Research Article

Imin Chen* and David Sun

The dihedral hidden subgroup problem

https://doi.org/10.1515/jmc-2022-0029
received October 06, 2022; accepted August 23, 2023

Abstract: The hidden subgroup problem (HSP) is a cornerstone problem in quantum computing, which captures many problems of interest and provides a standard framework algorithm for their study based on Fourier sampling, one class of techniques known to provide quantum advantage, and which succeeds for some groups but not others. The quantum hardness of the HSP problem for the dihedral group is a critical question for post-quantum cryptosystems based on learning with errors and also appears in subexponential algorithms for constructing isogenies between elliptic curves over a finite field. In this article, we give an updated overview of the dihedral hidden subgroup problem as approached by the "standard" quantum algorithm for HSP on finite groups, detailing the obstructions for strong Fourier sampling to succeed and summarizing other known approaches and results. In our treatment, we "contrast and compare" as much as possible the cyclic and dihedral cases, with a view to determining bounds for the success probability of a quantum algorithm that uses m coset samples to solve the HSP on these groups. In the last sections, we prove a number of no-go results for the dihedral coset problem (DCP), motivated by a connection between DCP and cloning of quantum states. The proofs of these no-go results are then adapted to give nontrivial upper bounds on the success probability of a quantum algorithm that uses m coset samples to solve DCP.

Keywords: quantum computation, hidden subgroup problem

MSC 2020: 81P94, 68Q12, 20C05, 14H52

1 Introduction

Let G be a finite group and H a hidden subgroup of G. A function $f : G \to S$, where S is a known finite set, which is constant on left H-cosets and takes distinct values on distinct left H-cosets, is called a separating function for the subgroup H.

The hidden subgroup problem (HSP) is the problem of finding generators for the hidden subgroup H, given access to evaluations of a separating function f for H. The HSP can be solved in polynomial time using a quantum computer when G is an abelian group [1,2] and has been extensively studied for many classes of finite groups [3,4].

Many problems can be cast in terms of the HSP and there is a "natural" standard quantum algorithm based on the quantum Fourier transform, which is typically used to study the HSP. For instance, Shor's integer factorization algorithm [5] can be described in terms of the HSP for cyclic groups [6], and in this case, the standard algorithm succeeds in yielding an efficient quantum algorithm. Another example is the HSP on the symmetric group, which can be used to solve the graph isomorphism problem [7–10], but here the standard algorithm fails to provide an efficient quantum algorithm [11].

A polynomial time quantum algorithm for solving the HSP on dihedral groups would imply a polynomial time quantum algorithm to solve certain hard lattice problems that are considered intractable using classical

*** Corresponding author: Imin Chen**, Department of Mathematics, Simon Fraser University Burnaby, BC V5A 1S6, Canada, e-mail: ichen@sfu.ca
David Sun: Department of Mathematics, Simon Fraser University Burnaby, BC V5A 1S6, Canada, e-mail: david_sun_2@sfu.ca

Open Access. © 2024 the author(s), published by De Gruyter. This work is licensed under the Creative Commons Attribution 4.0 International License.

computers [12]. Though the dihedral group is one of the simplest non-abelian groups, from the point of view of the HSP, it has remained a difficult case in terms of definitive results about its hardness. The best known quantum algorithms for the dihedral hidden subgroup problem (DHSP) are currently subexponential [13–16].

A problem closely related to DHSP is the dihedral coset problem (DCP), which is the problem of determining a hidden subgroup H of the dihedral group D_N from uniform coset samples obtained from evaluations of the separating function for H.

The latter subexponential algorithms have applications to constructing isogenies between elliptic curves over a finite field [17,18], though we note that the recent breakthroughs in [19–21] have changed the landscape on this problem in the presence of torsion point information.

In the study by Brakersk et al. [22], it is shown that the learning with errors (LWE) problem is the quantum polynomial time equivalent to an extrapolated version of the dihedral coset problem (EDCP). The LWE problem forms the basis for many proposed post-quantum key exchanges; therefore, the quantum hardness of the HSP for groups like the dihedral group becomes a critical question.

In the first sections of this article, we review the standard HSP algorithm as it applies to the dihedral groups D_N and detail the obstructions for this algorithm to succeed in this case. We also provide an overview describing other approaches to the HSP for dihedral groups, such as optimal measurements and its relations to the subset sum problem, complementing previous surveys of known results on DHSP [4,23,24] (see also [25] for a self-contained account). In our treatment, we "contrast and compare" as much as possible the cyclic and dihedral cases, with a view to determining bounds for the success probability of a quantum algorithm, which uses m coset samples to solve the HSP on these groups.

In the last sections of this article, we prove a number of a no-go theorems for DCP. The results yield an upper bound on the success probability of any quantum algorithm, which uses a unitary operation and then one measurement to determine the angle a of a hidden reflection in the dihedral group. Viewed in terms of positive operator-valued measurements (POVM), this gives a nontrivial upper bound on the success probability of the optimal measurement using m coset samples to solve DCP in the case when the density $\nu = m/\log_2 N \geq 1$ and the order of the dihedral group is $2N$.

Finally, we describe a connection between DCP and cloning of quantum states that helped motivate the proofs of the no-go results for DCP.

2 Quantum Fourier transform (QFT) for finite groups

Let G be a finite group and $\hat{G}$ denote a complete set of representatives for the isomorphism classes of irreducible representations of G over $\mathbb{C}$. For a representation $\rho \in \hat{G}$, let d_ρ be the dimension of ρ. Recall the QFT on G is defined as the linear transformation:

$$\begin{aligned} &F_G : V \to \hat{V}, \\ &F_G \coloneqq \sum_{g\in G}\sum_{\rho\in\hat{G}}\sum_{i,j=1}^{d_\rho}\sqrt{\frac{d_\rho}{|G|}}\,\rho(g)_{i,j}|\rho,i,j\rangle\langle g|, \end{aligned} \tag{2.1}$$

where V is the $\mathbb{C}$-vector space generated by $|g\rangle$, $g \in G$, and $\hat{V}$ is the $\mathbb{C}$-vector space generated by $|\rho,i,j\rangle$, $\rho \in \hat{G}, 1 \leq i,j \leq d_\rho$. Picking an isomorphism $V \cong \hat{V}$, it is a unitary operator that can be efficiently approximated using quantum circuits for many finite groups: abelian [26], meta-cyclic [10], symmetric group [7].

2.1 Cyclic group case

Suppose that $G = C_N \cong \mathbb{Z}/N\mathbb{Z}$ is the cyclic group of order N. There are N irreducible representations that are one-dimensional and given by:

$$\begin{aligned} \mathbb{Z}/N\mathbb{Z} &\to \mathbb{C}^\times, \\ t &\mapsto \zeta_N^t, \end{aligned}$$

where ζ_N is a choice of Nth root of unity.

2.2 Dihedral group case

Suppose that $G = D_N$ is the dihedral group of order $2n$, which can be presented as:

$$D_N = \langle x, y : x^n = e, y^2 = e, yxy^{-1} = x^{-1} \rangle.$$

If n is even, there are four one-dimensional representations given by:

$$\phi_{u,v} : x \mapsto (-1)^u, y \mapsto (-1)^v, \tag{2.2}$$

where $u, v \in \mathbb{Z}/2\mathbb{Z}$. These are pullbacks of the four one-dimensional representations of $D_N/\langle x^2 \rangle \cong C_2 \times C_2$ under the quotient homomorphism $D_N \to D_N/\langle x^2 \rangle$, where C_m denotes the cyclic group of order m.

If n is odd, there are two one-dimensional representations given by $\phi_{0,v}$ where $v \in \mathbb{Z}/2\mathbb{Z}$. These are pullbacks of the two one-dimensional representations of $D_N/\langle x \rangle \cong C_2$ under the quotient homomorphism $D_N \to D_N/\langle x \rangle$.

There are $\lfloor \frac{n-1}{2} \rfloor$ irreducible representations of dimension 2 given by:

$$\begin{aligned} \rho_k : D_N &\to \mathrm{GL}_2(\mathbb{C}), \\ x &\mapsto \begin{pmatrix} \omega_N^k & 0 \\ 0 & \omega_N^{-k} \end{pmatrix}, \\ y &\mapsto \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \end{aligned} \tag{2.3}$$

for $0 < k < \frac{n}{2}$, where $\omega_N = e^{2\pi i/n}$. These are the induction of the representation $\psi_k : C_n \to \mathbb{C}^\times$ given by $\psi_k(x) = \omega_N^k$ from C_n to D_N.

The representations $\phi_{u,v}$ and ρ_k form the complete list of irreducible representations of D_N up to isomorphism.

3 Standard HSP algorithm

In the standard algorithm for finding hidden subgroups from a separating function, we perform the following steps:

We form the state

$$\frac{1}{\sqrt{|G|}} \sum_{g \in G} |g\rangle |f(g)\rangle, \tag{3.1}$$

where $f : G \to S$ is the given separating function.

This can be achieved by starting with the state $|e_G\rangle|0\rangle$, where e_G is the identity element of G, then performing the following computations:

$$\begin{aligned} |e_G\rangle|0\rangle &\mapsto \frac{1}{\sqrt{|G|}} \sum_{g \in G} |g\rangle|0\rangle, \quad \text{apply the QFT over } G \leftrightarrow \mathbb{Z}/|G|\mathbb{Z}, e_G \leftrightarrow 0 \text{ to first register.} \\ &\mapsto \frac{1}{\sqrt{|G|}} \sum_{g \in G} |g\rangle|f(g)\rangle, \quad \text{compute } f \text{ into second register.} \end{aligned} \tag{3.2}$$

Measuring the second register and discarding it, we obtain a state of the form:

$$\frac{1}{\sqrt{|H|}} \sum_{h \in H} |ch\rangle, \tag{3.3}$$

where $c \in G$.

We apply the QFT to the aforementioned state to obtain

$$\sum_{\rho \in \hat{G}} \sqrt{\frac{d_\rho}{|G||H|}} \sum_{i,j=1}^{d_\rho} \sum_{h \in H} \rho(ch)_{i,j} |\rho, i, j\rangle. \tag{3.4}$$

In the case of G being an abelian group, measuring ρ gives sufficient information to determine H efficiently after running this process repeatedly and using post-processing [1].

3.1 Cyclic group case

Fix an integer $N > 1$. Let S be a finite set, and $G \coloneqq (\mathbb{Z}/N\mathbb{Z}, +)$. Suppose that we have a function $f : G \to S$, which separates a subgroup $H \subset G$, where $H = \langle d \rangle$. Let $M \coloneqq \#H$. Assume that we have a quantum machine capable of computing the unitary transformation on two registers $U_f : |x\rangle|y\rangle \to |x\rangle|f(x) \oplus y\rangle$ (recall that we can take $|x\rangle|y\rangle$ as $|x\rangle \otimes |y\rangle$).

Suppose that we do not know M, d, nor H and we only know G and have a machine computing f. We want to determine a generating set for H, calling the "black-box" function f as few times as possible.

Let F_N be the QFT for the cyclic group G. Explicitly, this is an operator on a register with $n \geq \log_2 N$ qubits given by:

$$F_N \coloneqq \frac{1}{\sqrt{N}} \sum_{j,k=0}^{N-1} \exp\left(\frac{2\pi i j k}{N}\right) |k\rangle\langle j|.$$

The F_N is a unitary transformation. If we let $\omega \coloneqq \exp\left(\frac{2\pi i}{N}\right)$ be the primitive Nth root of unity, then

$$F_N = \frac{1}{\sqrt{N}} \begin{pmatrix} 1 & 1 & 1 & \cdots & 1 \\ 1 & \omega & \omega^2 & \cdots & \omega^{N-1} \\ 1 & \omega^2 & \omega^4 & \cdots & \omega^{2(N-1)} \\ \vdots & \vdots & \vdots & & \vdots \\ 1 & \omega^{N-1} & \omega^{2(N-1)} & \cdots & \omega^{(N-1)(N-1)} \end{pmatrix}.$$

One can check that $F_N \cdot F_N^* = I_N$, where I_N is the $N \times N$ identity matrix.

We map $G = \{0, 1, \ldots, N-1\}$ onto the basis of the quantum state $\{|0\rangle, |1\rangle, \ldots, |N-1\rangle\}$. Suppose that the hidden subgroup is given by $H = \{|0\rangle, |d\rangle, |2d\rangle, \ldots, |(M-1)d\rangle\}$.

Computing on two registers:

$$\begin{aligned} |0\rangle|0\rangle &\xrightarrow{F_N \text{ on 1st}} \frac{1}{\sqrt{N}} \sum_{j=0}^{N-1} |j\rangle|0\rangle \\ &\xrightarrow{\text{apply } f} \frac{1}{\sqrt{N}} \sum_{j=0}^{N-1} |j\rangle|f(j)\rangle. \end{aligned}$$

Note that we put $|f(j)\rangle$ inside the sum since tensor product is distributive. Measuring in $|f(j_0)\rangle$ on the second register for some $0 \leq j_0 \leq N-1$ collapses our state, leaving only those values $g \in G$ such that $f(g) = f(j_0)$ in the first register. Since f separates cosets of H, we obtain (for simplicity, we now drop our second register that remains $|f(j_0)\rangle$):

$$\begin{aligned}
&\xrightarrow{\text{measure}} \frac{1}{\sqrt{M}} \sum_{h\in H} |j_0 + h\rangle = \frac{1}{\sqrt{M}} \sum_{s=0}^{M-1} |j_0 + sd\rangle \\
&\xrightarrow{\text{apply } F_N} \frac{1}{\sqrt{M}} \sum_{s=0}^{M-1} \frac{1}{\sqrt{N}} \sum_{k=0}^{N-1} \exp\left(\frac{2\pi i(j_0 + sd)k}{N}\right)|k\rangle \\
&= \frac{1}{\sqrt{MN}} \sum_{k=0}^{N-1} \exp\left(\frac{2\pi i j_0 k}{N}\right)|k\rangle \sum_{s=0}^{M-1} \exp\left(\frac{2\pi i s d k}{N}\right) \\
&= \frac{1}{\sqrt{d}} \sum_{t=0}^{d-1} \exp\left(\frac{2\pi i j_0 tM}{N}\right)|tM\rangle,
\end{aligned}$$

using the fact that

$$\sum_{s=0}^{M-1} \exp\left(\frac{2\pi i s d k}{N}\right) = \sum_{s=0}^{M-1} \exp\left(\frac{2\pi i k}{M}\right)^s = \begin{cases} 0, & \text{if } M \nmid k, \\ M, & \text{if } M | k, \end{cases}$$

for $0 \le k \le N-1$ and that $\frac{M}{N} = \frac{1}{d}$.

Now, measurement at this point gives a multiple of M in $\{0, M, \ldots, (d-1)M\}$ with uniform probability. We repeat this whole process many times to obtain a collection of multiples of M and take the greatest common divisor (GCD) to obtain M with high probability.

To estimate how many trials $m \ge 2$ we need, suppose that we have $t_1, \ldots, t_m \in \{0, 1, \ldots, d-1\}$. We want to estimate the probability that $\gcd(t_1, \ldots, t_m) = 1$; in particular, we have the lower bound:

$$\mathbf{P}(\gcd(t_1, \ldots, t_m) = 1) \ge \zeta(m)^{-1} + O(\log d/d), \tag{3.5}$$

where $\zeta(s)$ is the Riemann zeta function by Nymann [27]. Thus, a few runs of this algorithm determine H with high probability for any N and "most" d.

Lemma 3.6. *We have that*

$$\zeta(m)^{-1} > 1 - 3 \cdot 2^{-m},$$

for every $m \ge 2$.

Proof. We first recall that

$$\zeta(s) = \sum_{n=1}^{\infty} n^{-s} \quad \text{and} \quad \zeta(s)^{-1} = \sum_{n=1}^{\infty} \frac{\mu(n)}{n^s},$$

for real $s \ge 2$, where μ is the Möbius function. Then,

$$\begin{aligned}
\frac{1 - \zeta(s)^{-1}}{2^{-s}} &= \sum_{n\ge 2} \frac{-\mu(n)}{(n/2)^s} \\
&\le \sum_{n\ge 2} \frac{1}{(n/2)^2} \\
&= 4(\zeta(2) - 1) < 3.
\end{aligned}$$

□

We may therefore view the standard algorithm for HSP on the cyclic group G as producing a quantum state of the form:

$$\sum_{t_1, \ldots, t_m} \alpha_{t_1, \ldots, t_m} |t_1 M\rangle \ldots |t_m M\rangle. \tag{3.7}$$

We may compute the greatest common divisor of the aforementioned registers into a blank register:

$$\sum_{t_1, \ldots, t_m} \alpha_{t_1, \ldots, t_m} |t_1 M\rangle \ldots |t_m M\rangle |0\rangle \mapsto \sum_{t_1, \ldots, t_m} \alpha_{t_1, \ldots, t_m} |t_1 M\rangle \ldots |t_m M\rangle |\gcd(t_1 M, \ldots t_m M)\rangle. \tag{3.8}$$

Thus, the standard HSP algorithm for G can be viewed as a unitary operation of the form:

$$|A\rangle|\psi_d^1\rangle\cdots|\psi_d^m\rangle|0\rangle \mapsto \sum_e |\Sigma_e\rangle|N/e\rangle \mapsto \sum_e |\Sigma_e\rangle|e\rangle, \tag{3.9}$$

satisfying

$$|\Sigma_d|^2 \geq \lambda(m, d),$$

for every m. We remark the second map sending $e \mapsto N/e$ in the last register is unitary (if $e \nmid N$, the map leaves e alone).

Remark 3.10. Assume that for any guess for d, there is a quantum circuit that can decide whether d is correct. For a fixed m, we can improve the success probability above by the following method. Let us, instead, consider the probability of achieving a multiple kM of M for $1 \leq k \leq C$ for some $C \in \mathbb{N}$. For the given guess of M and hence for d, we can check if it is the correct value, and if not, adjust it to the correct value because we know the true value is d/k for some $1 \leq k \leq C$ and d/k being an integer. This increases the success parameter for a fixed number of samples $m \geq 2$. For instance, if $m = 2$ and $C = 10$, then the success parameter improves from ≈0.6079 to ≈0.9892.

The aforementioned example motivates the next definition.

Definition 3.11. Let $\mathcal{I}_d = \{|\psi_d\rangle\}$ be a collection of possible input states depending on the parameter d. The problem of determining d from a list of m samples in $\mathcal{I}_d$ is *unitarily solvable with success parameter* $\lambda(m, d)$ if there is a unitary operator that has the effect:

$$|A\rangle|\psi_d^1\rangle\cdots|\psi_d^m\rangle|0\rangle \mapsto |\Sigma_d(\psi_d)\rangle|d\rangle + \sum_{e\neq d} |\Sigma_e(\psi_d)\rangle|e\rangle, \tag{3.12}$$

where

$$|\Sigma_d(\psi_d)|^2 \geq \lambda(m, d),$$

for every m and d.

We may view (3.12) as computing a main term:

$$|\Sigma_d(\psi_d)\rangle|d\rangle, \tag{3.13}$$

with error term

$$\sum_{e\neq d} |\Sigma_e(\psi_d)\rangle|e\rangle. \tag{3.14}$$

The next theorem is stated for completeness and for later comparison to the case of DCP. It summarizes the well-known standard algorithm for HSP on a finite cyclic group in terms of the aforementioned definitions.

Theorem 3.15. *The problem of determining a generator for a hidden subgroup of a finite cyclic group, given a list of m HSP coset samples, is unitarily solvable with success parameter* $\zeta(m)^{-1} + O(\log d/d)$*, where* $\zeta(s)$ *is the Riemann zeta function.*

Remark 3.16. Here, $M = N/d$, so we may view the standard quantum algorithm as producing uniform samples in $H^\perp = \langle d\rangle^\perp = \langle N/d\rangle$, where $H^\perp = \cap_{h\in H} \ker\chi_h = \ker\chi_d$ (see [28, §4.1, §6.1] for more details). For a general finite abelian group G, the uniform samples in $H^\perp$ from the standard quantum algorithm are used to determine H by a classical probabilistic algorithm.

3.2 Dihedral group case

In [29], it is shown that the HSP for $G = D_N$ for a general subgroup H is reduced to the case of a single reflection subgroup $H = H_a$.

For $H = H_a = \langle yx^a \rangle$, the probability of obtaining $|\rho, i, j\rangle$ is $\frac{1}{|G|}$ when $d_\rho = 2$, which does not allow us to distinguish the groups H_a. Explicitly, in the complex basis (2.3):

If $\rho = \rho_k$, then

$$\sum_{h\in H} \rho(x^a h) = \begin{pmatrix} \omega_N^{ak} & \omega_N^{-(a-\alpha)k} \\ \omega_N^{(a-\alpha)k} & \omega_N^{-ak} \end{pmatrix}, \tag{3.17}$$

$$\sum_{h\in H} \rho(yx^a h) = \begin{pmatrix} \omega_N^{(a-\alpha)k} & \omega_N^{-ak} \\ \omega_N^{ak} & \omega_N^{-(a-\alpha)k} \end{pmatrix}. \tag{3.18}$$

If $\rho = \phi_{u,v}$, then

$$\sum_{h\in H} \rho(x^a h) = (-1)^{au} + (-1)^{v+(a-\alpha)u} = (-1)^{au}(-1 + (-1)^{v+au}), \tag{3.19}$$

$$\sum_{h\in H} \rho(yx^a h) = (-1)^{(a-\alpha)u} + (-1)^{v+au} = (-1)^{(a-\alpha)u}(-1 + (-1)^{v+au}). \tag{3.20}$$

If one changes to the real basis, we obtain a probability distribution dependent on a, but it is very flat, making it hard to distinguish the subgroups H_a.

More generally, in order for the QFT to be an unitary operator, we require that $\rho_k(g)$ be unitary for every k and $g \in D_N \Rightarrow |\rho_k(g)_{i,j}| \le 1$ for $1 \le i, j \le 2$. In particular, for any set of two-dimensional irreducible representations ρ_k, we have that

$$\begin{aligned} \mathbf{P}(\rho_k, i, j) &= \frac{1}{n}|(\rho_k(cyx^a) + \rho_k(c))_{i,j}|^2 \\ &\le \frac{1}{N}(|\rho_k(cyx^a)_{i,j}| + |\rho_k(c)_{i,j}|)^2 \\ &\le \frac{4}{N}, \end{aligned} \tag{3.21}$$

where $\mathbf{P}(\rho_k, i, j)$ is the probability of observing the state $|\rho_k, i, j\rangle$. Although the choice of basis may result in probability distributions of states that depend on a, if N is very large, the aforementioned inequalities show that the probabilities will always be very flat.

In the study by Moore and Russell [30], it is shown that the POVM to determine a from a single DCP sample exists and is given by the the pretty good measurement (PGM), and this has success probability:

$$\mathbf{P}_{\text{success}} = \frac{2}{N}\left(1 - \frac{1}{2N}\right). \tag{3.22}$$

Theorem 3.23. *The standard algorithm for DHSP cannot implement the optimal measurement using one coset sample.*

Proof. This follows because (3.21) and (3.22) are incompatible. □

3.3 Dihedral coset sampling

In the standard HSP algorithm, after the first step, we are left with random coset samples as in (3.3). In the case of $G = D_N$, the dihedral group of order n, and $H = H_a = \langle yx^a \rangle$, this is explicitly of the form:

$$\begin{aligned}\frac{1}{\sqrt{|H|}} \sum_{h \in H} |ch\rangle &= \frac{1}{\sqrt{2}}(|c\rangle + |cyx^a\rangle) \\ &= \frac{1}{\sqrt{2}}(|y^\beta x^\alpha\rangle + |y^\beta x^\alpha y x^a\rangle) \\ &= \frac{1}{\sqrt{2}}(|y^\beta x^\alpha\rangle + |y^{\beta+1} x^{\alpha-a}\rangle) \\ &= \begin{cases} \frac{1}{\sqrt{2}}(|x^\alpha\rangle + |yx^{\alpha-a}\rangle), & \text{if } \beta = 0 \\ \frac{1}{\sqrt{2}}(|x^{\alpha-a}\rangle + |yx^\alpha\rangle), & \text{if } \beta = 1 \end{cases}\end{aligned} \tag{3.24}$$

where $c = y^\beta x^\alpha$.

Remark 3.25. The second case is reduced to the first by the transformation $\alpha \to a - \alpha$ if this transformation leaves the distribution of α invariant.

Given samples of the form

$$\psi_a = \psi_{a;\alpha} = \frac{1}{\sqrt{2}}(|x^\alpha\rangle + |yx^{\alpha-a}\rangle), \tag{3.26}$$

the DCP is the problem of finding generators for the hidden subgroup $H = H_a$. The states $\psi_a = \psi_{a;\alpha}$ are called the DCP samples for a.

For HSP samples produced from the standard algorithm, where α is from the uniform distribution, we may view HSP samples as DCP samples by Remark 3.25.

Remark 3.27. We can encode a DCP sample $\psi_{a;\alpha}$ as:

$$\frac{1}{\sqrt{2}}(|0\rangle|\alpha\rangle + |1\rangle|a - \alpha\rangle).$$

Using the fact that $yx^\alpha = x^{-\alpha}y$, this can be encoded (after negating a) as:

$$\frac{1}{\sqrt{2}}(|0\rangle|\alpha\rangle + |1\rangle|a + \alpha\rangle),$$

which is another commonly used form used in the literature, especially in the context of the "hidden shift problem".

3.4 Generalizations of DCP

Let A be a finite abelian group (written multiplicatively) with identity e. Let $f_0, f_1 : A \to S$ be the injective functions to a finite set S and such that $f_1(x) = f_0(xs)$ for all $x \in A$ and some $s \in A$. The hidden shift problem is the problem of determining the hidden shift $s \in A$ from evaluations of the functions f_0 and f_1. We refer the reader to previous studies [17, §5] for a discussion of the hidden shift problem without the condition of injectivity.

Let $G = C_2 \ltimes A$, where $C_2 = \langle \tau \rangle$ acts from the left on A by inversion. In the resulting semi-direct product, let $y = (\tau, e)$. Then, for elements in the normal subgroup $\tilde{A}$ of G corresponding to A, i.e., elements of the form $\tilde{a} = (e, a)$, we have that $y\tilde{a}y^{-1} = \tilde{a}^{-1}$, and every element in G can be written in the form $y\tilde{a}$ or $\tilde{a}$.

The HSP for subgroups of the form:

$$H = \langle (\tau, s) \rangle = \{(e, e), (\tau, s)\} \tag{3.28}$$

is equivalent to the hidden shift problem with shift $s \in A$. Indeed, an element in $\mathbb{Z}[G]$, which is constant on left H-cosets, is a linear combination of sums of the form:

$$\sum_{h \in H} y^\beta \tilde{a} h = y^\beta \tilde{a} + y^{\beta+1} \tilde{s} \tilde{a}^{-1}, \tag{3.29}$$

which are (up to scaling) left H-coset samples that result from applying the standard algorithm to the group G to find hidden subgroups of the form (3.28). After mapping S injectively into $\mathbb{Z}$, the function $f_0 + f_1$ corresponds to an element in $\mathbb{Z}[G]$, which is distinctly constant on left H-cosets so it is a separating function for H. Hence, a solution to the HSP on G implies a solution to the hidden shift problem on A. For a complete proof of equivalence, see [13, Proposition 6.1].

In the application to constructing isogenies between elliptic curves over finite fields with identical endomorphism ring (i.e., horizontal isogenies), the hidden shift problem is applied in the following way [17]. Let $\mathcal{O}$ be an order in an imaginary quadratic field K, and consider the set

$$\mathrm{Ell}_{q,n}(\mathcal{O}) = \{\text{elliptic curves } E/\mathbb{F}_q, \text{ with } \#E(\mathbb{F}_q) = n \text{ and } \mathrm{End}_{\mathbb{F}_q}(E) \cong \mathcal{O}\}/\cong_{\mathbb{F}_q}. \tag{3.30}$$

The class group $\mathrm{Cl}(\mathcal{O})$ of the order $\mathcal{O}$ acts freely on $\mathrm{Ell}_{q,n}(\mathcal{O})$ assuming this set is non-empty for the given q a power of a prime p, with one orbit if p is not inert in $\mathcal{O}$ and two orbits otherwise (see [31, Theorem 4.5]). Given $[\mathfrak{b}] \in \mathrm{Cl}(\mathcal{O})$ and $E \in \mathrm{Ell}_{q,n}(\mathcal{O})$ denote the action by $[\mathfrak{b}] \star E$. Now, suppose $E_1 = [\mathfrak{s}] \star E_0$ for some hidden $\mathfrak{s} \in \mathrm{Cl}(\mathcal{O})$. For $i = 0, 1$, define $f_i([\mathfrak{b}]) = [\mathfrak{b}] \star E_i$. Then, it holds that $f_1([\mathfrak{b}]) = f_0([\mathfrak{b}][\mathfrak{s}])$ and both f_0 and f_1 are the injective functions to $S = \mathrm{Ell}_{q,n}(\mathcal{O})$. A solution to the hidden shift problem on $\mathrm{Cl}(\mathcal{O})$ allows us to recover the hidden shift $\mathfrak{s}$. The subexponential algorithm of Childs et al. [17] consists of applying a subexponential algorithm for the hidden shift problem to find $\mathfrak{s}$ and applying a subexponential algorithm to compute the $\star$ operator.

Remark 3.31. In the study by Childs et al. [17], the set in (3.30) is taken up to isomorphism over $\overline{\mathbb{F}}_q$. We use isomorphism over $\mathbb{F}_q$ because this is the form used in the reference we cite for the class group action.

In the EDCP, one considers the infinite abelian group:

$$G = \mathbb{Z} \times (\mathbb{Z}/N\mathbb{Z})^n,$$

where $n \in \mathbb{N}$, and hidden subgroups of the form:

$$H = \langle (1, s) \rangle. \tag{3.32}$$

Let $y = (1, 0) \in G$ and $f_\beta : H \to \mathbb{C}$ be a weight function satisfying

$$\sum_{h \in H} |f_\beta(h)|^2 = 1. \tag{3.33}$$

A weighted left H-coset sample has the form:

$$\begin{aligned} y^\beta x^\alpha \sum_{h \in H} f_\beta(h) h &= y^\beta x^\alpha \sum_{j \in \mathbb{Z}} f_\beta(j) y^j x^{sj} \\ &= \sum_{j \in \mathbb{Z}} f_\beta(j) y^{\beta+j} x^{\alpha+sj} \\ &= \sum_{j \in \mathbb{Z}} f(j) y^j x^{\alpha - s\beta + sj}, \end{aligned}$$

where we assume $f_\beta(j-\beta) = f(j)$ and we use the multiplicative notation to make the comparison with the hidden shift and dihedral cases more clearly. Here, x^α is the multiplicative notation for the element corresponding to $\alpha = (\alpha_1, \ldots, \alpha_n) \in (\mathbb{Z}/N\mathbb{Z})^n$.

If the transformation $\alpha \to \alpha + s\beta$ leaves the distribution of α invariant, we are left with samples of the form:

$$\sum_{j\in\mathbb{Z}} f(j) y^j x^{\alpha+sj}. \tag{3.34}$$

The problem of recovering s from such samples (EDCP) is shown to be equivalent to LWE up to polynomial loss in parameters [22].

Remark 3.35. When f is the normalized indicator function of $\{0,1\}$, an extrapolated dihedral coset sample reverts to a dihedral coset sample. Also, the study by Brakerski et al. [22] imposes a stronger condition than (3.33).

4 Other approaches to DHSP and DCP

4.1 Subexponential algorithms

The first row of (3.17) can be encoded as:

$$\frac{1}{\sqrt{2N}}\sum_k (\omega_N^{ak}|k\rangle|0\rangle + \omega_N^{(a-\alpha)k}|k\rangle|1\rangle) = \frac{1}{\sqrt{N}}\sum_k \omega_N^{ak} \otimes \frac{1}{\sqrt{2}}(|0\rangle + \omega_N^{ak}|1\rangle).$$

Measuring the first register yields the samples of the form:

$$|\Psi_k\rangle = \frac{1}{\sqrt{2}}(|0\rangle + \omega_N^{ak}|1\rangle), \tag{4.1}$$

where k is known from the measurement.

Let $N = 2^t$ for simplicity and $m = \lceil \sqrt{t-1} \rceil$. The idea behind the subexponential algorithm in the study by Kuperberg [13] is to combine states of the form (4.1). In particular, we see that

$$|\Psi_p\rangle|\Psi_q\rangle = \frac{1}{\sqrt{2}}(|\Psi_{p+q}\rangle|0\rangle + \omega_N^{aq}|\Psi_{p-q}\rangle|1\rangle). \tag{4.2}$$

If p and q have the same mj least significant bits, then $p \pm q$ strictly increases the number of least significant bits p and q share.

With sufficiently many samples of the form Ψ_p that have mj common least significant bits, it is shown in the study by Kuperberg [13] that combining the states as in (4.2) produces enough states with $m(j+1)$ common least significant bits. Thus, sieving from enough samples at the outset, we eventually produce states of the form:

$$\Psi_{2^{t-1}} = |0\rangle + (-1)^a|1\rangle,$$

which are sufficient to determine the parity of a. It is shown in the study by Kuperberg [13] that the aforementioned method yields an algorithm that requires $2^{O(\sqrt{\log N})}$ time, space, and queries. In the study by Regev [12], a modified algorithm is given that requires $2^{O(\sqrt{\log N \log\log N})}$ time and $\mathrm{poly}(\log N)$ space. An abstract description of this sieving process is given in previous studies [13, §9] and further improvements and generalizations can be found in the study by Kuperberg [15], in particular the so-called "collimation sieve."

Remark 4.3. In the study by Kuperberg [13], it is shown that HSP for D_{2^t} reduces to determining the parity of a.

4.2 Query complexity

In the study by Ettinger and Høyer [29], it is shown that a polynomial number of HSP samples is sufficient to recover H_a using exponential time post-processing. A related result in the study by Ettinger and Høyer [32] using different methods shows that the HSP problem in a general finite group has polynomial quantum query complexity.

Transposing $i \leftrightarrow j$, and applying a Hadamard gate to the state in (3.17), gives the state

$$\frac{1}{\sqrt{2}}\begin{pmatrix}1 & 1\\ 1 & -1\end{pmatrix}\begin{pmatrix}\omega_N^{ak} & \omega_N^{(a-\alpha)k}\\ \omega_N^{-(a-\alpha)k} & \omega_N^{-ak}\end{pmatrix} = \frac{1}{\sqrt{2}}\begin{pmatrix}\omega_N^{ak}(1+\omega_N^{-ak}) & \omega_N^{-ak}(1+\omega_N^{ak})\\ \omega_N^{ak}(1-\omega_N^{-ak}) & \omega_N^{-ak}(1-\omega_N^{ak})\end{pmatrix}. \tag{4.4}$$

The probability of observing the first row is

$$\frac{1}{2n}(1+\cos(2\pi ak/N)) = \frac{1}{n}\cos^2(\pi ak/N). \tag{4.5}$$

For the second row, it is

$$\frac{1}{2n}(1-\cos(2\pi ak/N)) = \frac{1}{n}\sin^2(\pi ak/N). \tag{4.6}$$

We are now in the situation of the study by Ettinger and Høyer [29] and can apply the post-processing algorithm described (which is exponential in time) to determine a with high probability, for large N.

4.3 Relation to the subset sum problem

Given $x = (x_1, \ldots, x_m) \in (\mathbb{Z}/N\mathbb{Z})^m$ and $r \in \mathbb{Z}/N\mathbb{Z}$, the problem of finding $b \in \{0,1\}^m$ such that $b \cdot x = r$ is called the subset sum problem over $\mathbb{Z}/N\mathbb{Z}$.

The vector b corresponds to specifying a subset of the $x_1, \ldots, x_m$ that sum to r. Denote by

$$S_r^x = \{b \in \{0,1\}^m : b \cdot x = r\}$$

the set of subset sums for (x, r).

If such a b exists, then (x, r) is called a legal instance. In the decision version of the subset sum problem, the problem is to determine whether a given (x, r) is a legal instance.

In the study by Regev [12], it is shown that the ability to efficiently find an element $b \in S_r^x$ for a large fraction of legal instances gives an efficient algorithm to solve DHSP. Furthermore, the study by Bacon et al. [24] shows that the ability to quantum sample from S_r^x allows us to efficiently implement an optimal measurement to determine a from m DCP samples.

The subset sum problem over $\mathbb{Z}$ is known to be an NP-complete problem. Since one can reduce the subset sum problem over $\mathbb{Z}$ to the subset sum problem over $\mathbb{Z}/N\mathbb{Z}$, by choosing a large enough modulus N, it follows that the subset sum problem over $\mathbb{Z}/N\mathbb{Z}$ is also NP-complete.

4.4 Optimal measurements

It is shown in the study by Ettinger and Høyer [33] that efficient elimination observables do not exist for the dihedral group. Further results can be found in the study by Bacon et al. [24]. In particular, let

$$\nu = m/\log_2 N$$

be the density defined in the study by Bacon et al. [24].

It is shown in [24, Theorem 2] that if $\nu > 1 + 4/\log_2 N$, the probability of determining a using the optimal measurement on m DCP samples is $\geq 1/8$. Furthermore, for any N and m, the probability of determining a is

$$\leq 2^m/N = 2^{(\nu-1)\log_2 N}, \tag{4.7}$$

which is exponentially small in $\log_2 N$ for any fixed $\nu < 1$, and gives a trivial upper bound when $\nu \geq 1$.

More general results on optimal measurements to distinguish conjugate hidden subgroups in certain groups can be found in the study by Moore and Russell [30].

In the study by Bacon et al. [24], the success probability of the optimal measurement is determined as:

$$p_{m,N} = \frac{1}{2^m N^{m+1}} \sum_{x \in (\mathbb{Z}/N\mathbb{Z})^m} \left(\sum_{r \in \mathbb{Z}/N\mathbb{Z}} \sqrt{\eta_r^x} \right)^2,$$

where $\eta_r^x \coloneqq |S_r^x|$.

Remark 4.8. For example, let $m = 2$, $N = 2^m$, and $\nu = 1$. Computer calculations show that $p_{m,N} \approx 0.6665$. On the other hand, we saw in Remark 3.10 that we can achieve a success probability of ≈0.9892 for $m = 2$ in the cyclic group case.

In the study by Moore and Russell [30], it is shown that the optimal POVM measurement to determine a from m DCP samples exists and is given by PGM. The theorem of Naimark states that a POVM measurement on a system can be realized by augmenting the system with ancilla registers, applying a unitary operator, and then a projection-valued measurement on the ancilla. Seen in this light, the result in the study by Ettinger and Høyer [29] implies that the success probability of the optimal measurement is $>1 - \frac{1}{2N}$ if $\nu > 89$, though no efficient implementation is known.

Remark 4.9. In the classical world, if we have a probabilistic algorithm that succeeds with probability $>\frac{1}{2}$, we can run the algorithm multiple times on the same input to make the success probability arbitrarily close to 1. In the quantum world, we cannot, in general, reuse inputs that are quantum states, so running the quantum algorithm multiple times requires more quantum samples, unless one can clone the input samples. However, we will see in the last section that for some problems such as DCP, cloning the input samples is essentially equivalent to solving the original problem.

5 A probabilistic no-go result for DCP

First, a unitary no-go result for DCP.

Theorem 5.1. *There is no unitary operation to compute the value of a into a register from a list of DCP samples for a.*

Proof. Suppose there is a unitary operator U, which has the effect:

$$U|A\rangle|\psi_a^1\rangle\cdots|\psi_a^m\rangle|0\rangle = |\Sigma_a(\psi_a)\rangle|a\rangle, \tag{5.2}$$

for every a, i.e., U takes a list of DCP samples for fixed but unknown a, a blank initialization state $|0\rangle$, and an ancilla state $|A\rangle$, and then computes a into the blank register.

For any other $b \neq a$, we must also have

$$U|A\rangle|\psi_b^1\rangle\cdots|\psi_b^m\rangle|0\rangle = |\Sigma_b(\psi_b)\rangle|b\rangle. \tag{5.3}$$

There are choices of ψ_c^i for $i = 1,\ldots, m$ such that

$$\langle\psi_a^i|\psi_b^i\rangle = \frac{1}{2}, \tag{5.4}$$

for all $a \neq b$ and $i = 1, \ldots m$. To see this, recall the states

$$\psi_a = \frac{1}{\sqrt{2}}(|x^\alpha\rangle + |yx^{a-\alpha}\rangle),$$
$$\psi_b = \frac{1}{\sqrt{2}}(|x^\beta\rangle + |yx^{b-\beta}\rangle),$$

which have possible inner product $\langle\psi_a|\psi_b\rangle \in \left\{0, \frac{1}{2}, 1\right\}$, and there are choices of ψ_a and ψ_b such that

$$\langle\psi_a|\psi_b\rangle \neq 0, 1, \tag{5.5}$$

for instance, if $a \neq b$ and $a - \alpha = b - \beta$ or $\alpha = \beta$. In particular, taking

$$\psi_c^i = |x^c\rangle + |yx^0\rangle,$$

for $c \in \mathbb{Z}/N\mathbb{Z}$, satisfies (5.4).

Taking the inner product of (5.2) and (5.3), we obtain

$$\langle\psi_a^1|\psi_b^1\rangle \cdots \langle\psi_a^m|\psi_b^m\rangle = \langle\Sigma_a(\psi_a)|\Sigma_b(\psi_b)\rangle\langle a|b\rangle = 0, \tag{5.6}$$

a contradiction as we have shown that there are choices of ψ_a^i and ψ_b^i, making the left-hand side of (5.6) non-zero. □

We will give yet another proof of Theorem 5.1 in Theorem 6.15. The proof of Theorem 5.1 mirrors the proof of the no cloning theorem [34] and precludes unitary operations, but not more general quantum algorithms, which may allow for approximate outputs, probabilistic processes, or post-processing. Indeed, computing the exact value of a into a register is rather strong: even in the finite cyclic group case, the standard algorithm only determines a generator for the hidden subgroup using a process of the type given in Theorem 3.15.

The following is a probabilistic no-go result for DCP based on modifying the proof of the unitary no-go result for DCP.

Theorem 5.7. *The problem of determining a, given a list of m DCP samples for unknown a, is not unitarily solvable with a success parameter independent of a, i.e.,* $\geq 1 - \frac{1}{9} \cdot 2^{-2m}$.

Proof. To ease notation, we let

$$\psi_a = |\psi_a^1\rangle \cdots |\psi_a^m\rangle, \tag{5.8}$$

$$\psi_b = |\psi_b^1\rangle \cdots |\psi_b^m\rangle. \tag{5.9}$$

Suppose that there is a unitary operator U, which has the effect:

$$U|A\rangle|\psi_a^1\rangle \cdots |\psi_a^m\rangle|0\rangle = |\Sigma_a(\psi_a)\rangle|a\rangle + \sum_{c \neq a} |\Sigma_c(\psi_a)\rangle|c\rangle, \tag{5.10}$$

$$U|A\rangle|\psi_b^1\rangle \cdots |\psi_b^m\rangle|0\rangle = |\Sigma_b(\psi_b)\rangle|b\rangle + \sum_{c \neq b} |\Sigma_c(\psi_b)\rangle|c\rangle, \tag{5.11}$$

where

$$|\Sigma_a(\psi_a)|^2 \geq 1 - 2^{-\delta}, \quad |\Sigma_b(\psi_b)|^2 \geq 1 - 2^{-\delta}, \tag{5.12}$$

and δ is to be chosen.

Because of (5.12), we have that

$$\sum_{c \neq a} |\Sigma_c(\psi_a)|^2 < 2^{-\delta}, \quad \sum_{c \neq b} |\Sigma_c(\psi_b)|^2 < 2^{-\delta}. \tag{5.13}$$

Taking the inner product of (5.10) and (5.11), we obtain

$$\begin{aligned}\langle\psi_a^1|\psi_b^1\rangle\cdots\langle\psi_a^m|\psi_b^m\rangle &\le |\langle\Sigma_a(\psi_a)|\Sigma_a(\psi_b)\rangle| + |\langle\Sigma_b(\psi_a)|\Sigma_b(\psi_b)\rangle| + \sum_{c\neq a,b} |\langle\Sigma_c(\psi_a)|\Sigma_c(\psi_b)\rangle| \\ &\le |\langle\Sigma_a(\psi_a)|\Sigma_a(\psi_b)\rangle| + |\langle\Sigma_b(\psi_a)|\Sigma_b(\psi_b)\rangle| + 2^{-\delta} \\ &\le 2^{-\delta} + 2\cdot 2^{-\delta/2} \\ &< 3\cdot 2^{-\delta},\end{aligned} \tag{5.14}$$

using Cauchy–Schwartz repeatedly. Arrange the left-side of (5.14) to be 2^{-m} as in (5.5), and we see that choosing $\delta \ge 2(m + \log_2 3)$ gives a contradiction to the aforementioned inequality. □

Remark 5.15. At fixed $\nu = m/\log_2 N$, Theorem 5.7 gives an upper bound on the success parameter of

$$1 - 2^{-2(\nu\log_2 N + \log_2 3)} = 1 - \frac{1}{9}N^{-2\nu}. \tag{5.16}$$

Although the bound in (5.16) seems far from optimal (see Remark 4.8), it is still stronger than trivial bounds, which result from (4.7) [24, Theorem 2] or [30] when $\nu \ge 1$.

6 Quantum cloning and DCP

In this section, we explain a connection between DCP and quantum cloning. Although the topics in this section are not needed for the results of the previous section, the connection with quantum cloning helped motivate the proofs of the previous section, so we have included it for completeness.

By copying a state $|\psi\rangle$, we mean forming the composite state $|A\rangle|\psi\rangle|0\rangle$ for a blank initialization state $|0\rangle$ and ancilla state $|A\rangle$, and applying a quantum algorithm to produce the state $|\Sigma(\psi)\rangle|\psi\rangle|\psi\rangle$.

The no cloning theorem asserts that there is no unitary operation that can copy a general unknown quantum state. However, if the states are chosen from a known set of mutually orthogonal states, it is well known that cloning is possible, as shown for completeness in the following proposition.

Proposition 6.1. *Let $|\psi_{a;1}\rangle,\dots,|\psi_{a;m}\rangle$ be a set of mutually orthogonal states that depend on a parameter a. Suppose $|\psi\rangle = |\psi_{a,i}\rangle$ for some index i (which is unknown).*

If the value of a is known and we can encode a unitary operator U_a such that $U_a|\psi_{a;i}\rangle = |i\rangle$, then there is a unitary operation that copies $|\psi\rangle$.

Proof. First, note that we can copy any state $|i\rangle$ of the computational basis. Start with

$$|i\rangle|0\rangle = |i_n\rangle\dots|i_0\rangle|0\rangle\dots|0\rangle,$$

where we have encoded the last two registers into n qubits, for n large enough.

Applying a CNOT gate to the jth and $(j + n + 1)$th qubits $|i_j\rangle|0\rangle$ produces $|i_j\rangle|i_j\rangle$ for every j. Hence, we can produce the state

$$|i_n\rangle\dots|i_0\rangle|i_n\rangle\dots|i_0\rangle = |i\rangle|i\rangle.$$

The unitary operator U_a has the effect:

$$U_a|\psi_{a;i}\rangle = |i\rangle.$$

Starting with

$$|\psi_{a;i}\rangle|0\rangle,$$

apply U_a to the first register to obtain

$$|i\rangle|0\rangle.$$

Copy the state $|i\rangle$ to obtain

$$|i\rangle|i\rangle.$$

Applying U_a^{-1} to both pairs of registers gives

$$|\psi_{a;i}\rangle|\psi_{a;i}\rangle. \qquad \square$$

Remark 6.2. Since a is known, to encode U_a, we can use universality results (c.f. [35, §4.5] or the Solovay–Kitaev Theorem [35, Appendix 3] for a fault-tolerant version).

Later, we will need a slightly stronger version of Proposition 6.1.

Proposition 6.3. *Let $|\psi_{a;1}\rangle,\ldots,|\psi_{a;m}\rangle$ be a set of mutually orthogonal states that depend on a parameter a, and assume that we can encode a unitary operator T such that $T|a\rangle|\psi_{a;i}\rangle = |a\rangle|i\rangle$.*

Suppose $|\psi\rangle = |\psi_{a;i}\rangle$ for some index i (which is unknown). If we have the value of a in a register, then there is a unitary operation that copies $|\psi\rangle$.

Proof. Starting with

$$|a\rangle|\psi_{a,i}\rangle|0\rangle|0\rangle,$$

apply T to obtain

$$|a\rangle|i\rangle|0\rangle|0\rangle.$$

Copy the states $|a\rangle$ and $|i\rangle$ to obtain

$$|a\rangle|i\rangle|a\rangle|i\rangle.$$

Applying T^{-1} to both pairs of registers gives

$$|a\rangle|\psi_{a,i}\rangle|a\rangle|\psi_{a,i}\rangle,$$

which we can permute to obtain

$$|a\rangle|\psi_{a,i}\rangle|\psi_{a,i}\rangle|a\rangle. \qquad \square$$

Proposition 6.4. *If we can copy any given DCP sample*

$$\psi_{a;\alpha} = \frac{1}{\sqrt{2}}(|x^\alpha\rangle + |yx^{a-\alpha}\rangle), \tag{6.5}$$

to produce a state of the form:

$$\psi_{a;\alpha} \otimes \psi_{a;\alpha} = \frac{1}{\sqrt{2}}(|x^\alpha\rangle + |yx^{a-\alpha}\rangle) \otimes \frac{1}{\sqrt{2}}(|x^\alpha\rangle + |yx^{a-\alpha}\rangle), \tag{6.6}$$

then we can determine the value of a from DCP samples for a.

If a is known, then we can copy any given DCP sample for a using a unitary operation.

Proof. Given samples of the form (6.6), we measure both registers, and with probability 1/2, we obtain

$$|x^\alpha\rangle|yx^{a-\alpha}\rangle \quad \text{or} \quad |yx^{a-\alpha}\rangle|x^\alpha\rangle. \tag{6.7}$$

The sum of the observed exponents of the two registers gives a.

If a is known, then DCP samples for a,

$$\psi_{a;\alpha} = \frac{1}{\sqrt{2}}(|x^\alpha\rangle + |yx^{a-\alpha}\rangle),$$

are chosen from a set of mutually orthogonal states depending on the parameter a. By Proposition 6.1, for each sample of the form (6.5), we can copy it to produce a sample of the form (6.6). □

Remark 6.8. Copying a DCP sample up to parity would allow us to determine the parity of a, and *vice versa*.

Theorem 6.9. *If a is unknown, there is no unitary operation, which from a list of DCP samples for a, copies an additional DCP sample for the same a, while leaving the list of DCP samples alone.*

Proof. Suppose that there is a unitary operator U that transforms

$$U|A\rangle|\psi_a^1\rangle\cdots|\psi_a^m\rangle|\psi_a\rangle|0\rangle = |\Sigma_a(\psi_a)\rangle|\psi_a^1\rangle\cdots|\psi_a^m\rangle|\psi_a\rangle|\psi_a\rangle, \tag{6.10}$$

where $\psi_a = \psi_{a;\alpha} = \frac{1}{\sqrt{2}}(|x^\alpha\rangle + |yx^{a-\alpha}\rangle)$ is a DCP sample for a fixed, and α randomly chosen for each such state. We are supposing that U performs the aforementioned operation for *any* (unknown) a. Thus, we also have that

$$U|A\rangle|\psi_b^1\rangle\cdots|\psi_b^m\rangle|\psi_b\rangle|0\rangle = |\Sigma_b(\psi_b)\rangle|\psi_b^1\rangle\cdots|\psi_b^m\rangle|\psi_b\rangle|\psi_b\rangle, \tag{6.11}$$

for any other b.

Taking the inner product of both sides of (6.10) and (6.11), we deduce

$$\langle\psi_a^1|\psi_b^1\rangle\cdots\langle\psi_a^m|\psi_b^m\rangle\langle\psi_a|\psi_b\rangle = \langle\psi_a^1|\psi_b^1\rangle\cdots\langle\psi_a^m|\psi_b^m\rangle\langle\psi_a|\psi_b\rangle^2\langle\Sigma_a(\psi_a)|\Sigma_b(\psi_b)\rangle. \tag{6.12}$$

However, there are choices of ψ_a^i, ψ_b^i for $i = 1,\ldots,m$, and ψ_a, ψ_b, which do not satisfy (6.12) from (5.5).

We may thus suppose without loss of generality that $\langle\psi_a^i|\psi_b^i\rangle \neq 0, 1$ for all $i = 1,\ldots,N$, and hence, (6.12) becomes

$$\langle\psi_a|\psi_b\rangle = \langle\psi_a|\psi_b\rangle^2\langle\Sigma_a(\psi_a)|\Sigma_b(\psi_b)\rangle.$$

We obtain a contradiction again by choosing ψ_a and ψ_b so that $\langle\psi_a|\psi_b\rangle \neq 0, 1$ as then

$$|\langle\psi_a|\psi_b\rangle| = \frac{1}{2}, \tag{6.13}$$

$$|\langle\psi_a|\psi_b\rangle^2\langle\Sigma_a(\psi_a)|\Sigma_b(\psi_b)\rangle| \leq \frac{1}{4}. \tag{6.14}$$

□

The following is another proof of Theorem 5.1 using the connection with quantum cloning.

Theorem 6.15. *There is no unitary operation to compute the value of a into a register from a list of DCP samples for a.*

Proof. Suppose that there is a unitary operator U, which has the effect

$$U|A\rangle|\psi_a^1\rangle\cdots|\psi_a^m\rangle|0\rangle = |\Sigma_a(\psi_a)\rangle|a\rangle, \tag{6.16}$$

i.e., U takes a list of DCP samples for fixed but unknown a, a blank initialization state $|0\rangle$, and an ancilla state $|A\rangle$, and then computes a into the blank register.

Using an additional blank register and copying $|a\rangle$, there is a unitary operator U' with the effect:

$$U'|A\rangle|\psi_a^1\rangle\cdots|\psi_a^m\rangle|0\rangle|0\rangle = |\Sigma_a(\psi_a)|a\rangle|a\rangle. \tag{6.17}$$

Use U^{-1} and permute $|a\rangle$ and $|0\rangle$ to obtain

$$|A\rangle|\psi_a^1\rangle\cdots|\psi_a^m\rangle|a\rangle|0\rangle. \tag{6.18}$$

Thus, without loss of generality, we may assume the unitary operator U has the effect:

$$U|A\rangle|\psi_a^1\rangle\cdots|\psi_a^m\rangle|0\rangle = |A\rangle|\psi_a^1\rangle\cdots|\psi_a^m\rangle|a\rangle,$$

i.e., U takes a list of DCP samples for fixed but unknown a, a blank initialization state $|0\rangle$, and an ancilla state $|A\rangle$, and then computes a into the blank register, while leaving the list of DCP samples alone.

Now, note that DCP samples $\psi_{a;\alpha}$ can be encoded using two registers as:

$$\frac{1}{\sqrt{2}}(|0\rangle|\alpha\rangle + |1\rangle|a - \alpha\rangle).$$

The unitary operator V, which sends

$$V|a\rangle|0\rangle|\alpha\rangle = |a\rangle|0\rangle|\alpha\rangle,$$
$$V|a\rangle|1\rangle|\alpha\rangle = |a\rangle|1\rangle|a - \alpha\rangle,$$

will have the effect:

$$V|a\rangle|\psi_{a;\alpha}\rangle = |a\rangle\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)|\alpha\rangle.$$

Using a Hadamard gate, we can encode a unitary operator U_0 such that

$$U_0\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)|\alpha\rangle = |0\rangle|\alpha\rangle,$$
$$U_0\frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)|\alpha\rangle = |1\rangle|\alpha\rangle.$$

Then, the unitary operator $(I \otimes U_0)V$ has the effect:

$$(I \otimes U_0)V|a\rangle|\psi_{a;\alpha}\rangle = |a\rangle|0\rangle|\alpha\rangle.$$

We can thus apply Proposition 6.3 to copy an additional DCP sample for the same a using a unitary operation, while leaving the list of DCP samples alone. This contradicts Theorem 6.9. □

7 Conclusion

The quantum hardness of the HSP for the dihedral group has important implications for post-quantum cryptography due to connections with LWE. The standard algorithm fails to provide a quantum speedup compared to cyclic groups, and there is a qualitative difference in the success probabilities of the optimal measurements for solving the coset sampling problem on the two groups.

Further work on understanding obstructions to efficiently implementing optimal measurements or new methods for quantum speed up are needed to resolve the question of its quantum hardness. On the other hand, it may be possible that DCP is harder than LWE, so EDCP could be the more relevant problem to study.

Acknowledgements: We would like to thank P. Høyer for helpful comments and bringing to our attention [13]. We also thank R. Goenka and N. de Silva for stimulating discussions related to the topics of this article.

Funding information: This work was supported by a NSERC Discovery Grant RGPIN-2017-03892 (Imin Chen) NSERC USRA 2020 (David Sun).

Author contributions: The authors contributed equally to the conception, design, execution, or interpretation of the reported study.

Conflict of interest: The authors state no conflict of interest.

Data availability statement: All data generated or analyzed during this study are included in this published article.

References

[1] Kitaev A. Quantum computations: Algorithms and error correction. Russian Math Surveys. 1997;52:1191–249.
[2] Mosca M. The abelian hidden subgroup problem. in: Encyclopedia of Algorithms M.-Y. Kao, (Ed.), New York: Springer; 2016. p. 1–4.
[3] Grigni M, Schulman L, Vazirani M, Vazirani U. Quantum mechanical algorithms for the nonabelian hidden subgroup problem. Combinatorica. 2004;24(1):137–54.
[4] Hallgren S, Russell A, Ta-Shma A. The hidden subgroup problem and quantum computation using group representations. SIAM J Comput. 2003;32(4):916–834.
[5] Shor P. Algorithms for quantum computation: discrete logarithms and factoring. Proceedings of the 35th Annual Symposium on Foundations of Computer Science (FOCS);1994. p. 124–34.
[6] Jozsa R. Quantum factoring, discrete logarithms, and the hidden subgroup problem. Comput Sci Eng. March–April 2001;3(2):34–43, doi: https://doi.org/10.1109/5992.909000.
[7] Beals R. Quantum computation of Fourier transforms over symmetric groups. in: Proceedings 29th Annual ACM Symposium on Theory of Computing (El Paso, Texas), ACM Press, 1997.
[8] Boneh R, Lipton R. Quantum cryptoanalysis of hidden linear functions. Advances in Cryptology - Crypto '95, Lecture Notes in Computer Science. vol. 963, Berlin: Springer-Verlag; 1995. p. 424–37.
[9] Ettinger M, Høyer P. A quantum observable for the graph isomorphism problem. 1999, https://arxiv.org/abs/quant-ph/9901029.
[10] Høyer P. Efficient quantum transforms. 1997. https://arxiv.org/abs/quant-ph/9702028.
[11] Moore C, Russell A, Schulman L. The symmetric group defies strong Fourier sampling. SIAM J Comput. 2008;37(6):1842–64.
[12] Regev O. Quantum computation and lattice problems. SIAM J Comput. 2004;33(3):738–60.
[13] Kuperberg G. A subexponential-time quantum algorithm for the dihedral hidden subgroup problem. SIAM J Comput. 2005;35(1):170–88.
[14] Regev O. A subexponential time algorithm for the dihedral hidden subgroup problem with polynomial space. 2004. https://arxiv.org/abs/:quant-ph/0406151.
[15] Kuperberg G. Another subexponential-time quantum algorithm for the dihedral hidden subgroup problem, 8th Conference on the theory of quantum computation. Communication and Cryptography. vol. 22, 2013, p. 20–34.
[16] Castryck W, Dooms A, Emerencia C, Lemmens A. A fusion algorithm for solving the hidden shift problem in finite Abelian groups, post-quantum cryptography. Lecture Notes in Computer Science. vol. 12841, Cham: Springer; 2021. p. 133–53.
[17] Childs A, Jao D, Soukharev V. Constructing elliptic curve isogenies in quantum subexponential time. J Math Cryptol. 2014;8:1–29.
[18] Biasse JF, Jao D, Sankar A. A quantum algorithm for computing isogenies between supersingular elliptic curves. in: Meier, W., Mukhopadhyay, D. (eds) Progress in Cryptology - INDOCRYPT 2014, Lecture Notes in Computer Science, vol. 8885, Cham: Springer; 2014.
[19] Castryck W, Decru T. An efficient key recovery attack on SIDH, Cryptology ePrint Archive. 2022, https://eprint.iacr.org/2022/975.
[20] Maino L, Martindale C. An attack on SIDH with arbitrary starting curve. Cryptology ePrint Archive. 2022. https://eprint.iacr.org/2022/1026.
[21] Robert D. Breaking SIDH in polynomial time. Cryptology ePrint Archive. 2022. https://eprint.iacr.org/2022/1038.
[22] Brakerski Z, Kirshanova E, Stehlé D, Wen W. Learning with errors and extrapolated dihedral cosets. In: Abdalla, M., Dahab, R. (eds) Public-Key Cryptography - PKC 2018. Lecture Notes in Computer Science, vol. 10770, Cham: Springer; 2018.
[23] Kobayashi H, Le Gall F. Dihedral hidden subgroup problem: a survey. IPSJ J. 2005;46(10):2409–16.
[24] Bacon D, Childs A, van Dam W. Optimal measurements for the dihedral hidden subgroup problem. Chicago J Theoretical Comp Sci. 2006;2006:2.
[25] Lomont C. The hidden subgroup problem - review and open problems. 2004, https://arxiv.org/abs/quant-ph/0411037.
[26] Hales L, Hallgren S. Improved quantum Fourier transform algorithm and applications. in: Proceedings of the 41st Annual Symposium on Foundations of Computer Science (Redondo Beach, California), FOCS, 2000.
[27] Nymann J. On the probability that k positive integers are relatively prime. J Number Theory. 1972;4:469–73.
[28] Childs A. Lecture Notes on Quantum Algorithms. https://www.cs.umd.edu/amchilds/qa/.
[29] Ettinger M, Høyer P. On quantum algorithms for noncommutative hidden subgroups. Adv Appl Math. 2000;25:239–251.
[30] Moore C, Russell A. For distinguishing conjugate hidden subgroups, the pretty good measurement is as good as it gets. Quantum Inform Comput. 2007;7:752–65.
[31] Schoof R. Nonsingular plane cube curves over finite fields. J Comb Theory Series A. 1987;46(2):183–211.
[32] Ettinger M, Høyer P. The quantum query complexity of the hidden subgroup problem is polynomial. Inform Process Lett. 2004;91(1):43–8.
[33] Ettinger M, Høyer P. Quantum state detection via elimination. 1999. https://arxiv.org/abs/quant-ph/9905099.
[34] Wootters WK, Zurek WH. A single quantum cannot be cloned. Nature. 1982;299:802–3.
[35] Nielsen M, Chuang I. Quantum computation and quantum information. Cambridge: Cambridge University Press; 2000.