Quantum Information and Computation, Vol. 8, No. 5 (2008) 0438–0467

HOW A CLEBSCH-GORDAN TRANSFORM HELPS TO SOLVE THE HEISENBERG HIDDEN SUBGROUP PROBLEM

D. BACON

Department of Computer Science & Engineering, University of Washington, Box 352350
Seattle, WA 98109, USA

Received April 3, 2007
Revised October 31, 2007

It has recently been shown that quantum computers can efficiently solve the Heisenberg hidden subgroup problem, a problem whose classical query complexity is exponential. This quantum algorithm was discovered within the framework of using pretty good measurements for obtaining optimal measurements in the hidden subgroup problem. Here we show how to solve the Heisenberg hidden subgroup problem using arguments based instead on the symmetry of certain hidden subgroup states. The symmetry we consider leads naturally to a unitary transform known as the Clebsch-Gordan transform over the Heisenberg group. This gives a new representation theoretic explanation for the pretty good measurement derived algorithm for efficiently solving the Heisenberg hidden subgroup problem and provides evidence that Clebsch-Gordan transforms over finite groups are a new primitive in quantum algorithm design.

Keywords: Quantum computing, quantum algorithms, hidden subgroup problem

Communicated by: M Mosca & R Cleve

1 Introduction

In 1994 Peter Shor discovered that a quantum computer could efficiently factor integers [1], a problem which is widely suspected to be intractable on a classical computer [2]. Since this discovery an intrepid group of researchers have been attempting to discover quantum algorithms which lie beyond Shor's factoring algorithm with mixed success. On the one hand, a great deal of success has been had in achieving polynomial speedups over classical algorithms within the framework of Grover's quantum search algorithm [3]. On the other hand, quantum algorithms which, like Shor's algorithm, perform exponentially faster than the best classical algorithms have been harder to come by. To be sure, notable successes have been achieved, including Hallgren's efficient quantum algorithm for solving Pell's equation [4] and exponential speedups in certain quantum random walks [5], but so far there have been no new efficient quantum algorithms which move quantum computing's killer application beyond factoring integers and the resulting breaking of multiple public key cryptosystems.

Among the most tempting problems which might be exponentially sped up on a quantum computer is the graph isomorphism problem. The reasons for this are two-fold. First, graph isomorphism belongs to a complexity class very much like that which contains integer factoring. In particular, the decision version of factoring is known to be in the complexity $NP \cap coNP$ while graph isomorphism is known to be in the similar complexity class

$NP \cap coAM$. Algorithms in these complexity classes are unlikely to be NP-complete. Further, both integer factoring and graph isomorphism are not known to have classical polynomial time algorithms despite considerable effort to find such algorithms. Thus graph isomorphism is, like integer factoring, of Goldilocks-like classical complexity, not too hard such that efficiently solving it would revolutionize our notion of tractable, but not so easy as to have already fallen into P (or BPP.) The second reason for attempting to find efficient quantum algorithms for the graph isomorphism problem is that this problem can be solved if there was an efficient quantum algorithm for the non-Abelian version of the problem which lies at the heart of Shor's algorithm, the hidden subgroup problem [6, 7]. Thus motivated, a great deal of effort has been expended in the last few years attempting to solve the non-Abelian hidden subgroup problem efficiently on a quantum computer. Towards this end, a series of efficient quantum algorithms for certain non-Abelian hidden subgroups have been developed [8, 9, 10, 11, 12, 13, 14, 15]. At the same time a series of negative results towards the standard approach to solving the hidden subgroup problem on a quantum computer have also appeared [8, 9, 16, 17, 18, 19, 20]. Viewed pessimistically, these results cast doubt on whether quantum computers can be used to efficiently solve non-Abelian hidden subgroup problems. An alternative optimistic view is also possible. In this view what these results show is that any efficient quantum algorithm for the non-Abelian hidden subgroup problem (in what is known as the standard method) must have a particular form. Specifically such algorithms must perform quantum circuits across many separate quantum queries to a hidden subgroup oracle. If we are to view these results in a positive manner, then this tells us that what is needed, if we are going to efficiently solve non-Abelian hidden subgroup problems, are new quantum transforms which can act across many such queries.

In this paper we provide some evidence in favor of this optimistic view. Recently, Bacon, Childs, and van Dam [14] have shown that quantum computers can efficiently solve the hidden subgroup problem for certain semidirect product groups. One such group which admits an efficient quantum algorithm is the Heisenberg group, $\mathbb{Z}_p^2 \rtimes \mathbb{Z}_p$. The algorithm of Bacon, Childs, and van Dam [14] was discovered using the framework of pretty good measurements [21, 14, 22, 15, 19]. This yields a particular algorithm for solving the Heisenberg hidden subgroup problem which is optimal and which is related to the solution of certain algebraic equations. Here we show that the structure of this quantum algorithm can be derived almost solely from symmetry arguments. These symmetry arguments give rise to a transform, the Clebsch-Gordan transform over the Heisenberg group, which can be used to help efficiently solve the Heisenberg hidden subgroup problem. Previously a Clebsch-Gordan transform was used by Kuperberg to find a subexponential algorithm for the dihedral hidden subgroup problem [23, 24]. Clebsch-Gordan transforms over the unitary group [25, 26] and a certain form of measurement were demonstrated to not help in solving a hidden subgroup problem by Childs, Harrow, and Wocjan [27]. Further Moore, Russell, and Śniady have recently shown that a certain form of Clebsch-Gordan transform used to perform a quantum algorithm which mimics Kuperberg's algorithm cannot do better than known classical algorithms for the hidden subgroup problem relevant to graph isomorphism [28]. In particular they obtain a subexponential lower bound for this approach matching the best known classical algorithms for the problem. Here we show that these negative results can be overcome, at least for the Heisenberg group, by performing a Clebsch-Gordan transform over the relevant finite group instead of over the unitary group, and further, and in direct contrast to the work of Moore, Russell, and Sniady [28], by working with a particular register, known as the multiplicity space register, which is output from a Clebsch-Gordan transform. This is the first time a Clebsch-Gordan transform and its multiplicity register has been identified as a key component in producing a polynomial time algorithm for a hidden subgroup problem.

Our motivation for using a Clebsch-Gordan transform in the hidden subgroup problem arises from considering a slight variant of the standard hidden subgroup problem. The setup for this variant is identical to the hidden subgroup problem, but now the task is not to identify the hidden subgroup but to return which set of conjugate subgroups the hidden subgroup belongs. It is this latter problem, which we call the hidden subgroup conjugacy problem, which endows our system with extra symmetries which allow us to exploit a Clebsch-Gordan transform. An essential step in our use of the Clebsch-Gordan transform is a demonstration that the hidden subgroup problem and the hidden subgroup conjugacy problem are quantum polynomial time equivalent for the Heisenberg group.

The outline of our paper is as follows. In Section 2 we introduce the hidden subgroup problem and discuss relevant prior work on quantum algorithms for this problem. In Section 3 we introduce a variant of the hidden subgroup problem which we call the hidden subgroup conjugacy problem. In Section 4 we review arguments for why the symmetry of hidden subgroup states leads one (in the standard approach to the hidden subgroup problem) to perform a quantum Fourier transform over the relevant group. In Section 5 we present our first new results in showing that for the hidden subgroup conjugacy problem, symmetry arguments lead one to perform (in addition to the quantum Fourier transform) a transform known as the Clebsch-Gordan transform over the relevant group. In Section 6 we introduce the Heisenberg group and show how solving the hidden subgroup conjugacy problem for this group leads to an algorithm for the hidden subgroup problem for this group. In Section 7 we discuss the Clebsch-Gordan transform over the Heisenberg group and show how to efficiently implement this transform with a quantum circuit. Finally in Section 8 we put this Clebsch-Gordan transform to use on the Heisenberg hidden subgroup conjugacy problem and show how this allows one to efficiently solve the Heisenberg hidden subgroup problem.

2 The Hidden Subgroup Problem

Here we define the hidden subgroup problem and give a brief history of attempts to solve this problem efficiently on a quantum computer.

The hidden subgroup problem (HSP) is as follows. Suppose we are given a known group $\mathcal{G}$ and a function f from this group to a set S, $f : \mathcal{G} \to S$. This function is promised to be constant and distinct on left cosets of a subgroup $\mathcal{H}$, i.e. $f(g_1) = f(g_2)$ iff g_1 and g_2 are members of the same left coset $g\mathcal{H}$. We do not know the subgroup $\mathcal{H}$. The goal of the HSP is to identify the hidden subgroup $\mathcal{H}$ by querying the function f. An algorithm for the HSP is said to be efficient if the subgroup can be identified with an algorithm which runs polynomially in the logarithm of the size of the group, $O(\mathrm{polylog}|\mathcal{G}|)$. We will assume, throughout this paper that the group $\mathcal{G}$ is finite and its elements, along with elements of $\mathcal{S}$, have an efficient representation ($O(\mathrm{polylog}|\mathcal{G}|)$) in terms of bitstrings.

In the quantum version of the HSP we are given access to an oracle which queries the function, U_f. This oracle is assumed to have been constructed using classical reversible gates in such a way that applying it to $|g\rangle \otimes |0\rangle$ results in $|g\rangle \otimes |f(g)\rangle$. In the standard query model of the HSP, one inputs a superposition over all group elements into the first register of the quantum oracle and $|0\rangle$ into the second register of oracle. This produces the state

$$U_f \frac{1}{\sqrt{|\mathcal{G}|}} \sum_{g \in \mathcal{G}} |g\rangle \otimes |0\rangle = \frac{1}{\sqrt{|\mathcal{G}|}} \sum_{g \in \mathcal{G}} |g\rangle \otimes |f(g)\rangle. \tag{1}$$

Suppose we now disregard (measure, throw away) the second register. Due to the promise on f, the state of the first register is then a mixed state whose exact form depends on the hidden subgroup $\mathcal{H}$,

$$\rho_{\mathcal{H}} = \frac{|\mathcal{H}|}{|\mathcal{G}|} \sum_{g = \text{coset representative}} |g\mathcal{H}\rangle\langle g\mathcal{H}|, \tag{2}$$

where we have defined the left coset states

$$|g\mathcal{H}\rangle = \frac{1}{\sqrt{|\mathcal{H}|}} \sum_{h \in \mathcal{H}} |gh\rangle. \tag{3}$$

We will call $\rho_{\mathcal{H}}$ the hidden subgroup state. In this paper we will restrict ourselves to algorithms which use the above standard procedure (with one slight variation of not querying over the entire group but only querying over a subgroup of the group.)

The HSP, when the group is Abelian, can be solved efficiently on a quantum computer [1, 29, 6]. The vast majority of early efficient quantum algorithms which demonstrated speedups over classical algorithms, including Simon's algorithm [30], the Deutsch-Jozsa algorithm [31], the non-recursive Bernstein-Vazirani algorithm [32], and Shor's algorithm [1] can all be recast as Abelian HSPs. Given the central nature of this problem to these algorithms, a natural generalization was to consider the hidden subgroup problem for non-Abelian groups. It was quickly noted that if one could efficiently solve the HSP for the symmetric group (or for the wreath product group, $\mathcal{S}_n \wr \mathcal{S}_2$) then one would immediately have an efficient algorithm for the graph isomorphism problem [6, 33, 7, 34]. There is no known efficient algorithm for the graph isomorphism problem despite a considerable amount of effort to solve this problem classically [35]. Adding considerably to the interest in the non-Abelian HSP was the discovery by Regev [16] that a quantum polynomial time algorithm for the dihedral hidden subgroup problem could be used to solve certain unique shortest vector in a lattice problems. Solving either of these two problems would represent a significant breakthrough in quantum algorithms and thus a great deal of research has been aimed at understanding the non-Abelian HSP.

Work on the non-Abelian HSP can be roughly divided into two categories: progress in finding efficient new quantum algorithms for the problem and attempts to elucidate the reason that the standard approach fails to efficiently solve the HSP. For the former, a small, but significant amount of success has been achieved with a general trend of finding algorithms which are, loosely, more and more non-Abelian. The latter has recently culminated in showing that for HSP relevant to the graph isomorphism problem will require a new class of measurements across multiple hidden subgroup states if an efficient quantum algorithm is possible. Here we review the progress in both of these categories.

An early positive result for the non-Abelian hidden subgroup problem was the discovery that the problem had an efficient quantum algorithm when the hidden subgroups are normal and there exists an efficient quantum algorithm for the quantum Fourier transform over the relevant group [9]. Further it was shown that there is an efficient quantum algorithm for the HSP when the groups are "almost Abelian" [8] or, a bit more generally, when the group is "near Hamiltonian" [13] (these conditions mean roughly that the intersection of the normalizers of all the subgroups of the group is large.) Groups with small commutator subgroup [10] along with solvable groups of bounded exponent and of bounded derived series [11] also admit HSPs which can be efficiently solved on a quantum computer. Further, a series of efficient quantum algorithms for the non-Abelian HSP over semidirect product groups have been discovered. Among these are certain groups of the form $\mathbb{Z}_{p^k}^n \rtimes \mathbb{Z}_2$ for a fixed power of a prime p^k [11], q-hedral groups with sufficiently large q [12], and certain metacyclic groups as well as groups of the form $\mathbb{Z}_p^r \rtimes \mathbb{Z}_p$ for fixed r [14]. This last work includes an efficient quantum algorithm for the Heisenberg HSP ($r = 2$) which is the main subject of this paper. Further, some non-Abelian groups can be solved using a classical reduction and only the Abelian version of the HSP [14], including the groups $\mathbb{Z}_2^n \wr \mathbb{Z}_2$ [36] and particular semidirect products of the form $\mathbb{Z}_{p^k} \rtimes \mathbb{Z}_p$ with p an odd prime [37]. Subexponential, but not polynomial, time quantum algorithms for the dihedral group were discovered by Kuperberg [23] and subsequently improved to use only polynomial space by Regev [24]. Finally, a subexponential time quantum algorithm for hidden subgroup problems over direct product groups was recently discovered by Alagic, Moore, and Russell [38]

In addition to the explicit efficient quantum algorithms for the non-Abelian HSP given above, a great deal of work has also been performed examining the query complexity of the problem and the related questions of what is needed in order to information theoretically reconstruct the hidden subgroup. A positive result along these lines was the result of Ettinger, Hoyer, and Knill who showed that the query complexity of the HSP is polynomial [39]. Thus it is known that if one is given a polynomial number of copies of the HSP state $\rho_{\mathcal{H}}$ then there exists a quantum measurement which can correctly identify $\mathcal{H}$ with high probability. However, no efficient quantum algorithm implementing this measurement is known to exist, except for the cases of the efficient quantum algorithms for the HSP described above.

Given that the query complexity of the HSP is polynomial, it is natural to ask how tight this query complexity is. In particular it is natural to ask how many copies of the hidden subgroup state must be supplied in order for there to exist a measurement on these copies which we can efficiently perform and which provides enough information to reconstruct the hidden subgroup. In such cases we say that the hidden subgroup can be information theoretically reconstructed. Note, however that being information theoretically reconstructible does not mean that there is an efficient algorithm for the problem because the classical post processing required to reconstruct the hidden subgroup may not be tractable. What is known about the number of copies needed to information theoretically construct the hidden subgroup? On the one hand it is known that for certain groups, in particular for the dihedral [40], affine [12], and Heisenberg groups [41], measurements on a single register of the hidden subgroup state is sufficient for information theoretic reconstruction of the hidden subgroup, and these quantum measurements can be efficiently enacted. However, it was shown by Moore, Russell and Schulman [17] that for the particular case of the symmetric group HSP, measurements on a single register of the hidden subgroup state reveal only an exponentially small amount of information about the identity of the hidden subgroup. In particular this means that if one makes measurements which act only on a single hidden subgroup problem state at a time, one cannot efficiently solve the HSP. This was subsequently extended to two registers of the hidden subgroup state by Moore and Russell [18], and then, in the culmination of this line of inquiry, Hallgren *et al.* [20] showed that this extends all the way up to the upper bound of Ettinger, Hoyer, and Knill. In other words, for the hidden subgroup problem relevant to the graph isomorphism problem, if one attempts to efficiently solve this problem on a quantum computer one is required to perform a measurement on $O(\log|\mathcal{G}|)$ registers containing the hidden subgroup state in order to solve the problem. In particular if this measurement can be implemented, even adaptively, on less than k registers, then this measurement will not be able to solve the HSP.

The results of Hallgren *et al.* imply that in order to solve the HSP, one must perform measurements on multiple copies of the hidden subgroup state in order to efficiently solve the non-Abelian hidden subgroup problem. One important consequence of this is that the standard method combined with performing a quantum Fourier transform on the hidden subgroup state, a paradigm which works for many of the efficient quantum algorithms for the hidden subgroup problem, cannot be used to find an efficient quantum algorithm for the non-Abelian hidden subgroup problem. In particular the above discussion makes it clear that if there is hope for an efficient quantum algorithm for the non-Abelian hidden subgroup problem in the standard paradigm, then measurements across multiple copies of the hidden subgroup state must be used. So far only a small number of quantum algorithms have used such measurements. The first such algorithm for HSPs was in Kuperberg's subexponential time algorithm for the dihedral hidden subgroup problem [23]. Recently Bacon, Childs, and van Dam [14] showed that for the Heisenberg group a measurement across two hidden subgroup states could be used to efficiently solve this HSP (and a measurement across r hidden subgroup states could be used to solve the hidden subgroup problem over $\mathbb{Z}_p^r \ltimes \mathbb{Z}_p$.) This result gives the first indication that while the results of Hallgren *et al.* put a damper on traditional attempts to solve the HSP using only the standard approach and the quantum Fourier transform, all is not lost, and if there is any hope for efficient quantum algorithm for the full non-Abelian HSP, techniques for efficient quantum measurements across multiple copies of the hidden subgroup state must be developed.

The efficient algorithm for the Heisenberg hidden subgroup problem was discovered by identifying an *optimal measurement* for the hidden subgroup problem. In the optimal measurement approach to the hidden subgroup problem, one assumes that one has been given k copies of the hidden subgroup state $\rho_{\mathcal{H}}$ with an a priori probability $p_{\mathcal{H}}$. One then wishes to find the generalized measurement (which we hereafter refer to only as a measurement) on these k copies which maximizes the probability of successfully identifying the hidden subgroup $\mathcal{H}$, averaged over the a priori probabilities $p_{\mathcal{H}}$. Cast in this form, the hidden subgroup problem becomes a problem of optimal state discrimination. A set of necessary and sufficient conditions for a measurement to be such an optimal measurement was discovered over thirty years ago by Holevo [42] and Yuen, Kennedy, and Lax [43].

Ip [21] was the first to consider the optimal measurement for the HSP. In particular he examined the optimal measurement for the HSP when all of the subgroups are given with equal a prior probability. Ip showed that in this case, for the Abelian HSP, the standard approach to solving the HSP is optimal. Further Ip showed that for the dihedral hidden subgroup problem, the optimal measurement was not to perform a quantum Fourier transform over the dihedral group followed by a projective measurement (sometimes called a von Neumann measurement) on this state space, but instead to use a more general measurement (a POVM.) Continuing on in this line of inquiry, Bacon, Childs, and van Dam derived an exact expression for the optimal measurement for the dihedral hidden subgroup problem [19]. This measurement turned out to be the so-called pretty good measurement [44]. Further the optimal measurement on k copies of the hidden subgroup state was discovered to be a nontrivial measurement across multiple copies of the hidden subgroup states. Thus, even though measurement on a single register containing the hidden subgroup state is enough to information theoretically reconstruct the hidden subgroup state, a measurement across many registers containing the hidden subgroup is optimal for solving this problem. Unfortunately, it is not known how to efficiently implement the optimal measurement described in [19]. However, building upon the optimal measurement approach, Bacon, Childs, and van Dam then applied the apparatus of optimal measurements for the HSP to the HSP for certain semidirect product groups of the for $\mathbb{Z}_p^k \rtimes \mathbb{Z}$, for a fixed k and prime p [14]. Again it was discovered that the optimal measurement on multiple registers containing the hidden subgroup states required a measurement over multiple registers containing the hidden subgroup state (in fact over r registers.) However this time the authors were able to find an efficient quantum algorithm implementing this measurement. Thus a non-trivial quantum algorithm for a non-Abelian hidden subgroup problem was discovered which was optimally solved by a measurement across multiple copies of the hidden subgroup state.

However, in spite of this success, there is much that remains mysterious about the efficient quantum algorithm for the Heisenberg HSP. In particular, why is there an efficient algorithm for implementing the optimal measurement in this case? Is there any *structure* behind measurements across multiple copies of hidden subgroup states which can be used to solve the hidden subgroup problem? In this paper we present partial answers to these questions and highlight the role of an important transform over many registers containing the hidden subgroup state, the Clebsch-Gordan transform, in providing an efficient algorithm. We believe that this is an important insight, first of all because it gives a new explanation for the efficient pretty good measurement based algorithm of Bacon, Childs, and van Dam. Further we believe that our result highlights the important role of that Clebsch-Gordan transform can play in quantum algorithms. Clebsch-Gordan transforms, like quantum Fourier transforms over finite groups, suffer from there not being a canonical choice for the bases of registers output by these transforms. In this work we show that by a judicious choice of this arbitrary basis, we can use Clebsch-Gordan transforms to solve a non-Abelian HSP. In particular, we show that there is a way to label the multiple appearances of a given irreducible representation appearing in the diagonal product of these irreducible representations which allow us to make a measurement which extracts the relevant information abut the hidden subgroup. Thus beyond showing that Clebsch-Gordan transforms can be useful for solving HSPs we can focus the search for efficient HSP algorithm to an even smaller problem of understanding the choose of basis for these transforms.

3 The Hidden Subgroup Conjugacy Problem

In this section we present a variant of the hidden subgroup problem which we label the hidden subgroup conjugacy problem.

Two subgroups $\mathcal{H}_1 \subset \mathcal{G}$ and $\mathcal{H}_2 \subset \mathcal{G}$ are said to be conjugate to each other if there exists an element of $g \in \mathcal{G}$ such that

$$\mathcal{H}_1 = \{gh_2g^{-1}, \forall h_2 \in \mathcal{H}_2\}. \tag{4}$$

The notion of conjugate subgroups forms an equivalence relationship. Therefore we can classify all subgroups into distinct sets of subgroups which are all conjugate to each other. The hidden subgroup conjugacy problem (HSCP) is exactly like the HSP, but instead of requiring that we correctly identify the hidden subgroup $\mathcal{H}$ of a function, we only require than one correctly identify which set of conjugate subgroups $\mathcal{H}$ belongs to.

Clearly solving the HSP in its original form will allow one to solve the hidden subgroup conjugacy problem, but less is known about the reverse relation. First it is clear that when the group is Abelian or when the subgroups are normal, the HSP is equivalent to the HSCP (since subgroups of Abelian groups and normal subgroups are conjugate only to themselves.) Recently Fenner and Zhang [45] have examined the difference between the search and decision version of the HSP (in the search problem one is required to return the hidden subgroup and in the decision subgroup one is required to distinguish whether the hidden subgroup is trivial or not.) Their results imply that the HSCP and the HSP over permutation groups are polynomial time equivalent. Similarly they show that for the dihedral group, when the order of the group is the product of many small primes, then the HSCP and the HSP are polynomial time equivalent. In Section 6 we will show that the HSCP and HSP are quantum polynomial time equivalent for the Heisenberg group.

Finally we note that the HSP can be decomposed into the HSCP along with what we call the hidden conjugate subgroup problem (HCSP). In the hidden conjugate subgroup, one is given a function which hides one of a set of subgroups all of which are conjugate to each other and one desires to identify the conjugate subgroup. For the single copy HCSP, the pretty good measurement was shown to be optimal by Moore and Russell [22].

4 Symmetry Considerations and the HSP

The hidden subgroup state of Eq. (2) possess a set of symmetries which allow us to, without loss of generality, perform a change of basis which exploits these symmetries. These symmetries are related to the regular representations of the group.

4.1 Regular Representations

There are two regular representations of the group $\mathcal{G}$ which we will be interested in, the left regular representation and the right regular representation. Both of these representations act on a Hilbert space with a basis labeled by the elements of the group $\mathcal{G}$. Define the left regular representation via its action on basis states of this Hilbert space,

$$R_L(g)|g'\rangle = |gg'\rangle, \tag{5}$$

where $gg\prime$ is the element of $\mathcal{G}$ obtained by multiplying g and g'. Similarly, define the right regular representation via

$$R_R(g)|g'\rangle = |g'g^{-1}\rangle. \tag{6}$$

The regular representations are, in general, reducible representations of $\mathcal{G}$. In fact these representations are particularly important in the representation theory of finite groups. The reason for this is that the regular representations are reducible into a direct sum of all irreducible representations (irreps) of the group $\mathcal{G}$. Thus, it is possible to find a basis in which R_L acts as

$$R_L(g) = \bigoplus_\mu I_{d_\mu} \otimes D_\mu(g), \tag{7}$$

where the direct sum is over all irreps of the group $\mathcal{G}$, $D_\mu(g)$ is the μth irrep evaluated at group element g and d_μ is the dimension of the irrep μ. A similar decomposition occurs for the right regular representation,

$$R_R(g) = \bigoplus_\mu D_\mu(g) \otimes I_{d_\mu}. \tag{8}$$

In fact, an elementary result of finite group representation theory tells us that the basis in which $R_L(g)$ acts as the above direct sum is also the basis in which $R_R(g)$ acts as the above direct sum. In other words, in this basis,

$$R_L(g)R_R(g') = R_R(g')R_L(g) = \bigoplus_\mu D_\mu(g') \otimes D_\mu(g). \tag{9}$$

Let us call the basis described above the $|\mu, r_\mu, l_\mu\rangle$ basis where $r_\mu = 1, \ldots, d_\mu$ and $l_\mu = 1, \ldots, d_\mu$.

4.2 Symmetry of Hidden Subgroup States

Why are the left and right regular representations relevant to the hidden subgroup problem? Well it is easy to check that the hidden subgroup state $\rho_\mathcal{H}$, for all hidden subgroups $\mathcal{H}$, are invariant under conjugation by the left regular representation,

$$D_L(g)\rho_\mathcal{H} D_L(g^{-1}) = \rho_\mathcal{H}, \tag{10}$$

for all $g \in \mathcal{G}$. The reason for this is that left multiplication by a fixed group element acts as a permutation on left cosets. Notice that $\rho_\mathcal{H}$ is not invariant under a similar transform using the right regular representation unless the representation being used is an element of the subgroup $\mathcal{H}$.

What is the consequence of the invariance of the $\rho_\mathcal{H}$ with respect to the left regular representation? Schur's lemma [46] tells us that if an operator is invariant with respect to the operators enacting a representation of the group, then that operator has support only over the commutant of the this representation. For our purposes the commutant is simply the algebra of matrices which commute with the left regular representation, $MD_L(g) = D_L(g)M$. Since $\rho_\mathcal{H}$ is invariant with respect to the operators enacting the left regular representation of $\mathcal{G}$ and the commutant of the left regular representation of $\mathcal{G}$ is the right regular representation of $\mathcal{G}$, this means that in the $|\mu, r_\mu, l_\mu\rangle$ basis, $\rho_\mathcal{H}$ can be expressed as

$$\rho_\mathcal{H} = \bigoplus_\mu \sigma_{\mu,\mathcal{H}} \otimes I_{d_\mu}, \tag{11}$$

where $\sigma_{\mu,\mathcal{H}}$ is an operator with support only on the space acted upon by the irreducible representation of the right regular representation.

What does this mean for obtaining information about $\mathcal{H}$ by making measurements on $\rho_{\mathcal{H}}$? For now we focus on a single measurement on a single copy of $\rho_{\mathcal{H}}$. Recall that a generalized measurement is described by a set of positive operators $\{P_1, \dots, P_k\}$ which sum to identity $\sum_{\alpha=1}^{k} P_\alpha = I$. Outcomes of the measurement correspond to the indices, with the probability of getting outcome α when measuring state ρ given by $p_\alpha = \mathrm{Tr}\left[\rho P_\alpha\right]$. So if we make a measurement with operator P_α on $\rho_{\mathcal{H}}$ the probability of getting outcome α is

$$p_\alpha = \mathrm{Tr}\left[\rho_{\mathcal{H}} P_\alpha\right]. \tag{12}$$

But, since ρ_H is invariant under the left regular representation operators, we can express this probability as

$$p_\alpha = \mathrm{Tr}\left[\frac{1}{|\mathcal{G}|}\sum_{g\in\mathcal{G}} R_L(g)\rho_{\mathcal{H}} R_L^\dagger(g) P_\alpha\right] = \mathrm{Tr}\left[\rho_{\mathcal{H}}\tilde{P}_\alpha\right], \tag{13}$$

where

$$\tilde{P}_\alpha = \frac{1}{|\mathcal{G}|}\sum_{g\in\mathcal{G}} R_L^\dagger(g) P_\alpha R_L(g). \tag{14}$$

But $\tilde{P}_\alpha$ is P_α symmetrized over left regular representation of the group $\mathcal{G}$. This implies that $\tilde{P}_\alpha$ commutes with the left regular representation,

$$R_L(g')^\dagger \tilde{P}_\alpha R_L^\dagger(g') = \tilde{P}_\alpha, \tag{15}$$

for all g in $\mathcal{G}$. Thus using Schur's lemma, $\tilde{P}_\alpha$ has support over only the commutant of the left regular representation, our good friend the right regular representation. In other words, in the $|\mu, r_\mu, l_m u\rangle$ basis, we find that

$$\tilde{P}_\alpha = \bigoplus_\mu Q_{\mu,\alpha} \otimes I_{d_\mu}. \tag{16}$$

This means that given some measurement P_α, the probabilities of the different outcomes α depends only on the symmetrized version of P_α. Therefore without loss of generality we can deal with measurement which are already symmetrized and therefore have the above decomposition over the regular representation decomposition.

As a final consequence of the symmetry of the $\rho_{\mathcal{H}}$, we recall that the algebra formed by a representation of a group (the group algebra) is a complete basis for operators which have support on the space the irreducible representations of the representation of the group act. For the right regular representation, where we have

$$R_R(g) = \bigoplus_\mu D_\mu(g) \otimes I_{d_\mu}, \tag{17}$$

then if we look at linear combinations of the operators $R_R(g)$ for all $g \in \mathcal{G}$, we find that these operators span the space of operators defined by

$$\bigoplus_\mu M_\mu \otimes I_{d_\mu}, \tag{18}$$

for all choices of M_μ. Because $\rho_{\mathcal{H}}$ has support over the irreducible representations of the right regular representations of $\mathcal{G}$ this means that

$$\rho_{\mathcal{H}} = \sum_{g\in\mathcal{G}} c_g(\mathcal{H}) R_R(g). \tag{19}$$

Notice that the dependence of the subgroup $\mathcal{H}$ is only in the coefficients of the expansion.

Furthermore we can directly evaluate these coefficients in the expansion over the right regular representation. Recall that for the right regular representation,

$$\mathrm{Tr}\left[R_R(g)\right] = |\mathcal{G}|\delta_{g,e}, \tag{20}$$

where e is the identity element of the group. This implies

$$\mathrm{Tr}\left[R_R(g)R_R^\dagger(g')\right] = \mathrm{Tr}\left[R_R((g')^{-1}g)\right] = |\mathcal{G}|\delta_{g,g'}. \tag{21}$$

Hence we find that

$$c_g(\mathcal{H}) = \frac{1}{|\mathcal{G}|}\mathrm{Tr}\left[\rho_{\mathcal{H}}R_R(g)\right]. \tag{22}$$

Next we get tricky and note that

$$\begin{aligned}
\mathrm{Tr}\left[\rho_{\mathcal{H}}R_R(g)\right] &= \mathrm{Tr}\left[\frac{|\mathcal{H}|}{|\mathcal{G}|}\sum_{g'=\text{coset rep.}} |g'\mathcal{H}\rangle\langle g'\mathcal{H}|R_R(g)\right] = \frac{|\mathcal{H}|}{|\mathcal{G}|}\sum_{g'=\text{coset rep.}}\langle g'\mathcal{H}|R_R(g)|g'\mathcal{H}\rangle \\
&= \frac{|\mathcal{H}|}{|\mathcal{G}|}\sum_{g'=\text{coset rep.}}\langle g'\mathcal{H}|g'\mathcal{H}g^{-1}\rangle,
\end{aligned} \tag{23}$$

where

$$|g'\mathcal{H}g^{-1}\rangle = \frac{1}{\sqrt{|\mathcal{H}|}}\sum_{h\in\mathcal{H}}|g'hg^{-1}\rangle. \tag{24}$$

But

$$\langle g'\mathcal{H}|g'\mathcal{H}g\rangle = \frac{1}{|\mathcal{H}|}\sum_{h_1,h_2\in\mathcal{H}}\langle g'h_1|g'h_2g\rangle = \frac{1}{|\mathcal{H}|}\sum_{h_1,h_2\in\mathcal{H}}\langle h_1|h_2g\rangle = \delta_{g,\mathcal{H}}, \tag{25}$$

where $\delta_{g,\mathcal{H}}$ is shorthand for 1 if $g\in\mathcal{H}$ and 0 otherwise. Therefore

$$c_g(\mathcal{H}) = \frac{1}{|\mathcal{G}|}\mathrm{Tr}\left[\rho_{\mathcal{H}}R_R(g)\right] = \frac{1}{|\mathcal{G}|}\delta_{g,\mathcal{H}}. \tag{26}$$

In terms of $\rho_{\mathcal{H}}$ this implies the neat little expression

$$\rho_{\mathcal{H}} = \frac{1}{|\mathcal{G}|}\sum_{h\in\mathcal{H}} R_R(h). \tag{27}$$

(An alternative way to see this is to look at $\rho_{\mathcal{H}}$ in the group basis, $\{|g\rangle\}$. Summing $R_R(h)$ over the subgroup $\mathcal{H}$ implies that $\rho_{\mathcal{H}}$ will be made up of uniform blocks, properly normalized, each block corresponding to a coset. This is nothing more than the uniform sample over coset states.)

For example, using the above expression, it is easy to check that $\rho_{\mathcal{H}}$ is a proportional to a projector,

$$\rho_{\mathcal{H}}^2 = \frac{1}{|\mathcal{G}|^2} \sum_{h_1,h_2 \in \mathcal{H}} R_R(h_1 h_2) = \frac{|\mathcal{H}|}{|\mathcal{G}|^2} \sum_{h \in \mathcal{H}} R_R(h) = \frac{|\mathcal{G}|}{|\mathcal{H}|} \rho_{\mathcal{H}}. \tag{28}$$

Finally it is useful to write $\rho_{\mathcal{H}}$ in $|\mu, r_\mu, l_\mu\rangle$ basis

$$\rho_{\mathcal{H}} = \bigoplus_\mu \left[\frac{1}{|\mathcal{G}|} \sum_{h \in \mathcal{H}} D_\mu(h) \right] \otimes I_{d_\mu}. \tag{29}$$

To recap this subsection, we have shown that the hidden subgroup state is invariant under conjugation by the left regular representation. This implies that the state is block diagonal in a basis where the right and left regular representations are fully reduced. Finally this leads to an expression for the hidden subgroup state in terms of the right regular representation, Eq. (27), which in turn leads to a simple block-diagonal expression for the hidden subgroup state, Eq. (29).

4.3 The Quantum Fourier Transform over a Finite Group

Having identified the symmetries of a hidden subgroup state $\rho_{\mathcal{H}}$ and shown that this leads to $\rho_{\mathcal{H}}$ being block diagonal in a particular basis, an important question is whether one can actually perform this basis change efficiently on a quantum computer. Indeed we haven't even identified what this change of basis is. In fact, the basis change is nothing more than the quantum fourier transform over the group $\mathcal{G}$. The quantum fourier transform (QFT) over the group $\mathcal{G}$ is defined as the unitary transform

$$Q_{\mathcal{G}} = \sum_{g \in \mathcal{G}} \sum_\mu \sum_{i,j=1}^{d_\mu} \sqrt{\frac{d_\mu}{|\mathcal{G}|}} [D_\mu(g)]_{i,j} |\mu, i, j\rangle\langle g|. \tag{30}$$

It is then easy to check, using the orthogonality relationships of irreducible representations that

$$\begin{aligned} Q_{\mathcal{G}} R_L(g) Q_{\mathcal{G}}^\dagger &= \bigoplus_\mu \sum_{i=1}^{d_\mu} |i\rangle\langle i| \otimes \sum_{j,j'=1}^{d_\mu} [D_\mu]_{j,j'}[g] |j\rangle\langle j'| \\ Q_{\mathcal{G}} R_R(g) Q_{\mathcal{G}}^\dagger &= \bigoplus_\mu \sum_{j,j'=1}^{d_\mu} [D_\mu]_{j,j'}[g] |j\rangle\langle j'| \otimes \sum_{i=1}^{d_\mu} |i\rangle\langle i|. \end{aligned} \tag{31}$$

Or in other words the QFT performs exactly the change of basis which block diagonalizes the left and right regular representations into a irreducible irreps as described in Eqs. (9).

When can the QFT over a group $\mathcal{G}$ be enacted by a quantum circuit of size polynomial in $\log |\mathcal{G}|$? While the full answer to this question is not known, for many groups, including the important symmetric group [33] and dihedral group [34, 47], efficient quantum circuits for the QFT are known. A quite general method for performing efficient QFTs over a large class of finite groups, including the Heisenberg group, has been derived by Moore, Rockmore, and

Russell [48]. We refer the reader to the latter paper for more details on the quantum Fourier transform.

Finally we should note that the QFT is a change of basis which is defined only up to the choice of the basis for the D_μ irreps. For our purposes, this basis choice will not matter, and we refer the reader to the paper of Moore, Rockmore, and Russell [48] for details on choices of this basis for different groups.

4.4 The Symmetry Exploiting Protocol

We have now seen that the QFT over the group $\mathcal{G}$ is the transform which block diagonalizes all hidden subgroup states. By the symmetry argument above we can, without loss of generality, apply this transform. Further, since $\rho_{\mathcal{H}}$ is an incoherent sum over the different irreps μ, we can, without a loss of generality, perform this transform, and then measure the irrep index $|\mu\rangle$. Further, since $\rho_{\mathcal{H}}$ acts trivially over the space where the left regular representation acts, the $|l_\mu\rangle$ index contains no information about the hidden subgroup and can also be measured (resulting in a uniformly random number between 1 and d_μ.)

Hence we see that we can recast the single copy hidden subgroup problem as

1. Perform a QFT on the hidden subgroup state.
2. Measure the irrep label μ and throw away the register where the left regular representation irrep acts. The probability of obtaining irrep μ is given by
$$p_\mu[\mathcal{H}] = \frac{d_\mu}{|\mathcal{G}|} \sum_{h\in\mathcal{H}} \chi_\mu(h), \tag{32}$$
where $\chi_\mu(g) = \mathrm{Tr}[D_\mu(g)]$ is the character of the group element g in the irrep μ.
3. One is then left with a state with support over only the space where the right regular representation irrep acts. This state is given by
$$\rho_\mu[\mathcal{H}] = \frac{d_\mu}{|\mathcal{G}|p_\mu} \sum_{h\in\mathcal{H}} D_\mu(h), \tag{33}$$
or, in other words the state
$$\rho_\mu[\mathcal{H}] = \frac{\sum_{h\in\mathcal{H}} D_\mu(h)}{\sum_{h'\in\mathcal{H}} \chi_\mu(h')}, \tag{34}$$
assuming that $p_\mu \neq 0$.

The above protocol has been designed based solely on the symmetry arguments of the hidden subgroup state. Notice that there are two locations already where information can appear about the hidden subgroup. The first is in the probabilities $p_\mu[\mathcal{H}]$. The second is in the state $\rho_\mu[\mathcal{H}]$.

Finally note that if the hidden subgroup is the trivial group $\mathcal{T} = \{e\}$, then we obtain the μth irrep with probability
$$p_\mu[\mathcal{H}] = \frac{d_\mu^2}{|\mathcal{G}|}. \tag{35}$$
Note as a sanity check that this is indeed a probability since the sum of the squares of the dimensions of all irreps is the order of the group.

4.5 Multiple Copies of the Hidden Subgroup State

As we have discussed in the introduction to the HSP, we are most interested in the setting where we use multiple copies of a hidden subgroup state to determine the hidden subgroup. In this setting, each of the hidden subgroup states will retain the symmetry we have described above. In other words our state is $\rho_{\mathcal{H}}^{\otimes m}$ and each of these can be reduced into irreducible irreps as described in Eq. (29). In this case we immediately note that there is another symmetry of these multiple copies of the same state. This symmetry is a permutation symmetry. If we permute the different copies, $\rho_{\mathcal{H}}^{\otimes m}$ is invariant. In this paper we will not explore this symmetry, only noting that measuring the irrep label alone for the Schur transform cannot be used to solve either the HSP or the HSCP. For a negative result in using permutation symmetries in trying to solve the HSP, we point the reader to the work of Childs, Harrow, and Wocjan [27].

5 Symmetry Considerations and the HSCP

We have seen in the previous section that the symmetry of the hidden subgroup states means that these states have a structure related to the left and right regular representations of the group being considered. This led, in turn, to the idea that one should exploit this structure by performing a quantum Fourier transform over the group $\mathcal{G}$ on the hidden subgroup states. In this section we will turn our attention to the hidden subgroup conjugacy problem and show how similar arguments lead one to a different transform than the QFT, the Clebsch-Gordan transform over the group $\mathcal{G}$.

5.1 Single Copy Case

Recall that in the HSCP we wish to determine which set of conjugate subgroups a hidden subgroup state belongs to. In particular we wish to design an algorithm which learns the set of conjugate subgroups hidden by f, but not the particular hidden subgroup. Stated in this manner, it is then natural to define the hidden subgroup conjugacy states which are an incoherent sum of all of the hidden subgroup states which belong to a given set of conjugate subgroups. Thus we can define the single copy hidden subgroup conjugacy state as

$$\rho_{[\mathcal{H}]} = \frac{1}{|\mathcal{G}|} \sum_{g \in \mathcal{G}} \rho_{g\mathcal{H}g^{-1}}, \tag{36}$$

where $g\mathcal{H}g^{-1} = \{ghg^{-1}, h \in \mathcal{H}\}$. Notice that we have summed over the entire group here and not over simply the different conjugate subgroups. However, since every conjugate subgroup will appear an equal number of times in this sum, and we have normalized by the size of the group, our state is equivalent to a state which is one of the elements of the set of conjugate subgroups with equal probability.

Using Eq. (27) it is easy to see that

$$\rho_{g\mathcal{H}g^{-1}} = \frac{1}{|\mathcal{G}|} \sum_{h \in \mathcal{H}} R_R(ghg^{-1}) = \frac{1}{|\mathcal{G}|} \sum_{h \in \mathcal{H}} R_R(g) R_R(h) R_R(g^{-1}), \tag{37}$$

so that

$$\rho_{[\mathcal{H}]} = \frac{1}{|\mathcal{G}|} \sum_{g \in \mathcal{G}} R_R(g) \rho_{\mathcal{H}} R_R(g^{-1}). \tag{38}$$

In other words, $\rho_{[\mathcal{H}]}$ is $\rho_{\mathcal{H}}$ symmetrized over the right regular representation. This implies that $\rho_{[\mathcal{H}]}$ is invariant under the right regular representation:

$$R_R(g)\rho_{[\mathcal{H}]}R_R(g^{-1}) = \rho_{[\mathcal{H}]}. \tag{39}$$

Thus we can again apply Schur's lemma. Since $\rho_{\mathcal{H}}$ already only has support over the right regular representation, we can deduce that $\rho_{[\mathcal{H}]}$ in the $|\mu, r_\mu, l_\mu\rangle$ basis is given by

$$\rho_{[\mathcal{H}]} = \bigoplus_{\mu} c_\mu(\mathcal{H}) I_{d_\mu} \otimes I_{d_\mu}. \tag{40}$$

We can use the character projection operator, which projects onto a given irrep space,

$$C_\mu = \frac{d_\mu}{|\mathcal{G}|} \sum_{g \in \mathcal{G}} \chi_\mu(g)^* R_R(g), \tag{41}$$

where $\chi_\mu(g)$ is the character of the μth irrep evaluated at group element g, along with our simple representation of the hidden subgroup states to determine $c_\mu(\mathcal{H})$,

$$c_\mu(\mathcal{H}) = \frac{1}{d_\mu^2} \mathrm{Tr}[C_\mu \rho_{[\mathcal{H}]}] = \frac{1}{|\mathcal{G}| d_\mu} \sum_{h \in \mathcal{H}} \chi_\mu(h)^*. \tag{42}$$

What does our expression for the hidden subgroup conjugacy state in Eq. 40 tell us? It tells us that, without loss of generality, if we are trying to solve the HSCP using a single copy of the hidden subgroup state, then we can without loss of generality perform a QFT over $\mathcal{G}$ and then measuring the irrep label μ. Further all of the information, if there is enough information to reconstruct which set of conjugate subgroups the hidden subgroup belongs to, can be derived from such a measurement.

5.2 Multiple Copy Case

Above we have argued that $\rho_{[\mathcal{H}]}$ is the appropriate state to consider when attempting to describe algorithms for the HSCP when we have been given a single copy of the hidden subgroup state. What state should we consider for the multiple copy case? At first glance one is tempted to answer $\rho_{[\mathcal{H}]}^{\otimes m}$. However, this is the state which is relevant if we are attempting to identify the which set of conjugate subgroups a hidden subgroup state belongs to, and we are given different hidden subgroup states from this set every time we produce a hidden subgroup state. This however, is not our case, as we are still querying a function which hides a fixed subgroup $\mathcal{H}$. Thus the actual state which is relevant for the HSCP is instead

$$\rho_{[\mathcal{H}],m} = \frac{1}{|\mathcal{G}|} \sum_{g \in \mathcal{G}} \rho_{g\mathcal{H}g^{-1}}^{\otimes m}. \tag{43}$$

The nice thing about Eq. (43) is that it represents a state which has been averaged over the m-fold direct product of the right regular representation. The m-fold direct product of right regular representations is the representation of $\mathcal{G}$ given by $R_R(g)^{\otimes m}$, i.e. the tensor product of the right regular representation acting in the same manner on every tensor product space. Indeed we can express Eq. (43) as

$$\rho_{[\mathcal{H}],m} = \frac{1}{|\mathcal{G}|} \sum_{g \in \mathcal{G}} R_R(g)^{\otimes m} \rho_{\mathcal{H}}^{\otimes m} R_R(g^{-1})^{\otimes m}. \tag{44}$$

This implies, as before, that $\rho_{[\mathcal{H}],m}$ is invariant under the m-fold direct product of the right regular representation,

$$R_R(g)^{\otimes m}\rho_{[\mathcal{H}],m}R_R(g^{-1})^{\otimes m}=\rho_{[\mathcal{H}],m}. \tag{45}$$

This is the symmetry which we are going to exploit to help us solve the multicopy HSCP.

In order to exploit the symmetry in Eq. (45) we begin by first rewriting $\rho_{[\mathcal{H}]}$. In particular we rewrite $\rho_{[\mathcal{H}]}$ in a basis in which every hidden subgroup state copy has been subject to a QFT over $\mathcal{G}$:

$$\rho_{[\mathcal{H}],m}=\frac{1}{|\mathcal{G}|^{k+1}}\sum_{g\in\mathcal{G}}\bigotimes_{i=1}^{m}\left[\bigoplus_{\mu_i}\sum_{h_i\in\mathcal{H}}D_{\mu_i}(gh_ig^{-1})\otimes I_{d_{\mu_i}}\right]. \tag{46}$$

Exchanging the product of sums with a sum of products we can express this as

$$\frac{1}{|\mathcal{G}|^{k+1}}\sum_{g\in\mathcal{G}}\bigoplus_{\mu_1,\dots,\mu_m}\bigotimes_{i=1}^{m}\left[\sum_{h_i\in\mathcal{H}}D_{\mu_i}(gh_ig^{-1})\otimes I_{d_{\mu_i}}\right]. \tag{47}$$

From this expression we see that we should focus on terms for a fixed choice of single copy irrep labels, $\mu_1,\dots,\mu_m$.

$$M=\frac{1}{|\mathcal{G}|^{k+1}}\sum_{g\in\mathcal{G}}\bigotimes_{i=1}^{m}\sum_{h\in\mathcal{H}}D_{\mu_i}(gh_ig^{-1}), \tag{48}$$

which is invariant under the direct product action of the group,

$$\left[\bigotimes_{i=1}^{m}D_{\mu_i}(g)\right]M\left[\bigotimes_{i'=1}^{m}D_{\mu_{i'}}(g^{-1})\right]=M. \tag{49}$$

Thus we see that the symmetry of Eq. (43) is going to be related to direct product representation $\bigotimes_{i=1}^{m}D_{\mu_i}(g)$.

What can we say about the direct product of m irreps represented by $\bigotimes_{i=1}^{m}D_{\mu_i}(g)$? Well fist we note that this is, of course, a representation of the group $\mathcal{G}$. As such it will be, in general, reducible. In other words, there exists a basis under which

$$\bigotimes_{i=1}^{m}D_{\mu_i}(g)=\bigoplus_{\mu}I_{n_\mu}\otimes D_\mu(g). \tag{50}$$

where n_μ is the multiplicity of the μth irrep in this decomposition. We will call the transform which enacts the above basis change the m-fold Clebsch-Gordan transform. Notice that this transform can act in a non-separable manner over the different spaces where each irrep lives. It is this transform which we will use below to solve the Heisenberg hidden subgroup problem.

Given the basis change described in Eq. (50), and the fact that $R_R(g)^{\otimes m}$ commutes with $\rho_{[\mathcal{H}],m}$, we can use Schur's lemma to show that $\rho_{[\mathcal{H}],m}$ has a block diagonal form related to this basis change. In particular, the state $\rho_{[\mathcal{H}],m}$ will act trivially on the space where the μ irrep acts in Eq. (50) and will only have non-trivial support over the space arising from the multiplicity of the μth irrep, i.e. the space where identity acts in Eq. (50).

To be concrete let us define the m-fold Clebsch-Gordan transform. It is the transform which takes as input the m irrep labels $|\mu_1\rangle\otimes\cdots\otimes|\mu_m\rangle$ along with the spaces upon which these irreps act $|v_i\rangle \otimes \cdots \otimes |v_m\rangle$ where $v_i = 1, \ldots, d_{\mu_i}$ and transforms this into the basis given by Eq. (50), which is $|\mu_1\rangle \otimes \cdots \otimes |\mu_m\rangle \otimes |\mu\rangle \otimes |w\rangle \otimes |v\rangle$, where μ is the irrep label for the subspaces on the RHS of Eq. (50), $w = 1, \ldots, n_\mu$ labels the multiplicity of these irreps, and $v = 1, \ldots, d_\mu$ labels the space where the μth irrep acts. Notice that we have not defined an explicit basis for $|v\rangle$ and $|w\rangle$, just as in the QFT over a finite group, there was a choice in the basis for the $|r_\mu\rangle$ and $|l_\mu\rangle$ basis. When we encounter the Clebsch-Gordan transform relevant for the Heisenberg HSCP, we will pick a particular basis.

Given the above definitions we can now describe the a protocol for distinguishing the hidden subgroup conjugacy states when we are given m copies of the hidden subgroup state just as we did for the hidden subgroup states:

1. Perform QFTs over $\mathcal{G}$ over all m hidden subgroup states.
2. Measure the irrep label for each of these m states, resulting in outcome μ_i.
3. Throw away the registers where the left regular representation irreps act.
4. Take the m remaining spaces where the right regular representation acts, along with the irrep labels, and perform a Clebsch-Gordan transform.
5. Measure the irrep label after the Clebsch-Gordan transform, μ.
6. Throw away the space where the irrep acts nontrivially in the Clebsch-Gordan transform, leaving only the space arising from the multiplicity of the μth irrep. From this remaining state, a measurement should be performed which solves the hidden subgroup conjugacy problem, along with the classical data resulting from the irrep measurements, $\mu_1, \ldots, \mu_m$ and μ.

The above protocol can be used without a loss of computational power for distinguishing hidden subgroup conjugacy states. Notice that the final step is a step which requires a judicious choice of measurement. Since the Clebsch-Gordan transform does not yield a canonical basis for the multiplicity space, this implies that all of the difficulty of distinguishing HSCP states lies in determining a basis for this measurement. Further, getting into this basis should be accomplished using an efficient quantum circuit. At this point one might wonder whether it is ever possible to make such a judicious choice. It is precisely the goal of this paper to show that such a choice is possible and can lead to efficient quantum algorithms for the HSCP.

6 The Heisenberg Group

We now turn from our general discussion of the HSP and the HSCP for any finite group $\mathcal{G}$ to a discussion where $\mathcal{G}$ is a particular group, the Heisenberg group. In this section we collect much of the relevant information about this group, its subgroups, and its irreps.

We are interested in the Heisenberg group $\mathcal{H}_p = \mathbb{Z}_p \rtimes \mathbb{Z}_p^2$ where p is prime. This is the group of upper right triangular 3×3 matrices with multiplication and addition over the field $\mathbb{F}_p$. We denote elements of this group by a three-tuple of numbers, (x, y, z), with $x, y, z \in \mathbb{Z}_p$. The multiplication rule for this group is then given by

$$(x, y, z)(x', y', z') = (x + x', y + y' + xz', z + z'). \tag{51}$$

The inverse of an element (x, y, z) of the Heisenberg group is the element $(-x, -y + xz, -z)$.

The representation theory and subgroup structure of the Heisenberg group is easily deduced. We will use (up to a change in the way we write group elements) the notation from the paper of Radhakrishnan, Rötteler and Sen [41]. The Heisenberg group has p^2 different one dimensional irreps and $p-1$ different p dimensional irreps. The one dimensional irreps are given by

$$\chi_{a,b}((x,y,z)) = \omega^{ax+bz}, \tag{52}$$

where $a, b \in \mathbb{Z}_p$ and $\omega = \exp\left(\frac{2\pi i}{p}\right)$. The p dimensional irreps are given by

$$\sigma_k((x,y,z)) = \omega^{ky} \sum_{r \in \mathbb{Z}_p} \omega^{kzr} |r + x\rangle\langle r|, \tag{53}$$

where $k \in \mathbb{Z}_p^*$. For future reference, the characters of the p dimensional irreps are

$$\chi_k((x,y,z)) = \delta_{x,0}\delta_{z,0}\omega^{ky} p. \tag{54}$$

The Heisenberg group has subgroups of four different orders, p^3, p^2, p, and 1. In Table 1 we catalog these subgroups and describe their generators. The first two groups in Table 1 are the self-explanatory group itself and the trivial subgroup. The $\mathcal{N}_\alpha$ groups are normal subgroups of $\mathcal{H}_p$, $\mathcal{N}_\alpha \lhd \mathcal{H}_p$. These groups are isomorphic to the Abelian group $\mathbb{Z}_p \times \mathbb{Z}_p$. The $\mathcal{A}_{i,j}$ subgroups are Abelian subgroups which are not normal in $\mathcal{H}_p$. These groups are subgroups of the appropriate $\mathcal{N}_i$ subgroups. In particular they are normal subgroups $\mathcal{A}_{\alpha,j} \lhd \mathcal{N}_\alpha$ and $\mathcal{A}_{\alpha,j} \not\subseteq \mathcal{N}_{\alpha'}$ if $i \neq i'$. Finally, $\mathcal{C}$ is the center of $\mathcal{H}_p$.

Table 1. Subgroups of the Heisenberg group, $\mathcal{H}_p$.

Subgroup	Generators	Label Range	Order
$\mathcal{H}_p$	$(1,0,0),(0,1,0),(0,0,1)$	-	p^3
$\mathcal{T}$	$(0,0,0)$	-	1
$\mathcal{N}_i$	$(1,0,i),(0,1,0)$	$i \in \mathbb{Z}_p$	p^2
$\mathcal{N}_\infty$	$(0,1,0),(0,0,1)$	-	p^2
$\mathcal{A}_{i,j}$	$(1,j,i)$	$i,j \in \mathbb{Z}_p$	p
$\mathcal{A}_{\infty,j}$	$(0,j,1)$	$j \in \mathbb{Z}_p$	p
$\mathcal{C}$	$(0,1,0)$	-	p

We are interested in solving the HSP for $\mathcal{H}_p$. A theorem from Bacon, Childs, and van Dam [14] (which is an extension of a theorem of Ettinger and Hoyer [40]) shows that this problem can be reduced to the hidden subgroup problem when the subgroup being hidden is restricted to being either $\mathcal{A}_{i,j}$, $i,j \in \mathbb{Z}_p$ or the trivial subgroup $\mathcal{T}$. In fact we really only need to identify $\mathcal{A}_{i,j}$, $i,j \in \mathbb{Z}_p$ since if our algorithm returns a subgroup of this form, we can easily verify, by querying the function twice, whether it is constant on left cosets this subgroup. If the hidden subgroup were a trivial subgroup, we would find the function is not constant and would then return that the hidden subgroup is trivial.

We are interested in the hidden subgroup conjugacy problem. Thus we would like to know which of the subgroups $\mathcal{A}_{i,j}$ are conjugate to each other. Conjugating a general element of $\mathcal{H}_p$ about a generator of $\mathcal{A}_{i,j}$, we find that

$$(x,y,z)(1,j,i)(x,y,z)^{-1} = (1, j + xi - z, i). \tag{55}$$

Thus we see that subgroups $\mathcal{A}_{i,j}$ and $\mathcal{A}_{i,k}$ for all $j,k \in \mathbb{Z}_p$ are conjugate to each other. Thus in the HSCP we are required to identify not i and j of $A_{i,j}$ (and distinguish this from the trivial subgroup), but only to identify i.

We are now in a position to show that solving the HSCP for the Heisenberg group can be used to solve the HSP for this group. Suppose that we have solved the HSCP for the Heisenberg group restricted to subgroups $\mathcal{A}_{i,j}$ and we identified i. We can identify j via the observation that $\mathcal{A}_{i,j}$ is normal in $\mathcal{N}_i$. In other words, if we now work with a HSP where our original f is restricted to the subgroup $\mathcal{N}_i$ (which we explicitly know, since we have identified i) then we can use the efficient quantum algorithm for finding normal subgroups [9] to find j. Alternatively (and equivalently, really) we could have simply noted that $\mathcal{N}_i$ is Abelian and thus we can run the standard HSP algorithm for Abelian groups over this to identify $\mathcal{A}_{i,j}$. Thus we see that for the Heisenberg group, the HSP and the HSCP are actually quantum polynomial time equivalent to each other.

6.1 *The Heisenberg Hidden Subgroup State*

The hidden subgroup state for the Heisenberg group can be readily calculated using the formalism described in the previous sections. We are particularly interested in a procedure where we perform a QFT over the Heisenberg group on the hidden subgroup state, and measure the irrep label, as prescribed in the protocol of Sec. 4.4.

We need to study the case that the hidden subgroup is one of the $\mathcal{A}_{i,j}$ subgroups. It is useful to note that this subgroup is the set $\mathcal{A}_{i,j} = \{(l, 2^{-1}l(l-1)i + lj, li), \forall l \in \mathbb{Z}_p\}$. Thus the probability of observing the one dimensional irrep with character $\chi_{a,b}$ is

$$p_{a,b} = \frac{1}{p^3}\sum_{l\in\mathbb{Z}_p}\omega^{al+bli} = \frac{1}{p^2}\delta_{a,bi}. \tag{56}$$

The probability of observing the p-dimensional irrep with character χ_k is

$$p_k = \frac{1}{p^3}\sum_{l\in\mathbb{Z}_p}\chi_k((l, 2^{-1}l(l-1)i + lj, li)) = \frac{1}{p}. \tag{57}$$

Comparing this to the results from the trivial subgroup, we see that the trivial and the $\mathcal{A}_{i,j}$ subgroups cannot be distinguished by the above measurement of μ. Indeed, the probability of getting one of the p dimensional irreps in both cases is uniformly $\frac{1}{p}$ and there are $p-1$ such irreps such that with probability $1-\frac{1}{p}$ the states are not distinguished using just this irrep measurement.

We are interested in more than just the probability of the measurement of the irrep label. In particular we are most interested in the exact form of the hidden subgroup state produced after we measure the irrep index and throw away the index of the left regular representation. When the hidden subgroup is $\mathcal{A}_{i,j}$ and we measure the one dimensional irrep with character $\chi_{a,b}$ we obtain, of course, a single one dimensional space with density matrix just the scalar 1. However when obtain one of the p dimensional irreps, we obtain the state over the p dimensional irrep space,

$$\rho_k(\mathcal{A}_{i,j}) = \frac{1}{p}\sum_{l\in\mathbb{Z}_p}\omega^{2^{-1}l(l-1)ik+ljk}\sum_{r\in\mathbb{Z}_p}\omega^{klir}|r+l\rangle\langle r|. \tag{58}$$

This latter state is the state which we will be most relevant to our efficient algorithm for the HSCP over the Heisenberg group.

7 The Clebsch-Gordan Transform over the Heisenberg Group

If we have two irreps of the Heisenberg group, D_{μ_1} and D_{μ_2} what are the rules for decomposing the direct product of these irreps into irreps? In other words, what is the Clebsch-Gordan transform for the 2-fold direct product of two irreps of the Heisenberg group?

The first case to consider is when both of the irreps are one dimensional irreps with characters χ_{a_1,b_1} and χ_{a_2,b_2}. In this case it is easy to see that the direct product of these irreps is a one dimensional irrep with corresponding character $\chi_{a_1+a_2,b_1+b_2}$ where the addition is done over $\mathbb{Z}_p$. The second case to consider is when one of the irreps is the one dimensional irrep with character χ_{a_1,b_1} and the second irrep is a p dimensional irrep with character χ_{k_2}. In this case it is easy to see that

$$\chi_{a_1,b_1}((x,y,z)) \otimes \sigma_{k_2}((x,y,z)) = \omega^{k_2 y + a_1 x + b_1 z} \sum_{r \in \mathbb{Z}_p} \omega^{k_2 z r} |r+x\rangle\langle r| \tag{59}$$

Now making the change of basis described by the unitary matrix

$$V = \sum_{t \in \mathbb{Z}_p} |t + k_2^{-1} b_1\rangle\langle t| \left[\sum_{s \in \mathbb{Z}_p} \omega^{-a_1 s} |s\rangle\langle s| \right], \tag{60}$$

to $\chi_{a_1,b_1}((x,y,z)) \otimes \sigma_{k_2}((x,y,z))$ yields

$$V \omega^{k_2 y + a_1 x + b_2 z} \sum_{r \in \mathbb{Z}_p} \omega^{k_2 z r} |r+x\rangle\langle r| V^\dagger = \sigma_k((x,y,z)). \tag{61}$$

Thus the direct product of the one dimensional irrep with character $\chi_{a,b}$ and the p dimensional irrep σ_k is a p dimensional irrep σ_k. Further there is no multiplicity in this decomposition.

The final case to consider is the case where one is forming the direct product of two p dimensional irreps, σ_{k_1} and σ_{k_2}. In this case we see that

$$\sigma_{k_1}((x,y,z)) \otimes \sigma_{k_2}((x,y,z)) = \sum_{r_1 \in \mathbb{Z}_p} \omega^{k_1 y + k_1 z r_1} |r_1 + x\rangle\langle r_1| \otimes \sum_{r_2 \in \mathbb{Z}_p} \omega^{k_2 y + k_2 z r_2} |r_2 + x\rangle\langle r_2|. \tag{62}$$

Consider first the case where $k_1 + k_2 \neq 0 \bmod p$. Define the following change of basis

$$W = \sum_{a,b \in \mathbb{Z}_p} |a-b\rangle\langle a| \otimes |(k_1 a + k_2 b)(k_1 + k_2)^{-1}\rangle\langle b|. \tag{63}$$

Applying this transform to $\sigma_{k_1}((x,y,z)) \otimes \sigma_{k_2}((x,y,z))$ results in

$$\sum_{r_1, r_2 \in \mathbb{Z}_p} \omega^{k' y + (k_1 r_1 + k_2 r_2) z} |r_1 - r_2\rangle\langle r_1 - r_2| \otimes |(k_1 r_2 + k_2 r_2)(k')^{-1} + x\rangle\langle (k_1 r_2 + k_2 r_2)(k')^{-1}| \tag{64}$$

where $k' = k_1 + k_2$. Using $u = r_1 - r_2$ and $v = (k_1 r_1 + k_2 r_2)(k')^{-1}$ this can be reexpressed as

$$\omega^{k' y} \sum_{u \in \mathbb{Z}_p} |u\rangle\langle u| \otimes \sum_{v \in \mathbb{Z}_p} \omega^{k' v z} |v + x\rangle\langle v| \tag{65}$$

which we see is just $I \otimes \sigma_{k'}((x,y,z))$. Thus when $k_1+k_2 \neq 0 \bmod p$, the direct product of the two p dimensional irreps labeled by k_1 and k_2 is reducible to the irrep labeled by $k' = k_1 + k_2$ with multiplicity p.

When $k_1 + k_2 = 0 \bmod p$ we find that

$$\sigma_{k_1}((x,y,z)) \otimes \sigma_{k_2}((x,y,z)) = \sum_{r_1,r_2 \in \mathbb{Z}_p} \omega^{k_1 z(r_1-r_2)} |r_1+x\rangle\langle r_1| \otimes |r_2+x\rangle\langle r_2|. \tag{66}$$

Consider the unitary change of basis

$$X = \frac{1}{\sqrt{p}} \sum_{a,b,c \in \mathbb{Z}_p} \omega^{(a+b)c} |a-b\rangle\langle a| \otimes |c\rangle\langle b|. \tag{67}$$

This transforms the direct product to

$$\sum_{u,c \in \mathbb{Z}_p} \omega^{k_1 uz + 2xc} |u\rangle\langle u| \otimes |c\rangle\langle c|, \tag{68}$$

where $u = r_1 - r_2$ and we have summed over $v = r_1 + r_2$. This we recognize as p^2 one dimensional irreps, with every such irrep appearing exactly once.

Above we have derived how the direct product of two irreps of the Heisenberg group decomposes into new irreps. Can we efficiently enact a quantum algorithm to achieve this transformation? Certainly. In fact we have done most of the heavy lifting already in identifying the transforms V, W, and X. Recall the Clebsch-Gordan transform will act from a space with two irrep label registers, $|\mu_1\rangle \otimes |\mu_2\rangle$, along with the space upon which these irreps act $|v_1\rangle \otimes |v_2\rangle$ and have an output space where the two irrep labels, $|\mu_1\rangle \otimes |\mu_2\rangle$ are kept around and the direct product irrep label $|\mu\rangle$ is produced along with the space where this irrep acts, $|v\rangle$ and the space of the multiplicity of this irrep, $|w\rangle$.

The algorithm for efficiently enacting the Clebsch-Gordan then proceeds as follows. First, notice the transform above can all be done conditionally on the $|\mu_1\rangle$ and $|\mu_2\rangle$ registers. In other words, given that we can efficiently enact the appropriate transform for fixed classical labels of $|\mu_1\rangle$ and $|\mu_2\rangle$, then we can efficiently implement the full Clebsch-Gordan transform by using the appropriate conditional gates. Thus we can divide up our algorithm into the cases we described above.

1. (μ_1 and μ_2 are one dimensional irreps) In this case it is easy to efficiently classically compute the new irrep label from the old irrep labels, a_1, b_1 and a_2, b_2. Indeed the new label is simply $a = a_1 + a_2$ and $b = b_1 + b_2$ done with addition modulo p. Further the spaces involved are all one dimensional, ($|v_1\rangle$ and $|v_2\rangle$, along with $|v\rangle \otimes |w\rangle$) Thus no other work needs to be done for this portion of the Clebsch-Gordan transform.

2. (μ_1 and μ_2 are one and p dimensional irreps) In this case the new irrep label μ is nothing more than the original irrep label of the p dimensional irrep, so such a label should be copied (reversibly added) into this register. If we wish for the irreps to be expressed in the same basis that we express the σ_k irreps, then we will need to apply the V gate to the $|v_1\rangle \otimes |v_2\rangle$ register. It is easy to see that V can be enacted using classical reversible computation plus a diagonal phase gate which is easy to efficiently enact.

3. ($\mu_1 = k_1$ and $\mu_2 = k_2$ are both p dimensional irreps, such that $k_1 + k_2 \neq 0 \bmod p$) In this case the new irrep label register will need to hold $k_1 + k_2 \bmod p$ which can easily be calculated using reversible classical circuits. In addition, the transform W must be enacted on the vector space $|v_1\rangle \otimes |v_2\rangle$. W can also be enacted using classical reversible circuits efficiently.

4. ($\mu_1 = k_1$ and $\mu_2 = k_2$ are both p dimensional irreps, such that $k_1 + k_2 = 0 \bmod p$) This is the only case where we cannot compute the new irrep label directly. Instead we must enact X, which can be done with a combination of classical reversible circuits and an efficient QFT (or approximation thereof) over $\mathbb{Z}_p$. Once this is done, the new irrep label can be found by simply transforming the vector basis labels $|u\rangle$ and $|c\rangle$ to the irrep label $|a = 2u, b = k_1 c\rangle$.

Thus we have shown that the Clebsch-Gordan transform for the Heisenberg group can be implemented with an efficient quantum circuit.

7.1 Clebsch-Gordan Transform Methods

For completeness we would here like to discuss prior results for obtaining transforms very similar to Clebsch-Gordan transforms over finite groups, and why it is important to distinguish these methods from the method we have described above. The first such method was used by Kuperberg in his algorithm for the dihedral HSP [23] (where it is listed as Proposition 9.1). The basic idea of Kuperberg's method is as follows. Suppose that we are working on the tensor product of two Hilbert spaces, each of which is spanned by a basis of the group elements $\{|g\rangle, g \in \mathcal{G}\}$. Define the following unitary operator on this tensor product space:

$$U_l(|a\rangle \otimes |b\rangle) = (|b^{-1}a\rangle \otimes |b\rangle) \tag{69}$$

where $a, b \in \mathcal{G}$. Then this operation takes left multiplication by the direct product of the group to left multiplication on the right Hilbert space. In other words,

$$U_l[R_L(g) \otimes R_L(g)]U_l^\dagger = I \otimes R_L(g). \tag{70}$$

Kuperberg uses this observation to perform *summand extraction.* Summand extraction works by taking two spaces which carry irreducible representations of a group $\mathcal{G}$ and then performing quantum Fourier transforms (and inverses) over $\mathcal{G}$ along with U_l (we refer the reader to [23] for details.) The net effect of these transforms is to perform a partial Clebsch-Gordan transform. In particular whereas in a full Clebsch-Gordan transform, one has access to registers containing the total irrep label $|\mu\rangle$, the space where the multiplicity acts $|w\rangle$, and the space where the total irrep acts $|v\rangle$ (see Section 5.2), in using this method the multiplicity register is not directly accessible. In other words there is no canonical basis which could be used to reveal the information stored in the multiplicity register. Or, to put it differently, none of the algorithmic uses to which this method has been applied use the multiplicity register and therefore arbitrary basis for this information was defined. The method of *summand extraction* is also explored by Moore, Russell and Sniady [28] where again the multiplicity register is not used.

Finally, a similar method for getting access to the irrep register and the space where the irrep acts for a Clebsch-Gordan transform is described in an early preprint version of [25], as well as in the the thesis of Harrow [49]. This method uses a generalizes the quantum phase estimation algorithm to non-Abelian groups. For our purposes, just as in summand extraction, these circuits can be used to enact a Clebsch-Gordan transform in its entirety, but there has been no work describing how information in the multiplicity register can be accessed.

We can now put our Clebsch-Gordan transform into perspective. While previous work has focused on the role of the irrep label register and the register where the representation lives, none has focused on the multiplicity space register. Our circuit for a Clebsch-Gordan transform picks out a particular basis for this multiplicity space, and, as we shall show in the next section, this will lead to an efficient algorithm for the Heisenberg HSCP.

8 The Clebsch-Gordan Transform over the Heisenberg Group and Solving the Heisenberg HSCP

We will now show that applying the Clebsch-Gordan transform over the Heisenberg group allows us, with a little bit more work, to solve the Heisenberg HSCP using two copies of the Heisenberg hidden subgroup state.

Consider the general procedure described in Section 5.2 for two copies of the Heisenberg hidden subgroup state. This procedure begins by prepare two copies of the hidden subgroup states, perform a QFT over the Heisenberg group on these states, measuring the irrep index, and throwing away the registers where the left regular representation acts. With probability $\frac{(p-1)(p-2)}{p^2}$ we will obtain a label corresponding to the k_1st irrep, σ_{k_1}, in the first hidden subgroup and k_2st irrep, σ_{k_2}, in the second hidden subgroup state with k_1 and k_2 satisfying $k_1 + k_2 \neq 0 \bmod p$. This probability is exponentially close to 1, so we can assume that we are exactly in this situation. Further we may assume that k_1 and k_2 are given with uniform probability assuming that $k_1 + k_2 \neq 0$. In this case, the two registers where the right regular representation acts contains the state

$$\rho_{k_1}(\mathcal{A}_{i,j}) \otimes \rho_{k_2}(\mathcal{A}_{i,j}). \tag{71}$$

We will now apply the Clebsch-Gordan transform to this state. Given that we are dealing with two p dimensional irreps, with $k_1 + k_2 \neq 0 \bmod p$, this is equivalent to applying the W of Eq. (63) to the state and to adding the two irrep labels k_1 and k_2 together to obtain the new irrep label, $k' = k_1 + k_2 \neq 0$. Using Eq. (58) we find that

$$\begin{aligned}
&W(\rho_{k_1}(\mathcal{A}_{i,j}) \otimes \rho_{k_2}(\mathcal{A}_{i,j}))W^\dagger \\
&= \frac{1}{p^2} \sum_{l_1,l_2,r_1,r_2 \in \mathbb{Z}_p} \omega^{k_1(2^{-1}l_1(l_1-1)i + l_1 r_1 i + l_1 k_1 j) + k_2(2^{-1}l_2(l_2-1)i + l_2 r_2 i + l_2 k_2 j)} \\
&|r_1 + l_1 - r_2 - l_2\rangle\langle r_1 - r_2| \otimes |[k_1(r_1+l_1) + k_2(r_1+l_1)](k')^{-1}\rangle\langle [k_1 r_1 + k_2 r_2](k')^{-1}|
\end{aligned} \tag{72}$$

Next, we measure the second of these registers since this register will contain no information about which conjugate subgroup the state belongs to. If we make this measurement, then we will obtain outcome $|m\rangle$ with probability p and the state remaining in the first register will be

$$\frac{1}{p} \sum_{l_1, u \in \mathcal{Z}_p} \omega^{k_1 2^{-1} l_1^2 (1 + k_1 k_2^{-1}) i + k_1 l_1 i u} |u + l_1(1 + k_2^{-1} k_1)\rangle\langle u|, \tag{73}$$

where we have used the fact that only on diagonal elements of the second register matter, so $k_1 l_1 + k_2 l_2 = 0 \bmod p$, and relabeled $u = r_1 - r_2$. Further relabeling this sum by $s_1 = u + l_1(1 + k_2^{-1} k_1)$ and $s_2 = u$ this becomes

$$\frac{1}{p} \sum_{s_1, s_2 \in \mathbb{Z}_p} \omega^{i 2^{-1} \frac{k_1 k_2}{k_1 + k_2} (s_1^2 - s_2^2)} |s_1\rangle\langle s_2|. \tag{74}$$

At this point we can begin to see how to solve the Heisenberg HSCP by making a particular measurement on this state. In particular the above density matrix corresponds to the pure state

$$\frac{1}{\sqrt{p}} \sum_{s \in \mathbb{Z}_p} \omega^{i 2^{-1} \frac{k_1 k_2}{k_1 + k_2} s^2} |s\rangle. \tag{75}$$

Suppose, now, that we had a transform U_2 which acted upon an equal superposition of the two square roots (done modulo p) of a number in the computational basis, $\frac{1}{\sqrt{2}}(|\sqrt{t}\rangle + |-\sqrt{t}\rangle)$ and produced the square number $|t\rangle$, and if it acts on the state $|0\rangle$ it produces the state $|0\rangle$. Applying this transform to the our state results in

$$\sqrt{\frac{2}{p}} \sum_{v \in \mathbb{Z}_p | v \text{ is square} \neq 0} \omega^{i 2^{-1} \frac{k_1 k_2}{k_1 + k_2} v} |v\rangle + \frac{1}{\sqrt{p}} |0\rangle \tag{76}$$

It is convenient to rewrite this state as

$$\sqrt{\frac{2}{p}} \left[\sum_{v \in \mathbb{Z}_p} \omega^{av} |v\rangle - \sum_{v \in \mathbb{Z}_p | v \text{ not square}} \omega^{av} |v\rangle \right] + \frac{1 - \sqrt{2}}{\sqrt{p}} |0\rangle \tag{77}$$

where $a = i \frac{k_1 k_2}{2(k_1 + k_2)}$. Now performing an inverse QFT over $\mathbb{Z}_p$ results in the state

$$\sqrt{2}|a\rangle - \frac{\sqrt{2}}{p} \sum_{v \in \mathbb{Z}_p | v \text{ not square}} \sum_{x \in \mathbb{Z}_p} \omega^{av - vx} |x\rangle + \frac{1 - \sqrt{2}}{p} \sum_{x \in \mathbb{Z}_p} |x\rangle \tag{78}$$

If we now measure this state, then the probability that we obtain the outcome a is

$$Pr(a) = \left[\sqrt{2} - \sum_{v \in \mathbb{Z}_p | v \text{ not square}} \frac{\sqrt{2}}{p} + \frac{1 - \sqrt{2}}{p} \right]^2 = \left[\frac{1}{\sqrt{2}} + (1 - \frac{1}{\sqrt{2}}) \frac{1}{p} \right]^2 = \frac{1}{2} + O(\frac{1}{p^2}) \tag{79}$$

Thus we see that with probability approximately $\frac{1}{2}$ a measurement will result in the outcome $|a\rangle$. Since we know k_1 and k_2, we can easily invert $a = i \frac{k_1 k_2}{k_1 + k_2}$ and find i, the label of the hidden subgroup conjugacy.

Thus, assuming that we can efficiently implement U_2 we have shown how to efficiently solve the HSCP for the Heisenberg. How do we implement U_2 efficiently? This can be done via an idea explained in the original Heisenberg hidden subgroup algorithm of Bacon, Childs, and van Dam [14]. In particular we note that there exists an efficient deterministic classical algorithm which can compute the two square roots of a number modulo p. In particular this algorithm can work with a control bit labeling which of the two square roots is returned: $V|x\rangle|b\rangle = |(-1)^b \sqrt{x}\rangle|b\rangle$ where $\sqrt{x}$ is a canonical square root. Then consider running this algorithm backwards, with an input on the control bit being $|b\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$. After this operation, perform a Hadamard on the control qubit and measure this register. If the outcome is $|0\rangle$, the applied transform on the non-control register will take $\frac{1}{\sqrt{2}}(|\sqrt{x}\rangle + |-\sqrt{x}\rangle)$ to $|x\rangle$ as desired. This will occur with probability $\frac{1}{2}$. Thus we find that we can efficiently implement U_2 with probability $\frac{1}{2}$ such that our full algorithm succeeds with probability $\frac{1}{4} + O(\frac{1}{p^2})$.

We have thus shown that the Clebsch-Gordan transform plus some quantum post processing can be used to efficiently solve the Heisenberg HSCP and hence solve the full Heisenberg HSP.

8.1 Comparison to the PGM Heisenberg HSP Algorithm

Above we have shown how to arrive at an efficient quantum algorithm for the Heisenberg HSP by focusing on the Heisenberg HSCP and using the extra symmetry of this state to motivate the use of a Clebsch-Gordan transform to solve the problem. How does this compare with the algorithm derived by Bacon, Childs, and van Dam [14]? We will see below that we have, in effect, derived a portion of this algorithm.

Here we will briefly recap the hidden subgroup algorithm of [14] for the Heisenberg group which was derived using the PGM formalism. The first step begins with the standard preparation of the hidden subgroup state. One then performs a quantum Fourier transform over $\mathbb{Z}_p^2$ on the registers corresponding to the y and z values of the group element (x, y, z), and measuring these registers, obtaining with uniform probability over $\mathbb{Z}_p$ a value for y and z. After this one is left with the state,

$$|i, j, y, z\rangle = \frac{1}{\sqrt{p}} \sum_{b \in \mathbb{Z}_p} \omega^{bjy + 2^{-1}b(b-1)iy + bjz} |b\rangle, \tag{80}$$

for random uniform $y, z \in \mathbb{Z}_p$. When we have two copies of this state, we obtain four uniform random variables, $y_1, z_1, y_2, z_2 \in \mathbb{Z}_p$ along with the state

$$|i, j, y_1, z_1\rangle \otimes |i, j, y_2, z_2\rangle = \frac{1}{p} \sum_{b_1, b_2 \in \mathbb{Z}_p} \omega^{j(b_1 y_1 + b_2 y_2) + 2^{-1} i(b_1(b_1-1)y_1 + b_2(b_2-1)y_2) + j(b_1 z_1 + b_2 z_2)} |b_1, b_2\rangle \tag{81}$$

If we rewrite this state in terms of the variables

$$\begin{aligned} u &= b_1 y_1 + b_2 y_2 \\ v &= 2^{-1} b_1 (b_1 - 1) y_1 + 2^{-1} b_2 (b_2 - 1) y_2 + b_1 z_1 + b_2 z_2 \end{aligned} \tag{82}$$

the two copy state can be expressed as

$$\frac{1}{p} \sum_{b_1, b_2 \in \mathbb{Z}_p} \omega^{uj + vi} |b_1, b_2\rangle \tag{83}$$

One then notes that if one could enact a transformation which takes the an equal superposition of the two solutions to the quadratic equations defined in Eq. (82) to their respective values of u and v, then we could transform the above state into

$$\frac{1}{p} \sum_{u, v \in \mathbb{Z}_p | u, v \text{ solve the quadratic equation}} \omega^{uj + vi} |u, v\rangle \tag{84}$$

from which a QFT over $\mathbb{Z}_p^2$ will immediately reveal i and j with probability $\frac{1}{2}$ (since only half of the values of u and v will arise in solving the quadratic equations.) Exactly such a transform which can take the superposition over two solutions to the u and v values which can be enacted using the same basic idea as that used to construct U_2 efficiently above, but instead of taking the two square roots to the square, a more complicated quadratic equation must be solved. We refer the reader to [14] for details of this construction.

From the above description of the PGM derived quantum algorithm for the Heisenberg HSP we can use the insights we have gained in solving this using a Clebsch-Gordan transform to explain the structure of the above algorithm. First note that by performing a QFT over $\mathbb{Z}_p^2$, measuring the outcome and producing the state in Eq. (80) we are, in effect, doing something very near to that of the QFT over the Heisenberg group. Notice however that in the original PGM state, we keep around two indices y and z. However it is easy to see that if we average the pure states in Eq. (80) over y, we obtain another pure state

$$|i,j,y\rangle = \frac{1}{\sqrt{p}} \sum_{b\in\mathbb{Z}_p} \omega^{2^{-1}b(b-1)iy+bjy}|b\rangle. \tag{85}$$

In other words the y register contains no information about the hidden subgroup. This is like the register containing the left regular representation not containing information about the hidden subgroup. Thus in the two copy state we could have equally well have dealt with the state

$$\frac{1}{p} \sum_{b_1,b_2\in\mathbb{Z}_p} \omega^{uj+wi}|b_1,b_2\rangle \tag{86}$$

where

$$\begin{aligned} u &= b_1y_1 + b_2y_2 \\ w &= 2^{-1}b_1(b_1-1)y_1 + 2^{-1}b_2(b_2-1)y_2. \end{aligned} \tag{87}$$

We see that y_1 and y_2 serve as the k_1 and k_2 irreps labels (with an exception occurring when $k_1 = k_2 = 0$.)

Next we note that, following our observation about how to perform the Clebsch-Gordan transform on the Heisenberg group, we can perform a basis change

$$\sum_{s,t\in\mathbb{Z}_p} |s-t\rangle\langle s| \otimes |(sy_1+ty_2)(y_1+y_2)^{-1}\rangle\langle t|, \tag{88}$$

assuming $y_1 + y_2 \neq 0$. This transforms $|i,j,y_1\rangle|i,j,y_2\rangle$ into

$$\frac{1}{p} \sum_{b_1,b_2\in\mathbb{Z}_p} \omega^{uj+wi}|b_1-b_2, (y_1b_1+y_2b_2)(y_1+y_2)^{-1}\rangle. \tag{89}$$

In the Clebsch-Gordan method we can measure the register of the irrep label. In the PGM case this corresponds to performing the above transform and then measuring the second register. Let this second register measurement outcome yield value m. This will produce the state

$$\frac{1}{\sqrt{p}} \sum_{r\in\mathbb{Z}_p} \omega^{wi}|r\rangle \otimes |m\rangle, \tag{90}$$

where $r = b_1 - b_2$ and we can express w as

$$w = 2^{-1}\left[\frac{y_1 y_2}{y_1+y_2} r^2 + m^2(y_1+y_2) + m\right], \tag{91}$$

which, factoring out the global phase dependent on m is

$$\frac{1}{\sqrt{p}} \sum_{r \in \mathbb{Z}_p} \omega^{\frac{y_1 y_2}{2(y_1+y_2)} i r^2} |r\rangle \otimes |m\rangle. \tag{92}$$

Which we see is exactly the state we obtained after the Clebsch-Gordan transform, Eq. (75). Thus we see that the effect of the Clebsch-Gordan transform is to transform the quadratic equation in Eq. (82) into a form in which the resulting quadratic equation contains only a quadratic term. So, in quite a real sense, our derivation of solving the HSCP using the Clebsch-Gordan transform leads directly to at least part of the measurement for the HSP used in the PGM approach. Thus it is best to view our Clebsch-Gordan derivation as giving a representation theoretic derivation of the PGM algorithm. Note however that it does not give quite the same derivation, since in the PGM based algorithm one obtains the full subgroup while in the Clebsch-Gordan based algorithm we only obtain which set of conjugate subgroups the hidden subgroup belongs to, using the hidden subgroup algorithm for normal subgroups to completely identify the hidden subgroup.

9 Conclusion

By focusing on the HSCP instead of the HSP, we have been able to apply symmetry arguments to multicopy algorithms which do not hold for the HSP. These symmetry arguments were exploited by performing a Clebsch-Gordan transform over the Heisenberg group and then applying an appropriate post processing measurement on the multiplicity register output of this transform. The algorithm we describe bears a great deal in common with the algorithm of Bacon, Childs, and van Dam [14]. However, we believe that an important insight into *why* the algorithm of Bacon, Childs, and van Dam works has been discovered. In particular, instead of relying on the optimal measurement criteria, we instead see that focusing on the symmetry of the multicopy HSCP leads to a change of basis which great facilitates solving the Heisenberg HSCP. In a real sense this implies that the Clebsch-Gordan transform over the Heisenberg group naturally arises for our problem. We believe that this is a new insight whose significance bears further investigation for other non-Abelian HSPs.

At this point it is useful to draw an analogy with a previous stage in quantum algorithms research. For a long time it was known that a quantum Fourier transform could be used on hidden subgroup states without destroying any of their coherence. However quantum Fourier transforms do not have a canonical basis choice for their registers and thus the question investigated by researchers was whether a good choice of basis exists for solving the hidden subgroup problem. The Clebsch-Gordan transform similarly does not have a canonical basis choice for the multiplicity register. We have shown, however, that for the case of the Heisenberg group, a particular choice of basis, *can* be used to eventually solve the hidden subgroup problem. The work performed on choosing a proper basis for the quantum Fourier transform led to the result of Hallgren *et al* [20] that multi-register measurements are needed to solve the HSP. We are hopeful that understanding the proper basis choice for the multiplicity space register of the Clebsch-Gordan transform will not yield such negative results and this work represents evidence in favor of this view.

Numerous obvious open problems arise from the above investigations. An immediate question is for what other groups can Clebsch-Gordan transforms be used to efficiently solve the HSCP or HSP? Another important open issue is over what groups can the Clebsch-Gordan transform be efficiently enacted in such a way as there is some natural basis choice for the multiplicity space register. In this regard, the subgroup adapted basis techniques of Bacon, Harrow, and Chuang [25, 26] along with those of Moore, Rockmore, and Russell [12], should be of great help. Finally, an interesting question is whether the HSP and HSCP are (classical or quantum) polynomial time equivalent to each other. While this is known to be true for certain finite groups, the general case remains open. Of particular significance is this question for the dihedral groups when the group order is not smooth.

Acknowledgements

DB is supported under ARO/NSA quantum algorithms grant number W911NSF-06-1-0379, NSF grant number 0523359, and NSF grant number 0621621. DB would like to acknowledge Andrew Childs for suggesting many improvements in an earlier draft of this paper and Thomas Decker for discussions on the Heisenberg group. DB would also like to thank John Preskill for repeatedly asking DB whether Clebsch-Gordan transforms are actually useful for quantum algorithms.

References

1. P. W. Shor. Algorithms for quantum computation: Discrete log and factoring. In S. Goldwasser, editor, *Proceedings of the 35th Annual Symposium on the Foundations of Computer Science*, pages 124–134, Los Alamitos, CA, 1994. IEEE Computer Society.
2. R. Crandall and C. Pomerance. *Prime Numbers: A Computational Perspective*. Springer, 2001.
3. L. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the 28th Annual ACM Symposium on the the Theory of Computation*, pages 212–219, New York, 1996. ACM Press.
4. S. Hallgren. Polynomial-time quantum algorithms for pell's equation and the principal ideal problem. In *Proceedings of the 34lth Annual ACM Symposium on the Theory of Computation*, pages 653–658, New York, 2002. ACM Press.
5. A. M. Childs, R. Cleve, E. Deotto, E. Farhi, S. Gutmann, and D. A. Spielman. Exponential algorithmic speedup by quantum walk. In *Proc. 35th ACM Symposium on Theory of Computing*, pages 59–68, 2003, arXiv:quant-ph/0209131.
6. R. Boneh and R. Lipton. Quantum cryptoanalysis of hidden linear functions. *Lecture Notes in Computer Science*, 963:424–437, 1995.
7. M. Ettinger and P. Høyer. A quantum observable for the graph isomorphism problem. 1999, arXiv:quant-ph/9901029.
8. M. Grigni, L. Schulman, M. Vazirani, and U. Vazirani. Quantum mechanical algorithms for the nonabelian hidden subgroup problem. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing*, pages 68–74, New York, NY, 2001. ACM Press.
9. S. Hallgren, A. Russell, and A. Ta-Shma. Normal subgroup reconstruction and quantum computing using group representations. In *Proceedings of the 32nd annual ACM symposium on Theory of computing*, pages 627–635, New York, NY, 2000. ACM Press.
10. G. Ivanyos, F. Magniez, and Miklos Santha. Efficient quantum algorithms for some instances of the non-abelian hidden subgroup problem. In *Proceedings of the 13th Annual ACM Symposium*

on Parallel Algorithms and Architectures, pages 263–270, New York, NY, 2001. ACM Press.
11. K. Friedl, G. Ivanyos, F. Magniez, M. Santha, and P. Sen. Hidden translation and orbit coset in quantum computing. In *Proceedings of 35th ACM Symposium on Theory of Computing*, pages 1–9, New York, NY, 2003. ACM Press.
12. C. Moore, D. N. Rockmore, A. Russell, and L. J. Schulman. The power of basis selection in Fourier sampling: hidden subgroup problems in affine groups. In *Proceedings of the 15th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1113–1122, Philadelphia, PA, 2004. ACM.
13. D. Gavinsky. Quantum solution to the hidden subgroup problem for poly-near-hamiltonian groups. *Quantum Inform. Compu.*, 4:229–235, 2004.
14. D. Bacon, A. M. Childs, and W. van Dam. From optimal measurement to efficient quantum algorithms for the hidden subgroup problem over semidirect product groups. In *Proceedings of the 46th Annual Symposium on Foundations of Computer Science*, page 469, Los Alamitos, California, 2005. IEEE.
15. A. M. Childs and W. van Dam. Quantum algorithm for a generalized hidden shift problem. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 1225–1232, Philadelphia, PA, USA, 2007. Society for Industrial and Applied Mathematics, arXiv:quant-ph/0507190.
16. O. Regev. Quantum computation and lattice problems. In *43rd Symposium on Foundations of Computer Science*, pages 520–529. IEEE Computer Society, 2002.
17. C. Moore, A. Russell, and L. J. Schulman. The symmetric group defies strong fourier sampling. In *Proc. of the 46th Ann. IEEE Symp. on Found. of Comp. Sci.*, pages 479–490, Washington, DC, USA, 2005. IEEE Computer Society.
18. C. Moore and A. Russell. The symmetric group defies strong fourier sampling: Part ii, 2005, arXiv:quant-ph/0501066.
19. D. Bacon, A.M. Childs, and W. vam Dam. Optimal measurements for the dihedral hidden subgroup problem. *Chicago Journal of Theoretical Computer Science*, 2, 2006, arXiv:quant-ph/0501044.
20. S. Hallgren, C. Moore, M. Rötteler, A. Russell, and P. Sen. Limitations of quantum coset states for graph isomorphism. In *Proc. of the 38th Ann. ACM Symp. on Theory of Comp.*, pages 604–617, New York, NY, 2006.
21. L. Ip. Shor's algorithm is optimal, 2003. `http://lawrenceip.com/papers/hspsdpabstract.html`.
22. C. Moore and A. Russell. For distinguishing conjugate hidden subgroups, the pretty good measurement is as good as it gets, 2005, arXiv:quant-ph/0507190.
23. G. Kuperberg. A subexponential-time quantum algorithm for the dihedral hidden subgroup. *SIAM Journal of Computing*, 35:170–188, 2005, arXiv:quant-ph/0302112.
24. O. Regev. A subexponential time algorithm for the dihedral hidden subgroup problem with polynomial space. 2004, arXiv:quant-ph/0406151.
25. D. Bacon, I. L. Chuang, and A. W. Harrow. Efficient quantum circuits for Schur and Clebsch-Gordan transforms. *Phys. Rev. Lett.*, 97:170502, 2006, arXiv:quant-ph/0407082.
26. D. Bacon, I. L. Chuang, and A. W. Harrow. The quantum Schur transform: I. Efficient qudit circuits. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 1235–1244, Philadelphia, PA, 2007, SIAM, arXiv:quant-ph/0601001.
27. A. M. Childs, A. W. Harrow, and P. Wocjan. Weak fourier-schur sampling, the hidden subgroup problem, and the quantum collision problem, 2006, arXiv:quant-ph/0609110. to appear in STACS 2007.
28. C. Moore, A. Russell, and P. Śniady. On the impossibility of a quantum sieve algorithm for graph isomorphism: unconditional results. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 536–545, New York, NY, 2007, ACM, arXiv:quant-ph/0612089.
29. A. Kitaev. Quantum measurements and the abelian stabilizer problem, 1995, arXiv:quant-ph/9511026.
30. D. Simon. On the power of quantum computation. In S. Goldwasser, editor, *Proceedings of the 35th Annual Symposium on the Foundations of Computer Science*, pages 116–123, Los Alamitos,

CA, 1994. IEEE Computer Society.

31. D. Deutsch and R. Jozsa. Rapid solution of problems by quantum computation. *Proc. Roy. Soc. London Ser. A*, 439:553–558, 1992.
32. E. Bernstein and U. Vazirani. Quantum complexity theory. In *SIAM J. Comput.*, volume 26, pages 1411–1473, 1997.
33. R. Beals. Quantum computation of Fourier transforms over symmetric groups. In *Proceedings of the 29th Annual ACM Symposium on Theory of Computing*, pages 48–53, New York, NY, May 1997. ACM Press.
34. P. Høyer. Efficient quantum transforms. 1997, arXiv:quant-ph/9702028.
35. J. Köbler, U. Schöning, and J. Torán. *The Graph Isomorphism Problem: Its Structural Complexity.* Birkhäuser, 1993.
36. M. Rötteler and T. Beth. Polynomial-time solution to the hidden subgroup problem for a class of non-abelian groups. *Lecture Notes in Computer Science*, 1719:148–159, 1998.
37. Y. Inui and F. Le Gall. An efficient algorithm for the hidden subgroup problem over a class of semi-direct product groups, 2004, arXiv:quant-ph/0412033.
38. G. Alagic, C. Moore, and A. Russell. Subexponential-time algorithms for hidden subgroup problems over product groups. In *Proc. 18th ACM-SIAM Symposium on Discrete Algorithms*, 2006, arXiv:quant-ph/0603251.
39. M. Ettinger, P. Høyer, and E. Knill. Hidden subgroup states are almost orthogonal. 1997, arXiv:quant-ph/9901034.
40. M. Ettinger and P. Høyer. On quantum algorithms for noncommutative hidden subgroups. *Advances in Applied Mathematics*, 25:239–251, 2000.
41. J. Radhakrishnan, M. Röetteler, and P. Sen. On the power of random bases in fourier sampling: Hidden subgroup problem in the heisenberg group. In *Lecture Notes in Computer Science*, volume 3580, pages 1399–1411, Berlin, 2005. Springer Verlag, arXiv:quant-ph/0503114.
42. A. S. Holevo. Bounds for the quantity of information transmittable by a quantum communications channel. *Problemy Peredači Informacii*, 9:3–11, 1973. English translation: A. S. Holevo, Problems of Information Transmission **9**, 177, (1973).
43. H. P. Yuen, R. S. Kennedy, and M. Lax. Optimum testing of multiple hypotheses in quantum detection theory. *IEEE Trans. Inform. Theory*, 21:125–134, 1975.
44. P. Hausladen and W. K. Wootters. A pretty good measurement for distinguishing quantum states. *J. Mod. Optics*, 41:2385–2390, 1994.
45. S. Fenner and Y. Zhang. The central nature of the hidden subgroup problem. 2006, arXiv:quant-ph/0610086.
46. J. F. Cornwell. *Group Theory in Physics: An Introduction.* Academic Press, San Diego, CA, 1997.
47. M. Puschel, M. Rotteler, and T. Beth. Fast quantum fourier transforms for a class of non-abelian groups. *Lecture Notes in Computer Science*, 1719:148–159, 1999.
48. C. Moore, D. Rockmore, and A. Russell. Generic quantum fourier transforms. In *Proceedings of the 15th annual ACM-SIAM symposium on Discrete algorithms*, pages 778–787, Philadelphia, PA, 2004. SIAM.
49. A. W. Harrow. *Applications of coherent classical communication and the Schur transform to quantum information theory.* PhD thesis, Massachusetts Institute of Technology, Cambridge, MA, 2005.