

Analysing the Quantum Fourier Transform for finite groups through the Hidden Subgroup Problem

Jean-Noël Murphy

School of Computer Science

McGill University, Montreal

November 2001

A Thesis submitted to the Faculty of Graduate Studies and
Research in partial fulfillment of the requirements of the degree
of Master of Science

©Jean-Noël Murphy 2001

Abstract

We present an in-depth study of the Quantum Fourier Transform for finite groups and the underlying mathematics. The study includes a look at the most salient results linking the Quantum Fourier Transform to the Hidden Subgroup Problem. This provides a useful context for determining the extent to which the Fourier transform can serve to recognize periodicity of a function on a finite group.

Résumé

Cette thèse vise à analyser en profondeur la transformée de Fourier quantique pour groupes finis et la théorie mathématique sur laquelle elle est construite. Elle comporte également une étude des résultats les plus importants parmi ceux qui font un lien entre la transformée de Fourier et le problème du sous-groupe caché. Ce dernier fournit un contexte utile pour déterminer l'aptitude de la transformée de Fourier à reconnaître une fonction périodique définie sur un groupe fini.

Acknowledgments

My deepest gratitude to my supervisors Gilles Brassard and Claude Crépeau. From them I received unique guidance and support. Their efforts on my behalf were truly extraordinary. Thank you for all your help and patience. It was a great honour to work with you both.

My thanks also to all the members of the Cryptography and Quantum Information Laboratory of McGill University.

Contents

1	Introduction	5
2	Theoretical Background	8
2.1	Representation Theory	8
2.1.1	Decomposing Representations	10
2.1.2	Unitary Representations	12
2.1.3	Characters	12
2.1.4	Two Important Representations	16
2.1.5	Fourier Coefficients	18
2.2	Representations of Abelian Groups	19
2.3	The Group Algebra	24
3	The Quantum Fourier Transform	26
4	The Hidden Subgroup Problem	30
4.1	Abelian Groups	33
4.2	Groups in General	35
4.2.1	Measuring the Representation Name	36
4.2.2	Measuring the Fourier Coefficients	41

5	Dihedral Groups and Their Hidden Subgroups	44
5.1	The Representations of the Dihedral Group	46
5.2	The QFT for the Dihedral Group	48
5.3	Examples for Dihedral Hidden Subgroup Problem	49
5.3.1	Finding Normal Subgroups	50
5.3.2	Distinguishing Between $ H = 1$ and $ H = 2$	51
6	Conclusion	56
A	Group Theory	62
A.1	Subgroups	63
A.2	Isomorphisms and Homomorphisms	65
A.3	Group Actions	67
B	Concepts from Linear Algebra	69
B.1	Direct Sums	69
B.2	Unitary Matrices	70

Chapter 1

Introduction

In the manipulation of quantum information, the obstacle to be overcome is always the creation of suitable interference between logical paths of the computation. Perhaps the most important mechanism of all for accomplishing this to date is the Quantum Fourier Transform for finite groups (QFT). Of course, the QFT is not a single transformation but a family of these with the form depending on the group upon which it is acting. For example, the initiated reader will surely recognize the omnipresent Walsh-Hadamard transform which is in fact the QFT for the group \mathbb{Z}_2 . Considering its prevalence, the current literature in quantum computing deals somewhat inadequately with the derivation from representation theory of the Quantum Fourier Transform particularly in its general form. A lack of details is aggravated further by the absence of an established convention. Mathematical textbooks on the subject of representation theory date from nearly forty years ago and use notation which differs from current conventions. Furthermore, they make no reference to the QFT which is a more contemporary subject. Other available sources are those which discuss the *classical* Fourier

Transform. The result is that complete information about the QFT can only be obtained by laboriously piecing together obscure mathematical theories and casual remarks.

In this thesis we seek to remedy this by rigorously defining the Quantum Fourier Transform for finite groups in general and reconciling the different definitions which appear in current research. With this done, we will move on to analysing the QFT. We will do this in a unique way by considering the Hidden Subgroup Problem (HSP). The Hidden Subgroup Problem is one which has been studied extensively and in every instance where an efficient quantum algorithm has been found, the mechanism solving it is the Quantum Fourier Transform. We submit that the relationship between the two is even stronger, that all the power of the QFT for finite groups is embodied in the ability to deal with instances of the Hidden Subgroup Problem. Therefore, it is by studying the HSP that we arrive at a deep understanding of the limits of this power. For there do exist many instances of the HSP for which the Quantum Fourier Transform appears to be inadequate.

This thesis is organized in the following way. As promised, a great deal of time will be spent discussing the mathematical theory upon which this tool is built. This is very important if new applications are to be found, and is the content of chapter 2 in this thesis. Chapter 3 will give the general definition of the Quantum Fourier Transform based on this theory, as well as some initial observations. It is in chapter 4 that we will introduce and define the Hidden Subgroup Problem. A thorough look at how the QFT can sometimes solve the problem and also at where it fails will convince the reader of the interplay between the two. The final chapter concentrates on the special case the dihedral Hidden Subgroup Problem. By introducing

a new concrete discussion of how the Quantum Fourier Transform fails to resolve the problem, the reader will see the limitations of the QFT in an illuminating context.

Chapter 2

Theoretical Background

The Fourier transform for finite groups is based upon a branch of mathematics called representation theory. It was originally conceived as a means of making the tools of linear algebra available to the study of groups. Particularly, the aim is to study the structure of a group indirectly through its various homomorphic images in matrix groups. Much study was done on this topic in the mid twentieth century, but it is no longer fashionable amongst mathematicians. Nonetheless, a relation to quantum mechanics has kept the subject alive primarily in the work of theoretical physicists. The topics covered are not overly complicated mathematically, although the reader may like to look to the appendices for some additional material.

2.1 Representation Theory

We begin by establishing the notion of a representation.

Definition Let G be a group and V a vector space over a field \mathbb{K} . A homomorphism $\rho : G \rightarrow GL_{\mathbb{K}}(V)$ is called a *representation* of G . The

representation ρ is said to have *representation space* V and *dimension* d_ρ equal to the dimension of V .

By the above definition we see that a representation is a group homomorphism that sends every element of a group G to a matrix of a linear transformation from V to V . For our purposes we will always take the field \mathbb{K} to be the complex numbers. This deceptively simple idea is the basis of the entire theory. The remainder of this section will be entirely devoted to exploring the consequences of this one definition.

As ρ is a homomorphism, it is immediate that $\text{im } \rho$ is a subgroup of $GL_{\mathbb{C}}(V)$ and that $\ker \rho$ is a subgroup of G (readers not familiar with this notation can find an explanation in the appendices). Any information about the structure of G is contained in these two groups. There are an infinite number of representations of a given finite group, fortunately we need not consider them all. To begin with, given that the image of every representation is a matrix group, we wish to reduce to a single representative all those representations whose images are isomorphic. There is nothing to be learned by examining multiple copies of the same group. So we collect representations into classes according to this property.

Two representations $\rho : G \rightarrow GL_{\mathbb{C}}(V)$ and $\tau : G \rightarrow GL_{\mathbb{C}}(W)$ are said to be *equivalent* if there exists a linear transformation $T : V \rightarrow W$ such that,

$$T \circ \rho(g) = \tau(g) \circ T \quad \forall g \in G.$$

As T is a linear transformation, it represents a change from some basis B_V of V to a basis B_W of W . As a consequence we can say that two representations are equivalent if they differ only by a change of basis, with the understanding that this is not limited to within a single vector space.

It will be possible to further limit the field of distinct representations of a group to a number of “elementary” representations. Showing how this is done and that it is sufficient is the work of the next subsection.

2.1.1 Decomposing Representations

Let $\rho : G \rightarrow GL_{\mathbb{C}}(V)$ be a group representation. We may treat ρ as defining an action of G on the elements of V and write $g \cdot v$ to represent the matrix product $\rho(g)v$, for all $g \in G$ and $v \in V$. This point of view leads to the concept of a G -invariant subspace, which is a subspace $W \subseteq V$ such that,

$$\forall g \in G \quad g \cdot W = \rho(g)W \subseteq W$$

Clearly, for any representation with representation space V , both V and the trivial subspace of V are G -invariant. However, when the representation space V of ρ possesses a proper G -invariant subspace W , then by restricting the action of ρ to one on W only, we define a new representation, denoted $\rho \downarrow_W^V$. This representation will now send elements of G to the group $GL_{\mathbb{C}}(W)$. As, W is G -invariant, each $\rho(g)$ is indeed a linear transformation from W to W , as it was one before from V to V .

In addition, given such a proper G -invariant subspace $W \subset V$, there exists another proper G -invariant subspace W' , called the orthogonal complement of W , which is independent of W . This allows us to define a second representation, the restriction of ρ to W' , $\rho \downarrow_{W'}^V$. As W and W' are independent and furthermore span V , then we can conclude that V is the direct sum of the subspaces W and W' . That is, $V = W \oplus W'$. Now with $\rho_1 = \rho \downarrow_W^V$ and $\rho_2 = \rho \downarrow_{W'}^V$, it is possible to choose a basis for which ρ will have the *block*

diagonal form,

$$\rho(g) = \left[\begin{array}{c|c} \rho_1(g) & 0 \\ \hline 0 & \rho_2(g) \end{array} \right] \quad \forall g \in G$$

Therefore, when the representation space of $\rho : G \rightarrow GL_{\mathbb{C}}(V)$ is the direct sum of G -invariant subspaces W and W' , then we say that ρ is the direct sum of the *representations* $\rho \downarrow_W^V$ and $\rho \downarrow_{W'}^V$, and write,

$$\rho = \rho \downarrow_W^V \oplus \rho \downarrow_{W'}^V$$

So, if we are able to find a proper G -invariant subspace of the representation space, we may decompose a representation into ones of smaller dimension. Obviously this process cannot continue indefinitely. At some point we must arrive at representations over vector spaces which contain no proper G -invariant subspaces. This brings us to the notion of an *irreducible* representation. A representation $\rho : G \rightarrow GL_{\mathbb{C}}(V)$ is said to be irreducible if, for any subspace $W \subseteq V$,

W is G -invariant $\Leftrightarrow W$ is either equal to V or is the trivial subspace of V .

In particular, any one-dimensional representation must be irreducible.

The full consequence of what we have just seen cannot be yet apparent. However, the following theorem will help to convince the reader of the importance of irreducible representations.

Theorem 1 (Mascke) Every representation can be decomposed into a direct sum of irreducible representations.

We will later add to the above by showing that there are only a finite number of irreducible representations and that these have many remarkable properties.

2.1.2 Unitary Representations

A *unitary* representation is a homomorphism $\rho : G \rightarrow GL_{\mathbb{C}}(V)$ such that for each $g \in G$, $\rho(g)$ is a unitary matrix. Now, for any group representation τ there exists a basis for V for which every matrix $\tau(g)$ will be unitary. Thus every representation is *equivalent* to a unitary representation. This is a consequence of the fact that every subgroup of the general linear group is isomorphic to a subgroup of the unitary group. This will allow us to prove a number of interesting general properties of representations by using the characteristics of unitary matrices. Particularly so when we discuss *characters*. For the moment we limit ourselves to noticing that as every unitary matrix is diagonalizable, by the stated equivalence, so must every matrix $\tau(g)$ of any group representation τ .

2.1.3 Characters

Definition Let $\rho : G \rightarrow GL_{\mathbb{C}}(V)$ be a representation of the group G . Define a function $\chi_{\rho} : G \rightarrow \mathbb{C}$ such that,

$$\chi_{\rho}(g) = \text{trace}(\rho(g)) \quad \forall g \in G$$

χ_{ρ} is known as the *character* of the representation ρ .

As the trace of a matrix is independent of the choice of basis, equivalent representations will have the same character. This will allow us to pick the most convenient basis for every representation, and indeed that which is most appropriate for each individual matrix $\rho(g)$, whenever we are dealing with characters. This allows for many nice properties of characters of which we will make extensive use.

As every representation $\rho : G \rightarrow GL_{\mathbb{C}}(V)$ is by definition a homomorphism, then $\rho(e)$ must be the identity matrix in the vector space V . It follows then that for every representation ρ , $\chi_{\rho}(e) = d_{\rho}$.

Recall that for every $g \in G$ and representation $\rho : G \rightarrow GL_{\mathbb{C}}(V)$ it is possible to choose a basis B_g of V so that $\rho(g)$ is diagonal. That is,

$$[\rho(g)]_{B_g} = \begin{bmatrix} \lambda_1 & & & \\ & \lambda_2 & & \\ & & \ddots & \\ & & & \lambda_n \end{bmatrix}$$

where $\lambda_1, \lambda_2, \dots, \lambda_n$ are the eigenvalues of the matrix $\rho(g)$ (the eigenvalues do not depend on the basis). Since the trace is independent of the basis, we have $\chi_{\rho}(g) = \text{tr}([\rho(g)]_{B_g})$ and so,

$$\chi_{\rho}(g) = \lambda_1 + \lambda_2 + \dots + \lambda_n$$

A further property of ρ as a homomorphism is that $\rho(hgh^{-1}) = \rho(h)\rho(g)\rho(h)^{-1}$. As we have assumed that representations are unitary, $\rho(h)\rho(g)\rho(h)^{-1}$ can be thought of as a change of basis for the representation space of the matrix $\rho(g)$. Therefore, $\chi_{\rho}(hgh^{-1}) = \chi_{\rho}(g)$, $\forall h, g \in G$. This means that characters are what are known as *class functions* or in other words they are constant on the conjugacy classes of the group. Although not a particularly captivating property at first, it will be returned to later.

As G is a finite group, each element $g \in G$ has finite order. This means that for some integer r , $g^r = 1$. As ρ is a homomorphism, $\rho(g)^r$ is the identity

matrix. Thus in an appropriate basis,

$$\rho(g)^r = \begin{bmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_n \end{bmatrix}^r = \begin{bmatrix} \lambda_1^r & & \\ & \ddots & \\ & & \lambda_n^r \end{bmatrix} = \begin{bmatrix} 1 & & \\ & \ddots & \\ & & 1 \end{bmatrix}$$

Which tells us that for $i = 1, \dots, n$, $\lambda_i^r = 1$. Put another way, the eigenvalues of the matrices $\rho(g)$ are all r^{th} -roots of unity. Now, as $\chi_\rho(g^{-1}) = \text{tr}(\rho(g^{-1})) = \text{tr}(\rho(g)^{-1})$, then again the diagonalisability of $\rho(g)$ allows us to say that $\chi_\rho(g^{-1}) = \lambda_1^{-1} + \dots + \lambda_n^{-1}$. Finally as each λ_i is a root of unity, $\lambda_i^{-1} = \bar{\lambda}_i$. We conclude that

$$\chi_\rho(g^{-1}) = \bar{\lambda}_1 + \dots + \bar{\lambda}_n = \overline{\chi_\rho(g)}$$

This is useful as there is a standard inner product for complex functions over a group. If $f : G \rightarrow \mathbb{C}$ and $g : G \rightarrow \mathbb{C}$ then

$$\langle f|g \rangle = \frac{1}{|G|} \sum_{x \in G} f(x) \overline{g(x)}$$

For characters, which are also complex functions over the group, we then have,

$$\langle \chi_\rho | \chi_\tau \rangle = \frac{1}{|G|} \sum_{x \in G} \chi_\rho(x) \overline{\chi_\tau(x)}$$

This inner product is crucial in manipulating both characters and their representations. Its usefulness stems mostly from the following theorem:

Theorem 2 Characters of irreducible representations are orthonormal with respect to the given inner product. That is, let χ_ρ and χ_τ be the characters of the inequivalent irreducible representations ρ and τ respectively. Then,

$$(i) \quad \langle \chi_\rho | \chi_\rho \rangle = \langle \chi_\tau | \chi_\tau \rangle = 1$$

$$(ii) \langle \chi_\rho | \chi_\tau \rangle = \langle \chi_\tau | \chi_\rho \rangle = 0$$

This is a beautiful result with very important consequences. The first is that the above leads to the conclusion that characters form a basis for the space of all class functions. Another basis for this space would be the set of indicator functions for the conjugacy classes of the group. So, the number of distinct irreducible characters is exactly the number of conjugacy classes. Therefore we have fixed the number of inequivalent irreducible representations to a precise finite quantity. Secondly, it provides us with a precise way to decompose representations into direct sums of irreducibles.

Let $\{\rho_1, \rho_2, \dots, \rho_k\}$ be a complete set of inequivalent irreducible representations of a group G and let χ_i be the character of ρ_i . Then, any representation τ of G can be decomposed in a unique way as,

$$\tau = \langle \chi_\tau | \chi_1 \rangle \rho_1 \oplus \langle \chi_\tau | \chi_2 \rangle \rho_2 \oplus \dots \oplus \langle \chi_\tau | \chi_k \rangle \rho_k$$

Thus, the inner product $n_i = \langle \chi_\tau | \chi_i \rangle$ gives the number of times that the irreducible representation ρ_i appears in the decomposition of τ . Furthermore, the same relationship applies to the characters. By taking the trace on both sides of the previous equation, we obtain,

$$\chi_\tau = n_1 \chi_1 + n_2 \chi_2 + \dots + n_k \chi_k$$

We can use this to show $\langle \chi | \chi \rangle = 1$ is not only necessary for irreducible characters, but sufficient.

$$\begin{aligned} \chi &= n_1 \chi_1 + n_2 \chi_2 + \dots + n_k \chi_k \\ \Leftrightarrow \langle \chi | \chi \rangle &= n_1^2 + n_2^2 + \dots + n_k^2 \end{aligned}$$

This last equation being a consequence of the orthogonality of the irreducible characters. As each n_i is a non-negative integer, the inner product can only

equal 1 if there exist a single $n_i = 1$ and all the rest are zero. We are now equipped with a much easier way of testing irreducibility than the previous method of searching for G -invariant subspaces.

Now, take two representations ρ and ρ' with identical characters and let them be decomposed into irreducibles such that,

$$\begin{aligned}\rho &= n_1\rho_1 \oplus \dots \oplus n_k\rho_k \\ \rho' &= n'_1\rho_1 \oplus \dots \oplus n'_k\rho_k\end{aligned}$$

The assumption that the two representations have identical characters implies,

$$n_1\chi_1 + \dots + n_k\chi_k = n'_1\chi_1 + \dots + n'_k\chi_k$$

The orthogonality of irreducible characters then requires that $n'_i = n_i$ for all $i = 1, \dots, k$. Note that this is only true because the decomposition is over the set of inequivalent irreducible representations. The representations decompose identically into a direct sum of irreducibles and hence they must be equivalent. With the property that characters are independent of changes of basis we conclude,

Theorem 3 Two representations are equivalent if and only if their characters are identical.

This theorem tells us that characters completely determine representations up to equivalence. Together with the stunning properties they possess we have a greatly simplified way of dealing with group representations.

2.1.4 Two Important Representations

One of the most useful representations of a group is also the simplest. This is the function $1_G : G \rightarrow GL_{\mathbb{C}}$, called the *trivial representation*, which sends

every element of the group to the one-dimensional matrix (1). Of course, χ_{1_G} is identically 1 over all group elements, which makes it clear that the trivial representation is always irreducible. As the title of this section suggests there is another representation that we will define here, but this will require a somewhat longer explanation.

Previously we hinted at the possibility of considering a representation as the action of a group on a vector space. We now return to the notion of group action in order to define the so-called *permutation* representations.

Suppose that the group G acts on some set $S = \{s_1, s_2, \dots, s_n\}$ by permutation. We define a vector space on S considering all complex linear combinations of the elements of S ,

$$v = \sum_{i=1}^n \alpha_i s_i$$

Now let us define a representation of G , $perm_S : G \rightarrow V(S)$ such that $perm_S(g)$ sends this basis (s_1, s_2, \dots, s_n) to the basis $(gs_1, gs_2, \dots, gs_n)$. That is,

$$perm_S(g)v = \sum_{i=1}^n \alpha_i (gs_i) \quad \forall v \in V(S), g \in G$$

In other words, each $perm_S(g)$ is a $n \times n$ permutation matrix. Further study of this representation reveals that its character χ_S has the property that $\chi_S(g)$ is the number of elements of S fixed by the action of g . Particular choices of the set S lead to different representations. Among these, if we set $S = G$ and let G act on itself by left multiplication, we obtain what is known as the *regular representation*, ρ_{reg} . Note that for the action of left multiplication, the identity element fixes all elements of the group and all

other group elements fix none. We deduce that,

$$\chi_{reg}(g) = \begin{cases} |G| & \text{if } g = e \\ 0 & \text{otherwise} \end{cases}$$

As a consequence, for any group representation ρ with character χ_ρ ,

$$\langle \chi_{reg} | \chi_\rho \rangle = \frac{1}{|G|} \sum_{g \in G} \chi_{reg}(g) \chi_\rho(g^{-1}) = \frac{1}{|G|} (|G| \chi_\rho(e)) = d_\rho$$

Therefore, if $\rho_1, \rho_2, \dots, \rho_k$ is a complete set of inequivalent irreducibles of the group G with characters χ_i and dimensions d_i , when we decompose the regular representation in the standard way, we arrive at,

$$\rho_{reg} = \langle \chi_{reg} | \chi_1 \rangle \rho_1 \oplus \dots \oplus \langle \chi_{reg} | \chi_k \rangle \rho_k$$

which, by what we have just seen, gives

$$\rho_{reg} = d_1 \rho_1 \oplus d_2 \rho_2 \oplus \dots \oplus d_k \rho_k$$

Naturally, this also implies a similar relationship between the character of the regular representation and the irreducible characters. This is particularly interesting in the values assumed at the identity.

$$\begin{aligned} \chi_{reg}(e) &= d_1 \chi_1(e) + d_2 \chi_2(e) + \dots + d_k \chi_k(e) \\ |G| &= \sum_{i=1}^k d_i^2 \end{aligned}$$

Thus we make the general statement that for any complete set R of inequivalent irreducible representations of a group G , $\sum_{\rho \in R} d_\rho^2 = |G|$.

2.1.5 Fourier Coefficients

Let $R = \{\rho_1, \rho_2, \dots, \rho_k\}$ be a complete set of inequivalent irreducible representations of the group G . For each representation $\rho \in R$ and $g \in G$, $\rho(g)$

is a $d_\rho \times d_\rho$ matrix with complex entries. Therefore, it is possible to define a family of functions $\rho_{ij} : G \rightarrow \mathbb{C}$, called the *Fourier coefficients* which will send every element g to the $(i, j)^{th}$ entry of the matrix $\rho(g)$. That is,

$$\rho_{ij}(g) = (\rho(g))_{i,j} \quad \text{for all } g \in G \text{ and } 1 \leq i, j \leq d_\rho$$

A fundamental result in the theory of representations is *Schur's Lemma*. We will not state it here as it goes beyond the intended scope of this survey. We will however use one of its corollaries.

Theorem 4 Let ρ and τ be two unitary irreducible representations of the group G . Then, for all i, j, r, s ,

$$\langle \rho_{ij} | \tau_{rs} \rangle = 0$$

and

$$\langle \rho_{ij} | \rho_{rs} \rangle = \begin{cases} \frac{1}{d_\rho} & \text{if } i = r \text{ and } j = s \\ 0 & \text{otherwise} \end{cases}$$

Note that the inner product $\langle | \rangle$ is the standard one for complex functions we used previously for characters. Thus we see that the Fourier coefficients defined from a set inequivalent irreducible representations are pairwise orthogonal. It is also from this theorem that the orthonormality of irreducible characters is proved.

2.2 Representations of Abelian Groups

We saw that the number of inequivalent irreducibles of a group is equal to the number of conjugacy classes. In an Abelian group, $g^{-1}hg = hg^{-1}g = h$ and so, every element forms a conjugacy class onto itself. Therefore, for a complete set of irreducibles $\{\rho_i | 1 \leq i \leq k\}$ of an Abelian group we have, $k = |G|$

and, $\sum_{i=1}^{|G|} d_{\rho_i}^2 = |G|$. The last equation is only satisfied if $d_{\rho_i} = 1, \quad \forall i$. Hence we conclude that for any Abelian group, all irreducibles are one-dimensional. As every finite Abelian group is the direct product of cyclic groups we will first find the irreducibles of these to arrive at a general expression for the irreducibles of all Abelian groups.

Claim 1 For the cyclic group \mathbb{Z}_n of order n , the following functions $\rho_k : G \rightarrow GL_{\mathbb{C}}$ form a complete set of irreducible representations of the group,

$$\rho_k(a^r) = (\omega_n^{kr}) \text{ for } k = 0, \dots, n-1$$

Proof This proof and several others which will follow, use the following lemma,

Lemma 1 Let $d \in \mathbb{N}$, $m \in \mathbb{Z}$ and $\omega_d = e^{\frac{2\pi i}{d}}$ be the standard d^{th} root of unity. Then,

$$\sum_{r=0}^{d-1} \omega_d^{mr} = \begin{cases} d & \text{if } d|m \\ 0 & \text{otherwise} \end{cases}$$

(Proof omitted)

Let a^r and a^s be two elements of $\mathbb{Z}_n = \langle a \rangle$. Then, for all $k = 0, \dots, n-1$,

$$\rho_k(a^r)\rho_k(a^s) = \omega_n^{kr}\omega_n^{ks} = \omega_n^{k(r+s)} = \rho_k(a^{r+s}) = \rho_k(a^r a^s)$$

Furthermore as, $\rho_k(e = a^0) = \omega_n^{0k} = 1$, it is clear that ρ_k is a homomorphism. All that remains is to use the inner product for characters to show that all representations ρ_k are irreducible and inequivalent. Irreducibility is already given as each ρ_k is one-dimensional.

Now, without loss of generality we compute the inner product of the characters of the representations ρ_k and ρ_l where $l < k$.

$$\begin{aligned}\langle \chi_k | \chi_l \rangle &= \frac{1}{n} \sum_{r=0}^{n-1} \chi_k(a^r) \chi_l(a^{-r}) \\ &= \frac{1}{n} \sum_{r=0}^{n-1} \omega_n^{r(k-l)}\end{aligned}$$

By Lemma 1, as $0 \leq l, k \leq n - 1$, we have that n does not divide $k - l$ and so the summation must be zero. Therefore, as ρ_k and ρ_l are irreducible they must be inequivalent. \square

With the next result we will be able to extend what we know of the representations of cyclic groups to all Abelian groups.

Proposition 1 Let A be an Abelian group with identity e and let $\{\rho_i | 0 \leq i \leq k\}$ be a complete set of irreducible representations of A with respective characters χ_i . Now consider the group $G = A \times \mathbb{Z}_n$. Then, the functions, $\tau_{ij} : G \rightarrow GL_{\mathbb{C}}$ given by,

$$\tau_{ij}((a, r)) = (\chi_{\rho_i}(a) \omega_n^{rj}) \quad \forall a \in A, 0 \leq r \leq n - 1$$

for all $1 \leq i \leq k, 0 \leq j \leq n - 1$, form a complete set of irreducible representations of the group G .

Proof A is an Abelian group, hence every irreducible representation of A is one-dimensional. Consequently, we know the following are true:

- (i) Each χ_i is a homomorphism from A to $GL_{\mathbb{C}}$. Hence, for all elements $a, b \in A$,

$$\chi_i(a) \chi_i(b) = \chi_i(a + b) \quad \text{and} \quad \chi_i(e) = 1$$

(ii) By irreducibility and inequivalence,

$$\langle \chi_i | \chi_i \rangle = 1 \quad \text{and} \quad \langle \chi_i | \chi_j \rangle = 0$$

(iii) The number k of inequivalent irreducible representations of A is equal to the order of A .

In order to show that each τ_{ij} is a homomorphism, let (a, r) and (b, s) be two elements of the group G . Then,

$$\begin{aligned} \tau_{ij}((a, r)) \cdot \tau_{ij}((b, s)) &= (\chi_i(a)\omega_n^{rj}) (\chi_i(b)\omega_n^{sj}) \\ &= (\chi_i(a)\chi_i(b)\omega_n^{rj}\omega_n^{sj}) \\ &= (\chi_i(a+b)\omega_n^{j(r+s)}) \\ &= \tau_{ij}((a+b, r+s)) \\ &= \tau_{ij}((a, r) + (b, s)) \end{aligned}$$

Furthermore it is clear that each function τ_{ij} sends the identity element of G , $(e, 0)$ to the identity in $GL_{\mathbb{C}}$, (1) . Therefore, for every $1 \leq i \leq k$ and $0 \leq j \leq n-1$, the function τ_{ij} is a one-dimensional representation of the group G and hence irreducible as well.

Showing that they are inequivalent requires a bit more work for there are two cases to consider. Take the representations, τ_{is} and τ_{jt} where $1 \leq i, j \leq k$ and $0 \leq s, t \leq n-1$.

Case 1 ($i \neq j$) We reorganize the expression for the inner product of

$\chi_{\tau_{is}}$ and $\chi_{\tau_{jt}}$ as follows,

$$\begin{aligned} \langle \chi_{\tau_{is}} | \chi_{\tau_{jt}} \rangle &= \frac{1}{nk} \sum_{a \in A} \sum_{r=0}^{n-1} (\chi_i(a)\omega_n^{rs}) (\chi_j(a^{-1})\omega_n^{-rt}) \\ &= \frac{1}{n} \sum_{r=0}^{n-1} \langle \chi_i | \chi_j \rangle \omega_n^{r(s-t)} \end{aligned}$$

However, $i \neq j$ by assumption means χ_i and χ_j are inequivalent characters and so $\langle \chi_i | \chi_j \rangle = 0$. Therefore in this case, $\langle \chi_{\tau_{is}} | \chi_{\tau_{jt}} \rangle = 0$ and we conclude that the representations τ_{is} and τ_{jt} are inequivalent.

Case 2 ($s \neq t$) Without loss of generality suppose that $s > t$. This time we change the order of the sum so that,

$$\langle \chi_{\tau_{is}} | \chi_{\tau_{jt}} \rangle = \frac{1}{nk} \sum_{a \in A} \left(\sum_{r=0}^{n-1} \omega_n^{r(s-t)} \right) \chi_i(a) \chi_j(a^{-1})$$

Invoking Lemma 1 and remarking that $s > t \Rightarrow n \nmid (s-t)$ we see that the sum, $\sum_{r=0}^{n-1} \omega_n^{r(s-t)} = 0$. Thus we again conclude that τ_{is} and τ_{jt} are inequivalent irreducible representations of the group G .

That the set of all representations τ_{ij} is complete can be verified by simply stating that there are $nk = |G|$ of them. \square

As a corollary to this consider the general form of an Abelian group as a direct product of cyclic groups, $G = \mathbb{Z}_{n_1} \times \mathbb{Z}_{n_2} \times \dots \times \mathbb{Z}_{n_k}$. The elements of such a group can be treated as k -tuples $(a_1^{r_1}, \dots, a_k^{r_k})$ where a_i generates \mathbb{Z}_{n_i} . It follows that then, the function

$$\rho_{(i_1, \dots, i_k)}((a_1^{r_1}, \dots, a_k^{r_k})) = \left(\prod_{j=1}^k \omega_{n_j}^{i_j r_j} \right)$$

from G to $GL_{\mathbb{C}}$ is an irreducible representation of G for every

$$0 \leq i_1 \leq n_1 - 1, \dots, 0 \leq i_k \leq n_k - 1.$$

2.3 The Group Algebra

Given any finite group G , it is natural to consider the set of complex-valued functions over G . This set, denoted $\mathbb{C}[G]$ has the structure of an algebra over the complex numbers. A complex algebra is a ring $(A, +, \cdot, 0, 1)$ with the additional property that the set A and the addition operation “+” form a complex vector space. In the case of the complex algebra $\mathbb{C}[G]$ the addition and scalar multiplication are the obvious operations for functions. Multiplication between functions is the composition operation.

For every $g \in G$, we can define the *point-mass* or *delta* function δ_g ,

$$\delta_g(x) = \begin{cases} 1 & \text{if } x = g \\ 0 & \text{otherwise} \end{cases}$$

The set $\Gamma = \{\delta_g | g \in G\}$ of all such point-mass functions forms a basis for $\mathbb{C}[G]$. For if f is any complex-valued function over the group,

$$f = \sum_{g \in G} c_g \delta_g \quad \text{where } f(g) = c_g, \forall g \in G$$

The above tells us that in the basis Γ we can write any function f as a complex vector $(c_{g_1}, c_{g_2}, \dots, c_{g_n})$ and that the vector space $\mathbb{C}[G]$ has dimension $|G|$.

There is another important basis for the group algebra which we have already encountered. This is the set of all Fourier coefficient functions drawn from a complete set of inequivalent irreducibles of the group G . For each representation $\rho \in R$ we have d_ρ^2 such functions. As $\sum_{\rho \in R} d_\rho^2 = |G|$, then there are clearly $|G|$ such functions in total. That they form a linearly independent set follows from Theorem 4 which we saw when Fourier coefficients were first discussed. Thus, the set $\{\rho_{ij} | \rho \in R \text{ and } 1 \leq i, j \leq d_\rho\}$ forms a basis for the

group algebra. Although it has taken quite a bit of work we now have all the raw materials needed to define the Quantum Fourier Transform over any finite group.

Chapter 3

The Quantum Fourier Transform

In simplest terms a Fourier transform over a group G is a change of basis in the group algebra from the basis Γ of point-mass functions to the basis of Fourier coefficient functions. It may of course not be completely obvious at first what this implies in a quantum setting. If the Fourier transform is a change of basis, then it is natural to relate the basis vectors to basis states. Hence, we adopt a naming convention for the elements of both bases. First, we write the state $|g\rangle$ to represent the function δ_g . Second, we represent the function ρ_{ij} by a state in three registers $|\rho, i, j\rangle$. With this we will redefine the Quantum Fourier Transform in terms of the effect on an arbitrary vector in $\mathbb{C}[G]$.

Let G be a finite group and R a complete set of inequivalent irreducible representations of G .

Definition Let f be a complex-valued function on the group G and $\rho \in R$. The *Fourier transform of the function f at the representation ρ* is

denoted by $\hat{f}(\rho)$ and is given by

$$\hat{f}(\rho) = \sqrt{\frac{d_\rho}{|G|}} \sum_{g \in G} f(g) \rho(g)$$

$\hat{f}(\rho)$ is a $d_\rho \times d_\rho$ matrix with complex entries. The *Fourier transform of f at R* is composed of the $|G|$ entries of the matrices $\hat{f}(\rho)$ for each $\rho \in R$.

Consider the case when f is function of unit norm in the group algebra with respect to the basis Γ . (Of course, if the function does not have unit norm, we take the normalized form.) Then by extending the logic we used previously, we associate to the function a state $|f\rangle$ such that,

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

By the above definition, the Quantum Fourier Transform (QFT) sends the state $|f\rangle$ to,

$$\sum_{\rho \in R} \sum_{1 \leq i, j \leq d_\rho} \hat{f}(\rho)_{i,j} |\rho, i, j\rangle$$

This allows us to translate the definition of the QFT into more useful forms. In particular, let f be identical to the point-mass function δ_g for some element $g \in G$ and so $|f\rangle = |g\rangle$ according to our naming convention. It is clear that,

$$\hat{f}(\rho) = \sqrt{\frac{d_\rho}{|G|}} \rho(g) \quad \forall \rho \in R$$

Then, the QFT is a function F such that for every $g \in G$,

$$F(|g\rangle) = \sum_{\rho \in R} \sum_{1 \leq i, j \leq d_\rho} \sqrt{\frac{d_\rho}{|G|}} \rho_{i,j}(g) |\rho, i, j\rangle$$

Let us briefly consider the matrix form of the QFT. Let us choose an ordering of the group elements, (g_1, g_2, \dots, g_n) , with $n = |G|$ and have the state $|g_i\rangle$ correspond to the i^{th} n -dimensional elementary vector e_i . We also define some bijection μ which sends each integer t between 1 and n to a triple (ρ, i, j) . This allows us to write f_t to denote the function $\sqrt{\frac{d_\rho}{|G|}}\rho_{ij}$ where $\mu(t) = (\rho, i, j)$. With this definition, as the QFT was established to be a change of basis, we can describe it by a *matrix* F with,

$$F = \begin{bmatrix} f_1(g_1) & f_1(g_2) & \dots & f_1(g_n) \\ f_2(g_1) & f_2(g_2) & \dots & f_2(g_n) \\ \vdots & \vdots & \ddots & \vdots \\ f_n(g_1) & f_n(g_2) & \dots & f_n(g_n) \end{bmatrix}$$

Naturally, we may also consider the conjugate transpose of this matrix,

$$F^\dagger = \begin{bmatrix} \overline{f_1(g_1)} & \overline{f_2(g_1)} & \dots & \overline{f_n(g_1)} \\ \overline{f_1(g_2)} & \overline{f_2(g_2)} & \dots & \overline{f_n(g_2)} \\ \vdots & \vdots & \ddots & \vdots \\ \overline{f_1(g_n)} & \overline{f_2(g_n)} & \dots & \overline{f_n(g_3)} \end{bmatrix}$$

Looking at the product of the two matrices FF^\dagger , we notice that the $(i, j)^{\text{th}}$ entry of this product is of the form $\sum_{k=1}^n f_i(g_k)\overline{f_j(g_k)}$ which by Theorem 4 we know to be 0 for $i \neq j$ and 1 when $i = j$. Thus, the product FF^\dagger is the identity matrix and hence, F is unitary. Therefore, the Quantum Fourier Transform is a unitary transformation.

We may now also define the inverse of the Quantum Fourier Transform, QFT^{-1} . Computing $QFT^{-1}|\rho, i, j\rangle$ is equivalent to performing the multiplication $F^\dagger e_t$ where $\mu(t) = (\rho, i, j)$.

$$F^\dagger e_t = \overline{f_t(g_1)}e_1 + \overline{f_t(g_2)}e_2 + \dots + \overline{f_t(g_n)}e_n$$

By the definition $\overline{f_t(g_k)} = \sqrt{\frac{d_\rho}{|G|}} \overline{\rho_{ij}(g_k)}$ and e_k corresponds to the state $|g_k\rangle$.

In summary we have

$$QFT^{-1}(|\rho, i, j\rangle) = \sqrt{\frac{d_\rho}{|G|}} \sum_{g \in G} \overline{\rho_{ij}(g)} |g\rangle$$

In particular, notice that $QFT^{-1}(|1_G, 1, 1\rangle) = \frac{1}{|G|} \sum_{g \in G} |g\rangle$ providing a simple way of producing an equal superposition over all group elements.

We will attempt now to study the properties of the Quantum Fourier Transform, somewhat indirectly, in its application to a problem which seems to capture its essential nature.

Chapter 4

The Hidden Subgroup Problem

The field of quantum computing is still in its infancy. Yet a few spectacular results have been achieved. Deutsch's seminal algorithm which launched the entire field. Shor's algorithm ([17]), which factors integers and solves discrete logarithm problems in polynomial-time, drastically outperforms all known classical algorithms and presents serious problems for cryptography. Simon's algorithm ([18]), the first to solve a problem in polynomial time which provably requires exponential time to solve classically with any bounded-error probabilistic algorithm. The remarkable thing is that Deutsch's problem, the discrete logarithm problem, Simon's problem and a few others in this select group of successes are all special cases of the Hidden Subgroup Problem ([12]). This would suggest that other special cases of the HSP which are difficult problems classically might also have efficient quantum algorithms. Consider that even Graph Isomorphism is a special case of the Hidden Subgroup Problem.

Historically discoveries are rarely made so neatly as they appear in retrospect. Indeed, in this case, the Hidden Subgroup Problem arose as a gener-

alization, first formulated by Brassard and Høyer ([2]), of Simon’s Problem. The realization that it contained so many others came later. It is stated as follows:

Given a group G and a function $f : G \rightarrow R$ such that f is contained in a black box and R is some range. Also it is promised that the function f is constant and distinct on the cosets of some unknown subgroup $H \leq G$.

Find a generating set for the subgroup H .

The mathematical elegance of the above problem has generated a good deal of interest and a lot of work has been done in the attempt to find the general solution. At this point, efficient algorithms have only been found for restricted classes of groups. What is interesting to note is that in all cases of efficient algorithms, the main tool at work is the Quantum Fourier Transform. Not only that, but the quantum parts of these algorithms are in essence identical. This is not accidental. The Hidden Subgroup Problem is nothing more than a generalization to groups of the notion known as “periodicity of a function”, to which classical Fourier transforms have long been associated. A discrete function is *periodic* if there exists a an integer α such that

$$\forall x \in \mathbb{Z}, f(x + \alpha) = f(x)$$

In a group setting we are forced to replace integer addition in the above relation by the group operation. Thus, a function on a group G is periodic if for some fixed $\alpha \in G$, $f(x\alpha) = f(x)$ for all group elements x . Of course, this definition has the consequence that if the element α has order n ,

$$f(x) = f(x\alpha) = f(x\alpha^2) = \dots = f(x\alpha^{n-1})$$

Therefore, f is automatically constant on the cosets of the subgroup generated by α in G . By extending the concept to arbitrary subgroups and adding the restriction that f be distinct on different cosets, we recover the premise of the Hidden Subgroup Problem.

Suppose we are given a black box which contains a circuit, which we will call U_f such that, $U_f(|g\rangle|b\rangle) = |g\rangle|b \oplus f(g)\rangle$, for all $g \in G$. The first step is to create an equal superposition over a random coset cH of the hidden subgroup $H \leq G$ as follows:

1. Create an equal superposition over all group elements. That is,

$$\frac{1}{\sqrt{|G|}} \sum_{g \in G} |g\rangle$$

In general this is done by calculating $QFT^{-1}|1_G, 1, 1\rangle$. However, if the order of the group is 2^n , for some $n \in \mathbb{N}$ this is accomplished by applying a stack of Walsh-Hadamard transforms to a sufficient number of qubits all initialized to $|0\rangle$.

2. Apply the unitary transformation U_f to the previous state in the first registers and a second register containing $|0^m\rangle$. This calculates the value of the function f on every group element. Thus the resulting state is,

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

3. Measure the second register. The result of the measurement is a random value of the function f , $f(c)$ for some $c \in G$. As a consequence the state of the first register becomes,

$$|cH\rangle = \frac{1}{\sqrt{|H|}} \sum_{h \in G} |ch\rangle$$

All that remains is to compute the QFT on the state $|cH\rangle$ and measure the outcome. The first step presents a problem in that efficient implementations of QFT^{-1} are not known for all groups. However as an issue outside the scope of this thesis, we will assume that by some means the desired outcome can be achieved in deterministic polynomial time. Step 2 is straightforward within the quantum computing model. It is in the last step which interests us primarily and distinctions need to be made amongst groups starting from the application of the QFT. We will follow the historical progression and consider the case of Abelian groups first.

4.1 Abelian Groups

Let G be a finite Abelian group and $|cH\rangle = \frac{1}{\sqrt{|H|}} \sum_{h \in H} |ch\rangle$ be the equal superposition over a random coset of $H \leq G$. Recall that the representations of an Abelian group are all one-dimensional. Therefore, the general form of the QFT on an element g becomes,

$$\begin{aligned} F(|g\rangle) &= \sqrt{\frac{1}{|G|}} \sum_{\rho \in R} \rho_{1,1}(g) |\rho, 1, 1\rangle \\ &= \sqrt{\frac{1}{|G|}} \sum_{\rho \in R} \chi_{\rho}(g) |\rho\rangle \end{aligned}$$

It is now immediate that here, by linearity,

$$\begin{aligned} F(|cH\rangle) &= \frac{1}{\sqrt{|H|}} \sum_{h \in H} F(|ch\rangle) \\ &= \frac{1}{\sqrt{|H|}} \sum_{h \in H} \left(\frac{1}{\sqrt{|G|}} \sum_{\rho \in R} \chi_{\rho}(ch) |\rho\rangle \right) \\ &= \frac{1}{\sqrt{|H||G|}} \sum_{\rho \in R} \left(\sum_{h \in H} \chi_{\rho}(ch) \right) |\rho\rangle \end{aligned}$$

Now, as each $\rho \in R$ is one-dimensional, the character $\chi_{\rho} : G \rightarrow \mathbb{C}^{\times}$ is also a homomorphism. This means that $\chi_{\rho}(ch) = \chi_{\rho}(c)\chi_{\rho}(h)$. Therefore,

the amplitude of the state $|\rho\rangle$ is $\frac{\chi_\rho(c)}{\sqrt{|H||G|}} \sum_{h \in H} \chi_\rho(h)$. We saw that the sum in this expression had an interesting property. It is $|H|$ if ρ is the trivial representation over H and is zero for any other irreducible representation. Now, ρ is one-dimensional, therefore not only is it irreducible over G but also over H . The difference is that even if ρ is not the trivial representation over G , it may yet be over H .

We can rewrite $\sum_{h \in H}$ in a more compact form by letting 1_H be the trivial representation of H and noticing that,

$$\begin{aligned} \langle \chi_\rho | \chi_{1_H} \rangle &= \frac{1}{|H|} \sum_{h \in H} \chi_\rho(h) \chi_{\tau_1}(h^{-1}) \\ \Leftrightarrow \sum_{h \in H} \chi_\rho(h) &= \langle \chi_\rho | \chi_{1_H} \rangle |H| \end{aligned}$$

and remark that $\langle \chi_\rho | \chi_{1_H} \rangle$ is here only ever 0 or 1. The result is that the probability of observing ρ is,

$$\begin{aligned} \left| \frac{\chi_\rho(c)}{\sqrt{|G||H|}} \sum_{h \in H} \chi_\rho(h) \right|^2 &= \frac{|\chi_\rho(c)|^2}{|H||G|} |H|^2 \langle \chi_\rho | \chi_{1_H} \rangle \\ &= \sqrt{|H||G|} \langle \chi_\rho | \chi_{1_H} \rangle \end{aligned}$$

Let us focus on the case where $G = \langle a \rangle$, $G \cong \mathbb{Z}_n$. The hidden subgroup $H \leq G$ must also be cyclic generated by $\langle a^t \rangle$ for some $0 \leq t \leq n-1$. That is, $H = \{1, a^t, a^{2t}, \dots, a^{(d-1)t}\}$ where d is the order of H . Therefore, for each representation ρ_k of the cyclic group G ,

$$\sum_{h \in H} \chi_{\rho_k}(h) = \sum_{r=0}^{d-1} \chi_{\rho_k}(a^{tr}) = \sum_{r=0}^{d-1} \omega_n^{ktr}$$

This sum, by Lemma 1, is non-zero if and only if $n|kt$. If this is the case it is observed with probability $\frac{|H|}{|G|}$. In other words, if a representation is obtained as the result of a measurement on the state $F(|cH\rangle)$ then we know that $\frac{n}{\gcd(n,k)}|t$. Repeating the procedure m times, we obtain values k_1, k_2, \dots, k_m with this property. We would then conclude that

$t = \text{lcm}\left(\frac{n}{\gcd(n,k_1)}, \frac{n}{\gcd(n,k_2)}, \dots, \frac{n}{\gcd(n,k_m)}\right)$ with probability dependent on m . In fact for $m \in O(\log |G|)$ this probability can be made arbitrarily close to 1.

Returning to the general Abelian case, notice that the procedure would correspond to solving for the hidden subgroup with the function restricted to each of the cyclic groups in the direct product.

There is a significant problem in the implementation of this algorithm. Unfortunately, according to existing techniques, efficient implementations of the QFT only exist for Abelian groups which are of *smooth* order. The order n of a group is said to be smooth when all the prime factors of n are of size $O(\log n)$. It has been suggested by Jozsa ([10]) that this problem might be overcome by taking the smallest m such that 2^m is greater than n and using an implementation of the QFT for \mathbb{Z}_{2^m} as an approximation. Nothing to this effect exists yet.

4.2 Groups in General

We now move on to more general results. As was mentioned the algorithm is essentially the same. However, as in general a group will have irreducible representations which are not all one-dimensional, we have a choice in what we measure. Intuitively, the natural choices are to measure either simply the representation name or the entire Fourier coefficient names. The first case is drawn from the work of Hallgren, Russell and Ta-Shma ([7]) and the second from the work of Grigni, Schulman, Vazirani and Vazirani ([6]).

4.2.1 Measuring the Representation Name

If we choose to measure only the representation name, the state $F|g\rangle$ is best expressed in the form:

$$\sum_{\rho \in R} \sqrt{\frac{d_\rho}{|G|}} |\rho\rangle \left(\sum_{1 \leq i, j \leq d_\rho} \rho_{ij}(g) |i, j\rangle \right)$$

By linearity, for a random coset cH we have,

$$F(|cH\rangle) = \sum_{\rho \in R} \sqrt{\frac{d_\rho}{|G|}} |\rho\rangle \left(\sum_{1 \leq i, j \leq d_\rho} \left(\sum_{h \in H} \rho_{ij}(ch) \right) |i, j\rangle \right)$$

If A is a matrix with complex entries we may define the *euclidean* norm of A , $\|A\|$ such that,

$$\|A\|^2 \equiv \sum_{i, j} |a_{ij}|^2 = \text{tr}(A^\dagger A)$$

For a unitary matrix U we have the property that $\|UA\|^2 = \|A\|^2$. Note that this implies that if U is a unitary matrix, then $\|UAU^\dagger\|^2 = \|A\|^2$ and so the euclidean norm is invariant under any change of basis for unitary matrices.

Armed with this we see that the probability of observing a representation name when we measure the state $(F|cH\rangle)$ is exactly $\frac{d_\rho}{|H||G|} \|\sum_{h \in H} \rho(ch)\|^2$. As each representation ρ is unitary and a homomorphism we have $\|\sum_{h \in H} \rho(ch)\|^2 = \|\rho(c) \sum_{h \in H} \rho(h)\|^2 = \|\sum_{h \in H} \rho(h)\|^2$. Therefore, the distribution on the representation names is independent of the coset.

Now we seek to simplify $\|\sum_{h \in H} \rho(h)\|^2$. We know that for any irreducible representation which is not the trivial one, $\sum_{g \in G} \chi_\rho(g) = 0$. This property is again useful when combined with the following lemma,

Lemma 2 Let K be a finite subgroup of $GL_{\mathbb{C}}(V)$ for some complex vector space V of dimension n . Then,

$$\sum_{k \in K} \text{tr}(k) = 0 \implies \sum_{k \in K} k = 0_{n \times n}$$

That is, if the sum of the traces of all the matrices in a subgroup is zero, then the sum of the matrices must be the zero matrix. The image of a representation is a subgroup of a general linear group and so, by Lemma 2, whenever ρ is an irreducible representation which is not the trivial one,

$$\sum_{g \in G} \chi_\rho(g) = 0 \implies \sum_{g \in G} \rho(g) = 0_{d_\rho \times d_\rho}$$

On the other hand, for the trivial representation 1_G , $\sum_{g \in G} 1_G(g) = |G|$.

We now apply a reasoning similar to that which we used for Abelian groups. The problem here is that an irreducible representation ρ of G is not in general irreducible for H . Hence first we decompose each irreducible representation ρ of G into a direct sum of irreducibles of H .

$$\rho = n_1 \tau_1 \oplus n_2 \tau_2 \oplus \dots \oplus n_l \tau_l$$

Therefore, there exists a certain basis for the representation space so that each matrix $\rho(g)$ will have block diagonal form,

$$\rho(g) = \begin{bmatrix} \tau_1(g) & & \\ & \ddots & \\ & & \tau_l(g) \end{bmatrix}$$

with each representation τ_i appearing $n_i = \langle \chi_\rho | \chi_{\tau_i} \rangle$ times. Thus by Lemma 2, when we take the sum of the $\rho(h)$ over all $h \in H$, all blocks which do not correspond to the trivial representation will disappear. Let us assume without loss of generality that τ_1 is the trivial representation of H . Therefore, the only entries which will not be zero are the n_1 diagonal entries which correspond to each of the occurrences of τ_1 . Necessarily, these will each have the value $|H|$. The conclusion is that, as the euclidean norm is independent

of the basis, the probability of observing ρ is,

$$\begin{aligned} \frac{d_\rho}{|H||G|} \left\| \sum_{h \in H} \rho(h) \right\|^2 &= \frac{d_\rho}{|H||G|} \sum_{1 \leq i, j \leq d_\rho} \left| \left(\sum_{h \in H} \rho(h) \right)_{ij} \right|^2 \\ &= \frac{d_\rho}{|H||G|} \langle \chi_\rho | \chi_{\tau_1} \rangle |H|^2 \\ &= \frac{d_\rho}{|H||G|} |H| \sum_{h \in H} \chi_\rho(h) \end{aligned}$$

which is exactly the same form we obtained when we restricted ourselves to Abelian groups. Despite the appeal of this symmetry, it now implies a serious drawback. When characters were introduced we saw that they were class functions. As a consequence, we see that if we only measure the representation name we will not be able to distinguish between conjugate subgroups. That is, the distribution on representation names given by $F(|H\rangle)$ is identical to that of $F(|g^{-1}|Hg\rangle)$. This is very disturbing, for in noncommutative groups there may be many subgroups which are conjugate.

Happily though, it will now be shown that if the hidden subgroup is promised to be self-conjugate (i.e. normal), sampling the QFT on random cosets will be sufficient to determine subgroup uniquely. That is consistent with what is known of the Abelian case where every subgroup is normal.

The essential observation is contained in the following lemma:

Lemma 3 If H is a normal subgroup, the probability of observing a representation ρ as a result of measuring $F(|cH\rangle)$ is non-zero if and only if $H \subseteq \ker \rho$.

Proof We know that under the given circumstances, the probability of observing a representation name ρ is,

$$\frac{d_\rho}{|G||H|} \left\| \sum_{h \in H} \rho(h) \right\|^2 = \frac{d_\rho}{|G||H|} |H|^2 \langle \chi_\rho | \chi_{1_H} \rangle = d_\rho \frac{|H|}{|G|} \langle \chi_\rho | \chi_{1_H} \rangle$$

Let us introduce a new representation of the group, $\tau : G \rightarrow GL_{\mathbb{C}}(V)$. We will define τ as the permutation representation on the set G/H on which the group acts by left multiplication. Thus, $\tau(g')$ is a $\frac{|G|}{|H|} \times \frac{|G|}{|H|}$ matrix which sends basis vector corresponding to gH to the one for $g'gH$. As H is normal in G , $hgH = gH$ for all $h \in H$. This means that the action of any $h \in H$ leaves every coset fixed. Since each $\tau(g)$ is a permutation, it must then be that $\tau(h)$ is the identity matrix for all $h \in H$.

On the other hand, every element of the group not in the subgroup H will fix no elements by this action. This allows us to describe the character of the representation completely:

$$\chi_{\tau}(g) = \begin{cases} \frac{|G|}{|H|} & \text{if } g \in H \\ 0 & \text{otherwise} \end{cases}$$

Now let ρ be any irreducible representation of the group G .

$$\begin{aligned} \langle \chi_{\rho} | \chi_{\tau} \rangle &= \frac{1}{|G|} \sum_{g \in G} \chi_{\rho}(g) \chi_{\tau}(g^{-1}) \\ &= \frac{1}{|G|} \frac{|G|}{|H|} \sum_{h \in H} \chi_{\rho}(h) \\ &= \langle \chi_{\rho} | \chi_{1_H} \rangle \end{aligned}$$

Recall that $\langle \chi_{\rho} | \chi_{\tau} \rangle$ corresponds to the number of times the irreducible representation ρ appears in the decomposition of τ . That is, if $R = \{\rho_1, \dots, \rho_k\}$ is a complete set of inequivalent irreducibles of G ,

$$\tau = \langle \chi_1 | \chi_{\tau} \rangle \rho_1 \oplus \dots \oplus \langle \chi_k | \chi_{\tau} \rangle \rho_k$$

However we also know that in the appropriate basis τ will have block

diagonal form,

$$\tau(g) = \begin{bmatrix} \rho_1(g) & & \\ & \ddots & \\ & & \rho_k(g) \end{bmatrix}$$

It is clear from this form, that for $\tau(h)$ to be the identity matrix for every $h \in H$, we require the same of every representation ρ which appears in the decomposition of τ over G , for which the condition is that $\langle \chi_\rho | \chi_\tau \rangle \neq 0$. In summary then,

$$\langle \chi_\rho | \chi_\tau \rangle \neq 0 \Leftrightarrow \rho(h) = 1_{d_\rho \times d_\rho} \forall h \in H$$

which we can restate as,

$$\langle \chi_\rho | \chi_{1_H} \rangle \neq 0 \Leftrightarrow H \subseteq \ker \rho$$

This completes the proof. \square

Another way of seeing this lemma is that if a representation ρ of G , contains the trivial representation of a normal subgroup $H \leq G$ when decomposed over H , it can contain no other irreducible representation of H in this decomposition. As a last remark, notice that when $H \subseteq \ker \rho$, then the probability of observing the representation is $d_\rho^2 \frac{|H|}{|G|}$.

Given Lemma 3 we see that the procedure to determine a hidden subgroup is to sample a series of representations and take the intersection of the kernels of these. After a sufficient number of samples, with high probability the intersection of the kernels of the sampled representations will be the hidden subgroup H . The number of such samples required will be linear in the size of the input. However, intersecting the kernels of the representations is a problem with complexity related to the underlying group and thus little can

be said of this in general. This last operation will typically be performed classically.

4.2.2 Measuring the Fourier Coefficients

One might hope that by not restricting ourselves only to the representation names, the obstacle on non-normal subgroups might be overcome. There is so far no indication that this might be the case. One of the main problems is that whereas previously the distribution of $F|cH\rangle$ was independent of the coset, when we measure the Fourier coefficients individually this is no longer true. This follows from the fact that the Fourier coefficients ρ_{ij} are not homomorphisms in general. Surprisingly, the consequence of this will be that measuring the rows provides no additional information whatever. But there are further problems. The probability of observing (ρ, i, j) is a function of

$$\sum_{h \in H} \rho(gh)_{ij} = \sum_{k=1}^{d_\rho} \rho(g)_{ik} \left(\sum_{h \in H} \rho(h)_{kj} \right)$$

We may now reason as for representation names. The representation ρ can be decomposed into irreducibles over H , and a basis can be chosen so that $\rho(h)$ will have block diagonal form for all $h \in H$. Once in this form, we know that $\sum_{h \in H} \rho(h)_{kj} = 0$ unless (k, j) is the coordinate of a diagonal entry which corresponds to a manifestation of the trivial representation of H . In the latter case $\sum_{h \in H} \rho(h)_{kj} = 1$. It would seem then that despite the difficulty of dependence on the coset there are nonetheless interesting properties to be exploited similar to those which served us so well in the last section. However, here again we are foiled for the actual expression for the

probability of observing (ρ, i, j) for a fixed coset gH is,

$$\frac{d_\rho}{|G||H|} \left| \sum_{h \in H} \rho_{ij}(gh) \right|^2$$

which very sadly depends on the basis of the representation space. The loss of freedom in choosing bases is a severe obstacle and in particular means that the tidy conclusions made about the sum $\sum_{h \in H} \rho_{ij}(gh)$ cannot be used.

The dependence of the above probability on the coset implies that we are actually sampling from the statistical mixture,

$$\left\{ \left(F(|gH\rangle), \frac{1}{|G|} \right) \right\}_{g \in G}$$

Therefore, the resulting probability of obtaining (ρ, i, j) is

$$\frac{d_\rho}{|H||G|^2} \sum_{g \in G} \left| \sum_{h \in H} \rho_{ij}(gh) \right|^2$$

For the sake of brevity let us write $\rho(H)$ to signify $\sum_{h \in H} \rho(h)$. Let us further define an ordering of the group elements, (g_1, g_2, \dots, g_n) , where $n = |G|$. This allows us to restate the sum $\sum_{g \in G} \left| \sum_{h \in H} \rho_{ij}(gh) \right|^2$ as the norm squared of the complex vector,

$$v = (\rho(g_1 H)_{ij}, \rho(g_2 H)_{ij}, \dots, \rho(g_n H)_{ij})$$

As ρ is a homomorphism, we know that $\rho(gH)_{ij} = \sum_{k=1}^{d_\rho} \rho(g)_{ik} \rho(H)_{kj}$ and so,

$$v = \sum_{k=1}^{d_\rho} (\rho(g_1)_{ik}, \rho(g_2)_{ik}, \dots, \rho(g_n)_{ik}) \rho(H)_{kj} \equiv \sum_{k=1}^{d_\rho} v_{ik} (\rho(H)_{kj})$$

It turns out that the vectors v_{ik} are orthogonal with respect to the standard inner product $(X, Y) = XY^\dagger$. This follows from the orthogonality relations of the Fourier coefficients. For notice that,

$$(v_{ik}, v_{il}) \equiv v_{ik} v_{il}^\dagger = \sum_{j=1}^n \rho_{ik}(g_j) \overline{\rho_{il}(g_j)} = n \langle \rho_{ik} | \rho_{il} \rangle$$

Therefore,

$$(v_{ik}, v_{il}) = \begin{cases} 0 & \text{if } k \neq l \\ \frac{n}{d_\rho} & \text{if } k = l \end{cases}$$

The conclusion is that the probability of observing the triple (p, i, j) by measuring the given statistical mixture is,

$$\frac{d_\rho}{|H||G|^2} \sum_{g \in G} \left| \sum_{h \in H} \rho_{ij}(gh) \right|^2 = \frac{1}{|G||H|} \sum_{k=1}^{d_\rho} |\rho(H)_{kj}|^2$$

This shows that the probability of observing (ρ, i, j) is a linear function of $\rho(H)$ and completely independent of the row. Thus if we are to extend the type of measurement from the representation name, we need at most to measure the name and columns. Please note that $\|\rho(H)_j\|^2$ is also a quantity which is dependent on the basis. The second general result is that if the basis is chosen randomly, only a negligible amount of additional information can be obtained. This does not exclude the possibility that a more deliberate choice of bases might prove more fruitful.

Chapter 5

Dihedral Groups and Their Hidden Subgroups

As a summary of the results we have just seen, we will look at a specific family of groups, the dihedral group. This provides an interesting context because it is a group which is noncommutative yet with a structure otherwise near to that of an Abelian group. On the other hand it is complex enough that the obstacles anticipated in the previous section cannot be avoided.

Definition The *dihedral group* of order $2N$, denoted by D_N is the group of symmetries of the regular N -sided polygon in the plane. It is generated by two elements a and b subject to the following relations,

$$(i) \quad a^N = e$$

$$(ii) \quad b^2 = e$$

$$(iii) \quad ab = ba^{-1}$$

Thus, the generator a corresponds to a rotation about the center of the polygon of $\frac{2\pi}{N}$ and b to any one of its reflection symmetries. Note that

every element of the group D_N can be written as $b^s a^r$ for some choice of s and r .

We will limit our attention only to the case where $N = 2^n$ or in other words, dihedral groups of order 2^{n+1} . Not only is this a natural first step in quantum computing, but again it is only for groups of such orders that efficient implementations of the QFT currently exist. The reason for this might be found in the fact that the dihedral groups of order 2^n are exactly those which are nilpotent, although this requires more investigation.

We may also view the group D_N as the *semi-direct product* between the two cyclic groups \mathbb{Z}_N and \mathbb{Z}_2 . This means that every element of D_N may be viewed as a pair (x, y) such $x \in \mathbb{Z}_N$ and $y \in \mathbb{Z}_2$ and with multiplication defined by,

$$(x_1, y_1) \cdot (x_2, y_2) = (x_1 + (-1)^{y_1} x_2, y_1 \oplus y_2)$$

We can relate the two definitions by setting $a = (1, 0)$ and $b = (0, 1)$.

Let us now look at the subgroups of D_N . It is clear that $K = \langle a \rangle = \{(x, 0) | x \in \mathbb{Z}_N\}$ is a subgroup of D_N of order N , isomorphic to \mathbb{Z}_N . It is also normal in D_N . All other normal subgroups of the dihedral group are also subgroups of K . In fact they can all be assembled in a chain of inclusion,

$$K_1 \triangleleft K_2 \triangleleft \dots \triangleleft K_n = K \triangleleft D_N$$

where for each $0 \leq i \leq n$ is the cyclic subgroup of D_N of order 2^i generated by $a^{2^{n-i}}$. There are no other normal subgroups of D_N apart from those listed.

Besides these normal subgroups there is one more family which will complete the list of all cyclic subgroups of D_N . Each has order two and is of the form $H_t = \{e, ba^t\}$ with $0 \leq t \leq N - 1$. Not only are these not normal but the conjugacy class of H_t contains $\frac{N}{2}$ subgroups. Every other subgroup of D_N

can be described as the product of at most two of those we have just seen. In other words they are of the form $K_i H_t$ for some $0 \leq i \leq n$, $0 \leq j \leq N-1$. In particular notice that $K_i H_0$ is a subgroup of D_N isomorphic to $D_{2^{i-1}}$ for every admissible i . We conclude that every subgroup of the dihedral group can be written as $K_i L$ with $L = e$ or $L = H_t$ for some t .

5.1 The Representations of the Dihedral Group

Again Lemma 1 will be extremely useful in proving that the following representations are irreducible.

Claim 2 Let $\rho_k : D_N \rightarrow GL_{\mathbb{C}}(\mathbb{C}^2)$ be the map given by

$$\rho_k(a^r) = \begin{pmatrix} \omega_N^{kr} & 0 \\ 0 & \omega_N^{-kr} \end{pmatrix} \quad \rho_k(ba^r) = \begin{pmatrix} 0 & \omega_N^{-kr} \\ \omega_N^{kr} & 0 \end{pmatrix}$$

Then, ρ_k is an irreducible representation for each $k = 1, \dots, N-1$ and $k \neq \frac{N}{2}$. Moreover, each ρ_k is equivalent to only one other representation in this list, $\rho_{k^{-1}}$.

Proof Let ρ_k be defined as above and let χ_k be the associated character.

Clearly, $\chi_k(a^r) = \omega_N^{kr} + \omega_N^{-kr}$ and $\chi_k(ba^r) = 0$. It is also easily verified that ρ_k is a homomorphism for each $k = 1, \dots, N-1$.

(i) ρ_k is irreducible if and only if $\langle \chi_k | \chi_k \rangle = 1$. By definition,

$$\begin{aligned}
\langle \chi_k | \chi_k \rangle &= \frac{1}{2N} \sum_{r=0}^{N-1} \sum_{s=0}^1 \chi_k(b^s a^r) \overline{\chi_k(b^s a^r)} \\
&= \frac{1}{2N} \sum_{r=0}^{N-1} \chi_k(a^r) \chi_k(a^{-r}) \\
&= \frac{1}{2N} \sum_{r=0}^{N-1} \left(\omega_N^{2kr} + 2\omega_N^{k(r-r)} + \omega_N^{-2kr} \right) \\
&= 1 + \frac{1}{2N} \sum_{r=0}^{N-1} \omega_N^{2kr} + \frac{1}{2N} \sum_{r=0}^{N-1} \omega_N^{-2kr}
\end{aligned}$$

For each $k = 1, \dots, N-1$, $k \neq \frac{N}{2}$, we have $N \nmid 2k$ and $N \nmid (-2k)$.

Then by Lemma 1,

$$\sum_{r=0}^{N-1} \omega_N^{2kr} = 0 = \sum_{r=0}^{N-1} \omega_N^{-2kr}$$

Therefore, $\langle \chi_k | \chi_k \rangle = 1$ and the representation ρ_k is irreducible for all such k .

- (ii) ρ_{k_1} and ρ_{k_2} are inequivalent irreducible representations iff $\langle \chi_{k_1} | \chi_{k_2} \rangle = 0$. Now by definition, for $1 \leq k_1 < k_2 \leq N-1$ and $k_1, k_2 \neq \frac{N}{2}$, we have

$$\begin{aligned}
\langle \chi_{k_1} | \chi_{k_2} \rangle &= \frac{1}{2N} \sum_{r=0}^{N-1} \sum_{s=0}^1 \chi_{k_1}(b^s a^r) \overline{\chi_{k_2}(b^s a^r)} \\
&= \frac{1}{2N} \sum_{r=0}^{N-1} \chi_{k_1}(a^r) \chi_{k_2}(a^{-r}) \\
&= \frac{1}{2N} \sum_{r=0}^{N-1} \left(\omega_N^{r(k_1+k_2)} + \omega_N^{-r(k_1+k_2)} + \omega_N^{r(k_1-k_2)} + \omega_N^{-r(k_1-k_2)} \right)
\end{aligned}$$

As $k_1 < k_2$ it follows immediately that $N \nmid (k_1 - k_2)$ and that $N \nmid -(k_1 - k_2)$. Now, for $1 \leq k_1, k_2 \leq N-1$, if $N \mid (k_1 + k_2)$ then $k_2 = k_1^{-1}$. Therefore, if $k_2 \neq k_1^{-1}$ then the sum of each of the terms is zero. In this case ρ_{k_1} and ρ_{k_2} are inequivalent. \square

Therefore, by taking $k = 1, \dots, \frac{N}{2} - 1$ we have found a set of $\frac{N}{2} - 1$ inequivalent irreducible representations of the group D_N , each of dimension two. The set can be made complete by adding the following one-dimensional representations:

- (i) $\tau_0 : G \rightarrow GL_{\mathbb{C}}$ with $\tau_0(a^r) = (1)$ and $\tau_0(ba^r) = (1)$
- (ii) $\tau_1 : G \rightarrow GL_{\mathbb{C}}$ with $\tau_1(a^r) = (1)$ and $\tau_1(ba^r) = (-1)$
- (iii) $\tau_2 : G \rightarrow GL_{\mathbb{C}}$ with $\tau_2(a^r) = (-1)^r$ and $\tau_2(ba^r) = (-1)^r$
- (iv) $\tau_3 : G \rightarrow GL_{\mathbb{C}}$ with $\tau_3(a^r) = (-1)^r$ and $\tau_3(ba^r) = (-1)^{r+1}$

5.2 The QFT for the Dihedral Group

Now that we have pinpointed a precise complete set of irreducible representations of the dihedral group D_N , we may write down explicitly the form of the QFT. Recall that we had the following expression in general,

$$F(|g\rangle) = \sum_{\rho \in R} \sum_{1 \leq i, j \leq d_{\rho}} \sqrt{\frac{d_{\rho}}{|G|}} \rho_{i,j}(g) |\rho, i, j\rangle$$

Consider the expression now with the representations we know,

$$F(|g\rangle) = \frac{1}{\sqrt{2N}} \left(\sum_{q=0}^3 \chi_{\tau_q}(g) |\tau_q\rangle + \sum_{k=1}^{\frac{N}{2}-1} \sum_{1 \leq i, j \leq 2} \sqrt{2} \rho_k(g)_{i,j} |\rho_k, i, j\rangle \right)$$

By setting $|g\rangle = |a^r\rangle$ we have,

$$F(|a^r\rangle) = \frac{1}{\sqrt{2N}} \left(\sum_{q=0}^3 \chi_{\tau_q}(a^r) |\tau_q\rangle + \sum_{k=1}^{\frac{N}{2}-1} \sum_{1 \leq i, j \leq 2} \sqrt{2} \rho_k(a^r)_{i,j} |\rho_k, i, j\rangle \right)$$

Which simplifies to,

$$F(|a^r\rangle) = \frac{1}{\sqrt{2N}} \left(\sum_{q=0}^3 \chi_{\tau_q}(a^r) |\tau_q\rangle + \sqrt{2} \sum_{k=1}^{\frac{N}{2}-1} (\omega_N^{kr} |\rho_k, 1, 1\rangle + \omega_N^{-kr} |\rho_k, 2, 2\rangle) \right)$$

On the other hand suppose we set $|g\rangle = |ba^r\rangle$, then we obtain,

$$F(|ba^r\rangle) = \frac{1}{\sqrt{2N}} \left(\sum_{q=0}^3 \chi_{\tau_q}(g) |\tau_q\rangle + \sqrt{2} \sum_{k=1}^{\frac{N}{2}-1} (\omega_N^{-kr} |\rho_k, 2, 1\rangle + \omega_N^{kr} |\rho_k, 1, 2\rangle) \right)$$

Notice that in the expression for the state $F|a^r\rangle$ all the basis states $|\rho_k, i, j\rangle$ have amplitudes of norm 0 when $i \neq j$ and amplitudes of norm 1 when $i = j$.

In contrast in the state $F|ba^t\rangle$ the situation is the reverse.

5.3 Examples for Dihedral Hidden Subgroup Problem

Let D_N be the dihedral group of order $2N$ and let f be a function which fulfills the hidden subgroup promise on D_N with respect to some subgroup $H \leq D_N$.

We saw that any subgroup of D_N can be written as the product $K_i L$ where K_i is the cyclic subgroup $\langle a^{2^{n-i}} \rangle$ and $L = \{e\}$ or $L = H_t = \{e, ba^t\}$. So in particular, without loss of generality we say that the hidden subgroup $H = K_j L$. By restricting f to K_n we create an instance of the *Abelian* Hidden Subgroup Problem for this restriction of f fulfills the hidden subgroup promise for $K_j \leq K_n$. As K_n is both Abelian and smooth, we have a quantum algorithm for finding K_j efficiently.

Now, once the generator of K_j has been found, and as K_j is a normal subgroup, we will be able to construct a homomorphism $\varphi : D_N \rightarrow D_{2^{n-i}}$ with kernel K_i . Then, the composition $f \circ \varphi$ will hide L in $D_{2^{n-i}}$ just as before f hid $K_j L$. In this way the problem of finding an arbitrary hidden subgroup in D_N is reduced to finding one which is either trivial or H_t for

some $0 \leq t \leq N - 1$. Of course, while the situation has become much simpler to describe, it remains nontrivial. For not only are there an exponential number in the size of the input n of such subgroups in D_N , but also none of the H_t are normal. Hence we immediately run into a serious obstacle, the method of measuring only the representation name cannot distinguish between conjugate subgroups. Even if this method could be used to determine a conjugacy class, the size of the conjugacy class of each H_t is still exponential in n . It would be impossible then by inspection alone to determine the hidden subgroup within its conjugacy class in polynomial time.

5.3.1 Finding Normal Subgroups

We have seen how all normal subgroups of D_N can be found by restricting the function f to the subgroup of D_N which is isomorphic to the cyclic group \mathbb{Z}_N . However, since our interest lies more in examining the properties of the QFT, than in solving the Hidden Subgroup Problem, let us briefly see how we may find hidden normal subgroups using the QFT for D_N , rather than with the previous shortcut.

Let $|cH\rangle$ be a random coset state of a subgroup $H \triangleleft D_N$ where H is generated by an element a^t for some $0 \leq t \leq N - 1$. Little can be learned in this case if the result of the measurement is one of the τ_q . This would occur with probability $\frac{|H|}{2^{N-2}}$. When the result of a measurement is one of the ρ_k we learn a great deal more. The probability of observing ρ_k from a measurement of $F(|cH\rangle)$ is,

$$\frac{1}{|H|2^{N-1}} \sum_{i,j} \left| \sum_{h \in H} \rho_{i,j}(ch) \right|^2$$

No matter the coset we are considering, for some $0 \leq r \leq N - 1$ the above

can be written as,

$$\begin{aligned}
&= \frac{2}{|H|2^{N-1}} \left| \sum_{s=0}^{|H|-1} \omega_N^{(r+st)k} \right|^2 \\
&= \frac{2}{|H|2^{N-1}} |\omega_N^{rk}|^2 \left| \sum_{s=0}^{|H|-1} \omega_N^{(tk)s} \right|^2 \\
&= \frac{2}{|H|2^{N-1}} \left| \sum_{s=0}^{|H|-1} \omega_N^{(tk)s} \right|^2
\end{aligned}$$

This probability is non-zero if and only if $n|(kt)$ in which case it is $\frac{|H|}{2^{N-2}}$. The procedure for determining t so that $H = \langle a^t \rangle$ is the same as was used for Abelian groups. Note that if we wish to accomplish the reduction from a hidden subgroup K_jL to L only, we cannot use this last method and must first restrict the function f to K_n .

5.3.2 Distinguishing Between $|H| = 1$ and $|H| = 2$

Let us return now to the reduced problem of determining a hidden subgroup $H \leq D_N$ when H is promised to be either trivial or one of the $H_t = \{e, ba^t\}$. We know that there is no hope of finding the hidden subgroup by measuring only the representation name. However, perhaps we would be more lucky if the entire Fourier coefficient name were measured instead. Unfortunately this is not the case and we show this by considering an equivalent problem.

One of the results of Hallgren, Russell and Ta-Shma ([7]), which we did not cover, is that one cannot use the Quantum Fourier Transform to distinguish between a subgroup of order 1 and one of order 2 in the Symmetric group S_n . Specifically, their result applies to the case where only the representation name is measured. Here we will see a similar result for Dihedral groups with the difference that we will be measuring the indices as well. Note that we are not including the normal subgroup of order 2 in D_N , $\{e, a^{\frac{N}{2}}\}$ as

one of the instances which the algorithm need distinguish. This is important to determine for the two problems are equivalent. If we had an algorithm which finds a hidden subgroup $H \leq D_N$ of the required form, then we necessarily learn its order. On the other hand, an efficient means of determining the order of a hidden subgroup H in a dihedral would enable us to determine H explicitly using the approach we will now describe.

Let $\phi : D_N \rightarrow D_{\frac{N}{2}}$ be the group homomorphism given by,

$$\phi(b^s a^r) = \begin{cases} b^s a^r & \text{if } r < \frac{N}{2} \\ b^s & \text{if } r = \frac{N}{2} \\ b^s a^{r-\frac{N}{2}} & \text{if } r > \frac{N}{2} \end{cases}$$

Suppose that $H = \{e, ba^t\}$. It is easily verified that if $t > \frac{N}{2}$ then, $f \circ \phi$ hides the trivial subgroup in $D_{\frac{N}{2}}$. Whereas if $t < \frac{N}{2}$, the function $f \circ \phi$ would still hide $H \leq D_{\frac{N}{2}}$.

Therefore, if one had an oracle O for deciding whether a hidden subgroup K was of order 1 or of order 2 in a dihedral group, one could construct a recursive algorithm A as follows:

$A(D_K, g)$

1. If $g(ba^{\frac{K}{2}}) = g(0)$ then return $(ba^{\frac{K}{2}})$
2. Otherwise, apply the oracle O to D_K and $g \circ \phi$.
3. If the oracle answers that the hidden subgroup has order two then return $(A(D_{\frac{K}{2}}, g))$. Otherwise the subgroup was trivial. So, return $(A(D_{\frac{K}{2}}, g \circ \xi)a^{\frac{K}{2}})$. Where $\xi : D_{\frac{K}{2}} \rightarrow D_K$ is given by,

$$\xi(b^s a^r) = \begin{cases} a^r & \text{if } s = 0 \\ ba^{r+\frac{K}{2}} & \text{if } s = 1 \end{cases} \quad \text{for all } 0 \leq r \leq \frac{K}{2} - 1$$

We precede algorithm A by a call to the oracle on D_N and f to eliminate the possibility that the hidden subgroup H is initially trivial. Then, by running $A(D_N, f)$, at each step of the recursion we learn the current most significant bit of t such that $H = \{e, ba^t\}$. Thus finding the value of t in a linear number of recursive steps.

It is clear that when H is the trivial subgroup, a random coset of H is simply a random element of the group. Therefore, when we are sampling from the Fourier transform of these, we are in fact sampling from the statistical mixture,

$$\left\{ \left(F(|g\rangle), \frac{1}{2N} \right) \right\}_{g \in G}$$

From earlier findings we know that the probability of observing the triple (ρ_k, i, j) from this statistical mixture is,

$$\frac{1}{|2N|} \|\rho(e)_j\|^2 = \frac{1}{|2N|} (|1|^2 + |0|^2) = \frac{1}{|2N|}$$

where $\rho_k(e)_j$ is the j^{th} column of $\rho_k(e)$. The remaining representations τ_0, τ_1, τ_2 , and τ_3 each being one-dimensional are also each observed with this same probability. The result is that all of the Fourier coefficients are observed with equal probability $\frac{1}{2N}$.

Suppose now that $H = \{e, ba^t\}$. Then the form of a random coset of H is $\{a^r, ba^{t-r}\}$. As each coset has two possible representatives, we are sampling now from the statistical mixture,

$$\left\{ \left(\frac{F|a^r\rangle + F|ba^{t-r}\rangle}{\sqrt{2}}, \frac{1}{N} \right) \right\}_{r=0}^{N-1}$$

In this case the probabilities of observing a triple (ρ, i, j) is still $\frac{1}{2N}$ exactly as when H was trivial. For, if we take $j = 1$ (and the other case $j = 2$ gives the same value), the probability of observing $(\rho, i, 1)$ is,

$$\frac{1}{2N} \|\rho_k(e)_1 + \rho_k(ba^t)\|^2 = \frac{1}{2N} (|1|^2 + |\omega_n^{kt}|^2) = \frac{1}{2N}$$

The only differences appear in the distribution of the τ_q 's. Specifically,

- τ_0 is observed with probability $\frac{1}{N}$, hence twice as often as when H is trivial.
- τ_1 is observed with probability 0.
- Depending on the value of t , one of τ_2 and τ_3 is observed with probability $\frac{1}{N}$ and the other not at all.

The differences are so small that we can see how a polynomial number of samples would be insufficient to distinguish between the two cases with any reasonable degree of confidence. Of course, as this was based on a specific choice of basis for each representation, the evidence is not conclusive. In contrast if we only perform a partial measurement of the representation name, then there is absolutely no way of distinguishing directly between some H_t and the trivial subgroup efficiently. The underlying problem is that for any hidden subgroup of the form H_t , every representation ρ_k decomposes identically over H_t as the direct sum of both irreducibles of the subgroup. Explicitly recall that when we perform a partial measurement, the probability of observing ρ is,

$$\frac{d_\rho}{|G|} \sum_{h \in H} \chi_\rho(h)$$

independent of the basis on which ρ is expressed. Therefore, when $H = H_t$, the probability of observing the name ρ_k is,

$$\frac{2}{2N} (\chi_k(e) + \chi_k(ba^t)) = \frac{1}{N}(2 + 0) = \frac{2}{N}$$

If, however $H = \{e\}$, the probability of observing ρ_k is still $\frac{1}{N}(\chi_k(e)) = \frac{2}{N}$. The relative probabilities of observing the τ_i 's are exactly the same as for

complete measurements, for they are all one-dimensional. In summary, we conclude that the Fourier Transform does not allow us to distinguish between subgroups of order 1 and order 2 no matter the basis if we only measure the representation name.

Chapter 6

Conclusion

This thesis has sought to provide a thorough study of the Quantum Fourier Transform for finite groups of which there is a lack in the existing literature. Previously, the mathematics concerned in the construction of the QFT were only briefly summarised in papers whose focus was the presentation of an algorithm that made use of it. Yet, for the purposes of further research in the area of quantum algorithms, a true understanding is necessary. Thus, whereas before the uninitiated would be required to combine elements from many sources in several different fields, now these can be found integrated in a single source.

The relationship of the Quantum Fourier Transform and the Hidden Subgroup Problem was also firmly established. It shows the mechanism by which the QFT creates interference and the extent to which, according to current knowledge, this can be used. We learned that the partial measurement of representation names has well defined limits. Subsequently, we saw that complete measurements might be able to extend these but that they require careful choice of bases for the representations. A choice which would depend

very much on the group and the type of interference desired. These considerations were made explicit by the example of the dihedral Hidden Subgroup Problem.

As a good summary of what we learned of the Quantum Fourier Transform, consider the function given by,

$$|gH\rangle = \sum_{\rho \in R} \left(\sqrt{\frac{d_\rho}{|G|}} \sum_{h \in H} \chi_\rho(h) \right) |\rho\rangle$$

The above gives exactly the same distribution on the representation names as the QFT on any group and for any subgroup H . It allows us to see at a glance the advantages the QFT possess when measurement is restricted to representation names: The right-hand side is independent of the coset and independent of the choice of basis for the representation space of each $\rho \in R$. This last characteristic is the most important. It is the crucial element which is lost when moving from a partial measurement to a complete one.

Earlier we alluded to other problems which have been shown to be special cases of the Hidden Subgroup Problem. Besides Simons problem, these include Shors factorization and discrete logarithm algorithms ([17]) and Kitaevs Abelian stabilizer problem ([11]). Regarding the non-Abelian Hidden Subgroup Problem, there are a few results which were not included in the main body that need mentioning. These were excluded on the basis that although they almost all make use of the Quantum Fourier Transform, the algorithms used contain important variations. This renders the results less interesting for use as illustrative examples of the Quantum Fourier Transform at work. The first step was made by Ettinger and Høyer ([3]). They were able to devise an algorithm for solving the Hidden Subgroup Problem for the dihedral case which used only a linear number of samples from coset states.

However, the algorithm remained inefficient as determining the subgroup required an exponential-time classical algorithm to process the information obtained from the samples. Furthermore, the algorithm uses the Quantum Fourier Transform for the Abelian group $\mathbb{Z}_N \times \mathbb{Z}_2$, rather than the one for dihedral group proper.

Very soon after Rötteler and Beth ([15]) developed the first polynomial-time algorithm for solving the HSP in a noncommutative case. The group they considered was the wreath product group $\mathbb{Z}_N \wr \mathbb{Z}_2$. This was followed by another paper by Ettinger, Høyer and Knill ([5]) which shows that information theoretically, it is possible to determine an arbitrary hidden subgroup for any fixed group with only a linear number of samples. Unfortunately they were unable to discover how this could be done and again their result required an exponential-time classical algorithm to analyse the samples. The results of Hallgren, Russell and Ta-Shma ([7]), as well as those of Grigni, Schulman, Vazirani ([6]) we have already covered. Lately, Ivanyos, Magniez and Santha ([9]) were able to find polynomial time solutions to the Hidden Subgroup Problem for a few select classes of noncommutative groups with the added virtue that they allowed for the group to be given in a black box.

Bibliography

- [1] M. Artin. *Algebra*. (Prentice-Hall, New Jersey), 1991.
- [2] G. Brassard, P. Høyer. *An Exact Quantum Polynomial-Time Algorithm for Simon's Problem*. Proceedings of the Fifth Israeli Symposium on Theory of Computing and Systems. pages 12-33. ISTCS IEEE Computer Society Press, 1997. LANL preprint quant-ph/9704027, April 1997.
- [3] M. Ettinger, P. Høyer. *On Quantum Algorithms for Noncommutative Hidden Subgroups*. LANL preprint quant-ph/9807029, July 1998.
- [4] M. Ettinger, P. Høyer. *Quantum State Detection via Elimination*. LANL preprint quant-ph/9905099, May 1999.
- [5] M. Ettinger, P. Høyer, E. Knill. *Hidden Subgroup States are Almost Orthogonal*. LANL preprint quant-ph/9901034, 1999.
- [6] M. Grigni, L. Schulman, M. Vazirani and U. Vazirani. *Quantum Mechanical Algorithms for the Nonabelian Hidden Subgroup Problem*. Proc. 33rd Annual ACM Symposium on the Theory of Computing-STOC 2001. (ACM Press, New York), 68-74.
- [7] S. Hallgren, A. Russell, A. Ta-Shma. *Normal Subgroup Reconstruction and Quantum Computation Using Group Representations*. Proc. 32nd

- Annual ACM Symposium on the Theory of Computing-STOC 2000. (ACM Press, New York), 627-35.
- [8] P. Høyer. *Efficient Quantum Transforms*. LANL preprint quant-ph/9702028, February 1997.
- [9] G. Ivanyos, F. Magniez, M. Santha. *Efficient Quantum Algorithms for some Instances of the Non-Abelian Hidden Subgroup Problem*. LANL preprint quant-ph/0102014, February 2001.
- [10] R. Jozsa. *Quantum Factoring Discrete Logarithms, and the Hidden Subgroup Problem*. LANL preprint quant-ph/0012084, 2000.
- [11] A. Yu. Kitaev. *Quantum Measurement and the Abelian Stabilizer Problem*. LANL preprint quant-ph/9511026, November, 1995.
- [12] M. Mosca, A. Ekert. *The Hidden Subgroup Problem and Eigenvalue Estimation on a Quantum Computer*. LANL preprint quant-ph/9903071, 1999.
- [13] M. Püschel, M. Rötteler, T. Beth. *Fast Quantum Fourier Transforms for a Class of Non Abelian Groups*. LANL preprint quant-ph/9807064, 1998.
- [14] D. Rockmore. *Some Applications of Generalized FFTs*. DIMACS Series in Discrete Mathematics and Theoretical Computer Science, Volume 28, 1997.
- [15] M. Rötteler, T. Beth. *Polynomial-Time Solution to the Hidden Subgroup Problem for a Class of Non-Abelian Groups*. LANL preprint quant-ph/9812070, December 1998.

- [16] J.-P. Serre. *Linear Representations of Finite Groups*. Graduate Texts in Mathematics, volume 42. (Springer-Verlag, New York), 1977.
- [17] P. W. Shor. *Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer*. SIAM Journal of Computing, Vol. 26, No. 5, 1484-1509, October 1997
- [18] D. R. Simon. *On the Power of Quantum Computing*. SIAM Journal of Computing, Vol. 26, No. 5, 1474-1483, October 1997.

Appendix A

Group Theory

Formally, a *group* is a triple (G, \cdot, e) composed of a set G equipped with an associative operation “ \cdot ” on the members of the set such that,

- (i) There is a unique element e known as the *identity* with the property that $\forall g \in G, e \cdot g = g \cdot e = g$.
- (ii) For every element g of the group there is an *inverse* element g^{-1} with respect to the operation “ \cdot ”. That is, $\forall g \in G, \exists g^{-1} \in G$ such that $g \cdot g^{-1} = g^{-1} \cdot g = e$.

Note that in the above definition, there is no requirement that the group operation be commutative. When this is the case the group is said to be *Abelian*. One should note that it follows from the definition that the inverse of an element is unique.

Typically the group (G, \cdot, e) is referred to simply by G . In agreement with set notation, the number of elements in a group is called the *order* of the group and is denoted $|G|$. We will only deal with those groups of finite order. The structure and properties of groups can vary enormously depending on the choice of set and operation. The following are some common examples:

- The *symmetric* group of order n , S_n , is the set of all permutations of n objects, together with the composition operation. It has order $n!$. Interestingly, it can be shown that every group of order n is equivalent to a subgroup of S_n .
- The group \mathbb{Z}_N is defined by taking, for some fixed integer n , the set $\{0, 1, \dots, n-1\}$. The operation on this set is then addition modulo n . Thus, for $a, b \in \mathbb{Z}_N$, $a \cdot b \equiv a+b \pmod{n}$. It follows that 0 is the identity element of the group. The group \mathbb{Z}_n has order n and is Abelian. It is often taken as the representative of the *cyclic group of order n* .
- For any vector space V over a field \mathbb{K} , the set of all linear transformations from V to V forms a group under composition. It is known as the *general linear group* over V and is denoted $GL_{\mathbb{K}}(V)$. If the vector space V has dimension n then, by fixing a basis of V , the group can be associated to all $n \times n$ invertible matrices with entries in the field \mathbb{K} in the chosen basis. In this case the group operation is matrix multiplication.
- For any vector space V over a field \mathbb{K} , the set of all *unitary* transformations from V to V also forms a group under composition. It is known as the *unitary group* over V and is denoted $U_{\mathbb{K}}(V)$.

A.1 Subgroups

A non empty subset $H \subseteq G$ which fulfills the properties of a group with respect to the operation defined on G is called a *subgroup* of G which we denote by $H \leq G$. Precisely, $H \subseteq G$ is subgroup if we have:

- (i) Closure: $a, b \in H \Rightarrow a \cdot b \in H$

(iii) $a \in H \Rightarrow a^{-1} \in H$.

Suppose we choose some $g \in G$, and consider the set of all powers of g , $\{g, g^2, \dots, g^k, \dots\}$. As each g^i must be an element of the group, which is finite, then there must exist repetitions of group elements in the set. That is, for some integers $i > j$, we must have $g^i = g^j$. It follows by the definition of the group that $g^{i-j} = e$. We conclude that for every element g of a finite group, there exists an integer n such that $g^n = e$. The smallest such positive integer is known as the *order of g* . Now let n be the order of $g \in G$, (also denoted $|g|$), and let $\langle g \rangle = \{e, g, g^2, \dots, g^{n-1}\}$. The set $\langle g \rangle$ forms a subgroup of G of order n . It is said to be the subgroup *generated* by the element g as it consists only of products of g with itself. Indeed, for any subset S of elements of a group we may consider the subgroup generated by S , $\langle S \rangle$. In accordance with what we have already seen, this subgroup consists of all distinct products between elements of S .

If $H \leq G$ is a subgroup then consider the sets $gH = \{gh | h \in H\}$ for all $g \in G$. These are known as the *cosets* of H in G . They define an equivalence relation \equiv_H on the group elements by the property $x \equiv_H y \Leftrightarrow xH = yH \Leftrightarrow y^{-1}x \in H$. As with any equivalence relation, this creates a partition of G into equivalence classes. The definition of the relation implies that the number of partitions is exactly the number of distinct cosets of the subgroup $H \leq G$. Suppose we denote this quantity by $[G : H]$, called the index of H in G , and note that the number of elements in each coset gH is exactly $|H|$. Then this leads to the following theorem,

Theorem 5 (Lagrange) Let G be a group and let $H \leq G$ be a subgroup of G . Then,

$$|G| = [G : H]|H|$$

As a corollary, remark that the order of a subgroup divides the order of the group. Furthermore we saw that the subgroup generated by an element g contained precisely $|g|$ elements. Therefore, by Lagrange, the order of an element must also divide the order of the group. In general it is important to distinguish between *left* cosets gH and *right* cosets Hg when the group operation is not commutative. The properties of these two are the same but they are not interchangeable once a choice has been made. There may exist subgroups for which the left and right cosets are all equal. These are all called *normal* subgroups. By definition, if $N \leq G$ is a normal subgroup of a group G then,

$$gN = Ng \quad \forall g \in G \iff gNg^{-1} = N \quad \forall g \in G$$

If N is a normal subgroup of G , then the set of all distinct cosets of N also forms a group. This is called the *quotient* or *factor* group of N and is written G/N . The operation in this group is defined as follows,

$$\forall aN, bN \in G/N \quad (aN) \cdot (bN) = (ab)N$$

The identity element of the group is the coset $1N = N$. The order of the group G/N is clearly the index of N in G . Therefore for a normal subgroup N , $|G| = |G/N||N|$.

A.2 Isomorphisms and Homomorphisms

Two groups G_1 and G_2 are said to be *isomorphic* if they are equivalent up to a renaming of their elements. When confronted with two groups, to show that this is the case, we require a well defined function $\varphi : G_1 \rightarrow G_2$ which

will establish this correspondence appropriately. Formally, this implies that the function φ must have the following characteristics:

(i) φ is one-to-one. That is,

$$\begin{aligned} \forall y \in G_2 \quad \exists x \in G_1 \quad \text{such that } y = \varphi(x) \\ \varphi(x_1) = \varphi(x_2) \Rightarrow x_1 = x_2 \text{ for all } x_1, x_2 \in G \end{aligned}$$

(ii) φ preserves the group operation for G_1 . Mathematically,

$$\varphi(x)\varphi(y) = \varphi(xy) \quad \text{for all } x, y \in G_1$$

Such a function is known as an *isomorphism* between the groups G_1 and G_2 . Suppose that we remove the first condition. In that case, the function is known as a *homomorphism*. This of course no longer serves to establish equivalence of groups, but is nonetheless very useful. To any homomorphism ϕ we can define two sets, $\text{im } \phi = \{y \in G_2 \mid \phi(x) = y \text{ for some } x \in G_1\}$ called the *image* of ϕ , and $\text{ker } \phi = \{x \in G_1 \mid \phi(x) \text{ is the identity in } G_2\}$ called the *kernel* of ϕ . The image of ϕ is a subgroup of G_2 whereas the kernel of ϕ is a subgroup of G_1 . It is the kernel of a homomorphism which is particularly interesting. It can be shown that the kernel of any homomorphism is always a *normal* subgroup. Furthermore, for every normal subgroup $N \trianglelefteq G$ there exists a homomorphism on G with kernel N . Finally, the quotient group $G/\text{ker } \phi$ is isomorphic to the image of ϕ .

Now that we are familiar with group isomorphisms we can introduce an important fact used many times in this thesis. Let (G, \cdot, e) and (G', \circ, e') be two groups. From these we can form a new group $G \times G'$ called the *direct product* of G and G' . Each element of $G \times G'$ is an ordered pair (g, g') where $g \in G$ and $g' \in G'$. The identity element is (e, e') and the operation between

the elements takes place as follows,

$$(x, x')(y, y') = (x \cdot y, x' \circ y')$$

It is easily proved that every Abelian group is isomorphic to a direct product of cyclic groups.

A.3 Group Actions

A final concept to which we refer in the body of the thesis is that of the *action* of a group on a set. Let G be a group and let S be some set of objects. Suppose that to each $g \in G$ we associate an operation ϕ_g from S to itself. Then, if the operation meets certain conditions, the group is said to *act* on the set S . The common notation for group actions is unfortunately somewhat misleading, for it is customary to denote $\phi_g(s)$ by gs . However, as long as we bear in mind that gs is an element of S and does not represent the product of g and s in general, confusion can be avoided.

We require two properties of the action of G on S : First that the action of the identity of the group be the identity operation. That is, $es = s$ for all $s \in S$. The second is the associativity of the action $(hg)s = h(gs)$ for all $h, g \in G$ and $s \in S$.

Every group actions defines two main structures, the orbit of $s \in S$, O_s , and the stabilizer of s , G_s . The former is the set of all images of s under the action of G . That is,

$$O_s = \{gs | g \in G\}$$

Reminiscent of cosets, O_s is a subset of S and defines an equivalence relation in S by which two elements of S are consider equivalent if the belong to the

same orbit. The stabilizer of an element $s \in S$ is a subgroup of G consisting of all those group elements which fix s . Hence,

$$G_s = \{g \in G \mid gs = s\}$$

A common example of a group action occurs when we set $S = G$ and define the operation to be that of left multiplication. Formally, to each $g \in G$ we associated the map $g_L : G \rightarrow G$ such that,

$$\forall x \in G \quad g_L(x) = g \cdot x$$

In this case, the orbit of each group element is the entire group itself and the stabilizer is the identity. Much more interesting is to have a group G act on the cosets of some subgroup $H \leq G$, again by left multiplication. Thus, for all $xH \in G/H$, $g(xH) = (gx)H$. It is easily verified that $O_H = G/H$ and $G_H = H$, where $H = eH \in G/H$. It turns out that every action of a group on a set is equivalent to the action by left multiplication on the cosets of some subgroup of H , as stated in the following:

Theorem 6 Suppose G acts on some set S . Let G_s be the stabilizer for some $s \in S$. As G_s is a subgroup of G , there is a bijective map $\psi : G/G_s \rightarrow O_s$ such that,

$$\psi(xG_s) = xs$$

Not only does this provide a simple general form for describing all group actions but it also leads to a result inherited from the theorem of Lagrange,

$$|G| = |G_s| |O_s| \quad \forall s \in S$$

for every group action on a set S .

Appendix B

Concepts from Linear Algebra

B.1 Direct Sums

Let V be a vector space with subspaces W_1, W_2, \dots, W_n . The *span* of these subspaces is the set of all vectors v that can be written as a sum of the following form,

$$v = w_1 + w_2 + \dots + w_n \quad \text{such that } w_i \in W_i \text{ for all } i = 1, \dots, n$$

This span is denoted by $W_1 + W_2 + \dots + W_n$. These subspaces are said to be *independent* if every set of vectors $\{w_1, \dots, w_n | w_i \in W_i\}$ is an independent set in the usual way for vectors. Now, if W_1, W_2, \dots, W_n are independent subspaces of V and additionally they span the entire vector space V , then we say that V is the *direct sum* of the subspaces W_1, \dots, W_n and write

$$V = W_1 \oplus W_2 \oplus \dots \oplus W_n$$

B.2 Unitary Matrices

Let A be an $n \times n$ matrix with complex entries. The *adjoint* matrix, A^\dagger , of A is its conjugate transpose. That is,

$$\text{If } A = (a_{ij}) \text{ then } A^\dagger = (a'_{ij}) \text{ where } a'_{ij} = \overline{a_{ji}}$$

Note that especially in references with a more mathematical focus, it is common to see the adjoint denoted A^* . A matrix is said to be *unitary* if its adjoint is also its inverse. In other words, an $n \times n$ matrix A is unitary if and only if,

$$AA^\dagger = A^\dagger A = 1_{n \times n}$$

By a result known as the Spectral Theorem, it can be shown that every unitary matrix is diagonalizable. But there are other more powerful properties. The columns of a unitary matrix are mutually orthonormal, as are the rows. The converse statement is that if the columns of an $n \times n$ matrix are orthonormal then the matrix is unitary. Additionally, the product of unitary matrices is also unitary and the determinant of a unitary matrix is always 1. If A is a unitary matrix then so are A^t , \overline{A} , A^{-1} and the matrix,

$$\begin{bmatrix} I & 0 \\ 0 & A \end{bmatrix}$$

Another fact is that unitary matrices preserve the norm of a vector and more generally the inner product $(X, Y) = \overline{X}^t Y$, which is a necessary and sufficient condition for unitary matrices.