The Symmetric Group Defies Strong Fourier Sampling: Part II

Cristopher Moore
moore@cs.unm.edu
Department of Computer Science
University of New Mexico

Alexander Russell
acr@cse.uconn.edu
Department of Computer Science and Engineering
University of Connecticut

January 18, 2005

Abstract

Part I of this paper showed that the hidden subgroup problem over the symmetric group—including the special case relevant to Graph Isomorphism—cannot be efficiently solved by strong Fourier sampling, even if one may perform an arbitrary POVM on the coset state. In this paper, we extend these results to entangled measurements. Specifically, we show that the hidden subgroup problem on the symmetric group cannot be solved by any POVM applied to pairs of coset states. In particular, these hidden subgroups cannot be determined by any polynomial number of one- or two-register experiments on coset states.

1 Introduction: the hidden subgroup problem

Many problems of interest in quantum computing can be reduced to an instance of the *Hidden Subgroup Problem* (HSP). This is the problem of determining a subgoup H of a group G given oracle access to a function $f : G \to S$ with the property that

$$f(g) = f(hg) \Leftrightarrow h \in H \quad .$$

Equivalently, f is constant on the cosets of H and takes distinct values on distinct cosets.

All known efficient solutions to the problem rely on the *standard method* or the method of *Fourier sampling* [3], described below.

Step 1. Prepare two registers, the first in a uniform superposition over the elements of G and the second with the value zero, yielding the state

$$\psi_1 = \frac{1}{\sqrt{|G|}} \sum_{g \in G} |g\rangle \otimes |0\rangle \quad .$$

Step 2. Query (or calculate) the function f defined on G and XOR it with the second register. This entangles the two registers and results in the state

$$\psi_2 = \frac{1}{\sqrt{|G|}} \sum_{g \in G} |g\rangle \otimes |f(g)\rangle \quad .$$

Step 3. Measure the second register. This puts the first register in a uniform superposition over one of f's level sets, i.e., one of the cosets of H, and disentangles it from the second register. If we observe the value $f(c)$, we have the state $\psi_3 \otimes |f(c)\rangle$ where

$$\psi_3 = |cH\rangle = \frac{1}{\sqrt{|H|}} \sum_{h \in H} |ch\rangle \quad .$$

Step 4. Carry out the quantum Fourier transform on ψ_3 and measure the result.

The result of Step 3 above is the *coset state* $|cH\rangle$, where c is chosen uniformly from G. Expressing this as a mixed state, let

$$\rho_H = \frac{1}{|G|}\sum_{c\in G} |cH\rangle \langle cH| \quad .$$

We shall focus on the *hidden conjugate problem*, where the hidden subgroup is a random conjugate $H^g = g^{-1}Hg$ of a known (non-normal) subgroup H. It was shown in the first part of this paper that when $G = S_{2n}$, the symmetric group on $2n$ letters, and H is a the subgroup generated by the involution $(1\ 2)\ldots(2n-1\ 2n)$, the outcome of *any* measurement on ρ_H is nearly independent of the random choice of $g \in S_n$. In particular, no polynomial number of coset state experiments can determine such a hidden subgroup with non-negligible probability.

It is known, however, that a measurement exists to determine hidden subgroups of a group G from $k = \operatorname{poly}\log|G|$ independent copies of ρ_H. In light of the discussion above, this measurement cannot, in general, be a product measurement: it must involve *entangled* measurement operators. In this paper, we extend the framework of part I to such entangled measurements, showing that for the subgroup H of S_{2n} described above, the result of any measurement of the two-coset state $\rho_{H^g} \otimes \rho_{H^g}$ is nearly independent of g. In particular, no polynomial number of two-register coset-state experiments can determine H with non-negligible probability.

Related work. Both Simon's and Shor's seminal algorithms rely on the standard method over an Abelian group. In Simon's problem [29], $G = \mathbb{Z}_2^n$ and f is an oracle such that, for some y, $f(x) = f(x+y)$ for all x; in this case $H = \{0, y\}$ and we wish to identify y. In Shor's factoring algorithm [28] G is the group $\mathbb{Z}_n^*$ where n is the number we wish to factor, $f(x) = r^x \bmod n$ for a random $r < n$, and H is the subgroup of $\mathbb{Z}_n^*$ whose index is the multiplicative order of r. (Note that in Shor's algorithm, since $|\mathbb{Z}_n^*|$ is unknown, the Fourier transform is performed over $\mathbb{Z}_q$ for some $q = \operatorname{poly}(n)$; see [28] or [10, 11].)

For such abelian instances; it is not hard to see that a polynomial number (i.e., polynomial in $\log|G|$) of experiments of this type determine H. In essence, each experiment yields a random element of the dual space $H^\perp$ perpendicular to H's characteristic function, and as soon as these elements span $H^\perp$ they, in particular, determine H.

While the *nonabelian* hidden subgroup problem appears to be much more difficult, it has very attractive applications. In particular, solving the HSP for the symmetric group S_n would provide an efficient quantum algorithm for the Graph Automorphism and Graph Isomorphism problems (see e.g. Jozsa [17] for a review). Another important motivation is the relationship between the HSP over the dihedral group with hidden shift problems [4] and cryptographically important cases of the Shortest Lattice Vector problem [23].

So far, algorithms for the HSP are only known for a few families of nonabelian groups, including wreath products $\mathbb{Z}_2^k \wr \mathbb{Z}_2$ [24]; more generally, semidirect products $K \ltimes \mathbb{Z}_2^k$ where K is of polynomial size, and groups whose commutator subgroup is of polynomial size [16]; "smoothly solvable" groups [7]; and some semidirect products of cyclic groups [14]. Ettinger and Høyer [5] provided another type of result, by showing that Fourier sampling can solve the HSP for the dihedral groups D_n in an *information-theoretic* sense. That is, a polynomial number of experiments gives enough information to reconstruct the subgroup, though it is unfortunately unknown how to determine H from this information in polynomial time.

To discuss Fourier sampling for a nonabelian group G, one needs to develop the Fourier transform over G. For abelian groups, the Fourier basis functions are homomorphisms $\phi : G \to \mathbb{C}$ such as the familiar exponential function $\phi_k(x) = e^{2\pi i k x/n}$ for the cyclic group $\mathbb{Z}_n$. In the nonabelian case, there are not enough such homomorphisms to span the space of all $\mathbb{C}$-valued functions on G; to complete the picture, one introduces *representations* of the group, namely homomorphisms $\rho : G \to \mathsf{U}(V)$ where $\mathsf{U}(V)$ is the group of unitary matrices acting on some $\mathbb{C}$-vector space V of dimension d_ρ. It suffices to consider *irreducible* representations, namely those for which no nontrivial subspace of V is fixed by the various operators $\rho(g)$. Once a basis for each irreducible ρ is chosen, the matrix elements ρ_{ij} provide an orthogonal basis for the vector space of all $\mathbb{C}$-valued functions on G.

The quantum Fourier transform then consists of transforming (unit-length) vectors in $\mathbb{C}[G] = \{\sum_{g \in G} \alpha_g |g\rangle \mid \alpha_g \in \mathbb{C}\}$ from the basis $\{|g\rangle \mid g \in G\}$ to the basis $\{|\rho, i, j\rangle\}$ where ρ is the name of an irreducible representation and $1 \leq i, j \leq d_\rho$ index a row and column (in a chosen basis for V). Indeed, this transformation can be carried out efficiently for a wide variety of groups [2, 13, 21]. Note, however, that a nonabelian group G does not distinguish any specific basis for its irreducible representations which necessitates a rather dramatic choice on the part of the transform designer. Indeed, careful basis selection appears to be critical for obtaining efficient Fourier transforms for the groups mentioned above.

Perhaps the most fundamental question concerning the hidden subgroup problem is whether there is always a basis for the representations of G such that measuring in this basis (in Step 4, above) provides enough information to determine the subgroup H. This framework is known as *strong Fourier sampling*. Part I of this article answers this question in the negative, showing that natural subgroups of S_n cannot be determined by this process; in fact, it shows this for an even more general model, where we perform an arbitrary positive operator-valued measurement (POVM) on coset states $|cH\rangle$. We emphasize that this result includes the most important special cases of the nonabelian HSP, as they are those to which Graph Isomorphism naturally reduces. Namely, as in [12] we focus on order-2 subgroups of the form $\{1, m\}$, where m is an involution consisting of $n/2$ disjoint transpositions; then if we fix two rigid connected graphs of size $n/2$ and consider permutations of their disjoint union, then the hidden subgroup is of this form if the graphs are isomorphic and trivial if they are not.

The next logical step is to consider *multi-register* algorithms, in which we prepare multiple coset states and subject them to *entangled* measurements. Ettinger, Høyer and Knill [6] showed that the HSP on arbitrary groups can be solved information-theoretically with a polynomial number of registers, although their algorithm takes exponential time for most groups of interest. Kuperberg [20] devised a subexponential ($2^{O(\sqrt{\log n})}$) algorithm for the HSP on the dihedral group D_n that works by performing entangled measurements on two registers at a time, and Bacon, Childs, and van Dam [1] have determined the optimal multiregister measurement for the dihedral group.

Whether a similar approach can be taken for the symmetric group is a major open question. In this paper, we take a step towards answering this question by showing that if we perform arbitrary entangled measurements over *pairs* of registers, distinguishing $H = \{1, m\}$ from the trivial group in S_n requires a superpolynomial number (specifically, $e^{\Omega(\sqrt{n}/\log n)}$) of experiments.

2 Two combinatorial representations

With apologies to the reader, we will rely on the introductory sections of Part I rather than repeating them here. However, here we introduce two combinatorial representations that will be extremely useful to us.

For a group G, we let $\widehat{G}$ denote a collection of unitary representations of G consisting of exactly one from each isomorphism class. We let $\mathbb{C}[G]$ denote the *group algebra* of G; this is the $|G|$-dimensional vector space of formal sums

$$\left\{ \sum_g \alpha_g \cdot g \mid \alpha_g \in \mathbb{C} \right\}$$

equipped with the unique inner product for which $\langle g, h\rangle$ is equal to one when $g = h$ and zero otherwise. (Note that $\mathbb{C}[G]$ is precisely the Hilbert space of a single register containing a superposition of group elements.)

We introduce two combinatorial representations related to the group algebra. The first is the *regular* representation R, given by the permutation action of G on itself. Then R is the representation $\mathrm{R} : G \to \mathsf{U}(\mathbb{C}[G])$ given by linearly extending left multiplication, $\mathrm{R}(g) : h \mapsto gh$. It is not hard to see that its character χ_{R} is given by

$$\chi_{\mathrm{R}}(g) = \begin{cases} |G| & g = 1 \ , \\ 0 & g \neq 1 \ , \end{cases}$$

in which case we have $\langle \chi_{\mathrm{R}}, \chi_\rho \rangle_G = d_\sigma$ for each $\rho \in \widehat{G}$. Thus $\mathbb{R}$ contains d_ρ copies of each irreducible $\rho \in \widehat{G}$, and counting dimensions on each side of this decomposition implies $|G| = \sum_{\rho \in \widehat{G}} d_\rho^2$.

The other combinatorial representation we will rely on is the *conjugation* representation C, given by the conjugation action of G on $\mathbb{C}[G]$. Specifically, $\mathrm{C} : G \to \mathsf{U}(\mathbb{C}[G])$ is the map obtained by linearly extending the rule $\mathrm{C}(g) : h \mapsto ghg^{-1}$. While the decomposition of C into irreducibles is, in general, unknown, one does have the decomposition

$$\mathrm{C} = \bigoplus_{\rho \in \widehat{G}} \rho \otimes \rho^* \quad \text{and therefore} \quad \chi_{\mathrm{C}}(g) = \sum_{\rho \in \widehat{G}} \chi_\rho(g) \chi_\rho(g)^* \ . \tag{2.1}$$

Here ρ^* denotes the complex conjugate representation of ρ, which acts on vectors $\mathbf{u}^*$ as $\rho^*(g)\mathbf{u} = (\rho(g)\mathbf{u})^*$. We also note that an elementary argument shows that

$$\chi_{\mathrm{C}}(g) = \frac{|G|}{|[g]|} \ ,$$

where $[g] = \{h^{-1}gh \mid h \in G\}$ denotes the conjugacy class of g.

3 Background from Part I

3.1 The structure of the optimal measurement

As in Part I, we focus on the special case of the hidden subgroup problem called the *hidden conjugate problem* in [22]. Here there is a (non-normal) subgroup H, and we are promised that the hidden subgroup is one of its conjugates, $H^g = g^{-1}Hg$ for some $g \in G$; the goal is to determine which.

The most general possible measurement in quantum mechanics is a positive operator-valued measurement (POVM). Part I of this paper establishes that the optimal POVM for the Hidden Subgroup Problem on a single coset state consists of measuring the name ρ of the irreducible representation, followed by a POVM on the vector space V on which ρ acts. In the special case of a von Neumann measurement, this corresponds to measuring the row of ρ in some orthonormal basis; in general it consists of measuring according to some over-complete basis, or *frame*, $B = \{\mathbf{b}\}$ with positive real weights $a_{\mathbf{b}}$ that obeys the completeness condition

$$\sum_{\mathbf{b}} a_{\mathbf{b}} \pi_{\mathbf{b}} = \mathbb{1} \ , \tag{3.1}$$

where $\pi_{\mathbf{b}}$ denotes the projection onto the unit length vector $\mathbf{b}$. We remark that the frame B weighted according to a is energy-conserving in the sense that

$$\|\mathbf{x}\|^2 = \langle \mathbf{x}, \mathbb{1}\mathbf{x} \rangle = \langle \mathbf{x}, \sum_{\mathbf{b}} a(\mathbf{b}) \pi_{\mathbf{b}}(\mathbf{x}) \rangle = \sum_{\mathbf{b}} a_{\mathbf{b}} \|\pi_{\mathbf{b}}(\mathbf{x})\|^2 \ .$$

During Fourier sampling, the probability we observe ρ, and the conditional probability that we observe a given $\mathbf{b} \in B$, are given by

$$P(\rho) = \frac{d_\rho |H|}{|G|} \, \mathbf{rk}\, \Pi_H \tag{3.2}$$

$$P(\rho, \mathbf{b}) = a_j \frac{\|\Pi_H \mathbf{b}\|^2}{\mathbf{rk}\, \Pi_H} \tag{3.3}$$

where Π_H is the projection operator $1/|H| \sum_{h \in H} \rho(h)$. In the case where H is the trivial subgroup, $\Pi_H = \mathbb{1}_{d_\rho}$ and $P(\rho, \mathbf{b}_j)$ is given by

$$P(\rho, \mathbf{b}) = \frac{a_{\mathbf{b}}}{d_\rho} \ . \tag{3.4}$$

We call this the *natural distribution* on the frame $B = \{\mathbf{b}\}$. In the case that B is an orthonormal basis, $a_{\mathbf{b}} = 1$ and this is simply the uniform distribution. This probability distribution over B changes for a conjugate H^g in the following way:

$$P(\rho, \mathbf{b}) = a_j \frac{\|\Pi_H g \mathbf{b}\|^2}{\mathbf{rk}\, \Pi_H}$$

where we write $g\mathbf{b}$ for $\rho(g)\mathbf{b}$. It is not hard to show that, for any $\mathbf{b} \in V$, the *expected* value of $\|\Pi_H(g\mathbf{b})\|^2$, over the choice of $g \in G$, is $\mathbf{rk}\,\Pi_H/d_\rho$.

3.2 The expectation and variance of an involution projector

The following lemmas are proved in Part I; we repeat them here for convenience.

Lemma 1. *Let ρ be a representation of a group G acting on a space V and let $\mathbf{b} \in V$. Let m be an element chosen uniformly from a conjugacy class I of involutions. If ρ is irreducible, then*

$$\mathrm{Exp}_m \langle \mathbf{b}, m\mathbf{b} \rangle = \frac{\chi_\rho(I)}{\dim \rho} \|\mathbf{b}\|^2 \quad .$$

If ρ is reducible, then

$$\mathrm{Exp}_m \langle \mathbf{b}, m\mathbf{b} \rangle = \sum_{\sigma \prec \rho} \frac{\chi_\sigma(I)}{\dim \sigma} \|\Pi_\sigma^\rho \mathbf{b}\|^2 \quad .$$

Lemma 2. *Let ρ be a representation of a group G acting on a space V and let $\mathbf{b} \in V$. Let m be an element chosen uniformly at random from a conjugacy class I of involutions. Then*

$$\mathrm{Exp}_m |\langle \mathbf{b}, m\mathbf{b} \rangle|^2 = \sum_{\sigma \prec \rho \otimes \rho^*} \frac{\chi_\sigma(I)}{\dim \sigma} \left\| \Pi_\sigma^{\rho \otimes \rho^*} (\mathbf{b} \otimes \mathbf{b}^*) \right\|^2 \quad .$$

Given an involution m and the hidden subgroup $H = \{1, m\}$, let $\Pi_m = \Pi_H$ denote the projection operator given by

$$\Pi_m \mathbf{v} = \frac{\mathbf{v} + m\mathbf{v}}{2} \quad .$$

Then the expectation and variance of $\|\Pi_m \mathbf{b}\|^2$ are given by the following lemma.

Lemma 3. *Let ρ be an irreducible representation acting on a space V and let $\mathbf{b} \in V$. Let m be an element chosen uniformly at random from a conjugacy class I of involutions. Then*

$$\mathrm{Exp}_m \|\Pi_m \mathbf{b}\|^2 = \frac{1}{2} \|\mathbf{b}\|^2 \left(1 + \frac{\chi_\rho(I)}{\dim \rho} \right) \tag{3.5}$$

$$\mathrm{Var}_m \|\Pi_m \mathbf{b}\|^2 \leq \frac{1}{4} \sum_{\sigma \prec \rho \otimes \rho^*} \frac{\chi_\sigma(I)}{\dim \sigma} \left\| \Pi_\sigma^{\rho \otimes \rho^*} (\mathbf{b} \otimes \mathbf{b}^*) \right\|^2 \quad . \tag{3.6}$$

Finally, we point out that since

$$\mathrm{Exp}_m \|\Pi_m \mathbf{b}\|^2 = \|\mathbf{b}\|^2 \frac{\mathbf{rk}\,\Pi_m}{\dim \rho}$$

we have

$$\frac{\mathbf{rk}\,\Pi_m}{\dim \rho} = \frac{1}{2} \left(1 + \frac{\chi_\rho(I)}{\dim \rho} \right) \quad . \tag{3.7}$$

3.3 The representation theory of the symmetric group

We will use several specific properties of the symmetric group S_n and its asymptotic representation theory; we refer the reader to Section 5 of Part I for more background and notation. Recall that the irreducible representations S^λ of S_n are labeled by Young diagrams λ, and that the number of irreducible representations is the partition number $p(n)$. We denote the dimension and character of S^λ as d^λ and χ^λ, respectively. Recall also that the *Plancherel distribution* assigns the probability $d_\rho^2/|G|$ to each irreducible representation ρ. Then we will rely on the following results of Vershik and Kerov.

Theorem 4 ([30]). *Let S^λ be chosen from $\widehat{S_n}$ according to the Plancherel distribution. Then there exist positive constants c_1 and c_2 for which*

$$\lim_{n\to\infty} \Pr\left[e^{-c_1\sqrt{n}}\sqrt{n!} \le d^\lambda \le e^{-c_2\sqrt{n}}\sqrt{n!}\right] = 1 \ .$$

Theorem 5 ([30]). *There exist positive constants $\check{c}$ and $\hat{c}$ such that for all $n \ge 1$,*

$$e^{-\check{c}\sqrt{n}}\sqrt{n!} \le \max_{S^\lambda \in \widehat{S_n}} d^\lambda \le e^{-\hat{c}\sqrt{n}}\sqrt{n!} \ .$$

In Part I we prove the following:

Lemma 6. *Let S^λ be chosen according to the Plancherel distribution on $\widehat{S_n}$.*

1. *Let $\delta = \pi\sqrt{2/3}$. Then for sufficiently large n, $\Pr\left[d^\lambda \le e^{-\delta\sqrt{n}}\sqrt{n!}\right] < e^{-\delta\sqrt{n}}$.*
2. *Let $0 < c < 1/2$. Then $\Pr[d^\lambda \le n^{cn}] = n^{-\Omega(n)}$.*

Finally, we will also apply Roichman's [25] estimates for the characters of the symmetric group:

Definition 1. *For a permutation $\pi \in S_n$, define the* support *of π, denoted $\mathrm{supp}(\pi)$, to be the cardinality of the set $\{k \in [n] \mid \pi(k) \neq k\}$.*

Theorem 7 ([25]). *There exist constants $b > 0$ and $0 < q < 1$ so that for $n > 4$, for every conjugacy class C of S_n, and every irreducible representation S^λ of S_n,*

$$\left|\frac{\chi^\lambda(C)}{d^\lambda}\right| \le \left(\max\left(q, \frac{\lambda_1}{n}, \frac{\lambda_1'}{n}\right)\right)^{b\cdot\mathrm{supp}(C)} \quad ,$$

where $\mathrm{supp}(C) = \mathrm{supp}(\pi)$ for any $\pi \in C$.

In our application, we take n to be even and consider involutions m in the conjugacy class of elements consisting of $n/2$ disjoint transpositions, $M = M_n = \{\sigma\left((12)(34)\cdots(n-1\ n)\right)\sigma^{-1} \mid \sigma \in S_n\}$. Note that each $m \in M_n$ is associated with one of the $(n-1)!!$ perfect matchings of n things, and that $\mathrm{supp}(m) = n$.

3.4 Strong Fourier sampling on one register

The main result of Part I is the following.

Theorem 8. *Let $B = \{\mathbf{b}\}$ be a frame with weights $\{a_\mathbf{b}\}$ satisfying the completeness condition (3.1) for an irreducible representation S^λ. Given the hidden subgroup $H = \{1, m\}$ where m is chosen uniformly at random from M, let $P_m(\mathbf{b})$ be the probability that we observe the vector $\mathbf{b}$ conditioned on having observed the representation name S^λ, and let N be the natural distribution (3.4) on B. Then there is a constant $\delta > 0$ such that for sufficiently large n, with probability at least $1 - e^{-\delta n}$ in m and λ, we have*

$$\|P_m - N\|_1 < e^{-\delta n} \ .$$

The proof strategy is to bound $\mathrm{Var}_m \|\Pi_m \mathbf{b}\|^2$ using Lemma 3, and apply Chebyshev's inequality to conclude that it is almost certainly close to its expectation. Recall, however, that our bounds on the variance of $\|\Pi_m \mathbf{b}\|^2$ depend on the decomposition of $S^\lambda \otimes (S^\lambda)^*$ is into irreducibles and, furthermore, on the projection of $\mathbf{b} \otimes \mathbf{b}^*$ into these irreducible subspaces. Matters are somewhat complicated by the fact that certain S^μ appearing in $S^\lambda \otimes (S^\lambda)^*$ may contribute more to the variance than others. While Theorem 7 allows us to bound the contribution of those constituent irreducible representations S^μ for which μ_1 and μ_1' are much smaller than n, those which violate this condition could conceivably contribute large terms to the variance estimates. Fortunately, in this single coset case, the total fraction of the space $S^\lambda \otimes (S^\lambda)^*$, dimensionwise, consisting of such S^μ is small with overwhelming probability. Despite this, we cannot preclude the possibility that for a *specific* vector $\mathbf{b}$, the quantity $\operatorname{Var}\|\Pi_m \mathbf{b}\|^2$ is large, as $\mathbf{b}$ may project solely into spaces of the type described above. On the other hand, as these troublesome spaces amount to a small fraction of $S^\lambda \otimes (S^\lambda)^*$, only a few $\mathbf{b}$ can have this property.

Specifically, let $0 < c < 1/4$ be a constant, and let $\Lambda = \Lambda_c$ denote the collection of Young diagrams μ with the property that either $\mu_1 \geq (1-c)n$ or $\mu'_1 \geq (1-c)n$. Then Part I establishes the following upper bounds on the cardinality of Λ and the dimension of any S^μ with $\mu \in \Lambda$:

Lemma 9. *Let $p(n)$ denote the number of integer partitions of n. Then $|\Lambda| \leq 2cnp(cn)$, and $d^\mu < n^{cn}$ for any $\mu \in \Lambda$.*

As a result, the representations associated with diagrams in Λ constitute a negligible fraction of $\widehat{S_n}$; specifically, from Lemma 6, part 2, the probability that a λ drawn according to the Plancherel distribution falls into Λ is $n^{-\Omega(n)}$. The following lemma shows that this is also true for the distribution $P(\rho)$ induced on $\widehat{S_n}$ by weak Fourier sampling the coset state $|H\rangle$.

Lemma 10. *Let $d < 1/2$ be a constant and let n be sufficiently large. Then there is a constant $\gamma > 0$ such that we observe a representation S^λ with $d^\lambda \geq n^{dn}$ with probability at least $1 - n^{-\gamma n}$.*

On the other hand, for a representation S^μ with $\mu \notin \Lambda$, Theorem 7 implies that

$$\left|\frac{\chi^\mu(M)}{d^\mu}\right| \leq \big(\max(q, 1-c)\big)^{bn} \leq e^{-\alpha n} \tag{3.8}$$

for a constant $\alpha \geq bc > 0$. Thus the contribution of such an irreducible to the variance estimate of Lemma 3 is exponentially small. The remainder of the proof of Theorem 8 uses a combination of Chebyshev's and Markov's inequalities to bound the total variation distance between P_m and the natural distribution.

4 Variance and decomposition for multiregister experiments

We turn now to the multi-register case, where Steps 1, 2 and 3 are carried out on k independent registers. This yields a state in $\mathbb{C}[G^k]$, i.e.,

$$|c_1 H\rangle \otimes \cdots \otimes |c_k H\rangle$$

where the c_i are uniformly random coset representatives. The symmetry argument of Section 3 of Part I applies to each register, so that the optimal measurement is consistent with first measuring the representation name in each register. However, the optimal measurement generally does not consist of k independent measurements on this tensor product state; rather, it is entangled, consisting of measurement in a basis whose basis vectors $\mathbf{b}$ are not of the form $\mathbf{b}_1 \otimes \cdots \otimes \mathbf{b}_k$. As mentioned above, for the dihedral groups in particular, a fair amount is known: Ip [15] showed that the optimal measurement for two registers is entangled, Kuperberg [20] showed that an entangled measurement yields a subexponential-time algorithm for the hidden subgroup problem, and Bacon, Childs and van Dam [1] have calculated the optimal measurement on k registers.

Extending the results of part I to this case involving multiple coset states will proceed in three steps:

- In Section 4.1, we generalize the expectation and variance bounds of Lemma 3 to the algebra $\mathbb{C}[G^k]$, viewed as a representation of G.
- As in the single register proof, we must control the decomposition of the representations that appear in the expressions for expectation and variance. Unfortunately, the naive bounds applied in part I (relying on the fact that $\langle \chi_\sigma, \chi_\rho \chi_\tau \rangle_G \leq d_\sigma$ for irreducible representations ρ, σ, and τ) are insufficient for controlling these decompositions. In Section 4.2, we show how to bound the decomposition of these representations on average.
- Finally, in Section 5, we show how to apply these results to eliminate the possibility of solving the HSP over S_n with any polynomial number of two-register experiments on coset states.

4.1 Variance for Fourier sampling product states

We begin by generalizing Lemmas 1, 2, and 3 of Part I to the multi-register case. The reasoning is analogous to that of Section 4 of Part I; the principal difficulty is notational, and we ask the reader to bear with us.

We assume we have measured the representation name on each of the registers, and that we are currently in an irreducible representation of G^k labeled by $\rho_1 \otimes \cdots \otimes \rho_k$. Given a subset $I \subseteq \{1, \ldots, k\}$, we can separate this tensor product into the registers inside I and those outside, and then decompose the product of those inside I into irreducibles σ:

$$\begin{aligned} \rho_1 \otimes \cdots \otimes \rho_k &= \bigotimes_{i \in I} \rho_i \otimes \bigotimes_{i \notin I} \rho_i \\ &= \left(\bigoplus_{\sigma \prec \otimes_{i \in I} \rho_i} a_\sigma^I \sigma \right) \otimes \bigotimes_{i \notin I} \rho_i \end{aligned}$$

where a_σ^I is the multiplicity of σ in $\otimes_{i \in I} \rho_i$. Now given an irrep σ, let Π_σ^I denote the projection operator onto the subspace acted on by

$$a_\sigma^I \sigma \otimes \bigotimes_{i \notin I} \rho_i \ .$$

In other words, Π_σ^I projects the registers in I onto the subspaces isomorphic to σ, and leaves the registers outside I untouched. Note that in the case where I is a singleton we have $\Pi_{\rho_i}^{\{i\}} = \mathbb{1}$.

As before, the hidden subgroup is $H = \{1, m\}$ for an involution m chosen at random from a conjugacy class M. However, we now have, in effect, the subgroup $H^k \subset G^k$, and summing over the elements of H^k gives the projection operator $\Pi_{H^k} = \Pi_m^{\otimes k}$. The probability we observe an (arbitrarily entangled) basis vector $\mathbf{b} \in \rho_1 \otimes \cdots \otimes \rho_k$ is then

$$P_m(\mathbf{b}) = \frac{\left\| \Pi_m^{\otimes k} \mathbf{b} \right\|^2}{\mathbf{rk}\, \Pi_m^{\otimes k}} \ . \tag{4.1}$$

When we calculate the expectation of this over m, we will find ourselves summing the following quantity over the subsets $I \subseteq \{1, \ldots, k\}$:

$$E^I(\mathbf{b}) = \sum_{\sigma \prec \otimes_{i \in I} \rho_i} \frac{\chi^\sigma(M)}{\dim \sigma} \left\| \Pi_\sigma^I \mathbf{b} \right\|^2 \tag{4.2}$$

with $E^\emptyset(\mathbf{b}) = \|\mathbf{b}\|^2$ (since an empty tensor product gives the trivial representation).

For the variance, we will find ourselves dealing with pairs of subsets $I_1, I_2 \subseteq \{1, \ldots, k\}$ and decompositions of the form

$$\begin{aligned} (\rho_1 \otimes \cdots \otimes \rho_k) \otimes (\rho_1^* \otimes \cdots \otimes \rho_k^*) &= \left(\bigotimes_{i \in I_1} \rho_i \otimes \bigotimes_{i \in I_2} \rho_i^* \right) \otimes \left(\bigotimes_{i \notin I_1} \rho_i \otimes \bigotimes_{i \notin I_2} \rho_i^* \right) \\ &= \left(\bigoplus_{\sigma \prec \bigotimes_{i \in I} \rho_i \otimes \bigotimes_{i \in I_2} \rho_i^*} a_\sigma^{I_1, I_2} \sigma \right) \otimes \left(\bigotimes_{i \notin I_1} \rho_i \otimes \bigotimes_{i \notin I_2} \rho_i^* \right) \end{aligned}$$

just as we considered $\rho \otimes \rho^*$ in the one-register case. We can then define a projection operator $\Pi_\sigma^{I_1, I_2}$ onto the subspace acted on by

$$a_\sigma^{I_1, I_2} \sigma \otimes \left(\bigotimes_{i \notin I_1} \rho_i \otimes \bigotimes_{i \notin I_2} \rho_i^* \right)$$

and we define the following quantity,

$$E^{I_1, I_2}(\mathbf{b}) = \sum_{\sigma \prec \bigotimes_{i \in I} \rho_i \otimes \bigotimes_{i \in I_2} \rho_i^*} \frac{\chi^\sigma(M)}{\dim \sigma} \left\| \Pi_\sigma^{I_1, I_2} (\mathbf{b} \otimes \mathbf{b}^*) \right\|^2 \tag{4.3}$$

with $E^{\emptyset,\emptyset}(\mathbf{b}) = \|\mathbf{b}\|^4$.

We can now state the following lemma. The reader can check that (4.5) corresponds exactly to Equation (4.3) of part I in the one-register case.

Lemma 11. *Let $\mathbf{b} \in \rho_1 \otimes \cdots \otimes \rho_k$ and let m be an element chosen uniformly from a conjugacy class M of involutions. Then*

$$\begin{aligned}
\mathrm{Exp}_m \left\|\Pi_m^{\otimes k}\mathbf{b}\right\|^2 &= \frac{1}{2^k}\left(1 + \sum_{I\subseteq\{1,\ldots,k\}:I\neq\emptyset} E^I(\mathbf{b})\right) & (4.4)\\
\mathrm{Var}_m \left\|\Pi_m^{\otimes k}\mathbf{b}\right\|^2 &= \frac{1}{4^k}\sum_{I_1,I_2\subseteq\{1,\ldots,k\}:I_1,I_2\neq\emptyset} E^{I_1,I_2}(\mathbf{b}) - E^{I_1}(\mathbf{b})E^{I_2}(\mathbf{b})^* \ . & (4.5)
\end{aligned}$$

Proof. Let m^I denote the operator that operates on the ith register by m for each $i \in I$ and leaves the other registers unchanged. This acts on $\mathbf{b}$ as $\tau(m)$ where $\tau = \bigotimes_{i\in I}\rho_i(m)$, and Lemma 1 implies that

$$\mathrm{Exp}_m\langle \mathbf{b}, m^I\mathbf{b}\rangle = E^I(\mathbf{b}) \ .$$

Then (4.4) follows from the observation that

$$\Pi_m^{\otimes k}\mathbf{b} = \frac{1}{2^k}\sum_{I\subseteq\{1,\ldots,k\}} m^I\mathbf{b}$$

and so

$$\mathrm{Exp}_m \left\|\Pi_m^{\otimes k}\mathbf{b}\right\|^2 = \mathrm{Exp}_m\langle \mathbf{b}, \Pi_m^{\otimes k}\mathbf{b}\rangle = \frac{1}{2^k}\sum_{I\subseteq\{1,\ldots,k\}}\mathrm{Exp}_m\langle \mathbf{b}, m^I\mathbf{b}\rangle = \frac{1}{2^k}\sum_{I\subseteq\{1,\ldots,k\}} E^I(\mathbf{b}) \ .$$

Separating out the term $E^{\emptyset}(\mathbf{b}) = \|\mathbf{b}\|^2$ completes the proof of (4.4).

Similarly, let the operator $m^{I_1}\otimes m^{I_2}$ act on $\mathbf{b}\otimes\mathbf{b}^*$ by multiplying the ith register of $\mathbf{b}$ by m whenever $i \in I_1$, and multiplying the ith register of $\mathbf{b}^*$ whenever $i \in I_2$. Then it acts as $\tau(m)$ where $\tau = \bigotimes_{i\in I_1}\rho_i(m)\otimes \bigotimes_{i\in I_2}\rho_i^*(m)$, and Lemma 1 implies

$$\mathrm{Exp}_m\langle \mathbf{b}\otimes\mathbf{b}^*, (m^{I_1}\otimes m^{I_2})(\mathbf{b}\otimes\mathbf{b}^*)\rangle = E^{I_1,I_2}(\mathbf{b}) \ .$$

Then analogous to Lemmas 2 and 3, the second moment is

$$\begin{aligned}
\mathrm{Exp}_m \left\|\Pi_m^{\otimes k}\mathbf{b}\right\|^4 &= \mathrm{Exp}_m\langle \mathbf{b}, \Pi_m^{\otimes k}\mathbf{b}\rangle\langle \mathbf{b}^*, \Pi_m^{\otimes k}\mathbf{b}^*\rangle\\
&= \mathrm{Exp}_m\langle \mathbf{b}\otimes\mathbf{b}^*, (\Pi_m^{\otimes k}\otimes\Pi_m^{\otimes k})(\mathbf{b}\otimes\mathbf{b}^*)\rangle\\
&= \frac{1}{4^n}\sum_{I_1,I_2\subseteq\{1,\ldots,k\}}\mathrm{Exp}_m\langle \mathbf{b}\otimes\mathbf{b}^*, (m^{I_1}\otimes m^{I_2})(\mathbf{b}\otimes\mathbf{b}^*)\rangle\\
&= \frac{1}{4^n}\sum_{I_1,I_2\subseteq\{1,\ldots,k\}} E^{I_1,I_2}(\mathbf{b})
\end{aligned}$$

and so the variance is

$$\begin{aligned}
\mathrm{Var}_m \left\|\Pi_m^{\otimes k}\mathbf{b}\right\|^2 &= \mathrm{Exp}_m \left\|\Pi_m^{\otimes k}\mathbf{b}\right\|^4 - \left(\mathrm{Exp}_m \left\|\Pi_m^{\otimes k}\mathbf{b}\right\|^2\right)^2\\
&= \frac{1}{4^k}\sum_{I_1,I_2\subseteq\{1,\ldots,k\}} E^{I_1,I_2}(\mathbf{b}) - E^{I_1}(\mathbf{b})E^{I_2}(\mathbf{b})^* \ .
\end{aligned}$$

Finally, (4.5) follows from the fact that the two terms in the sum cancel whenever I_1 or I_2 is empty. □

4.2 The associated Clebsch-Gordan problem

The expressions $E^I(\mathbf{b})$ and $E^{I_1,I_2}(\mathbf{b})$ above depend on the decomposition of representations of the form

$$\bigotimes_{i\in I_1}\rho_i \otimes \bigotimes_{i\in I_2}\rho_i^* = \left(\bigotimes_{i\in I_1\setminus I_2}\rho_i \otimes \bigotimes_{i\in I_2\setminus I_1}\rho_i^*\right) \otimes \bigotimes_{i\in I_1\cap I_2}(\rho_i\otimes\rho_i^*)\ .$$

Moreover, since the Plancherel distribution is symmetric with respect to conjugation, this is a tensor product of $|I_1 \triangle I_2|$ representations ρ_i with $|I_1\cap I_2|$ representations $\sigma_j\otimes\sigma_j^*$, where both the ρ_i and the σ_j are chosen according to the Plancherel distribution. This motivates the following definition.

Definition 2. *For non-negative integers k and ℓ and $\rho = (\rho_1,\ldots,\rho_k)\in\widehat{G}^k$ and $\sigma=(\sigma_1,\ldots,\sigma_\ell)\in\widehat{G}^\ell$, let $V(\rho,\sigma)$ denote the representation*

$$\bigotimes_{i=1}^{k}\rho_i\otimes\bigotimes_{j=1}^{\ell}(\sigma_j\otimes\sigma_j^*)\ .$$

Of particular interest is the dimensionwise fraction of such representations consisting of low-dimensional irreducibles. For these representations, the naive decomposition results of Equation (6.4) of Part I no longer suffice to obtain nontrivial estimates. Fortunately, the combinatorial representations R and C discussed in Section 2 can be used to control the structural properties of these tensor products *on average*. We will apply this machinery in Section 5 to control general two-register experiments.

Recall that the multiplicity of an irreducible representation τ in the decomposition of a representation V into irreducibles is the inner product $\langle\chi_\tau,\chi_V\rangle_G$, and that $[g]$ denotes the conjugacy class of g.

Lemma 12. *Fix $\tau\in\widehat{G}$ and let ρ and σ be random variables taking values in $\widehat{G}^k$ and $\widehat{G}^\ell$, respectively, so that each ρ_i and σ_j is independently distributed according to the Plancherel distribution. Then*

$$\begin{aligned}
\operatorname{Exp}_{\rho,\sigma}\frac{\left\langle\chi_\tau,\chi_{V(\rho,\sigma)}\right\rangle_G}{\dim V(\rho,\sigma)} &= \frac{d_\tau}{|G|} \quad \textit{if } k>0\ ,\\
\operatorname{Exp}_{\rho,\sigma}\frac{\left\langle\chi_\tau,\chi_{V(\rho,\sigma)}\right\rangle_G}{\dim V(\rho,\sigma)} &\le \frac{d_\tau}{|G|}\sum_g\frac{1}{|[g]|^\ell} \quad \textit{if } k=0\ .
\end{aligned}$$

Proof. The two permutation representations R and C will play a special role in the analysis: in particular, we will see that the expectation of interest can be expressed in terms of these combinatorial characters. Specifically,

$$\operatorname{Exp}_{\rho,\sigma}\left[\frac{\left\langle\chi_\tau,\chi_{V(\rho,\sigma)}\right\rangle_G}{\dim V(\rho,\sigma)}\right] = \sum_{\rho\in\widehat{G}^k}\sum_{\sigma\in\widehat{G}^\ell}\left(\prod_i\frac{d_{\rho_i}^2}{|G|}\right)\left(\prod_j\frac{d_{\sigma_j}^2}{|G|}\right)\frac{\langle\chi_\tau,\chi_{V(\rho,\sigma)}\rangle_G}{\dim V(\rho,\sigma)} \tag{4.6}$$

$$= \frac{1}{|G|^{k+\ell}}\sum_{\rho\in\widehat{G}^k}\sum_{\sigma\in\widehat{G}^\ell}\left(\prod_i d_{\rho_i}\right)\langle\chi_\tau,\chi_{V(\rho,\sigma)}\rangle_G \tag{4.7}$$

$$= \frac{1}{|G|^{k+\ell}}\left\langle\chi_\tau,\sum_{\rho\in\widehat{G}^k}\sum_{\sigma\in\widehat{G}^\ell}\left(\prod_i d_{\rho_i}\chi_{\rho_i}\right)\left(\prod_j\chi_{\sigma_j}\chi_{\sigma_j}^*\right)\right\rangle_G \tag{4.8}$$

$$= \frac{1}{|G|^{k+\ell}}\left\langle\chi_\tau,\chi_{\mathrm{R}}^k\chi_{\mathrm{C}}^\ell\right\rangle_G \tag{4.9}$$

where the equality of line (4.7) follows from the fact that the dimension of $V(\rho,\sigma)$ is $\prod_i d_{\sigma_i}\cdot\prod_j d_{\sigma_j}^2$ and the equality of line (4.8) follows from the fact that the character of $V(\rho,\sigma)$ is $\prod_i\chi_{\rho_i}\prod_j\chi_{\sigma_j}\chi_{\sigma_j}^*$.

Recall that for any representation υ we have $\chi_\upsilon(1) = d_\upsilon$. As χ_{R} is a multiple of the delta function δ_g, whenever $k > 1$ we have $\langle \chi_\tau, \chi_{\mathrm{R}}^k \chi_{\mathrm{C}}^\ell \rangle_G = d_\tau |G|^{k+\ell-1}$ and

$$\mathrm{Exp}_{\rho,\sigma}\left[\frac{\left\langle \chi_\tau, \chi_{V(\rho,\sigma)} \right\rangle_G}{\dim V(\rho,\sigma)}\right] = \frac{d_\tau}{|G|} \ .$$

On the other hand, when $k = 0$ we have

$$\begin{aligned}
\mathrm{Exp}_{\rho,\sigma}\left[\frac{\left\langle \chi_\tau, \chi_{V(\rho,\sigma)} \right\rangle_G}{\dim V(\rho,\sigma)}\right] &= \frac{1}{|G|^\ell} \left\langle \chi_\tau, \chi_{\mathrm{C}}^\ell \right\rangle_G = \frac{1}{|G|^{\ell+1}} \sum_g \chi_\tau^*(g) \chi_{\mathrm{C}}^\ell(g) \\
&= \frac{1}{|G|^{\ell+1}} \sum_g \chi_\tau^*(g) \frac{|G|^\ell}{|[g]|^\ell} = \frac{1}{|G|} \sum_g \chi_\tau^*(g) \frac{1}{|[g]|^\ell} \\
&\le \frac{d_\tau}{|G|} \sum_g \frac{1}{|[g]|^\ell} \ ,
\end{aligned}$$

where the last inequality follows from the fact that $|\chi_\tau(g)| \le d_\tau$ for all g. □

Now note that the sum $\sum_g 1/\left|[g]\right|^\ell$ can also be written as a sum over the conjugacy classes C. In particular, if $\ell \ge 2$,

$$\sum_g \frac{1}{|[g]|^\ell} = \sum_C \frac{1}{|C|^{\ell-1}} \le \sum_C \frac{1}{|C|} \ .$$

In the case of the symmetric group S_n, the next lemma shows that this quantity is in fact $1 + o(1)$.

Lemma 13. *Given a partition $\lambda = (\lambda_1, \ldots, \lambda_t)$ with $\sum_i \lambda_i = n$ and $\lambda_i \ge \lambda_{i+1}$ for all i, let C_λ denote the conjugacy class of S_n consisting of permutations with cycle structure λ. Then*

$$\sum_\lambda \frac{1}{|C_\lambda|} = 1 + o(1) \ .$$

Proof. First note that if we group the λ_i into blocks consisting of τ_1 1s, τ_2 2s, and so on (such that $\sum_i \tau_i i = n$) then the size of the conjugacy class is given by

$$|C_\lambda| = \frac{n!}{\left(\prod_i \tau_i!\right)\left(\prod_i \lambda_i\right)}$$

since we can cyclically permute the elements of each cycle, and permute cycles of the same size with each other. Thus

$$\sum_\lambda \frac{1}{|C_\lambda|} = \sum_\lambda \frac{1}{n!} \left(\prod_i \tau_i!\right) \left(\prod_i \lambda_i\right) \ . \tag{4.10}$$

Now suppose that the conjugacy class consists of elements with support s, i.e., with $\tau_1 = n - s$ fixed points. Since we can specify such a partition with a partition of s objects, the number of such partitions is at most $p(s)$. Moreover we have

$$\prod_i \tau_i! = (n-s)! \prod_{i \ge 2} \tau_i! \le (n-s)!\,(s/2)!$$

and

$$\prod_i \lambda_i = \prod_{\lambda > 1} \lambda \le e^{s/e} \ .$$

since this is true for any set of reals $\lambda \geq 0$ such that $\sum \lambda = s$. Then (4.10) becomes

$$\sum_\lambda \frac{1}{|C_\lambda|} \leq \sum_{s=0}^{n} \frac{(n-s)!\,(s/2)!}{n!}\, p(s) e^{s/e} \tag{4.11}$$

Now, for $s > \sqrt{n}$, we have

$$\frac{(n-s)!\,(s/2)!}{n!} = \frac{1}{\binom{n}{s}} \frac{(s/2)!}{s!} \leq \frac{(s/2)!}{s!} < \left(\frac{2s}{e}\right)^{-s/2} \leq n^{-s/4} \tag{4.12}$$

and for $s \leq \sqrt{n}$, for sufficiently large n a stronger bound holds,

$$\frac{(n-s)!\,(s/2)!}{n!} \leq \frac{(\sqrt{n}/2)^{s/2}}{(n-s)^s} \leq n^{-3s/4} \;\; .$$

Thus (4.12) holds for all s. Using the absurdly crude bound $p(s)\, e^{s/e} < 4^s$, (4.11) then becomes

$$\sum_\lambda \frac{1}{|C_\lambda|} \leq \sum_{s=0}^{n} n^{-s/4} 4^s < \frac{1}{1-4n^{-1/4}} = 1 + O(n^{-1/4}) \;\; .$$

□

On the other hand, if $\ell = 1$ then the sum $\sum_g 1/\,|[g]|$ is simply the number of conjugacy classes. Therefore, in the case of the symmetric group, we have the following corollary of Lemma 12.

Corollary 14. *Let $G = S_n$ and k, ℓ, ρ, and σ as in Lemma 12. Then*

$$\mathrm{Exp}_{\rho,\sigma} \frac{\left\langle \chi_\tau, \chi_{V(\rho,\sigma)} \right\rangle_{S_n}}{\dim V(\rho,\sigma)} \leq (1 + o(1))\, \frac{d_\tau}{n!}$$

unless $\ell = 1$ and $k = 0$, in which case

$$\mathrm{Exp}_{\rho,\sigma} \frac{\left\langle \chi_\tau, \chi_{V(\rho,\sigma)} \right\rangle_{S_n}}{\dim V(\rho,\sigma)} \leq \frac{d_\tau p(n)}{n!} \;\; .$$

5 Two registers are insufficient for the symmetric group

In this section we show that no polynomial number of two-register experiments can distinguish the involutions we have been considering in S_n from each other or from the trivial subgroup. As in Section 4, we assume we have measured the representation name on each of the two registers, and that we observed the irreducible representations S^λ and S^μ. For simplicity we present the proof for von Neumann measurements; the generalization to arbitrary frames $\{\mathbf{b}\}$ proceeds exactly as in the proof of Theorem 8.

Theorem 15. *Let $B = \{\mathbf{b}\}$ be an orthonormal basis for $S^\lambda \otimes S^\mu$. Given the hidden subgroup $H = \{1, m\}$ where m is chosen uniformly at random from M, let $P_m(\mathbf{b})$ be the probability that we observe the vector $\mathbf{b}$ conditioned on having observed the representation names S^λ and S^μ, and let U be the uniform distribution on B. Then there is a constant $\delta > 0$ such that for sufficiently large n, with probability at least $1 - e^{-\delta\sqrt{n}/\log n}$ in m, λ and μ, we have*

$$\|P_m - U\|_1 < e^{-\delta\sqrt{n}/\log n} \;\; .$$

Proof. For $k=2$, Lemma 11 specializes to the following:

$$\mathrm{Exp}_m \left\|\Pi_m^{\otimes 2}\mathbf{b}\right\|^2 = \frac{1}{4}\left(1+\sum_{I\subseteq\{\lambda,\mu\}:I\neq\emptyset} E^I(\mathbf{b})\right) \tag{5.1}$$

$$\mathrm{Var}_m \left\|\Pi_m^{\otimes 2}\mathbf{b}\right\|^2 = \frac{1}{16}\sum_{I_1,I_2\subseteq\{\lambda,\mu\}:I_1,I_2\neq\emptyset} E^{I_1,I_2}(\mathbf{b}) - E^{I_1}(\mathbf{b})E^{I_2}(\mathbf{b})^* \ . \tag{5.2}$$

As before, S^λ and S^μ are chosen with the distribution $P(\rho)$. Since this is exponentially close to the Plancherel distribution [12], we can use Lemma 12 to calculate the expectations over λ and μ of $E^I(\mathbf{b})$ and $E^{I_1,I_2}(\mathbf{b})$ with negligible error. We will then show that $E^I(\mathbf{b})$ and $E^{I_1,I_2}(\mathbf{b})$ are superpolynomially small with the stated probability, for all but a small fraction of basis vectors $\mathbf{b}$, namely those that project into low-dimensional representations. As in Theorem 8, we will then use Markov's inequality to control the number of these basis vectors and use Chebyshev's inequality to control the rest, and thus bound the total distance $\|P_m - U\|_1$.

However, the analysis, at least when $|I_1|=|I_2|=2$, is more delicate than for the one-register case. As before, we exclude a set of low-dimensional representations Λ, but now we restrict Λ to Young diagrams with width or height extremely close to n. Specifically, let $c>0$ be a constant to be determined below, and let $\Lambda=\Lambda_c$ be the set of Young diagrams ν such that

$$\max(\nu_1,\nu_1') \geq n - c\sqrt{n}/\ln n \ .$$

Analogously to (3.8), Theorem 7 provides the following bound on the characters χ^ν for $\nu\notin\Lambda$,

$$\left|\frac{\chi^\nu(M)}{d^\nu}\right| \leq \left(1-\frac{c}{\sqrt{n}\ln n}\right)^{bn} < e^{-\alpha\sqrt{n}/\ln n} \tag{5.3}$$

where $\alpha = bc > 0$. The size and dimension of Λ is bounded by the following lemma.

Lemma 16. *$|\Lambda| = e^{o(\sqrt{n})}$ and $d^\nu < e^{c\sqrt{n}}$ for any $\nu\in\Lambda$. Therefore, $\sum_{\nu\in\Lambda}(d^\nu)^2 < e^{2c\sqrt{n}+o(\sqrt{n})}$.*

Proof. The proof of Lemma 9 applies, except now $|\Lambda| < 2xp(x)$ where $x = c\sqrt{n}/\ln n$. □

Then the next lemma shows that with high probability in λ and μ, $E^I(\mathbf{b})$ is superpolynomially small for *all* $\mathbf{b}\in B$. (Indeed, it is exponentially small for all but a few $\mathbf{b}$, but we give this statement for simplicity.)

Lemma 17. *Let S^μ and S^λ be distributed according to the Plancherel distribution in $\widehat{S_n}$. Let $I\subseteq\{\lambda,\mu\}$, $I\neq\emptyset$. There is a constant $\gamma>0$ such that for sufficiently large n, with probability $1-e^{-\gamma\sqrt{n}}$, $\left|E^I(\mathbf{b})\right| \leq e^{-\alpha\sqrt{n}/\ln n}$ for all $\mathbf{b}\in B$.*

Proof. The case when $|I|=1$ is identical to the one-register case, since then $E^I(\mathbf{b}) = \chi^\lambda(M)/d^\lambda$. Lemma 6 implies $\lambda\notin\Lambda$ with probability $1-e^{-\delta\sqrt{n}}$, and (5.3) completes the proof of this case.

For the case $|I|=2$, it suffices to ensure that $S^\lambda\otimes S^\mu$ contains no low-dimensional representations. Let $\nu\in\Lambda$; then by Lemma 12 and Lemma 16, the expected multiplicity of S^ν in $S^\lambda\otimes S^\mu$ is

$$\mathrm{Exp}_{\lambda,\mu}\langle\chi^\nu,\chi^\lambda\chi^\mu\rangle_{S_n} = \mathrm{Exp}_{\lambda,\mu}\, d^\lambda d^\mu \frac{\langle\chi^\nu,\chi^\lambda\chi^\mu\rangle_{S_n}}{d^\lambda d^\mu} \leq e^{-2\hat{c}\sqrt{n}} d^\nu \leq e^{(c-2\hat{c})\sqrt{n}}$$

where $\hat{c}$ is the constant appearing in Theorem 5. Thus if $c<\hat{c}$, Lemma 16 and Markov's inequality imply that the probability *any* S^ν with $\nu\in\Lambda$ appears in $S^\lambda\otimes S^\mu$ is at most $e^{-\hat{c}\sqrt{n}}$. If none do, then (5.3) and the fact that $\left|E^I\right| \leq \max_{\nu\notin\Lambda}|\chi^\nu(M)/d^\nu|$ complete the proof with $\gamma=\hat{c}$. □

For the variance estimates, for each $S^\lambda, S^\mu \in \widehat{S_n}$ and $I_1, I_2 \subset \{\lambda,\mu\}$, recall Definition 2 and let

$$V[I_1,I_2] = V(I_1 \triangle I_2, I_1\cap I_2) \ ,$$

where $I_1 \triangle I_2$ is the symmetric difference. (We abuse notation here, allowing, e.g., the set $I_1 \triangle I_2$ to stand for the tuple of representations S^λ with $\lambda \in I_1 \triangle I_2$.) For the variance calculation, as in the single-register case, let $L[I_1, I_2] \subset V(I_1, I_2)$ be the subspace consisting of copies of representations S^ν with $\nu \in \Lambda$, and let $\Pi_{L[I_1,I_2]}$ be the projection operator onto this subspace. We will abbreviate $L = L[I_1, I_2]$ and $V = V[I_1, I_2]$ when the parameters are clear from context. Then the following lemma bounds the dimension of this subspace.

Lemma 18. *Let S^λ and S^μ be distributed according to the Plancherel distribution in $\widehat{S_n}$ and let $I_1, I_2 \subseteq \{\lambda, \mu\}$, $I_1, I_2 \neq \emptyset$. There is a constant $\beta > 0$ such that for sufficiently large n, with probability at least $1 - e^{-\beta\sqrt{n}}$, $\dim L[I_1, I_2] \leq e^{-\beta\sqrt{n}}|B|$.*

Proof. If $|I_1| = |I_2| = 1$ and $I_1 \neq I_2$, then the proof of the previous lemma shows that L is in fact empty with probability $1 - e^{-\Omega(\sqrt{n})}$. When $I_1 = I_2$ and $|I_1| = 1$, however, this is not true; for instance, for any λ, $S^\lambda \otimes (S^\lambda)^*$ contains exactly one copy of the trivial representation. However, since $\dim V = (d^\lambda)^2$ and $|B| = d^\lambda d^\mu$, Corollary 14 gives

$$\begin{aligned}
\operatorname{Exp}_{\lambda,\mu} \frac{\dim L}{|B|} &= \sum_{\nu \in \Lambda} d^\nu \operatorname{Exp}_{\lambda,\mu} \frac{d^\lambda}{d^\mu} \frac{\langle \chi^\nu, (\chi^\lambda)^2 \rangle_{S_n}}{\dim V} \\
&\leq \frac{e^{(\delta - \hat{c})\sqrt{n}}\, p(n)}{n!} \sum_{\nu \in \Lambda} (d^\nu)^2 \leq \frac{e^{(2c + 2\delta - \hat{c})\sqrt{n} + o(\sqrt{n})}}{n!} = n^{-\Omega(n)}
\end{aligned}$$

where we assume that the event of Lemma 6 occurs and δ is the constant defined there.

When $|I_1| = 2$ and $|I_2| = 1$, e.g. $I_1 = \{\lambda, \mu\}$ and $I_2 = \{\lambda\}$, then $\dim V = (d^\lambda)^2 d^\mu$ and Lemma 12 yields

$$\begin{aligned}
\operatorname{Exp}_{\lambda,\mu} \frac{\dim L}{|B|} &= \sum_{\nu \in \Lambda} d^\nu \operatorname{Exp}_{\lambda,\mu} d^\lambda \frac{\langle \chi^\nu, (\chi^\lambda)^2 \rangle_{S_n}}{\dim V} \\
&\leq \frac{e^{-\hat{c}\sqrt{n}}}{\sqrt{n!}} \sum_{\nu \in \Lambda} (d^\nu)^2 \leq \frac{e^{(2c - \hat{c})\sqrt{n} + o(\sqrt{n})}}{\sqrt{n!}} = n^{-\Omega(n)} \quad .
\end{aligned}$$

The case when $|I_2| = 2$ and $|I_1| = 1$ is identical.

Finally, the case when $|I_1| = |I_2| = 2$ is the most delicate. Now $\dim V = (d^\lambda)^2 (d^\mu)^2$, and Corollary 14 gives

$$\begin{aligned}
\operatorname{Exp}_{\lambda,\mu} \frac{\dim L}{|B|} &= \sum_{\nu \in \Lambda} d^\nu \operatorname{Exp}_{\lambda,\mu} d^\lambda d^\mu \frac{\langle \chi^\nu, (\chi^\lambda)^2 \rangle_{S_n}}{\dim V} \\
&\leq (1 + o(1)) e^{-2\hat{c}\sqrt{n}} \sum_{\nu \in \Lambda} (d^\nu)^2 \leq e^{(2c - 2\hat{c})\sqrt{n} + o(\sqrt{n})} < e^{-\hat{c}\sqrt{n}} \quad .
\end{aligned}$$

if we set $c < \hat{c}/2$.

Thus with probability at least $1 - e^{-\delta\sqrt{n}}$ we have $\operatorname{Exp}[\dim L/|B|] = e^{-\hat{c}\sqrt{n}}$. By Markov's inequality, the probability that $\dim L > e^{-(\hat{c}/2)\sqrt{n}}|B|$ is at most $e^{-(\hat{c}/2)\sqrt{n}}$. Thus setting $\beta < \min(\delta, \hat{c}/2)$ completes the proof. $\square$

Now, let E_0 denote the following event:

1. $\max\left(|\chi^\lambda(M)/d^\lambda|, |\chi^\mu(M)/d^\mu|\right) \leq e^{-\Omega(n)}$,

2. $\left|E^I(\mathbf{b})\right| = e^{-\alpha\sqrt{n}/\ln n}$ for all $\mathbf{b} \in B$ and all $I \subset \{\lambda, \mu\}$, and

3. $\dim L[I_1, I_2] \leq e^{-\beta\sqrt{n}}|B|$ for each $I_1, I_2 \subset \{\lambda, \mu\}$ with $I_1 \neq \emptyset$ and $I_2 \neq \emptyset$.

As a consequence of (3.8) and Lemmas 10, 17 and 18, E_0 occurs with probability $1-e^{-\Omega(\sqrt{n})}$. In what follows we condition on E_0. This will allow us to control the three principal parameters that determine the total variation distance between P_m and the uniform distribution: $\mathbf{rk}\ \Pi_m^{\otimes 2}$, $\mathrm{Exp}_m[\Pi_m^{\otimes 2}(\mathbf{b})]$, and $\mathrm{Var}_m[\Pi_m^{\otimes 2}(\mathbf{b})]$.

Considering $\mathbf{rk}\ \Pi_m^{\otimes 2}$, note that the rank of $\Pi_m^{\otimes 2}$ restricted to a representation $S^\lambda \otimes S^\mu$ is the product of the ranks of Π_m restricted to S^λ and S^μ; then (3.7), and item 1 of E_0 give

$$\mathbf{rk}\ \Pi_m^{\otimes 2} = \frac{d^\mu d^\lambda}{4}\left(1+\frac{\chi^\mu(M)}{d^\mu}\right)\left(1+\frac{\chi^\lambda(M)}{d^\lambda}\right) = \frac{|B|}{4}\left(1+e^{-\Omega(n)}\right) \quad . \tag{5.4}$$

As for the expectation $\mathrm{Exp}_m[\Pi_m^{\otimes 2}(\mathbf{b})]$, in light of (5.1) and item 2 of E_0, we conclude that for each $\mathbf{b}\in B$,

$$\left|\mathrm{Exp}_m \left\|\Pi_m^{\otimes 2}\mathbf{b}\right\|^2 - \frac{1}{4}\right| \leq 3e^{-\alpha\sqrt{n}/\ln n} \quad . \tag{5.5}$$

Finally, we focus on the variance. Define $B_L \subset B$ to be the set of basis vectors $\mathbf{b}$ such that for some nontrivial $I_1, I_2 \subset \{\lambda,\mu\}$, $\left\|\Pi_{L[I_1,I_2]}\mathbf{b}\right\|^2 \geq e^{-(\beta/2)\sqrt{n}}$. Then since item 3 of E_0 holds for each of the $3^2 = 9$ pairs of nonempty subsets I_1, I_2, we have

$$|B_L| \leq e^{(\beta/2)\sqrt{n}} \sum_{I_1,I_2} \dim L[I_1,I_2] \leq 9e^{-(\beta/2)\sqrt{n}}|B| \quad .$$

Observe that for any $\mathbf{b}\in B\setminus B_L$, Equations (5.2), (5.3), and item 2 of E_0 give

$$\mathrm{Var}_m \left\|\Pi_m^{\otimes 2}\mathbf{b}\right\|^2 \leq \frac{9}{16}\left(e^{-\alpha\sqrt{n}/\ln n} + e^{-2\alpha\sqrt{n}/\ln n} + e^{-(\beta/2)\sqrt{n}}\right) < e^{-\alpha\sqrt{n}/\ln n} \quad . \tag{5.6}$$

Then Chebyshev's inequality gives

$$\Pr\left[\left|\left\|\Pi_m^{\otimes 2}\mathbf{b}\right\|^2 - \mathrm{Exp}_m\left\|\Pi_m^{\otimes 2}\mathbf{b}\right\|^2\right| \geq e^{-(\alpha/3)\sqrt{n}/\ln n}\right] \leq e^{-(\alpha/3)\sqrt{n}/\ln n} \quad . \tag{5.7}$$

Analogous to Theorem 8, let $B_{\mathrm{bad}} \subset B\setminus B_L$ denote the subset of basis vectors for which the event of (5.7) is violated. (As in the one-register case, while B_L depends only on the choice of λ and μ, B_{bad} depends also on m.) Let E_1 denote the event

$$|B_{\mathrm{bad}}| < e^{-(\alpha/6)\sqrt{n}/\ln n}|B| \quad .$$

Then (5.7) and Markov's inequality imply that E_1 occurs with probability $1 - e^{-(\alpha/6)\sqrt{n}/\ln n}$.

So, finally, recall that $P_m(\mathbf{b}) = \left\|\Pi_m^{\otimes 2}(\mathbf{b})\right\|^2/\mathbf{rk}\ \Pi_m^{\otimes 2}$ and let $\overline{P}(\mathbf{b})$ denote the distribution $\overline{P}(\mathbf{b}) = \mathrm{Exp}_m[P_m(\mathbf{b})]$. We separate $\left\|P_m - \overline{P}\right\|_1$ into contributions from basis vectors outside and inside $B_L\cup B_{\mathrm{bad}}$:

$$\left\|P_m - \overline{P}\right\|_1 = \sum_{\mathbf{b}\notin B_L\cup B_{\mathrm{bad}}} \left|P_m(\mathbf{b}) - \overline{P}(\mathbf{b})\right| + \sum_{\mathbf{b}\in B_L\cup B_{\mathrm{bad}}} \left|P_m(\mathbf{b}) - \overline{P}(\mathbf{b})\right| \quad . \tag{5.8}$$

The first sum is taken only over vectors $\mathbf{b}$ for which

$$\left|\left\|\Pi_m^{\otimes 2}(\mathbf{b})\right\|^2 - \mathrm{Exp}_m\left\|\Pi_m^{\otimes 2}(\mathbf{b})\right\|^2\right| < e^{-(\alpha/3)\sqrt{n}/\ln n} \quad .$$

Then conditioning on E_0 and E_1, the rank estimate of (5.4) implies that

$$\sum_{\mathbf{b}\notin B_L\cup B_{\mathrm{bad}}} \left|P_m(\mathbf{b}) - \overline{P}(\mathbf{b})\right| \leq \frac{e^{-(\alpha/3)\sqrt{n}/\ln n}}{\mathbf{rk}\ \Pi_m^{\otimes 2}}\cdot|B| = \frac{4e^{-(\alpha/3)\sqrt{n}/\ln n}}{1+e^{-\Omega(n)}} < 8e^{-(\alpha/3)\sqrt{n}/\ln n} \quad . \tag{5.9}$$

On the other hand, conditioning on E_0 and E_1 we have

$$|B_L\cup B_{\mathrm{bad}}| \leq \left(9e^{-(\beta/2)\sqrt{n}} + e^{-(\alpha/6)\sqrt{n}/\ln n}\right)|B| < 2e^{-(\alpha/6)\sqrt{n}/\ln n}|B| \quad ,$$

and then (5.4) and (5.5) imply that the total expected probability of the basis vectors in $B_L \cup B_{\text{bad}}$ is

$$\sum_{\mathbf{b}\in B_L\cup B_{\text{bad}}} \overline{P}(\mathbf{b}) = \sum_{\mathbf{b}\in B_L\cup B_{\text{bad}}} \frac{\operatorname{Exp}_m \left\|\Pi_m^{\otimes 2}(\mathbf{b})\right\|^2}{\mathbf{rk}\,\Pi_m^{\otimes 2}} \leq \frac{|B_L \cup B_{\text{bad}}|}{\mathbf{rk}\,\Pi_m^{\otimes 2}} \cdot \left(\frac{1}{4} + 3e^{-\alpha\sqrt{n}/\ln n}\right)$$
$$\leq 2e^{-(\alpha/6)\sqrt{n}/\ln n}(1+o(1)) < 3e^{-(\alpha/6)\sqrt{n}/\ln n} \ . \tag{5.10}$$

Then we must have

$$\sum_{\mathbf{b}\notin B_L\cup B_{\text{bad}}} \overline{P}(\mathbf{b}) > 1 - 3e^{-(\alpha/6)\sqrt{n}/\ln n}$$

and hence, by (5.9),

$$\sum_{\mathbf{b}\notin B_L\cup B_{\text{bad}}} P_m(\mathbf{b}) > 1 - 3e^{-(\alpha/6)\sqrt{n}/\ln n} - 8e^{-(\alpha/3)\sqrt{n}/\ln n} > 1 - 4e^{-(\alpha/6)\sqrt{n}/\ln n}$$

and so

$$\sum_{\mathbf{b}\in B_L\cup B_{\text{bad}}} P_m(\mathbf{b}) < 4e^{-(\alpha/6)\sqrt{n}/\ln n} \ .$$

Combining this with (5.10) bounds the second sum in (5.8),

$$\sum_{\mathbf{b}\in B_L\cup B_{\text{bad}}} \left|P_m(\mathbf{b}) - \overline{P}(\mathbf{b})\right| < 7e^{-(\alpha/6)\sqrt{n}/\ln n} \ . \tag{5.11}$$

Then combining (5.8), (5.9) and (5.11),

$$\left\|P_m - \overline{P}\right\|_1 < 8e^{-(\alpha/6)\sqrt{n}/\ln n}$$

with probability at least $\Pr[E_0 \wedge E_1] \geq 1 - e^{-\Omega(\sqrt{n})} - e^{-(\alpha/6)\sqrt{n}/\ln n} > 1 - 2e^{-(\alpha/6)\sqrt{n}/\ln n}$.

Finally, it remains to be proved that $\overline{P}$ is, with high probability, close to the uniform distribution U on B. But this follows from (5.4) and (5.5); conditioning on E_0, we have

$$\left\|\overline{P} - U\right\|_1 \leq \sum_{\mathbf{b}\in B} \left|\frac{\operatorname{Exp}_m \left\|\Pi_m^{\otimes 2}\mathbf{b}\right\|^2}{\mathbf{rk}\,\Pi_m^{\otimes 2}} - \frac{1}{|B|}\right| < 12e^{-\alpha\sqrt{n}/\ln n}(1 + e^{-\Omega(n)}) \ .$$

We complete the proof by setting $\delta < \alpha/6$ and invoking the triangle inequality. □

6 Conclusion

The reader will notice that our current machinery cannot extend to three or more registers when applied to the symmetric group, as the representations of S_n have typical dimension equal to $(n!)^{1/2-o(1)}$. However, we have been very pessimistic in our analysis; in particular, we have assumed that vectors of the form $\mathbf{b}\otimes\mathbf{b}$ project into low-dimensional representations, S^ν with $\nu \in \Lambda$, as much as possible. Perhaps a more detailed understanding of how these vectors lie inside the decomposition of $V(\rho,\sigma)$ into irreducibles would allow one to prove that this hidden subgroup problem requires entangled measurements over $\Omega(\log|G|) = \Omega(n\log n)$ coset states. Therefore, we make the following conjecture.

Conjecture 1. *Let $B = \{\mathbf{b}\}$ with weights $\{a_\mathbf{b}\}$ be a complete frame for $S^{\lambda_1} \otimes \cdots \otimes S^{\lambda_k}$. Given the hidden subgroup $H = \{1, m\}$ where m is chosen uniformly at random from M, and a coset state $|c_1 H\rangle \otimes \cdots \otimes |c_k H\rangle$ on k registers, let $P_m(\mathbf{b})$ be the probability that we observe the vector $\mathbf{b}$ conditioned on having observed the representation names $\{S^{\lambda_i}\}$, and let U be the natural distribution on B. Then for all $c > 0$, with probability $1 - o(n^{-c})$ in m and $\{S^{\lambda_i}\}$, we have*

$$\|P_m - U\|_1 = o(n^{-c})$$

unless $k = \Omega(n\log n)$.

Acknowledgments.

This work was supported by NSF grants CCR-0093065, PHY-0200909, EIA-0218443, EIA-0218563, CCR-0220070, and CCR-0220264. We are grateful to Denis Thérien, McGill University, and Bellairs Research Institute for organizing a workshop at which this work began; to Dorit Aharonov, Daniel Rockmore, Leonard Schulman, and Umesh Vazirani for helpful conversations; and to Tracy Conrad and Sally Milius for their support and tolerance. C.M. also thanks She Who Is Not Yet Named for her impending arrival, and for providing a larger perspective.

References

[1] David Bacon, Andrew Childs, and Wim van Dam. Optimal measurements for the dihedral hidden subgroup problem. Preprint, quant-ph/0501044 (2005).

[2] Robert Beals. Quantum computation of Fourier transforms over symmetric groups. *Proc. 29th Annual ACM Symposium on the Theory of Computing*, pages 48–53, 1997.

[3] Ethan Bernstein and Umesh Vazirani. Quantum complexity theory (preliminary abstract). *Proc. 25th Annual ACM Symposium on the Theory of Computing*, pages 11–20, 1993.

[4] Wim van Dam, Sean Hallgren, and Lawrence Ip. Quantum algorithms for some hidden shift problems. *Proc. 14th ACM-SIAM Symposium on Discrete Algorithms*, pages 489–498, 2003.

[5] Mark Ettinger and Peter Høyer. On quantum algorithms for noncommutative hidden subgroups. Preprint, quant-ph/9807029 (1998).

[6] Mark Ettinger and Peter Høyer and Emmanuel Knill. The quantum query complexity of the hidden subgroup problem is polynomial. *Information Processing Letters*, to appear.

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

[8] William Fulton and Joe Harris. *Representation Theory: A First Course*. Number 129 in Graduate Texts in Mathematics. Springer-Verlag, 1991.

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

[10] Lisa Hales and Sean Hallgren. Quantum Fourier sampling simplified. *Proc. 31st Annual ACM Symposium on Theory of Computing*, 1999.

[11] Lisa Hales and Sean Hallgren. An improved quantum Fourier transform algorithm and applications. *Proc. 41st Annual Symposium on Foundations of Computer Science*, 2000.

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

[13] Peter Høyer. Efficient quantum transforms. Preprint, quant-ph/9702028 (1997).

[14] Yoshifumi Inui and François Le Gall. An efficient algorithm for the hidden subgroup problem over a class of semi-direct product groups. Proc. EQIS 2004.

[15] Lawrence Ip. Shor's algorithm is optimal. Preprint, 2004.

[16] Gábor Ivanyos, Frédéric Magniez, and Miklos Santha. Efficient quantum algorithms for some instances of the non-abelian hidden subgroup problem. *Int. J. Found. Comput. Sci.* 14(5): 723–740, 2003.

[17] Richard Jozsa. Quantum factoring, discrete logarithms and the hidden subgroup problem. Preprint, quant-ph/0012084 (2000).

[18] Julia Kempe and Aner Shalev. The hidden subgroup problem and permutation group theory. Preprint, quant-ph/0406046 (2004).

[19] S. V. Kerov. *Asymptotic representation theory of the symmetric group and its applications in analysis.* Translated by N. V. Tsilevich. Volume 219 in Translations of Mathematical Monographs. American Mathematical Society, 2003.

[20] Greg Kuperberg. A subexponential-time quantum algorithm for the dihedral hidden subgroup problem. Preprint, quant-ph/0302112 (2003).

[21] Cristopher Moore, Daniel Rockmore, and Alexander Russell. Generic quantum Fourier transforms. *Proc. 15th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 778–787, 2004.

[22] Cristopher Moore, Daniel Rockmore, Alexander Russell, and Leonard Schulman. The value of basis selection in Fourier sampling: hidden subgroup problems for affine groups. *Proc. 15th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1113–1122, 2004.

[23] Oded Regev. Quantum computation and lattice problems. *Proc. 43rd Symposium on Foundations of Computer Science*, pages 520–530, 2002.

[24] Martin Roetteler and Thomas Beth. Polynomial-time solution to the hidden subgroup problem for a class of non-abelian groups. Preprint, quant-ph/9812070 (1998).

[25] Yuval Roichman. Upper bound on the characters of the symmetric groups. *Inventiones Mathematicae*, 125:451–485, 1996.

[26] Steven Roman. *Advanced Linear Algebra*. Number 135 in Graduate Texts in Mathematics. Springer, 1992.

[27] Jean-Pierre Serre. *Linear Representations of Finite Groups*. Number 42 in Graduate Texts in Mathematics. Springer-Verlag, 1977.

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

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

[30] A. M. Vershik and S. V. Kerov. Asymptotic behavior of the maximum and generic dimensions of irreducible representations of the symmetric group. Funk. Anal. i Prolizhen, 19(1):25–36, 1985; English translation, Funct. Anal. Appl., 19:21–31, 1989.