

Quantum Fourier transforms for extracting hidden linear structures in finite fields*

J. Niel de Beaudrap[†] Richard Cleve[†] John Watrous[§]

Department of Computer Science
University of Calgary
Calgary, Alberta, Canada T2N 1N4

Abstract

We propose a definition for quantum Fourier transforms in settings where the algebraic structure is that of a finite field, and show that they can be performed efficiently by a quantum computer. Using these finite field quantum Fourier transforms, we obtain the strongest separation between quantum and classical query complexity known to date—specifically, we define a problem that requires $\Omega(2^{r/2})$ queries in the classical (bounded error) case, but can be solved exactly with a single query in the quantum case using a polynomial number (in n) of auxiliary operations. Finally, we consider quantum Fourier transforms over arbitrary finite rings, and give efficient quantum circuits for implementing quantum Fourier transforms for the particular case of rings of matrices over finite fields.

1 Introduction

Quantum Fourier transforms are conventionally associated with finite groups. We consider suitable definitions for quantum Fourier transforms (QFTs) in settings where there is a different algebraic structure, namely, finite rings. Our focus is on the case of finite fields, for which we give an explicit definition of a QFT, show that it can be implemented efficiently by quantum circuits, and demonstrate an application regarding query complexity.

In order to discuss what a suitable definition for a QFT over a finite field (or arbitrary finite ring) would be, let us briefly digress and consider properties of the standard quantum Fourier transform modulo m (for given $m > 1$). Recall that the QFT over \mathbb{Z}_m is defined as a unitary mapping

$$|x\rangle \mapsto \frac{1}{\sqrt{m}} \sum_{y=0}^{m-1} (e^{2\pi i/m})^{xy} |y\rangle$$

for each $x, y \in \mathbb{Z}_m$. The simplest case is the Hadamard transform, which can be viewed as the QFT over \mathbb{Z}_2 . The Hadamard transform has an interesting property in relation to the two-qubit controlled-NOT operation: conjugating a controlled-NOT gate with $H \otimes H$ (a Hadamard transform on each of two qubits) results in a controlled-NOT gate with its orientation inverted. In the language of quantum circuits, this property is expressed in Figure 1.

*Research partially supported by Canada's NSERC.

[†]Email: jn@cpsc.ucalgary.ca

[‡]Email: cleve@cpsc.ucalgary.ca

[§]Email: jwatrous@cpsc.ucalgary.ca




Figure 1: The control/target inversion property.

A natural extension of the controlled-NOT gate acting on quantum registers whose classical states correspond to elements of \mathbb{Z}_m for arbitrary $m > 1$ is the controlled-ADD $_r$ gate, acting as follows:

$$|x\rangle|y\rangle \mapsto |x\rangle|y + rx\rangle.$$

Here, the first register is the *control*, and the second register is the *target*. Each value for r gives a different gate. It is straightforward to verify that a similar property to that displayed in Figure 1 holds when we conjugate a controlled-ADD $_r$ gate appropriately by QFTs. Specifically, conjugating a controlled-ADD $_r$ by $F_m \otimes F_m^\dagger$ for F_m denoting the QFT modulo m switches the control and target.

Now, let us generalize this *control/target inversion* property from \mathbb{Z}_m to arbitrary finite rings. Let $(R, +, \cdot, 0, 1)$ be any ring, and assume we have a unique classical representation for each element of R . First we define a natural generalization to R -valued data of controlled-ADD gates. For $r \in R$, define the *controlled-ADD $_r$* (*c-ADD $_r$* for short) gate as the unitary transformation that maps the pair of R -valued quantum states $|x\rangle|y\rangle$ ($x, y \in R$) to $|x\rangle|y + rx\rangle$. If the ring is not commutative, we also refer to the above as the *left c-ADD $_r$* gate, and define a *right c-ADD $_r$* gate as the mapping $|x\rangle|y\rangle \mapsto |x\rangle|y + xr\rangle$. We introduce notation for these gates in quantum circuits in Figure 2.



Figure 2: Notation for left controlled-ADD $_r$ gates and right controlled-ADD $_r$ gates (respectively).

Call a unitary operation F on R -valued quantum states a *quantum Fourier transform (QFT) over R* if it satisfies the following property. Conjugating a left c-ADD $_r$ gate with $F \otimes F^\dagger$ results in a right c-ADD $_r$ gate with an inverted orientation. Such an F is not generally unique. In pictures, we call any F that satisfies the property in Figure 3 a QFT with respect to R .




Figure 3: The control/target inversion property for rings.

While it may be natural to consider the above property as being a defining property of a QFT, this by itself is of limited utility—to make the definition useful we must explicitly describe the

action of the transform on any given input. We do this first for the case that R is a finite field $GF(p^n)$. We show that, for any fixed prime p , the QFT for $GF(p^n)$ can be computed with $O(n^2)$ basic operations, and for large p that the QFT for $GF(p^n)$ can be approximated to within accuracy ε with $O(n^2(\log p)^2) + nC(p, \varepsilon/n)$ operations, where $C(p, \varepsilon)$ denotes the circuit size required to implement the QFT modulo p to within accuracy ε . Recently, Hales and Hallgren [9] have shown that $C(p, \varepsilon) \in O(\log p \log \log p + \log p \log 1/\varepsilon)$ for a reasonable range of ε . We also consider QFTs for the ring of $m \times m$ matrices over a given finite field. In this case, the cost of performing the QFT is $O(m^2)$ times the cost of performing the QFT on the finite field.

Our main application of these new QFTs is to obtain the strongest quantum vs. classical separation known to date for the black-box setting. We define the *hidden linear structure* black-box problem and show that:

- In the classical setting, $\Omega(2^{n/2})$ queries to the black-box are necessary to solve the problem, even allowing bounded-error probabilistic techniques.
- In the quantum setting, a single query suffices to solve the problem exactly, and the auxiliary operations are very simple; they consist of $O(n)$ Hadamard gates followed by $O(n^2)$ classical gate operations after a measurement is made.

We have recently learned that W. van Dam and S. Hallgren [6] have independently proposed the definition for QFTs over finite fields that appears in this paper. They have applied such transforms in the context of black-box problems that are different from ours, called the “shifted quadratic character problems”.

The remainder of this paper has the following organization. In Section 2 we propose an explicit definition for the QFT over any finite field, prove that it satisfies the control/target inversion property, and analyze the complexity of the transform. In Section 3 we consider the consequences of our circuits for the QFT over a finite field for the query complexity model, specifically in regard to the hidden linear structure problem. In Section 4 we extend the definition of the QFT to rings of matrices over finite fields. We conclude with Section 5, which compares our quantum vs. classical query separations with others.

2 Quantum Fourier transforms for finite fields

In this section we propose a definition for quantum Fourier transforms over finite fields, and prove that the definition satisfies the properties discussed in Section 1.

We assume the reader is familiar with basic concepts regarding finite fields and computations over finite fields (see, for instance, [5, 7, 10]). As usual, we let $GF(q)$ denote the finite field having $q = p^n$ elements for some prime p . We assume that a monic, irreducible polynomial $f(Z) = Z^n - \sum_{j=0}^{n-1} a_j Z^j \in GF(p)[Z]$ is fixed, and that elements of $GF(q)$ are represented as polynomials in $GF(p)[Z]$ modulo $f(Z)$ in the usual way. We will write $x = (x_0, \dots, x_{n-1})$ to denote the field element corresponding to $x_0 + x_1 Z + \dots + x_{n-1} Z^{n-1}$, and we write \vec{x} to denote the column vector $[x_0, \dots, x_{n-1}]^T$.

Definition 2.1 Let $\phi : GF(q) \rightarrow GF(p)$ be any nonzero linear mapping (viewing $GF(q)$ as an n dimensional vector space over $GF(p)$ as above). Then we define the *quantum Fourier transform*

(QFT) over $GF(q)$ relative to ϕ (denoted $F_{q,\phi}$) as follows. For each $x \in GF(q)$,

$$F_{q,\phi} : |x\rangle \mapsto \frac{1}{\sqrt{q}} \sum_{y \in GF(q)} \omega^{\phi(xy)} |y\rangle,$$

where $\omega = e^{2\pi i/p}$ (and we extend $F_{q,\phi}$ to arbitrary quantum states by linearity).

If ϕ is not made explicit then we assume ϕ refers to the trace.

Theorem 1 For $q = p^n$ and any nonzero linear mapping $\phi : GF(q) \rightarrow GF(p)$, $F_{q,\phi}$ is unitary and satisfies the control/target inversion property of Fig. 3.

Proof: It suffices to show that $F_{q,\phi}^\dagger F_{q,\phi} |x\rangle = |x\rangle$ for every $x \in GF(q)$. We have

$$\begin{aligned} F_{q,\phi}^\dagger F_{q,\phi} |x\rangle &= F_{q,\phi}^\dagger \frac{1}{\sqrt{q}} \sum_{y \in GF(q)} \omega^{\phi(xy)} |y\rangle = \frac{1}{q} \sum_{y \in GF(q)} \sum_{z \in GF(q)} \omega^{\phi(xy) - \phi(yz)} |z\rangle \\ &= \sum_{z \in GF(q)} \left(\frac{1}{q} \sum_{y \in GF(q)} \omega^{\phi(y(x-z))} \right) |z\rangle = |x\rangle, \end{aligned}$$

following from the fact that $\phi(w)$ must be uniformly distributed over $GF(p)$ as w ranges over $GF(q)$ (since ϕ is linear and not identically zero).

Next let us verify that the control/target inversion property discussed in Section 1 holds, namely that for A_r and B_r defined by $A_r|x\rangle|y\rangle = |x\rangle|y+rx\rangle$ and $B_r|x\rangle|y\rangle = |x+ry\rangle|y\rangle$ we have

$$(F_{q,\phi}^\dagger \otimes F_{q,\phi}) A_r (F_{q,\phi} \otimes F_{q,\phi}^\dagger) = B_r.$$

(See Figure 3.) To prove this relation holds, let us define

$$|\psi_x\rangle = F_{q,\phi}|x\rangle = \frac{1}{\sqrt{q}} \sum_{y \in GF(q)} \omega^{\phi(xy)} |y\rangle$$

for each $x \in GF(q)$, and note that for S_w defined by $S_w|x\rangle = |x+w\rangle$ we have

$$S_w|\psi_{-x}\rangle = \frac{1}{\sqrt{q}} \sum_{y \in GF(q)} \omega^{-\phi(xy)} |y+w\rangle = \frac{1}{\sqrt{q}} \sum_{y \in GF(q)} \omega^{-\phi(xy-xw)} |y\rangle = \omega^{\phi(xw)} |\psi_{-x}\rangle.$$

Now, for each $x, y \in GF(q)$ we have

$$\begin{aligned} (F_{q,\phi}^\dagger \otimes F_{q,\phi}) A_r (F_{q,\phi} \otimes F_{q,\phi}^\dagger) |x\rangle|y\rangle &= (F_{q,\phi}^\dagger \otimes F_{q,\phi}) A_r \left(\frac{1}{\sqrt{q}} \sum_{z \in GF(q)} \omega^{\phi(xz)} |z\rangle |\psi_{-y}\rangle \right) \\ &= (F_{q,\phi}^\dagger \otimes F_{q,\phi}) \left(\frac{1}{\sqrt{q}} \sum_{z \in GF(q)} \omega^{\phi(xz)} \omega^{\phi(yrz)} |z\rangle |\psi_{-y}\rangle \right) \\ &= (F_{q,\phi}^\dagger \otimes F_{q,\phi}) |\psi_{x+ry}\rangle |\psi_{-y}\rangle \\ &= |x+ry\rangle|y\rangle \\ &= B_r|x\rangle|y\rangle \end{aligned}$$

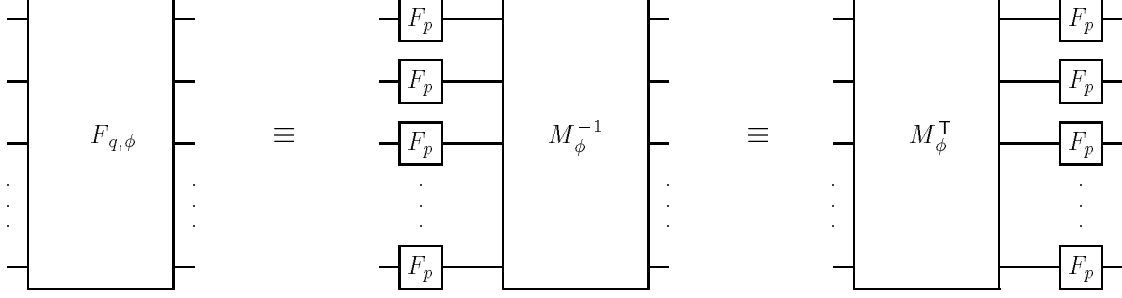


Figure 4: Equivalent circuits for $F_{q,\phi}$

as required. ■

Next we describe quantum circuits for performing $F_{q,\phi}$ and analyze their complexity. Let $C(p, \varepsilon)$ denote the minimum size of a quantum circuit approximating the quantum Fourier transform modulo p to within accuracy ε . Note that $C(p, 0) \in O(p^2 \log p)$ [1] and, for $\varepsilon > 0$, $C(p, \varepsilon) \in O(\log p \log \log p + \log p \log 1/\varepsilon)$ when $\varepsilon \in \Omega(1/p)$ [9].

Theorem 2 *For $q = p^n$ and any nonzero linear mapping $\phi : GF(q) \rightarrow GF(p)$, $F_{q,\phi}$ can be explicitly computed with accuracy ε by a quantum circuit of size $O(n^2(\log p)^2) + nC(p, \varepsilon/n)$.*

It should be noted that, when $p = 2$ (or any constant), the QFT circuit size is $O(n^2)$.

Proof of Theorem 2: First, we note that for any choice of ϕ (linear and nonzero), there exists an $n \times n$ matrix M_ϕ over $GF(p)$ such that $\phi(xy) = \vec{x}^\top M_\phi \vec{y}$. For any given ϕ we show how to obtain such a matrix M_ϕ explicitly below. We note that M_ϕ may be computed classically and of course is fixed for any choice of ϕ . Thus, computation of M_ϕ does not contribute to the size of a circuit for $F_{q,\phi}$ (although it can easily be computed in polynomial time as shown below). It is straightforward to prove that M_ϕ must be invertible. Now we have

$$F_{q,\phi}|x\rangle = \frac{1}{\sqrt{q}} \sum_{y \in GF(q)} \omega^{\vec{x}^\top M_\phi \vec{y}} |y\rangle = \frac{1}{\sqrt{q}} \sum_{y \in GF(q)} \omega^{\vec{x}^\top \vec{y}} |M_\phi^{-1} y\rangle = \frac{1}{\sqrt{q}} \sum_{y \in GF(q)} \omega^{(M_\phi^\top \vec{x})^\top \vec{y}} |y\rangle$$

Thus, we have

$$F_{q,\phi} = M_\phi^{-1} (F_p \otimes \cdots \otimes F_p) = (F_p \otimes \cdots \otimes F_p) M_\phi^\top,$$

where F_p denotes the usual quantum Fourier transform modulo p and, for $A \in \{M_\phi^{-1}, M_\phi^\top\}$, we identify A with the reversible operation that maps $|x\rangle$ to $|Ax\rangle$. This relation is illustrated in Figure 4.

Now let us return to the question of determining the matrix M_ϕ corresponding to a given ϕ . First, note that multiplication of field elements satisfies

$$(z_0, \dots, z_{n-1}) = (x_0, \dots, x_{n-1}) \cdot (y_0, \dots, y_{n-1})$$

where

$$z_i = \vec{x}^\top B_i \vec{y} \tag{1}$$

for a certain sequence of $n \times n$ matrices B_0, \dots, B_{n-1} over $GF(p)$. We now construct a sequence B_0, \dots, B_{n-1} that satisfies Eq. 1. To do this, it will be helpful to review the notion of *Hankel matrices*. An $n \times n$ Hankel matrix A is a matrix of the form

$$A = \begin{bmatrix} t_0 & t_1 & t_2 & \cdots & t_{n-1} \\ t_1 & t_2 & t_3 & \cdots & t_n \\ t_2 & t_3 & t_4 & \cdots & t_{n+1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ t_{n-1} & t_n & t_{n+1} & \cdots & t_{2n-2} \end{bmatrix}. \quad (2)$$

That is, the “anti-diagonals” each contain only one number (or, equivalently, $A[i, j]$ depends only on $i + j$). The Hankel matrix in Eq. 2 will be denoted $\text{Hankel}(t_0, t_1, \dots, t_{2n-2})$.

We have

$$Z^n \equiv \sum_{j=0}^{n-1} a_j Z^j \pmod{f(Z)}.$$

Write $a_j^{(0)} = a_j$ for $j = 0, \dots, n-1$. We will actually need numbers $a_j^{(k)}$ (for $j = 0, \dots, n-1$, $k = 0, \dots, n-2$) such that

$$Z^{n+k} \equiv \sum_{j=0}^{n-1} a_j^{(k)} Z^j \pmod{f(Z)}.$$

These numbers are easy to obtain. Define an $n \times n$ matrix V as follows:

$$V = \begin{bmatrix} 0 & 0 & \cdots & 0 & a_0 \\ 1 & 0 & \cdots & 0 & a_1 \\ 0 & 1 & \cdots & 0 & a_2 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 1 & a_{n-1} \end{bmatrix}$$

Then

$$\left[a_0^{(k)}, \dots, a_{n-1}^{(k)} \right]^\top = V^k [a_0, \dots, a_{n-1}]^\top = V^{k+1} [0, \dots, 0, 1]^\top.$$

Finally, we can describe the matrices B_0, \dots, B_{n-1} . For each $i = 0, \dots, n-1$,

$$B_i = \text{Hankel} \left(\delta_{0,i}, \delta_{1,i}, \dots, \delta_{n-1,i}, a_i^{(0)}, a_i^{(1)}, \dots, a_i^{(n-2)} \right).$$

(Here, $\delta_{i,j}$ is the Kronecker- δ symbol.) A straightforward computation reveals that this choice for B_0, \dots, B_{n-1} satisfies Eq. 1. It is also not hard to show that these matrices B_0, \dots, B_{n-1} are the only matrices satisfying Eq. 1, and that each B_i is necessarily invertible.

Now, since $\phi : GF(q) \rightarrow GF(p)$ is linear and not identically zero, we must have $\phi(x) = \sum_{i=0}^{n-1} \lambda_i x_i$ for each $x \in GF(q)$. At this point we see that $\phi(xy) = \vec{x}^\top M_\phi \vec{y}$ for $M_\phi = \sum_{i=0}^{n-1} \lambda_i B_i$. Equivalently, we have

$$M_\phi = \text{Hankel} \left(\lambda_0, \dots, \lambda_{n-1}, \sum_{i=0}^{n-1} \lambda_i a_i^{(0)}, \dots, \sum_{i=0}^{n-1} \lambda_i a_i^{(n-2)} \right).$$

■

3 The hidden linear structure problem

Define the *hidden linear structure* problem¹ over a ring R as follows. In the classical version, one is given a black-box that maps $(x, y) \in R \times R$ to $(x, \pi(y + rx))$, where π is an arbitrary permutation on the elements of R and $r \in R$. Analogously, in the quantum case, one is given a black-box performing the unitary transformation that maps $|x\rangle|y\rangle$ ($x, y \in R$) to $|x\rangle|\pi(y + rx)\rangle$. In both cases, the goal is to extract the value of r .

We show that, in the classical case, whenever R is a field, $\Omega(\sqrt{|R|})$ queries are necessary to solve this problem, even with bounded error. On the other hand, we show that, in the quantum case, for any ring R (including the case of a field), if F and F^\dagger operations (satisfying the control/target inversion property) can be performed then a single quantum query is sufficient to solve this problem exactly. As shown in Section 2, for the specific field $GF(2^n)$, the QFT can be performed exactly with only $O(n^2)$ basic operations (which are Hadamard gates and controlled-NOT gates). Moreover, for the specific field $GF(2^n)$, the algorithm solving the hidden linear structure problem can be streamlined so as to consist of $O(n)$ Hadamard gates, one query and $O(n^2)$ classical post-processing after a measurement is made.

We first note that when R is *not* a field the classical query complexity of the hidden linear structure problem can be much less than $\Omega(\sqrt{|R|})$. For example, when $R = \mathbb{Z}_{2^n}$, there is a simple procedure solving the hidden linear structure problem with only $n+1$ queries. It begins by querying $(0, 0)$ and $(2^{n-1}, 0)$, yielding $\pi(0)$ and $\pi(r2^{n-1})$ respectively. If $\pi(0) = \pi(r2^{n-1})$ then r is even; otherwise r is odd. Thus, two queries reduce the number of possibilities for r by a factor of 2. If r is even then the next query is $(2^{n-2}, 0)$, yielding $\pi(r2^{n-2})$, which determines whether $r \bmod 4$ is 0 or 2. If r is odd then the next query is $O(2^{n-2}, 2^n - 2^{n-2})$, yielding $\pi(2^n - 2^{n-2} + r2^{n-2})$, which determines whether $r \bmod 4$ is 1 or 3. This process can be continued so as to deduce r after $n+1$ queries.

Theorem 3 *When R is a field, $\Omega(\sqrt{|R|})$ queries are necessary to solve the hidden linear structure problem within error probability $\frac{1}{2}$.*

Proof: The lower bound proof is similar to that for Simon's problem [12]. Set both $r \in R$ and π (a permutation on R) randomly, according to the uniform distribution. Consider the information obtained after k queries $(x_1, y_1), \dots, (x_k, y_k)$ (without loss of generality, the queries are all distinct). If, for some $i \neq j$, the outputs of the i^{th} and j^{th} queries collide in that $\pi(y_i + rx_i) = \pi(y_j + rx_j)$, then $y_i + rx_i = y_j + rx_j$, which implies that the value of r can be determined as

$$r = \frac{y_i - y_j}{x_j - x_i} \quad (3)$$

(note that $x_j - x_i \neq 0$, since this would imply that $(x_i, y_i) = (x_j, y_j)$). On the other hand, if there are no collisions among the outputs of all k queries then all that can be deduced about r is that

$$r \neq \frac{y_i - y_j}{x_j - x_i} \quad (4)$$

for all $1 \leq i < j \leq k$. This leaves $|R| - \frac{k(k-1)}{2}$ values for r , which are equally likely by symmetry. Thus, the probability of a collision occurring at the k^{th} query given that no collisions have occurred

¹Note that this problem is *not* the same as the *hidden linear structure* problem considered by Boneh and Lipton [3]. In our problem, the term *linear structure* refers to the ring properties of R , whereas Boneh and Lipton's problem concerns the structure of certain periodic functions from the additive group \mathbb{Z}^k to some arbitrary range S .

in previous queries is $\frac{k-1}{|R|-k+1}$. It follows that the probability of a collision occurring at all during the first l queries is bounded above by

$$\sum_{k=1}^l \frac{k-1}{|R|-k+1} \leq \frac{l(l-1)}{2(|R|-l+2)}. \quad (5)$$

If this probability is to be greater than or equal to $\frac{1}{2}$ then $l \geq \sqrt{|R|+2} \in \Omega(\sqrt{|R|})$ must hold. ■

Theorem 4 *For any finite ring R , if F and F^\dagger can be performed then a single query is sufficient to solve the hidden linear structure problem exactly.*

Proof: The quantum procedure is to initialize the state of two R -valued registers to $|0\rangle|1\rangle$ (where 0 and 1 are respectively the additive and multiplicative identities of the ring) and perform the following operations:

1. Apply $F \otimes F^\dagger$.
2. Query the black box.
3. Apply $F^\dagger \otimes F$.

Then the state of the first register is measured.

Let us trace through the evolution of the state of the registers during the execution of the above algorithm. To facilitate this, define $|\psi_x\rangle = F|x\rangle$ and U_π as the unitary transformation that maps $|x\rangle$ to $|\pi(x)\rangle$ ($x \in R$). Then the state after each step is:

1. $|\psi_0\rangle|\psi_{-1}\rangle$
2. $|\psi_r\rangle U_\pi |\psi_{-1}\rangle$
3. $|r\rangle F U_\pi |\psi_{-1}\rangle$

Therefore, the output of the algorithm will be r . ■

In order to obtain the maximum classical query complexity via Theorem 3, it is desirable to set R to a field. The quantum query complexity will always be 1; however, the number of auxiliary operations necessary may vary, depending on the cost of performing the QFTs. When $R = \mathbb{Z}_p$ (where p is prime) the best procedure that we are aware of for performing the QFT exactly is $O(p^2 \log p)$, which is exponential in the number of bits of p . On the other hand, when $R = GF(2^n)$, the QFT can be computed exactly with only $O(n)$ Hadamard gates plus $O(n^2)$ c-NOT gates. Furthermore, using the structure of the decompositions given in Figure 4, the algorithm solving the hidden linear structure problem can be streamlined so as to consist of $O(n)$ Hadamard gates, one query, $O(n)$ Hadamard gates, a measurement, followed by $O(n^2)$ classical post-processing.

The streamlined algorithm is to initialize the state of two $GF(2^n)$ -valued registers to $|0\rangle|M_\phi\bar{1}\rangle$ and perform the following operations:

1. Apply a Hadamard transform to each qubit of each register.
2. Query the black-box.

3. Apply a Hadamard transform to each qubit of each register.
4. Measure the first register, yielding an n -bit string z .
5. Classically, compute $(M_\phi^{-1})^\top \bar{z}$.

It is straightforward to show that the result will be r .

4 Extension to matrices over a finite field

Having defined the Fourier transform over finite fields in general, it seems a natural step to ask if it can be usefully defined for more general algebraic constructs. In this section we consider matrices over finite fields.

Before we proceed, let us introduce some notation. Given m^2 quantum registers $|x_{ij}\rangle$ ($1 \leq i, j \leq m$), let

$$\bigotimes_{i=1}^m \bigotimes_{j=1}^m |x_{ij}\rangle = |x_{11}\rangle |x_{12}\rangle \cdots |x_{1m}\rangle |x_{21}\rangle \cdots |x_{mm}\rangle.$$

We identify this state with an $m \times m$ matrix $X = (x_{ij})$ in the natural way.

Definition 4.1 *Let $F_{q,\phi}$ be the quantum Fourier transform over $GF(q)$ relative to ϕ . Then the quantum Fourier transform over $GF(q)^{m \times m}$ is defined by the following mapping for each matrix $A = (x_{ij}) \in GF(q)^{m \times m}$:*

$$F_{q,m,\phi} : |A\rangle \mapsto \bigotimes_{i=1}^m \bigotimes_{j=1}^m F_{q,\phi} |x_{ji}\rangle.$$

That is, the quantum Fourier transform of A relative to ϕ is performed by applying the Fourier transform relative to ϕ independently on all the entries of A , and then transposing the resulting matrix.

Now let us denote the left controlled-add gate with parameter R as

$$C_{R*} : |X\rangle|Y\rangle \mapsto |X\rangle|Y + RX\rangle$$

and the right controlled-add gate with parameter R as

$$C_{*R} : |X\rangle|Y\rangle \mapsto |X\rangle|Y + XR\rangle$$

The target/control inversion property of Figure 3 for the particular ring $GF(q)^{m \times m}$ may be described as follows:

$$\begin{aligned} (F_{q,m,\phi}^\dagger \otimes F_{q,m,\phi}) C_{R*} (F_{q,m,\phi} \otimes F_{q,m,\phi}^\dagger) |X\rangle|Y\rangle &= |X + YR\rangle|Y\rangle \\ (F_{q,m,\phi}^\dagger \otimes F_{q,m,\phi}) C_{*R} (F_{q,m,\phi} \otimes F_{q,m,\phi}^\dagger) |X\rangle|Y\rangle &= |X + RY\rangle|Y\rangle. \end{aligned}$$

It is reasonably straightforward to verify that this property holds. We omit the details from this abstract.

Finally, we mention how this result may be applied to a hidden linear structure problem similar to the one described in Section 3. In this case, we have a black box of either of the following forms:

$$|X\rangle|Y\rangle \mapsto |X\rangle|\pi(Y + RX)\rangle = (I_m \otimes P_\pi)C_{R*}|X\rangle|Y\rangle$$

or

$$|X\rangle|Y\rangle \mapsto |X\rangle|\pi(Y + XR)\rangle = (I_m \otimes P_\pi)C_{*R}|X\rangle|Y\rangle$$

Now, analogously to the hidden linear structure problem considered in Section 3, the parameter R can be found by setting X to the additive identity ($0_{m \times m}$), and Y to the multiplicative identity ($I_{m \times m}$).

5 Conclusion

A comparison with other known quantum vs. classical query separations in the bounded-error model is given in Table 1.

References	number of bits	quantum upper bound	classical lower bound
Bernstein & Vazirani [2]	$n+1$	$O(1)$	$\Omega(n)$
Bernstein & Vazirani [2]	$\Theta(n)$	$O(n \log n)$	$n^{\Omega(\log n)}$
Simon [12]	$2n$	$O(n)$	$\Omega(2^{n/2})$
Grover [8]	$n+1$	$O(2^{n/2})$	$\Omega(2^n)$
Shor [11] / Cleve [4]	$3n$	$O(1)$	$\Omega(2^{n/3}/\sqrt{n})$
Present result	$2n$	1	$\Omega(2^{n/2})$

Table 1: Comparison of quantum vs. classical separations for query problems.

Although not indicated in Table 1, the auxiliary operations required for an exact computation in the present result are very simple (as explained in Section 3): $O(n)$ Hadamard gates plus $O(n^2)$ classical post-processing after measuring the state.

Acknowledgments

R.C. gratefully acknowledges the University of California at Berkeley and the California Institute of Technology where some of the writing and revisions to this paper occurred.

References

- [1] A. Barenco, C. H. Bennett, R. Cleve, D. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. Smolin, and H. Weinfurter. Elementary gates for quantum computation. *Physical Review A*, 52:3457–3467, 1995.
- [2] E. Bernstein and U. Vazirani. Quantum complexity theory. *SIAM Journal on Computing*, 26(5):1411–1473, 1997.

- [3] D. Boneh and R. Lipton. Quantum cryptanalysis of hidden linear functions. In *Advances in Cryptology – Crypto’95*, volume 963 of *Lecture Notes in Computer Science*, pages 242–437. Springer-Verlag, 1995.
- [4] R. Cleve. The query complexity of order-finding. In *Proceedings of the Fifteenth IEEE Conference on Computational Complexity*, pages 54–59, 2000.
- [5] H. Cohen. *A Course in Computational Algebraic Number Theory*. Springer-Verlag, 1993.
- [6] W. van Dam and S. Hallgren. Efficient quantum algorithms for shifted quadratic character equations. To appear on the quant-ph archive, November, 2000.
- [7] J. von zur Gathen and J. Gerhard. *Modern Computer Algebra*. Cambridge University Press, 1999.
- [8] L. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, pages 212–219, 1996.
- [9] L. Hales and S. Hallgren. An improved quantum Fourier transform algorithm and applications. In *Proceedings of the Thirty-Third Annual ACM Symposium on Theory of Computing*, 2000. to appear.
- [10] R. Lidl and H. Niederreiter. *Introduction to Finite Fields and Their Applications*. Cambridge University Press, revised edition, 1994.
- [11] P. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5):1484–1509, 1997.
- [12] D. Simon. On the power of quantum computation. In *Proceedings of the 35th Annual Symposium on Foundations of Computer Science*, pages 116–123, 1994.