

The McEliece Cryptosystem Resists Quantum Fourier Sampling Attacks

Hang Dinh
Indiana University South Bend
hdinh@cs.iusb.edu

Cristopher Moore
University of New Mexico
and Santa Fe Institute
moore@cs.unm.edu

Alexander Russell
University of Connecticut
acr@cse.uconn.edu

Considering that common public-key cryptosystems such as RSA and El Gamal are insecure against quantum attacks, the susceptibility of other well-studied public-key systems to such attacks is naturally of fundamental interest. In this article we present evidence for the strength of the McEliece cryptosystem against quantum attacks, demonstrating that the quantum Fourier sampling attacks that cripple RSA and El Gamal do not apply to the McEliece system coupled with *well-permuted*, *well-scrambled* linear codes. While our results do not rule out other quantum (or classical) attacks, they do demonstrate security against the hidden subgroup methods that have proven so powerful for computational number theory. Additionally, we partially extend results of Kempe et al. [7] concerning the subgroups of S_n reconstructible by quantum Fourier sampling.

The McEliece cryptosystem. This public-key cryptosystem was proposed by McEliece in 1978 [9], and is typically built over Goppa codes. There are two basic types of attacks known against the McEliece cryptosystem: ciphertext only attacks, and attacks on the private key. The former is unlikely to work because it relies on solving the general decoding problem, which is NP-hard. The latter can be successful on certain classes of linear codes, and is our focus. In the McEliece cryptosystem, the private key of a user Alice consists of three matrices: a $k \times n$ generator matrix M of a hidden q -ary $[n, k]$ -linear code, an invertible $k \times k$ matrix A over the finite field \mathbb{F}_q , and an $n \times n$ permutation matrix P . Both matrices A and P are selected randomly. Alice’s public key includes the matrix $M^* = AMP$, which is a generator matrix of a linear code equivalent to the secret code. An adversary may attack the private key by first computing the secret generator matrix M , and then computing¹ the secret row “scrambler” A and the secret permutation P .

There have been some successful attacks on McEliece-type public-key systems. A notable one is Sidelnokov and Shestakov’s attack [14], which efficiently computes the matrices A and MP from the public matrix AMP , in the case that the secret code is a generalized Reed-Solomon (GRS) code. Note that this attack does not reveal the secret permutation. An attack in which the secret permutation is revealed was proposed by Loidreau and Sendrier [8]. However, this attack only works with a very limited subclass of classical binary Goppa codes, namely those with a binary generator polynomial.

Although the McEliece cryptosystem is efficient and still considered (classically) secure [3], it is rarely used in practice because of the comparatively large public key (see remark 8.33 in [10]). The discovery of successful quantum attacks on RSA and El Gamal, however, have changed the landscape: as suggested by Ryan [12] and Bernstein et al. [1], the McEliece cryptosystem could become a “post-quantum” alternative to RSA.

¹Recovering 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.

Quantum Fourier sampling. Quantum Fourier Sampling (QFS) is a key ingredient in most efficient algebraic quantum algorithms, including Shor’s algorithms for factorization and discrete logarithm [13] and Simon’s algorithm [15]. In particular, 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 irrep² 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. 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 [4, 11, 5] have suggested that the approach is inherently limited when the underlying groups are rich enough. In particular, Moore, Russell, and Schulman [11] 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 the GRAPH ISOMORPHISM cannot be solved by the naive reduction to strong Fourier sampling. Hallgren et al. [5] strengthened these results, demonstrating that even entangled measurements on $o(\log n!)$ coset states result in essentially information-free outcome distributions. Kempe and Shalev [6] showed that weak Fourier sampling single coset states in S_n cannot distinguish the trivial subgroup from larger subgroups H with polynomial size and non-constant minimal degree.³ They conjectured, conversely, that if a subgroup $H < S_n$ can be distinguished from the trivial subgroup by weak Fourier sampling, then the minimal degree of H must be constant. Their conjecture was later proved by Kempe, Pyber, and Shalev [7].

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. 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 [11], we first develop a general framework 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 $(\text{GL}_k(\mathbb{F}_q) \times S_n) \wr \mathbb{Z}_2$, bounding the distinguishability for the HSP corresponding to the private-key attack on the McEliece cryptosystem, i.e., the problem of determining A and P from M^* and M . Our bound depends on the minimal degree and the size of the automorphism group of the secret code, as well as on the column rank of the secret generator matrix. In

²Throughout the paper, we write “irrep” as short for “irreducible representation”.

³The minimal degree of a permutation group H is the minimal number of points moved by a non-identity element of H .

particular, the rational Goppa codes have good values for these quantities, i.e., they have small automorphism groups with large minimal degree, and have generator matrices of full rank. In general, our result indicates that the McEliece cryptosystem resists all known attacks based on strong Fourier sampling if its secret q -ary $[n, k]$ -code (i) is *well-permuted*, i.e., its automorphism group has minimal degree $\Omega(n)$ and size $e^{o(n)}$, and (ii) is *well-scrambled*, i.e., it has a generator matrix of rank at least $k - o(\sqrt{n})$. Here, we assume $q^{k^2} \leq n^{0.2n}$, which implies $\log |\mathrm{GL}_k(\mathbb{F}_q)| = O(n \log n)$, so that Alice only needs to flip $O(n \log n)$ bits to pick a random matrix A from $\mathrm{GL}_k(\mathbb{F}_q)$. Thus she needs only $O(n \log n)$ coin flips overall to generate her private key.

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 $\mathrm{GL}_2(\mathbb{F}_q)$. For the symmetric group, we extend the results of [11] 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 . In other words, if the conjugates of H can be distinguished from each other—or from the trivial subgroup—by strong Fourier sampling, then the minimum degree of H must be $O(\log |H|) + O(\log n)$. This partially extends the results of Kempe et al. [7], which apply only to weak Fourier sampling.

We go on to demonstrate another application of our general framework for the general linear group $\mathrm{GL}_2(\mathbb{F}_q)$, giving the first negative result regarding the power of strong Fourier sampling over $\mathrm{GL}_2(\mathbb{F}_q)$. We show that any subgroup $H < \mathrm{GL}_2(\mathbb{F}_q)$ that does not contain non-identity scalar matrices and has order $|H| \leq q^\delta$ for some $\delta < 1/2$ is indistinguishable by strong Fourier sampling. Examples of such subgroups are those generated by a constant number of triangular unipotent matrices.

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 [11], 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 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 $\rho \in \text{LARGE}$. 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 there are such two sets SMALL and LARGE, and if the order of H is sufficiently small, then H is indistinguishable by strong Fourier sampling over G .

In the case $G = S_n$, we choose SMALL as the set Λ_c of all Young diagrams with at least $(1 - c)n$ rows or at least $(1 - c)n$ columns, and set $D = n^{dn}$, for reasonable constants $0 < c, d < 1$. For this case, we use the same techniques as in [11].

For the case $G = (\mathrm{GL}_k(\mathbb{F}_q) \times S_n) \wr \mathbb{Z}_2$ corresponding to the McEliece cryptosystem, the normalized characters on the hidden subgroup K depend on the minimal degree of the automorphism group $\mathrm{Aut}(C)$, where C is the secret code. Moreover, $|K|$ depends on $|\mathrm{Aut}(C)|$ and the column rank of the secret generator matrix. Now we can choose SMALL as the set of all irreps constructed from tensor product representations $\tau \times \lambda$ of $\mathrm{GL}_k(\mathbb{F}_q) \times S_n$ with $\lambda \in \Lambda_c$. Then the “small” features of Λ_c will induce the “small” features of this set SMALL. To show that any irrep outside SMALL has small normalized characters on K , we show that any Young diagram λ outside Λ_c has large dimension. See [2] for a full technical version.

References

- [1] Daniel J. Bernstein, Tanja Lange, and Christiane Peters. Attacking and defending the McEliece cryptosystem. In *PQCrypto '08: Proceedings of the 2nd International Workshop on Post-Quantum Cryptography*, pages 31–46, Berlin, Heidelberg, 2008. Springer-Verlag. ISBN 978-3-540-88402-6.
- [2] Hang Dinh, Cristopher Moore, and Alexander Russell. The McEliece cryptosystem resists quantum Fourier sampling attacks, 2010. URL <http://arxiv.org/abs/1008.2390>. Preprint.
- [3] D. Engelbert, R. Overbeck, and A. Schmidt. A summary of McEliece-type cryptosystems and their security. *J. Math. Crypt.*, 1:151199, 2007.
- [4] Michelangelo Grigni, J. Schulman, Monica Vazirani, and Umesh Vazirani. Quantum mechanical algorithms for the nonabelian hidden subgroup problem. *Combinatorica*, 24(1):137–154, 2004.
- [5] Sean Hallgren, Cristopher Moore, Martin Rötteler, Alexander Russell, and Pranab Sen. Limitations of quantum coset states for graph isomorphism. In *STOC '06: Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 604–617, 2006.
- [6] Julia Kempe and Aner Shalev. The hidden subgroup problem and permutation group theory. In *SODA '05: Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 1118–1125, 2005.
- [7] Julia Kempe, Laszlo Pyber, and Aner Shalev. Permutation groups, minimal degrees and quantum computing. *Groups, Geometry, and Dynamics*, 1(4):553–584, 2007. URL <http://xxx.lanl.gov/abs/quant-ph/0607204>.
- [8] Pierre Loidreau and Nicolas Sendrier. Weak keys in the McEliece public-key cryptosystem. *IEEE Transactions on Information Theory*, 47(3):1207–1212, 2001.
- [9] R.J. McEliece. A public-key cryptosystem based on algebraic coding theory. *JPL DSN Progress Report*, pages 114–116, 1978.
- [10] A.J. Menezes, P.C. van Oorschot, and S.A. Vanstone. *Handbook of applied cryptography*. CRC Press, 1996.
- [11] Cristopher Moore, Alexander Russell, and Leonard J. Schulman. The symmetric group defies strong quantum Fourier sampling. *SIAM Journal of Computing*, 37:1842–1864, 2008.
- [12] John A. Ryan. Excluding some weak keys in the McEliece cryptosystem. In *Proceedings of the 8th IEEE Africon*, pages 1–5, 2007.
- [13] Peter. W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26:1484–1509, 1997.
- [14] V. M. Sidelnikov and S. O. Shestakov. On insecurity of cryptosystems based on generalized Reed-Solomon codes. *Discrete Mathematics and Applications*, 2(4):439–444, 1992.
- [15] Daniel R. Simon. On the power of quantum computation. *SIAM J. Comput.*, 26(5):1474–1483, 1997.