McEliece and Niederreiter Cryptosystems That Resist Quantum Fourier Sampling Attacks

Hang Dinh[1], Cristopher Moore[2,⋆], and Alexander Russell[3,⋆⋆]

[1] Indiana University South Bend
hdinh@cs.iusb.edu
[2] University of New Mexico, and Santa Fe Institute
moore@cs.unm.edu
[3] University of Connecticut
acr@cse.uconn.edu

Abstract. Quantum computers can break the RSA, El Gamal, and elliptic curve public-key cryptosystems, as they can efficiently factor integers and extract discrete logarithms. This motivates the development of *post-quantum* cryptosystems: classical cryptosystems that can be implemented with today's computers, that will remain secure even in the presence of quantum attacks.

In this article we show that the McEliece cryptosystem over *rational Goppa codes* and the Niederreiter cryptosystem over *classical Goppa codes* resist precisely the attacks to which the RSA and El Gamal cryptosystems are vulnerable—namely, those based on generating and measuring coset states. This eliminates the approach of strong Fourier sampling on which almost all known exponential speedups by quantum algorithms are based. Specifically, we show that the natural case of the Hidden Subgroup Problem to which McEliece-type cryptosystems reduce cannot be solved by strong Fourier sampling, or by any measurement of a coset state. To do this, we extend recent negative results on quantum algorithms for Graph Isomorphism to subgroups of the automorphism groups of linear codes.

This gives the first rigorous results on the security of the McEliece-type cryptosystems in the face of quantum adversaries, strengthening their candidacy for post-quantum cryptography. We also strengthen some results of Kempe, Pyber, and Shalev on the Hidden Subgroup Problem in S_n.

1 Introduction

If and when quantum computers are built, common public-key cryptosystems such as RSA, El Gamal, and elliptic curve cryptography will no longer be secure. Given that fact, the susceptibility or resistance of other well-studied public-key cryptosystems to quantum attacks is of fundamental interest. We present evidence for the strength of McEliece-type cryptosystems against quantum attacks, demonstrating that the quantum Fourier sampling attacks that cripple RSA and El Gamal do not apply to the McEliece or Niederreiter cryptosystems as long as the underlying code satisfies certain algebraic

⋆ This work was supported by the NSF under grants CCF-0829931, 0835735, and 0829917 and by the DTO under contract W911NF-04-R-0009.

⋆⋆ This work was supported by the NSF under grants 1117427 and 0835735 and by the DTO under contract W911NF-04-R-0009.

P. Rogaway (Ed.): CRYPTO 2011, LNCS 6841, pp. 761–779, 2011.

properties. While there are known classical attacks on these systems for the case of rational Goppa codes, our results also apply to the Niederreiter cryptosystem with classical Goppa codes, which to our knowledge is still believed to be classically secure. While our results do not rule out other quantum (or classical) attacks, they do demonstrate security precisely against the types of quantum algorithms that have proven so powerful for number-theoretic problems. We also strengthen some results of Kempe et al. [9] on subgroups of S_n reconstructible by Fourier sampling.

McEliece-type cryptosystems. The McEliece cryptosystem is a public-key cryptosystem proposed by McEliece in 1978 [13], conventionally built over Goppa codes. A dual variant of the system, proposed by Niederreiter [16], can provide slightly improved efficiency with equivalent security [10]. This dual system can additionally be used to construct a digital signature scheme [2], a shortcoming of the original system.

There are two basic types of attacks known against the McEliece-type cryptosystems: decoding attacks, and direct attacks on the private key. The former appears challenging, considering that the general decoding problem is NP-hard; indeed, historical confidence in the security of the McEliece system relies on the idea that this hardness can be retained for scrambled version of specific codes. This same intuition applies to quantum attacks: NP-hard problems are believed to be intractable, in general, for quantum computers and no significant quantum algorithmic developments appear to be directly relevant to these decoding problems. The latter—direct attacks on the key—can be successful on certain classes of linear codes, and is our focus. In a McEliece-type cryptosystem, the private key of a user Alice consists of three matrices: a $k \times n$ matrix M over a finite field $\mathbb{F}_{q^\ell}$, a $k \times k$ invertible matrix S over the field $\mathbb{F}_q$, and an $n \times n$ permutation matrix P. In the McEliece version, M is a generator matrix of a q-ary $[n,k]$-linear code (hence, $\ell = 1$), while in Niederreiter's dual system, M is a parity check matrix of a q-ary linear code of length n. The matrices S and P are selected randomly. Alice's public key consists of the matrix $M^* = SMP$. An adversary may attack the private key, attempting to recover the secret row "scrambler" S and the secret permutation P from M^* and M, assuming he already knew M.[1] As pointed out in [4], it crucial to keep S and P secret for the security of the McEliece system.

The security of these McEliece-type systems have received considerable attention in the literature, often focusing on particular choices for the underlying codes. Various classes of Goppa codes have received the greatest attention: along these lines, Sidelnokov and Shestakov's attack [23] can efficiently compute the matrices S and MP from the public matrix $M^* = SMP$ if the underlying code is a generalized Reed-Solomon code.[2] While this attack can reveal the structure of an alternative code, it does not reveal the secret permutation. An attack in which the secret permutation is revealed was proposed by Loidreau and Sendrier [11], using the Support Splitting Algorithm [21]. However, this attack only works with a very limited subclass of classical binary Goppa codes, namely those with a binary Goppa polynomial.

[1] Recovering the secret scrambler and the secret permutation is different from the Code Equivalence problem. The former finds a transformation between two equivalent codes, while the latter decides whether two linear codes are equivalent.

[2] We remark that the class of generalized Reed-Solomon codes is essentially equal to the class of rational Goppa codes.

Although the McEliece-type cryptosystems are efficient and still considered classically secure, at least with classical binary Goppa codes [4], they are rarely used in practice because of their comparatively large public key (see remark 8.33 in [14]). The discovery of successful quantum attacks on RSA and El Gamal, however, has changed the landscape: as suggested by Ryan [20] and Bernstein et al. [1], if "post-quantum" security guarantees can be made for the McEliece cryptosystem, this may compensate for its comparatively expensive computational demands.

Quantum Fourier sampling. Quantum Fourier Sampling (QFS) is the key ingredient in nearly all known efficient quantum algorithms for algebraic problems, including Shor's algorithms for factorization and discrete logarithm [22] and Simon's algorithm [24]. Shor's algorithm relies on quantum Fourier sampling over the cyclic group $\mathbb{Z}_N$, while Simon's algorithm uses quantum Fourier sampling over $\mathbb{Z}_2^n$. In general, these algorithms solve instances of the *Hidden Subgroup Problem* (HSP) over a finite group G. Given a function f on G whose level sets are left cosets of some unknown subgroup $H < G$, i.e., such that f is constant on each left coset of H and distinct on different left cosets, they find a set of generators for the subgroup H.

The standard approach to this problem treats f as a black box and applies f to a uniform superposition over G, producing the coset state $|cH\rangle = (1/\sqrt{|H|}) \sum_{h \in H} |ch\rangle$ for a random c. We then measure $|cH\rangle$ in a Fourier basis $\{|\rho, i, j\rangle\}$ for the space $\mathbb{C}[G]$, where ρ is an irrep[3] of G and i, j are row and column indices of a matrix $\rho(g)$. In the *weak* form of Fourier sampling, only the representation name ρ is measured, while in the *strong* form, both the representation name and the matrix indices are measured, the latter in a chosen basis. This produces probability distributions from which classical information can be extracted to recover the subgroup H. Moreover, since $|cH\rangle$ is block-diagonal in the Fourier basis, the optimal measurement of the coset state can always be described in terms of strong Fourier sampling.

Understanding the power of Fourier sampling in nonabelian contexts has been an ongoing project, and a sequence of negative results [6, 15, 7] have suggested that the approach is inherently limited when the underlying groups are rich enough. In particular, Moore, Russell, and Schulman [15] showed that over the symmetric group, even the strong form of Fourier sampling cannot efficiently distinguish the conjugates of most order-2 subgroups from each other or from the trivial subgroup. That is, for any $\sigma \in S_n$ with large support, and most $\pi \in S_n$, if $H = \{1, \pi^{-1}\sigma\pi\}$ then strong Fourier sampling, and therefore any measurement we can perform on the coset state, yields a distribution which is exponentially close to the distribution corresponding to $H = \{1\}$. This result implies that GRAPH ISOMORPHISM cannot be solved by the naive reduction to strong Fourier sampling. Hallgren et al. [7] strengthened these results, demonstrating that even entangled measurements on $o(\log n!)$ coset states yield essentially no information.

Kempe and Shalev [8] showed that weak Fourier sampling of single coset states in S_n cannot distinguish the trivial subgroup from larger subgroups H with polynomial size and non-constant minimal degree.[4] They conjectured, conversely, that if a subgroup $H < S_n$ can be distinguished from the trivial subgroup by weak Fourier sampling, then

[3] Throughout the paper, we write "irrep" as short for "irreducible representation."

[4] The minimal degree of a permutation group H is the minimal number of points moved by a non-identity element of H.

the minimal degree of H must be constant. Their conjecture was later proved by Kempe, Pyber, and Shalev [9].

Our contributions. To state our results, we say that a subgroup $H < G$ is *indistinguishable by strong Fourier sampling* if the conjugate subgroups $g^{-1}Hg$ cannot be distinguished from each other (or from the trivial subgroup) by measuring the coset state in an arbitrary Fourier basis. A precise definition is presented in Section 3.2. Since the optimal measurement of a coset state can always be expressed as an instance of strong Fourier sampling, these results imply that no measurement of a single coset state yields any useful information about H. Based on the strategy of Moore, Russell, and Schulman [15], we first develop a general framework, formalized in Theorem 1, to determine indistinguishability of a subgroup by strong Fourier sampling. We emphasize that their results cover the case where the subgroup has order two. Our principal contribution is to show how to extend their methods to more general subgroups.

We then apply this general framework to a class of semi-direct products $(\mathsf{GL}_k(\mathbb{F}_q) \times S_n) \wr \mathbb{Z}_2$, bounding the distinguishability for the HSP corresponding to the private-key attack on a McEliece-type cryptosystem, i.e., the problem of determining a secret scrambler S and a secret permutation P from $M^* = SMP$ and M. Our bound, given in Corollary 1 of Theorem 4, depends on the column rank[5] of the matrix M as well as the minimal degree and the size of the *automorphism group* $\mathrm{Aut}(M)$, where $\mathrm{Aut}(M)$ is defined in Subsection 4.2 as the set of all permutations P on the columns of M such that $M = SMP$ for some $S \in \mathsf{GL}_k(\mathbb{F}_q)$. In general, our result indicates that McEliece-type cryptosystems resist known attacks based on strong Fourier sampling if M has column rank at least $k - o(\sqrt{n})/\ell$, and the automorphism group $\mathrm{Aut}(M)$ has minimal degree $\Omega(n)$ and size $\mathrm{e}^{o(n)}$. In particular, generator matrices of rational Goppa codes and canonical parity check matrices of classical Goppa codes have good values for these quantities (see Lemma 3). The result is most interesting for classical Goppa codes, which are considered classically secure; the McEliece system over *rational* Goppa codes is subject to the Sidelnokov-Shestakov [23] attack.

While our main application is the security of the McEliece cryptosystem, we show in addition that our general framework is applicable to other classes of groups with simpler structure, including the symmetric group and the finite general linear group[6] $\mathsf{GL}_2(\mathbb{F}_q)$. For the symmetric group, we extend the results of [15] to larger subgroups of S_n. Specifically, we show that any subgroup $H < S_n$ with minimal degree $m \geq \Theta(\log|H|) + \omega(\log n)$ is indistinguishable by strong Fourier sampling over S_n. This partially extends the results of Kempe et al. [9], which apply only to weak Fourier sampling.

Remark 1. Our results show that the natural reduction of McEliece to a hidden subgroup problem yields negligible information about the secret key. Thus they rule out the direct analogue of the quantum attack that breaks, for example, RSA. Of course, our results do not rule out other quantum (or classical) attacks. Neither do they establish that a quantum algorithm for the McEliece cryptosystem would violate a natural hardness

[5] The column rank of M is understood to be over the field F_{q^ℓ}. Recall that the entries of the matrix M are in F_{q^ℓ}.

[6] The case of $\mathsf{GL}_2(\mathbb{F}_q)$ is omitted in this version for lack of space.

assumption, as do recent lattice cryptosystem constructions whose hardness is based on the Learning With Errors problem (e.g. Regev [18]). Nevertheless, they indicate that any such algorithm would have to involve significant new ideas beyond than those that have been proposed so far.

Summary of technical ideas. Let G be a finite group. We wish to establish general criteria for indistinguishability of subgroups $H < G$ by strong Fourier sampling. We begin with the general strategy, developed in [15], that controls the resulting probability distributions in terms of the representation-theoretic properties of G. In order to handle richer subgroups, however, we have to overcome some technical difficulties. Our principal contribution here is a "decoupling" lemma that allows us to handle the cross terms arising from pairs of nontrivial group elements.

Roughly, the approach (presented in Section 3.2) identifies two disjoint subsets, SMALL and LARGE, of irreps of G. The set LARGE consists of all irreps whose dimensions are no smaller than a certain threshold D. While D should be as large as possible, we also need to choose D small enough so that the set LARGE is large. In contrast, the representations in SMALL must have small dimension (much smaller than $\sqrt{D}$), and the set SMALL should be small or contain few irreps that appear in the decomposition of the tensor product representation $\rho \otimes \rho^*$ for any $\rho \in$ LARGE. In addition, any irrep ρ outside SMALL must have small normalized character $|\chi_\rho(h)|/d_\rho$ for any nontrivial element $h \in H$. If two such sets exist, and if $|H|$ is sufficiently small, we establish that H is indistinguishable by strong Fourier sampling over G.

In the case $G = S_n$, as in [15] we define SMALL as the set Λ_c of all Young diagrams whose top row or left column has length at least $(1-c)n$, and define LARGE by setting $D = n^{dn}$, for appropriate constants $0 < c, d < 1$. We show that any irrep outside SMALL has large dimension and therefore small normalized characters.

For the case $G = (\mathsf{GL}_k(\mathbb{F}_q) \times S_n) \wr \mathbb{Z}_2$ corresponding to McEliece-type cryptosystems, the normalized characters on the hidden subgroup K depend on the minimal degree of the automorphism group $\mathrm{Aut}(M) < S_n$. If we choose SMALL as the set of all irreps constructed from tensor product representations $\tau \times \lambda$ of $\mathsf{GL}_k(\mathbb{F}_q) \times S_n$ with $\lambda \in \Lambda_c$, then the "small" features of Λ_c will induce the "small" features of this set SMALL. Finally, $|K|$ depends on $|\mathrm{Aut}(M)|$ and the column rank of M. When M is a generator matrix of a rational Goppa code or a canonical parity check matrix of a classical Goppa code, $\mathrm{Aut}(M)$ lies inside the automorphism group of a rational Goppa code, which can be controlled using Stichtenoth's Theorem [25].

2 Hidden Subgroup Attacks on McEliece-type Cryptosystems

As mentioned in the Introduction, we consider an attack attempting to recover the secret scrambler S and permutation P from M and M^*. We frame the problem such an attacker needs to solve as follows:

Definition 1 (Scrambler-Permutation Problem). *Given two $k \times n$ matrices M and M^* with entries in a finite field containing $\mathbb{F}_q$ such that $M^* = SMP$ for some $S \in \mathsf{GL}_k(\mathbb{F}_q)$ and some $n \times n$ permutation matrix P, find such a pair (S,P).*

In the case where the matrix M is a generator matrix of a linear code over $\mathbb{F}_q$, the decision version of this problem is known as the CODE EQUIVALENCE problem, which is at least as hard as GRAPH ISOMORPHISM, although it is unlikely to be NP-complete [17]. This problem can be immediately recast as a Hidden Subgroup Problem (described below). We begin with a presentation of the problem as a Hidden Shift Problem:

Definition 2 (Hidden Shift Problem). *Let G be a finite group and Σ be a finite set. Given two functions $f_0 : G \to \Sigma$ and $f_1 : G \to \Sigma$ with the promise that there is an element $s \in G$ for which $f_1(x) = f_0(sx)$ for all $x \in G$, the problem is to determine such s by making queries to f_0 and f_1. An element s with this property is called a* left shift *from f_0 to f_1 (or, simply, a* shift*).*

The Scrambler-Permutation Problem can be immediately reduced to the Hidden Shift Problem over the group $G = \mathsf{GL}_k(\mathbb{F}_q) \times S_n$ by defining functions f_0 and f_1 on $\mathsf{GL}_k(\mathbb{F}_q) \times S_n$ so that for all $(S,P) \in \mathsf{GL}_k(\mathbb{F}_q) \times S_n$,

$$f_0(S,P) = S^{-1}MP\,, \qquad f_1(S,P) = S^{-1}M^*P\,. \tag{1}$$

Here and from now on, we identify each $n \times n$ permutation matrix with its corresponding permutation in S_n. Evidently, $SMP = M^*$ if and only if (S^{-1},P) is a shift from f_0 to f_1.

Next, following the standard approach to developing quantum algorithms for such problems, we reduce this Hidden Shift Problem on a group G to the Hidden Subgroup Problem on the wreath product $G \wr \mathbb{Z}_2 = G^2 \rtimes \mathbb{Z}_2$. Given two functions f_0 and f_1 on G, we define the function $f : G \wr \mathbb{Z}_2 \to \Sigma \times \Sigma$ as follows: for $(x,y) \in G^2$ and $b \in \mathbb{Z}_2$,

$$f((x,y),b) \stackrel{\text{def}}{=} \begin{cases} (f_0(x), f_1(y)) & \text{if } b = 0 \\ (f_1(y), f_0(x)) & \text{if } b = 1 \end{cases} \tag{2}$$

Now we would like to see that the Hidden Shift Problem is equivalent to determining the subgroup whose cosets are distinguished by f. Recall that a function f on a group G *distinguishes the right cosets* of a subgroup $H < G$ if for all $x,y \in G$, $f(x) = f(y) \iff yx^{-1} \in H$.

Definition 3. *Let f be a function on a group G. We say that f is* injective under right multiplication *if for all $x,y \in G$, $f(x) = f(y) \iff f(yx^{-1}) = f(1)$. Define the subset $G|_f \subseteq G$ as the level set containing the identity,*

$$G|_f \stackrel{\text{def}}{=} \{g \in G \mid f(g) = f(1)\}\,.$$

Proposition 1. *Let f be a function on a group G. If f distinguishes the right cosets of a subgroup $H < G$, then f must be injective under right multiplication and $G|_f = H$. Conversely, if f is injective under right multiplication, then $G|_f$ is a subgroup and f distinguishes the right cosets of the subgroup $G|_f$.*

Hence, the function f defined in (2) can distinguish the right cosets of some subgroup if and only if it is injective under right multiplication.

Lemma 1. *The function f defined in* (2) *is injective under right multiplication if and only if (1) f_0 is injective under right multiplication and (2) $f_1(x) = f_0(sx)$ for some s.*

The proof of this lemma is straightforward, so we omit it here.

Proposition 2. *Assume f_0 is injective under right multiplication. Let $H_0 = G|_{f_0}$ and s be a shift. Then the function f defined in* (2) *distinguishes right cosets of the following subgroup of $G \wr \mathbb{Z}_2$:*

$$G \wr \mathbb{Z}_2|_f = \big((H_0, s^{-1}H_0 s), 0\big) \cup \big((H_0 s, s^{-1}H_0), 1\big),$$

which has size $2|H_0|^2$. The set of all shifts from f_0 to f_1 is $H_0 s$.

If we can determine the hidden subgroup $K = G \wr \mathbb{Z}_2|_f$, we can find a shift by selecting an element of the form $((g_1, g_2), 1)$ from K. Then g_1 must belong to $H_0 s$, and so is a shift from f_0 to f_1.

Application to the Scrambler-Permutation problem. Returning to the Hidden Shift Problem over $G = \mathsf{GL}_k(\mathbb{F}_q) \times S_n$ corresponding to the Scrambler-Permutation problem, it is clear that the function f_0 defined in (1) is injective under right multiplication, and that

$$H_0 = \mathsf{GL}_k(\mathbb{F}_q) \times S_n|_{f_0} = \left\{(S, P) \in \mathsf{GL}_k(\mathbb{F}_q) \times S_n \mid S^{-1}MP = M\right\}.$$

The automorphism group of M is the projection of H_0 onto S_n, i.e.,

$$\mathrm{Aut}(M) = \left\{P \in S_n \mid \exists S : S^{-1}MP = M\right\}.$$

Note that each $P \in \mathrm{Aut}(M)$ has the same number of preimages $S \in \mathsf{GL}_k(\mathbb{F}_q)$ in this projection.

3 Quantum Fourier sampling (QFS)

3.1 Preliminaries and Notation

Fix a finite group G, abelian or non-abelian, and let $\widehat{G}$ denote the set of irreducible unitary representations, or "irreps" for short, of G. For each irrep $\rho \in \widehat{G}$, let V_ρ denote a vector space over $\mathbb{C}$ on which ρ acts so that ρ is a group homomorphism from G to the general linear group over V_ρ, and let d_ρ denote the dimension of V_ρ. For each ρ, we fix an orthonormal basis $B_\rho = \left\{\mathbf{b}_1, \ldots, \mathbf{b}_{d_\rho}\right\}$ for V_ρ. Then we can represent each $\rho(g)$ as a $d_\rho \times d_\rho$ unitary matrix whose j^{th} column is the vector $\rho(g)\mathbf{b}_j$.

Viewing the vector space $\mathbb{C}[G]$ as the regular representation of G, we can decompose $\mathbb{C}[G]$ into irreps as the direct sum $\bigoplus_{\rho \in \widehat{G}} V_\rho^{\oplus d_\rho}$. This has a basis $\{|\rho, i, j\rangle : \rho \in \widehat{G}, 1 \le i, j \le d_\rho\}$, where $\{|\rho, i, j\rangle \mid 1 \le i \le d_\rho\}$ is a basis for the j^{th} copy of V_ρ. Up to normalization, $|\rho, i, j\rangle$ corresponds to the i, j entry of the irrep ρ.

Definition 4. *The* Quantum Fourier transform *over G is the unitary operator, denoted F_G, that transforms a vector in $\mathbb{C}[G]$ from the point-mass basis $\{|g\rangle \mid g \in G\}$ into the basis given by the decomposition of $\mathbb{C}[G]$. For all $g \in G$,*

$$F_G |g\rangle = \sum_{\rho, i, j} \sqrt{\frac{d_\rho}{|G|}}\, \rho(g)_{i,j} \, |\rho, i, j\rangle,$$

where $\rho(g)_{ij}$ is the (i, j)-entry of the matrix $\rho(g)$. Alternatively, we can view $F_G|g\rangle$ as a block diagonal matrix consisting of the block $\sqrt{d_\rho/|G|}\,\rho(g)$ for each $\rho \in \widehat{G}$.

Notation. For each subset $X \subseteq G$, define $|X\rangle = (1/\sqrt{|X|}) \sum_{x \in X} |x\rangle$, which is the uniform superposition over X. For each $X \subseteq G$ and $\rho \in \widehat{G}$, define the operator $\Pi_X^\rho \stackrel{\text{def}}{=} \frac{1}{|X|} \sum_{x \in X} \rho(x)$, and let $\widehat{X}(\rho)$ denote the $d_\rho \times d_\rho$ matrix block at ρ in the quantum Fourier transform of $|X\rangle$, i.e.,

$$\widehat{X}(\rho) \stackrel{\text{def}}{=} \sqrt{\frac{d_\rho}{|G||X|}} \sum_{x \in X} \rho(x) = \sqrt{\frac{d_\rho |X|}{|G|}} \Pi_X^\rho .$$

Fact. If X is a subgroup of G, then Π_X^ρ is a projection operator. That is, $(\Pi_X^\rho)^\dagger = \Pi_X^\rho$ and $(\Pi_X^\rho)^2 = \Pi_X^\rho$.

Quantum Fourier Sampling (QFS) is a standard procedure based on the Quantum Fourier Transform to solve the Hidden Subgroup Problem (HSP) (see [12] for a survey). An instance of the HSP over G consists of a black-box function $f : G \to \{0,1\}^*$ such that $f(x) = f(y)$ if and only if x and y belong to the same left coset of H in G, for some subgroup $H \leq G$. The problem is to recover H using the oracle $O_f : |x, y\rangle \mapsto |x, y \oplus f(x)\rangle$. The general QFS procedure for this is the following:

1. Prepare a 2-register quantum state, the first in a uniform superposition of the group elements and the second with the value zero: $|\psi_1\rangle = (1/\sqrt{|G|}) \sum_{g \in G} |g\rangle |0\rangle$.
2. Query f, i.e., apply the oracle O_f, resulting in the state

$$|\psi_2\rangle = O_f |\psi_1\rangle = \frac{1}{\sqrt{|G|}} \sum_{g \in G} |g\rangle |f(g)\rangle = \frac{1}{\sqrt{|T|}} \sum_{\alpha \in T} |\alpha H\rangle |f(\alpha)\rangle$$

   where T is a transversal of H in G.
3. Measure the second register of $|\psi_2\rangle$, resulting in the state $|\alpha H\rangle |f(\alpha)\rangle$ with probability $1/|T|$ for each $\alpha \in T$. The first register of the resulting state is then $|\alpha H\rangle$ for some uniformly random $\alpha \in G$.
4. Apply the quantum Fourier transform over G to the coset state $|\alpha H\rangle$ observed at step 3:

$$F_G |\alpha H\rangle = \sum_{\rho \in \widehat{G}, 1 \leq i,j \leq d_\rho} \widehat{\alpha H}(\rho)_{i,j} |\rho, i, j\rangle .$$

5. (Weak) Observe the representation name ρ. (Strong) Observe ρ and matrix indices i, j.
6. Classically process the information observed from the previous step to determine the subgroup H.

Probability distributions produced by QFS. For a particular coset αH, the probability of measuring the representation ρ in the state $F_G |\alpha H\rangle$ is

$$P_{\alpha H}(\rho) = \|\widehat{\alpha H}(\rho)\|_F^2 = \frac{d_\rho |H|}{|G|} \mathrm{Tr}\left((\Pi_{\alpha H}^\rho)^\dagger \Pi_{\alpha H}^\rho \right) = \frac{d_\rho |H|}{|G|} \mathrm{Tr}\left(\Pi_H^\rho \right)$$

where $\mathrm{Tr}(A)$ denotes the trace of a matrix A, and $\|A\|_F := \sqrt{\mathrm{Tr}(A^\dagger A)}$ is the Frobenius norm of A. The last equality is due to the fact that $\Pi_{\alpha H}^\rho = \rho(\alpha) \Pi_H^\rho$ and that Π_H^ρ is an orthogonal projector.

Since there is no point in measuring the rows [6], we are only concerned with measuring the columns. As pointed out in [15], the optimal von Neumann measurement on a coset state can always be expressed in this form for some basis B_ρ. Conditioned on observing ρ in the state $F_G|\alpha H\rangle$, the probability of measuring a given $\mathbf{b} \in B_\rho$ is $\|\widehat{\alpha H}(\rho)\mathbf{b}\|^2$. Hence the conditional probability that we observe the vector $\mathbf{b}$, given that we observe the representation ρ, is then

$$P_{\alpha H}(\mathbf{b} \mid \rho) = \frac{\|\widehat{\alpha H}(\rho)\mathbf{b}\|^2}{P_{\alpha H}(\rho)} = \frac{\|\Pi^\rho_{\alpha H}\mathbf{b}\|^2}{\mathrm{Tr}\left(\Pi^\rho_H\right)} = \frac{\|\Pi^\rho_H\mathbf{b}\|^2}{\mathrm{Tr}\left(\Pi^\rho_H\right)}$$

where in the last equality, we use the fact that as $\rho(\alpha)$ is unitary, it preserves the norm of the vector $\Pi^\rho_H\mathbf{b}$.

The coset representative α is unknown and is uniformly distributed in T. However, both distributions $P_{\alpha H}(\rho)$ and $P_{\alpha H}(\mathbf{b} \mid \rho)$ are independent of α and are the same as those for the state $F_G|H\rangle$. Thus, in Step 5 of the QFS procedure above, we observe $\rho \in \widehat{G}$ with probability $P_H(\rho)$, and conditioned on this event, we observe $\mathbf{b} \in B_\rho$ with probability $P_H(\mathbf{b} \mid \rho)$.

If the hidden subgroup is trivial, $H = \{1\}$, the conditional probability distribution on B_ρ is uniform,

$$P_{\{1\}}(\mathbf{b} \mid \rho) = \frac{\|\Pi^\rho_{\{1\}}\mathbf{b}\|^2}{\mathrm{Tr}\left(\Pi^\rho_{\{1\}}\right)} = \frac{\|\mathbf{b}\|^2}{d_\rho} = \frac{1}{d_\rho}\,.$$

3.2 Distinguishability by QFS

We fix a finite group G and consider quantum Fourier sampling over G in the basis given by $\{B_\rho\}$. For a subgroup $H < G$ and for $g \in G$, let H^g denote the conjugate subgroup $g^{-1}Hg$. Since $\mathrm{Tr}\left(\Pi^\rho_H\right) = \mathrm{Tr}\left(\Pi^\rho_{H^g}\right)$, the probability distributions obtained by QFS for recovering the hidden subgroup H^g are

$$P_{H^g}(\rho) = \frac{d_\rho|H|}{|G|}\mathrm{Tr}\left(\Pi^\rho_H\right) = P_H(\rho) \quad \text{and} \quad P_{H^g}(\mathbf{b} \mid \rho) = \frac{\|\Pi^\rho_{H^g}\mathbf{b}\|^2}{\mathrm{Tr}\left(\Pi^\rho_H\right)}\,.$$

As $P_{H^g}(\rho)$ does not depend on g, weak Fourier sampling can not distinguish conjugate subgroups. Our goal is to point out that for certain nontrivial subgroup $H < G$, strong Fourier sampling can not efficiently distinguish the conjugates of H from each other or from the trivial one. Recall that the distribution $P_{\{1\}}(\cdot \mid \rho)$ obtained by performing strong Fourier sampling on the trivial hidden subgroup is the same as the uniform distribution U_{B_ρ} on the basis B_ρ. Thus, our goal can be boiled down to showing that the probability distribution $P_{H^g}(\cdot \mid \rho)$ is likely to be close to the uniform distribution U_{B_ρ} in total variation, for a random $g \in G$ and an irrep $\rho \in \widehat{G}$ obtained by weak Fourier sampling.

Definition 5. *We define the* distinguishability *of a subgroup H (using strong Fourier sampling over G), denoted $\mathscr{D}_H$, to be the expectation of the squared L_1-distance between $P_{H^g}(\cdot \mid \rho)$ and U_{B_ρ}:*

$$\mathscr{D}_H \stackrel{\text{def}}{=} \mathbb{E}_{\rho,g}\left[\|P_{H^g}(\cdot \mid \rho) - U_{B_\rho}\|_1^2\right]\,,$$

where ρ is drawn from $\widehat{G}$ according to the distribution $P_H(\rho)$, and g is chosen from G uniformly at random. We say that the subgroup H is indistinguishable *if $\mathscr{D}_H \leq \log^{-\omega(1)}|G|$.*

Note that if $\mathscr{D}_H$ is small, then the total variation distance between $P_{Hg}(\cdot \mid \rho)$ and U_{B_ρ} is small with high probability due to Markov's inequality: for all $\varepsilon > 0$,

$$\Pr_g\left[\|P_{Hg}(\cdot \mid \rho) - U_{B_\rho}\|_{t.v.} \geq \varepsilon/2\right] = \Pr_g\left[\|P_{Hg}(\cdot \mid \rho) - U_{B_\rho}\|_1^2 \geq \varepsilon^2\right] \leq \mathscr{D}_H/\varepsilon^2 .$$

In particular, if the subgroup H is indistinguishable by strong Fourier sampling, then for all constant $c > 0$,

$$\|P_{Hg}(\cdot \mid \rho) - U_{B_\rho}\|_{t.v.} < \log^{-c}|G|$$

with probability at least $1 - \log^{-c}|G|$ in both g and ρ. Our notion of indistinguishability is the direct analogue of that of Kempe and Shalev [8]. Focusing on weak Fourier sampling, they say that H is indistinguishable if $\|P_H(\cdot) - P_{\{1\}}(\cdot)\|_{t.v.} < \log^{-\omega(1)}|G|$.

Our main theorem below will serve as a general guideline for bounding the distinguishability of H. For this purpose we define, for each $\sigma \in \widehat{G}$, the *maximal normalized character of σ on H* as

$$\overline{\chi}_\sigma(H) \stackrel{\text{def}}{=} \max_{h \in H \setminus \{1\}} \frac{|\chi_\sigma(h)|}{d_\sigma} .$$

For each subset $S \subset \widehat{G}$, let

$$\overline{\chi}_{\overline{S}}(H) = \max_{\sigma \in \widehat{G} \setminus S} \overline{\chi}_\sigma(H) \quad \text{and} \quad d_S = \max_{\sigma \in S} d_\sigma .$$

In addition, for each reducible representation ρ of G, we let $I(\rho)$ denote the set of irreps of G that appear in the decomposition of ρ into irreps.

Theorem 1. *(*Main Theorem*) Suppose S is a subset of $\widehat{G}$. Let $D > d_S^2$ and $L = L_D \subset \widehat{G}$ be the set of all irreps of dimension at least D. Let*

$$\Delta = \Delta_{S,L} = \max_{\rho \in L} \left|S \cap I(\rho \otimes \rho^*)\right| . \tag{3}$$

Then the distinguishability of H is bounded by $\mathscr{D}_H \leq 4|H|^2 \left(\overline{\chi}_{\overline{S}}(H) + \Delta \frac{d_S^2}{D} + \frac{|\overline{L}|D^2}{|G|}\right)$.

Intuitively, the set S consists of irreps of small dimension, and L consists of irreps of large dimension. Moreover, we wish to have that the size of S is small while the size of L is large, so that most irreps are likely in L. In the cases where there are relatively few irreps, i.e., $|S| \ll D$ and $|\widehat{G}| \ll |G|$, we can simply upper bound Δ by $|S|$ and upper bound $|\overline{L}|$ by $|\widehat{G}|$.

We discuss the proof of this theorem in Section 5. Most details are relegated to the Appendix A.

4 Applications of the Main Theorem

In this section, we present applications of Theorem 1 to analyze strong Fourier sampling over certain non-abelian groups, including the symmetric group and the wreath product corresponding to the McEliece-type cryptosystems. Another application to the HSP over the groups $\mathsf{GL}_2(\mathbb{F}_q)$ is omitted for lack of space.

4.1 Strong Fourier Sampling over S_n

We focus now on the case where G is the symmetric group S_n. Recall that each irrep of S_n is in one-to-one correspondence to an integer partition $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_t)$ of n often given by a *Young diagram* of t rows in which the i^{th} row contains λ_i columns. The conjugate representation of λ is the irrep corresponding to the partition $\lambda' = (\lambda'_1, \lambda'_2, \ldots, \lambda'_{t'})$, obtained by flipping the Young diagram λ about the diagonal.

As in [15], we shall apply Roichman's upper bound [19] on normalized characters:

Theorem 2 (Roichman's Theorem [19]). *There exist constant $b > 0$ and $0 < q < 1$ so that for $n > 4$, for every $\pi \in S_n$, and for every irrep λ of S_n,*

$$\left|\frac{\chi_\lambda(\pi)}{d_\lambda}\right| \leq \left(\max\left(q, \frac{\lambda_1}{n}, \frac{\lambda'_1}{n}\right)\right)^{b\cdot \mathrm{supp}(\pi)}$$

where $\mathrm{supp}(\pi) = \#\{k \in [n] \mid \pi(k) \neq k\}$ *is the support of* π.

This bound works well for unbalanced Young diagrams. In particular, for a constant $0 < c < 1/4$, let Λ_c denote the collection of partitions λ of n with the property that either $\frac{\lambda_1}{n} \geq 1 - c$ or $\frac{\lambda'_1}{n} \geq 1 - c$, i.e., the Young diagram λ contains at least $(1-c)n$ rows or contains at least $(1-c)n$ columns. Then, Roichman's upper bound implies that for every $\pi \in S_n$ and $\lambda \notin \Lambda_c$, and a universal constant $\alpha > 0$,

$$\left|\frac{\chi_\lambda(\pi)}{d_\lambda}\right| \leq \mathrm{e}^{-\alpha\cdot\mathrm{supp}(\pi)}. \tag{4}$$

On the other hand, both $|\Lambda_c|$ and the maximal dimension of representations in Λ_c are small, as shown in the following Lemma of [15].

Lemma 2 (Lemma 6.2 in [15]). *Let $p(n)$ denote the number of integer partitions of n. Then $|\Lambda_c| \leq 2cn \cdot p(cn)$, and $d_\mu < n^{cn}$ for any $\mu \in \Lambda_c$.*

To give a more concrete bound for the size of Λ_c, we record the asymptotic formula for the partition function [5, pg. 45]: $p(n) \approx \mathrm{e}^{\pi\sqrt{2n/3}}/(4\sqrt{3}n) = \mathrm{e}^{O(\sqrt{n})}/n$ as $n \to \infty$.

Now we are ready to prove the main result of this section, an application of Theorem 1.

Theorem 3. *Let H be a nontrivial subgroup of S_n with minimal degree m, i.e., $m = \min_{\pi\in H\setminus\{1\}} \mathrm{supp}(\pi)$. Then for sufficiently large n, $\mathscr{D}_H \leq O(|H|^2\mathrm{e}^{-\alpha m})$.*

Proof. Let $2c < d < 1/2$ be constants. We will apply Theorem 1 by setting $S = \Lambda_c$ and $D = n^{dn}$. By Lemma 2, we have $d_S \leq n^{cn}$. Hence, the condition $2c < d$ guarantees that $D > d_S^2$. First, we need to bound the maximal normalized character $\overline{\chi}_{\overline{S}}(H)$. By (4), we have $\overline{\chi}_\mu(H) \leq \mathrm{e}^{-\alpha m}$ for all $\mu \in \widehat{S_n} \setminus S$. Hence, $\overline{\chi}_{\overline{S}}(H) \leq \mathrm{e}^{-\alpha m}$. To bound the second term in the upper bound of Theorem 1, as $\Delta \leq |S|$, it suffices to bound:

$$|S| \cdot \frac{d_S^2}{D} \leq 2cn \cdot p(cn) \cdot \frac{n^{2cn}}{n^{dn}} \leq \mathrm{e}^{O(\sqrt{n})} \cdot n^{(2c-d)n} \leq n^{-\gamma n}/2$$

for sufficiently large n, so long as $\gamma < d - 2c$. Now bounding the last term in the upper bound of Theorem 1: Since $|\overline{L_D}| \leq |\widehat{S_n}| = p(n)$ and $n! > n^n \mathrm{e}^{-n}$ by Stirling's approximation,

$$\frac{|\overline{L_D}|D^2}{|S_n|} \leq \frac{p(n)n^{2dn}}{n!} \leq \frac{\mathrm{e}^{O(\sqrt{n})}n^{2dn}}{n^n \mathrm{e}^{-n}} \leq \mathrm{e}^{O(n)}n^{(2d-1)n} \leq n^{-\gamma n}/2$$

for sufficiently large n, so long as $\gamma < 1 - 2d$. By Theorem 1, $\mathscr{D}_H \leq 4|H|^2(\mathrm{e}^{-\alpha m} + n^{-\gamma n})$.

Theorem 3 generalizes Moore, Russell, and Schulman's result [15] on strong Fourier sampling over S_n, which only applied in the case $|H| = 2$. To relate our result to the results of Kempe et al. [9], observe that since $\log|S_n| = \Theta(n \log n)$, the subgroup H is indistinguishable by strong Fourier sampling if $|H|^2 \mathrm{e}^{-\alpha m} \leq (n \log n)^{-\omega(1)}$ or, equivalently, if $m \geq (2/\alpha)\log|H| + \omega(\log n)$.

4.2 Applications to McEliece-type Cryptosystems

Our main application of Theorem 1 is to show the limitations of strong Fourier sampling in attacking the McEliece-type cryptosystems. Throughout this section, we fix parameters n, k, q of a McEliece-type cryptosystem, and fix the underlying $k \times n$ matrix M of the system. Here, M can be a generator matrix or a parity check matrix of the q-ary linear code used in the cryptosystem. Note that the entries of M are in a finite field $\mathbb{F}_{q^\ell} \supset \mathbb{F}_q$ (when M is a generator matrix of a q-ary linear code, we must have $\ell = 1$).

Recall that the canonical quantum attack against this McEliece cryptosystem involves the HSP over the wreath product group $(\mathsf{GL}_k(\mathbb{F}_q) \times S_n) \wr \mathbb{Z}_2$; the hidden subgroup in this case is

$$K = ((H_0, s^{-1}H_0 s), 0) \cup ((H_0 s, s^{-1}H_0), 1) \tag{5}$$

for some hidden element $s \in \mathsf{GL}_k(\mathbb{F}_q) \times S_n$. Here, H_0 is a subgroup of $\mathsf{GL}_k(\mathbb{F}_q) \times S_n$ given by

$$H_0 = \left\{ (A, P) \in \mathsf{GL}_k(\mathbb{F}_q) \times S_n \mid A^{-1}MP = M \right\}. \tag{6}$$

To understand the structure of the subgroup H_0, we define the *automorphism group* of M as $\mathrm{Aut}(M) \stackrel{\text{def}}{=} \left\{ P \in S_n \mid SMP = M \text{ for some } S \in \mathsf{GL}_k(\mathbb{F}_q) \right\}$. Note that $\mathrm{Aut}(M)$ is a subgroup of the symmetric group S_n and each element $(A, P) \in H_0$ must have $P \in \mathrm{Aut}(M)$. This allows us to control the maximal normalized characters on K through the minimal degree of $\mathrm{Aut}(M)$. Then applying Theorem 1, we show that

Theorem 4. *Assume $q^{k^2} \leq n^{an}$ for some constant $0 < a < 1/4$. Let m be the minimal degree of the automorphism group $\mathrm{Aut}(M)$. Then for sufficiently large n, the subgroup K defined in* (5) *has $\mathscr{D}_K \leq O(|K|^2 \mathrm{e}^{-\delta m})$, where $\delta > 0$ is a constant.*

The proof of Theorem 4 follows the technical ideas discussed in the Introduction. The details can be found in [3].

As $q^{k^2} \leq n^{an}$, we have $\log\left|(\mathsf{GL}_k(\mathbb{F}_q) \times S_n) \wr \mathbb{Z}_2\right| = O(\log n! + \log q^{k^2}) = O(n \log n)$. Hence, the subgroup K is indistinguishable if $|K|^2 \mathrm{e}^{-\delta m} \leq (n \log n)^{-\omega(1)}$. The size of the subgroup K is given by $|K| = 2|H_0|^2$, and $|H_0| = |\mathrm{Aut}(M)| \times |\mathrm{Fix}(M)|$, where

$\mathrm{Fix}(M) \stackrel{\text{def}}{=} \{S \in \mathsf{GL}_k(\mathbb{F}_q) \mid SM = M\}$ is the set of scramblers fixing M. To bound the size of $\mathrm{Fix}(M)$, we record an easy fact which can be obtained by the orbit-stabilizer formula:

Fact. Let r be the column rank of M. Then $|\mathrm{Fix}(M)| \le q^{\ell k(k-r)}$.

Proof. WLOG, assume the first r columns of M are $\mathbb{F}_{q^\ell}$-linearly independent, and each remaining column is an $\mathbb{F}_{q^\ell}$-linear combination of the first r columns. Let N be the $k \times r$ matrix consisting of the first r columns of M. Then we can decompose M as $M = (N \mid NA)$, where A is an $r \times (n-r)$ matrix with entries in $\mathbb{F}_{q^\ell}$. Clearly, $\mathrm{Fix}(M) = \mathrm{Fix}(N)$. Consider the action of $\mathsf{GL}_k(\mathbb{F}_{q^\ell})$ on the set of $k \times r$ matrices over $\mathbb{F}_{q^\ell}$. Under this action, the stabilizer of N contains $\mathrm{Fix}(N)$, and the orbit of the matrix N, denoted $\mathrm{Orb}(N)$, consists of all $k \times r$ matrices over $\mathbb{F}_{q^\ell}$ whose columns are $\mathbb{F}_{q^\ell}$-linearly independent. Thus, $|\mathrm{Orb}(N)| = (q^{\ell k} - 1)(q^{\ell k} - q^\ell) \dots (q^{\ell k} - q^{\ell(r-1)})$. By the orbit-stabilizer formula, we have

$$|\mathrm{Fix}(N)| \le \frac{|\mathsf{GL}_k(\mathbb{F}_{q^\ell})|}{|\mathrm{Orb}(N)|} = \frac{(q^{\ell k} - 1)(q^{\ell k} - q^\ell) \dots (q^{\ell k} - q^{\ell(k-1)})}{(q^{\ell k} - 1)(q^{\ell k} - q^\ell) \dots (q^{\ell k} - q^{\ell(r-1)})}$$
$$= (q^{\ell k} - q^{\ell r})(q^{\ell k} - q^{\ell(r+1)}) \cdots (q^{\ell k} - q^{\ell(k-1)}) \le q^{\ell k(k-r)}.$$

Corollary 1. *Assume $q^{k^2} \le n^{0.2n}$ and the automorphism group $\mathrm{Aut}(M)$ has minimal degree $\Omega(n)$. Let r be the column rank of M. Then the subgroup K defined in* (5) *has $\mathscr{D}_K \le |\mathrm{Aut}(M)|^4 q^{4\ell k(k-r)} \mathrm{e}^{-\Omega(n)}$. In particular, the subgroup K is indistinguishable if, further, $|\mathrm{Aut}(M)| \le \mathrm{e}^{o(n)}$ and $r \ge k - o(\sqrt{n})/\ell$.*

The constraint $q^{k^2} \le n^{0.2n}$ implies $\log|\mathsf{GL}_k(\mathbb{F}_q)| = O(n \log n)$, so Alice only needs to flip $O(n \log n)$ bits to pick a random S from $\mathsf{GL}_k(\mathbb{F}_q)$. Thus she needs only $O(n \log n)$ coin flips overall to generate her private key.

Application to the McEliece cryptosystem. Consider a McEliece cryptosystem using a q-ary linear $[n,k]$-code C, with parameters satisfying $q^{k^2} \le n^{0.2n}$. Since the automorphism group of the code C equals the automorphism group of its generator matrix, we can conclude that this McEliece cryptosystem resists the standard quantum Fourier sampling attack if the code C is *(i) well-scrambled*, i.e., it has a generator matrix of rank at least $k - o(\sqrt{n})$, and is *(ii) well-permuted*, i.e., its automorphism group has minimal degree at least $\Omega(n)$ and has size at most $\mathrm{e}^{o(n)}$. Recall that in terms of security, the Niederreiter system using $(n-k) \times n$ parity check matrices over $\mathbb{F}_q$ of the same code C is equivalent to the McEliece system using the code C [10].

Application to Goppa codes. We would like to point out that if M is a generator matrix of a rational Goppa code or a canonical parity check matrix of a classical Goppa code, it will give good bounds in Corollary 1. Specifically, we consider a matrix M over a finite field $\mathbb{F}_{q^\ell} \supset \mathbb{F}_q$ of the following form:

$$M = \begin{pmatrix} v_1 f_1(\alpha_1) & v_2 f_1(\alpha_2) & \cdots & v_n f_1(\alpha_n) \\ v_1 f_2(\alpha_1) & v_2 f_2(\alpha_2) & \cdots & v_n f_2(\alpha_n) \\ \vdots & \vdots & \ddots & \vdots \\ v_1 f_k(\alpha_1) & v_2 f_k(\alpha_2) & \cdots & v_n f_k(\alpha_n) \end{pmatrix} \tag{7}$$

where $v_1, \ldots, v_n$ are nonzero elements in the field $\mathbb{F}_{q^\ell}$, $(\alpha_1, \ldots, \alpha_n)$ is a list of distinct points in the projective line $\mathbb{F}_{q^\ell} \cup \{\infty\}$, and $f_1, \ldots, f_k$ are $\mathbb{F}_{q^\ell}$-linearly independent polynomials in $\mathbb{F}_{q^\ell}[X]$ of degree less than k (by convention, $f_i(\infty)$ is the X^{k-1}-coefficient of $f_i(X)$). Note that such a matrix M is a generator matrix of a rational Goppa $[n,k]$-code over the field $\mathbb{F}_{q^\ell}$, and is also a parity check matrix of a classical Goppa $[n, \geq n-\ell k]$-code over $\mathbb{F}_q$. To apply Corollary 1, we show the following properties of the matrix M:

Lemma 3. *The matrix M in the form of* (7) *has full rank (i.e., its column rank equals k), and* $\mathrm{Aut}(M)$ *has minimal degree at least $n-2$, and* $|\mathrm{Aut}(M)| \leq q^{3\ell}$.

Proof. We can show that M has full rank directly by decomposing M as $M = AVD$, where $A = (a_{ij})$ is an $k \times k$ invertible matrix with entry a_{ij} being the the X^{j-1}-coefficient of polynomial $f_i(X)$; V is a $k \times n$ Vandermonde matrix with (i,j)-entry being α_j^{i-1}; and D is an $n \times n$ diagonal matrix with v_i in the (i,i)-entry. Then the rank of M equals the rank of the Vandermonde matrix V, which has full rank.

Now we can view M as a generator matrix of a rational Goppa $[n,k]$-code R over the field $\mathbb{F}_{q^\ell}$. Then we have $\mathrm{Aut}(M) \subset \mathrm{Aut}(R)$, where $\mathrm{Aut}(R)$ is the automorphism group of the code R, that is, $\mathrm{Aut}(R) = \left\{ P \in S_n \mid SMP = M \text{ for some } S \in \mathsf{GL}_k(\mathbb{F}_{q^\ell}) \right\}$. Now we can apply Stichtenoth's Theorem [25] to control the automorphism group $\mathrm{Aut}(R)$.

Theorem 5 (Stichtenoth [25]). *Let $2 \leq k \leq n-2$. Then the automorphism group of any rational Goppa $[n,k]$-code over a field F is isomorphic to a subgroup of* $\mathrm{Aut}(F(x)/F)$.

On the other hand, we also have the useful fact that $\mathrm{Aut}(F(x)/F) \simeq \mathsf{PGL}_2(F)$. Hence, $\mathrm{Aut}(R)$ is isomorphic to a subgroup of the projective linear group $\mathsf{PGL}_2(\mathbb{F}_{q^\ell})$, which implies that $|\mathrm{Aut}(M)| \leq |\mathrm{Aut}(R)| \leq |\mathsf{PGL}_2(\mathbb{F}_{q^\ell})| \leq q^{3\ell}$.

To show that the minimal degree of $\mathrm{Aut}(M)$ is at least $n-2$, we view $\mathrm{Aut}(M) \subset \mathsf{PGL}_2(\mathbb{F}_{q^\ell})$, and observe that any transformation in $\mathsf{PGL}_2(\mathbb{F}_{q^\ell})$ that fixes at least three distinct projective lines must be the identity. Q.E.D.

Hence, classical Goppa codes or rational Goppa codes are good choices for the security of McEliece-type cryptosystems against standard quantum Fourier sampling attacks. Since the rational Goppa codes are broken (classically) by the Sidelnokov-Shestakov [23] structural attack, we shall focus on the classical Goppa codes, which remain secure given suitable choice of parameters.

Application to Neiderreiter systems with classical Goppa codes. Consider a classical q-ary Goppa code C constructed by a support list of distinct points $\alpha_1, \ldots, \alpha_n \in \mathbb{F}_{q^\ell}$ and a Goppa polynomial $g(X) \in \mathbb{F}_{q^\ell}[X]$ of degree k. This code has dimension $k' \geq n - \ell k$. More importantly, it has $k \times n$ parity check matrices in the form of (7) in which $v_j = 1/g(\alpha_j)$ (see [26]), we refer to those matrices as *canonical parity check matrices* of the classical Goppa code C. By Corollary 1 and Lemma 3, *the Niederreiter cryptosystem using $k \times n$ canonical parity check matrices of this code C resists the known quantum attack, provided $q^{k^2} \leq n^{0.2n}$ and $q^{3\ell} \leq \mathrm{e}^{o(n)}$.* As pointed out in [4], this Niederreiter system is secure under the Sidelnokov-Shestakov attack. We remark, however, that the

security of this Niederreiter cryptosystem may *not* be equivalent to that of the McEliece cryptosystem using the same code C (unless $k' = n - \ell k$ as discussed below), since the equivalence showed in [10] only applies to the Niederreiter cryptosystem using a parity check matrix over the subfield $\mathbb{F}_q$.

Setting the parameters. We discuss the parameters for classical Goppa codes that meet our security requirement. Traditionally, the code length is chosen as $n = q^\ell$, then our parameter setting requires only one constraint, $k^2 \le 0.2n\ell$, which imposes that the code C must have large dimension, i.e., $k' \ge n - \ell k \ge n - \sqrt{0.2n}(\log_q n)^{3/2}$.

Now we compare our parameter setting with practical parameters suggestion. In most McEliece cryptosystems considered in practice, classical binary Goppa codes are used, that is $q = 2$ and $n = 2^\ell$. The code is also designed so that it has dimension $k' = n - \ell k$ and minimal distance $d \ge 2t + 1$, where $t \ll n$ is a predetermined parameter indicating the number of errors the code can correct. For those systems, the original parameters suggested by McEliece were $(n = 1024, k' \ge 524, t = 50)$, which would meet our requirement as long as the dimension k' is chosen to be slightly larger ($k' \ge 572$). The parameters $(n = 1024, k' = 524, t = 50)$, which can be broken in just 7 days by a cluster of 200 CPUs under Bernstein et al.'s attack [1], clearly do not meet our requirement. An optimal choice of parameters for the Goppa code which maximizes the adversary's work factor was recommended to be $(n = 1024, k' \ge 644, t = 38)$ (see Note 8.32 in [14]). Bernstein et al. [1] suggested two other sets of parameters, $(n = 2048, k' = 1751, t = 27)$ and $(n = 1632, k' = 1269, t = 34)$, that achieve the standard security against all known (classical) attacks. All of these parameters meet our requirement. Well, of course, these parameters were recommended for the original McEliece, or for the equivalent Neiderreiter system that uses parity check matrices over the subfield $\mathbb{F}_2$ with $n - k' = \ell k$ rows. However, if we view each element in $\mathbb{F}_{q^\ell}$ as a vector of dimension ℓ over the subfield $\mathbb{F}_q$, then a $k \times n$ canonical parity check matrix over $\mathbb{F}_{q^\ell}$ can be viewed as a $\ell k \times n$ parity check matrix over $\mathbb{F}_q$.

5 Bounding Distinguishability

We now sketch the proof for the main theorem (Theorem 1). Fixing a nontrivial subgroup $H < G$, we want to upper bound $\mathscr{D}_H$. Let us start with bounding the expectation over the random group element $g \in G$, for a fixed irrep $\rho \in \widehat{G}$:

$$E_H(\rho) \stackrel{\text{def}}{=} \mathbb{E}_g \left[\| P_{Hg}(\cdot \mid \rho) - U_{B_\rho} \|_1^2 \right] .$$

Obviously we always have $E_H(\rho) \le 4$. More interestingly, we have

$$\begin{aligned} E_H(\rho) &= \mathbb{E}_g \left[\left(\sum_{\mathbf{b} \in B_\rho} \left| P_{Hg}(\mathbf{b} \mid \rho) - \frac{1}{d_\rho} \right| \right)^2 \right] \\ &\le \mathbb{E}_g \left[d_\rho \sum_{\mathbf{b} \in B_\rho} \left(P_{Hg}(\mathbf{b} \mid \rho) - \frac{1}{d_\rho} \right)^2 \right] \quad \text{(by Cauchy-Schwarz)} \\ &= d_\rho \sum_{\mathbf{b} \in B_\rho} \mathrm{Var}_g[P_{Hg}(\mathbf{b} \mid \rho)] \qquad (\text{since } \mathbb{E}_g[P_{Hg}(\mathbf{b} \mid \rho)] = \frac{1}{d_\rho}) \end{aligned}$$

$$= \frac{d_\rho}{\mathrm{Tr}(\Pi_H^\rho)^2} \sum_{\mathbf{b}\in B_\rho} \mathrm{Var}_g\left[\|\Pi_{Hg}^\rho \mathbf{b}\|^2\right]. \tag{8}$$

The equation $\mathbb{E}_g[P_{Hg}(\mathbf{b} \mid \rho)] = 1/d_\rho$ can be shown using *Schur's lemma*.

From (8), we are motivated to bound the variance of $\|\Pi_{Hg}^\rho \mathbf{b}\|^2$ when g is chosen uniformly at random. We provide an upper bound that depends on the projection of the vector $\mathbf{b}\otimes\mathbf{b}^*$ onto irreducible subspaces of $\rho\otimes\rho^*$, and on maximal normalized characters of σ on H for all irreps σ appearing in the decomposition of $\rho\otimes\rho^*$. Recall that the representation $\rho\otimes\rho^*$ is typically reducible and can be written as an orthogonal direct sum of irreps $\rho\otimes\rho^* = \bigoplus_{\sigma\in\widehat{G}} a_\sigma\sigma$, where $a_\sigma \geq 0$ is the multiplicity of σ. Then $I(\rho\otimes\rho^*)$ consists of σ with $a_\sigma > 0$, and we let $\Pi_\sigma^{\rho\otimes\rho^*}$ denote the projection operator whose image is $a_\sigma\sigma$, that is, the subspace spanned by all copies of σ. Our upper bound given in Lemma 4 below generalizes the bound given in Lemma 4.3 of [15], which only applies to subgroups H of order 2.

Lemma 4. *(*DECOUPLING LEMMA*) Let ρ be an irrep of G. Then for any vector $\mathbf{b}\in V_\rho$,*

$$\mathrm{Var}_g\left[\|\Pi_{Hg}^\rho \mathbf{b}\|^2\right] \leq \sum_{\sigma\in I(\rho\otimes\rho^*)} \overline{\chi}_\sigma(H) \left\|\Pi_\sigma^{\rho\otimes\rho^*}(\mathbf{b}\otimes\mathbf{b}^*)\right\|^2.$$

Back to our goal of bounding $E_H(\rho)$ using the bound in Lemma 4, the strategy will be to separate irreps appearing in the decomposition of $\rho\otimes\rho^*$ into two groups, those with small dimension and those with large dimension, and treat them differently. If d_σ is large, we shall rely on bounding $\overline{\chi}_\sigma(H)$. If d_σ is small, we shall control the projection given by $\Pi_\sigma^{\rho\otimes\rho^*}$ using the following lemma which was proved implicitly in [15]:

Lemma 5. *For any irrep σ, we have $\sum_{\mathbf{b}\in B_\rho} \left\|\Pi_\sigma^{\rho\otimes\rho^*}(\mathbf{b}\otimes\mathbf{b}^*)\right\|^2 \leq d_\sigma^2$.*

The method discussed above for bounding $E_H(\rho)$ is culminated into Lemma 6 below.

Lemma 6. *Let $\rho\in\widehat{G}$ be arbitrary and $S\subset\widehat{G}$ be any subset of irreps that does not contain ρ. Then*

$$E_H(\rho) \leq 4|H|^2\left(\overline{\chi}_{\overline{S}}(H) + |S\cap I(\rho\otimes\rho^*)|\,\frac{d_S^2}{d_\rho}\right).$$

To apply this lemma, we should choose the subset S such that $d_S^2 \ll d_\rho$, that is, S should consist of small dimensional irreps. Then applying Lemma 6 for all irreps ρ of large dimension, we can prove our general main theorem straightforwardly.

The detailed proofs of the main theorem and the decoupling lemma are put in Appendix A. The proof for Lemma 6 is omitted for lack of space. See [3] for a full technical version.

References

1. Bernstein, D.J., Lange, T., Peters, C.: Attacking and defending the mcEliece cryptosystem. In: Buchmann, J., Ding, J. (eds.) PQCrypto 2008. LNCS, vol. 5299, pp. 31–46. Springer, Heidelberg (2008)

2. Courtois, N., Finiasz, M., Sendrier, N.: How to achieve a mcEliece-based digital signature scheme. In: Boyd, C. (ed.) ASIACRYPT 2001. LNCS, vol. 2248, pp. 157–174. Springer, Heidelberg (2001)
3. Dinh, H., Moore, C., Russell, A.: The McEliece cryptosystem resists quantum Fourier sampling attacks, preprint (2010), http://arxiv.org/abs/1008.2390
4. Engelbert, D., Overbeck, R., Schmidt, A.: A summary of McEliece-type cryptosystems and their security. J. Math. Crypt. 1, 151–199 (2007)
5. Fulton, W., Harris, J.: Representation Theory - A First Course. Springer-Verlag, New York Inc., Heidelberg (1991)
6. Grigni, M., Schulman, J., Vazirani, M., Vazirani, U.: Quantum mechanical algorithms for the nonabelian hidden subgroup problem. Combinatorica 24(1), 137–154 (2004)
7. Hallgren, S., Moore, C., Rötteler, M., Russell, A., Sen, P.: Limitations of quantum coset states for graph isomorphism. In: STOC 2006: Proceedings of the Thirty-Eighth Annual ACM Symposium on Theory of Computing, pp. 604–617 (2006)
8. Kempe, J., Shalev, A.: The hidden subgroup problem and permutation group theory. In: SODA 2005: Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 1118–1125 (2005)
9. Kempe, J., Pyber, L., Shalev, A.: Permutation groups, minimal degrees and quantum computing. Groups, Geometry, and Dynamics 1(4), 553–584 (2007), http://xxx.lanl.gov/abs/quant-ph/0607204
10. Li, Y.X., Deng, R.H., Wang, X.M.: On the equivalence of McElieces and Niederreiters public-key cryptosystems. IEEE Transactions on Information Theory 40(1), 271–273 (1994)
11. Loidreau, P., Sendrier, N.: Weak keys in the McEliece public-key cryptosystem. IEEE Transactions on Information Theory 47(3), 1207–1212 (2001)
12. Lomont, C.: The hidden subgroup problem - review and open problems (2004), http://arXiv.org:quantph/0411037
13. McEliece, R.J.: A public-key cryptosystem based on algebraic coding theory. JPL DSN Progress Report, 114–116 (1978)
14. Menezes, A.J., van Oorschot, P.C., Vanstone, S.A.: Handbook of applied cryptography. CRC Press, Boca Raton (1996)
15. Moore, C., Russell, A., Schulman, L.J.: The symmetric group defies strong quantum Fourier sampling. SIAM Journal of Computing 37, 1842–1864 (2008)
16. Niederreiter, H.: Knapsack-type cryptosystems and algebraic coding theory. Problems of Control and Information Theory. Problemy Upravlenija i Teorii Informacii 15(2), 159–166 (1986)
17. Petrank, E., Roth, R.M.: Is code equivalence easy to decide? IEEE Transactions on Information Theory 43(5), 1602–1604 (1997), doi:10.1109/18.623157
18. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. In: STOC 2005: Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing, pp. 84–93 (2005)
19. Roichman, Y.: Upper bound on the characters of the symmetric groups. Invent. Math. 125(3), 451–485 (1996)
20. Ryan, J.A.: Excluding some weak keys in the McEliece cryptosystem. In: Proceedings of the 8th IEEE Africon, pp. 1–5 (2007)
21. Sendrier, N.: Finding the permutation between equivalent linear codes: the support splitting algorithm. IEEE Transactions on Information Theory 46(4), 1193–1203 (2000)
22. Shor, P.W.: Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. SIAM Journal on Computing 26, 1484–1509 (1997)
23. Sidelnikov, V.M., Shestakov, S.O.: On insecurity of cryptosystems based on generalized Reed-Solomon codes. Discrete Mathematics and Applications 2(4), 439–444 (1992)

24. Simon, D.R.: On the power of quantum computation. SIAM J. Comput. 26(5), 1474–1483 (1997)
25. Stichtenoth, H.: On automorphisms of geometric Goppa codes. Journal of Algebra 130, 113–121 (1990)
26. van Lint, J.H.: Introduction to coding theory, 2nd edn. Springer, Heidelberg (1992)

Appendix A Proofs for the Main Theorem

Proof of the Decoupling Lemma

Proof (Proof of Lemma 4). Fix a vector $\mathbf{b} \in V_\rho$. To simplify notations, we shall write Π_g as shorthand for Π^ρ_{Hg}, and write $g\mathbf{b}$ for $\rho(g)\mathbf{b}$. For any $g \in G$, we have

$$\begin{aligned}\|\Pi_g \mathbf{b}\|^2 &= \langle \Pi_g \mathbf{b}, \Pi_g \mathbf{b} \rangle = \langle \mathbf{b}, \Pi_g \mathbf{b} \rangle \\ &= \frac{1}{|H|} \left(\langle \mathbf{b}, \mathbf{b} \rangle + \sum_{h \in H \setminus \{1\}} \langle \mathbf{b}, g^{-1} h g \mathbf{b} \rangle \right).\end{aligned}$$

Let $S_g = \sum_{h \in H \setminus \{1\}} \langle \mathbf{b}, g^{-1} h g \mathbf{b} \rangle$. Then

$$\mathrm{Var}_g \left[\|\Pi_g \mathbf{b}\|^2 \right] = \frac{\mathrm{Var}_g[S_g]}{|H|^2} = \frac{\mathbb{E}_g\left[S_g^2\right] - \mathbb{E}_g[S_g]^2}{|H|^2}.$$

To bound the variance, we upper bound S_g^2 for all $g \in G$. Since S_g is real, applying Cauchy-Schwarz inequality, we have

$$S_g^2 = \left| \sum_{h \in H \setminus \{1\}} \langle \mathbf{b}, g^{-1} h g \mathbf{b} \rangle \right|^2 \leq (|H| - 1) \left(\sum_{h \in H \setminus \{1\}} \left| \langle \mathbf{b}, g^{-1} h g \mathbf{b} \rangle \right|^2 \right).$$

As in Lemma 4.2 of [15], one can express the second moment of the inner product $\langle \mathbf{b}, g^{-1} h g \mathbf{b} \rangle$ in terms of the projection of $\mathbf{b} \otimes \mathbf{b}^*$ into the irreducible constituents of the tensor product representation $\rho \otimes \rho^*$. Specifically, for any $h \in G$, we have

$$\mathbb{E}_g \left[|\langle \mathbf{b}, g^{-1} h g \mathbf{b} \rangle|^2 \right] = \sum_{\sigma \in I(\rho \otimes \rho^*)} \frac{\chi_\sigma(h)}{d_\sigma} \left\| \Pi_\sigma^{\rho \otimes \rho^*} (\mathbf{b} \otimes \mathbf{b}^*) \right\|^2.$$

It follows that

$$\begin{aligned}\mathrm{Var}_g \left[\|\Pi^\rho_{Hg} \mathbf{b}\|^2 \right] &\leq \frac{|H| - 1}{|H|^2} \sum_{h \in H \setminus \{1\}} \mathbb{E}_g \left[\left| \langle \mathbf{b}, g^{-1} h g \mathbf{b} \rangle \right|^2 \right] \\ &\leq \mathbb{E}_{h \in H \setminus \{1\}} \left[\sum_{\sigma \in I(\rho \otimes \rho^*)} \frac{\chi_\sigma(h)}{d_\sigma} \left\| \Pi_\sigma^{\rho \otimes \rho^*} (\mathbf{b} \otimes \mathbf{b}^*) \right\|^2 \right] \\ &\leq \sum_{\sigma \in I(\rho \otimes \rho^*)} \overline{\chi}_\sigma(H) \left\| \Pi_\sigma^{\rho \otimes \rho^*} (\mathbf{b} \otimes \mathbf{b}^*) \right\|^2.\end{aligned}$$

Proof of the Main Theorem

Proof (Proof of Theorem 1:). For any $\rho \in L$, since $d_\rho \geq D > d_S^2$, we must have $\rho \notin S$. By Lemma 6,

$$E_H(\rho) \leq 4|H|^2 \left(\overline{\chi}_{\overline{S}}(H) + \Delta \frac{d_S^2}{D} \right) \quad \text{for all } \rho \in L.$$

Combining this with the fact that $E_H(\rho) \leq 4$ for all $\rho \notin L$, we obtain

$$\mathscr{D}_H = \mathbb{E}_\rho[E_H(\rho)] \leq 4|H|^2 \left(\overline{\chi}_{\overline{S}}(H) + \Delta \frac{d_S^2}{D} \right) + 4\mathrm{Pr}_\rho[\rho \notin L] .$$

To complete the proof, it remains to bound $\mathrm{Pr}_\rho[\rho \notin L]$. Since $\mathrm{Tr}(\Pi_H^\rho) \leq d_\rho$, we have

$$P(\rho) = \frac{d_\rho |H|}{|G|} \mathrm{Tr}(\Pi_H^\rho) \leq \frac{d_\rho^2 |H|}{|G|} .$$

Since $d_\rho < D$ for all $\rho \in \widehat{G} \setminus L$, it follows that

$$\mathrm{Pr}_\rho[\rho \notin L] = \sum_{\rho \notin L} P(\rho) \leq \frac{|\overline{L}| D^2 |H|}{|G|} \leq \frac{|\overline{L}| D^2 |H|^2}{|G|} .$$