Shor's Algorithm is Optimal

Lawrence Ip*

November 5, 2003

Abstract

We show that the 'standard' quantum algorithm for the abelian hidden subgroup problem is not only efficient but is optimal in the information theoretic sense, in that it maximizes the probability of correctly identifying the hidden subgroup. The proof uses semidefinite programming to show that the standard algorithm implements the optimal set of measurements.

The arguments break down for the nonabelian hidden subgroup problem, and for the special case of the dihedral group, we give explicit expressions for the optimal measurement to distinguish between the subgroups given one random coset state. This measurement cannot be expressed in terms of the Fourier basis, which suggests that to find a quantum algorithm for the nonabelian hidden subgroup problem we may have to look beyond the Fourier transform.

1 Introduction

The greatest success in quantum algorithms to date has been Shor's algorithm for the order finding problem (and thus for factoring) [20]. The natural generalization of the order finding problem is the hidden subgroup problem:

Definition 1 (Hidden Subgroup Problem). *Given a group G, an unknown subgroup H, and oracle access to a function f that is constant and distinct on cosets of H, find H.*

The abelian hidden subgroup problem (when G is abelian) can be solved using a natural generalization of Shor's algorithm in what has sometimes been called the 'standard' algorithm for the hidden subgroup problem.

Algorithm 1 (Abelian Hidden Subgroup Problem). [1]

1. *Compute the superposition*

$$\sum_{g \in G} |g\rangle |f(g)\rangle,$$

and measure the second register to obtain a superposition over cH, a random coset of H, in the first register.

$$\sum_{h \in H} |ch\rangle |f(c)\rangle.$$

*University of California, Berkeley. Supported by DARPA Grant F30602-01-2-2054. lip@cs.berkeley.edu

[1] Here as for the rest of the paper, we ignore normalization constants where convenient.

2. *Compute the Fourier transform of the coset state to obtain*

$$\sum_{\chi \in \hat{G}} \sum_{h \in H} \chi(ch)|\chi\rangle,$$

where $\hat{G}$ is the dual group of G, the group of characters of G (homomorphisms from G to $(\mathbb{C}, \times)$).

3. *Measure the character χ and repeat the first 3 steps n times.*

4. *Output the subgroup of G whose 'perp'[2] is equal to the subgroup of $\hat{G}$ generated by the observed characters $\chi_1, \ldots, \chi_n$.*

The outstanding open question in quantum algorithms has been to find an efficient quantum algorithm to solve the *nonabelian* hidden subgroup problem. In particular, graph isomorphism can be formulated as a hidden subgroup problem over the symmetric group [12]. Recently Regev showed that a lattice problem, the $n^{2.5}$-unique shortest vector problem, has a quantum reduction to the dihedral coset problem (a problem intimately related to the dihedral hidden subgroup problem) [18]. Both these problems are believed to be neither in P nor NP-hard.

All attempts to solve the nonabelian hidden subgroup problem have approached the problem in a similar way: generate random coset states as in Step 1 of the standard algorithm and then try to distinguish between random coset states of different subgroups. The problem of distinguishing between these states has been approached from two different directions, one information theoretic, the other computational.

The information theoretic question: Is there a measurement that distinguishes between the different subgroups? This was answered by Ettinger, Høyer and Knill who showed that for general groups only a polynomial number of coset states are required to *information theoretically* determine the hidden subgroup [6]. However, the measurements they give do not appear to be efficiently implementable.

The computational question: Is there an *efficiently implementable* measurement that distinguishes between the different subgroups? All attempts to answer this question have *started* with an efficiently computable basis and then tried to analyze what class of groups this basis is able to handle [5, 7, 8, 9, 11, 15, 16, 19]. In particular, *all* of these attempts measure in the Fourier basis, that is, they compute the nonabelian analogue of Step 2 of the standard algorithm. This approach has had only limited success. Most of the results have been negative, and the positive results have tended to be only applicable to very specific classes of groups.

The main difficulty is that the Fourier basis is the only 'natural' basis we know and we have no other candidate bases to try. This motivates the question:

> For the nonabelian hidden subgroup problem, is the Fourier basis the right one? And if not, what is the right basis?

Our work is an attempt to answer this question.

1.1 Our Results

To have any hope of coming up with an efficiently implementable basis we need to look for a basis that has some structure. We impose structure by looking for an *information theoretically optimal* basis. We consider the problem of determining the optimal set of measurements to distinguish between random coset states. If we make n queries to f, we want to distinguish between the different subgroups given n tensored random coset states (from the same subgroup). In general, the optimal measurement will be some *joint* measurement over all the tensored states. More precisely, we consider the following question:

[2] $H^{\perp}$, the 'perp' of a subgroup H is the set of characters χ whose kernel contains H, that is characters χ that satisfy $\chi(h) = 1$ for all $h \in H$.

Definition 2 (Optimal Measurement). *Suppose the hidden subgroup H is chosen uniformly at random from all subgroups of G. Given n tensored random coset states (cosets of H), what is the measurement that maximizes the probability of correctly identifying H?*

The most general quantum measurement can be expressed in terms of a POVM (positive operator-valued measure). The optimal POVM is a solution to a semidefinite program, where the variables are the matrices representing the POVM and the density matrices corresponding to the random coset states appear in the objective function.

For the abelian hidden subgroup problem, we give an explicit solution to this semidefinite program and show that the resulting optimal measurement precisely corresponds to the measurement implemented by the standard algorithm. Thus not only is the standard algorithm efficient for the abelian hidden subgroup problem, it is also optimal in an information theoretic sense.

For the nonabelian subgroup problem, the arguments used in the abelian case break down. We are able to explicitly solve this semidefinite program when the group is a dihedral group (group of symmetries of a regular p-gon) and we have *one* random coset state. The resulting optimal measurement *cannot* be implemented using measurements in the Fourier basis, even if we allow the freedom to choose the basis in the irreducible representations. This shows that the measurements[3] given by Ettinger and Høyer for the dihedral group are not optimal [5]. This suggests that, in general, if we want to measure in the optimal basis for the nonabelian hidden subgroup problem, we must look beyond the Fourier transform.

1.2 Related Work

For general groups, Ettinger, Høyer and Knill showed that information theoretically, only a polynomial number of coset states are required to determine the hidden subgroup [6].

The standard algorithm arose from work by Simon [21], Shor [20] and Kitaev [14]. Hallgren, Russell, Ta-Shma [9] and Grigni, Schulman, Vazirani and Vazirani [8] considered the natural generalization of the standard algorithm to nonabelian groups using group representations, where they defined the *weak* standard method (only measuring the name of the representation) and the *strong* standard method (measuring the name of the representation as well as the row and column). Moore, Rockmore, Russell and Schulman showed that there are groups for which the strong standard method gives an advantage over the weak standard method [16].

There are various results for specific classes of groups. Ettinger and Høyer gave an algorithm for the dihedral group which makes polynomially many quantum queries, but requires exponential classical postprocessing [5]. Kuperberg gave an algorithm that required a subexponential number of quantum queries [15]. Rötteler and Beth gave an algorithm for the wreath product $\mathbb{Z}_2^n \wr \mathbb{Z}_2$ [19]. Ivanyos, Magniez and Santha gave algorithms for groups with with small commutator subgroup and for groups having an abelian normal 2-subgroup of small index [11]. Friedl, Ivanyos, Magniez, Santha and Sen gave an algorithm for $\mathbb{Z}_p^n \rtimes \mathbb{Z}_2$ [7].

The results presented here are similar in flavor to that of Ettinger and Høyer, who defined a notion of an 'efficient elimination observable', showed that this was consistent with the Fourier transform for the abelian hidden subgroup problem, and that such observables do not exist for the dihedral group [4]. Their elimination observables are a restricted class of measurements which eliminate incorrect subgroups with constant probability for each measurement on *one* copy of a coset state, and thus will eliminate *all* incorrect subgroups with constant probability after polynomially (in $\log |G|$) many repetitions. In contrast, rather than restricting ourselves to a particular class of measurements that operate on only one coset state, we optimize over *all joint measurements* over k tensored coset states.

[3] Ettinger and Høyer use an abelian Fourier transform on $\mathbb{Z}_p \times \mathbb{Z}_2$ but their measurements are equivalent to taking the nonabelian Fourier transform on D_{2p} with choice of basis given by a Hadamard.

Semidefinite programming has appeared in quantum information in various contexts. Kitaev used semidefinite programming duality to prove the impossibility of quantum coin flipping [13], Doherty, Parrilo and Spedalieri used semidefinite programming to give criteria for separability [2], and Rains gave bounds on distillable entanglement using semidefinite programming [17].

In the context of quantum computation, Barnum, Saks and Szegedy reformulated quantum query complexity in terms of a semidefinite program [1].

The problem of finding the optimal measurement to distinguish between a set of quantum states was first formulated as a semidefinite program in 1972 by Holevo[4], who gave optimality conditions equivalent to the complementary slackness conditions [10]. Recently, Eldar, Megretski and Verghese showed that the optimal measurements can be found efficiently by solving the dual followed by the use of linear programming [3].

2 Background

2.1 Density Matrices and POVMs

Density matrices allow us to describe quantum systems whose state is not known. If a system is the state $|\psi_i\rangle$ with probability p_i (an ensemble $\{p_i, |\psi_i\rangle\}$), the corresponding density matrix for the system is

$$\rho = \sum_i p_i |\psi_i\rangle\langle\psi_i|.$$

Density matrices capture everything we can say about a quantum system. If two different ensembles have the same density matrix, the ensembles are completely indistinguishable. Any set of measurements we make will have exactly the same statistics for both ensembles[5].

The most general quantum measurement can be expressed as a POVM (positive operator-valued measure). A POVM can be described as a collection of matrices $\{\Pi_i\}_i$ that are positive semidefinite, and complete ($\sum_i \Pi_i = I$). If a system is in the state ρ and we perform the measurement described by $\{\Pi_i\}_i$, the probability that we obtain outcome i is $\mathrm{tr}(\Pi_i\rho)$. Projections are special cases of POVMs. If we want to project onto the basis $\{|\phi_i\rangle\}_i$, the corresponding POVM is given by $\{\Pi_i = |\phi_i\rangle\langle\phi_i|\}_i$. The probability of measuring outcome i given state $|\psi\rangle$ is then $\mathrm{tr}(\Pi_i\rho) = \mathrm{tr}(|\phi_i\rangle\langle\phi_i||\psi\rangle\langle\psi|) = |\langle\phi_i||\psi\rangle|^2$.

We will use ρ_H to denote a random coset state of subgroup H ($|G|/|H|$ cosets, each occurring with probability $|H|/|G|$,

$$\rho_H = \frac{|H|}{|G|} \sum_c |cH\rangle\langle cH|,$$

where $|cH\rangle$ is the uniform superposition over the coset cH and the sum is over coset representatives of H.

2.2 Optimal Set of Measurements as a Semidefinite Program

The optimal set of measurements to distinguish a set of quantum states can be expressed as a semidefinite program. Given a set of density matrices ρ_i with prior probabilities p_i, the POVM $\{\Pi_i\}_i$ that maximizes

[4]This was before the advent of semidefinite programming! In fact, the idea for this paper arose when I read Helstrom's book and observed that Holevo's formulation was a semidefinite program and that his optimality conditions followed immediately (and much less painfully) from standard theorems in semidefinite programming. That Holevo's formulation was a semidefinite program was independently observed by Eldar, Megretski and Verghese [3].

[5]One way to think about this is to observe that measurement probabilities are quadratic functions of the amplitudes and since the density matrix contains all degree 2 terms, it completely determines the measurement statistics.

the probability of identifying the correct state is the solution to the following semidefinite program:

$$\begin{aligned} &\max \operatorname{tr}(\sum_i p_i \rho_i \Pi_i) \\ &\text{subject to } \Pi_i \succcurlyeq 0 \\ &\sum_i \Pi_i = I, \end{aligned}$$

where the $\succcurlyeq$ means that the matrix is positive semidefinite. The dual program is

$$\begin{aligned} &\min \operatorname{tr}(X) \\ &\text{subject to } X \succcurlyeq p_i \rho_i \\ &X \succcurlyeq 0. \end{aligned}$$

Because both the primal and dual are strictly feasible, there is no duality gap and a necessary and sufficient condition for optimality is that the complementary slackness conditions are satisfied:

$$(X - p_i \rho_i)\Pi_i = \Pi_i (X - p_i \rho_i) = 0$$

If $\{\Pi_i\}_i$ and X are optimal solutions to the primal, they satisfy the complementary slackness conditions and by summing over all the conditions and observing that $\sum_i \Pi = I$ we see that $X = \sum_i p_i \rho_i \Pi_i$. So to prove that $\{\Pi_i\}_i$ is an optimal solution for the primal, all we need to do is show that $X = \sum_i p_i \rho_i \Pi_i$ is dual feasible and check the complementary slackness conditions.

2.3 Complex Semidefinite Programming

In general, POVMs may be complex valued. It turns out that our density matrices are real, and so we can restrict our POVMs to being real. For suppose $\Pi = \Pi_R + i\Pi_I$ where Π_R is symmetric and Π_I is antisymmetric. Then $(\rho\Pi_I)^T = \Pi_I^T \rho^T = -\Pi_I \rho$. So $\operatorname{tr}(\rho\Pi_I) = \operatorname{tr}(\Pi_I \rho) = -\operatorname{tr}(\rho\Pi_I) = 0$. Thus the imaginary part of Π does not contribute towards the measurement and without loss of generality we can restrict Π to being real.

3 The Abelian Hidden Subgroup Problem

We first give the optimal measurements for the abelian hidden subgroup problem and show that the these are the measurements implemented by the standard algorithm. We then give a proof of optimality for the measurements when we have one coset state. The proof extends naturally to k tensored coset states.

3.1 The Optimal Measurement and the Standard Algorithm

Let G be a finite abelian group, and let H' be a subgroup chosen uniformly at random from the set of all subgroups of G. Suppose we are given n coset states, each a superposition over a random coset of H'. A set of measurements that maximizes the probability of correctly identifying H' is

$$\left\{ \Pi_H^{(n)} = P_H^{\otimes n} - \sum_{J > H} \Pi_J^{(n)} \right\}_H,$$

where P_H is the projection onto the subspace spanned by uniform superpositions of cosets of H for one copy of the coset state.

To see that the optimal POVM corresponds to measuring in the Fourier basis, recall the standard algorithm. In the standard algorithm, we measure in the Fourier basis for each copy of the coset state to obtain a list of characters $\chi_1, \ldots, \chi_n$, all from $H^\perp$. We take the smallest subgroup of the dual group $\hat{G}$ containing all the $\chi_1, \ldots, \chi_n$ as our estimate of $H^\perp$. This is equivalent to taking the largest subgroup H of G such that $\chi_1, \ldots, \chi_n \in H^\perp$. In other words, we choose the largest subgroup consistent with our measurement outcomes.

First we examine what happens when we have one coset state. Consider P_H, the projection onto the subspace spanned by uniform superpositions over cosets of H. This subspace is spanned by $|\chi\rangle = \sum_x \chi(x)|x\rangle$ for all $\chi \in H^\perp$. Thus

$$\Pi_H = P_H - \sum_{J>H} \Pi_J$$

is the projection onto the subspace spanned by uniform superpositions over cosets of H *and not spanned by uniform superpositions over cosets of any subgroup J containing H*. This is equivalent to measuring in the Fourier basis to obtain χ and then choosing our estimate to be the largest subgroup H such that $\chi \in H^\perp$.

When we have n tensored copies of a random coset state,

$$\Pi_H^{(n)} = P_H^{\otimes n} - \sum_{J>H} \Pi_J^{(n)}$$

is the projection onto the subspace where each coordinate is in the subspace spanned by uniform superpositions over cosets of H *and no coordinate is in the subspace spanned by uniform superpositions over cosets of any subgroup J containing H*. This is equivalent to measuring in the Fourier basis in each coordinate to obtain $\chi_1, \ldots, \chi_n$ and taking the largest subgroup H of G such that $\chi_1, \ldots, \chi_n \in H^\perp$.

3.2 One coset state

We now show that the measurement is optimal given one coset state.

Theorem 1. *Let G be a finite abelian group, and let H' be a subgroup chosen uniformly at random from the set of all subgroups of G. Suppose we are given the coset state that is the superposition over a random coset of H'. A set of measurements that maximizes the probability of correctly identifying H' is*

$$\left\{ \Pi_H = P_H - \sum_{J>H} \Pi_J \right\}_H ,$$

where P_H is the projection onto the subspace spanned by uniform superpositions of cosets of H.

We will need the following key lemma:

Lemma 1. *Given two subgroups H, J of G we have*

(a) $P_H P_J = P_{\langle H,J\rangle}$, where $\langle H, J\rangle$ is the smallest subgroup containing H, J,

(b) $P_H \Pi_J = \Pi_J$ if $J \geq H$,

(c) $P_H \Pi_J = 0$ if $J \not\geq H$,

(d) $\Pi_H \Pi_J = 0$ if $H \neq J$,

(e) $\Pi_H^2 = \Pi_H$.

Proof. (a) Let $K = H \cap J$ and $L = \langle H, J\rangle$. Because G is abelian, K and H are normal in H and L and so we can write $L \cong K \times H/K \times L/H$. Moreover, by the second isomorphism theorem we have $J/K = J/(H \cap J) \cong HJ/H = L/H$. Thus $L \cong K \times H/K \times J/K$ and $G \cong K \times H/K \times J/K \times G/L$. We can now write

$$P_H = \frac{1}{|H|} E_K \otimes E_{H/K} \otimes I_{J/K} \otimes I_{G/L}$$
$$P_J = \frac{1}{|J|} E_K \otimes I_{H/K} \otimes E_{J/K} \otimes I_{G/L},$$

where E_A is the $|A| \times |A|$ all ones matrix and I_B is the $|B| \times |B|$ identity matrix. Thus

$$\begin{aligned} P_H P_J &= \left(\frac{1}{|H|} E_K \otimes E_{H/K} \otimes I_{J/K} \otimes I_{G/L}\right)\left(\frac{1}{|J|} E_K \otimes I_{H/K} \otimes E_{J/K} \otimes I_{G/L}\right) \\ &= \frac{|K|}{|H||J|} E_K \otimes E_{H/K} \otimes E_{J/K} \otimes I_{G/L} = \frac{1}{|L|} E_L \otimes I_{G/L} = P_L, \end{aligned}$$

where we have used the fact that $E_K^2 = |K| E_K$.

(b) Suppose $H \leq J$. Observe that Π_J is a linear combination of P_J and elements of $\{\Pi_{J'}\}_{J'>J}$ and thus a linear combination of $\{P_{J'}\}_{J' \geq J}$. Since $J' \geq J \geq H$, $P_H P_{J'} = P_{J'}$ and so $P_H \Pi_J = \Pi_J$.

(c) Suppose $J \not\geq H$. We use induction on J. Fix H and suppose that $P_H \Pi_{J'} = 0$ for all $J' > J$, $J' \not\geq H$. The base case follows by observing that if there is no such J' then (c) is true by default. Then

$$\begin{aligned} P_H \Pi_J &= P_H \left(P_J - \sum_{J'>J} \Pi_{J'}\right) = P_H P_J - \sum_{J'>J} P_H \Pi_{J'} \\ &= P_{\langle H,J\rangle} - \sum_{J' \geq \langle H,J\rangle} \Pi_{J'} = P_{\langle H,J\rangle} - P_{\langle H,J\rangle} = 0, \end{aligned}$$

where the third line comes from our inductive hypothesis and the fourth line from the definition of $\Pi_{\langle H,J\rangle}$.

(d) Suppose $J \neq H$. Because Π_H, P_H are symmetric we have $P_H \Pi_H = \Pi_H = \Pi_H^T = (P_H \Pi_H)^T = \Pi_H^T P_H^T = \Pi_H P_H$. Thus

$$\Pi_H \Pi_J = \Pi_H P_H P_J \Pi_J = \Pi_H P_{\langle H,J\rangle} \Pi_J = 0,$$

because $J \neq H$ implies that we cannot have both $H, J \geq \langle H, J\rangle$.

(e)

$$\Pi_H^2 = \left(P_H - \sum_{J>H} \Pi_J\right) \Pi_H = \Pi_H - 0 = \Pi_H.$$

□

We now prove optimality of the measurement.

Proof (Theorem 1). Let ρ_H be the uniform mixed state over uniform superpositions of cosets of H, observe that $\rho_H = \frac{|H|}{|G|} P_H$. Let p_H be the probability of choosing subgroup H (here uniform) and $X = \sum_H p_H \rho_H \Pi_H$. Then we need to show that

$$\sum_H \Pi_H = I \tag{1}$$
$$\Pi_H \succcurlyeq 0 \tag{2}$$
$$X - p_H \rho_H \succcurlyeq 0 \tag{3}$$
$$(X - p_H \rho_H)\Pi_H = \Pi_H (X - p_H \rho_H) = 0. \tag{4}$$

(1) is trivially satisfied because $I = P_{\{0\}} = \sum_{H>\{0\}} \Pi_H = \sum_H \Pi_H$.

From Lemma 1(e) we see that Π_H is a projection and thus its eigenvalues are either 0 or 1. This gives us (2).

Now,

$$X = \sum_H p_H \rho_H \Pi_H = \sum_H p_H \frac{|H|}{|G|} P_H \Pi_H = \sum_H p_H \frac{|H|}{|G|} \Pi_H.$$

Then if $p_H \leq \frac{|J|}{|H|} p_J$ for all $J > H$ we have

$$X - p_H \rho_H = \sum_J p_J \frac{|J|}{|G|} \Pi_J - p_H \frac{|H|}{|G|} P_H \succcurlyeq \sum_{J \geq H} p_H \frac{|H|}{|G|} \Pi_J - p_H \frac{|H|}{|G|} P_H = p_H \frac{|H|}{|G|} (P_H - P_H) = 0.$$

If the H is chosen uniformly at random then p_J is constant and we have $p_H \leq \frac{|J|}{|H|} p_J$ for all $J > H$. This gives (3).

To prove (4) we use Lemma 1(d) to show that

$$(X - p_H \rho_H)\Pi_H = \left(\sum_J p_J \frac{|J|}{|G|} \Pi_J - p_H \frac{|H|}{|G|} P_H \right) \Pi_H = \sum_{J \neq H} p_J \frac{|J|}{|G|} \Pi_J \Pi_H + p_H \frac{|H|}{|G|} (\Pi_H^2 - \Pi_H) = 0.$$

Similarly, $\Pi_H (X - p_H \rho_H) = 0$. □

3.3 n copies of coset state

The result for one coset state extends naturally to n tensored copies of the coset state.

Theorem 2. *Let G be a finite abelian group, and let H' be a subgroup chosen uniformly at random from the set of all subgroups of G. Suppose we are given n coset states, each a superposition over a random coset of H'. A set of measurements that maximizes the probability of correctly identifying H' is*

$$\left\{ \Pi_H^{(n)} = P_H^{\otimes n} - \sum_{J>H} \Pi_J^{(n)} \right\}_H,$$

where P_H is the projection onto the subspace spanned by uniform superpositions of cosets of H for one copy.

Proof. The proof is analogous to the proof for the case where we had one copy of the coset state. The proof for Lemma 1 still works when we replace P_H with $P_H^{\otimes n}$. The proof for Theorem 1 becomes a proof for Theorem 2 when we replace $\rho_H = \frac{|H|}{|G|} P_H$ with $\rho_H^{\otimes n} = \left(\frac{|H|}{|G|} \right)^n P_H^{\otimes n}$. □

4 The Nonabelian Hidden Subgroup Problem

The proof of optimality for the abelian hidden subgroup problem does not generalize to the nonabelian case. The point at which it breaks down is Lemma 1(a), where the proof relies on the fact that $HJ = \langle H, J \rangle$. This is true when one of H, J is normal (which is always true if G is abelian) but in general is false if G is nonabelian.

We give the optimal measurement for the hidden subgroup problem over the dihedral group when given one coset state. This measurement cannot be expressed in terms of the Fourier basis.

4.1 Optimal Measurements for the Dihedral Group

For simplicity, we only consider the dihedral group D_{2p} where p is prime. The subgroups of D_{2p} are

$$\begin{aligned} G &= D_{2p} = \{1, r, \ldots, r^{p-1}, s, sr, \ldots, sr^{p-1}\}, \\ C &= \{1, r, \ldots, r^{p-1}\}, \\ R_l &= \{1, sr^l\}, \qquad\qquad \text{for } l = 0, \ldots, p-1, \\ E &= \{1\}. \end{aligned}$$

Given one random coset state, it can be shown that the optimal measurements is

$$\begin{aligned} \Pi_G &= P_G, \\ \Pi_C &= P_C - \Pi_G, \\ \Pi_{R_l} &= \frac{2}{p}(P_{R_l} - \Pi_G), \qquad\qquad \text{for } l = 0, \ldots, p-1, \\ \Pi_E &= 0. \end{aligned}$$

4.2 Optimal Measurements not in Fourier Basis

For D_{2p}, the Fourier basis is not a refinement of the optimal POVM, even if we allow a change of basis through unitary equivalence. This is shown by writing down the POVM corresponding to the Fourier transform (with variables that represent the unitary change of basis) and then showing that no linear combination of these will yield the Π_{R_l}s.

We show the details of the proof for D_6. Analogous arguments work for D_{2p}. The Fourier basis in nonabelian groups is given by the irreducible representations of the group. For $D_6 = \{1, r, r^2, s, sr, sr^2\}$ there are two 1-dimensional and one 2-dimensional irreducible representations. The 1-dimensional representations are the trivial representation $\theta_1(\cdot)$ and the representation $\theta_2(\cdot)$ given by

$$\theta_2(r^l) = 1, \qquad \theta_2(sr^l) = -1.$$

The 2-dimensional irreducible representation ρ is given by

$$\rho(r^l) = \begin{pmatrix} \omega^l & 0 \\ 0 & \omega^{-l} \end{pmatrix}, \qquad \rho(sr^l) = \begin{pmatrix} 0 & \omega^{-l} \\ \omega^l & 0 \end{pmatrix},$$

where $\omega = e^{2\pi i/3}$ is a cube root of unity. The choice of basis for the higher dimensional representations is not unique. Given a unitary matrix U, $U\rho U^\dagger$ is also an irreducible representation. In fact, all unitary irreducible representations can be written in this form. Write U in the form

$$\begin{pmatrix} a & b \\ b^* & -a^* \end{pmatrix},$$

where $|a|^2 + |b|^2 = 1$. Then we have

$$U^\dagger \rho(r^l) U = \begin{pmatrix} a^* a\omega^l + b^* b\omega^{-l} & a^* b(\omega^l - \omega^{-l}) \\ ab^*(\omega^l - \omega^{-l}) & a^* a\omega^{-l} + b^* b\omega^l \end{pmatrix}$$

and

$$U^\dagger \rho(sr^l) U = \begin{pmatrix} ab\omega^l + a^* b^* \omega^{-l} & -a^{*2}\omega^{-l} + b^2\omega^l \\ -a^2\omega^l + b^{*2}\omega^{-l} & -ab\omega^l - a^* b^* \omega^{-l} \end{pmatrix}.$$

So measuring in the Fourier basis is equivalent to projecting into the subspace spanned by each of

$$
\begin{aligned}
v_1 &= \begin{pmatrix} 1 & 1 & 1 & 1 & 1 & 1 \end{pmatrix}/\sqrt{6},\\
v_2 &= \begin{pmatrix} 1 & 1 & 1 & -1 & -1 & -1 \end{pmatrix}/\sqrt{6},\\
v_{11} &= \begin{pmatrix} 1 & a^*a\omega + b^*b\omega^{-1} & a^*a\omega^2 + b^*b\omega^{-2} & ab + a^*b^* & ab\omega + a^*b^*\omega^{-1} & ab\omega^2 + a^*b^*\omega^{-2} \end{pmatrix}/\sqrt{3},\\
v_{22} &= \begin{pmatrix} 1 & a^*a\omega^{-1} + b^*b\omega & a^*a\omega^{-2} + b^*b\omega^2 & -ab - a^*b^* & -ab\omega - a^*b^*\omega^{-1} & -ab\omega^2 - a^*b^*\omega^{-2} \end{pmatrix}/\sqrt{3},\\
v_{12} &= \begin{pmatrix} 0 & a^*b(\omega - \omega^{-1}) & a^*b(\omega^2 - \omega^{-2}) & -a^{*2} + b^2 & -a^{*2}\omega^{-1} + b^2\omega & -a^{*2}\omega^{-2} + b^2\omega^2 \end{pmatrix}/\sqrt{3},\\
v_{21} &= \begin{pmatrix} 0 & ab^*(\omega - \omega^{-1}) & ab^*(\omega^2 - \omega^{-2}) & -a^2 + b^{*2} & -a^2\omega + b^{*2}\omega^{-1} & -a^2\omega^2 + b^{*2}\omega^{-2} \end{pmatrix}/\sqrt{3}.
\end{aligned}
$$

We now show that Π_{R_0}, Π_{R_1}, Π_{R_2} cannot be obtained from the Fourier basis, that is, they cannot be expressed as a positive linear combination of $v_1^\dagger v_1$, $v_2^\dagger v_2$, $v_{11}^\dagger v_{11}$, $v_{22}^\dagger v_{22}$, $v_{12}^\dagger v_{12}$ and $v_{21}^\dagger v_{21}$. Suppose that they can be expressed in this form. The first row of each of Π_{R_0}, Π_{R_1}, Π_{R_2} is

$$
\begin{aligned}
&\begin{pmatrix} +1/3 & -1/6 & -1/6 & +1/3 & -1/6 & -1/6 \end{pmatrix},\\
&\begin{pmatrix} +1/3 & -1/6 & -1/6 & -1/6 & +1/3 & -1/6 \end{pmatrix},\\
&\begin{pmatrix} +1/3 & -1/6 & -1/6 & -1/6 & -1/6 & +1/3 \end{pmatrix}
\end{aligned}
$$

respectively. This is orthogonal to the first row of each of $v_1^\dagger v_1$, $v_2^\dagger v_2$, $v_{12}^\dagger v_{12}$ and $v_{21}^\dagger v_{21}$, so the first row of each of Π_{R_0}, Π_{R_1}, Π_{R_2} must be a positive linear combination of the first row of each of $v_{11}^\dagger v_{11}$ and $v_{22}^\dagger v_{22}$ (note that the first row of $v_{11}^\dagger v_{11}$ and $v_{22}^\dagger v_{22}$ is a multiple of v_{11} and v_{22}).

Now consider the restriction to the last 3 coordinates of the first row. In the last 3 coordinates of v_{11} and v_{22}, one is the negative of the other, and thus they span a 1-dimensional subspace. The last 3 coordinates of the first row of Π_{R_0}, Π_{R_1}, Π_{R_2} span a 2-dimensional subspace. Thus the first row of each of Π_{R_0}, Π_{R_1}, Π_{R_2} cannot be expressed as a positive linear combination of the first row of each of $v_{11}^\dagger v_{11}$ and $v_{22}^\dagger v_{22}$. This shows that Π_{R_0}, Π_{R_1}, Π_{R_2} cannot be expressed in terms of the Fourier basis, even up to unitary equivalence.

This argument can easily be extended to D_{2p}. The only difference is that now we consider the last p coordinates of the first row of Π_{R_l}. These span a subspace of dimension $p-1$. The corresponding coordinates in the Fourier basis span a subspace of dimension $(p-1)/2$ and so the optimal basis cannot be expressed in terms of the Fourier basis.

5 Further Work

The most pressing open question is: What are the optimal measurements for the dihedral group when we are given n copies of the coset state and can we implement these measurements?

The optimal measurements for the dihedral group with one coset state were discovered by numerically solving the semidefinite program and guessing the analytic form of the solutions. As yet we have been unable to extend this method to more than one copy of the coset state because the size of the problems rapidly exceeded the resources of the computers available to us. It may be possible to obtain additional numerical results by using computers with more memory or by finding ways to reformulate the semidefinite program to reduce the computational requirements.

The techniques used to prove optimality of the standard algorithm for the hidden subgroup problem may be applicable to proving optimality of other quantum algorithms. The idea of using semidefinite programming to find candidate bases is also generic and may be useful in other contexts.

6 Acknowledgements

We thank Umesh Vazirani, Julia Kempe, Oded Regev, Iordanis Kerenidis, Sean Hallgren, Zeph Landau and Dave Bacon for valuable discussions. Laurent El Ghaoui and David Williamson gave much appreciated advice on semidefinite program solvers.

7 Bibliography

References

[1] Howard Barnum, Michael Saks, and Mario Szegedy. Quantum query complexity and semi-definite programming. In *Proceedings of the 18th IEEE Annual Conference on Computational Complexity*, pages 179–193, 2003.

[2] Andrew C. Doherty, Pablo A. Parillo, and Federico M. Spedalieri. Distinguishing separable and entangled states. *Physical Review Letters*, 88(187904), 2002.

[3] Yonina C. Eldar, Alexandre Megretski, and George C. Verghese. Designing optimal quantum detectors via semidefinite programming. *IEEE Transactions on Information Theory*, 49(4):1007–1012, April 2003.

[4] Mark Ettinger and Peter Høyer. Quantum state detection via elimination. *quant-ph 9905099*, 1999.

[5] Mark Ettinger and Peter Høyer. On quantum algorithms for noncommutative hidden subgroups. *Advances in Applied Mathematics*, 25(3):239–251, 2000.

[6] Mark Ettinger, Peter Høyer, and Emanuel Knill. Hidden subgroup states are almost orthogonal. *quant-ph 9901034*, 1999.

[7] Katalin Friedl, Gábor Ivanyos, Frédéric Magniez, Miklos Santha, and Pranab Sen. Hidden translation and orbit coset in quantum computing. In *Proceedings of the 35th Annual ACM Symposium on the Theory of Computing*, 2003.

[8] Michelangelo Grigni, Leonard J. Schulman, Monica Vazirani, and Umesh Vazirani. Quantum mechanical algorithms for the nonabelian hidden subgroup problem. In *Proceedings of the 33rd Annual ACM Symposium on the Theory of Computing*, pages 68–74, 2001.

[9] Sean Hallgren, Alexander Russell, and Amnon Ta-Shma. Normal subgroup reconstruction and quantum computation using group representations. In *Proceedings of the 32nd Annual ACM Symposium on the Theory of Computing*, pages 627–635, 2000.

[10] Carl W. Helstrom. *Quantum Detection and Estimation Theory.* Academic Press, 1976.

[11] Gábor Ivanyos, Frédéric Magniez, and Miklos Santha. Efficient quantum algorithms for some instances of the non-abelian hidden subgroup problem. In *Proceedings of the 13th ACM Symposium on Parallel Algorithms and Architectures*, 2001.

[12] Uwe Schöning Johannes Köbler and Jacobo Tóran. *The Graph Isomorphism Problem: Its Structural Complexity.* Birkhäuser, 1993.

[13] Alexei Kitaev. Quantum coin flipping. Talk at Quantum Information Processing 2003, slides and video at `http://www.msri.org`, December 2002.

[14] Alexey Y. Kitaev. Quantum measurements and the abelian stabilizer problem. *quant-ph 9511026*, 1995.

[15] Greg Kuperberg. A subexponential-time quantum algorithm for the dihedral hidden subgroup problem. *quant-ph 0302112*, 2003.

[16] Cristopher Moore, Daniel Rockmore, Alexander Russell, and Leonard Schulman. The power of basis selection in fourier sampling: Hidden subgroup problems in affine groups. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*, 2004. to appear.

[17] Eric M. Rains. A semidefinite program for distillable entanglement. *IEEE Transactions on Information Theory*, 47(7):2921–2933, November 2001.

[18] Oded Regev. Quantum computation and lattice problems. In *Proceedings of the 43rd Annual Symposium on Foundations of Computer Science*, pages 520–529, 2002.

[19] Martin Rötteler and Thomas Beth. Polynomial-time solution to the hidden subgroup problem for a class of non-abelian groups. *quant-ph 9812070*, 1998.

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

[21] Daniel R. Simon. On the power of quantum computation. *SIAM Journal on Computing*, 26(5):1474–1483, October 1997.