

REPORT DOCUMENTATION PAGE			Form Approved OMB NO. 0704-0188		
<p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA, 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p> <p>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</p>					
1. REPORT DATE (DD-MM-YYYY) 02-05-2012		2. REPORT TYPE Conference Proceeding		3. DATES COVERED (From - To) -	
4. TITLE AND SUBTITLE On the impossibility of a quantum sieve algorithm for graph isomorphism			5a. CONTRACT NUMBER W911NF-05-1-0347		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER ODXXX1		
6. AUTHORS Cristopher Moore, Alexander Russell, Piotr Sniady			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAMES AND ADDRESSES University of New Mexico Albuquerque Regents of The University of New Mexico MSC01 1247 Albuquerque, NM 87131 -0001			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211			10. SPONSOR/MONITOR'S ACRONYM(S) ARO		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S) 47976-PH-QC.2		
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.					
13. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.					
14. ABSTRACT It is known that any quantum algorithm for Graph Isomorphism that works within the framework of the hidden subgroup problem (HSP) must perform highly entangled measurements across $2^{(n \log n)}$ coset states. One of the only known models for how such a measurement could be carried out efficiently is Kuperberg's algorithm for the HSP in the dihedral group, in which quantum states are adaptively combined and measured according to the decomposition of tensor products into irreducible representations. This "quantum sieve" starts with coset states,					
15. SUBJECT TERMS quantum algorithms, graph isomorphism					
16. SECURITY CLASSIFICATION OF:		17. LIMITATION OF ABSTRACT UU	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON Cristopher Moore	
a. REPORT UU	b. ABSTRACT UU			c. THIS PAGE UU	19b. TELEPHONE NUMBER 505-277-3112

Report Title

On the impossibility of a quantum sieve algorithm for graph isomorphism

ABSTRACT

It is known that any quantum algorithm for Graph Isomorphism that works within the framework of the hidden subgroup problem (HSP) must perform highly entangled measurements across $\Theta(n \log n)$ coset states. One of the only known models for how such a measurement could be carried out efficiently is Kuperberg's algorithm for the HSP in the dihedral group, in which quantum states are adaptively combined and measured according to the decomposition of tensor products into irreducible representations. This "quantum sieve" starts with coset states, and works its way down towards representations whose probabilities differ depending on, for example, whether the hidden subgroup is trivial or nontrivial.

In this paper we show that no such approach can produce a polynomial-time quantum algorithm for Graph Isomorphism. Specifically, we consider the natural reduction of Graph Isomorphism to the HSP over the wreath product $S_n \wr \mathbb{Z}_2$. Using a recently proved bound on the irreducible characters of S_n , we show that no algorithm in this family can solve Graph Isomorphism in less than $e^{\Theta(n)}$ time, no matter what adaptive rule it uses to select and combine quantum states. In particular, algorithms of this type can offer essentially no improvement over the best known classical algorithms, which run in time $e^{O(n \log n)}$.

Conference Name: the thirty-ninth annual ACM symposium

Conference Date: June 10, 2007

ON THE IMPOSSIBILITY OF A QUANTUM SIEVE ALGORITHM FOR GRAPH ISOMORPHISM

CRISTOPHER MOORE, ALEXANDER RUSSELL, AND PIOTR ŚNIADY

ABSTRACT. It is known that any quantum algorithm for Graph Isomorphism that works within the framework of the hidden subgroup problem (HSP) must perform highly entangled measurements across $\Omega(n \log n)$ coset states. One of the only known models for how such a measurement could be carried out efficiently is Kuperberg’s algorithm for the HSP in the dihedral group, in which quantum states are adaptively combined and measured according to the decomposition of tensor products into irreducible representations. This “quantum sieve” starts with coset states, and works its way down towards representations whose probabilities differ depending on, for example, whether the hidden subgroup is trivial or nontrivial.

In this paper we show that no such approach can produce a polynomial-time quantum algorithm for Graph Isomorphism. Specifically, we consider the natural reduction of Graph Isomorphism to the HSP over the wreath product $S_n \wr \mathbb{Z}_2$. Using a recently proved bound on the irreducible characters of S_n , we show that no algorithm in this family can solve Graph Isomorphism in less than $e^{\Omega(\sqrt{n})}$ time, no matter what adaptive rule it uses to select and combine quantum states. In particular, algorithms of this type can offer essentially no improvement over the best known classical algorithms, which run in time $e^{O(\sqrt{n \log n})}$.

1. INTRODUCTION

The discovery of Shor’s and Simon’s algorithms began a frenzied charge to uncover the full algorithmic potential of a general purpose quantum computer. Creative invocations of the order-finding primitive yielded efficient quantum algorithms for a number of other number-theoretic problems [Hal02, Hal05]. As the field matured, these algorithms were roughly unified under the general framework of the *hidden subgroup problem*, where one must determine a subgroup H of a group G by querying an oracle $f : G \rightarrow S$ known to have the property that $f(g) = f(gh) \Leftrightarrow h \in H$. Solutions to this general problem are the foundation for almost all known superpolynomial speedups offered by quantum algorithms over their classical counterparts (see [AJL06] for an important exception).

The algorithms of Simon and Shor essentially solve the hidden subgroup problem on abelian groups, namely \mathbb{Z}_2^n and \mathbb{Z}_n^* respectively. Since then,

non-abelian hidden subgroup problems have received a great deal of attention (e.g. [HRTS00, GSVV01, FIM⁺03, MRS04, BCvD05, HMR⁺06]). A major motivation for this work is the fact that we can reduce Graph Isomorphism for rigid graphs of size n to the case of the hidden subgroup problem over the symmetric group S_{2n} , or more specifically the wreath product $S_n \wr \mathbb{Z}_2$, where the hidden subgroup is promised to be either trivial or of order two. The standard approach to these problems is to prepare “coset states” of the form

$$\rho_H = \frac{1}{|G|} \sum_c |cH\rangle \langle cH| \quad ,$$

where $|S\rangle$, for a subset $S \subset G$, denotes the uniform superposition $(1/\sqrt{|S|}) \sum_{g \in S} |g\rangle$. In the abelian case, one proceeds by computing the quantum Fourier transform of such coset states, measuring the resulting states, and appropriately interpreting the results. In the case of the symmetric group, however, determining H from a quantum measurement of coset states is far more difficult. In particular, no *product measurement* (that is, a measurement which treats each coset state independently) can efficiently determine a hidden subgroup over S_n [MRS05]; in fact, any successful measurement must be *entangled* over $\Omega(n \log n)$ coset states at once [HMR⁺06].

One of the few proposed methods for building such an entangled measurement comes from Kuperberg’s algorithm for the hidden subgroup problem in the dihedral group [Kup05]. It starts by generating a large number of coset states and subjecting each one to *weak Fourier sampling*, so that it lies inside a known irreducible representation. It then proceeds with an adaptive “sieve” process, at each step of which it judiciously selects pairs of states and measures them in a basis consistent with the *Clebsch-Gordan* decomposition of their tensor product into irreducible representations. This sieve continues until we obtain a state lying in an “informative” representation: namely, one from which information about the hidden subgroup can be easily extracted. We can visualize a run of the sieve as a forest, where leaves consist of the initial coset states, each internal node measures the tensor product of its parents, and the informative representations lie at the roots.

This approach is especially attractive in cases like Graph Isomorphism, where all we need to know is whether the hidden subgroup is trivial or nontrivial. Specifically, suppose that the hidden subgroup H is promised to be either the trivial subgroup $\{1\}$ or a conjugate of a known subgroup H_0 . Assume further that there is an irreducible representation σ of G with the property that $\sum_{h \in H_0} \sigma(h) = 0$; that is, a “missing harmonic” in the sense of [MR05a]. In this case, if H is nontrivial then the probability of observing σ under weak Fourier sampling of the coset state ρ_H is zero.

More generally, as we discuss below, the irrep σ cannot appear at any time in the sieve. If, on the other hand, one can guarantee that the sieve *does* observe σ with significant probability when the hidden subgroup is trivial and the corresponding states are completely mixed, it gives us an algorithm to distinguish the two cases.

For example, if we consider the case of the hidden subgroup problem in the dihedral group D_n where H is either trivial or a conjugate of $H_0 = \{1, m\}$ where m is an involution, then the sign representation π is a missing harmonic. Applying Kuperberg's sieve, we observe π with significant probability after $e^{O(\sqrt{n})}$ steps if H is trivial, while we can never observe it if H is of order 2. A similar approach was applied to groups of the form G^n by Alagić et al. [AMR06].

We show here, however, that the hidden subgroup problem related to Graph Isomorphism cannot be solved efficiently by any algorithm in this family. Specifically, no matter what adaptive selection rule it uses to choose pairs of states to combine and measure, such a sieve cannot distinguish the isomorphic and nonisomorphic cases unless it takes $e^{\Omega(\sqrt{n})}$ time (and uses this many coset states). In comparison, the best known classical algorithms for Graph Isomorphism run in time $e^{O(\sqrt{n} \log n)}$ for general graphs [Bab80, Bab83] and $e^{O(n^{1/3} \log^2 n)}$ for strongly regular graphs [Spi96]. Therefore, quantum algorithms of this kind can offer no meaningful improvement over their classical counterparts.

Our proof relies on several ingredients. First, we give a formal definition of quantum sieve algorithms, and we derive a combinatorial description of the probability distributions of their observations in the trivial and nontrivial cases. We then focus on the case where the ambient group is a wreath product $G \wr \mathbb{Z}_2$, and show that no information is gained until the sieve observes a so-called inhomogeneous representation. Then, in the case where $G = S_n$, we rely on a bound on the characters of the symmetric group proved very recently by Rattan and Śniady [RS06] to show that the total variation distance between the trivial and nontrivial cases is at most $e^{-b\sqrt{n}}$ unless the sieve takes $e^{a\sqrt{n}}$ time, for some constants $a, b > 0$.

We note that two of the present authors gave this result in conditional form in [MR06], in which they presented a conjectured bound on the characters of S_n . Indeed, it was this conjecture which inspired the work of [RS06] who proved its weaker version, which, along with some additional arguments, allows us to prove the results of [MR06] unconditionally.

We refer the reader to [Ser77, JK81] for an introduction to the representation theory of finite groups, and in particular of the symmetric group S_n . One fact which we use repeatedly is that the τ -isotypic subspace, i.e., the subspace of a representation σ which consists of copies of an irrep τ , is

the image of the projection operator

$$\Pi_\tau = \frac{1}{|G|} \sum_{g \in G} d_\tau \chi_\tau(g)^* \sigma(g) .$$

These projection operators can be combined to create a measurement whose outcomes are names of irreducible representations. Applying such a measurement to coset states is known as weak Fourier sampling; we use the term *isotypic sampling* to refer to the more general case of applying an arbitrary group action to a multiregister state.

2. FOURIER ANALYSIS ON FINITE GROUPS

In this section we review the representation theory of finite groups. Our treatment is primarily for the purposes of setting down notation; we refer the reader to [Ser77] for a complete account. Let G be a finite group. A *representation* σ of G is a homomorphism $\sigma : G \rightarrow \mathbf{U}(V)$, where V is a finite-dimensional Hilbert space and $\mathbf{U}(V)$ is the group of unitary operators on V . The *dimension* of σ , denoted d_σ , is the dimension of the vector space V . Fixing a representation $\sigma : G \rightarrow \mathbf{U}(V)$, we say that a subspace $W \subset V$ is *invariant* if $\sigma(g) \cdot W = W$ for all $g \in G$. When σ has no invariant subspaces other than the trivial subspace $\{0\}$ and V itself, σ is said to be *irreducible*.

If two representations σ and σ' are the same up to a unitary change of basis, we say that they are *equivalent*. It is a fact that any finite group G has a finite number of distinct irreducible representations up to equivalence and, for a group G , we let \widehat{G} denote a set of representations containing exactly one from each equivalence class. We often say that each $\sigma \in \widehat{G}$ is the *name* of an irreducible representation, or an *irrep* for short.

The irreps of G give rise to the Fourier transform. Specifically, for a function $f : G \rightarrow \mathbb{C}$ and an element $\sigma \in \widehat{G}$, define the *Fourier transform of f at σ* to be

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

The leading coefficients are chosen to make the transform unitary, so that it preserves inner products:

$$\langle f_1, f_2 \rangle = \sum_g f_1^*(g) f_2(g) = \sum_{\sigma \in \widehat{G}} \mathbf{tr} \left(\hat{f}_1(\sigma)^\dagger \cdot \hat{f}_2(\sigma) \right) .$$

If σ is *not* irreducible, it can be decomposed into a direct sum of irreps τ_i , each of which acts on an invariant subspace, and we write $\sigma \cong \tau_1 \oplus \cdots \oplus \tau_k$. In general, a given τ can appear multiple times in this decomposition, in

the sense that σ may have an invariant subspace isomorphic to the direct sum of a_τ copies of τ . In this case a_τ is called the *multiplicity* of τ in the decomposition of σ .

There is a natural product operation on representations: if $\lambda : G \rightarrow \mathbf{U}(V)$ and $\mu : G \rightarrow \mathbf{U}(W)$ are representations of G , we may define a new representation $\lambda \otimes \mu : G \rightarrow \mathbf{U}(V \otimes W)$ as $(\lambda \otimes \mu)(g) : \mathbf{u} \otimes \mathbf{v} \mapsto \lambda(g)\mathbf{u} \otimes \mu(g)\mathbf{v}$. This representation corresponds to the *diagonal action* of G on $V \otimes W$, in which we apply the same group element to both parts of the tensor product. In general, the representation $\lambda \otimes \mu$ is not irreducible, even when both λ and μ are. This leads to the *Clebsch-Gordan problem*, that of decomposing $\lambda \otimes \mu$ into irreps.

Given a representation σ we define the *character* of σ , denoted χ_σ , to be the trace $\chi_\sigma(g) = \mathbf{tr} \sigma(g)$. As the trace of a linear operator is invariant under conjugation, characters are constant on the conjugacy classes of G . Characters are a powerful tool for reasoning about the decomposition of reducible representations. In particular, when $\sigma = \bigoplus_i \tau_i$ we have $\chi_\sigma = \sum_i \chi_{\tau_i}$ and, moreover, for $\sigma, \tau \in \widehat{G}$, we have the orthogonality conditions

$$\langle \chi_\sigma, \chi_\tau \rangle_G = \frac{1}{|G|} \sum_{g \in G} \chi_\sigma(g) \chi_\tau(g)^* = \begin{cases} 1 & \sigma = \tau \\ 0 & \sigma \neq \tau \end{cases} .$$

Therefore, given a representation σ and an irrep τ , the multiplicity a_τ with which τ appears in the decomposition of σ is $\langle \chi_\tau, \chi_\sigma \rangle_G$. For example, since $\chi_{\lambda \otimes \mu}(g) = \chi_\lambda(g) \cdot \chi_\mu(g)$, the multiplicity of τ in the Clebsch-Gordan decomposition of $\lambda \otimes \mu$ is $\langle \chi_\tau, \chi_\lambda \chi_\mu \rangle_G$.

A representation σ is said to be *isotypic* if the irreducible factors appearing in the decomposition are all isomorphic, which is to say that there is a single nonzero a_τ in the decomposition above. Any representation σ may be uniquely decomposed into maximal isotypic subspaces, one for each irrep τ of G ; these subspaces are precisely those spanned by all copies of τ in σ . In fact, for each τ this subspace is the image of an explicit projection operator Π_τ which can be written as

$$\Pi_\tau = \frac{1}{|G|} \sum_{g \in G} d_\tau \chi_\tau(g)^* \sigma(g) .$$

A useful fact is that Π_τ commutes with the group action; that is, for any $h \in G$ we have

$$\begin{aligned} \sigma(h)\Pi_\tau\sigma(h)^\dagger &= \frac{1}{|G|} \sum_{g \in G} d_\tau \chi_\tau(g)^* \sigma(hgh^{-1}) = \\ &= \frac{1}{|G|} \sum_{g \in G} d_\tau \chi_\tau(h^{-1}gh)^* \sigma(g) = \frac{1}{|G|} \sum_{g \in G} d_\tau \chi_\tau(g)^* \sigma(g) = \Pi_\tau . \end{aligned}$$

Our algorithms will perform measurements which project into these maximal isotypic subspaces and observe the resulting irrep name τ . For the particular case of coset states, this measurement is called *weak Fourier sampling* in the literature; however, since we are interested in a more general process which in fact performs a kind of strong multiregister sampling on the original coset states, we will use the term *isotypic sampling* instead. Finally, we discuss the structure of a specific representation, the (*right*) *regular* representation **reg**, which plays an important role in the analysis below. **reg** is given by the permutation action of G on itself. Specifically, let $\mathbb{C}[G]$ be the *group algebra* of G ; this is the $|G|$ -dimensional vector space of formal sums

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

(Note that $\mathbb{C}[G]$ is precisely the Hilbert space of a single register containing a superposition of group elements.) Then **reg** is the representation $\mathbf{reg} : G \rightarrow \mathbf{U}(\mathbb{C}[G])$ given by linearly extending right multiplication, $\mathbf{reg}(g) : h \mapsto hg$. It is not hard to see that its character $\chi_{\mathbf{reg}}$ is given by

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

in which case we have $\langle \chi_{\mathbf{reg}}, \chi_\sigma \rangle_G = d_\sigma$ for each $\sigma \in \widehat{G}$. Thus **reg** contains d_σ copies of each irrep $\sigma \in \widehat{G}$, and counting dimensions on each side of this decomposition implies

$$(1) \quad |G| = \sum_{\sigma \in \widehat{G}} d_\sigma^2 .$$

This equation suggests a natural probability distribution on \widehat{G} , the *Plancherel distribution*, which assigns to each irrep σ the probability $\mathcal{P}_{\text{planch}}^G(\sigma) = d_\sigma^2/|G|$. This is simply the dimensionwise fraction of $\mathbb{C}[G]$ consisting of copies of σ ; indeed, if we perform isotypic sampling on the completely mixed state on $\mathbb{C}[G]$, or equivalently the coset state where the hidden subgroup is trivial, we observe exactly this distribution.

In general, we can consider subspaces of $\mathbb{C}[G]$ that are invariant under left multiplication, right multiplication, or both; these subspaces are called *left-*, *right-*, or *bi-invariant* respectively. For each $\sigma \in \widehat{G}$, the maximal σ -isotypic subspace is a d_σ^2 -dimensional bi-invariant subspace; it can be broken up further into d_σ d_σ -dimensional left-invariant subspaces, or (transversely) d_σ d_σ -dimensional right-invariant subspaces. However, this decomposition is not unique. If σ acts on a vector space V , then choosing an orthonormal basis for V allows us to view $\sigma(g)$ as a $d_\sigma \times d_\sigma$ matrix. Then σ acts on the d_σ^2 -dimensional space of such matrices by left or right multiplication, and the columns and rows correspond to left- and right-invariant spaces respectively.

3. CLEBSCH-GORDAN SIEVES

Consider the hidden subgroup problem over a group G with the added promise that the hidden subgroup H is either the trivial subgroup, or a conjugate of some fixed nontrivial subgroup H_0 . We shall consider sieve algorithms for this problem that proceed as follows:

1. The oracle is used to generate $\ell = \ell(n)$ coset states ρ_H , each of which is subjected to weak Fourier sampling. This results in a set of states ρ_i , where ρ_i is a mixed state known to lie in the σ_i -isotypic subspace of $\mathbb{C}[G]$ for some irrep σ_i .

2. The following *combine-and-measure* procedure is then repeated as many times as we like. Two states ρ_i and ρ_j in the set are selected according to an arbitrary adaptive rule that may depend on the entire history of the computation (in existing algorithms of this type, this selection in fact depends only on the irreps σ_i and σ_j in which they lie). We then perform isotypic sampling on their tensor product $\rho_i \otimes \rho_j$: that is, we apply a measurement operator which observes an irrep σ in the Clebsch-Gordan decomposition of $\sigma_i \otimes \sigma_j$ (see [Kup05] or [MR05a] for how this measurement can actually be carried out by applying the diagonal action). This measurement destroys ρ_i and ρ_j , and results in a new mixed state ρ which lies in the maximal σ -isotypic subspace; we add this new state to the set.

3. Finally, depending on the sequence of observations obtained throughout this process, the algorithm guesses the hidden subgroup.

We set down some notation to discuss the result of applying such an algorithm. Fixing a group G and a subgroup H , let A be a sieve algorithm which initially generates ℓ coset states. As a bookkeeping tool, we will describe intermediate states of A 's progress as a *forest of labeled binary trees*. Throughout, we will maintain the invariant that the roots of the trees in this forest correspond to the current set of states available to the algorithm.

Initially, the state of the algorithm consists of a forest consisting of ℓ single-node trees, each of which is labeled with the irrep name σ_i that resulted from weak Fourier sampling a coset state, and is associated with the resulting state ρ_i . Then, each combine-and-measure step selects two root nodes, r_1 and r_2 , and applies isotypic sampling to the tensor product of their states. We associate the resulting state ρ with a new root node r , and place the nodes r_1 and r_2 below it as its children. We label this new node with the irrep name σ observed in this measurement.

Thus, every node of the forest corresponds to a state that existed at some point during the algorithm, and each node i is labeled with the name of the irrep σ_i observed in the isotypic measurement performed when that node was created. We call the resulting labeled forest the *transcript* of the algorithm: note that this transcript contains all the information the algorithm may use to determine the hidden subgroup.

We make several observations about algorithms of this type. First, it is easy to see that nothing is gained by combining $t > 2$ states at a time; we can simulate this with an algorithm which builds a binary tree with t leaves, and which ignores the results of all its measurements except the one at the root.

Second, the algorithm maintains the following kind of symmetry under the action of the subgroup H . Suppose we have a representation σ acting on a Hilbert space V . Given a subgroup H , we say that a state $\psi \in V$ is H -invariant if $\sigma(h) \cdot \psi = \psi$ for all $h \in H$. Similarly, given a mixed state ρ , we say that ρ is H -invariant if $\sigma(h) \cdot \rho \cdot \sigma(h)^\dagger = \rho$ or, equivalently, if $\sigma(h)$ and ρ commute. For instance, the coset state ρ_H is H -invariant under the right regular representation, since right-multiplying by any $h \in H$ preserves each left coset cH . Now, suppose that ρ_1 and ρ_2 are H -invariant; clearly $\rho_1 \otimes \rho_2$ is H -invariant under the diagonal action, and performing isotypic sampling preserves H -invariance since Π_τ commutes with the action of any group element. Thus the states produced by the algorithm are H -invariant throughout.

Third, it is important to note that while at each stage we observe only an irrep name, rather than a basis vector inside that representation, by iterating this process the sieve algorithm actually performs a kind of *strong multi-register Fourier sampling* on the original set of coset states. For instance, in the dihedral group, suppose that performing weak Fourier sampling on two coset states results in the two-dimensional irreps σ_j and σ_k , and that we then observe the irrep σ_{j+k} under isotypic sampling of their tensor product. We now know that the original coset states were in fact confined to a particular subspace, spanned by two entangled pairs of basis vectors. Finally, we note that the states produced by a sieve algorithm are quite different from coset states. In particular, they belong not to a maximal isotypic subspace

of $\mathbb{C}[G]$, but to a (typically much higher-dimensional) *non-maximal* isotypic subspace of $\mathbb{C}[G]^{\otimes \ell}$, where ℓ is the number of coset states feeding into that state (i.e., the number of leaves of the corresponding tree). Moreover, they have more symmetry than coset states, since each isotypic measurement implies a symmetry with respect to the diagonal action on the set of leaves descended from the corresponding internal node. In the next sections we will show how these states can be written in terms of projection operators applied to this high-dimensional space.

4. OBSERVED DISTRIBUTIONS FOR FIXED TOPOLOGIES

In general, the probability distributions arising from the combine-and-measure steps of a sieve algorithm depend on both the hidden subgroup and the entire history of previous measurements and observations (that is, the labeled forest, or transcript, describing the algorithm's history thus far). In this section and the next, we focus on the probability distribution induced by a *fixed forest topology* and subgroup H . We can think of this either as the probability distribution conditioned on the forest topology, or as the distribution of transcripts produced by some *non-adaptive* sieve algorithm, which chooses which states it will combine and measure ahead of time. We will show that for all forest topologies of sufficiently small size, the induced distributions on irrep labels fail to distinguish trivial and nontrivial subgroups. Then, in Section 7, we will complete the argument for adaptive algorithms. Clearly, in this non-adaptive case the distributions of irrep labels associated with different trees in the forest are independent. Therefore, we can focus on the distribution of labels for a specific tree. At the leaves, the labels are independent and identically distributed according to the distribution resulting from weak Fourier sampling a coset state [HRTS00]. However, as we move inside the tree and condition on the irrep labels observed previously, the resulting distributions are quite different from this initial one. To calculate the resulting joint probability distribution, we need to define projection operators acting on $\mathbb{C}[G]^{\otimes \ell}$ corresponding to the isotypic measurement at each node.

First, note that the coset state ρ_H can be written in the following convenient form:

$$\rho_H = \frac{1}{|G|} \sum_c |cH\rangle \langle cH| = \frac{1}{|G|} \sum_{h \in H} \mathbf{reg}(h)$$

where \mathbf{reg} is the *right regular representation*: that is, ρ_H is proportional to the projection operator which right-multiplies by a random element of H ,

$$\Pi_H = \frac{1}{|H|} \sum_{h \in H} \mathbf{reg}(h) .$$

If H is trivial, ρ_H is the completely mixed state $\rho_{\{1\}} = (1/|G|)\mathbb{1}$. On the other hand, if $H = \{1, m\}$ for an involution m , then $\rho_H = (2/|G|)\Pi_H$, where Π_H is the projection operator

$$\Pi_H = \frac{1}{2}(1 + \mathbf{reg}(m)) .$$

Now consider the tensor product of ℓ “registers,” each containing a coset state. Given a linear operator M on $\mathbb{C}[G]$ and a subset $I \subseteq [\ell] = \{1, \dots, \ell\}$, let M^I denote the operator on $\mathbb{C}[G^\ell] \cong \mathbb{C}[G]^{\otimes \ell}$ which applies M to the registers in I and leaves the other registers unchanged. Then the mixed state consisting of ℓ independent coset states is $\rho_H^{\otimes \ell} = (2/|G|)^\ell \Pi_H^{\otimes \ell}$, where

$$(2) \quad \Pi_H^{\otimes \ell} = \frac{1}{2^\ell} \prod_{j=1}^{\ell} (1 + \mathbf{reg}(m)^{\{j\}}) = \frac{1}{2^\ell} \sum_{I \subseteq [\ell]} \mathbf{reg}(m)^I .$$

Note the sum over subsets of registers, a theme which has appeared repeatedly in discussions of multiregister Fourier sampling [Reg02, BCvD06, HMR⁺06, Kup05, MR05a, MR05b]. Now consider a tree T with ℓ leaves corresponding to the ℓ initial registers, and k nodes including the leaves. We represent this tree as a set system, in which each node i is associated with the subset $I_i \subseteq [\ell]$ of leaves descended from it. In particular, $I_{\text{root}} = [\ell]$ and $I_j = \{j\}$ for each leaf j .

Performing isotypic sampling at a node i corresponds to applying the diagonal action to its children (or in terms of the algorithm, its parents) and inductively to the registers in I_i : that is, we multiply each register in I_i by the same element g and leave the others fixed. If σ_i is the irrep label observed at that node, let us denote its character and dimension by χ_i and d_i respectively, rather than the more cumbersome χ_{σ_i} and d_{σ_i} . Then the projection operator corresponding to this observation is

$$(3) \quad \Pi_i^T = \frac{1}{|G|} \sum_{g \in G} d_i \chi_i(g)^* \mathbf{reg}(g)^{I_i} .$$

Now consider a transcript of the sieve process which results in observing a set of irrep labels $\sigma = \{\sigma_i\}$ on the internal nodes of the tree. The projection operator associated with this outcome is

$$(4) \quad \Pi^T[\sigma] = \prod_{i=1}^k \Pi_i^T .$$

We will abbreviate this as Π^T whenever the context is clear. Note that the various Π_i^T in the product (4) pairwise commute, since for any two nodes i, j either I_i and I_j are disjoint, or one is contained in the other. In the former case a^{I_i} and b^{I_j} for all a, b . In the latter case, say if $I_i \subset I_j$, we have

$a^{I_i} b^{I_j} = b^{I_j} (b^{-1} a b)^{I_i}$, and since $\chi_i(b^{-1} a b) = \chi_i(a)$ it follows from (3) that $\Pi_i^T \Pi_j^T = \Pi_j^T \Pi_i^T$.

Given a tree T with k nodes, we write $P_T^{\{1\}}[\sigma]$ for the probability that we observe the set of irrep labels $\sigma = \{\sigma_i\}$ in the case where the hidden subgroup is trivial. Since the tensor product of coset states is then the completely mixed state in $\mathbb{C}[G^\ell]$, this is simply the dimensionwise fraction of $\mathbb{C}[G^\ell]$ consisting of the image of Π^T , or

$$P_T^{\{1\}}[\sigma] = \frac{1}{|G|^\ell} \mathbf{tr} \Pi^T .$$

Moreover, since measuring a completely mixed state results in the completely mixed state in the observed subspace, each state produced by the algorithm is completely mixed in the image of Π^T . In particular, if the irrep label at the root of a tree is σ , the corresponding state consists of a classical mixture across some number of copies of σ , in each of which it is completely mixed. Thus, when combining two parent states with irrep labels λ and μ , we observe each irrep τ with probability equal to the dimensionwise fraction of $\lambda \otimes \mu$ consisting of copies of τ , namely

$$(5) \quad \mathcal{P}_{\lambda \otimes \mu}(\tau) = \frac{d_\tau}{d_\lambda d_\mu} \langle \chi_\tau, \chi_\lambda \chi_\mu \rangle_G$$

(recall that $\langle \chi_\tau, \chi_\rho \rangle_G = (1/|G|) \sum_{g \in G} \chi_\tau(g) \chi_\rho^*(g)$ is the multiplicity of τ in the decomposition of a representation ρ into irreducibles). We will refer to this as the *natural distribution* in $\lambda \otimes \mu$.

Now let us consider the case where the hidden subgroup is nontrivial. Since the mixed state ρ_{H^ℓ} can be thought of as a pure state chosen randomly from the image of $\Pi_H^{\otimes \ell}$, the probability of observing a set of irrep labels σ in this case is

$$P_T^H[\sigma] = \frac{\mathbf{tr} \Pi^T \Pi_H^{\otimes \ell}}{\mathbf{tr} \Pi_H^{\otimes \ell}} = \frac{2^\ell}{|G|^\ell} \mathbf{tr} \Pi^T \Pi_H^{\otimes \ell}$$

where we use the fact that $\mathbf{tr} \Pi_H^{\otimes \ell} = [G : H]^\ell = (|G|/2)^\ell$. Below we abbreviate these distributions as $P_T^{\{1\}}$ and P_T^H whenever the context is clear. Our goal is to show that, until the tree T is deep enough, these two distributions are extremely close, so that the algorithm fails to distinguish subgroups of the form $\{1, m\}$ from the trivial subgroup.

Now let us derive explicit expressions for $P_T^{\{1\}}$ and P_T^H . First, we fix some additional notation. Given an assignment of group elements $\{a_i\}$ to the nodes, for each leaf j we let $\prod_{i \rightsquigarrow j} a_i$ denote the product of the elements along the path from the root to j :

$$\prod_{i \rightsquigarrow j} a_i = \prod_{i: j \in I_i} a_i$$

where the product is taken in order from the root to the leaf. Then using (3) and (4) we can write

$$(6) \quad \Pi^T = \frac{1}{|G|^k} \sum_{\{a_i\}} d_i \chi_i(a_i)^* \bigotimes_{j=1}^{\ell} \mathbf{reg} \left(\prod_{i \rightsquigarrow j} a_i \right)^{\{i\}}.$$

We say that an assignment $\{a_i\}$ is *trivial* if $\prod_{i \rightsquigarrow j} a_i = 1$ for every leaf j . Then, since $\mathbf{tr} \mathbf{reg}(g) = \chi_{\mathbf{reg}}(g) = |G|$ if $g = 1$ and 0 otherwise, we have

$$(7) \quad P_T^{\{1\}} = \frac{1}{|G|^\ell} \mathbf{tr} \Pi^T = \frac{1}{|G|^k} \sum_{\{a_i\} \text{ trivial}} \prod_{i=1}^k d_i \chi_i(a_i)^*.$$

To get a sense of how this expression scales, note that the particular trivial assignment where $a_i = 1$ for all i contributes $\prod_{i=1}^k d_i^2 / |G| = \prod_i \mathcal{P}_{\text{planch}}(\sigma_i)$, as if the σ_i were independent and Plancherel-distributed.

Now consider P_T^H . Combining (2) with (6) gives the following expression for $\Pi^T \Pi_H^{\otimes \ell}$:

$$(8) \quad \Pi^T \Pi_H^{\otimes \ell} = \frac{1}{2^\ell |G|^k} \sum_{\{a_i\}} d_i \chi_i(a_i)^* \bigotimes_{j=1}^{\ell} \mathbf{reg} \left(\left(\prod_{i \rightsquigarrow j} a_i \right) (1+m) \right)^{\{i\}}.$$

We say that an assignment $\{a_i\}$ is *legal* if $\prod_{i \rightsquigarrow j} a_i \in \{1, m\}$ for every leaf j . Then the trace of the term corresponding to $\{a_i\}$ is $|G|^\ell$ if $\{a_i\}$ is legal, and is 0 otherwise, and analogous to (7) we have

$$(9) \quad P_T^H = \frac{2^\ell}{|G|^\ell} \mathbf{tr} \Pi^T \Pi_H^{\otimes \ell} = \frac{1}{|G|^k} \sum_{\{a_i\} \text{ legal}} \prod_{i=1}^k d_i \chi_i(a_i)^*.$$

Thus these two distributions differ exactly by the terms corresponding to assignments which are legal but nontrivial. Our main result will depend on the fact that for most σ these terms are identically zero, in which case P_T^H and $P_T^{\{1\}}$ coincide.

5. THE IMPORTANCE OF BEING HOMOGENEOUS

For any group G , the wreath product $G \wr \mathbb{Z}_2$ is the semidirect product $(G \times G) \rtimes \mathbb{Z}_2$, where we extend $G \times G$ by an involution which exchanges the two copies of G . Thus the elements $((\alpha, \beta), 0)$ form a normal subgroup $K \cong G \times G$ of index 2, and the elements $((\alpha, \beta), 1)$ form its nontrivial coset. We will call these elements “non-flips” and “flips,” respectively. The Graph Isomorphism problem reduces to the hidden subgroup problem on $S_n \wr \mathbb{Z}_2$ in the following natural way. We consider the disjoint union of the two graphs, and consider permutations of their $2n$ vertices. Then $S_n \wr \mathbb{Z}_2$

is the subgroup of S_{2n} which either maps each graph onto itself (the non-flips) or exchanges the two graphs (the flips). We assume for simplicity that the graphs are rigid. Then if they are nonisomorphic, the hidden subgroup is trivial; if they are isomorphic, $H = \{1, m\}$ where m is a flip of the form $((\alpha, \alpha^{-1}), 1)$, where α is the permutation describing the isomorphism between them.

For any group G , the irreps of $G \wr \mathbb{Z}_2$ can be written in a simple way in terms of the irreps of G . It is useful to construct them by *inducing* upward from the irreps of $K \cong G \times G$ (see [Ser77] for the definition of an induced representation). First, each irrep of K is the tensor product $\lambda \otimes \mu$ of two irreps of G . Inducing this irrep from K up to G gives a representation

$$\sigma_{\{\lambda, \mu\}} = \text{Ind}_K^G(\lambda \otimes \mu)$$

of dimension $2d_\lambda d_\mu$. If $\lambda \not\cong \mu$, then this is irreducible, and $\sigma_{\{\lambda, \mu\}} \cong \sigma_{\{\mu, \lambda\}}$ (hence the notation). We call these irreps *inhomogeneous*. Their characters are given by

$$(10) \quad \chi_{\{\lambda, \mu\}}((\alpha, \beta), t) = \begin{cases} \chi_\lambda(\alpha)\chi_\mu(\beta) + \chi_\mu(\alpha)\chi_\lambda(\beta) & \text{if } t = 0 \\ 0 & \text{if } t = 1 \end{cases} .$$

In particular, the character of an inhomogeneous irrep is zero at any flip.

On the other hand, if $\lambda \cong \mu$, then $\sigma_{\{\lambda, \lambda\}}$ decomposes into two irreps of dimension d_λ^2 , which we denote $\sigma_{\{\lambda, \lambda\}}^+$ and $\sigma_{\{\lambda, \lambda\}}^-$. We call these irreps *homogeneous*. Their characters are given by

$$(11) \quad \chi_{\{\lambda, \lambda\}}^\pm((\alpha, \beta), t) = \begin{cases} \chi_\lambda(\alpha)\chi_\lambda(\beta) & \text{if } t = 0 \\ \pm\chi_\lambda(\alpha\beta) & \text{if } t = 1 \end{cases} .$$

In the next section, we will show that sieve algorithms obtain precisely zero information that distinguishes hidden subgroups of the form $\{1, m\}$ from the trivial subgroup until it observes at least one homogeneous representation.

Suppose that the irrep labels $\sigma = \{\sigma_i\}$ observed during a run of the sieve algorithm consist entirely of inhomogeneous irreps of $G \wr \mathbb{Z}_2$. Since the irreps have zero character at any flip, the only trivial or legal assignments $\{a_i\}$ that contribute to the sums (7) and (9) are those where each a_i is a non-flip, i.e., is contained in the subgroup $K \cong G \times G$. But the product of any string of such elements is also contained in K , so if this product is in $H = \{1, m\}$ where $m \notin K$, it is equal to 1. Thus any legal assignment of this kind is trivial, the sums (7) and (9) coincide, and the probability of observing σ is the same in the trivial and nontrivial cases. That is, so long as every σ_i in σ is inhomogeneous,

$$(12) \quad P_T^H[\sigma] = P_T^{\{1\}}[\sigma] .$$

Our strategy will be to show that observing even a single homogeneous irrep is unlikely, unless the tree generated by the sieve algorithm is quite large. Moreover, because the two distributions coincide unless this occurs, it suffices to show that this is unlikely in the case where H is trivial. Now, it is easy to see that the probability of observing a given representation in $G \wr \mathbb{Z}_2$, under either the Plancherel distribution or a natural distribution, factorizes neatly into the probabilities that we observe the corresponding pair of irreps, in either order, in a pair of similar experiments in G . First, the Plancherel measure of an inhomogeneous irrep $\sigma_{\{\lambda, \mu\}}$ is

$$(13) \quad \mathcal{P}_{\text{planch}}^{G \wr \mathbb{Z}_2}(\sigma_{\{\lambda, \mu\}}) = \frac{(2d_\lambda d_\mu)^2}{2|G|^2} = 2 \mathcal{P}_{\text{planch}}^G(\lambda) \mathcal{P}_{\text{planch}}^G(\mu) .$$

Similarly, the probability that we observe a homogeneous irrep $\sigma_{\{\lambda, \lambda\}}^\pm$ is the probability of observing λ twice under the Plancherel distribution in G , in which case the sign \pm is chosen uniformly:

$$(14) \quad \mathcal{P}_{\text{planch}}^{G \wr \mathbb{Z}_2}(\sigma_{\{\lambda, \lambda\}}^\pm) = \frac{d_\lambda^4}{|G|^2} = \mathcal{P}_{\text{planch}}^G(\lambda)^2 .$$

Now consider the natural distribution in the tensor product of two inhomogeneous irreps $\sigma_{\{\lambda, \lambda'\}}$ and $\sigma_{\{\mu, \mu'\}}$. The multiplicity of a given homogeneous irrep $\sigma_{\{\tau, \tau\}}^\pm$ in this tensor product, equal to

$$\left\langle \chi_{\{\tau, \tau\}}^\pm, \chi_{\{\lambda, \lambda'\}} \chi_{\{\mu, \mu'\}} \right\rangle_{G \wr \mathbb{Z}_2} ,$$

factorizes as follows

$$\frac{\langle \chi_\tau, \chi_\lambda \chi_{\lambda'} \rangle_G \langle \chi_\tau, \chi_{\mu'} \chi_\mu \rangle_G}{2} + \frac{\langle \chi_\tau, \chi_\lambda \chi_{\mu'} \rangle_G \langle \chi_\tau, \chi_{\lambda'} \chi_\mu \rangle_G}{2} .$$

Thus the probability of observing either $\sigma_{\{\tau, \tau\}}^+$ or $\sigma_{\{\tau, \tau\}}^-$ under the natural distribution is

$$(15) \quad \mathcal{P}_{\sigma_{\{\lambda, \lambda'\}} \otimes \sigma_{\{\mu, \mu'\}}}(\sigma_{\{\tau, \tau\}}^\pm) = \frac{1}{2} (\mathcal{P}_{\lambda \otimes \mu}(\tau) \mathcal{P}_{\lambda' \otimes \mu'}(\tau) + \mathcal{P}_{\lambda \otimes \mu'}(\tau) \mathcal{P}_{\lambda' \otimes \mu}(\tau)) .$$

In other words, the probability of observing a homogeneous irrep of $G \wr \mathbb{Z}_2$ is the probability of observing the same irrep in two natural distributions on G . Let us denote the probability that we observe the same irrep in the natural distributions in $\lambda \otimes \mu$ and $\lambda' \otimes \mu'$ —that is, that these two distributions collide—as

$$\mathcal{P}_{\lambda \otimes \mu, \lambda' \otimes \mu'}^{\text{coll}} = \sum_{\tau} \mathcal{P}_{\lambda \otimes \mu}(\tau) \mathcal{P}_{\lambda' \otimes \mu'}(\tau) .$$

Then (15) implies that the total probability of observing a homogeneous irrep is

$$(16) \quad \sum_{\tau} \mathcal{P}_{\sigma_{\{\lambda, \lambda'\}} \otimes \sigma_{\{\mu, \mu'\}}}(\sigma_{\{\tau, \tau\}}^{\pm}) = \frac{1}{2} (\mathcal{P}_{\lambda \otimes \mu, \lambda' \otimes \mu'}^{\text{coll}} + \mathcal{P}_{\lambda \otimes \mu', \lambda' \otimes \mu}^{\text{coll}}) \\ \leq \max(\mathcal{P}_{\lambda \otimes \mu, \lambda' \otimes \mu'}^{\text{coll}}, \mathcal{P}_{\lambda \otimes \mu', \lambda' \otimes \mu}^{\text{coll}}) .$$

In the next section, we show that if λ, μ, λ' and μ' are typical irreps of S_n , then no irrep τ occurs too often in any of these natural distributions, and so the probability of a collision is small.

6. COLLISIONS, SMOOTHNESS, AND CHARACTERS

Let us bound the probability $\mathcal{P}^{\text{coll}} = \mathcal{P}_{\lambda \otimes \mu, \lambda' \otimes \mu'}^{\text{coll}}$ that the natural distributions in $\lambda \otimes \mu$ and $\lambda' \otimes \mu'$ collide. The idea is that $\mathcal{P}^{\text{coll}}$ is small as long as both of either or both of these distributions is *smooth*, in the sense that they are spread fairly uniformly across many τ . The following lemmas show that this notion of smoothness can be related to bounds on the normalized characters of these representations. First, we present a lemma which relates the natural distribution in a representation ρ to the Plancherel distribution.

Lemma 1. *Let ρ be a (possibly reducible) representation of a group G , and let $\mathcal{P}_{\rho}(\tau)$ denote the probability of observing an irrep $\tau \in \widehat{G}$ under the natural distribution in ρ . Let $X \subseteq \widehat{G}$, and let $\mathcal{P}_{\rho}(X) = \sum_{\tau \in X} \mathcal{P}_{\rho}(\tau)$ and $\mathcal{P}_{\text{planch}}(X) = \sum_{\tau \in X} d_{\tau}^2/|G|$ denote the total probability of observing an irrep in X in the natural and Plancherel distributions respectively. Then*

$$\mathcal{P}_{\rho}(X) \leq \sqrt{\mathcal{P}_{\text{planch}}(X)} \sqrt{\sum_{g \in G} \left| \frac{\chi_{\rho}(g)}{d_{\rho}} \right|^2} .$$

Proof. In general, we have

$$\mathcal{P}_{\rho}(\tau) = \frac{d_{\tau}}{d_{\rho}} \langle \chi_{\tau}, \chi_{\rho} \rangle_G .$$

Therefore, if we define

$$\mathbb{1}_X = \sum_{\tau \in X} d_{\tau} \chi_{\tau} ,$$

then by Cauchy-Schwartz we have

$$\begin{aligned} \mathcal{P}_\rho(X) &= \left\langle \mathbb{1}_X, \frac{\chi_\rho}{d_\rho} \right\rangle_G \leq \sqrt{\langle \mathbb{1}_X, \mathbb{1}_X \rangle_G} \sqrt{\left\langle \frac{\chi_\rho}{d_\rho}, \frac{\chi_\rho}{d_\rho} \right\rangle_G} = \\ &= \sqrt{\frac{1}{|G|} \langle \mathbb{1}_X, \mathbb{1}_X \rangle_G} \sqrt{\sum_{g \in G} \left| \frac{\chi_\rho(g)}{d_\rho} \right|^2} \end{aligned}$$

and by Schur's lemma we have

$$\frac{1}{|G|} \langle \mathbb{1}_X, \mathbb{1}_X \rangle_G = \sum_{\tau \in X} \frac{d_\tau^2}{|G|} \langle \chi_\tau, \chi_\tau \rangle_G = \sum_{\tau \in X} \frac{d_\tau^2}{|G|} = \mathcal{P}_{\text{planch}}(X)$$

which completes the proof. \square

Now we bound the probability of a collision as follows.

Lemma 2. *Given a family of groups $\{G_n\}$, say that an irrep λ of G_n is $f(n)$ -smooth if*

$$\sum_{g \in G_n} \left| \frac{\chi_\lambda(g)}{d_\lambda} \right|^4 \leq f(n) .$$

Suppose that λ and μ are $f(n)$ -smooth. Then

$$\mathcal{P}^{\text{coll}} \leq \frac{\max_\tau d_\tau}{\sqrt{|G_n|}} \sqrt{f(n)} .$$

Proof. We write G for G_n to conserve ink. We have $\mathcal{P}^{\text{coll}} \leq \max_\tau \mathcal{P}_{\lambda \otimes \mu}(\tau)$. Setting $\rho = \lambda \otimes \mu$ and $X = \{\tau\}$ in Lemma 1 and applying Cauchy-Schwartz gives

$$\begin{aligned} \mathcal{P}^{\text{coll}} &\leq \sqrt{\max_\tau \mathcal{P}_{\text{planch}}(\tau)} \sqrt{\sum_{g \in G} \left| \frac{\chi_\lambda(g)}{d_\lambda} \right|^2 \left| \frac{\chi_\mu(g)}{d_\mu} \right|^2} \leq \\ &= \frac{\max_\tau d_\tau}{\sqrt{|G|}} \sqrt[4]{\sum_{g \in G} \left| \frac{\chi_\lambda(g)}{d_\lambda} \right|^4 \sum_{g \in G} \left| \frac{\chi_\mu(g)}{d_\mu} \right|^4} \end{aligned}$$

which completes the proof. \square

Now let us focus on the case relevant to Graph Isomorphism, where $G = S_n$. Here we recall that each irrep of the symmetric group S_n corresponds to a *Young diagram*, or equivalently an integer partition $\lambda_1 \geq \lambda_2 \geq \dots$ where $\sum_i \lambda_i = n$. The maximum dimension of any irrep is bounded by the following result of Vershik and Kerov:

Theorem 3 ([VK85]). *There is a constant $\hat{c} > 0$ such that $\max_{\tau} d_{\tau} \leq e^{-(\hat{c}/2)\sqrt{n}}\sqrt{n!}$.*

In this case, Lemma 2 gives

$$(17) \quad \mathcal{P}^{\text{coll}} \leq e^{-(\hat{c}/2)\sqrt{n}}\sqrt{f(n)} .$$

Therefore, our goal is to show that typical irreps of S_n are $f(n)$ -smooth where $f(n)$ grows slowly enough with n , and to show inductively that with high probability all the irreps we observe throughout the sieve are typical. We do this by defining a typical irrep as follows.

Definition 4. *Let $D > e$ be a fixed constant, and say that an irrep λ of S_n is typical if the following two conditions hold true:*

- *the height and width of its Young diagram are less than $D\sqrt{n}$ or, in other words, if the Young diagram is D -balanced [Bia98],*
- *the dimension d_{λ} fulfills*

$$d_{\lambda} > e^{-\frac{1}{2}\sqrt{n}\log n}\sqrt{n!} .$$

To motivate this definition, and to provide the base case for our induction, we show the following.

Lemma 5. *There are constants $c > 0$ and n_0 such that, if λ has n boxes with $n > n_0$ and λ is chosen according to the Plancherel distribution, then λ is typical with probability at least $1 - e^{-c\sqrt{n}}$.*

Proof. Firstly, we bound the probability that λ is not D -balanced. The Robinson-Schensted correspondence [Ful97] maps permutations to Young diagrams in such a way that the uniform measure on S_n maps to the Plancherel measure. In addition, the width (resp. height) of the Young diagram is equal to the length of the longest increasing (resp. decreasing) subsequence. Therefore, the probability in the Plancherel measure that an irrep is not typical is at most twice the probability that a random permutation has an increasing subsequence of length $w = D\sqrt{n}$.

The problem of determining the typical size of the longest increasing subsequence is known as Ulam's problem; it can be solved using representation theory [Ker03] or by a beautiful hydrodynamic argument [AD95], and indeed this Lemma holds even if we take $D > 2$ in Definition 4. Here we content ourselves with an elementary bound for $D > e$. By Markov's inequality, the probability an increasing subsequence of length $w = D\sqrt{n}$ is at most the expected number of such subsequences, which is

$$(18) \quad \binom{n}{w} \frac{1}{w!} < \left(\frac{e^2 n}{w^2}\right)^w = \left(\frac{e^2}{D^2}\right)^{D\sqrt{n}}$$

where we used Stirling's approximation $w! > w^w e^{-w}$.

Secondly, we shall bound the probability that

$$(19) \quad d_\lambda \leq e^{-\frac{1}{2}\sqrt{n}\log n} \sqrt{n!} .$$

The number of irreps is the partition number

$$p(n) = (1 + o(1)) \frac{1}{4\sqrt{3} \cdot n} e^{\delta\sqrt{n}} < e^{\delta\sqrt{n}}$$

where

$$\delta = \sqrt{2/3} \pi ;$$

therefore the Plancherel measure of the set of irreps λ of S_n for which (19) holds true is at most the number of irreps times the measure of a single such λ , so this probability is at most

$$(20) \quad p(n) \frac{d_\lambda^2}{n!} < e^{\delta\sqrt{n}} e^{-\sqrt{n}\log n} = e^{-\omega(\sqrt{n})} .$$

The sum of the probabilities (18) and (20) is bounded from above by $e^{-c\sqrt{n}}$ for sufficiently small $c > 0$ and for n sufficiently large. \square

Given a permutation π , let $t(\pi)$ denote the length of the shortest sequence of transpositions whose product is π ; for instance, if π is a single k -cycle, then $t(\pi) = k - 1$.

Lemma 6. *There is a constant A such that, for n sufficiently large, the normalized character of all typical λ obeys*

$$\left| \frac{\chi_\lambda(\pi)}{d_\lambda} \right| \leq \left(\frac{A}{\sqrt{n}} \right)^{t(\pi)}$$

for all $\pi \in S_n$ with $t(\pi) > \sqrt{n} \log n$.

Proof. We use the Murnaghan-Nakayama formula for the character [JK81]. A *ribbon tile* of length k is a polyomino of k cells, arranged in a path where each step is up or to the right. Given a Young diagram λ and a permutation π with cycle structure $k_1 \geq k_2 \geq \dots$, a *consistent tiling* consists of removing a ribbon tile of length k_1 from the boundary of λ , then one of length k_2 , and so on, with the requirement that the remaining part of λ is a Young diagram at each step. Let h_i denote the height of the ribbon tile corresponding to the i th cycle: then the Murnaghan-Nakayama formula states that

$$(21) \quad \chi_\lambda(\pi) = \sum_T \prod_i (-1)^{h_i+1}$$

where the sum is over all consistent tilings T .

Clearly the number of consistent tilings is an upper bound on $|\chi_\lambda(\pi)|$. Now, we claim that for any fixed k , the number of possible locations for a ribbon tile of length k on the boundary of a Young diagram λ of size n is less than $\sqrt{2n}$. To see this, associate each one with the cell of λ which is




FIGURE 1. We associate each possible location for a ribbon tile of fixed length k with a cell (shaded) which is above the tile's lower end and to the left of its upper end. The resulting sequence of cells moves up and to the right at each step, implying that the number of locations is less than $\sqrt{2n}$. Here $k = 3$.

directly above the tile's lower end, and directly to the left of its upper end, as shown in Figure 1. A little thought reveals that the resulting sequence of cells has the property that each one is above and to the right of the previous one. Therefore, if there are ℓ locations, we have

$$n \geq \sum_{i=1}^{\ell} i > \ell^2/2 .$$

and so $\ell < \sqrt{2n}$. It follows that the number of ways to remove the ribbon tiles corresponding to the $c(\pi)$ nontrivial cycles is less than

$$(2n)^{c(\pi)/2} .$$

Moreover, after these ribbon tiles are removed, the number of consistent tilings of the remaining Young diagram is simply the dimension of the corresponding irrep of $S_{n-s(\pi)}$, which is less than $\sqrt{|S_{n-s(\pi)}|} = \sqrt{(n-s(\pi))!}$. Therefore, if λ is typical we have

$$\begin{aligned} \left| \frac{\chi_{\lambda}(\pi)}{d_{\lambda}} \right| &< \frac{(2n)^{c(\pi)/2} \sqrt{(n-s(\pi))!}}{e^{-\sqrt{n} \log n} \sqrt{n!}} \\ &< 2 \cdot e^{\sqrt{n} \log n} 2^{c(\pi)/2} e^{s(\pi)/2} n^{(c(\pi)-s(\pi))/2} \\ &\leq 2 \cdot e^{\sqrt{n} \log n} (\sqrt{2}e)^{t(\pi)} n^{-t(\pi)/2} . \end{aligned}$$

Here we used the bound $(n-s)!/n! < 4 \cdot n^{-s}e^s$, implied by Stirling's approximation, and the facts that $c(\pi) \leq t(\pi)$, $s(\pi) \leq 2t(\pi)$, and $t(\pi) = s(\pi) - c(\pi)$. Finally, if $t(\pi) > \sqrt{n} \log n$, the term $e^{\sqrt{n} \log n}$ can be absorbed into $A^{t(\pi)}$, and Lemma holds for any $A > \sqrt{2}e^2$. \square

Lemma 7 (Rattan and Śniady [RS06]). *For every $D > 0$ there exists a constant A' with the following property. If λ is a Young diagram with n boxes which has at most $D\sqrt{n}$ rows and columns and $\pi \in S_n$ is a permutation then*

$$(22) \quad \left| \frac{\chi_\lambda(\pi)}{d_\lambda} \right| < \left(\frac{A' \max(1, t(\pi)^2/n)}{\sqrt{n}} \right)^{t(\pi)}.$$

Lemma 8. *All typical irreps λ are $O(1)$ -smooth.*

Proof. If λ is typical, then Lemma 6 implies

$$\sum_{\substack{\pi \in S_n \\ t(\pi) > \sqrt{n} \log n}} \left| \frac{\chi_\lambda(\pi)}{d_\lambda} \right|^4 \leq \sum_{\pi \in S_n} (A^{t(\pi)} n^{-t(\pi)/2})^4 = \sum_{\pi \in S_n} z^{t(\pi)}$$

for $z = A^4/n^2$. Since each $\pi \in S_n$ appears exactly once in the product

$$[1 + (12)] [1 + (13) + (23)] \cdots [1 + (1n) + \cdots + (n-1, n)]$$

where (i, j) denotes the transposition interchanging i and j , and since each product of the summands provides a factorization of π into a minimal number of transpositions, we have

$$\begin{aligned} \sum_{\pi \in S_n} z^{t(\pi)} &= (1+z)(1+2z) \cdots (1+(n-1)z) < \\ &e^z e^{2z} \cdots e^{(n-1)z} < e^{zn^2/2} = e^{A^4/2} \end{aligned}$$

therefore

$$(23) \quad \sum_{\substack{\pi \in S_n \\ t(\pi) > \sqrt{n} \log n}} \left| \frac{\chi_\lambda(\pi)}{d_\lambda} \right|^4 < e^{A^4/2}.$$

Very similar but slightly more involved reasoning can be applied to the estimate from Lemma 7 (for details we refer to [RS06]) which shows that there exist constants $E > 0$ and E' (which depend only on D) with a property that if a Young diagram λ with n boxes has at most $D\sqrt{n}$ boxes in each row and column then

$$(24) \quad \sum_{\substack{\pi \in S_n, \\ t(\pi) \leq En^{4/7}}} \left| \frac{\chi_\lambda(\pi)}{d_\lambda} \right|^4 \leq E'.$$

The domains of the summations in the inequalities (23) and (24) cover the whole group S_n for sufficiently large n which finishes the proof. \square

Lemma 9. *There are constants $c' > 0$ and n_0 such that for all pairs of typical irreps λ and μ , if τ is chosen according to the natural distribution $\mathcal{P}_{\lambda \otimes \mu}(\tau)$, then τ is typical with probability at least $1 - e^{-c'\sqrt{n}}$ if $n > n_0$.*

Proof. Let X be the set of atypical representations, and let $\rho = \lambda \otimes \mu$. Then applying Lemma 1 and Lemma 5, using Cauchy-Schwartz as in the proof of Lemma 2, and finally applying Lemma 8 gives

$$\begin{aligned} \mathcal{P}_{\lambda \otimes \rho}(X) &\leq \sqrt{\mathcal{P}_{\text{planch}}(X)} \sqrt{\sum_{g \in G} \left| \frac{\chi_\lambda(g)}{d_\lambda} \right|^2 \left| \frac{\chi_\mu(g)}{d_\mu} \right|^2} \\ &\leq e^{-(c/2)\sqrt{n}} \sqrt[4]{\sum_{g \in G} \left| \frac{\chi_\lambda(g)}{d_\lambda} \right|^4 \sum_{g \in G} \left| \frac{\chi_\mu(g)}{d_\mu} \right|^4} \\ &\leq e^{-(c/2)\sqrt{n}} O(1) \end{aligned}$$

which completes the proof for any $c' < c/2$. \square

7. PROOF OF THE MAIN RESULT

We are now in a position to present our main result.

Theorem 10. *Let \hat{c}, c, c' be the constants defined above. Then for any constants a, b such that $a + b < \min(\hat{c}/2, c, c')$, no sieve algorithm which combines less than $e^{a\sqrt{n}}$ coset states can solve Graph Isomorphism with success probability greater than $e^{-b\sqrt{n}}$.*

Proof. We first consider the behavior of a sieve algorithm A in the case where the hidden subgroup $H \subset S_n \wr \mathbb{Z}_2$ is trivial. For convenience, let us say that a representation $\sigma_{\{\lambda, \mu\}}$ of $S_n \wr \mathbb{Z}_2$ is typical if both λ and μ are. We will establish that with overwhelming probability, all the irrep labels observed by A are both typical and inhomogeneous.

Let ℓ be the number of coset states initially generated by the algorithm. We begin by showing that with high probability, the irrep labels on the ℓ leaves, i.e., those resulting from weak Fourier sampling these coset states, are all both typical and homogeneous. If H is trivial, then these irrep labels are Plancherel-distributed; by (13) the probability that a given one fails to be typical is at most twice the probability that a Plancherel-distributed irrep of S_n fails to be, which by Lemma 5 is at most $e^{-c\sqrt{n}}$. Moreover, by (14) the probability that the label of a given leaf is homogeneous is the probability

that we observe the same irrep of S_n twice in two independent samples of the Plancherel distribution, which using Theorem 3 is

$$\sum_{\lambda} \left(\frac{d_{\lambda}^2}{n!} \right)^2 < \max_{\lambda} \frac{d_{\lambda}^2}{n!} \leq e^{-\hat{c}\sqrt{n}} .$$

Thus the combined probability that any of the ℓ leaves have a label which is not both typical and inhomogeneous is at most

$$(25) \quad \ell \left(2e^{-c\sqrt{n}} + e^{-\hat{c}\sqrt{n}} \right) .$$

Now, assume inductively that all the irreps observed by the algorithm before the i th combine-and-measure step are typical and inhomogeneous, and that the i th step combines states with two such labels $\sigma_{\{\lambda, \lambda'\}}$ and $\sigma_{\{\mu, \mu'\}}$. By (16), the probability this results in a homogeneous irrep is bounded by the probability $\mathcal{P}^{\text{coll}}$ of a collision between a pair of natural distributions in S_n . Then Theorem 3 and Lemmas 2 and 8 imply that this probability is bounded by

$$\mathcal{P}^{\text{coll}} \leq e^{-(\hat{c}/2)\sqrt{n}} O(1) .$$

In addition, Lemma 9 implies that the the probability the observed irrep fails to be typical is at most $e^{-c'\sqrt{n}}$. Since each combine-and-measure step reduces the number of states by one, there are less than ℓ such steps; taking a union bound over all of them, the probability that any of the observed irreps fail to be both homogeneous and typical is

$$(26) \quad \ell \left(e^{-(\hat{c}/2)\sqrt{n}} O(1) + e^{-c'\sqrt{n}} \right) .$$

Let us call a transcript inhomogeneous if all of its irrep labels are. Combining (25) and (26) and setting $\ell < e^{a\sqrt{n}}$, we see that, for n sufficiently large, A 's transcript is inhomogeneous with probability greater than $1 - e^{-b\sqrt{n}}$ for any $b < \min(\hat{c}/2, c, c') - a$.

Now consider A 's behavior in the case of a nontrivial hidden subgroup $H = \{1, m\}$. Inductively applying Equation (12) shows that the probability of observing any inhomogeneous transcript is exactly the same as it would have been if H were trivial. Thus the total variation distance between the distribution of transcripts generated by A in these two cases is less than $e^{-b\sqrt{n}}$, and the theorem is proved. \square

ACKNOWLEDGMENTS

We are very grateful to Philippe Biane, Persi Diaconis and Valentin Féray for helpful conversations about the character theory of the symmetric groups, and to Sally Milius, Tracy Conrad and Rosemary Moore for inspiration. C.M. and A.R. are supported by NSF grants CCR-0220070, EIA-0218563,

and CCF-0524613, and ARO contract W911NF-04-R-0009. P.Š. is supported by the MNiSW research grant 1-P03A-013-30 and by the EC Marie Curie Host Fellowship for the Transfer of Knowledge “Harmonic Analysis, Nonlinear Analysis and Probability,” contract MTKD-CT-2004-013389.

REFERENCES

- [AJL06] Dorit Aharonov, Vaughan Jones, and Zeph Landau. A polynomial quantum algorithm for approximating the Jones polynomial. *Proc. 38th Symposium on Theory of Computing*, 427–436, 2006.
- [AMR06] Gorjan Alagić, Cristopher Moore, and Alexander Russell. Quantum algorithms for Simon’s problem over general groups. *Proc. 18th Symposium on Discrete Algorithms (2007)*.
- [AD95] David Aldous and Persi Diaconis. Hammersley’s interacting particle process and longest increasing subsequences. *Probab. Theory Relat. Fields* **103** (1995), 199–213.
- [Bab80] László Babai. On the complexity of canonical labeling of strongly regular graphs. *SIAM J. Computing*, 9(1):212–216, 1980.
- [Bab83] László Babai and Eugene M. Luks. Canonical labeling of graphs. *Proc. 15th Symposium on Theory of Computing*, 171–183.
- [BCvD05] David Bacon, Andrew Childs, and Wim van Dam. From optimal measurement to efficient quantum algorithms for the hidden subgroup problem over semidirect product groups. *Proc. 46th Symposium on Foundations of Computer Science*, 469–478, 2005.
- [BCvD06] David Bacon, Andrew Childs, and Wim van Dam. Optimal measurements for the dihedral hidden subgroup problem. *Chicago Journal of Theoretical Computer Science*, 2006.
- [Bia98] Philippe Biane. Representations of symmetric groups and free probability. *Advances in Mathematics*, 138(1):126–181, 1998.
- [FIM⁺03] Katalin Friedl, Gábor Ivanyos, Frédéric Magniez, Miklos Santha, and Pranab Sen. Hidden translation and orbit coset in quantum computing. *Proc. 35th Symposium on Theory of Computing*, 1–9, 2003.
- [Ful97] William Fulton. *Young Tableaux: with Applications to Representation Theory and Geometry*, volume 35 of *Student Texts*. London Mathematical Society, 1997.
- [GSVV01] Michelangelo Grigni, Leonard Schulman, Monica Vazirani, and Umesh Vazirani. Quantum mechanical algorithms for the nonabelian hidden subgroup problem. *Proc. 33rd Symposium on Theory of Computing*, 68–74, 2001.
- [Hal02] Sean Hallgren. Polynomial-time quantum algorithms for Pell’s equation and the principal ideal problem. *Proc. 34th Symposium on Theory of Computing*, 653–658.
- [Hal05] Sean Hallgren. Fast quantum algorithms for computing the unit group and class group of a number field. *Proc. 37th Symposium on Theory of Computing*, 468–474, 2005.
- [HMR⁺06] Sean Hallgren, Cristopher Moore, Martin Rötteler, Alexander Russell, and Pranab Sen. Limitations of quantum coset states for graph isomorphism. *Proc. 38th Symposium on Theory of Computing*, 604–617, 2006.

- [HRTS00] Sean Hallgren, Alexander Russell, and Amnon Ta-Shma. Normal subgroup reconstruction and quantum computation using group representations. *Proc. 32nd Symposium on Theory of Computing*, 627–635, 2000.
- [JK81] Gordon James and Adalbert Kerber. *The representation theory of the symmetric group*, volume 16 of *Encyclopedia of mathematics and its applications*. Addison–Wesley, 1981.
- [Ker03] S. V. Kerov. *Asymptotic representation theory of the symmetric group and its applications in analysis*, volume 219 of *Translations of Mathematical Monographs*. American Mathematical Society, 2003. Translated by N. V. Tsilevich.
- [Kup05] Greg Kuperberg. A subexponential-time quantum algorithm for the dihedral hidden subgroup. *SIAM J. Computing* 35(1): 170–188, 2005.
- [MR05a] Christopher Moore and Alexander Russell. Explicit multiregister measurements for hidden subgroup problems; or, Fourier sampling strikes back. Preprint, quant-ph/0504067 (2005).
- [MR05b] Christopher Moore and Alexander Russell. The symmetric group defies strong Fourier sampling: part II. Preprint, quant-ph/0501066 (2005).
- [MR06] Christopher Moore and Alexander Russell. On the impossibility of a quantum sieve algorithm for Graph Isomorphism. Preprint, quant-ph/0609138 (2006).
- [MRS04] Christopher Moore, Alexander Russell, and Leonard Schulman. The power of basis selection in fourier sampling: hidden subgroup problems in affine groups. *Proc. 15th Symposium on Discrete Algorithms*, 1113–1122, 2004.
- [MRS05] Christopher Moore, Alexander Russell, and Leonard Schulman. The symmetric group defies Fourier sampling. *Proc. 46th Symposium on Foundations of Computer Science*, 479–488, 2005.
- [RŚ06] Amarpreet Rattan and Piotr Śniady. Upper bound on the characters of the symmetric groups for balanced Young diagrams and a generalized Frobenius formula. Preprint, math.RT/0610540 (2006).
- [Reg02] Oded Regev. Quantum computation and lattice problems. *Proc. 43rd Symposium on Foundations of Computer Science*, 520–529, 2002.
- [Ser77] Jean-Pierre Serre. *Linear Representations of Finite Groups*. Number 42 in Graduate Texts in Mathematics. Springer-Verlag, 1977.
- [Sho94] Peter W. Shor. Algorithms for quantum computation: discrete logarithms and factoring. *Proc. 35th Symposium on Foundations of Computer Science*, 124–134, 1994.
- [Sim94] D. R. Simon. On the power of quantum computation. *Proc. 35th Symposium on Foundations of Computer Science*, 116–123, 1994.
- [Spi96] Daniel A. Spielman. Faster isomorphism testing of strongly regular graphs. *Proc. 28th Symposium on Theory of Computing*, 576–584, 1996.
- [VK85] 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(1):21–31, 1989.

E-mail address: moore@cs.unm.edu

UNIVERSITY OF NEW MEXICO, DEPARTMENT OF COMPUTER SCIENCE, MAIL STOP:
MSC01 1130, 1 UNIVERSITY OF NEW MEXICO, ALBUQUERQUE, NM 87131-0001,
USA

SANTA FE INSTITUTE, 1399 HYDE PARK ROAD, SANTA FE, NEW MEXICO 87501,
USA

E-mail address: `acr@cse.uconn.edu`

DEPARTMENT OF COMPUTER SCIENCE & ENGINEERING UNIVERSITY OF CON-
NECTICUT 371 FAIRFIELD RD., U-2155 STORRS, CT 06269, USA

E-mail address: `piotr.sniady@math.uni.wroc.pl`

INSTITUTE OF MATHEMATICS, UNIVERSITY OF WROCLAW, PL. GRUNWALDZKI 2/4,
50-384 WROCLAW, POLAND