

On the Power of Random Bases in Fourier Sampling: Hidden Subgroup Problem in the Heisenberg Group

Jaikumar Radhakrishnan *

Martin Rötteler †

Pranab Sen ‡

Abstract

The hidden subgroup problem (HSP) provides a unified framework to study problems of group-theoretical nature in quantum computing such as order finding and the discrete logarithm problem. While it is known that Fourier sampling provides an efficient solution in the abelian case, not much is known for general non-abelian groups. Recently, some authors raised the question as to whether post-processing the Fourier spectrum by measuring in a random orthonormal basis helps for solving the HSP. Several negative results on the shortcomings of this *random strong* method are known. In this paper however, we show that the random strong method can be quite powerful under certain conditions on the group G . We define a parameter $r(G)$ for a group G and show that $O((\log |G|/r(G))^2)$ iterations of the random strong method give enough classical information to identify a hidden subgroup in G . We illustrate the power of the random strong method via a concrete example of the HSP over finite Heisenberg groups. We show that $r(G) = \Omega(1)$ for these groups; hence the HSP can be solved using polynomially many random strong Fourier samplings followed by a possibly exponential classical post-processing without further queries. The quantum part of our algorithm consists of a polynomial computation followed by measuring in a random orthonormal basis. This gives the first example of a group where random representation bases do help in solving the HSP and for which no explicit representation bases are known that solve the problem with $(\log G)^{O(1)}$ Fourier samplings. As an interesting by-product of our work, we get an algorithm for solving the *state identification problem* for a set of nearly orthogonal pure quantum states.

1 Introduction

The hidden subgroup problem (HSP) is defined as follows: We are given a function $f : G \rightarrow S$ from a group G to a set S with the promise that there exists a subgroup $H \leq G$ such that f is constant on the left cosets of H and takes distinct values on distinct cosets. In this paper, all groups and sets are finite and all vector spaces are finite dimensional over \mathbb{C} . The function f is given via a black box, i. e., given $x \in G$ as input, the black box outputs $f(x)$. The task is to find a set of generators for H while making as few queries to f as possible. We would also like our algorithm to be efficient in terms of total running time. The abelian HSP (i. e. G is abelian) encompasses several interesting problems such as finding the order of an element in a group and the discrete logarithm problem. Factoring an integer n can be reduced to order finding in the group \mathbb{Z}_n^* , the multiplicative group of integers modulo n which are coprime to n . The problems of graph isomorphism and graph automorphism can be cast as hidden subgroup problems over the non-abelian group S_n , the group of permutations on n symbols.

*School of Technology and Computer Science, Tata Institute of Fundamental Research, Mumbai, India and Toyota Technological Institute, Chicago, U.S.A. Email: jaikumar@tifr.res.in

†NEC Laboratories America, Inc., Princeton, U.S.A. Email: mroetteler@nec-labs.com

‡Institute for Quantum Computing, University of Waterloo, Canada. Email: p2sen@iqc.ca

The classical query complexity of the HSP is $|G|^{\Omega(1)}$ which is exponential in the input size $\log |G|$. This is true for many families of groups including several families of abelian groups. The biggest success of quantum algorithms so far has been a polynomial time (both query complexity as well as total running time) solution for the abelian HSP [Kit95, BH97, ME98]. The heart of this solution is Fourier sampling with respect to the abelian group G .

In sharp contrast to the abelian HSP, progress on the non-abelian HSP (i.e. G is non-abelian) has so far been quite limited. Ettinger, Høyer and Knill [EHK04] prove that the quantum query complexity of the non-abelian HSP is $O(\log |G|)$; however, their algorithm takes $2^{O(\log^2 |G|)}$ quantum operations. Ivanyos et al. [IMS03] and Friedl et al. [FIM⁺03] apply abelian Fourier transform methods to give polynomial quantum algorithms for the HSP for some special classes of non-abelian groups. Given the success of Fourier sampling in solving the abelian HSP, one can similarly ask whether Fourier sampling over the non-abelian group G helps in solving the HSP over G . The Fourier transform over a (in general, non-abelian) group G gives us a superposition over (ρ, i, j) where ρ is an irreducible unitary representation of G and i, j are the row and column indices of the matrix ρ . The choice of basis for ρ gives us a degree of freedom in defining the Fourier transform over G . This is in contrast to the abelian case, where all representations are one-dimensional and hence only their names ρ matter. The algorithm starts out with a tensor product of $t = O(\log |G|)$ superpositions over random cosets of the hidden subgroup H . Exploiting the symmetries in these states, one can show that (see e.g. [Kup03, Ip03, MRS05]) the optimal measurement to recover H consists of applying the Fourier transform to each coset state, measuring the names of the t irreducible representations, followed by a joint POVM on the column spaces of the resulting t states. In *strong Fourier sampling*, one measures each of the t column spaces using an orthonormal basis, i.e., one performs a tensor product of t complete von Neumann measurements instead of a joint POVM. In *weak Fourier sampling*, one measures the names of the t representations only.

Hallgren, Russell and Ta-Shma [HRTS03] showed that polynomially many iterations of weak Fourier sampling give enough information to reconstruct normal hidden subgroups. More generally, they show that the normal core $c(H)$ of the hidden subgroup H (i.e. the largest normal subgroup of G contained in H) can be reconstructed via the weak method. Grigni, Schulman, Vazirani and Vazirani [GSVV04] and Gavinsky [Gav04] extended the weak method to find a hidden subgroup H in G if $[G : \kappa(G)H] = (\log |G|)^{O(1)}$. Here, $\kappa(G)$ is the Baer subgroup of G defined as $\kappa(G) = \bigcap_{K:K \leq G} N(K)$, where $N(K)$ denotes the normaliser of K in G . The main shortcoming of the weak method is that it gives exactly the same probability distribution if the hidden subgroup is H or a conjugate gHg^{-1} of H . This leads us to consider the *strong* method. The amount of additional information about the hidden subgroup H that can be extracted by measuring the column space in an orthonormal basis depends, in general, on the particular basis. In a recent paper, Moore, Russell and Schulman [MRS05] showed that for the symmetric group S_n , for any choice of bases for the representations, there are order two subgroups that require exponential number of strong Fourier samplings in order to distinguish them from the identity subgroup. Grigni et al. [GSVV04] study the *random strong* method where a random measurement basis is used for each representation ρ . They define a group-theoretic parameter α depending on G and H and show that if α is exponentially large, the additional advantage of the random strong method over the weak method is exponentially small. In particular, this is case when $G = S_n$ and $H \leq S_n$, $|H| = 2^{O(n \log n)}$.

1.1 Our contributions

In this paper, we analyse the power of the random strong method and show, for the first time, that under certain (different) general conditions on G polynomially many iterations of the random strong method do give enough classical information to identify H . We illustrate the power of the random strong method via a

concrete example of the HSP over finite Heisenberg groups \mathcal{H}_p of order p^3 , where $p \geq 3$ is a prime. \mathcal{H}_p is defined as the following set of upper triangular matrices:

$$\mathcal{H}_p := \left\{ \begin{pmatrix} 1 & x & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{pmatrix} : x, y, z \in \mathbb{F}_p \right\}. \quad (1)$$

A convenient encoding for the elements of \mathcal{H}_p is to write (x, y, z) , where $x, y, z \in \mathbb{F}_p$ match the components in equation (1). The composition of two elements is then given by

$$(x_1, y_1, z_1) \circ (x_2, y_2, z_2) = (x_1 + x_2, y_1 + y_2, z_1 + z_2 + x_1 y_2),$$

and the inverse of an element is given by $(x, y, z)^{-1} = (-x, -y, xy - z)$. It is easy to see that the classical randomised query complexity of the HSP on \mathcal{H}_p is $\theta(p)$. The generic quantum algorithm of Ettinger, Høyer and Knill [EHK04] achieves $O(\log p)$ query complexity, but at the expense of $p^{O(\log p)}$ quantum operations. An algorithm with $2^{\theta(\sqrt{\log p})}$ quantum operations can be obtained by combining the ideas of [Kup03] and [FIM⁺03]. However, the query complexity of this algorithm is also $2^{\theta(\sqrt{\log p})}$. It seems non-trivial to design a quantum algorithm with $(\log p)^{O(1)}$ query complexity and total running time $p^{O(1)}$. In the following paragraphs, we indicate how various existing methods for non-abelian HSP fail to achieve this goal. After that, we show how the random strong method attains this goal, illustrating the power of random bases in Fourier sampling.

It can be shown that \mathcal{H}_p is a semidirect product of the form $\mathbb{Z}_p \ltimes (\mathbb{Z}_p \times \mathbb{Z}_p)$, where the normal subgroup is given by $N_\infty := \{(0, y, z) : y, z \in \mathbb{F}_p\}$ and the complement by $A_{0,0} := \{(x, 0, 0) : x \in \mathbb{F}_p\}$. The commutator subgroup of \mathcal{H}_p is given by $[\mathcal{H}_p, \mathcal{H}_p] = \{(0, 0, z) : z \in \mathbb{F}_p\}$, which is also the centre $\zeta(\mathcal{H}_p)$. The commutator subgroup is isomorphic to \mathbb{Z}_p ; hence it is abelian but not *smoothly abelian* (an abelian group G is said to be smoothly abelian [FIM⁺03] if it is the direct product of a subgroup of bounded exponent and a subgroup of size $(\log |G|)^{O(1)}$). The Baer subgroup turns out to be $\kappa(\mathcal{H}_p) = \zeta(\mathcal{H}_p)$. If $A \leq \mathcal{H}_p$, $|A| = p$, then $|\kappa(\mathcal{H}_p)A| \leq p^2$; therefore for such an A , $[G : \kappa(\mathcal{H}_p)A] \geq p$. In fact, we will see later that there are $(p^2 + p + 1)$ order p subgroups of \mathcal{H}_p . Thus, the methods of [GSVV04, Gav04, IMS03, FIM⁺03] are not applicable in order to solve the HSP for \mathcal{H}_p efficiently. For more details about the Heisenberg group, see Section 2.

The chief obstacle to finding hidden subgroups in \mathcal{H}_p arises from the order p subgroups of \mathcal{H}_p other than its centre. There are $(p^2 + p)$ such order p subgroups; we shall call them $A_{i,j}$, $i \in \mathbb{F}_p \cup \{\infty\}$, $j \in \mathbb{F}_p$. The *forgetful abelian method* (i. e. Fourier sampling over the abelian group $\mathbb{Z}_p \times (\mathbb{Z}_p \times \mathbb{Z}_p)$ instead of the non-abelian group $\mathcal{H}_p \cong \mathbb{Z}_p \ltimes (\mathbb{Z}_p \times \mathbb{Z}_p)$), weak Fourier sampling, strong Fourier sampling in the natural representation basis of \mathcal{H}_p (i. e. the representation basis adapted to the distinguished subgroup tower $\{1\} \triangleleft N_\infty \triangleleft \mathcal{H}_p$) as well as strong Fourier sampling in the \mathbb{Z}_p -Fourier transform of the natural representation basis give exponentially small information about the index i of $A_{i,j}$. For more details, see Section 2.4. For now, we give an intuitive description of the main difficulty posed by these subgroups. Suppose the hidden subgroup is $A_{i,j}$ for some $i \in \mathbb{F}_p \cup \{\infty\}$, $j \in \mathbb{F}_p$. With exponentially high probability, Fourier sampling over \mathcal{H}_p gives us a representation uniformly at random from one of the $(p-1)$ irreducible representations ρ_k of degree p for $k = 1, \dots, p-1$ of \mathcal{H}_p . Suppose one such representation ρ_k shows up. The state essentially collapses to a vector $|\psi_{k,i,j}\rangle \in \mathbb{C}^p$, i. e., $(\mathcal{H}_p, A_{i,j})$ is a Gelfand pair for all i, j (see also [MR05] for Gelfand pairs in the context of the HSP). The vectors $|\psi_{k,i,j}\rangle$ have the property that

$$|\langle \psi_{k,i,j} | \psi_{k,i',j'} \rangle| = \begin{cases} \frac{1}{\sqrt{p}} & : i \neq i', \text{ for all } j, j', \\ \delta_{j,j'} & : i = i', \end{cases}$$

i. e., they form a set of $(p + 1)$ mutually unbiased bases [WF89] of \mathbb{C}^p . The main difficulty is that it is not clear a priori that there is any orthonormal basis that can pairwise distinguish between these $(p^2 + p)$ vectors with inverse polynomial probability. Note that the so-called *hidden conjugate problem* [MRRS04] is easy to solve information-theoretically for \mathcal{H}_p ; the conjugacy classes of the order p subgroups are defined by i and the above property says that $\{|\psi_{k,i,j}\rangle\}_j$ is an orthonormal basis of \mathbb{C}^p , so given the conjugacy class i one can measure in this orthonormal basis to determine the actual hidden subgroup $A_{i,j}$. In view of this, the main challenge in solving the HSP for \mathcal{H}_p is to identify the conjugacy class i .

In this paper however, we show that a random representation basis for ρ_k does in fact pairwise distinguish between $|\psi_{k,i,j}\rangle$ with constant probability. In fact, we refine the method of random measurement bases to distinguish between families of nearly orthogonal subspaces. We combine the geometric ideas of random measurement bases together with representation-theoretic techniques to obtain a parameter $r(G; H_1, H_2)$ of a group G and subgroups $H_1, H_2 \leq G$. We show that $r(G; H_1, H_2)$ is a lower bound on the total variation distance between the distributions on pairs (ρ, j) of representation names and column indices obtained by the random strong method for candidate hidden subgroups H_1, H_2 . The parameter $r(G; H_1, H_2)$ is defined in terms of the ranks and overlaps of the projectors obtained by averaging representations ρ over H_1, H_2 . Define $r(G) := \min_{H_1, H_2} r(G; H_1, H_2)$, where H_1, H_2 range over all pairs of subgroups of G . We show that $O\left(\frac{\log s(G)}{r^2(G)}\right)$ iterations of the random strong method give sufficient classical information to identify the hidden subgroup H , where $s(G)$ denotes the number of distinct subgroups of G . Note that $s(G) \leq 2^{\log^2 |G|}$ for any group G .

We will see later in Section 2 that $s(\mathcal{H}_p) = O(p^2)$. In Section 4, we show that $r(\mathcal{H}_p) = \Omega(1)$, implying that $O(\log p)$ iterations of the random strong method give sufficient information to extract the hidden subgroup in \mathcal{H}_p . This gives us an algorithm solving the HSP over \mathcal{H}_p with $O(\log p)$ query complexity, $O(\log^3 p)$ quantum operations for implementing the non-abelian Fourier transforms (see Section 2.5), $\tilde{O}(p^2)$ quantum operations to measure in a random basis, and $\tilde{O}(p^4)$ classical post-processing operations. This gives the first example of a group where random representation bases do help in solving the HSP and for which no explicit representation bases are known that solve the problem with $(\log p)^{O(1)}$ Fourier samplings.

As an interesting by-product of our work, we get an algorithm for solving the following *quantum state identification* problem: Consider a set of pure quantum states $\{|\psi_1\rangle, \dots, |\psi_m\rangle\} \in \mathbb{C}^n$ with the property that $|\langle\psi_i|\psi_j\rangle| \leq \delta$ for all $i \neq j$, where δ is a sufficiently small constant (and typically $m \gg n$). We are given t independent copies of $|\psi_i\rangle$. The task is to identify the index i . We show that $t = O(\log m)$ independent random complete von Neumann measurements in \mathbb{C}^n suffice to identify i with high probability.

1.2 Relation to other work

Moore, Rockmore, Russell and Schulman [MRRS04] use non-abelian strong Fourier sampling to give an efficient algorithm for the HSP over the q -hedral group $\mathbb{Z}_q \times \mathbb{Z}_p$ when p, q are prime, $q \mid (p - 1)$ and $(p - 1)/q = (\log p)^{O(1)}$. Our techniques show that for p, q prime, $q \mid (p - 1)$, $q = \Omega(\sqrt{p})$, $r(\mathbb{Z}_q \times \mathbb{Z}_p) = \Omega(1)$, which proves that polynomially many random strong Fourier samplings suffice to find an arbitrary hidden subgroup of $\mathbb{Z}_q \times \mathbb{Z}_p$ in this case. For prime $p, q \mid (p - 1)$, $q = \Omega(p^{3/4})$, subgroups H_1, H_2 conjugate to $\mathbb{Z}_q \leq \mathbb{Z}_{p-1}$, our techniques show that $r(\mathbb{Z}_{p-1} \times \mathbb{Z}_p; H_1, H_2) = \Omega\left(\sqrt{\frac{q}{p}}\right)$. Moore et al. [MRRS04] prove a nearly matching upper bound of $r(\mathbb{Z}_{p-1} \times \mathbb{Z}_p; H_1, H_2) = O\left(\sqrt{\frac{q}{p} \log p}\right)$. Thus, a polynomial amount of random strong Fourier sampling can solve the hidden conjugate problem for subgroup $\mathbb{Z}_q \leq \mathbb{Z}_{p-1}$ of the affine group $\mathbb{Z}_{p-1} \times \mathbb{Z}_p$ if and only if $p/q = (\log p)^{O(1)}$.

In this paper, we confine ourselves to random strong Fourier sampling. Our quantum operations always

factor into a tensor product over the coset states obtained by querying the function oracle. This distinguishes the Heisenberg group from the symmetric group for which Moore, Russell and Schulman [MRS05] show that tensor product Fourier sampling is not sufficient to solve the HSP. The quantum part of our algorithm consists of a polynomial computation followed by measuring in a random orthonormal basis. In fact, if a suitable kind of pseudo-random unitary transformation can be generated and implemented efficiently, then the quantum part of the algorithm can be made fully polynomial. Various notions of pseudo-random unitary transformations have been studied (see e.g. [EWS⁺03, Eme04]), but it has to be investigated whether they are sufficient for our purposes.

2 Heisenberg groups over \mathbb{F}_p

The groups \mathcal{H}_p , where $p \geq 3$ is prime, are discrete versions of the continuous Heisenberg groups studied in physics in the context of conjugate observables. Abstractly, \mathcal{H}_p is isomorphic to the following group given in terms of generators and relations: $\mathcal{H}_p \cong \langle x, y, z : x^p = y^p = z^p = 1, xy = zy, xz = zx, zy = yz \rangle$.

2.1 The subgroup lattice

Since the order of \mathcal{H}_p is p^3 we can expect to find subgroups of order p and p^2 besides the trivial subgroup $\{1\}$ and \mathcal{H}_p . The centre of \mathcal{H}_p is given by

$$\zeta(\mathcal{H}_p) = \langle (0, 0, 1) \rangle = \{(0, 0, z) : z \in \mathbb{F}_p\}.$$

Note that $|\zeta(\mathcal{H}_p)| = p$. There are $p + 1$ subgroups N_i of order p^2 , where $i \in \mathbb{F}_p \cup \{\infty\}$. They are given by

$$N_i := \langle (1, i, 0), (0, 0, 1) \rangle = \{(x, xi, z) : x, z \in \mathbb{F}_p\}, \quad \forall i \in \mathbb{F}_p.$$

The group N_∞ is given by $N_\infty := \langle (0, 1, 0), (0, 0, 1) \rangle = \{(0, y, z) : y, z \in \mathbb{F}_p\}$; $N_\infty \cong \mathbb{Z}_p \times \mathbb{Z}_p$. It is easy to see that for all $i \in \mathbb{F}_p \cup \{\infty\}$, $\zeta(\mathcal{H}_p) \triangleleft N_i$. Furthermore the N_i are normal subgroups, $N_i \triangleleft \mathcal{H}_p$ and $N_i \cong \mathbb{Z}_p \times \mathbb{Z}_p$. For each $i \in \mathbb{F}_p \cup \{\infty\}$, we have that N_i contains p subgroups $A_{i,j}$ for $j \in \mathbb{F}_p$. The subgroups $A_{i,j}$ satisfy $|A_{i,j}| = p$, whence $A_{i,j} \cong \mathbb{Z}_p$. For $i, j \in \mathbb{F}_p$ we have the following explicit description of the elements of $A_{i,j}$:

$$A_{i,j} := \langle (1, i, j) \rangle = \left\{ \left(\mu, \mu i, \binom{\mu}{2} i + \mu j \right) : \mu \in \mathbb{F}_p \right\}.$$

For $i = \infty, j \in \mathbb{F}_p$ we obtain $A_{\infty,j} := \langle (0, 1, j) \rangle = \{(0, \mu, \mu j) : \mu \in \mathbb{F}_p\}$. It is easy to see that $A_{i,j} \not\triangleleft N_{i'}$ if $i \neq i'$, and the normaliser is given by $N_{\mathcal{H}_p}(A_{i,j}) = N_i$. The above groups form a complete list of distinct subgroups of \mathcal{H}_p . The following table summarizes the subgroup structure of \mathcal{H}_p .

Size	Subgroup	Number	Containment
p^3	\mathcal{H}_p	1	
p^2	$N_i, i \in \mathbb{F}_p \cup \{\infty\}$	$p + 1$	$N_i \triangleleft \mathcal{H}_p$
p	$\zeta(\mathcal{H}_p), A_{i,j}, i \in \mathbb{F}_p \cup \{\infty\}, j \in \mathbb{F}_p$	$p^2 + p + 1$	$A_{i,j} \triangleleft N_i, \zeta(\mathcal{H}_p) \triangleleft N_k, \forall k \in \mathbb{F}_p \cup \{\infty\}$
1	$\{1\}$	1	$\{1\} \triangleleft \zeta(\mathcal{H}_p), \{1\} \triangleleft A_{i,j}$

For $i, i' \in \mathbb{F}_p \cup \{\infty\}$ where $i \neq i'$ we have that $N_i \cap N_{i'} = \zeta(\mathcal{H}_p)$. This shows that $\kappa(\mathcal{H}_p) = \bigcap_{K:K \leq \mathcal{H}_p} N(K) = \zeta(\mathcal{H}_p)$. Also, it is easy to check that the commutator subgroup is given by $[\mathcal{H}_p, \mathcal{H}_p] = \zeta(\mathcal{H}_p)$.

2.2 The irreducible representations of \mathcal{H}_p

Since we want to perform Fourier analysis on the groups \mathcal{H}_p we have to determine the irreducible representations of \mathcal{H}_p . The reader not familiar with the standard notations of representation theory is referred to standard references like [CR62] or [Ser77]. Observe that $\mathcal{H}_p = A_{0,0} \times N_\infty \cong \mathbb{Z}_p \times (\mathbb{Z}_p \times \mathbb{Z}_p)$. This semidirect product structure can be used to construct the irreducible representations of \mathcal{H}_p . First, there are p^2 one-dimensional representations $\chi_{a,b}$ for $a, b \in \mathbb{F}_p$ which come from the factor group $\mathcal{H}_p/\zeta(\mathcal{H}_p) \cong \mathbb{Z}_p^2$. In the following, let ω denote a fixed p th root of unity in the complex numbers. Then the one-dimensional irreducible representations of \mathcal{H}_p are given by

$$\chi_{a,b}((x, y, z)) := \omega^{ax+by} \quad a, b \in \mathbb{F}_p.$$

Let \mathbb{F}_p^* denote the group of non-zero elements of \mathbb{F}_p under multiplication. There are $p-1$ irreducible representations ρ_k , $k \in \mathbb{F}_p^*$ of degree p . They are obtained in the following way: Take a nontrivial character of the centre $\zeta(\mathcal{H}_p)$, extend it to the abelian group N_∞ , and induce it to \mathcal{H}_p . Explicitly, we obtain the following representations: For each $k \in \mathbb{F}_p^*$, we have a nontrivial character ϕ_k of $\zeta(\mathcal{H}_p)$ given by $\phi_k((0, 0, z)) := \omega^{kz}$. Since $\zeta(\mathcal{H}_p) \triangleleft N_\infty$ and N_∞ is abelian, we can extend ϕ_k to a character $\bar{\phi}_k$ of N_∞ by simply defining $\bar{\phi}_k((0, y, 0)) := 1$. We choose the elements of $A_{0,0}$ as transversals for N_∞ in \mathcal{H}_p . Then ρ_k is defined to be the induction $\rho_k := \bar{\phi}_k \uparrow_{A_{0,0}} \mathcal{H}_p$. On the generators of \mathcal{H}_p , we find that ρ_k takes the following values: $\rho_k((1, 0, 0)) = \sum_{a \in \mathbb{F}_p} |a\rangle\langle a+1|$, $\rho_k((0, 1, 0)) = \sum_{a \in \mathbb{F}_p} \omega^{ka} |a\rangle\langle a|$ and $\rho_k((0, 0, 1)) = \omega^k \mathbf{1}_p$, where $\mathbf{1}_p$ denotes the identity operator in \mathbb{C}^p . Since $(x, y, z) = (0, 0, z)(0, y, 0)(x, 0, 0)$ for all $x, y, z \in \mathbb{F}_p$, we obtain that

$$\rho_k((x, y, z)) = \omega^{kz} \sum_{a \in \mathbb{F}_p} \omega^{kya} |a\rangle\langle a+x|.$$

It can be readily checked that the $\chi_{a,b}$, for $a, b \in \mathbb{F}_p$ and ρ_k , for $k \in \mathbb{F}_p^*$ form a complete set of inequivalent irreducible representations of \mathcal{H}_p .

2.3 Ranks and overlaps of various projectors

Define $P_{k;i,j} := \frac{1}{p} \sum_{a \in A_{i,j}} \rho_k(a)$. It is easy to see that $P_{k;i,j}$ is an orthogonal projection operator. In order to calculate the parameter $r(\mathcal{H}_p)$ (see Section 4 for the details of the calculation) we have to compute the ranks of $P_{k;i,j}$ and pairwise *overlaps* $\|P_{k;i,j} P_{k;i',j'}\|$ (the reason for the nomenclature of *overlap* will be made clear in Section 3.1). For $i, j \in \mathbb{F}_p$, we obtain by a straightforward computation that $P_{k;i,j} = \frac{1}{p} \sum_{\mu, \nu \in \mathbb{F}_p} \omega_p^{k \left(\binom{\mu}{2} i + \mu j - \binom{\nu}{2} i - \nu j \right)} |\nu\rangle\langle \mu|$. Hence, $P_{k;i,j} = |\psi_{k;i,j}\rangle\langle \psi_{k;i,j}|$, where

$$|\psi_{k;i,j}\rangle = \frac{1}{\sqrt{p}} \sum_{\mu \in \mathbb{F}_p} \omega^{-k \left(\binom{\mu}{2} i + \mu j \right)} |\mu\rangle, \quad i, j \in \mathbb{F}_p, k \in \mathbb{F}_p^*.$$

In the case $i = \infty, j \in \mathbb{F}_p$, we get $P_{k;\infty,j} = |\psi_{k;\infty,j}\rangle\langle \psi_{k;\infty,j}|$, where $|\psi_{k;\infty,j}\rangle = |-j\rangle$ $j \in \mathbb{F}_p, k \in \mathbb{F}_p^*$. Thus for all $k \in \mathbb{F}_p^*, i \in \mathbb{F}_p \cup \{\infty\}, j \in \mathbb{F}_p$, $\text{rank}(P_{k;i,j}) = 1$ and $P_{k;i,j}$ is an orthogonal projection onto $|\psi_{k;i,j}\rangle$. For $j, j' \in \mathbb{F}_p$, we get $\|P_{k;\infty,j} P_{k;\infty,j'}\| = \delta_{j,j'}$. For $i, i', j' \in \mathbb{F}_p$, we get $\|P_{k;i,j} P_{k;\infty,j'}\| = \frac{1}{\sqrt{p}}$. For $i, i', j, j' \in \mathbb{F}_p$, we get

$$\|P_{k;i,j} P_{k;i',j'}\| = |\langle \psi_{k;i,j} | \psi_{k;i',j'} \rangle| = \frac{1}{p} \sum_{\mu \in \mathbb{F}_p} \omega^{k \left(\binom{\mu}{2} (i-i') + \mu (j-j') \right)}.$$

To evaluate the last term above, we need the following fact about quadratic Weil sums in \mathbb{F}_p .

Fact 1 ([LN94, Theorem 5.37]) Let $h(X) \in \mathbb{F}_p[X]$ be a degree two polynomial. Then, $\left| \sum_{x \in \mathbb{F}_p} \omega^{h(x)} \right| = \sqrt{p}$.

By Fact 1, if $i \neq i'$, $|\langle \psi_{k;i,j} | \psi_{k;i',j'} \rangle| = \frac{1}{\sqrt{p}}$ irrespective of j and j' . If $i = i'$, it is easy to see that $|\langle \psi_{k;i,j} | \psi_{k;i',j'} \rangle| = \delta_{j,j'}$. To summarise, we have shown the following result:

Lemma 1 Suppose p is an odd prime. Let $i, i' \in \mathbb{F}_p \cup \{\infty\}$, $j, j' \in \mathbb{F}_p$ and $A_{i,j}, A_{i',j'}$ be two order p subgroups of \mathcal{H}_p other than the centre $\zeta(\mathcal{H}_p)$. Let ρ_k , where $k \in \mathbb{F}_p^*$, be an irreducible representation of \mathcal{H}_p of degree p . Let $P_{k;i,j}$ be defined by $P_{k;i,j} := \frac{1}{p} \sum_{a \in A_{i,j}} \rho_k(a)$ and let $P_{k;i',j'}$ be defined similarly. Then $P_{k;i,j}, P_{k;i',j'}$ are rank one orthogonal projections, and their overlap is given by

$$\|P_{k;i,j} P_{k;i',j'}\| = \begin{cases} \frac{1}{\sqrt{p}} & : i \neq i', \text{ for all } j, j', \\ \delta_{j,j'} & : i = i'. \end{cases}$$

Thus, for any $k \in \mathbb{F}_p^*$, the vectors $|\psi_{k;i',j'}\rangle$ form a set of $(p+1)$ mutually unbiased bases for \mathbb{C}^p .

2.4 Failure of existing methods to solve the HSP over \mathcal{H}_p

A straightforward classical randomised algorithm for the HSP over \mathcal{H}_p is as follows: Query $f : \mathcal{H}_p \rightarrow S$ at $O(p)$ random elements of \mathcal{H}_p . If we do not find $a_1, a_2 \in \mathcal{H}_p$, $a_1 \neq a_2$ such that $f(a_1) = f(a_2)$, we declare $\{1\}$ to be the HSP of f . Suppose we do find such a pair a_1, a_2 . Then there is a unique order p subgroup A of \mathcal{H}_p such that $a_1^{-1}a_2 \in A$. f can now be thought of as a function on \mathcal{H}_p/A . Query f at $O(\sqrt{p})$ random elements of \mathcal{H}_p/A . If we do not find $b_1, b_2 \in \mathcal{H}_p/A$, $b_1 \neq b_2$ such that $f(b_1) = f(b_2)$, we declare A to be the HSP of f . Suppose we do find such a pair b_1, b_2 . Let $B = \langle A, b_1^{-1}b_2 \rangle$. If $|B| = p^3$, declare the HSP to be \mathcal{H}_p . If $|B| = p^2$, query f at an element $c \in \mathcal{H}_p$, $c \notin B$. If $f(c) = f(B)$, declare the HSP to be \mathcal{H}_p , else declare the HSP to be B . The correctness of the algorithm follows from the subgroup structure of \mathcal{H}_p and the birthday paradox. A matching lower bound of $\Omega(p)$ for classical randomised algorithms can be proved using the subgroup structure of \mathcal{H}_p and Yao's minimax principle.

Suppose the HSP is $A_{i,j}$, for some $i \in \mathbb{F}_p^*$, $j \in \mathbb{F}_p$. It can be shown (see Section 3.3 for details) that Fourier sampling gives a p -dimensional representation with probability $1 - \frac{1}{p}$, and each p -dimensional representation has equal probability to show up. Suppose one such representation ρ_k , $k \in \mathbb{F}_p^*$ shows up. The natural representation basis $|a\rangle$, $a \in \mathbb{F}_p$ is the basis $|\psi_{k;\infty,j}\rangle$, where $j \in \mathbb{F}_p$. The \mathbb{Z}_p -Fourier transform of the natural representation basis is the basis $|\psi_{k;0,j}\rangle$, where $j \in \mathbb{F}_p$. By Lemma 1, the probability distribution obtained by measuring the columns of ρ_k in the natural representation basis or in the \mathbb{Z}_p -Fourier transform of the natural representation basis is the uniform distribution. This shows that weak Fourier sampling, strong Fourier sampling in the natural representation basis of \mathcal{H}_p as well as strong Fourier sampling in the \mathbb{Z}_p -Fourier transform of the natural representation basis give exponentially small information about the index i of $A_{i,j}$.

Recall that $\mathcal{H}_p = A_{0,0} \times N_\infty \cong \mathbb{Z}_p \times (\mathbb{Z}_p \times \mathbb{Z}_p)$. Suppose we try to perform Fourier sampling over the abelian group $A_{0,0} \times N_\infty \cong \mathbb{Z}_p \times \mathbb{Z}_p \times \mathbb{Z}_p$ (the *forgetful abelian method*) instead of the non-abelian group \mathcal{H}_p . Let F denote the Fourier transform over $\mathbb{Z}_p \times \mathbb{Z}_p \times \mathbb{Z}_p$, i. e.,

$$F = p^{-3/2} \sum_{a,b,c,x,y,z \in \mathbb{F}_p} \omega^{ax+by+cz} |a,b,c\rangle \langle x,y,z|.$$

For abelian groups G , the probability distributions obtained by Fourier sampling over G are independent of the actual coset of the hidden subgroup that arises on measuring the function value; however, they depend

of course on the hidden subgroup. But since now we are doing abelian Fourier sampling over a non-abelian group, we have to consider the effect of applying F to a coset of $tA_{i,j}$, where $t = (0, 0, \tau)$ and $\tau \in \mathbb{F}_p$. Note that $tA_{i,j} = \{(\mu, \mu i, \binom{\mu}{2}i + \mu j + \tau) : \mu \in \mathbb{F}_p\}$. We obtain

$$F|tA_{i,j}\rangle = \frac{1}{p^2} \sum_{a,b,c \in \mathbb{F}_p} \omega^{c\tau} \left(\sum_{\mu \in \mathbb{F}_p} \omega^{(a+bi+cj-\frac{ci}{2})\mu+ci\mu^2} \right) |a, b, c\rangle.$$

Hence, the probability of observing a particular triple (a, b, c) is $p^{-4} |\sum_{\mu \in \mathbb{F}_p} \omega^{(a+bi+cj-(ci)/2)\mu+ci\mu^2}|^2$. If $c \neq 0$, this is a quadratic Weil sum and we can use Fact 1 to conclude that the probability of observing (a, b, c) is given by p^{-3} , independent of i, j . The probability of observing (a, b, c) , $c \neq 0$ is $1 - \frac{1}{p}$. If $c = 0$, only terms of the form $(-bi, b, 0)$ show up. These terms do give information about i ; however, the probability of observing such a term is $\frac{1}{p}$. Thus, the forgetful abelian method gives exponentially small information about i .

2.5 Efficient quantum circuits for the Fourier Transform on \mathcal{H}_p

The fact that any QFT for any finite group is a unitary matrix (when properly normalized) makes this class of transformations an important source of transformations a quantum computer can carry out. The problem of finding efficient implementations of QFTs in terms of quantum circuits was studied previously, see [Hø97, Bea97, PRB99, HRTS03, MRR04]. From [MRR04, Theorem 2] it follows that for any prime p the QFT for the Heisenberg group \mathcal{H}_p can be computed in $\text{polylog}(p)$ operations. In the following we give an explicit description of an efficient quantum circuit which computes $\text{QFT}_{\mathcal{H}_p}$. First, note that we are interested in a realization on a quantum computer which works on qubits. This means that we have to embed the states and transformations into a register of size 2^n for some positive integer n . In the following we will assume that n is the smallest integer such that $p < 2^n$ and we will identify the group elements $(x, y, z) \in \mathcal{H}_p$ with a subset of the binary strings of length $3n$: in each of the three components we choose the basis vectors $|0\rangle, \dots, |p-1\rangle$ to represent the respective component of the element (x, y, z) . The following proposition shows that a QFT for \mathcal{H}_p can be implemented efficiently in terms of elementary quantum gates.

Proposition 1 *Let p be prime, let \mathcal{H}_p be the Heisenberg group of order p^3 and let $\text{Irr}(\mathcal{H}_p) = \{\chi_{a,b} : (a, b) \in \mathbb{F}_p^2\} \cup \{\rho_k : k = 1, \dots, p-1\}$ denote the irreducible representations of \mathcal{H}_p . Then the QFT for \mathcal{H}_p with respect to $\text{Irr}(\mathcal{H}_p)$ can be computed using $O(\log^3 p)$ elementary quantum gates.*

Proof: First we consider the normal subgroup $N_\infty \triangleleft \mathcal{H}_p$ and compute a Fourier transform for this abelian group. This group is isomorphic to a direct product of two cyclic groups, i. e., $N_\infty \cong \mathbb{Z}_p \times \mathbb{Z}_p$. The elements of N_∞ are given by $N_\infty = \{(0, y, z) : y, z \in \mathbb{Z}_p\}$, i. e., we can identify the elements of N_∞ with those binary strings of length $3n$ which have trivial support on the first n positions. Note that the irreducible representations of N_∞ are given by $\psi_{a,b}$ for $a, b \in \mathbb{Z}_p$, where

$$\psi_{a,b}(0, y, z) := \exp(2\pi i/p(ay + bz)) = \omega_p^{ay+bz}.$$

Since N_∞ is normal the group \mathcal{H}_p operates on the irreducible representations [CR62]. We denote this action by “*”, i. e., we have a map $*$: $\mathcal{H}_p \times \text{Irr}(\mathcal{H}_p) \rightarrow \text{Irr}(\mathcal{H}_p)$ which is explicitly given by $(x, y, z) * \psi_{a,b} = \psi_{a,b-ax}$.

Next, we choose as a transversal for $N_\infty \triangleleft \mathcal{H}_p$ the ordered list $T = [(x, 0, 0) : x \in \mathbb{Z}_p]$. We have to be able to efficiently implement the images of all irreducible representations of \mathcal{H}_p evaluated at the

elements of T . This is required for the so-called ‘twiddle factors’ in the decomposition of $\text{QFT}_{\mathcal{H}_p}$ along the subgroup tower $\{1\} \triangleleft N_\infty \triangleleft \mathcal{H}_p$. Indeed, we construct a QFT adapted to this subgroup tower, see also [PRB99, MRR04]. We now use the following formula for implementing a QFT_G which holds in the situation where we have an abelian normal subgroup N and an abelian factor group G/N :

$$\text{QFT}_{\mathcal{H}_p} = (\mathbb{1}_{|G/N|} \otimes \text{QFT}_N) \left(\bigoplus_{t \in T} \Phi(t) \right) \left(\text{QFT}_{G/N} \otimes \mathbb{1}_{|N|} \right).$$

Here Φ denotes an extension of the decomposition of the regular representation of N into irreducibles. Denoting this direct sum by Λ , i. e., $\Lambda := \bigoplus_{t \in T} \Phi(t)$, this means that we have to implement the following transformation:

$$\Lambda : |x\rangle|a\rangle|b\rangle \mapsto \begin{cases} |x\rangle|a\rangle|b - ax\rangle & : \text{ if } a \neq 0, \\ \omega_p^{xb} |x\rangle|0\rangle|b\rangle & : \text{ if } a = 0. \end{cases}$$

It is straightforward to implement Λ using classical efficient circuits for modular addition and multiplication. Hence Λ can be implemented using $O(\log^3(p))$ quantum gates. Note that $\text{QFT}_{N_\infty} = \text{QFT}_{\mathbb{Z}_p} \otimes \text{QFT}_{\mathbb{Z}_p}$ and $\text{QFT}_{G/N_\infty} = \text{QFT}_{\mathbb{Z}_p}$, both of which can be either implemented approximately [Kit95] or exactly [MZ04] on a quantum computer using $O(\log^2 p)$ many elementary quantum gates. Hence the claimed complexity for computing a quantum Fourier transform for \mathcal{H}_p follows. \blacksquare

3 Random bases and Fourier sampling

3.1 Nearly orthogonal vectors

In this subsection, we state some results about sets of nearly orthogonal unit vectors in a Hilbert space. We use $\|\cdot\|$ to denote the ℓ_2 -norm of vectors as well as the ℓ_2 -induced operator norm of matrices. We use $\|v\|_1$ to denote the ℓ_1 -norm of a vector v . We let $\|M\|_{\text{tr}} = \text{Tr} \sqrt{M^\dagger M}$ denote the trace norm of a matrix M . For subspaces V_1, V_2 having trivial intersection, their *overlap* is defined as $\text{ovlap}(V_1, V_2) = \max_{v_1, v_2} |\langle v_1 | v_2 \rangle|$, where v_i range over unit vectors in V_i . Let Π_i denote the orthogonal projection operator onto V_i . It is easy to see that $\text{ovlap}(V_1, V_2) = \|\Pi_1 \Pi_2\|$.

Proposition 2 *Let V_1, V_2 be subspaces of a Hilbert space having trivial intersection. Let σ_2 denote the totally mixed state in V_2 . Let V'_2 denote the orthogonal complement of V_1 in $V_1 + V_2$ and σ'_2 denote the totally mixed state in V'_2 . Let $\delta = \text{ovlap}(V_1, V_2)$. Then,*

$$\|\sigma_2 - \sigma'_2\|_{\text{tr}} \leq 2\delta^{1/2}(1 - \delta^2)^{-1/4}.$$

Proof: Let $d = \dim V_2$ and a_1, \dots, a_d be an orthonormal basis for V_2 . Let a'_1, \dots, a'_d be the Gram-Schmidt orthonormalisation of a_1, \dots, a_d with respect to V_1 . Hence, a'_1, \dots, a'_d is an orthonormal basis for V'_2 . We will show that $\| |a_i\rangle\langle a_i| - |a'_i\rangle\langle a'_i| \|_{\text{tr}} \leq 2\delta^{1/2}(1 - \delta^2)^{-1/4}$ for all $1 \leq i \leq d$. Since $\sigma_2 = \frac{1}{d} \sum_{i=1}^d |a_i\rangle\langle a_i|$ and $\sigma'_2 = \frac{1}{d} \sum_{i=1}^d |a'_i\rangle\langle a'_i|$, we will get

$$\|\sigma_2 - \sigma'_2\|_{\text{tr}} \leq \frac{1}{d} \sum_{i=1}^d \| |a_i\rangle\langle a_i| - |a'_i\rangle\langle a'_i| \|_{\text{tr}} \leq 2\delta^{1/2}(1 - \delta^2)^{-1/4}.$$

Fix some i , $1 \leq i \leq d$. Let $b_i + c_i$ denote the orthogonal projection of a_i onto the space spanned by V_1 and a_1, \dots, a_{i-1} , where $b_i \in V_1$ and $c_i \in \text{span}\{a_1, \dots, a_{i-1}\} \subseteq V_2$. Then,

$$\begin{aligned} 1 \geq \|b_i + c_i\|^2 &\geq \|b_i\|^2 + \|c_i\|^2 - 2|\langle b_i | c_i \rangle| \\ &\geq \|b_i\|^2 + \|c_i\|^2 - 2\delta \|b_i\| \|c_i\| \\ &= (1 - \delta^2) \|b_i\|^2 + (\delta \|b_i\| - \|c_i\|)^2 \\ &\geq (1 - \delta^2) \|b_i\|^2, \end{aligned}$$

i. e. $\|b_i\| \leq \frac{1}{\sqrt{1-\delta^2}}$. Now,

$$\|b_i + c_i\|^2 = \langle a_i | b_i + c_i \rangle = \langle a_i | b_i \rangle + \langle a_i | c_i \rangle = \langle a_i | b_i \rangle \leq \frac{\delta}{\sqrt{1-\delta^2}},$$

i. e. $\|b_i + c_i\| \leq \delta^{1/2}(1 - \delta^2)^{-1/4}$. The third equality above follows from the fact that a_1, \dots, a_{i-1}, a_i are pairwise orthogonal. Now $\langle a_i | a'_i \rangle = \|a_i - b_i - c_i\| = \sqrt{1 - \|b_i + c_i\|^2}$, and hence,

$$\| |a_i\rangle\langle a_i| - |a'_i\rangle\langle a'_i| \|_{\text{tr}} = 2\sqrt{1 - |\langle a_i | a'_i \rangle|^2} = 2\|b_i + c_i\| \leq 2\delta^{1/2}(1 - \delta^2)^{-1/4}.$$

This completes the proof of the proposition. ■

Proposition 3 Let v'_1, \dots, v'_n be unit vectors in a Hilbert space. Let $0 \leq \delta < \frac{1}{2n}$. Suppose for all $i, j, i \neq j$, $|\langle v_i | v_j \rangle| \leq \delta$. Let v_1, \dots, v_n be unit vectors obtained by Gram-Schmidt orthonormalising v'_1, \dots, v'_n . Then for any $i, 1 \leq i \leq n$,

$$\| |v_i\rangle\langle v_i| - |v'_i\rangle\langle v'_i| \|_{\text{tr}} < 2\sqrt{6} \cdot \delta\sqrt{n}.$$

Proof: Fix some $i, 1 \leq i < n$. Let $a_{i+1} = \sum_{j=1}^i \alpha_j v'_j$ be the orthogonal projection of v'_{i+1} onto the subspace spanned by v'_1, \dots, v'_i . Then for all $k, 1 \leq k \leq i$, $\langle v'_k | v'_{i+1} - a_{i+1} \rangle = 0$ i. e. $\langle v'_k | v'_{i+1} \rangle = \sum_{j=1}^i \alpha_j \langle v'_k | v'_j \rangle$. Suppose $k, 1 \leq k \leq i$ is such that $|\alpha_k| = \max_{j:1 \leq j \leq i} |\alpha_j|$. Then,

$$\begin{aligned} \delta \geq |\langle v'_k | v'_{i+1} \rangle| &\geq |\alpha_k| |\langle v'_k | v'_k \rangle| - \sum_{\substack{j:1 \leq j \leq i \\ j \neq k}} |\alpha_j| |\langle v'_k | v'_j \rangle| \\ &\geq |\alpha_k| - \sum_{\substack{j:1 \leq j \leq i \\ j \neq k}} |\alpha_k| |\langle v'_k | v'_j \rangle| \\ &\geq |\alpha_k| (1 - (i-1)\delta) \\ &> |\alpha_k| \cdot \frac{1}{2}, \end{aligned}$$

i. e. $\max_{j:1 \leq j \leq i} |\alpha_j| < 2\delta$. Now,

$$\begin{aligned} \|a_{i+1}\|^2 &\leq \sum_{j=1}^i |\alpha_j|^2 \|v_j\|^2 + \sum_{\substack{j,j':1 \leq j,j' \leq i \\ j \neq j'}} |\alpha_j| |\alpha_{j'}| |\langle v_j | v_{j'} \rangle| \\ &< 4\delta^2 n + 4\delta^3 n^2 \\ &< 4\delta^2 n + 2\delta^2 n \\ &= 6\delta^2 n. \end{aligned}$$

Reasoning as at the end of the proof of Proposition 2, we get

$$\begin{aligned}
\| |v_{i+1}\rangle\langle v_{i+1}| - |v'_{i+1}\rangle\langle v'_{i+1}| \|_{\text{tr}} &= 2\sqrt{1 - |\langle v_{i+1}|v'_{i+1}\rangle|^2} \\
&= 2\sqrt{1 - \|v'_{i+1} - a_{i+1}\|^2} \\
&= 2\sqrt{1 - (1 - \|a_{i+1}\|^2)} \\
&= 2\|a_{i+1}\| \\
&< 2\sqrt{6} \cdot \delta\sqrt{n}.
\end{aligned}$$

This completes the proof of the proposition. ■

3.2 Random orthonormal vectors

In this subsection, we state some facts about random orthonormal sets of vectors in \mathbb{C}^d . One way of generating a random unit vector in \mathbb{C}^d is as follows: Consider $(y_1, \dots, y_{2d}) \in \mathbb{R}^{2d}$, where each y_i is independently chosen according to the one dimensional Gaussian distribution with mean 0 and variance 1 (i. e. y_i is a real valued random variable with probability density function $\frac{1}{\sqrt{2\pi}} \exp(-y^2/2)$). Normalise to get the unit vector (x_1, \dots, x_{2d}) , where $x_i = \frac{y_i}{\sqrt{y_1^2 + \dots + y_{2d}^2}}$ (note that any $y_i = 0$ with zero probability). We thus get a random unit vector in \mathbb{R}^{2d} . Identifying a pair of real numbers with a single complex number, we get a random unit vector (z_1, \dots, z_d) in \mathbb{C}^d . To generate a random orthonormal ordered set $\{v_1, \dots, v_m\}$ of vectors in \mathbb{C}^d , we can first sample m unit vectors $\{v'_1, \dots, v'_m\}$ in \mathbb{C}^d and then do Gram-Schmidt orthonormalisation on them to get $\{v_1, \dots, v_m\}$ (note that with probability 1, $\{v'_1, \dots, v'_m\}$ are linearly independent).

The following fact can be proved by combining Theorem 14.3.2 and Proposition 14.3.3 of [Mat02, Chapter 14] and using the concavity of the square-root function.

Fact 2 *Let $t > 0$, and $|v\rangle, |w\rangle$ independent random unit vectors in \mathbb{C}^d . Then,*

$$\Pr \left[|\langle v|w\rangle| > t + \frac{10}{\sqrt{d}} \right] \leq 2 \exp(-t^2 d).$$

We will require the following upper and lower bounds on the tails of the chi-square distribution (the chi-square distribution with d degrees of freedom is the sum of squares of d independent Gaussians with mean 0 and variance 1). The upper bound can be proved via Chernoff-style arguments on the moment generating function of the chi-square distribution. The lower bound follows, for example, from the central limit theorem in probability theory. One can also give a direct proof of the lower bound using the probability density function of the chi-square distribution and estimating it via Stirling's approximation of the gamma function.

Fact 3 *Let (X_1, \dots, X_d) be independent random variables such that X_i is one-dimensional Gaussian with mean 0 and variance 1. Let $X^2 = X_1^2 + \dots + X_d^2$. Let $0 \leq \epsilon < 1/2$. There exists a universal constant $\gamma > 0$ such that*

1. $\Pr[|X^2 - d| > d\epsilon] < 2 \exp(-d\epsilon^2/6)$,
2. $\Pr[X^2 > d + \sqrt{d}] > \gamma, \Pr[X^2 < d - \sqrt{d}] > \gamma$.

The following result follows easily from Fact 3. A similar result appears as Lemma 2 in [MRRS04].

Fact 4 Let $V = \{a^1, \dots, a^p\}$ be a random orthonormal set of p vectors in \mathbb{C}^d . Let a_j^i denote the j th coordinate of vector a^i . Define the d -dimensional probability vector S as follows: $S_j = \frac{1}{p} \sum_{i=1}^p |a_j^i|^2$. Let $0 \leq \epsilon < 1/2$. Suppose $p = \Omega(\epsilon^{-2} \log d)$. Let U denote the uniform probability distribution on $\{1, \dots, d\}$. Then, with probability at least $1 - \exp(-\Omega(\epsilon^2 p))$ over the choice of V , $\|S - U\|_1 \leq \epsilon$.

We will also need the following Chernoff upper bounds on the tail of the sum of d independent identically distributed binary random variables.

Fact 5 ([AS00, Cor. A.7, Theorem A.13]) Let (X_1, \dots, X_d) be independent binary random variables such that $\Pr[X_i = 1] = p$. Let $X = X_1 + \dots + X_d$. Let $0 \leq \epsilon < 1/2$. Then,

1. $\Pr\left[\left|\frac{X}{d} - p\right| > \epsilon\right] < 2 \exp(-2\epsilon^2 d)$,
2. $\Pr[X < \frac{dp}{2}] < \exp(-dp/8)$.

3.3 Hidden subgroup problem and Fourier sampling

In this subsection, we recall the standard approach to solving the hidden subgroup problem based on Fourier sampling. A d -dimensional unitary representation of G is a group homomorphism $\rho : G \rightarrow \mathbf{U}(d)$, where $\mathbf{U}(d)$ is the group of $d \times d$ complex unitary matrices under multiplication. Let $\mathbb{C}[G]$ denote the group algebra; it is a $|G|$ -dimensional Hilbert space over \mathbb{C} with group elements $|g\rangle$, $g \in G$ as an orthonormal basis. Let $\mathcal{R}[G]$ denote the $|G|$ -dimensional Hilbert space over \mathbb{C} spanned by the orthonormal basis vectors $|\rho, i, j\rangle$, where ρ runs over inequivalent irreducible unitary representations of G and i, j run over the row and column indices of ρ . The quantum Fourier transform over G , QFT_G , is the following \mathbb{C} -linear map from $\mathbb{C}[G]$ to $\mathcal{R}[G]$ defined as follows:

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

where d_{ρ} denotes the dimension of ρ . QFT_G is an inner product preserving map from $\mathbb{C}[G]$ to $\mathcal{R}[G]$.

For a subset $T \subseteq G$, define $|T\rangle = \frac{1}{\sqrt{|T|}} \sum_{t \in T} |t\rangle$ to be the uniform superposition over elements of T . For a representation ρ , define the matrix $\rho(T) = \frac{1}{\sqrt{|T|}} \sum_{t \in T} \rho(t)$. If $H \leq G$, it can be shown (see e.g. [HRTS03]) that $\frac{1}{\sqrt{|H|}} \rho(H)$ is an orthogonal projection onto the subspace V_H^{ρ} of the representation space of ρ consisting of all vectors $|v\rangle$ such that $\rho(h)|v\rangle = |v\rangle$ for all $h \in H$. Thus, $\text{rank}(\rho(H)) = \dim V_H^{\rho}$.

In the strong Fourier sampling method for the hidden subgroup problem, we begin by forming the uniform superposition $\frac{1}{\sqrt{|G|}} \sum_{g \in G} |g\rangle |0\rangle$ and then query f to get the superposition $\frac{1}{\sqrt{|G|}} \sum_{g \in G} |g\rangle |f(g)\rangle$. We then measure the second register to get a uniform mixture over vectors $|gH\rangle$ in the first register. Assuming the first register is in state $|gH\rangle$ for some particular $g \in G$, its state after the application of QFT_G becomes

$$\frac{1}{\sqrt{|G||H|}} \sum_{\rho, i, j} \sqrt{d_{\rho}} \sum_{h \in H} \rho_{ij}(gh) |\rho, i, j\rangle.$$

If we now measure the representation name and column index, we sample (ρ, j) with probability

$$P_H^G(\rho, j) = \frac{d_{\rho}}{|G|} \sum_i |\rho_{ij}(gH)|^2 = \frac{d_{\rho}}{|G|} \|\rho(gH)|j\rangle\|^2 = \frac{d_{\rho}}{|G|} \|\rho(H)|j\rangle\|^2.$$

The third equality above follows from the fact that $\|\rho(gH)|j\rangle\| = \|\rho(g)\rho(H)|j\rangle\| = \|\rho(H)|j\rangle\|$, since $\rho(g)$ is unitary. Thus, as long as we measure just the representation name and column index (ρ, j) , the probabilities are independent of the actual coset gH that we find ourselves in. This fact can be viewed as the non-abelian generalisation of the fact that in abelian Fourier sampling the probability distribution on the characters is independent of the actual coset that we land up in. Also, it can be shown that (see [GSVV04])

$$P_H^G(\rho) = \sum_{j=1}^{d_\rho} \frac{d_\rho}{|G|} \|\rho(H)|j\rangle\|^2 = \frac{d_\rho|H|}{|G|} \text{rank}(\rho(H)) = \frac{d_\rho|H|}{|G|} \dim V_H^\rho.$$

In *weak Fourier sampling*, we only measure the names ρ of the representations and ignore the column indices j . It can be shown (see e.g. [HRTS03]) that for normal hidden subgroups H , no more information about H is contained in the column space of the resulting state after the measurement of ρ . Thus, weak Fourier sampling is the optimal measurement to recover a normal hidden subgroup starting from the uniform mixture of coset states.

Define a distance measure $w(G; H_1, H_2) = \sum_\rho |P_{H_1}^G(\rho) - P_{H_2}^G(\rho)|$ between subgroups $H_1, H_2 \leq G$. $w(G; H_1, H_2)$ is the total variation distance between the probability distributions, when the hidden subgroup is H_1 or H_2 , on the names of the representations obtained via weak Fourier sampling. [HRTS03, GSVV04] show that $O(\log |G|)$ weak Fourier samplings suffice to reconstruct the *normal core* $c(H)$ of the hidden subgroup H , where $c(H)$ is the largest normal subgroup of G contained in H . Adapting their arguments, we prove the following result.

Proposition 4 *Let $H_1, H_2 \leq G$. Suppose $c(H_1) \neq c(H_2)$. Then, $w(G; H_1, H_2) \geq 1/2$.*

Proof: Let $N_1 = c(H_1)$ and $N_2 = c(H_2)$. Without loss of generality, $N_1 \not\leq N_2$. Define the kernel of a representation $\ker \rho = \{g \in G : \rho(g) = \mathbb{1}_{d_\rho}\}$; $\ker(\rho) \triangleleft G$. It can be shown (see e.g. [HRTS03]) for an irreducible representation ρ and a subgroup $H \leq G$, that if $\text{rank}(\rho(H)) > 0$, $c(H) \triangleleft \ker \rho$. Hence,

$$1 = \sum_\rho P_{H_2}^G(\rho) = \sum_\rho \frac{d_\rho|H_2|}{|G|} \cdot \text{rank}(\rho H_2) = \sum_{\rho: N_2 \triangleleft \ker \rho} \frac{d_\rho|H_2|}{|G|} \cdot \text{rank}(\rho(H_2)).$$

Since $N_1 \triangleleft G$, $N_1 H_2$ is a subgroup of G . Hence, $N_1 \triangleleft c(N_1 H_2)$ and $N_2 \triangleleft c(N_1 H_2)$. Since $N_1 \not\leq N_2$, $|N_1 H_2| \geq 2 \cdot |H_2|$. For an irreducible representation ρ such that $N_1 \triangleleft \ker \rho$,

$$\text{rank}(\rho(N_1 H_2)) = \text{rank}(\rho(N_1)\rho(H_2)) = \text{rank}(\rho(H_2)).$$

Also,

$$1 = \sum_{\rho: N_1, N_2 \triangleleft \ker \rho} \frac{d_\rho|N_1 H_2|}{|G|} \cdot \text{rank}(\rho(N_1 H_2)) \geq 2 \cdot \sum_{\rho: N_1, N_2 \triangleleft \ker \rho} \frac{d_\rho|H_2|}{|G|} \cdot \text{rank}(\rho(H_2)),$$

i.e.

$$\sum_{\rho: N_1, N_2 \triangleleft \ker \rho} \frac{d_\rho|H_2|}{|G|} \cdot \text{rank}(\rho(H_2)) \leq \frac{1}{2}.$$

Finally,

$$w(G; H_1, H_2) = \sum_\rho \frac{d_\rho}{|G|} \cdot \left| |H_1| \text{rank}(\rho(H_1)) - |H_2| \text{rank}(\rho(H_2)) \right|$$

$$\begin{aligned}
&\geq \sum_{\rho: N_2 \triangleleft \ker \rho, N_1 \not\triangleleft \ker \rho} \frac{d_\rho}{|G|} \cdot |H_2| \operatorname{rank}(\rho(H_2)) \\
&\geq 1 - \sum_{\rho: N_1, N_2 \triangleleft \ker \rho} \frac{d_\rho}{|G|} \cdot |H_2| \operatorname{rank}(\rho(H_2)) \\
&\geq \frac{1}{2}.
\end{aligned}$$

This completes the proof. \blacksquare

For a normal subgroup $N \triangleleft G$, define the *normal core family* of N , $\operatorname{ncf}(N) = \{H : H \leq G, c(H) = N\}$. In view of Proposition 4, the remaining challenge is to distinguish between subgroups H_1, H_2 from the same normal core family.

The success of strong Fourier sampling depends on how much statistical information about H is present in the probability distribution $P_H^G(\rho, j)$. The amount of information, in general, depends on the choice of basis for each representation ρ , i. e., on the choice of basis for j ; see [MRRS04] for more details. Grigni et al. [GSVV04] show that under certain conditions on G and H , the *random strong* Fourier sampling method, where a random choice of basis is made for each representation, gives exponentially small information about distinguishing H from the identity subgroup. In the next section, we prove a complementary result viz. under different conditions on G , $(\log |G|)^{O(1)}$ random strong Fourier samplings do give enough information to reconstruct the hidden subgroup H with high probability.

4 Power of the random strong method

In this section, we define a parameter $r(G)$ on a group G which, if at least $(\log |G|)^{-O(1)}$, suffices for the random strong method to identify the hidden subgroup with $(\log |G|)^{O(1)}$ Fourier samplings. Let $H_1, H_2 \leq G$. We first define a distance measure $r(G; H_1, H_2)$ between H_1, H_2 . In what follows, we use the notation of Section 3.3.

Definition 1 ($r(G; H_1, H_2; \rho)$) *Suppose ρ is an irreducible d_ρ -dimensional unitary representation of G . Let Π_i denote the orthogonal projection onto $V_{H_i}^\rho$ i. e. $\Pi_i = \frac{1}{|H_i|} \sum_{h \in H_i} \rho(h)$. Let $\Pi_{1,2}$ denote the orthogonal projection onto $V_{H_1}^\rho \cap V_{H_2}^\rho$. It is easy to check that $V_{H_1}^\rho \cap V_{H_2}^\rho = V_{\langle H_1, H_2 \rangle}^\rho$, where $\langle H_1, H_2 \rangle$ denotes the subgroup of G generated by H_1 and H_2 . Thus, $\Pi_{1,2} = \frac{1}{|\langle H_1, H_2 \rangle|} \sum_{h \in \langle H_1, H_2 \rangle} \rho(h)$. Define $\Pi'_i = \Pi_i - \Pi_{1,2}$. Π'_i is the orthogonal projection onto the subspace V'_i defined as the orthogonal complement of $V_{H_1}^\rho \cap V_{H_2}^\rho$ in $V_{H_i}^\rho$. V'_1 and V'_2 have trivial intersection. Define $r_i = \operatorname{rank}(\Pi_i)$ and $r'_i = \operatorname{rank}(\Pi'_i)$. Define $\hat{h} = \max\{|H_1|r_1, |H_2|r_2\}$, $\tilde{h} = |(|H_1|r_1 - |H_2|r_2)|$ and $\delta = \|\Pi'_1 \Pi'_2\|$. Recall that $\delta = \operatorname{ovlap}(V'_1, V'_2)$. Consider the following three cases:*

1. When $\frac{\sqrt{d_\rho}}{\log |G|} = \Omega((r_1 + r_2)^{3/2})$. Loosely speaking, r_1, r_2 are both small. In this case, define

$$r(G; H_1, H_2; \rho) = \max \left\{ \frac{\hat{h}}{2} \left(\Omega \left(\frac{\sqrt{r'_1}}{r_1} + \frac{\sqrt{r'_2}}{r_2} \right) - 2\delta^{1/2}(1 - \delta^2)^{-1/4} \right), \tilde{h} \right\}.$$

2. When $\frac{\sqrt{d_\rho}}{\log |G|} = \Omega(r_1)$ and $\frac{r_2}{r_1} = \Omega(\log^2 |G|)$. Loosely speaking, r_1 is small and r_2 is relatively large with respect to r_1 . In this case, define

$$r(G; H_1, H_2; \rho) = \max \left\{ \frac{\hat{h}}{2} \cdot \Omega \left(\frac{1}{\sqrt{r_1}} \right), \tilde{h} \right\}.$$

3. Otherwise, define $r(G; H_1, H_2; \rho) = \tilde{h}$.

Definition 2 ($r(G; H_1, H_2), r(G)$) Let $H_1, H_2 \leq G$. Define $r(G; H_1, H_2) = \sum_{\rho} \frac{d_{\rho}}{|G|} \cdot r(G; H_1, H_2; \rho)$ and $r(G) = \min_{H_1, H_2} r(G; H_1, H_2)$.

From the above definition, it is easy to see that $r(G; H_1, H_2) \geq w(G; H_1, H_2)$.

Definition 3 ($P_{H, \mathcal{B}}^G$) Let \mathcal{B} be a set of orthonormal bases for the irreducible unitary representations of G . Suppose $H \leq G$. $P_H^{G, \mathcal{B}}$ denotes the probability distribution on the representation names and column indices (ρ, j) got by strong Fourier sampling the state $|H\rangle$ according to \mathcal{B} .

The significance of $r(G; H_1, H_2)$ arises from the following theorem.

Theorem 1 With probability at least $1 - \exp(-\Omega(\log^2 |G|))$ over the choice of random representation bases \mathcal{B} for Fourier sampling,

$$\|P_{H_1}^{G, \mathcal{B}} - P_{H_2}^{G, \mathcal{B}}\|_{\text{tr}} \geq r(G; H_1, H_2).$$

Using this theorem, we can apply a ‘minimum-finding-like’ algorithm to identify the hidden subgroup.

Corollary 1 Let $s(G)$ denote the number of distinct subgroups of G . With probability at least $2/3$ over the choice of random bases for representations of G , Fourier sampling $O\left(\frac{\log s(G)}{r^2(G)}\right)$ times in a random basis gives enough classical information to identify a hidden subgroup in G . In particular, $O\left(\left(\frac{\log |G|}{r(G)}\right)^2\right)$ random strong Fourier samplings suffice.

Proof: From Theorem 1, we get that for all pairs of subgroups $H_1, H_2 \leq G$, with probability at least $1 - \exp(-\Omega(\log^2 |G|))$ over the choice of random bases \mathcal{B} for representations of G , $\|P_{H_1}^{G, \mathcal{B}} - P_{H_2}^{G, \mathcal{B}}\|_{\text{tr}} \geq r(G)$. Call a set of representation bases \mathcal{B} *good* if $\|P_{H_1}^{G, \mathcal{B}} - P_{H_2}^{G, \mathcal{B}}\|_{\text{tr}} \geq r(G)$ for all pairs of subgroups $H_1, H_2 \leq G$. By the union bound on probabilities, a random choice of representation bases gives a good \mathcal{B} with probability at least $1 - s(G) \exp(-\Omega(\log^2 |G|)) = 1 - \exp(-\Omega(\log^2 |G|))$. Suppose we have such a good \mathcal{B} . Under the promise that the hidden subgroup is either H_1 or H_2 , \mathcal{B} recognises which one it is with confidence at least $1/2 + r(G)/4$ using Bayes’s rule. Using Fact 5, the confidence can be boosted to at least $1 - \frac{1}{4s(G)}$ by Fourier sampling $O\left(\frac{\log s(G)}{r^2(G)}\right)$ times with \mathcal{B} . We can now run a classical ‘minimum-finding-like’ algorithm on the measured samples, comparing two subgroups $H_1, H_2 \leq G$ at a time, to discover the actual hidden subgroup H in G with confidence at least $1 - \frac{s(G)}{4s(G)} = 3/4$. The overall confidence bound becomes $(1 - \exp(-\Omega(\log^2 |G|))) \cdot \frac{3}{4} \geq \frac{2}{3}$. The second bound follows from the fact that $s(G) \leq 2^{\log^2 |G|}$, since any group of size a has at most $\log a$ generators. ■

The rest of the section is devoted to proving Theorem 1. We first prove some necessary technical lemmas.

Lemma 2 Let $W = \{a^1, \dots, a^p\} \cup \{b^1, \dots, b^q\} \cup \{c^1, \dots, c^r\}$ be a random orthonormal set of $p + q + r$ vectors in \mathbb{C}^d . Let a_j^i denote the j th coordinate of vector a^i ; similar notations will be used for the vectors b^i, c^i too. Define two d -dimensional probability vectors S, T as follows:

$$S_j = \frac{1}{p+r} \left(\sum_{i=1}^p |a_j^i|^2 + \sum_{i=1}^r |c_j^i|^2 \right), T_j = \frac{1}{q+r} \left(\sum_{i=1}^q |b_j^i|^2 + \sum_{i=1}^r |c_j^i|^2 \right).$$

Then there exists $\delta = \theta((p+q+r)^{-3/2})$ such that the following holds: Define $\alpha = d\delta^2 - 2\log(p+q+r)$. Suppose $\alpha = \Omega(1)$. Then, with probability at least $1 - \exp(-\Omega(\alpha))$ over the choice of W ,

$$\|S - T\|_1 = \Omega\left(\frac{\sqrt{p}}{p+r} + \frac{\sqrt{q}}{q+r}\right).$$

Proof: Generate a set $W' = \{a^1, \dots, a^p\} \cup \{b^1, \dots, b^q\} \cup \{c^1, \dots, c^r\}$ of $p+q+r$ random independent unit vectors in \mathbb{C}^d as described in Section 3.2. Let $\alpha_j^i, \beta_j^i, \gamma_j^i, j = \{1, \dots, d\}$ denote the Gaussians used to generate the random unit vectors a^i, b^i, c^i respectively. Then,

$$a_j^i = \frac{\alpha_j^i}{\sum_{l=1}^d |\alpha_l^i|^2}, \quad b_j^i = \frac{\beta_j^i}{\sum_{l=1}^d |\beta_l^i|^2}, \quad c_j^i = \frac{\gamma_j^i}{\sum_{l=1}^d |\gamma_l^i|^2}.$$

By Fact 3, with probability at least $1 - 2\exp(-d\delta^2)$ over the choice of the Gaussians, the normalisation factor in the denominator of a given vector in W' is $\sqrt{(1 \pm \epsilon)d}$ where $\epsilon = O(\delta)$. Let E_0 be the event that the normalisation factors in the denominators of all vectors in W' are $\sqrt{(1 \pm \epsilon)d}$. By the union bound on probabilities, E_0 occurs with probability at least $1 - 2\exp(-\alpha)$ over the choice of the Gaussians.

Since $\alpha = \Omega(1)$, $\delta > \frac{10}{\sqrt{d}}$. By Fact 2, for any $w^i, w^j \in W', i \neq j$, $|\langle w^i | w^j \rangle| \leq \delta + \frac{10}{\sqrt{d}} < 2\delta$ with probability at least $1 - 2\exp(-d\delta^2)$ over the choice of the Gaussians. Let E_1 denote the event that $|\langle w^i | w^j \rangle| < 2\delta$ for all $w^i, w^j \in W', i \neq j$. By the union bound on probabilities, E_1 occurs with probability at least $1 - 2\exp(-\alpha)$ over the choice of the Gaussians.

Using Fact 3 we see that for any fixed coordinate j , with constant probability at least θ over the choice of the Gaussians, each of the following three events occurs :

$$\sum_{i=1}^p |\alpha_j^i|^2 > p + \sqrt{p}, \quad \sum_{i=1}^q |\beta_j^i|^2 < q - \sqrt{q}, \quad \sum_{i=1}^r |\gamma_j^i|^2 > r + 1.$$

Since these are independent events, all three of them hold at coordinate j simultaneously with constant probability at least θ^3 . Call such a coordinate j *good*. Let E_2 denote the event that more than $\frac{d\theta^3}{2}$ coordinates j are good. By Fact 5, E_2 occurs with probability at least $1 - \exp(-d\theta^3/8)$ over the choice of the Gaussians.

Now suppose that all three events E_0, E_1, E_2 occur. We Gram-Schmidt orthonormalise W' to get the random orthonormal set $W = \{a^1, \dots, a^p\} \cup \{b^1, \dots, b^q\} \cup \{c^1, \dots, c^r\}$. Let S', T' be the analogous probability vectors defined with respect to W' instead of W . From Proposition 3, we see that

$$\| |w\rangle\langle w| - |w'\rangle\langle w'| \|_{\text{tr}} < 20 \cdot \delta \cdot (p+q+r) = O((p+q+r)^{-1})$$

for corresponding vectors $w \in W, w' \in W'$. Define density matrices

$$\sigma = \frac{1}{p+r} \left(\sum_{i=1}^p |a^i\rangle\langle a^i| + \sum_{i=1}^r |c^i\rangle\langle c^i| \right), \quad \sigma' = \frac{1}{p+r} \left(\sum_{i=1}^p |a^i\rangle\langle a^i| + \sum_{i=1}^r |c^i\rangle\langle c^i| \right),$$

$$\tau = \frac{1}{q+r} \left(\sum_{i=1}^q |b^i\rangle\langle b^i| + \sum_{i=1}^r |c^i\rangle\langle c^i| \right), \quad \tau' = \frac{1}{q+r} \left(\sum_{i=1}^q |b^i\rangle\langle b^i| + \sum_{i=1}^r |c^i\rangle\langle c^i| \right).$$

Then, S, S', T, T' are the probability distributions got by measuring the states $\sigma, \sigma', \tau, \tau'$ in the standard basis of \mathbb{C}^d . By triangle inequality, $\|\sigma - \sigma'\|_{\text{tr}} = O((p+q+r)^{-1})$ and $\|\tau - \tau'\|_{\text{tr}} = O((p+q+r)^{-1})$. Hence, $\|S - S'\|_1 = O((p+q+r)^{-1})$ and $\|T - T'\|_1 = O((p+q+r)^{-1})$.

For a good coordinate j

$$\begin{aligned}
|S'_j - T'_j| &> \frac{p}{(p+r)(1+\epsilon)d} - \frac{q}{(q+r)(1-\epsilon)d} + \frac{r+1}{(1+\epsilon)d} \left(\frac{1}{p+r} - \frac{1}{q+r} \right) \\
&\quad + \frac{\sqrt{p}}{(p+r)(1+\epsilon)d} + \frac{\sqrt{q}}{(q+r)(1-\epsilon)d} \\
&= \frac{-2q\epsilon}{(q+r)(1-\epsilon^2)d} + \frac{q-p}{(p+r)(q+r)(1+\epsilon)d} + \frac{\sqrt{p}}{(p+r)(1+\epsilon)d} + \frac{\sqrt{q}}{(q+r)(1-\epsilon)d} \\
&> -O\left(\frac{1}{(p+q+r)^{1/2}(q+r)(1-O((p+q+r)^{-3}))d}\right) + \Omega\left(\frac{1}{d}\left(\frac{\sqrt{p}}{p+r} + \frac{\sqrt{q}}{q+r}\right)\right) \\
&= \Omega\left(\frac{1}{d}\left(\frac{\sqrt{p}}{p+r} + \frac{\sqrt{q}}{q+r}\right)\right).
\end{aligned}$$

The first, third and fourth steps above follow from the fact that $\epsilon = O(\delta) = O((p+q+r)^{-3/2})$ and $p \leq q$ without loss of generality.

Now,

$$\|S' - T'\|_1 \geq \sum_{j:j \text{ good}} |S'_j - T'_j| > \frac{d\theta^3}{2} \cdot \Omega\left(\frac{1}{d}\left(\frac{\sqrt{p}}{p+r} + \frac{\sqrt{q}}{q+r}\right)\right) = \Omega\left(\frac{\sqrt{p}}{p+r} + \frac{\sqrt{q}}{q+r}\right).$$

Finally,

$$\begin{aligned}
\|S - T\|_1 &\geq \|S' - T'\|_1 - \|S - S'\|_1 - \|T - T'\|_1 \\
&= \Omega\left(\frac{\sqrt{p}}{p+r} + \frac{\sqrt{q}}{q+r}\right) - 2 \cdot O\left(\frac{1}{p+q+r}\right) \\
&= \Omega\left(\frac{\sqrt{p}}{p+r} + \frac{\sqrt{q}}{q+r}\right).
\end{aligned}$$

The confidence bound is

$$\Pr[E_0 \wedge E_1 \wedge E_2] > 1 - 4 \exp(-\alpha) - \exp(d\theta^3/8) = 1 - \exp(-\Omega(\alpha)),$$

since $\delta = O(1)$. This completes the proof of the lemma. \blacksquare

We can prove the following lemma in a similar fashion as Lemma 2.

Lemma 3 Let $W = \{a^1, \dots, a^p\}$ be a random orthonormal set of p vectors in \mathbb{C}^d . Let a_j^i denote the j th coordinate of vector a^i . Define the d -dimensional probability vector S as follows: $S_j = \frac{1}{p} \sum_{i=1}^p |a_j^i|^2$. Then there exists $\delta = \theta(p^{-1})$ such that the following holds: Define $\alpha = d\delta^2 - 2 \log p$. Suppose $\alpha = \Omega(1)$. Let U denote the uniform probability distribution on $\{1, \dots, d\}$. Then, with probability at least $1 - \exp(-\Omega(\alpha))$ over the choice of V , $\|S - U\|_1 = \Omega(p^{-1/2})$.

Proof: (Sketch) Generate a set $W' = \{a'^1, \dots, a'^p\}$ of p random independent unit vectors in \mathbb{C}^d as described in Section 3.2. Let α_j^i , $j = \{1, \dots, d\}$ denote the Gaussians used to generate the random unit vectors a^i . Then, $a_j^i = \frac{\alpha_j^i}{\sum_{l=1}^d |\alpha_l^i|^2}$. We Gram-Schmidt orthonormalise W' to get the random orthonormal set $W = \{a^1, \dots, a^p\}$. Let E_0 be the event that the normalisation factors in the denominators of all vectors in W' are $\sqrt{(1 \pm \epsilon)d}$, where $\epsilon = O(\delta)$. E_0 occurs with probability at least $1 - 2 \exp(-\alpha)$ over the

choice of the Gaussians. Let E_1 denote the event that $|\langle w^i | w^j \rangle| < 2\delta$ for all $w^i, w^j \in W'$, $i \neq j$. E_1 occurs with probability at least $1 - 2\exp(-\alpha)$ over the choice of the Gaussians. Call a coordinate j *good* if $\sum_{i=1}^p |\alpha_j^i|^2 > p + \sqrt{p}$. Let E_2 denote the event that more than $\frac{d\theta}{2}$ coordinates j are good. E_2 occurs with probability at least $1 - \exp(d\theta/8)$ over the choice of the Gaussians.

Now suppose that all three events E_0, E_1, E_2 occur. Let S' be the analogous probability vector defined with respect to W' instead of W . Then, $\|S - S'\|_1 = O(p^{-1/2})$. For a *good* coordinate j

$$\begin{aligned} \left| S'_j - \frac{1}{d} \right| &= \frac{1}{(1+\epsilon)d} - \frac{1}{d} + \frac{1}{\sqrt{p}(1+\epsilon)d} \\ &= \frac{-\epsilon}{(1+\epsilon)d} + \frac{1}{\sqrt{p}(1+\epsilon)d} \\ &= -O\left(\frac{-1}{dp}\right) + \Omega\left(\frac{1}{d\sqrt{p}}\right) \\ &= \Omega\left(\frac{1}{d\sqrt{p}}\right). \end{aligned}$$

The third step above follows from the fact that $\epsilon = O(\delta) = O(p^{-1})$. Hence,

$$\|S' - U\|_1 \geq \sum_{j:j \text{ good}} |S'_j - T'_j| > \frac{d\theta}{2} \cdot \Omega\left(\frac{1}{d\sqrt{p}}\right) = \Omega\left(\frac{1}{\sqrt{p}}\right).$$

Finally,

$$\|S - U\|_1 \geq \|S' - U\|_1 - \|S - S'\|_1 = \Omega(p^{-1/2}) - O(p^{-1}) = \Omega(p^{-1/2}).$$

The confidence bound is

$$\Pr[E_0 \wedge E_1 \wedge E_2] > 1 - 4\exp(-\alpha) - \exp(d\theta/8) = 1 - \exp(-\Omega(\alpha)),$$

since $\delta = O(1)$. This completes the proof of the lemma. ■

We are now in a position to finally prove Theorem 1.

Proof: (of Theorem 1) Let ρ be an irreducible d_ρ -dimensional unitary representation of G . We follow the notation of Definition 1 for ρ . Let V_2'' denote the orthogonal complement of V_1' in $V_1' + V_2'$. Let σ_i denote the totally mixed state in V_i and σ_2'' denote the totally mixed state in $V_2'' + (V_1 \cap V_2)$. By Proposition 2, $\|\sigma_2 - \sigma_2''\|_{\text{tr}} < 2\delta^{1/2}(1 - \delta^2)^{-1/4}$. Let \mathcal{B}_ρ be a random orthonormal basis for ρ . Let $P_i := P_{H_i}^{G, \mathcal{B}_\rho}$ denote the probability distributions on the vectors of \mathcal{B}_ρ got by Fourier sampling the states $|H_i\rangle$ respectively, conditioned on ρ being observed. Then P_i is the probability distribution got by measuring σ_i in the basis \mathcal{B}_ρ . Let P_2'' denote the probability distribution got by measuring σ_2'' in the basis \mathcal{B}_ρ . Then, $\|P_2 - P_2''\|_1 < 2\delta^{1/2}(1 - \delta^2)^{-1/4}$. Define $r_{1,2} = \text{rank}(\Pi_{1,2})$. Note that $r_i = r'_i + r_{1,2}$.

Suppose case 1 of Definition 1 applies. Let W be a random orthonormal set of $r'_1 + r'_2 + r_{1,2}$ vectors in C^{d_ρ} . Define probability distributions S, T with respect to W as in Lemma 2. By symmetry, $P_1 = S$ and $P_2'' = T$. Note that $r'_1 + r'_2 + r_{1,2} \leq r_1 + r_2$ and $r'_1 + r'_2 + r_{1,2} \leq d_\rho < \sqrt{|G|}$. Hence, $d_\rho \cdot \theta((r'_1 + r'_2 + r_{1,2})^{-3}) - 2\log(r'_1 + r'_2 + r_{1,2}) = \Omega(\log^2 |G|)$. The conditions of Lemma 2 are satisfied, and we get, with probability at least $1 - \exp(-\Omega(\log^2 |G|))$ over the choice of \mathcal{B}_ρ , that $\|P_1 - P_2''\|_1 = \Omega\left(\frac{\sqrt{r'_1}}{r_1} + \frac{\sqrt{r'_2}}{r_2}\right)$.

Thus with probability at least $1 - \exp(-\Omega(\log^2 |G|))$ over the choice of \mathcal{B}_ρ ,

$$\|P_1 - P_2\|_1 \geq \|P_1 - P_2''\|_1 - \|P_2 - P_2''\|_1 \geq \Omega\left(\frac{\sqrt{r'_1}}{r_1} + \frac{\sqrt{r'_2}}{r_2}\right) - 2\delta^{1/2}(1 - \delta^2)^{-1/4}.$$

Suppose case 2 of Definition 1 applies. Let W be a random orthonormal set of r_1 vectors in C^{d_ρ} . Define probability distribution S with respect to W as in Lemma 3. By symmetry, $P_1 = S$. Note that $r_1 \leq d_\rho < \sqrt{|G|}$. Hence, $d_\rho \cdot \theta(r_1^{-2}) - 2 \log r_1 = \Omega(\log^2 |G|)$. The conditions of Lemma 3 are satisfied, and we get, with probability at least $1 - \exp(-\Omega(\log^2 |G|))$ over the choice of \mathcal{B}_ρ , that $\|P_1 - U\|_1 = \Omega\left(\frac{1}{\sqrt{r_1}}\right)$. Also, $r_2 = \Omega(r_1 \log d_\rho)$. The conditions of Fact 4 are satisfied, and we get, with probability at least $1 - \exp(-\Omega(\log^2 |G|))$ over the choice of \mathcal{B}_ρ , that $\|P_2 - U\|_1 = O\left(\frac{1}{\sqrt{r_1}}\right)$. Thus with probability at least $1 - \exp(-\Omega(\log^2 |G|))$ over the choice of \mathcal{B}_ρ ,

$$\|P_1 - P_2\|_1 \geq \|P_1 - U\|_1 - \|P_2 - U\|_1 = \Omega\left(\frac{1}{\sqrt{r_1}}\right).$$

Suppose $|H_1|r_1 \geq |H_2|r_2$ i. e. $\hat{h} = |H_1|r_1$. Then,

$$\|P_{H_1}^{G,\mathcal{B}} - P_{H_2}^{G,\mathcal{B}}\|_1 = \sum_{\rho} \|P_{H_1}^G(\rho)P_{H_1}^{G,\mathcal{B}_\rho} - P_{H_2}^G(\rho)P_{H_2}^{G,\mathcal{B}_\rho}\|_1 = \sum_{\rho} \frac{d_\rho}{|G|} \cdot \|H_1 r_1 P_{H_1}^{G,\mathcal{B}_\rho} - H_2 r_2 P_{H_2}^{G,\mathcal{B}_\rho}\|_1.$$

Now,

$$\begin{aligned} \|H_1 r_1 P_{H_1}^{G,\mathcal{B}_\rho} - H_2 r_2 P_{H_2}^{G,\mathcal{B}_\rho}\|_1 &= \|H_1 r_1 (P_{H_1}^{G,\mathcal{B}_\rho} - P_{H_2}^{G,\mathcal{B}_\rho}) + (H_1 r_1 - H_2 r_2) P_{H_2}^{G,\mathcal{B}_\rho}\|_1 \\ &\geq \frac{H_1 r_1}{2} \cdot \|P_{H_1}^{G,\mathcal{B}_\rho} - P_{H_2}^{G,\mathcal{B}_\rho}\|_1. \end{aligned}$$

The last step above follows from the facts $H_1 r_1 - H_2 r_2 \geq 0$, $\|v\|_1 \geq \|v_+\|_1$ where $(v_+)_i := v_i$ if $v_i \geq 0$, $(v_+)_i := 0$ otherwise, and $\|P_1 - P_2\|_1 = 2\|(P_1 - P_2)_+\|_1$ for probability vectors P_1, P_2 . Also note that for any choice of representation bases \mathcal{B} , $\|H_1 r_1 P_{H_1}^{G,\mathcal{B}_\rho} - H_2 r_2 P_{H_2}^{G,\mathcal{B}_\rho}\|_1 \geq |H_1 r_1 - H_2 r_2|$. Hence,

$$\begin{aligned} \|P_{H_1}^{G,\mathcal{B}} - P_{H_2}^{G,\mathcal{B}}\|_1 &\geq \sum_{\rho} \frac{d_\rho}{|G|} \cdot \max \left\{ \frac{H_1 r_1}{2} \cdot \|P_{H_1}^{G,\mathcal{B}_\rho} - P_{H_2}^{G,\mathcal{B}_\rho}\|_1, |H_1 r_1 - H_2 r_2| \right\} \\ &\geq \sum_{\rho} \frac{d_\rho}{|G|} r(G; H_1, H_2; \rho) \\ &= r(G; H_1, H_2). \end{aligned}$$

For each representation ρ , the confidence bound in applying the above random basis arguments is at least $1 - \exp(-\Omega(\log^2 |G|))$. Since there are at most $|G|$ representations, the total confidence bound is at least $1 - |G| \exp(-\Omega(\log^2 |G|)) = 1 - \exp(-\Omega(\log^2 |G|))$. This completes the proof of the theorem. \blacksquare

We now have all the tools to prove that $r(\mathcal{H}_p) = \Omega(1)$. In fact, we can now prove the following theorem.

Theorem 2 *The random strong method is sufficient to solve the hidden subgroup problem in the Heisenberg group \mathcal{H}_p . The query complexity of this algorithm is $O(\log p)$. The quantum part of the algorithm consists of a circuit of size $O(\log^4 p)$ followed by a circuit of size $\tilde{O}(p^2)$ for implementing the measurement in a random orthonormal basis. The classical post-processing does not make any queries and has a running time of $\tilde{O}(p^4)$.*

Proof: First, we characterize the normal core families in the Heisenberg group. We have that

$$\text{ncf}(\mathcal{H}_p) = \{\mathcal{H}_p\}, \quad \text{ncf}(\zeta(\mathcal{H}_p)) = \{\zeta(\mathcal{H}_p)\}, \quad \text{ncf}(N_i) = \{N_i\}, \quad \text{for } i \in \{0, \dots, p-1, \infty\}$$

are families of size 1 each. For the trivial group we get that

$$\text{ncf}(\{1\}) = \{A_{i,j} : i \in \{0, \dots, p-1, \infty\}, j \in \{0, \dots, p-1\}\} \cup \{\{1\}\}.$$

If H_1 and H_2 are candidate hidden subgroups from different normal core families, then by Proposition 4 we get that $r(\mathcal{H}_p; H_1, H_2) \geq w(\mathcal{H}_p; H_1, H_2) \geq 1/2$. We now consider the situation where both $H_1, H_2 \in \text{ncf}(\{1\})$. We fix an irreducible representation $\rho = \rho_k$ (for $k = 1, \dots, p-1$) of degree $\deg \rho = p$. Now, we distinguish two cases:

1. $|H_1| = |H_2| = p$, i. e., there are $i, j, i', j', (i, j) \neq (i', j')$ such that $H_1 = A_{i,j}$ and $H_2 = A_{i',j'}$. Using the notation of Definition 1 we have that $r_1 = r_2 = 1$, since by Lemma 1 the ranks of $P_{k;i,j}$ and $P_{k;i',j'}$ are both one. Also, $\|P_{k;i,j} P_{k;i',j'}\| \leq \frac{1}{\sqrt{p}}$. This also implies that $r'_1 = r_1 = 1$ and $r'_2 = r_2 = 1$, since the one-dimensional projectors $P_{k;i,j}, P_{k;i',j'}$ are linearly independent. Hence, $P'_{k;i,j} = P_{k;i,j}$ and $P'_{k;i',j'} = P_{k;i',j'}$. So, $\delta = \|P'_{k;i,j} P'_{k;i',j'}\| \leq \frac{1}{\sqrt{p}}$. Since $|H_1| r_1 = |H_2| r_2 = p$, $\hat{h} = p$ and $\tilde{h} = 0$. As

$$\frac{\sqrt{d_\rho}}{\log |\mathcal{H}_p|} = \frac{\sqrt{p}}{3 \log p} = \Omega((r_1 + r_2)^{3/2}) = \Omega(1),$$

we are in the first case of Definition 1 and obtain that

$$r(\mathcal{H}_p; H_1, H_2; \rho) = \frac{p}{2} \cdot \left(\Omega(1) - 2(p-1)^{-1/4} \right) = \Omega(p).$$

2. $|H_1| = p$ and $|H_2| = 1$, i. e., we have to distinguish $H_1 = A_{i,j}$ from the trivial subgroup $H_2 = \{1\}$. In this case $r_1 = 1$ and $r_2 = \text{rank}(\rho(\{1\})) = p$ which implies that $\hat{h} = p$, $\tilde{h} = 0$. Since

$$\frac{\sqrt{d_\rho}}{\log |\mathcal{H}_p|} = \Omega(r_1) \quad \text{and} \quad \frac{r_2}{r_1} = p = \Omega(\log^2 |\mathcal{H}_p|),$$

we are in the second case of Definition 1 and obtain that

$$r(\mathcal{H}_p; H_1, H_2; \rho) = \frac{p}{2} \cdot \Omega(1) = \Omega(p).$$

Overall we obtain that for $H_1, H_2 \in \text{ncf}(\{1\})$,

$$r(\mathcal{H}_p; H_1, H_2) = \sum_{\rho} \frac{d_\rho}{p^3} \cdot r(\mathcal{H}_p; H_1, H_2; \rho) \geq \sum_{k=1}^{p-1} \frac{p}{p^3} \cdot r(\mathcal{H}_p; H_1, H_2; \rho_k) \geq \frac{(p-1)p}{p^3} \cdot \Omega(p) = \Omega(1).$$

Hence, $r(\mathcal{H}_p) = \Omega(1)$. Recall that $s(\mathcal{H}_p) = O(p^2)$. Now Corollary 1 shows that with probability at least $2/3$ over the choice of random representation bases, the HSP for \mathcal{H}_p can be solved using $O(\log p)$ random strong Fourier samplings.

As shown in Proposition 1, the QFT over \mathcal{H}_p can be implemented using $O(\log^3 p)$ elementary quantum gates. Since there are $O(\log p)$ Fourier samplings, the initial part of the quantum circuit has size $O(\log^4 p)$. The claimed statements about the number of quantum operations necessary to implement a measurement in a random orthonormal basis follow from general upper bounds of $\tilde{O}(p^2)$ on the number of gates in a factorization of a unitary operation $U \in \mathbf{U}(p)$ into elementary gates. The classical time to generate the random U is $\tilde{O}(p^3)$ since we can start with a set of p random unit vectors and apply Gram-Schmidt orthonormalisation to obtain a random unitary matrix. For the classical post-processing we have to compute

a table of probability distributions with respect to the random measurement bases for all subgroups. Since there are $O(p^2)$ subgroups and each probability distribution computation takes time $\tilde{O}(p^2)$ we can upper bound this by $\tilde{O}(p^4)$. After this table has been precomputed the actual algorithm to find the hidden subgroup is ‘minimum-finding-like’ in which we ‘compare’ two subgroups at a time. This takes time $\tilde{O}(p^2)$. Overall, we obtain that the running time of the classical part of this algorithm can be upper bounded by $\tilde{O}(p^4)$. ■

Acknowledgments

We thank Frédéric Magniez, Leonard Schulman, Cris Moore, Alex Russell and Avery Miller for useful discussions.

References

- [AS00] N. Alon and J. Spencer. *The probabilistic method*. John Wiley and Sons, 2000.
- [Bea97] R. Beals. Quantum computation of Fourier transforms over the symmetric groups. In *Proceedings of the Symposium on Theory of Computing (STOC'97)*, pages 48–53, El Paso, Texas, 1997.
- [BH97] G. Brassard and P. Høyer. An exact polynomial-time algorithm for Simon’s problem. In *Proceedings of Fifth Israeli Symposium on Theory of Computing and Systems*, pages 12–33. ISTCS, IEEE Computer Society Press, 1997. See also ArXiv preprint quant-ph/9704027.
- [CR62] W. C. Curtis and I. Reiner. *Representation Theory of Finite Groups and Algebras*. Wiley and Sons, 1962.
- [EHK04] M. Ettinger, P. Høyer, and E. Knill. The quantum query complexity of the hidden subgroup problem is polynomial. *Information Processing Letters*, 91(1):43–48, 2004. See also ArXiv preprint quant-ph/0401083.
- [Eme04] J. Emerson. Random quantum circuits and pseudo-random operators: theory and applications. ArXiv preprint quant-ph/0410087, 2004.
- [EWS⁺03] J. Emerson, Y. Weinstein, M. Saraceno, S. Lloyd, and D. Cory. Pseudo-Random unitary operators for quantum information processing. *Science*, 302:2098–2100, 2003.
- [FIM⁺03] K. Friedl, G. Ivanyos, F. Magniez, M. Santha, and P. Sen. Hidden translation and orbit coset in quantum computing. In *Proceedings of the Symposium on Theory of Computing (STOC)*, pages 1–9, 2003.
- [Gav04] D. Gavinsky. Quantum solution to the hidden subgroup problem for poly-near-Hamiltonian groups. *Quantum Information and Computation*, 4(3):229–235, 2004.
- [GSVV04] M. Grigni, L. Schulman, M. Vazirani, and U. Vazirani. Quantum mechanical algorithms for the nonabelian hidden subgroup problem. *Combinatorica*, pages 137–154, 2004.
- [Høy97] P. Høyer. Efficient Quantum Transforms. ArXiv preprint quant-ph/9702028, February 1997.
- [HRTS03] S. Hallgren, A. Russell, and A. Ta-Shma. The Hidden Subgroup Problem and Quantum Computation Using Group Representations. *SIAM Journal on Computing*, 32(4):916–934, 2003.

- [IMS03] G. Ivanyos, F. Magniez, and M. Santha. Efficient quantum algorithms for some instances of the non-Abelian hidden subgroup problem. *International Journal of Foundations of Computer Science*, pages 723–740, 2003.
- [Ip03] L. Ip. Shor’s algorithm is optimal. Unpublished manuscript, 2003.
- [Kit95] A. Yu. Kitaev. Quantum measurements and the abelian stabilizer problem. ArXiv preprint quant-ph/9511026, 1995.
- [Kup03] G. Kuperberg. A subexponential-time quantum algorithm for the dihedral hidden subgroup problem. ArXiv preprint quant-ph/0302112, 2003.
- [LN94] R. Lidl and H. Niederreiter. *Introduction to finite fields and their applications*. Cambridge University Press, 2nd edition, 1994.
- [Mat02] J. Matoušek. *Lectures on Discrete Geometry*. Graduate Texts in Mathematics. Springer-Verlag, 2002.
- [ME98] M. Mosca and A. Ekert. The hidden subgroup problem and eigenvalue estimation on a quantum computer. In *Quantum Computing and Quantum Communications, QCQC’98, Palm Springs*, volume 1509 of *LNCS*, pages 174–188. Springer, 1998.
- [MR05] C. Moore and A. Russell. For distinguishing conjugate hidden subgroups, the pretty good measurement is as good as it gets. ArXiv preprint quant-ph/0501177, 2005.
- [MRR04] C. Moore, D. Rockmore, and A. Russell. Generic quantum Fourier transforms. In *Proceedings of the Fifteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA’04)*, pages 778–787, 2004. See also ArXiv preprint quant-ph/0304064.
- [MRRS04] C. Moore, D. Rockmore, A. Russell, and L. Schulman. The power of basis selection in Fourier sampling: hidden subgroup problems in affine groups. In *Proceedings of the Fifteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA’04)*, pages 1113–1122, 2004. Journal version in preparation. Also: arXiv preprint quant-ph/0503095.
- [MRS05] C. Moore, A. Russell, and L. Schulman. The symmetric group defies strong Fourier sampling: Part I. ArXiv preprint quant-ph/0501056, 2005.
- [MZ04] M. Mosca and Ch. Zalka. Exact quantum Fourier transforms and discrete logarithm algorithms. *International Journal of Quantum Information*, 2(1):91–100, 2004. See also ArXiv preprint quant-ph/0301093.
- [PRB99] M. Püschel, M. Rötteler, and Th. Beth. Fast quantum Fourier transforms for a class of non-abelian groups. In *Proceedings Applied Algebra, Algebraic Algorithms and Error-Correcting Codes (AAECC-13)*, volume 1719 of *LNCS*, pages 148–159. Springer, 1999. See also ArXiv preprint quant-ph/980764.
- [Ser77] J. P. Serre. *Linear Representations of Finite Groups*. Springer, 1977.
- [WF89] W. Wootters and B. Fields. Optimal state-determination by mutually unbiased measurements. *Ann. Physics*, 191(2):363–381, 1989.