin(\alpha)\sin(\tilde{\alpha}))\sqrt{1-(Q-1)^2 2^{-2k}} = \cos(\alpha-\tilde{\alpha})\sqrt{1-(Q-1)^2 2^{-2k}}$$

Therefore $D(|\xi_{\mu,\sigma,Q}\rangle, |\tilde{\xi}_{\mu,\sigma,Q}\rangle) = \sqrt{1-|\langle\xi_{\mu,\sigma,Q}|\tilde{\xi}_{\mu,\sigma,Q}\rangle|^2} \leq \sin(\alpha-\tilde{\alpha}) + (Q-1)2^{-k} \leq Q2^{-k}$.

Lemma 12. *Computing α with k-bits of precision in step 2 of Algorithm 2 can be done within $O(k^{3/2} \cdot \mathrm{polylog}(k))$ operations.*

Proof. Can be found in Appendix A.5.

Lemma 13. *Encoding μ and σ in $O(Q)$ bits, Algorithm 2 with input μ, σ, k and Q, runs on $O(Q+k)$ qubits and uses $O(Q \cdot k^{3/2} \cdot \mathrm{polylog}(k))$ quantum gates.*

Proof. The number of qubits used in Algorithm 2 equals $O(Q+k)$, because α is stored with k bits of precision and σ and μ with $O(Q)$ bits of precision.

For the number of gates, we go through the steps of Algorithm 2. Step 1 is the initial state. Step 2 (and step 4) computes α with precision 2^{-k}. By Lemma 12, we estimate that this costs $O(k^{3/2}\,\mathrm{polylog}(k))$ quantum gates. Step 3 applies the α-rotation, which costs k quantum gates, as a sequence of controlled $R_{\pi/2^j}$-gates. Step 5 (and step 7) is a parameter change, which costs a mere constant number of gates. Step 6 applies recursion, which, by induction, costs $O((Q-1) \cdot k^{3/2} \cdot \mathrm{polylog}(k))$ gates. Adding all up, gives a number of $O(Q \cdot k^{3/2} \cdot \mathrm{polylog}(k))$ gates.

Theorem 8. *For any positive integers Q, k and for any $\mu \in [0, 2^Q - 1]$ and any $\sigma > 1$, there exists an quantum algorithm that prepares the one-dimensional Gaussian state*

$$|\ddot{\rho}_{\mu,\sigma}\rangle = \sum_{j=0}^{2^Q-1} \ddot{\rho}_{\mu,\sigma}(j)|j\rangle \tag{14}$$

within trace distance $\beta^{(1)}_{\frac{d_\mu}{\sigma}} + Q \cdot 2^{-k}$ using $O(Q + k)$ qubits and $O(Q \cdot k^{3/2} \cdot \mathrm{polylog}(k))$ quantum gates.

Proof. The state in Equation (14) can be approximated by running Algorithm 2 with parameters μ, σ, Q, k. Combining Lemma 9 and Lemma 11 and using the fact that we can add trace distances [20, Ch. 9], this approximation is within trace distance $\beta^{(1)}_{\frac{d_\mu}{\sigma}} + Q2^{-k}$. For the running time, use Lemma 13 to conclude that Algorithm 2 with the mentioned parameters uses $O(Q+k)$ qubits and $O(Q \cdot k^{3/2})$ quantum gates, which proves the claim.

A.5 Proof of Lemma 12

Lemma 14. *The value $\rho_{\frac{\mu}{2}, \frac{\sigma}{2\sqrt{2}}}(\mathbb{Z})$ can be computed with relative precision 2^{-k} within time $O(k^{3/2} \, \mathrm{polylog}(k))$.*

Proof. We distinguish two cases.

- $\sigma < \sqrt{2}$. Then $\left| \rho_{\mu, \frac{\sigma}{\sqrt{2}}}(\mathbb{Z}) - \rho_{\lfloor \mu \rceil, \frac{\sigma}{\sqrt{2}}}(\{-h, \ldots, 0, \ldots h\}) \right| \leq \beta^{(1)}_{\sqrt{2}h/\sigma} \cdot \rho_{\mu, \frac{\sigma}{\sqrt{2}}}(\mathbb{Z})$, by Lemma 2.
- $\sigma > \sqrt{2}$. Applying the Poisson summation formula, we obtain $\rho_{\mu, \frac{\sigma}{\sqrt{2}}}(\mathbb{Z}) = \frac{\sigma}{\sqrt{2}} \sum_{t \in \mathbb{Z}} \rho_{0, \frac{\sqrt{2}}{\sigma}}(t) e^{-2\pi i t \mu}$. Therefore

$$\left| \rho_{\mu, \frac{\sigma}{\sqrt{2}}}(\mathbb{Z}) - \frac{\sigma}{\sqrt{2}} \sum_{t \in \{-h, \ldots, 0, \ldots h\}} \rho_{\frac{\sqrt{2}}{\sigma}}(t) e^{-2\pi i t \mu} \right| \leq \frac{\sigma}{\sqrt{2}} \beta^{(1)}_{\sigma h/\sqrt{2}} \cdot \rho_{0, \sqrt{2}/\sigma}(\mathbb{Z})$$

  which is bounded by $\beta^{(1)}_{\sigma h/\sqrt{2}} \cdot \rho_{0, \sigma/\sqrt{2}}(\mathbb{Z}) \leq 2\beta^{(1)}_{\sigma h/\sqrt{2}} \cdot \rho_{\mu, \frac{\sigma}{\sqrt{2}}}(\mathbb{Z})$, by the Poisson summation formula and by smoothing arguments (see Lemma 4), as $\rho_{\mu, \sigma/\sqrt{2}}(\mathbb{Z}) \geq (1 - 2\beta^{(1)}_{s/\sqrt{2}})\rho_{0, \sigma/\sqrt{2}} \geq \frac{1}{2}\rho_{0, \sigma/\sqrt{2}}$.

So the relative error is at most $2\beta^{(1)}_h \leq e^{-(h-1)^2}$ for $h > 2$. Therefore, choosing $h = k^{1/2} + 1$ is enough to compute $\rho_{\frac{\mu}{2}, \frac{\sigma}{2\sqrt{2}}}(\mathbb{Z})$ with relative error 2^{-k}. Because evaluating an exponential function takes $O(k \cdot \mathrm{polylog}(k))$ time [3], we arrive at the claim.

Lemma 15. *The fraction $\rho_{\frac{\mu}{2}, \frac{\sigma}{2\sqrt{2}}}(\mathbb{Z}) / \rho_{\mu, \frac{\sigma}{\sqrt{2}}}(\mathbb{Z})$ can be computed with precision 2^{-k} within time $O(k^{3/2} \cdot \mathrm{polylog}(k))$.*

Proof. Denote $a = \rho_{\frac{\mu}{2}, \frac{\sigma}{2\sqrt{2}}}(\mathbb{Z})$ and $b = \rho_{\mu, \frac{\sigma}{\sqrt{2}}}(\mathbb{Z})$. Suppose we have relative errors $|\tilde{a} - a| \leq 2^{-k}a/2 \leq 2^{-k}b/2$, $|\tilde{b} - b| \leq 2^{-k}b/2$ and $\tilde{a}/\tilde{b} < 1$, then $\left|\frac{\tilde{a}}{\tilde{b}} - \frac{a}{b}\right| \leq \frac{|\tilde{b}(a-\tilde{a}) - \tilde{a}(b-\tilde{b})|}{b\tilde{b}} \leq \frac{|a-\tilde{a}|}{b} + \frac{|b-\tilde{b}|}{b} \leq 2^{-k}$. By Lemma 14, we see that both a and b can be computed within relative precision $2^{-k}/2$ within time $O(k^{3/2}\operatorname{polylog}(k))$. Therefore, the fraction a/b can be computed with absolute precision 2^{-k} within time $O(k^{3/2}\operatorname{polylog}(k))$.

Lemma 16. *For $x \in [0, 1 - \epsilon]$ and $\epsilon < \frac{3}{4}$, we have*

$$|\arccos(\sqrt{x+\epsilon}) - \arccos(\sqrt{x})| \leq 8\sqrt{\epsilon}$$

Proof. The derivative of $\arccos(\sqrt{t})$ equals $w(t) = -\frac{2}{\sqrt{(1-t)t}}$. Therefore

$$|\arccos(\sqrt{x+\epsilon}) - \arccos(\sqrt{x})| \leq \left|\int_x^{x+\epsilon} w(t)dt\right| \leq \int_x^{x+\epsilon} |w(t)|dt \leq \int_0^{\epsilon} |w(t)|dt.$$

The last inequality follows from the fact that $w(t)$ is both strictly decreasing on $[0, 1/2]$ and symmetric around $t = 1/2$. The claim then follows from the bound $\int_0^{\epsilon} |w(t)|dt = \int_0^{\epsilon} \frac{2}{\sqrt{(1-x)x}} \leq 4\int_0^{\epsilon} \frac{dt}{\sqrt{t}} = 8\sqrt{\epsilon}$.

By combining Lemma 15 and Lemma 16, we obtain that the expression $\arccos\sqrt{\rho_{\frac{\mu}{2}, \frac{\sigma}{2\sqrt{2}}}(\mathbb{Z})\big/\rho_{\mu, \frac{\sigma}{\sqrt{2}}}(\mathbb{Z})}$ can be approximated with k bits of precision within $O(k^{3/2} \cdot \operatorname{polylog}(k))$ time, which proves Lemma 12.