On The Dihedral Coset Problem

Javad Doliskani*

Abstract

We propose an efficient quantum algorithm for a specific quantum state discrimination problem. An immediate corollary of our result is a polynomial time quantum algorithm for the Dihedral Coset Problem with a smooth modulus. This, in particular, implies that poly(n)-unique-SVP is in BQP.

1 Introduction

Let $\Sigma = \{0,1\}^n$ and let $\mathcal{X} = \mathbb{C}^\Sigma$ be the complex Euclidean space with basis Σ. Define two probability distributions $\mu_1, \mu_2 : \mathcal{S}(\mathcal{X}) \to [0,1]$ on the unit sphere $\mathcal{S}(\mathcal{X})$ as follows. The distribution μ_1 is defined by choosing (not necessarily independent) random $x, y \in \Sigma$ and outputting the state

$$\frac{1}{\sqrt{2}}(|0\rangle|x\rangle + |1\rangle|y\rangle). \tag{1}$$

The distribution μ_2 is defined by choosing random $(b, x) \in \{0,1\} \times \Sigma$ and outputting the state $|b\rangle|x\rangle$. In this note, we prove the following:

Theorem 1. *There is a quantum algorithm that distinguishes, with high probability, between the distributions μ_1 and μ_2 and runs in* poly(n) *operations.*

2 A Quantum Walk Algorithm

Our algorithm is based on a quantum walk algorithm introduced in [1]. Let p be an odd prime and d be a positive integer such that $2^n \leq p^d$. Let $f : \{0,1\}^n \hookrightarrow \mathbb{F}_p^d$ be any efficiently invertible injection of sets. Using f, we can assume that the outputs of the distributions μ_1 and μ_2 are in the space $\mathbb{C}^\Gamma$ where $\Gamma = \mathbb{F}_2 \times \mathbb{F}_p^d$.

For the sake of consistency, we follow the notations of [1]. Define $\Delta(x) = x_1^2 + \cdots + x_d^2$ for $x \in \mathbb{F}_p^d$. Let $\mathcal{S}_r$ be the sphere of radius r around zero in $\mathbb{F}_p^d$, that is, the points of $\mathbb{F}_p^d$ on the hypersurface $\Delta(x) - r = 0$. Note that we have $|\mathcal{S}_r| = \Theta(p^{d-1})$ [3, Section 6.2]. Assume for a moment that we could efficiently perform the (non-unitary) operation

$$\begin{array}{rccc} U: & \mathbb{C}^{\mathbb{F}_p^d} & \longrightarrow & \mathbb{C}^{\mathbb{F}_p^d} \\ & |x\rangle & \longmapsto & |\mathcal{S}_1 + x\rangle, \end{array} \tag{2}$$

where

$$|\mathcal{S}_1 + x\rangle = \frac{1}{\sqrt{|\mathcal{S}_1|}} \sum_{s \in \mathcal{S}_1} |s + x\rangle.$$

*Department of Computer Science, Ryerson University, (javad.doliskani@ryerson.ca).

Then we can distinguish between μ_1 and μ_2 as follows. Given an unknown distribution ρ that is one of the μ_1 or μ_2, obtain a sample state $|\psi\rangle$ from ρ. Then compute $(\mathbb{1} \otimes U)|\psi\rangle$ and measure the second register. Let us analyze the post-measurement state.

Case 1: $|\psi\rangle \in \mu_2$. The state $|\psi\rangle$ is of the form $|b\rangle|x\rangle$ for a random $(b, x) \in \mathbb{F}_2 \times \mathbb{F}_p^d$, so the post-measurement state is a random bit $|b\rangle$.

Case 2: $|\psi\rangle \in \mu_1$. The state $|\psi\rangle$ is of the form (1), so the outcome of the measurement is an element in $(\mathcal{S}_1 + x) \cap (\mathcal{S}_1 + y)$ with probability $\Theta(1/p)$. This is because for any $x, y \in \mathbb{F}_p^d$ we have $|(\mathcal{S}_1 + x) \cap (\mathcal{S}_1 + y)| \geq \Theta(p^{d-2})$ [3, Remark 6.28]. So the post-measurement state is $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ with probability at least $\Theta(1/p)$.

Next, we measure the remaining qubit in the Hadamard basis. If we observe $|+\rangle$ we return μ_1, otherwise we return μ_2. Therefore, given a poly(p) number of the states $|\psi\rangle$, we can distinguish, with overwhelming probability, between the two cases. The problem is that we do not know how to perform U efficiently. However, we show that we can replace U by a quantum walk unitary and still be able to perform the above steps and obtain a correct result! The following is adapted from Section 3 of [1].

Define the Cayley graph G with vertices the points in $\mathbb{F}_p^d$ and generating set the points in $\mathcal{S}_1$. The adjacency matrix of G is

$$A = \sum_{x \in \mathbb{F}_p^d} \sum_{s \in \mathcal{S}_1} |s + x\rangle\langle x|.$$

The eigenvectors of A, which are independent of the generating set $\mathcal{S}_1$, are $|\tilde{x}\rangle := \mathrm{F}_{p^d}|x\rangle$ where $x \in \mathbb{F}_p^d$ and F_{p^d} is the quantum Fourier transform over $\mathbb{F}_p^d$. Let $\omega_p = \exp(2\pi i/p)$. Then the eigenvalues of A corresponding to the eigenvectors $|\tilde{x}\rangle$ are

$$\lambda_x = \sum_{y \in \mathcal{S}_1} \omega_p^{\langle x, y\rangle} = \begin{cases} |\mathcal{S}_1| & x = 0, \\ \frac{G_1^d}{p} K_{\chi^d}\left(1, \frac{\Delta(x)}{4}\right) & \text{otherwise,} \end{cases}$$

where $G_1 = \sqrt{p}$ when $p = 1 \bmod 4$ and $G_1 = i\sqrt{p}$ when $p = 3 \bmod 4$, and where K_{χ^d} is the χ^d-twisted Kloosterman sum defined by

$$K_{\chi^d}(a, b) = \sum_{c \in \mathbb{F}_p} \chi^d(c) \omega_p^{ac + bc^{-1}}.$$

The eigenvalues λ_x can be computed in time poly(p). Define $\bar{A} = A - \lambda_0|\tilde{0}\rangle\langle\tilde{0}|$ so that $\|\bar{A}\| \leq 2\sqrt{p^{d-1}}$. Let $t = 1/\sqrt{p^{d-1}\log p}$, and let $x \in \mathbb{F}_p^d$. Define the operator U to be the continuous quantum walk with the Hamiltonian given by $\bar{A}$ for time t, i.e., $U = e^{i\bar{A}t}$. Then U leaves the subspace span$\{|x\rangle, |\mathcal{S}_0 + x\rangle, |\mathcal{S}_1 + x\rangle, \ldots, |\mathcal{S}_{p-1} + x\rangle\}$ invariant, so we can write

$$U|x\rangle = \alpha|x\rangle + \alpha_0|\mathcal{S}_0 + x\rangle + \alpha_1|\mathcal{S}_1 + x\rangle + \cdots + \alpha_{p-1}|\mathcal{S}_{p-1} + x\rangle.$$

Using the Taylor expansion of U we obtain

$$\alpha_1 = \langle \mathcal{S}_1 + x|U|x\rangle = it\sqrt{|\mathcal{S}_1|}(1 - O(p^{-1})) + O(\|\bar{A}^2\|t^2).$$

Note that α_1 is independent of the starting vertex x. If we measure $U|x\rangle$ in the vertex basis, we obtain an element of $\mathcal{S}_1 + x$ with probability

$$|\alpha_1|^2 = \frac{1}{\log p} + O(\log^{-3/2} p).$$

If we apply $\mathbb{1} \otimes U$ to the state

$$\frac{1}{\sqrt{2}}(|0\rangle|x\rangle + |1\rangle|y\rangle)$$

and measure the second register, the post-measurement state will be $(|0\rangle + |1\rangle)/\sqrt{2}$ with probability $\Theta(1/(p \log p))$. Therefore, replacing the non-unitary U in (2) with the unitary $U = e^{i\bar{A}t}$ in the above algorithm will only incur a $\Theta(1/\log p)$ loss in the distinguishing advantage.
Setting $p = O(n)$ in the above algorithm results in a $\text{poly}(d \log p) = \text{poly}(n)$ running time complexity, which is enough to prove Theorem 1. However, it is important to note that the number of samples required by the algorithm depends only on p, so the algorithm can be made sample-efficient by choosing small p. In any case, the running time complexity will remain $\text{poly}(n)$.

3 The Dihedral Coset Problem

Let N be a positive integer. A dihedral coset over the group $\mathbb{Z}_N$ is a state of the form

$$\frac{1}{\sqrt{2}}(|0\rangle|x\rangle + |1\rangle|x + s\rangle), \tag{3}$$

where $x \in \mathbb{Z}_N$ is uniformly random and $s \in \mathbb{Z}_N$ is fixed. Let $\Sigma = \mathbb{Z}_2 \times \mathbb{Z}_N$ and $\mathcal{X} = \mathbb{C}^{\Sigma}$. Define the distribution $\mu_s : \mathcal{S}(\mathcal{X}) \to [0, 1]$ by choosing $x \in \mathbb{Z}_N$ unifromly at random and outputting the state (3). The search Dihedral Coset Problem over $\mathbb{Z}_N$, denoted by DCP_N, is the problem of recovering s given outputs from μ_s. The decision-DCP_N is the problem of distinguishing between μ_s and the distribution $\mu : \mathcal{S}(\mathcal{X}) \to [0, 1]$ defined by choosing $(b, x) \in \Sigma$ uniformly at random and outputting the state $|b\rangle|x\rangle$.

Corollary 2. *There is a quantum algorithm for decision-DCP_N that tuns in $\text{poly}(\log N)$ operations.*

Proof. Set $\mu_1 = \mu_s$ and $\mu_2 = \mu$ in Theorem 1. □

When the modulus N has $\text{poly}(\log N)$-bounded prime factors, the search-DCP_N can be reduced to the decision-DCP_N in time $\text{poly}(\log N)$ [2]. Therefore, we have

Corollary 3. *For a modulus N with $\text{poly}(\log N)$-bounded prime factors, there is a quantum algorithm for search-DCP_N that runs in $\text{poly}(\log N)$ operations.*

It was shown in [4] that a polynomial time quantum algorithm for DCP_N with $N = 2^{\Theta(n^2)}$ implies a polynomial time quantum algorithm for $\text{poly}(n)$-unique-SVP. Therefore, we have

Corollary 4. *There is a polynomial time quantum algorithm for $\text{poly}(n)$-unique-SVP.*

References

[1] Andrew M Childs, Leonard J Schulman, and Umesh V Vazirani. Quantum algorithms for hidden nonlinear structures. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS'07)*, pages 395–404. IEEE, 2007.

[2] Javad Doliskani. Efficient quantum public-key encryption from learning with errors. 2020. https://eprint.iacr.org/2020/1557.

[3] Rudolf Lidl and Harald Niederreiter. *Finite fields*. Number 20. Cambridge university press, 1997.

[4] Oded Regev. Quantum computation and lattice problems. *SIAM Journal on Computing*, 33(3):738–760, 2004.