Quantum Computation and Lattice Problems

Oded Regev *

August 27, 2003

Abstract

We present the first explicit connection between quantum computation and lattice problems. Namely, our main result is a solution to the Unique Shortest Vector Problem (SVP) under the assumption that there exists an algorithm that solves the hidden subgroup problem on the dihedral group by coset sampling. Additionally, we present an approach to solving the hidden subgroup problem on the dihedral group by using an average case subset sum routine.

1 Introduction

Quantum computation is a computation model based on quantum physics. Assuming that the laws of nature as we know them are true, this might allow us to build computers that are able to perform tasks that classical computers cannot perform in any reasonable time. One task which quantum algorithms are known to perform much better than classical algorithm is that of factoring large integers. The importance of this problem stems from its ubiquitous use in cryptographic applications. While there are no known polynomial time classical algorithms for this problem, a groundbreaking result of Shor from 1994 [24] showed a polynomial time quantum algorithm for factoring integers. In the same paper, Shor showed an algorithm for finding the discrete log. However, despite enormous effort, we have only a few other problems for which quantum algorithms provide an exponential speedup (e.g., [11, 26]). Other notable quantum algorithms such as Deutsch and Jozsa's algorithm [5] and Simon's algorithm [25] operate in the black box model. Grover's algorithm [10] provides a square root speedup over classical algorithms.

The current search for new quantum algorithms concentrates on problems which are not known to be NP-hard. These include the graph isomorphism problem and lattice problems. In this paper we are interested in lattice problems or specifically, the unique shortest vector problem (SVP). A lattice is a set of all integral linear combinations of a set of n linearly independent vectors in $\mathbb{R}^n$. This set of n vectors is known as a basis of the lattice. In the SVP we are interested in finding the shortest nonzero vector in a lattice. In the $f(n)$-unique-SVP we are given the additional promise that the shortest vector is shorter by a factor of at least $f(n)$ from all other non parallel vectors. This problem also has important applications in cryptography. Namely, Ajtai and Dwork's cryptosystem [2] and the recent cryptosystem by Regev [22] are based on the hardness of this lattice problem.

A central problem in quantum computation is the hidden subgroup problem (HSP). Here, we are given a black box that computes a function on elements of a group G. The function is known to

*EECS Department, UC Berkeley, Berkeley, CA 94720. Email: odedr@cs.berkeley.edu. Most of this work was done while the author was at the Institute for Advanced Study, Princeton, NJ. Work supported by the Army Research Office grant DAAD19-03-1-0082 and NSF grant CCR-9987845.

be constant and distinct on left cosets of a subgroup $H \leqslant G$ and our goal is to find H. Interestingly, almost all known quantum algorithms which run super-polynomially faster than classical algorithms solve special cases of the HSP on Abelian groups. Also, it is known that solving the HSP on the symmetric group leads to a solution to graph isomorphism [14]. This motivated research into possible extensions of the HSP to noncommutative groups (see, e.g., [8, 12, 23, 7]). However, prior to this paper the HSP on groups other than the symmetric group and Abelian groups had no known applications.

In this paper we will be interested in the HSP on the dihedral group. The dihedral group of order $2N$, denoted D_N, is the group of symmetries of an N-sided regular polygon. It is isomorphic to the abstract group generated by the element ρ of order n and the element τ of order 2 subject to the relation $\rho\tau = \tau\rho^{-1}$. Although the dihedral group has a much simpler structure than the symmetric group, no efficient solution to the HSP on the dihedral group is known. Ettinger and Høyer [6] showed that one can obtain sufficient statistical *information* about the hidden subgroup with only a polynomial number of queries. However, there is no efficient algorithm that solves the HSP using this information. Currently, the best known algorithm is due to Kuperberg [17] and runs in subexponential time $2^{O(\sqrt{\log N})}$.

The following is the main theorem of this paper. The dihedral coset problem is described in the following paragraph.

Theorem 1.1 *If there exists a solution to the dihedral coset problem with failure parameter* f *then there exists a quantum algorithm that solves the* $\Theta(n^{\frac{1}{2}+2\mathsf{f}})$*-unique-SVP.*

The input to the dihedral coset problem (DCP) is a tensor product of a polynomial number of registers. Each register is in the state $|0,x\rangle + |1,(x+d) \bmod N\rangle$ for some arbitrary $x \in \{0,\ldots,N-1\}$ and d is the same for all registers. These can also be thought of as cosets of the subgroup $\{(0,0),(1,d)\}$ in D_N. Our goal is to find the value d. In addition, we say that the DCP has a failure parameter f if each of the registers with probability at most $\frac{1}{(\log N)^{\mathsf{f}}}$ is in the state $|b,x\rangle$ for arbitrary b,x instead of a coset state. We note that any algorithm that solves the dihedral HSP by sampling cosets also solves the DCP for some failure parameter f. The reason is that since the algorithm samples only a polynomial number of cosets, we can take f to be large enough such that with high probability all the registers are coset states. This is summarized in the following corollary.

Corollary 1.2 *If there exists a solution to the dihedral HSP that samples cosets (e.g., any solution using the 'standard method') then there exists a quantum algorithm that solves* $poly(n)$*-unique-SVP.*

The following is the second result of this paper. While still not an efficient solution, it shows a new way to approach the dihedral HSP. In the subset sum problem we are given two integers t, N and a set of numbers. We are asked to find a subset of the numbers that sums to t modulo N. A legal input is an input for which such a subset exists (a formal definition appears in Section 4) and we are interested in algorithms that solve a non-negligible part of the inputs:

Theorem 1.3 *If there exists an algorithm S that solves* $\frac{1}{poly(\log N)}$ *of the legal subset sum inputs with parameter N then there exists a solution to the DCP with failure parameter* $\mathsf{f} = 1$.

As shown in [6], the dihedral HSP can be reduced to the case where the subgroup is of the form $\{(0,0),(1,d)\}$. Then, by sampling cosets we obtain states of the form $|0,x\rangle + |1,(x+d) \bmod N\rangle$ with no error. Hence,

Corollary 1.4 *If there exists an algorithm S that solves* $\frac{1}{poly(\log N)}$ *of the legal subset sum inputs with parameter N then there exists a solution to the dihedral HSP.*

Finally, as a curiosity, let us comment that by combining the two previous theorems one can obtain the following corollary:

Corollary 1.5 *If there exists an algorithm that solves $\frac{1}{poly(\log N)}$ of the legal subset sum inputs with parameter N then there exists a quantum algorithm for the $\Theta(n^{2.5})$-unique-SVP.*

This result can be described as a worst case to average case quantum reduction. Such reductions are already known in the classical case [1, 3, 4, 19, 22]. The exponent 2.5 in our reduction is better than the one in [1, 3, 4, 19]. However, the reduction in [22], which appeared after the original publication of the current paper, further improves the exponent to 1.5 and hence subsumes our reduction. In addition, unlike the classical reductions, our subset sum problems have a density of one, i.e., the size of the input set is very close to $\log N$. Therefore, some cryptographic applications such as the one by Impagliazzo and Naor [13] cannot be used.

Intuitive overview

Before proceeding to the main part of the paper, we describe our methods in a somewhat intuitive way. First, let us describe the methods used in solving the unique-SVP. Recall that our solution is based on a solution to the DCP. We begin by showing how such a solution can be used to solve a slightly different problem which we call the two point problem. Instead of a superposition of two numbers with a fixed difference, our input consists of registers in a superposition of two n-dimensional vectors with a fixed difference. Then, the idea is to create an input to the two point problem in the following way. Start by creating a superposition of many lattice points and collapse the state to just two lattice points whose difference is the shortest vector. Repeating this procedure creates an input to the two point problem whose solution is the shortest vector.

Collapsing the state is performed by partitioning the space into cubes. Assume the partition has the property that in each cube there are exactly two lattice points whose difference is the shortest vector. Then, we compute the cube in which each point is located and measure the result. The state collapses to a superposition of just the two points inside the cube we measured. The important thing is to make sure that exactly two points are located in each cube. First, in order to make sure that the cubes are not aligned with the lattice, we randomly translate them. The length of the cubes is proportional to the length of the shortest vector. Although the exact length of the shortest vector is unknown, we can try several estimates until we find the right value. Since the lattice has a unique shortest vector, all other nonparallel vectors are considerably longer and do not fit inside a cube. Therefore we know that the difference between any two points inside the same cube is a multiple of the shortest vector. Still, this is not good enough since instead of two points inside each box we are likely to have more points aligned along the shortest vector. Hence, we space out the lattice: instead of creating a superposition of all the lattice points we create a superposition of a subset of the points. The set of points created by this technique has the property that along the direction of the shortest vector there are pairs of points whose difference is the shortest vector and the distance between two such pairs is much larger than the shortest vector. As before, this can be done without knowing the shortest vector by trying several possibilities.

The second part of the paper describes a solution to the DCP with failure parameter 1 which uses a solution to the average case subset sum problem. Recall that we are given registers of the form $|0,x\rangle + |1,(x+d) \bmod N\rangle$ where $x \in \{0,\ldots,N-1\}$ is arbitrary and we wish to find $d \in \{0,\ldots,N-1\}$. Consider one such register. We begin by applying the Fourier transform to the second part of the register (the one holding x and $x+d$) and then measuring it. If a is the value we measured, the state collapses to a combination of the basis states $|0\rangle$ and $|1\rangle$ such that their phase difference is $2\pi\frac{ad}{N}$. If we were lucky enough to measure $a = 1$, then the phase difference is $2\pi\frac{d}{N}$

and by measuring this phase difference we can obtain an estimation on d. This, however, happens with exponentially small probability. Since the phase is modulo 2π, extracting the value d is much harder when a is larger. Instead, we perform the same process on r registers and let $a_1, \dots, a_r$ be the values we measure. The resulting tensor state includes a combination of all 2^r different $0, 1$ sequences. The phase of each sequence can be described as follows. By ignoring a fixed phase, we can assume that the phase of the sequence $00\dots0$ is 0. Then, the phase of the sequence $100\dots0$ is $2\pi\frac{a_1 d}{N}$ and in general, the phase of the sequence $\alpha_1\alpha_2\dots\alpha_r$ is $2\pi\frac{d}{N}$ multiplied by the sum of the values a_i for which $\alpha_i = 1$. This indicates that we should try to measure the phase difference of two sequences whose sums differ by 1. However, although we can estimate the phase difference of one qubit, estimating the phase difference of two arbitrary sequences is not possible.

We proceed by choosing r to be very close to $\log N$. This creates a situation in which for almost every $t \in \{0, \dots, N-1\}$ there is a subset whose sum modulo N is t and in addition, there are not too many subsets that sum to the same t modulo N. Assume for simplicity that every t has exactly one subset that sums to t modulo N. We calculate for each sequence the value $\lfloor \frac{t}{2} \rfloor$ where t is its sum. After measuring the result, say s, we know that the state is a superposition of two sequences: one that sums to $2s$ and one that sums to $2s+1$. Notice that since $a_1, \dots, a_r$ are uniformly chosen between $\{0, \dots, N-1\}$ we can use them as an input to the subset sum algorithm. The key observation here is that the subset sum algorithm provides the reverse mapping, i.e., from a value t to a subset that sums to t. So, from s we can find the sequence α_1 that sums to $2s$ and the sequence α_2 that sums to $2s+1$. Since we know that the state is a superposition of $|\alpha_1\rangle$ and $|\alpha_2\rangle$ we can use a unitary transformation that transforms $|\alpha_1\rangle$ to $|0\rangle$ and $|\alpha_2\rangle$ to $|1\rangle$. Now, since the two states differ in one qubit, we can easily measure the phase difference and obtain an estimate on d. This almost completes the description of the DCP algorithm. The estimate on d is only polynomially accurate but in order to find d we need exponential accuracy. Hence, we repeat the same process with pairs whose difference is higher. So, instead of choosing pairs of difference 1 we choose pairs of difference 2 to get an estimate on $2d$, then 4 to get an estimate on $4d$ and so on[1].

Outline

The next section contains some notations that are used in this paper. The two main sections of this paper are independent. In Section 3 we prove Theorem 1.1 and Section 4 contains the proof of Theorem 1.3.

2 Preliminaries

We denote the imaginary unit by $\imath$ and use the notation $e(x) = e^{2\pi\imath x}$. Occasionally, we omit the normalization of quantum states. We use the term n-ball to refer to the n-dimensional solid body and the term sphere to refer to its surface. We denote the set $\{1, \dots, n\}$ by $[n]$. All logarithms are of base 2 unless otherwise specified. We use δ_{ij} to denote the Kronecker delta, i.e., 1 if $i = j$ and 0 otherwise. A sequence $\bar{\alpha} \in \{0,1\}^r$ is identified with the set $\{i \mid \alpha_i = 1\}$. Several constants appear in our proofs. To make it easier to follow, we denote constants with a subscript that is somewhat related to their meaning. Specifically, in Section 3, c_{cub} is related to the cubes that partition the space, c_{bal} is related to the radius of the balls, and c_{unq} appears in the guarantee of the unique shortest vector. Also, in Section 4 we use c_{r} in the definition of the parameter r, c_{s} in our assumptions on the subset sum subroutine and c_{m} when we prove the existence of matchings.

[1] This description is very similar to the method of exponentially accurate phase estimation used in Kitaev's algorithm [16]. Actually, our case is slightly more difficult because we cannot measure all the multiples 2^i. Nevertheless, we can measure enough multiples of the phase to guarantee exponential accuracy.

The following is the formal definition of the DCP:

Definition 2.1 *The input to the DCP with failure parameter* f *consists of poly*$(\log N)$ *registers. Each register is with probability at least* $1 - \frac{1}{(\log N)^{\mathsf{f}}}$ *in the state* $\frac{1}{\sqrt{2}}(|0, x\rangle + |1, (x + d) \bmod N\rangle)$ *on* $1 + \lceil \log N \rceil$ *qubits where* $x \in \{0, \ldots, N-1\}$ *is arbitrary and* d *is fixed. Otherwise, with probability at most* $\frac{1}{(\log N)^{\mathsf{f}}}$*, its state is* $|b, x\rangle$ *where* $b \in \{0, 1\}$ *and* $x \in \{0, \ldots, N-1\}$ *are arbitrary. We call such a register a 'bad' register. We say that an algorithm solves the DCP if it outputs* d *with probability poly*$(\frac{1}{\log N})$ *and time poly*$(\log N)$*.*

3 A Quantum Algorithm for unique-SVP

In this section we prove Theorem 1.1. We begin by showing a simple reduction from the two point problem to the DCP in Section 3.1. We then prove a weaker version of Theorem 1.1 with $\Theta(n^{1+2\mathsf{f}})$ instead of $\Theta(n^{\frac{1}{2}+2\mathsf{f}})$ in Section 3.2. We complete the proof of Theorem 1.1 in Section 3.3. Throughout this section, we use a failure parameter $\mathsf{f} > 0$ in order to make our results more general. The reader might find it easier to take $\mathsf{f} = 1$.

3.1 The Two Point Problem

Definition 3.1 *The input to the two point problem with failure parameter* f *consists of poly*$(n \log M)$ *registers. Each register is with probability at least* $1 - \frac{1}{(n \log(2M))^{\mathsf{f}}}$ *in the state* $\frac{1}{\sqrt{2}}(|0, \bar{a}\rangle + |1, \bar{a}'\rangle)$ *on* $1 + n\lceil \log M \rceil$ *qubits where* $\bar{a}, \bar{a}' \in \{0, \ldots, M-1\}^n$ *are arbitrary such that* $\bar{a}' - \bar{a}$ *is fixed. Otherwise, with probability at most* $\frac{1}{(n \log(2M))^{\mathsf{f}}}$*, its state is* $|b, \bar{a}\rangle$ *where* $b \in \{0, 1\}$ *and* $\bar{a} \in \{0, \ldots, M-1\}^n$ *are arbitrary. We say that an algorithm solves the two point problem if it outputs* $\bar{a}' - \bar{a}$ *with probability poly*$(\frac{1}{n \log M})$ *and time poly*$(n \log M)$*.*

Lemma 3.2 *If there exists an algorithm that solves the DCP with failure parameter* f *then there is an algorithm that solves the two point problem with failure parameter* f*.*

Proof: Consider the following mapping from $\{0, \ldots, M-1\}^n$ to $\{0, \ldots, (2M)^n - 1\}$:

$$f(a_1, \ldots, a_n) = a_1 + a_2 \cdot 2M + \ldots + a_n (2M)^{n-1}.$$

Given an input to the two point problem, we create an input to the DCP by using the above mapping on the last $n\lceil \log M \rceil$ qubits of each register. Hence, each register is with probability at least $1 - \frac{1}{(n(\log 2M))^{\mathsf{f}}}$ in state

$$\frac{1}{\sqrt{2}}(|0, f(\bar{a})\rangle + |1, f(\bar{a}')\rangle).$$

The difference $f(\bar{a}') - f(\bar{a})$ is $(a'_1 - a_1) + (a'_2 - a_2) \cdot 2M + \ldots + (a'_n - a_n)(2M)^{n-1}$ and is therefore fixed. Otherwise, with probability at most $\frac{1}{(n(\log 2M))^{\mathsf{f}}}$ the register is in the state $|b, f(\bar{a})\rangle$ for arbitrary $b, \bar{a}$. This is a valid input to the DCP with $N = (2M)^n$ since the probability of a bad register is at most $\frac{1}{(n \log(2M))^{\mathsf{f}}} = \frac{1}{(\log N)^{\mathsf{f}}}$.

Using the DCP algorithm with the above input we obtain the difference $b_1 + b_2 \cdot 2M + \ldots + b_n (2M)^{n-1}$ where $b_i = a'_i - a_i$. In order to extract the b_i's we add $M + M \cdot 2M + M(2M)^2 + \ldots + M(2M)^{n-1}$. Extracting b_i from $(b_1 + M) + (b_2 + M) \cdot 2M + \ldots + (b_n + M)(2M)^{n-1}$ is possible since each $b_i + M$ is an integer in the range 1 to $2M - 1$. The solution to the two point problem is the vector $(b_1, \ldots, b_n)$. ∎

3.2 A Weaker Algorithm

We recall several facts about an LLL-reduced basis. Such a basis can be found for any lattice by using a polynomial time algorithm [18]. Given a basis $\langle \bar{b}_1, \ldots, \bar{b}_n \rangle$, let $\langle \bar{b}_1^*, \ldots, \bar{b}_n^* \rangle$ be its Gram-Schmidt orthogonalization. That is, $\bar{b}_i^*$ is the component of $\bar{b}_i$ orthogonal to the subspace spanned by $\bar{b}_1, \ldots, \bar{b}_{i-1}$. An LLL reduced basis $\langle \bar{b}_1, \ldots, \bar{b}_n \rangle$ satisfies that $\|\bar{b}_i^*\| \leq \sqrt{2}\|\bar{b}_{i+1}^*\|$ and that for $i > j$, $|\langle \bar{b}_i, \bar{b}_j^* \rangle| \leq \frac{1}{2}\|\bar{b}_j^*\|^2$. In addition, recall that $\min_i \|\bar{b}_i^*\|$ is a lower bound on the length of the shortest vector. Since $\bar{b}_1^* = \bar{b}_1$ and $\|\bar{b}_1^*\| \leq 2^{(i-1)/2}\|\bar{b}_i^*\|$ we get that the vector $\bar{b}_1$ is at most $2^{(n-1)/2}$ times longer than the shortest vector. Consider the representation of the LLL basis in the orthonormal basis $\langle \frac{\bar{b}_1^*}{\|\bar{b}_1^*\|}, \ldots, \frac{\bar{b}_n^*}{\|\bar{b}_n^*\|} \rangle$. The vector $\bar{b}_i$ can be written as $(b_{i1}, b_{i2}, \ldots, b_{ii}, 0, \ldots, 0)$. Notice that $b_{ii} = \|\bar{b}_i^*\|$ and that $|b_{ij}| \leq \frac{1}{2}\|\bar{b}_j^*\|$ for every $i > j$. In the following, $\bar{u}$ denotes the shortest vector.

Lemma 3.3 *Consider the representation of the shortest vector $\bar{u}$ in the LLL-reduced lattice basis $\bar{u} = \sum_{i=1}^n u_i \bar{b}_i$. Then, $|u_i| \leq 2^{2n}$ for $i \in [n]$.*

Proof: Changing to the orthonormal basis, $\bar{u} = \sum_{i=1}^n u_i \bar{b}_i = \sum_{i=1}^n (\sum_{j=i}^n u_j b_{j,i}) \frac{\bar{b}_i^*}{\|\bar{b}_i^*\|}$. In addition, we know that $\|\bar{b}_i^*\| \geq 2^{-(i-1)/2}\|\bar{b}_1^*\| \geq 2^{-n}\|\bar{u}\|$. Hence, $|\sum_{j=i}^n u_j b_{j,i}| \leq 2^n \|\bar{b}_i^*\|$ for every $i \in [n]$. By taking $i = n$ we get that $|u_n|$ is at most 2^n. We continue inductively and show that $|u_k| \leq 2^{2n-k}$. Assume that the claim holds for $u_{k+1}, \ldots, u_n$. Then, $|\sum_{j=k+1}^n u_j b_{j,k}| \leq \frac{1}{2}|\sum_{j=k+1}^n u_j| \|\bar{b}_k^*\| \leq \frac{1}{2}(\sum_{j=k+1}^n 2^{2n-j})\|\bar{b}_k^*\| \leq \frac{1}{2} \cdot 2^{2n-k}\|\bar{b}_k^*\|$. By the triangle inequality, $|u_k b_{k,k}| \leq |\sum_{j=k+1}^n u_j b_{j,k}| + |\sum_{j=k}^n u_j b_{j,k}| \leq (\frac{1}{2}2^{2n-k} + 2^n)\|\bar{b}_k^*\| \leq 2^{2n-k}\|\bar{b}_k^*\|$ and the proof is completed. ■

Let $p > n^{2+2\mathsf{f}}$ be any fixed prime. The following is the main lemma of this section:

Lemma 3.4 *For any $\mathsf{f} > 0$ let $\bar{u} = \sum_{i=1}^n u_i \bar{b}_i$ be the shortest lattice vector in a $(c_{\mathsf{unq}} n^{1+2\mathsf{f}})$-unique lattice where $c_{\mathsf{unq}} > 0$ is a constant. If there exists a solution to the two point problem with failure parameter f then there exists a quantum algorithm that given this lattice and three integers l, m, i_0 returns $(u_1, \ldots, u_{i_0-1}, \frac{u_{i_0}-m}{p}, u_{i_0+1}, \ldots, u_n)$ with probability $1/poly(n)$ if the following conditions hold: $\|\bar{u}\| \leq l \leq 2\|\bar{u}\|$, $u_{i_0} \equiv m \pmod{p}$ and $1 \leq m \leq p-1$.*

We first show how this lemma implies Theorem 1.1 with $\Theta(n^{1+2\mathsf{f}})$ by describing the SVP algorithm. According to Lemma 3.2 and the assumption of the theorem, there exists a solution to the two point problem with failure parameter f. Hence, Lemma 3.4 implies that there exists an algorithm that given the right values of l, m, i_0 outputs $(u_1, \ldots, u_{i_0-1}, \frac{u_{i_0}-m}{p}, u_{i_0+1}, \ldots, u_n)$. The value l is an estimate of the length of the shortest vector $\bar{u}$. Because the LLL algorithm gives a $2^{(n-1)/2}$-approximation to the length of the shortest vector, one of $(n-1)/2$ different values of l is as required. In addition, since $\bar{u}$ is the shortest vector, $\bar{u}/p$ cannot be a lattice vector and therefore there exists an i_0 such that $u_{i_0} \not\equiv 0 \pmod{p}$. Hence, there are only $O(pn^2)$ possible values for l, m and i_0. With each of these values the SVP algorithm calls the algorithm of Lemma 3.4 a polynomial number of times. With high probability in one of these calls the algorithm returns the vector $(u_1, \ldots, u_{i_0-1}, \frac{u_{i_0}-m}{p}, u_{i_0+1}, \ldots, u_n)$ from which $\bar{u}$ can be extracted. The results of the other calls can be easily discarded because they are either longer lattice vectors or non-lattice vectors.

Proof: (of Lemma 3.4) We start by applying the LLL algorithm to the unique lattice in order to create a reduced basis. Denote the resulting basis by $\langle \bar{b}_1, \ldots, \bar{b}_n \rangle$. Let $\langle \bar{e}_1, \ldots, \bar{e}_n \rangle$ be the standard orthonormal basis of $\mathbb{R}^n$.

Let $w_1,\ldots,w_n$ be n real values in $[0,1)$ and let $M=2^{4n}$. Assume without loss of generality that $i_0=1$. The function f is defined as $f(t,\bar{a})=(a_1p+tm)\bar{b}_1+\sum_{i=2}^n a_i\bar{b}_i$ where $t\in\{0,1\}$ and $\bar{a}=(a_1,\ldots,a_n)\in\mathcal{A}=\{0,\ldots,M-1\}^n$. It maps the elements of $\{0,1\}\times\mathcal{A}$ to lattice points. In addition, consider a lattice vector $\bar{v}$ represented in the orthonormal basis $\bar{v}=\sum_{i=1}^n v_i\bar{e}_i$. The function g maps $\bar{v}$ to the vector $(\lfloor v_1/(c_{\mathsf{cub}}n^{\frac{1}{2}+2\mathsf{f}}\cdot l)-w_1\rfloor,\ldots,\lfloor v_n/(c_{\mathsf{cub}}n^{\frac{1}{2}+2\mathsf{f}}\cdot l)-w_n\rfloor)$ in $\mathbb{Z}^n$ where the constant $c_{\mathsf{cub}}>0$ will be specified later.

In the following, we describe a routine that creates one register in the input to the two point problem that hides the difference $(u_1,\ldots,u_{i_0-1},\frac{u_{i_0}-m}{p},u_{i_0+1},\ldots,u_n)$. We call the routine $poly(n\log M)=poly(n)$ times in order to create a complete input to the two point problem. We then call the two point algorithm and output its result. This completes the proof of the lemma since with probability $1/poly(n\log M)=1/poly(n)$ our output is correct.

The routine starts by choosing $w_1,\ldots,w_n$ uniformly from $[0,1)$. We create the state

$$\frac{1}{\sqrt{2M^n}}\sum_{t\in\{0,1\},\bar{a}\in\mathcal{A}}|t,\bar{a}\rangle.$$

Then, we compute the function $F=g\circ f$ and measure the result, say $r_1,\ldots,r_n$. The state collapses to (normalization omitted)

$$\sum_{\substack{t\,\in\,\{0,1\}\quad\bar{a}\,\in\,\mathcal{A}\\ F(t,\bar{a})\,=\,(r_1,\ldots,r_n)}}|t,\bar{a}\rangle|r_1,\ldots,r_n\rangle.$$

This completes the description of the routine. Its correctness is shown in the next two claims.

Claim 3.5 *For every $\bar{r}\in\mathbb{Z}^n$, there is at most one element of the form $(0,\bar{a})$ and at most one element of the form $(1,\bar{a}')$ that get mapped to $\bar{r}$ by F. Moreover, if both $(0,\bar{a})$ and $(1,\bar{a}')$ get mapped to $\bar{r}$ then $\bar{a}'-\bar{a}$ is the vector $(\frac{u_1-m}{p},u_2,\ldots,u_m)$.*

Proof: Consider two different lattice points in the image of f, $\bar{v}=f(t,\bar{a})$ and $\bar{v}'=f(t',\bar{a}')$, that get mapped to $\bar{r}$ by g. Let $\bar{v}=\sum_{i=1}^n v_i\bar{e}_i$ and $\bar{v}'=\sum_{i=1}^n v_i'\bar{e}_i$ be their representation in the orthonormal basis. If $\bar{v}'-\bar{v}$ is not a multiple of the shortest vector, then $\|\bar{v}'-\bar{v}\|>c_{\mathsf{unq}}n^{1+2\mathsf{f}}\|\bar{u}\|\geq\frac{1}{2}c_{\mathsf{unq}}n^{1+2\mathsf{f}}\cdot l$. Therefore, there exists a coordinate $i\in[n]$ such that $|v_i'-v_i|\geq\frac{1}{2}c_{\mathsf{unq}}n^{\frac{1}{2}+2\mathsf{f}}\cdot l$ and for $c_{\mathsf{unq}}>2c_{\mathsf{cub}}$ this implies $g(\bar{v})\neq g(\bar{v}')$ no matter how $w_1,\ldots,w_n$ are chosen. Hence, $\bar{v}'-\bar{v}=k\cdot\bar{u}$ for some integer $k\neq 0$. By considering the first coordinate of $\bar{v}'-\bar{v}$ in the lattice basis we get that $(a_1'p+t'm)-(a_1p+tm)\equiv k\cdot m\pmod{p}$. This implies that $k\equiv t'-t\pmod{p}$. If $t=t'$ then $k\equiv 0\pmod{p}$ which implies that $|k|\geq p$. Thus, $\|\bar{v}'-\bar{v}\|\geq p\|\bar{u}\|>c_{\mathsf{cub}}n^{1+2\mathsf{f}}\cdot l$ and again, $g(\bar{v})\neq g(\bar{v}')$. This proves the first part of the claim. For the second part, let $t=0$ and $t'=1$. Then, $k\equiv 1\pmod{p}$. As before, this can only happen when $k=1$ and hence the second part of the claim holds. ∎

Hence, it is enough to show that the probability that this register is bad is low enough. The probability of measuring $|r_1,\ldots,r_n\rangle$ equals $\frac{1}{2M^n}\cdot|\{(t,\bar{a})\mid F(t,\bar{a})=(r_1,\ldots,r_n)\}|$. Notice that this probability is the same as the probability that $F(t,\bar{a})=(r_1,\ldots,r_n)$ for randomly chosen t and $\bar{a}$. Hence, we consider a randomly chosen t and $\bar{a}$. If $t=0$, let $\bar{a}'=(a_1+\frac{u_1-m}{p},a_2+u_2,\ldots,a_n+u_n)$ and if $t=1$ let $\bar{a}'=(a_1-\frac{u_1-m}{p},a_2-u_2,\ldots,a_n-u_n)$.

Claim 3.6 *With probability at least $1-\frac{1}{(n\log(2M))^{\mathsf{f}}}$, for randomly chosen t and $\bar{a}$, $\bar{a}'$ is in $\mathcal{A}$ and $F(1-t,\bar{a}')=F(t,\bar{a})$.*

Proof: We assume that $t = 0$, the proof for $t = 1$ is similar. According to Lemma 3.3, $|u_i| < 2^{2n}$. Hence, unless there exists an i for which $a_i < 2^{2n}$ or $a_i > M - 2^{2n}$, $\bar{a}'$ is guaranteed to be in $\mathcal{A}$. This happens with probability at most $n2^{2n+1}/M$ because $\bar{a}$ is a random element of $\mathcal{A}$.

Notice that $f(1, \bar{a}') - f(0, \bar{a}) = \bar{u}$. Since $w_1, \ldots, w_n$ are randomly chosen, the probability that $F(1-t, \bar{a}')$ and $F(t, \bar{a})$ differ on the i'th coordinate is at most

$$\frac{|\langle \bar{u}, \bar{e}_i \rangle|}{c_{\mathsf{cub}} n^{\frac{1}{2}+2\mathsf{f}} \cdot l} \le \frac{|\langle \bar{u}, \bar{e}_i \rangle|}{c_{\mathsf{cub}} n^{\frac{1}{2}+2\mathsf{f}} \cdot \|\bar{u}\|}.$$

By the union bound, the probability that $F(1-t, \bar{a}') \ne F(t, \bar{a})$ is at most

$$\frac{\sum_i |\langle \bar{u}, \bar{e}_i \rangle|}{c_{\mathsf{cub}} n^{\frac{1}{2}+2\mathsf{f}} \cdot \|\bar{u}\|} \le \frac{1}{c_{\mathsf{cub}} n^{2\mathsf{f}}}$$

where we used the fact that the l_1 norm of a vector is at most $\sqrt{n}$ times its l_2 norm.

The sum of the two error probabilities $n\frac{2^{2n+1}}{M} + \frac{1}{c_{\mathsf{cub}} n^{2\mathsf{f}}}$ is at most $\frac{1}{(n \log(2M))^{\mathsf{f}}}$ for c_{cub} large enough. ■

This concludes the proof of Lemma 3.4. ■

3.3 An Improved Algorithm

In this section we complete the proof of Theorem 1.1. The algorithm we describe has many similarities with the one in the previous section. The main difference is that it is based on n-dimensional balls instead of cubes. The idea is to construct a ball of the right radius around lattice points and to show that if two lattice points are close then the two balls have a large intersection while for any two far lattice points the balls do not intersect. For technical reasons, we will assume in this section that the lattice is a subset of $\mathbb{Z}^n$. Any lattice with rational points can be scaled so that it is a subset of $\mathbb{Z}^n$. We begin with some technical claims:

Claim 3.7 *For any $R > 0$, let B_n be the ball of radius R centered around the origin in $\mathbb{R}^n$ and let $B'_n = B_n + \bar{d}$ for some vector $\bar{d}$ be a shifted ball. Then, the relative n-dimensional volume of their intersection is at least $1 - O(\sqrt{n}\|\bar{d}\|/R)$, i.e.,*

$$\frac{\mathrm{vol}(B_n \cap B'_n)}{\mathrm{vol}(B_n)} \ge 1 - O(\sqrt{n}\|\bar{d}\|/R).$$

Proof: Consider a point $\bar{x} \in \mathbb{R}^n$ such that $\langle \bar{x}, \bar{d} \rangle / \|\bar{d}\| \ge \|\bar{d}\|/2$, i.e., a point which is closer to the center of B'_n than to the center of B_n. Notice that $\bar{x} \in B_n$ implies $\bar{x} \in B'_n$. In other words, the cap C_n of B_n given by all such points $\bar{x}$ is contained in $B_n \cap B'_n$. By using a symmetric argument for points $\bar{x} \in \mathbb{R}^n$ such that $\langle \bar{x}, \bar{d} \rangle / \|\bar{d}\| < \|\bar{d}\|/2$ we get,

$$\mathrm{vol}(B_n \cap B'_n) = 2 \cdot \mathrm{vol}(C_n).$$

We can lower bound the volume of C_n by half the volume of B_n minus the volume of an n-dimensional cylinder of radius R and height $\|\bar{d}\|/2$:

$$\mathrm{vol}(C_n) \ge \frac{1}{2}\mathrm{vol}(B_n) - \frac{\|\bar{d}\|}{2}\mathrm{vol}(B_{n-1})$$

where B_{n-1} is the $n-1$-ball of radius R. We complete the proof by using the estimate $\mathrm{vol}(B_{n-1})/\mathrm{vol}(B_n) = O(\sqrt{n}/R)$,

$$\mathrm{vol}(C_n)/\mathrm{vol}(B_n) \ge \frac{1}{2} - O(\sqrt{n}\|\bar{d}\|/R).$$

■

In the algorithm we will actually represent the balls using points of a fine grid. Therefore, we would like to say that the above claim still holds if we consider the number of grid points inside B_n, B'_n and $B_n \cap B'_n$ instead of their volumes. The following claim is more than enough for our needs:

Claim 3.8 (Special case of Proposition 8.7 in [20]) *Let L be an integer and consider the scaled integer grid $\frac{1}{L}\mathbb{Z}^n$. Then, for any convex body Q that contains a ball of radius $r \geq \frac{1}{L}n^{1.5}$,*

$$\left| \frac{|\frac{1}{L}\mathbb{Z}^n \cap Q|}{L^n \mathrm{vol}(Q)} - 1 \right| < \frac{2n^{1.5}}{rL}.$$

Corollary 3.9 *Let $L = 2^n$ and consider the scaled integer grid $\frac{1}{L}\mathbb{Z}^n$. For any $R \geq 1$, let B_n be the ball of radius R centered around the origin in $\mathbb{R}^n$ and let $B'_n = B_n + \bar{d}$ for some vector $\bar{d}$ such that $R/poly(n) \leq \|\bar{d}\| \leq R$. Then, the relative number of grid points in their intersection is at least $1 - O(\sqrt{n}\|\bar{d}\|/R)$, i.e.,*

$$\frac{|\frac{1}{L}\mathbb{Z}^n \cap B_n \cap B'_n|}{|\frac{1}{L}\mathbb{Z}^n \cap B_n|} \geq 1 - O(\sqrt{n}\|\bar{d}\|/R).$$

Proof: We first note that B_n, B'_n and $B_n \cap B'_n$ all contain the ball of radius $R/2 \geq 1/2$ centered around $\bar{d}/2$. Using Claim 3.8 we obtain that the number of grid points in these bodies approximates their volume up to a multiplicative error of $\frac{2n^{1.5}}{L/2} = 2^{-\Omega(n)}$. We complete the proof by using Claim 3.7. ∎

Let $D(\cdot,\cdot)$ denote the trace distance between two quantum states [21]. It is known that the trace distance represents the maximum probability of distinguishing between the two states using quantum measurements. We need the following simple bound on the trace distance:

Claim 3.10 *For all $k > 0$ and density matrices $\sigma_1, \ldots, \sigma_k, \sigma'_1, \ldots, \sigma'_k$,*

$$D(\sigma_1 \otimes \ldots \otimes \sigma_k, \sigma'_1 \otimes \ldots \otimes \sigma'_k) \leq \sum_{i=1}^{k} D(\sigma_i, \sigma'_i)$$

Proof: Using the triangle inequality,

$$\begin{aligned}
D(\sigma_1 \otimes \ldots \otimes \sigma_k,\ \sigma'_1 \otimes \ldots \otimes \sigma'_k) \ &\leq\ D(\sigma_1 \otimes \ldots \otimes \sigma_k,\ \sigma'_1 \otimes \sigma_2 \otimes \ldots \otimes \sigma_k) + \\
&\ D(\sigma'_1 \otimes \sigma_2 \otimes \ldots \otimes \sigma_k,\ \sigma'_1 \otimes \sigma'_2 \otimes \sigma_3 \otimes \ldots \otimes \sigma_k) + \ldots \\
&\ D(\sigma'_1 \otimes \ldots \otimes \sigma'_{k-1} \otimes \sigma_k,\ \sigma'_1 \otimes \ldots \otimes \sigma'_k) \\
&=\ D(\sigma_1,\ \sigma'_1) + D(\sigma_2,\ \sigma'_2) + \ldots + D(\sigma_k,\ \sigma'_k).
\end{aligned}$$

∎

In addition, we will need the following lemma:

Lemma 3.11 *For any $1 \leq R \leq 2^{poly(n)}$, let*

$$|\eta\rangle = \frac{1}{\sqrt{|\frac{1}{L}\mathbb{Z}^n \cap B_n|}} \sum_{\bar{x} \in \frac{1}{L}\mathbb{Z}^n \cap B_n} |\bar{x}\rangle$$

be the uniform superposition on grid points inside a ball of radius R around the origin where $L = 2^n$. Then, for any $c > 0$, a state $|\tilde{\eta}\rangle$ whose trace distance from $|\eta\rangle$ is at most $1/n^c$ can be efficiently computed.

Proof: In order to bound the trace distance, we will use the fact that for any two pure states $|\psi_1\rangle, |\psi_2\rangle$,

$$D(|\psi_1\rangle, |\psi_2\rangle) = \sqrt{1 - |\langle\psi_1|\psi_2\rangle|^2} \leq \||\psi_1\rangle - |\psi_2\rangle\|_2. \quad (1)$$

The first equality appears in [21] and the inequality follows by a simple calculation.

Consider the (continuous) uniform probability distribution q over B_n. Then one can define its discretization q' to the grid $\frac{1}{L}\mathbb{Z}^n$ as

$$q'(\bar{x}) = \int_{\bar{x}+[0,1/L]^n} q(\bar{y})d\bar{y}$$

for $\bar{x} \in \frac{1}{L}\mathbb{Z}^n$. In other words, $q'(\bar{x})$ is proportional to the volume of the intersection of B_n with the cube $\bar{x} + [0, 1/L]^n$. Notice that for points $\bar{x}$ such that $\bar{x} + [0, 1/L]^n$ is completely contained in B_n, $q'(\bar{x}) = 1/(L^n \mathrm{vol}(B_n))$. We claim that the state

$$|\eta'\rangle = \sum_{\bar{x}\in\frac{1}{L}\mathbb{Z}^n} \sqrt{q'(\bar{x})}|\bar{x}\rangle$$

is exponentially close to $|\eta\rangle$. Intuitively, this holds since the two differs only on points which are very close to the boundary of the ball, namely, of distance $\sqrt{n}/L$ from the boundary. The number of such points is negligible compared to the number of points in the interior of the ball. More formally, define

$$|\eta''\rangle = \sqrt{\frac{L^n \mathrm{vol}(B_n)}{|\frac{1}{L}\mathbb{Z}^n \cap B_n|}}|\eta'\rangle.$$

Using Equation 1,

$$D(|\eta\rangle, |\eta'\rangle) \leq \||\eta'\rangle - |\eta\rangle\|_2 \leq \||\eta'\rangle - |\eta''\rangle\|_2 + \||\eta''\rangle - |\eta\rangle\|_2.$$

The first term is at most $2^{-\Omega(n)}$ according to Claim 3.8. For the second term, notice that the amplitudes of $|\eta''\rangle$ and $|\eta\rangle$ are the same except possibly on points $\bar{x}$ of distance $\sqrt{n}/L$ from the boundary. Using Claim 3.8 again we get that the fraction of such points is closely approximated by one minus the ratio of volumes of the ball of radius $R - \sqrt{n}/L$ and the ball of radius R. This ratio of volumes is $(1 - \sqrt{n}/(RL))^n \geq (1 - \sqrt{n}/L)^n \geq 1 - n^{1.5}/L = 1 - 2^{-\Omega(n)}$.

In the following we show how to approximate the state $|\eta'\rangle$. This idea is essentially due to Grover and Rudolph [9]. Let $m \in \mathbb{Z}$ be large enough so that B_n is contained in the cube $[-2^m, 2^m]^n$. Using our assumption on R, $m < n^{c_1}$ for some $c_1 \geq 1$. We represent $\bar{x}$ using $K = n(m+1+\log L) < 2n^{1+c_1}$ qubits, i.e., a block of $m + 1 + \log L$ qubits for each dimension. Hence, we can write $|\eta'\rangle$ as

$$|\eta'\rangle = \sum_{x_1,\ldots,x_K\in\{0,1\}} \sqrt{q'(x_1, \ldots, x_K)}|x_1, \ldots, x_K\rangle.$$

We now show an equivalent way of writing $|\eta'\rangle$. Let us extend the definition of q' in the following way: for any $k \leq K$ and any $x_1, \ldots, x_k \in \{0, 1\}$ define $q'(x_1, \ldots, x_k)$ as the sum of $q'(x_1, \ldots, x_k, x_{k+1}, \ldots, x_K)$ over all sequences $x_{k+1}, \ldots, x_K \in \{0, 1\}$. Notice that $q'(x_1, \ldots, x_k)$ corresponds to the volume of the intersection of B_n with a certain cuboid (also known as a rectangular parallelepiped). For example, $q'(0) = q'(1) = \frac{1}{2}$ since they represent the intersection of B_n with two halves of the cube $[-2^m, 2^m]^n$. Using the definition $s(x_1) = q'(x_1)$ and for $k > 1$, $s(x_1, \ldots, x_k) = q'(x_1, \ldots, x_k)/q'(x_1, \ldots, x_{k-1})$ we see that

$$|\eta'\rangle = \sum_{x_1\in\{0,1\}} \sqrt{s(x_1)} \sum_{x_2\in\{0,1\}} \sqrt{s(x_1, x_2)} \ldots \sum_{x_K\in\{0,1\}} \sqrt{s(x_1, \ldots, x_K)}|x_1, \ldots, x_K\rangle.$$

The algorithm starts with all K qubits in the state $|0\rangle$ and sets one qubit at a time. The first qubit is rotated to the state $\frac{1}{\sqrt{2}}(|0\rangle+|1\rangle)$. Assume we are now in the k'th step after setting the state of qubits $1,\ldots,k-1$. We use the fact that there exists a classical algorithm for approximating the volume of a convex body up to any $1/poly(n)$ error (see [15] and references therein). The body should be provided by a "well-guaranteed weak membership oracle", i.e., a sphere containing the body, a sphere contained in the body, both of non-zero radius and an oracle that given a point decides if it is inside the body or not. It is easy to construct such two spheres and an oracle for a body given by the intersection of a ball with a cuboid. Hence, we can compute two values $\tilde{s}(x_1,\ldots,x_{k-1},0)$ and $\tilde{s}(x_1,\ldots,x_{k-1},1)$ such that

$$\tilde{s}(x_1,\ldots,x_{k-1},0)+\tilde{s}(x_1,\ldots,x_{k-1},1)=1$$

and

$$\left|\frac{\tilde{s}(x_1,\ldots,x_{k-1},i)}{s(x_1,\ldots,x_{k-1},i)}-1\right|<n^{-c_2}$$

for $i=0,1$ and some constant c_2 which will be chosen later. Then, we rotate the i'th qubit to the state $\sqrt{\tilde{s}(x_1,\ldots,x_{k-1},0)}|0\rangle+\sqrt{\tilde{s}(x_1,\ldots,x_{k-1},1)}|1\rangle$. This completes the description of the procedure.

Notice that the amplitude of each basis state $|x_1,\ldots,x_K\rangle$ in the resulting state $|\tilde{\eta}\rangle$ is given by

$$\prod_{k=1}^{K}\sqrt{\tilde{s}(x_1,\ldots,x_k)}\geq(1-n^{-c_2})^K\prod_{k=1}^{K}\sqrt{s(x_1,\ldots,x_k)}.$$

Hence the inner product $\langle\tilde{\eta}|\eta'\rangle$ is at least

$$\begin{aligned}(1-n^{-c_2})^K\sum_{x_1,\ldots,x_K\in\{0,1\}}\prod_{k=1}^{K}s(x_1,\ldots,x_k) &= (1-n^{-c_2})^K\sum_{x_1,\ldots,x_K\in\{0,1\}}q'(x_1,\ldots,x_K)\\ &= (1-n^{-c_2})^K\geq 1-K\cdot n^{-c_2}\geq 1-2n^{1+c_1-c_2}.\end{aligned}$$

Using Equation 1,

$$D(|\eta'\rangle,|\tilde{\eta}\rangle)=\sqrt{1-|\langle\tilde{\eta}|\eta'\rangle|^2}<n^{-c}$$

for a large enough c_2. ∎

Let $p>n^{2+2\mathsf{f}}$ be any fixed prime. The following is the main lemma of this section. It essentially replaces Lemma 3.4 and hence implies Theorem 1.1.

Lemma 3.12 *For any $\mathsf{f}>0$ let $\bar{u}=\sum_{i=1}^{n}u_i\bar{b}_i$ be the shortest lattice vector in a $(c_{\mathsf{unq}}n^{\frac{1}{2}+2\mathsf{f}})$-unique lattice where $c_{\mathsf{unq}}>0$ is a constant. If there exists a solution to the two point problem with failure parameter f then there exists a quantum algorithm that given this lattice and three integers l,m,i_0 returns $(u_1,\ldots,u_{i_0-1},\frac{u_{i_0}-m}{p},u_{i_0+1},\ldots,u_n)$ with probability $1/poly(n)$ if the following conditions hold: $\|\bar{u}\|\leq l\leq 2\|\bar{u}\|$, $u_{i_0}\equiv m \pmod p$ and $1\leq m\leq p-1$.*

Proof: As before, let $\langle\bar{b}_1,\ldots,\bar{b}_n\rangle$ be an LLL reduced basis, let $M=2^{4n}$ and assume that $i_0=1$. We also define $f(t,\bar{a})$ as before. Assume that the number of registers needed by the two point algorithm is at most n^{c_1} for some constant $c_1>0$.

The algorithm starts by calling the routine of Claim 3.11 n^{c_1} times with accuracy parameter n^{-c_2} and $R=c_{\mathsf{bal}}n^{\frac{1}{2}+2\mathsf{f}}\cdot l$ for some constants $c_2,c_{\mathsf{bal}}>0$. The state we obtain is

$$|\tilde{\eta}_1\rangle\otimes\ldots\otimes|\tilde{\eta}_{n^{c_1}}\rangle \tag{2}$$

where each $|\tilde{\eta}_i\rangle$ has a trace distance of at most n^{-c_2} from $|\eta\rangle$. According to Claim 3.10, the above tensor product has a trace distance of at most $n^{c_1-c_2}$ from $|\eta\rangle^{\otimes n^{c_1}}$. In the following we show that the algorithm succeeds with probability at least n^{-c_3} for some $c_3 > 0$ given the state $|\eta\rangle^{\otimes n^{c_1}}$. This would complete the proof since given the state in Equation 2, the algorithm succeeds with probability at least $n^{-c_3} - n^{c_1-c_2} > \frac{1}{2}n^{-c_3}$ for large enough c_2.

We describe a routine that given the state $|\eta\rangle$ creates one register in the input to the two point problem. In order to produce a complete input to the two point problem, the algorithm calls this routine n^{c_1} times, each time with a new $|\eta\rangle$ register. It then calls the two point algorithm and outputs the result. As required, the success probability is $1/poly(n \log M) = n^{-c_3}$ for some $c_3 > 0$.

Given $|\eta\rangle$, the routine creates the state

$$\frac{1}{\sqrt{2M^n}} \sum_{t\in\{0,1\},\bar{a}\in\mathcal{A}} |t,\bar{a}\rangle \otimes |\eta\rangle,$$

or equivalently,

$$\sum_{t\in\{0,1\},\bar{a}\in\mathcal{A},\bar{x}\in\frac{1}{L}\mathbb{Z}^n\cap B_n} |t,\bar{a},\bar{x}\rangle$$

where B_n is the ball of radius R around the origin and $L = 2^n$. We add the value $f(t,\bar{a})$ to the last register,

$$\sum_{t\in\{0,1\},\bar{a}\in\mathcal{A},\bar{x}\in\frac{1}{L}\mathbb{Z}^n\cap B_n} |t,\bar{a},f(t,\bar{a})+\bar{x}\rangle.$$

Finally, we measure the last register and if $\bar{x}'$ denotes the result, the state collapses to

$$\sum_{t\in\{0,1\},\bar{a}\in\mathcal{A}|\bar{x}'\in f(t,\bar{a})+\frac{1}{L}\mathbb{Z}^n\cap B_n} |t,\bar{a},\bar{x}'\rangle.$$

Claim 3.13 *For every $\bar{x}'$, there is at most one element of the form $(0,\bar{a})$ and at most one element of the form $(1,\bar{a}')$ such that $\bar{x}' \in f(t,\bar{a}) + \frac{1}{L}\mathbb{Z}^n \cap B_n$. Moreover, if there are two such elements $(0,\bar{a})$ and $(1,\bar{a}')$ then $\bar{a}' - \bar{a}$ is the vector $(\frac{u_1-m}{p}, u_2, \ldots, u_m)$.*

Proof: Consider two different lattice points in the image of f, $\bar{v} = f(t,\bar{a})$ and $\bar{v}' = f(t',\bar{a}')$, such that $\bar{x}'$ is both in $\bar{v} + \frac{1}{L}\mathbb{Z}^n \cap B_n$ and $\bar{v}' + \frac{1}{L}\mathbb{Z}^n \cap B_n$. This implies that $\|\bar{v} - \bar{v}'\| \le c_{\mathsf{bal}} n^{\frac{1}{2}+2\mathsf{f}} \cdot l \le 2c_{\mathsf{bal}} n^{\frac{1}{2}+2\mathsf{f}} \cdot \|\bar{u}\|$. For $c_{\mathsf{unq}} > 2c_{\mathsf{bal}}$ this means that $\bar{v}' - \bar{v} = k \cdot \bar{u}$ for some integer $k \neq 0$. As before, by considering the first coordinate of $\bar{v}' - \bar{v}$ in the lattice basis we get that $(a_1'p + t'm) - (a_1 p + tm) \equiv k \cdot m \pmod p$. Hence, $k \equiv t' - t \pmod p$. If $t = t'$ then $k \equiv 0 \pmod p$ and therefore $|k| \ge p$ which contradicts the above upper bound on the distance between $\bar{v}$ and $\bar{v}'$. This proves the first part of the claim. For the second part, let $t = 0$ and $t' = 1$. Then, $k \equiv 1 \pmod p$. As before, this can only happen when $k = 1$ and hence the second part of the claim holds. ∎

Notice that the probability of measuring $\bar{x}'$ is the same as that obtained by first choosing random t and $\bar{a}$ and then choosing a random point in $f(t,\bar{a}) + \frac{1}{L}\mathbb{Z}^n \cap B_n$. Let us define for any t and $\bar{a}$ the vector $\bar{a}'$ as before.

Claim 3.14 *With probability at least $1 - \frac{1}{(n\log(2M))^{\mathsf{f}}}$, for randomly chosen t and $\bar{a}$ and a random point $\bar{x}'$ in $f(t,\bar{a}) + \frac{1}{L}\mathbb{Z}^n \cap B_n$, $\bar{a}'$ is in $\mathcal{A}$ and $\bar{x}'$ is also in $f(1-t,\bar{a}') + \frac{1}{L}\mathbb{Z}^n \cap B_n$.*

Proof: According to Lemma 3.3, $|u_i| < 2^{2n}$. Hence, unless there exists an i for which $a_i < 2^{2n}$ or $a_i > M - 2^{2n}$, $\bar{a}'$ is guaranteed to be in $\mathcal{A}$. This happens with probability at most $n2^{2n+1}/M$ because $\bar{a}$ is a random element of $\mathcal{A}$.

Fix $\bar{a}, \bar{a}' \in \mathcal{A}$. We would like to show that if $\bar{x}'$ is chosen uniformly from $f(t,\bar{a}) + \frac{1}{L}\mathbb{Z}^n \cap B_n$ then with high probability it is also in $f(1-t,\bar{a}') + \frac{1}{L}\mathbb{Z}^n \cap B_n$. By translating both sets by $-f(t,\bar{a})$ we get the equivalent statement that if $\bar{x}'$ is chosen uniformly from $\frac{1}{L}\mathbb{Z}^n \cap B_n$ then with high probability it is also in $(f(1-t,\bar{a}') - f(t,\bar{a})) + \frac{1}{L}\mathbb{Z}^n \cap B_n$. Since we assumed that our lattice is a subset of $\mathbb{Z}^n$, $f(1-t,\bar{a}') - f(t,\bar{a}) \in \mathbb{Z}^n$ and the latter set equals $\frac{1}{L}\mathbb{Z}^n \cap (f(1-t,\bar{a}') - f(t,\bar{a}) + B_n)$. Using Corollary 3.9 and the fact that $\|f(1-t,\bar{a}') - f(t,\bar{a})\| = \|\bar{u}\| \le l$, we get that the required probability is at least

$$1 - O(\sqrt{n}l/R) = 1 - O(\sqrt{n}l/(c_{\mathsf{bal}} n^{\frac{1}{2}+2\mathsf{f}} \cdot l)) = 1 - O(1/(c_{\mathsf{bal}} n^{2\mathsf{f}})).$$

The sum of the two error probabilities $n\frac{2^{2n+1}}{M} + O(1/(c_{\mathsf{bal}} n^{2\mathsf{f}}))$ is at most $\frac{1}{(n\log(2M))^{\mathsf{f}}}$ for c_{bal} large enough. ■

This concludes the proof of Lemma 3.12. ■

4 The Dihedral Coset Problem

We begin this section with a description of the average case subset sum problem. We describe our assumptions on the subroutine that solves it and prove some properties of such a subroutine. In the second subsection we present an algorithm that solves the DCP with calls to an average case subset sum subroutine.

4.1 Subset Sum

The subset sum problem is defined as follows. An input is a sequence of numbers $A = (a_1, \ldots, a_r)$ and two numbers t, N. The output is a subset $B \subseteq [r]$ such that $\sum_{i\in B} a_i \equiv t \pmod N$. Let a legal input be an input for which there exists a subset B with $\sum_{i\in B} a_i \equiv t \pmod N$. For a constant $c_{\mathsf{r}} > 0$, we fix r to be $\log N + c_{\mathsf{r}}$ since we will only be interested in such instances. First we show that there are many legal inputs:

Lemma 4.1 *For randomly chosen $a_1, \ldots, a_r, t$ in $\{0, \ldots, N-1\}$, the probability that there is no $B \subseteq [r]$ such that $\sum_{i\in B} a_i \equiv t \pmod N$ is at most $\frac{1}{2}$.*

Proof: Fix a value of t. Define a random variable $X_{\bar{b}}$ for every $\bar{b} \in \{0,1\}^r, \bar{b} \ne 0^r$ as 1 if $\sum_i b_i a_i \equiv t \pmod N$ and 0 otherwise. Since for every $\bar{b}$ the sum $\sum_i b_i a_i$ has any value modulo N with the same probability, the expectation of $X_{\bar{b}}$ is $\frac{1}{N}$ and its variance is $\frac{1}{N} - \frac{1}{N^2} < \frac{1}{N}$. Hence,

$$E[\sum_{\bar{b}} X_{\bar{b}}] = \sum_{\bar{b}} E[X_{\bar{b}}] = \frac{2^r - 1}{N}$$

Given two different sequences $\bar{b}, \bar{b}' \in \{0,1\}^r$ we show that $X_{\bar{b}}$ and $X_{\bar{b}'}$ are independent. Let i be such that $b_i \ne b'_i$ and assume without loss of generality that $b_i = 1, b'_i = 0$ and $i = 1$. Then,

$$\begin{aligned} Pr_{a_1,\ldots,a_r}[X_{\bar{b}} = 1 \ \wedge \ X_{\bar{b}'} = 1] &= E_{a_2,\ldots,a_r}[Pr_{a_1}[X_{\bar{b}} = 1 \ \wedge \ X_{\bar{b}'} = 1]] \\ &= E_{a_2,\ldots,a_r}[1/N \cdot \delta_{X_{\bar{b}'},1}] \\ &= Pr_{a_1,\ldots,a_r}[X_{\bar{b}} = 1] Pr_{a_1,\ldots,a_r}[X_{\bar{b}'} = 1] \end{aligned}$$

where the second equality holds because $X_{\bar{b}'}$ does not depend on a_1 and $X_{\bar{b}}$ is 1 with probability $1/N$ for any $a_2, \ldots, a_r$. A similar argument holds for other values of $X_{\bar{b}}$ and $X_{\bar{b}'}$. Therefore, the random variables are pairwise independent and by the Chebyshev bound,

$$Pr[\sum_{\bar{b}} X_{\bar{b}} < \frac{1}{2} \cdot \frac{2^r - 1}{N}] \leq 4 \cdot \frac{N}{2^r - 1} \leq \frac{8}{2^{c_r}}.$$

In particular, the probability of $\sum_{\bar{b}} X_{\bar{b}} = 0$, that is, the probability that there is no B such that $\sum_{i \in B} a_i \equiv t \pmod N$ is at most $\frac{8}{2^{c_r}} = \frac{1}{2}$ for $c_r = 4$. ∎

We assume that we are given a subroutine that answers a $\frac{1}{\log^{c_s} N}$ fraction of the legal subset sum inputs with parameter N where $c_s > 0$ is any constant. As can be seen from the previous lemma, this implies that the subroutine answers a non-negligible fraction of all inputs (and not just the legal inputs). In addition, we assume that the subroutine is deterministic. We denote by $S(A,t)$ the result of the subroutine S on the input $A = (a_1, \ldots, a_r), t$ and we omit N. This result can either be a set or an error. Let $S(A)$ denote the set of t's for which the subroutine returns a set and not an error, i.e., $S(A) = \{t \mid S(A,t) \neq error\}$.

Corollary 4.2 *For randomly chosen $a_1, \ldots, a_r$ in $\{0, \ldots, N-1\}$, $Pr_A[|S(A)| \geq \frac{N}{4\log^{c_s} N}] = \Omega(\frac{1}{\log^{c_s} N})$ where $A = (a_1, \ldots, a_r)$.*

Proof: Since $S(A,t) \neq error$ only when (A,t) is a legal input,

$$\begin{aligned} Pr_{A,t}[S(A,t) \neq error] &= Pr_{A,t}[\ S(A,t) \neq error \ \wedge \ (A,t) \text{ is legal }] \\ &= Pr_{A,t}[\ S(A,t) \neq error \mid (A,t) \text{ is legal }] \cdot Pr_{A,t}[\ (A,t) \text{ is legal }] \geq \frac{1}{2\log^{c_s} N}. \end{aligned}$$

In addition,

$$\begin{aligned} Pr_{A,t}[S(A,t) \neq error] &= E_A[\ \frac{|S(A)|}{N}\ ] \\ &\leq Pr_A[\ |S(A)| \geq \frac{N}{4\log^{c_s} N}\ ] + Pr_A[\ |S(A)| < \frac{N}{4\log^{c_s} N}\ ] \cdot \frac{1}{4\log^{c_s} N} \\ &\leq Pr_A[\ |S(A)| \geq \frac{N}{4\log^{c_s} N}\ ] + \frac{1}{4\log^{c_s} N}. \end{aligned}$$

By combining the two inequalities we obtain the corollary. ∎

Lemma 4.3 *Let $T \subseteq \{0, \ldots, N-1\}$ be a set such that $|T| > \frac{N}{s}$ for a certain s. Then, for any $q < \frac{N}{8s}$ there exists $q' \in \{q, 2q, \ldots, sq\}$ such that the number of pairs $t, t+q'$ that are both in T is $\Omega(\frac{N}{s^3})$.*

Proof: Define the partition of T into sets $T_0, \ldots, T_{q-1}$ as

$$T_k = \{i \mid i \in T, i \equiv k \pmod q\}.$$

At least $\frac{q}{2s}$ of the sets are of size at least $\frac{N}{2sq}$ since their union is T and $\frac{q}{2s}\frac{N}{q} + \frac{N}{2s} < |T|$. Let T_i be such a set and for $t \in T_i$ consider the values $t+q, t+2q, \ldots, t+4sq$. Therefore, the number of $t \in T_i$ such that none of these values is in T_i is less than $\frac{N}{4sq}$ because $|\{i \mid 0 \leq i < N, i \equiv k \pmod q\}| = \frac{N}{q}$. Therefore, more than $|T_i| - \frac{N}{4sq} \geq \frac{N}{4sq}$ of the elements $t \in T_i$ are such that one of $t+q, t+2q, \ldots, t+4sq$ is also in T_i. Summing over all sets T_i such that $|T_i| \geq \frac{N}{2sq}$, there are at least $\frac{N}{4sq} \cdot \frac{q}{2s} = \frac{N}{8s^2}$ elements $t \in T$ for which one of $t+q, t+2q, \ldots, t+4sq$ is also in T. Thus, there exists a $q' \in \{q, 2q, \ldots, 4sq\}$ such that the number of $t \in T$ for which $t + q' \in T$ is at least $\frac{N}{32s^3}$. ∎

Definition 4.4 *A partial function $f : \{0, \ldots, N-1\} \to \{0, \ldots, N-1\}$ is called a matching if for all i such that $f(i)$ is defined, $f(i) \neq i$ and $f(f(i)) = i$. A matching is a q-matching if for all i such that $f(i)$ is defined, $|f(i) - i| = q$. We define an equal partition of the domain of a matching f by $A_1(f) = \{i \mid f(i) \text{ defined } \wedge\ f(i) > i\}$ and $A_2(f) = \{i \mid f(i) \text{ defined } \wedge\ f(i) < i\}$. The intersection of a matching f and a set $T \subseteq \{0, \ldots, N-1\}$ is the set $\{i \mid i \in T\ \wedge\ f(i) \in T\}$.*

For any q we define the following q-matchings:

$$f^1_q(t) = \begin{cases} t+q & t \bmod 2q < q,\ t+q < N, \\ t-q & t \bmod 2q \geq q,\ t-q \geq 0, \\ undefined & \text{otherwise.} \end{cases} \quad f^2_q(t) = \begin{cases} t-q & t \bmod 2q < q,\ t-q \geq 0, \\ t+q & t \bmod 2q \geq q,\ t+q < N, \\ undefined & \text{otherwise.} \end{cases}$$

Lemma 4.5 *There exists a constant c_m such that for any integer $q < \frac{N}{\log^{c_\mathsf{m}} N}$ there exists a matching f among the $2\log^{c_\mathsf{m}} N$ matchings $f^1_q, f^1_{2q}, \ldots, f^1_{\log^{c_\mathsf{m}} N q}, f^2_q, f^2_{2q}, \ldots, f^2_{\log^{c_\mathsf{m}} N q}$ such that with probability at least $\frac{1}{\log^{c_\mathsf{m}} N}$ on the choice of A, the intersection of f and $S(A)$ is $\frac{N}{\log^{c_\mathsf{m}} N}$. We call such an f a* good *matching.*

Proof: According to Corollary 4.2, $\frac{1}{4\log^{c_\mathsf{s}} N}$ of the possible values of A satisfy $|S(A)| > \frac{N}{4\log^{c_\mathsf{s}} N}$. For such A, Lemma 4.3 with $s = 4\log^{c_\mathsf{s}} N$ implies that there exists a value $q' \in \{q, 2q, \ldots, 4\log^{c_\mathsf{s}} N \cdot q\}$ such that the number of pairs $t, t+q'$ that are both in $S(A)$ is $\Omega(\frac{N}{\log^{3c_\mathsf{s}} N})$. Therefore, for such A and q', the size of the intersection of one of the matchings $f^1_{q'}, f^2_{q'}$ and $S(A)$ is $\Omega(\frac{N}{\log^{3c_\mathsf{s}} N})$. This implies that one of the $8\log^{c_\mathsf{s}} N$ matchings considered must have an intersection of size $\Omega(\frac{N}{\log^{3c_\mathsf{s}} N})$ with at least $\frac{1}{32\log^{2c_\mathsf{s}} N}$ of the possible values of A. We conclude the proof by choosing $c_\mathsf{m} > 3c_\mathsf{s}$. ∎

4.2 The Quantum Algorithm

We begin with the following simple claim:

Claim 4.6 *For any two basis states $|a\rangle$ and $|b\rangle$, $a \neq b$, there exists a routine such that given the state $|a\rangle + e(\phi)|b\rangle$ outputs the state $|0\rangle + e(\phi)|1\rangle$.*

Proof: Consider the function f defined as $f(a) = 0, f(0) = a, f(b) = 1, f(1) = b$ and $f(i) = i$ otherwise. It is reversible and can therefore be implemented as a quantum routine. ∎

We now describe the main routine in the DCP algorithm.

Lemma 4.7 *There exist routines R_1, R_2 such that given a q-matching f and an input for the DCP with failure parameter 1, they either output a bit or they fail. Conditioned on non-failure, the probability of the bit being 1 is $\frac{1}{2} - \frac{1}{2}\cos(2\pi q \frac{d}{N})$ for R_1 and $\frac{1}{2} + \frac{1}{2}\sin(2\pi q \frac{d}{N})$ for R_2. Moreover, if f is a good matching, the success probability is $\Omega(\frac{1}{\log^{c_\mathsf{m}} N})$.*

Proof: The routines begin by performing a Fourier transform on the last $\log N$ qubits of each input register. Consider one register. Assuming it is a good register, the resulting state is

$$\frac{1}{\sqrt{2N}} \sum_{i=0}^{N-1} e(ix/N)|0, i\rangle + \frac{1}{\sqrt{2N}} \sum_{i=0}^{N-1} e(i(x+d)/N)|1, i\rangle =$$

$$\frac{1}{\sqrt{2N}} \sum_{i=0}^{N-1} e(ix/N)(|0\rangle + e(id/N)|1\rangle)|i\rangle.$$

We measure the last $\log N$ qubits and let $a \in \{0, \ldots, N-1\}$ be the result. The state collapses to

$$\frac{1}{\sqrt{2}} e(ax/N)(|0\rangle + e(ad/N)|1\rangle)|a\rangle.$$

If it is a bad register, it is in the state $|b, x\rangle$ where both b and x are arbitrary. After the Fourier transform the state is $\frac{1}{\sqrt{N}} \sum_{i=0}^{N-1} e(ix/N)|b, i\rangle$ and after measuring a in the last $\log N$ qubits, the state is $e(ax/N)|b, a\rangle$. Notice that in both cases any value a in $\{0, \ldots, N-1\}$ has an equal probability of being measured.

We choose the number of input registers to be r. Let $A = (a_1, \ldots, a_r)$ be the sequence of values measured in the above process. Notice that this sequence is uniform and hence can be used as an input to the average case subset sum algorithm. In the following, we assume that s of the r registers are bad. Later we will claim that with good probability, none of the registers is bad. Yet, we have to show that even if one of the registers is bad, the routine does not return erroneous results. Without loss of generality, assume that the first s registers are bad. The resulting state is:

$$\bigotimes_{i=1}^{s} [e(a_i x_i/N)|b_i, a_i\rangle] \bigotimes_{i=s+1}^{r} [\frac{1}{\sqrt{2}} e(a_i x_i/N)(|0\rangle + e(a_i d/N)|1\rangle)|a_i\rangle].$$

Or, by omitting the multiplication by the fixed phase and the $r \cdot \lceil \log N \rceil$ fixed qubits,

$$\bigotimes_{i=1}^{s} [|b_i\rangle] \bigotimes_{i=s+1}^{r} [\frac{1}{\sqrt{2}} (|0\rangle + e(a_i d/N)|1\rangle)].$$

Denote these r qubits by $\bar{\alpha} = (\alpha_1, \ldots, \alpha_r)$.

We add $r+1$ new qubits, $\bar{\beta} = (\beta_1, \ldots, \beta_r)$ and γ. Let $t_{\bar{\alpha}}$ denote the sum $\sum_{i=1}^{r} \alpha_i a_i$. Next, we perform the following operations:

if $S(A, t_{\bar{\alpha}}) \neq \bar{\alpha} \ \vee \ S(A, f(t_{\bar{\alpha}})) = \mathit{error}$
  then *exit*
if $t_{\bar{\alpha}} \in A_1(f)$
  then $\begin{cases} \bar{\beta} \leftarrow \bar{\alpha} \\ \gamma \leftarrow 1 \end{cases}$
  else if $t_{\bar{\alpha}} \in A_2(f)$
  then $\begin{cases} \bar{\beta} \leftarrow S(A, f(t_{\bar{\alpha}})) \\ \gamma \leftarrow 1 \end{cases}$
  else *exit*

In order to describe the state after the above procedure, we define the following subsets of $\{0,1\}^r$:

$$M = \{\bar{\alpha} \in \{0,1\}^r \mid \alpha_1 = b_1, \ldots, \alpha_s = b_s\}$$

$$L = \{\bar{\alpha} \in M \mid t_{\bar{\alpha}} \in A_1(f) \ \wedge \ S(A, t_{\bar{\alpha}}) = \bar{\alpha} \ \wedge S(A, f(t_{\bar{\alpha}})) \neq \mathit{error}\}$$

$$R = \{\bar{\alpha} \in M \mid t_{\bar{\alpha}} \in A_2(f) \ \wedge \ S(A, t_{\bar{\alpha}}) = \bar{\alpha} \ \wedge S(A, f(t_{\bar{\alpha}})) \neq \mathit{error}\}$$

Using the order $|\bar{\alpha}, \bar{\beta}, \gamma\rangle$, the resulting state is:

$$\frac{1}{\sqrt{2^{r-s}}} (\sum_{\bar{\alpha} \in M-L-R} e(\langle \bar{\alpha}, \bar{a} \rangle \frac{d}{N}) |\bar{\alpha}, \bar{0}, 0\rangle +$$

$$
\begin{aligned}
& \sum_{\bar{\alpha}\in L} e(\langle\bar{\alpha},\bar{a}\rangle\frac{d}{N})|\bar{\alpha},\bar{\alpha},1\rangle + \sum_{\bar{\alpha}\in R} e(\langle\bar{\alpha},\bar{a}\rangle\frac{d}{N})|\bar{\alpha},S(A,f(t_{\bar{\alpha}})),1\rangle) \\
= \; & \frac{1}{\sqrt{2^{r-s}}}(\sum_{\bar{\alpha}\in M-L-R} e(\langle\bar{\alpha},\bar{a}\rangle\frac{d}{N})|\bar{\alpha},\bar{0},0\rangle + \\
& \sum_{\bar{\alpha}\in L}(e(\langle\bar{\alpha},\bar{a}\rangle\frac{d}{N})|\bar{\alpha},\bar{\alpha},1\rangle + e(\langle S(A,f(t_{\bar{\alpha}})),\bar{a}\rangle\frac{d}{N})|S(A,f(t_{\bar{\alpha}})),\bar{\alpha},1\rangle)) \\
= \; & \frac{1}{\sqrt{2^{r-s}}}(\sum_{\bar{\alpha}\in M-L-R} e(\langle\bar{\alpha},\bar{a}\rangle\frac{d}{N})|\bar{\alpha},\bar{0},0\rangle + \\
& \sum_{\bar{\alpha}\in L} e(\langle\bar{\alpha},\bar{a}\rangle\frac{d}{N})(|\bar{\alpha}\rangle + e(q\cdot\frac{d}{N})|S(A,f(t_{\bar{\alpha}}))\rangle)|\bar{\alpha},1\rangle)
\end{aligned}
$$

Now we measure $\bar{\beta}$ and γ. If $\gamma = 0$, the routine failed. Otherwise, the state of $\bar{\alpha}$ is (omitting the fixed $\bar{\beta}$ and γ):

$$\frac{1}{\sqrt{2}}(|\bar{\beta}\rangle + e(q\cdot\frac{d}{N})|S(A,f(t_{\bar{\beta}}))\rangle).$$

Notice that since $\bar{\beta}$ is known and $S(A,f(t_{\bar{\beta}}))$ can be easily found by calling S, we can transform this state to the state

$$\frac{1}{\sqrt{2}}(|0\rangle + e(q\cdot\frac{d}{N})|1\rangle)$$

by using Claim 4.6. By omitting some qubits, we can assume that this is a state on one qubit. By using the Hadamard transform the state becomes

$$\frac{1}{2}((1+e(q\frac{d}{N}))|0\rangle + (1-e(q\frac{d}{N}))|1\rangle).$$

We measure the qubit and the probability of measuring 1 is

$$\frac{1}{4}|1-e(q\frac{d}{N})|^2 = \frac{1}{4}(2-2\cos(2\pi q\frac{d}{N})) = \frac{1}{2}-\frac{1}{2}\cos(2\pi q\frac{d}{N}).$$

This completes the description of R_1. The routine R_2 applies the transform

$$\begin{pmatrix} 1 & 0 \\ 0 & \imath \end{pmatrix}$$

before the Hadamard transform and thus the state becomes

$$\frac{1}{2}((1+e(1/4+q\frac{d}{N}))|0\rangle + (1-e(1/4+q\frac{d}{N}))|1\rangle)$$

and the probability of measuring 1 becomes $\frac{1}{2}-\frac{1}{2}\cos(\pi/2+2\pi q\frac{d}{N}) = \frac{1}{2}+\frac{1}{2}\sin(2\pi q\frac{d}{N})$.

From the previous description, it is clear that the probability of measuring 1 conditioned on a non-failure is correct. Thus, it remains to prove that when f is a good matching the failure probability is low. The success probability equals the probability of measuring $\gamma = 1$ which is $|L\cup R|/2^{r-s}$. Assume that none of the r registers is bad. Then, $|L\cup R|/2^{r-s} = |L\cup R|/2^r$ and $L\cup R$ becomes $\{\bar{\alpha}\in\{0,1\}^r \mid t_{\bar{\alpha}}\in A_1(f)\cup A_2(f) \ \wedge \ S(A,t_{\bar{\alpha}})=\bar{\alpha} \ \wedge \ S(A,f(t_{\bar{\alpha}}))\neq \mathit{error}\}$. Notice that the size of this set equals $|\{t \mid t\in S(A) \ \wedge \ f(t)\in S(A)\}|$ which, according to the definition of a good matching, is at least $\frac{N}{\log^{c_m} N}$. Therefore the probability of success conditioned on all of the registers being good is $|L\cup R|/2^r = \frac{1}{2^{c_r}\log^{c_m} N} = \Omega(\frac{1}{\log^{c_m} N})$. This concludes the proof since with probability at least $(1-\frac{1}{\log N})^r = (1-\frac{1}{\log N})^{\log N + c_r} = \Omega(1)$ none of the registers is bad. ■

Claim 4.8 *Given an approximation x of $\sin\phi$ and an approximation y of $\cos\phi$ with additive error ϵ, we can find $\phi \bmod 2\pi$ up to an additive error of $O(\epsilon)$.*

Proof: Assume $y \geq 0$ and let $z = \frac{x}{1+y}$. A simple calculation shows that z is an estimate of $\frac{\sin\phi}{1+\cos\phi}$ up to an additive error of at most 4ϵ. The estimate on ϕ is $2\arctan z$. Since the absolute value of the differential of arctan is at most 1, this is an estimate of $2\arctan(\frac{\sin\phi}{1+\cos\phi}) = \phi$ with an additive error of at most 8ϵ. When $y < 0$ we compute an estimate of $2\text{arccot}(\frac{\sin\phi}{1-\cos\phi}) = \phi$. ∎

Lemma 4.9 *There exists a routine R_3 such that with probability exponentially close to 1, given any $q < \frac{N}{\log^{c_m} N}$ finds a value $q' \in \{q, \ldots, \log^{c_m} N \cdot q\}$ and an estimate x such that $x \in [q'd - \frac{N}{\log^{c_m+1} N}, q'd + \frac{N}{\log^{c_m+1} N}] \pmod N$.*

Proof: Assume we are given a q'-matching f. We call routines R_1 and R_2 $\log^{3c_m+4} N$ times. If the number of successful calls to one of the routines is less than $\log^{2c_m+3} N$, we fail. Otherwise, let $x \in [0,1]$ be the average of the successful calls to R_1 and $y \in [0,1]$ be the average of the successful calls to R_2. According to the Chernoff bound,

$$Pr[|x - (\frac{1}{2} - \frac{1}{2}\cos(2\pi q' \frac{d}{N}))| > \frac{1}{c_e \log^{c_m+1} N}] < 2e^{-2\log^{2c_m+3} N/(c_e{}^2 \log^{2c_m+2} N)}$$

which is exponentially low in $\log N$ for any constant $c_e > 0$. A similar bound holds for y. Hence, we can assume that $x' = 1 - 2x$ and $y' = 2y - 1$ are approximations of $\cos(2\pi q' \frac{d}{N})$ and of $\sin(2\pi q' \frac{d}{N})$ respectively up to an additive error of $\frac{2}{c_e \log^{c_m+1} N}$. According to Claim 4.8, this translates to an estimate of $q' \frac{d}{N} \bmod 1$ with an additive error of $\frac{1}{\log^{c_m+1} N}$ for c_e large enough.

By repeating the above procedure with all the matchings that appear in Lemma 4.5, we are guaranteed to find a good matching. According to Lemma 4.7, a call to routine R_1 or to routine R_2 with a good matching succeeds with probability at least $c_g \frac{1}{\log^{c_m} N}$ for a certain $c_g > 0$. The probability that none of $\log^{c_m+1} N$ calls to the subroutine succeeds is $(1 - c_g \frac{1}{\log^{c_m} N})^{\log^{c_m+1} N}$ which is exponentially small. Thus, for one of the matchings, with probability exponentially close to 1 we have $\log^{2c_m+3} N$ successful calls to routines R_1 and R_2 and routine R_3 is successful. ∎

We conclude the proof of Theorem 1.3 with a description of the algorithm for finding d. We begin by using routine R_3 with the value 1 to obtain an estimate x_1 and a value $\hat{q} \leq \log^{c_m} N$ such that $x_1 \in [d' - \frac{N}{\log^{c_m+1} N}, d' + \frac{N}{\log^{c_m+1} N}] \pmod N$ where d' denotes $(d\hat{q} \bmod N)$. In the following we find d' exactly by calling R_3 with multiples of $\hat{q}$. The algorithm works in stages. In stage i we have an estimate x_i and a value q_i. The invariant we maintain is $x_i \in [q_i d' - \frac{N}{\log^{c_m+1} N},\ q_i d' + \frac{N}{\log^{c_m+1} N}] \pmod{q_i N}$. We begin with x_1 as above and $q_1 = 1$. Assume that the invariant holds in stage i. We use routine R_3 with the value $2q_i\hat{q}$ to obtain an estimate x with a value $q' \in \{2q_i\hat{q}, 4q_i\hat{q}, \ldots 2\log^{c_m} N \cdot q_i\hat{q}\}$ such that $x \in [q_{i+1}d' - \frac{N}{\log^{c_m+1} N},\ q_{i+1}d' + \frac{N}{\log^{c_m+1} N}] \pmod N$ where $q_{i+1} = q'/\hat{q}$. Notice that our previous estimate x_i satisfies $\frac{q_{i+1}}{q_i} x_i \in [q_{i+1}d' - \frac{2N}{\log N}, q_{i+1}d' + \frac{2N}{\log N}] \pmod{q_{i+1}N}$. Since this range is much smaller than N, we can combine the estimate x on $(q_{i+1}d' \bmod N)$ and the estimate $\frac{q_{i+1}}{q_i} x_i$ on $(q_{i+1}d' \bmod q_{i+1}N)$ to obtain x_{i+1} such that $x_{i+1} \in [q_{i+1}d' - \frac{N}{\log^{c_m+1} N},\ q_{i+1}d' + \frac{N}{\log^{c_m+1} N}] \pmod{q_{i+1}N}$. The last stage is when $q_i \geq \frac{4N}{\log^{c_m+1} N}$. Then, d' can be found by rounding $\frac{x_i}{q_i}$ to the nearest integer. Given d' there are at most $\hat{q} \leq \log^{c_m} N$ possible values for q. Since this is only a polynomial number of options we can output one randomly.

5 Acknowledgements

I would like to thank Dorit Aharonov, Noga Alon, Andris Ambainis, Irit Dinur, Sean Hallgren, Alexei Kitaev, Hartmut Klauck, Ashwin Nayak, Cliff Smyth and Avi Wigderson for many helpful discussions and comments.

References

[1] M. Ajtai. Generating hard instances of lattice problems. In *Proc. 28th ACM Symp. on Theory of Computing*, pages 99–108, 1996. Available from ECCC at `http://www.uni-trier.de/eccc/`.

[2] M. Ajtai and C. Dwork. A public-key cryptosystem with worst-case/average-case equivalence. In *Proc. 29th ACM Symp. on Theory of Computing*, pages 284–293, 1997. Available from ECCC at `http://www.uni-trier.de/eccc/`.

[3] J-Y. Cai. A new transference theorem and applications to Ajtai's connection factor. *Electronic Colloquium on Computational Complexity (ECCC)*, 5, 1998.

[4] J-Y. Cai and A. Nerurkar. An improved worst-case to average-case connection for lattice problems. In *Proc. 38th IEEE Symp. on Found. of Comp. Science*, pages 468–477, 1997.

[5] D. Deutsch and R. Jozsa. Rapid solution of problems by quantum computation. *Proc. Roy. Soc. London Ser. A*, 439(1907):553–558, 1992.

[6] M. Ettinger and P. Høyer. On quantum algorithms for noncommutative hidden subgroups. *Adv. in Appl. Math.*, 25(3):239–251, 2000.

[7] K. Friedl, G. Ivanyos, F. Magniez, M. Santha, and P. Sen. Hidden translation and orbit coset in quantum computing. In *Proc. 35th ACM Symp. on Theory of Computing*, 2003.

[8] M. Grigni, L. J. Schulman, M. Vazirani, and U. V. Vazirani. Quantum mechanical algorithms for the nonabelian hidden subgroup problem. In *Proc. 33rd ACM Symp. on Theory of Computing*, pages 68–74, 2001.

[9] L. Grover and T. Rudolph. Creating superpositions that correspond to efficiently integrable probability distributions. In *quant-ph/0208112, http://xxx.lanl.gov*, 2002.

[10] L. K. Grover. A fast quantum mechanical algorithm for database search. In *Proc. 28th ACM Symp. on Theory of Computing*, pages 212–219, 1996.

[11] S. Hallgren. Polynomial-time quantum algorithms for Pell's equation and the principal ideal problem. In *Proc. 34th ACM Symp. on Theory of Computing*, pages 653–658, 2002.

[12] S. Hallgren, A. Russell, and A. Ta-Shma. Normal subgroup reconstruction and quantum computation using group representations. In *Proc. 32nd ACM Symp. on Theory of Computing*, pages 627–635, 2000.

[13] R. Impagliazzo and M. Naor. Efficient cryptographic schemes provably as secure as subset sum. *J. Cryptology*, 9(4):199–216, 1996.

[14] K. Johannes, S. Uwe, and T. Jacobo. *The graph isomorphism problem: its structural complexity.* Birkhäuser Boston Inc., 1993.

[15] R. Kannan, L. Lovász, and M. Simonovits. Random walks and an $O^*(n^5)$ volume algorithm for convex bodies. *Random Structures Algorithms*, 11(1):1–50, 1997.

[16] A. Y. Kitaev. Quantum measurements and the abelian stabilizer problem. In *quant-ph/9511026, http://xxx.lanl.gov*, 1995.

[17] G. Kuperberg. A subexponential-time quantum algorithm for the dihedral hidden subgroup problem. In *quant-ph/0302112, http://xxx.lanl.gov*, 2003.

[18] A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász. Factoring polynomials with rational coefficients. *Math. Ann.*, 261(4):515–534, 1982.

[19] D. Micciancio. Improved cryptographic hash functions with worst-case/average-case connection. In *Proc. 34th ACM Symp. on Theory of Computing*, 2002.

[20] D. Micciancio and S. Goldwasser. *Complexity of Lattice Problems: a cryptographic perspective*, volume 671 of *The Kluwer International Series in Engineering and Computer Science*. Kluwer Academic Publishers, Boston, Massachusetts, March 2002.

[21] M. A. Nielsen and I. L. Chuang. *Quantum computation and quantum information*. Cambridge University Press, Cambridge, 2000.

[22] O. Regev. New lattice based cryptographic constructions. In *Proc. 35th ACM Symp. on Theory of Computing*, San Diego, CA, June 2003.

[23] M. Rötteler and T. Beth. Polynomial-time solution to the hidden subgroup problem for a class of non-abelian groups. In *quant-ph/9812070, http://xxx.lanl.gov*, 1998.

[24] P. W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5):1484–1509, 1997.

[25] D. R. Simon. On the power of quantum computation. *SIAM J. Comput.*, 26(5):1474–1483, 1997.

[26] W. van Dam, S. Hallgren, and I. Lawrence. Quantum algorithms for hidden coset problems. In *Proc. 14th ACM-SIAM Symp. on Discrete Algorithms*, 2003.