

Quantum Fourier Transform Revisited

Daan Camps^{1,*}, Roel Van Beeumen¹, Chao Yang¹

¹Computational Research Division, Lawrence Berkeley National Laboratory, CA, United States

Abstract

The fast Fourier transform (FFT) is one of the most successful numerical algorithms of the 20th century and has found numerous applications in many branches of computational science and engineering. The FFT algorithm can be derived from a particular matrix decomposition of the discrete Fourier transform (DFT) matrix. In this paper, we show that the quantum Fourier transform (QFT) can be derived by further decomposing the diagonal factors of the FFT matrix decomposition into products of matrices with Kronecker product structure. We analyze the implication of this Kronecker product structure on the discrete Fourier transform of rank-1 tensors on a classical computer. We also explain why such a structure can take advantage of an important quantum computer feature that enables the QFT algorithm to attain an exponential speedup on a quantum computer over the FFT algorithm on a classical computer. Further, the connection between the matrix decomposition of the DFT matrix and a quantum circuit is made. We also discuss a natural extension of a radix-2 QFT decomposition to a radix- d QFT decomposition. No prior knowledge of quantum computing is required to understand what is presented in this paper. Yet, we believe this paper may help readers to gain some rudimentary understanding of the nature of quantum computing from a matrix computation point of view.

1 Introduction

The fast Fourier transform (FFT) [3] is a widely celebrated algorithmic innovation of the 20th century [19]. The algorithm allows us to perform a discrete Fourier transform (DFT) of a vector of size N in $\mathcal{O}(N \log N)$ operations. This is much more efficient than the $\mathcal{O}(N^2)$ operations required in a brute force calculation in which the $N \times N$ matrix representation of the DFT is directly multiplied with the vector to be transformed. This efficiency improvement is quite remarkable and has enabled a wide range of applications such as signal processing and spectral methods for solving partial differential equations. However, the computational complexity of the transform can be further reduced to $\mathcal{O}((\log N)^2)$ by using the quantum Fourier transform (QFT) on a quantum computer. This is well known in the quantum computing community and is often touted as an example of the type of exponential speedup a quantum computer can achieve. So what makes the QFT so efficient?

In this paper, we will explain what the QFT exactly does, and why it can achieve an additional exponential factor of speedup. Unlike the derivation provided in other references [4, 14], our explanation of the QFT requires little knowledge of quantum computing. We use the language of matrix and tensor decompositions to describe and compare the operations performed in the QFT and FFT. Just like the FFT, the QFT relies on a special matrix factorization of the DFT matrix to attain its efficiency. This factorization produces a product of $\mathcal{O}((\log N)^2)$ simpler unitary matrices that can in turn be written as the sum of 2 Kronecker products of a 2×2 matrix with 2×2 identity matrices. This derivation of the QFT factorization only requires some knowledge of the elementary properties of Kronecker products.

We point out that a key distinction between a classical computer and a quantum computer makes the QFT decomposition well suited for a quantum computer. This is also what makes the QFT much more efficient on a quantum computer. On the other hand, we will also mention the limitation of QFT in terms of its application in classical algorithms such as fast convolution [14, 12].

This paper is mainly pedagogical in nature. Apart from giving a matrix and tensor decomposition oriented derivation of the QFT algorithm, presenting alternative quantum circuits for the QFT, and extending the QFT algorithm for qubit based quantum computers to algorithms for qudit based quantum computers, it does not contain new results. However, our presentation should make it easier for researchers in traditional numerical computation to understand the features and limitations of quantum algorithms and the potential of quantum computing.

*Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94720. Email: DCamps@lbl.gov.

The paper is structured as follows. We first review the definition of tensor and Kronecker products and its properties in section 2. We then examine a radix-2 FFT algorithm in section 3, which is applied to vectors of dimension $N = 2^n$, and describe the matrix factorization associated with the FFT algorithm. Similar approaches are used in [15, 18, 8, 20]. In section 4, we show how the matrix decomposition produced by the FFT algorithm can be modified to produce a matrix decomposition for the QFT algorithm. In section 5, we explain why QFT can achieve $\mathcal{O}((\log N)^2)$ complexity on a quantum computer and show how the QFT can be more conveniently described by a quantum circuit. In section 6, we generalize the radix-2 QFT algorithm to a radix- d QFT algorithm and present the quantum circuit for implementing such a QFT on a d -level quantum computer. The QFT on d -level quantum computers was also studied in [2, 13, 6].

Throughout the paper, we denote vectors by lowercase Roman characters and matrices by capital Roman characters, e.g., v and M . Matrices and vectors of exponential dimension 2^n , and in general d^n , are denoted by, e.g., \mathbf{M}_n and \mathbf{v}_n , respectively. The value of a (classical) bit is denoted as j and is either 0 or 1 in the radix-2 case, and $\{0, 1, \dots, d-1\}$ in the radix- d case.

2 Tensor and Kronecker products

In this section we briefly review the Kronecker product and its properties. We start by defining the Kronecker product of two matrices.

Definition 2.1 (Kronecker product). *The Kronecker product of the matrices $A \in \mathbb{C}^{n \times m}$ and $B \in \mathbb{C}^{p \times q}$ is defined as the following $np \times mq$ block matrix*

$$A \otimes B := \begin{bmatrix} a_{11}B & a_{12}B & \cdots & a_{1,m}B \\ a_{21}B & a_{22}B & \cdots & a_{2,m}B \\ \vdots & \vdots & \ddots & \vdots \\ a_{n,1}B & a_{n,2}B & \cdots & a_{n,m}B \end{bmatrix},$$

where a_{ij} is the (i, j) th element of A .

The Kronecker product is a special case of the tensor product, hence, it is bilinear

$$\begin{aligned} (\gamma A) \otimes B &= A \otimes (\gamma B) = \gamma(A \otimes B), \\ A \otimes (B + C) &= A \otimes B + A \otimes C, \\ (B + C) \otimes A &= B \otimes A + C \otimes A. \end{aligned} \tag{2.1}$$

Another important property is the *mixed-product* identity

$$(A \otimes B)(C \otimes D) = (AC) \otimes (BD) \tag{2.2}$$

for multiplying Kronecker products of matrices of compatible dimension. The *direct sum*, not to be confused by the Kronecker sum, is defined as follows.

Definition 2.2 (Direct sum). *The direct sum of the square matrices $A \in \mathbb{C}^{n \times n}$ and $B \in \mathbb{C}^{m \times m}$ is defined as the following $(n+m) \times (n+m)$ diagonal block matrix*

$$A \oplus B := \begin{bmatrix} A & \\ & B \end{bmatrix}.$$

Rank-1 tensors (of vectors) and tensor rank decompositions will play a key role in the remainder of the paper. We therefore formally define them and list some important properties.

Definition 2.3 (Tensor rank decomposition). *Let $v_1, v_2, \dots, v_n \in \mathbb{C}^2$. Then the 2^n dimensional vector*

$$\mathbf{v}_n = v_1 \otimes v_2 \otimes \cdots \otimes v_n,$$

is a rank-1 n -way tensor of size 2 in every mode. In general, \mathbf{v}_n is a vector with tensor rank- r if its minimal representation as a sum of vectors of tensor rank-1 requires r terms,

$$\mathbf{v}_n = \sum_{i=1}^r v_1^{(i)} \otimes v_2^{(i)} \otimes \cdots \otimes v_n^{(i)}. \tag{2.3}$$

The decomposition (2.3) is known as a tensor rank decomposition or canonical polyadic decomposition [10].

In this paper, we will use the *big-endian* binary convention, formalized by the following definition. We also state two lemmas related to binary representations of integers and tensors. These results will be used in some proofs later on.

Definition 2.4 (Binary representation). *We define the binary representation of $j \in \mathbb{N} : 0 \leq j \leq 2^n - 1$ as follows*

$$j = [j_1 j_2 \cdots j_{n-1} j_n] = j_1 \cdot 2^{n-1} + j_2 \cdot 2^{n-2} + \cdots + j_{n-1} \cdot 2^1 + j_n \cdot 2^0,$$

where $j_i \in \{0, 1\}$ for $i = 1, \dots, n$.

Lemma 2.5. *For every scalar $\alpha \in \mathbb{C}$ and $j \in \mathbb{N} : 0 \leq j \leq 2^n - 1$, we have*

$$\alpha^j = \alpha^{j_1 2^{n-1}} \alpha^{j_2 2^{n-2}} \cdots \alpha^{j_{n-1} 2^1} \alpha^{j_n 2^0}.$$

Proof. The proof directly follows from Definition 2.4. □

Lemma 2.6. *Let $v_1, v_2, \dots, v_n \in \mathbb{C}^2$. Then*

$$\mathbf{v}_n := v_1 \otimes v_2 \otimes \cdots \otimes v_n \quad \iff \quad \mathbf{v}_n(j) = v_1(j_1) v_2(j_2) \cdots v_n(j_n),$$

where $\mathbf{v}_n(j)$ denotes the j th element of the vector $\mathbf{v}_n \in \mathbb{C}^{2^n}$.

Proof. The proof directly follows from Definitions 2.1 and 2.4. □

We define two elementary matrices that are useful for our analysis.

Definition 2.7. *Let e_1 and e_2 be the first and second column of I_2 . Then we define the matrices:*

$$E_1 := e_1 e_1^\top = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, \quad E_2 := e_2 e_2^\top = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}.$$

We clearly have that $E_1 + E_2 = I_2$. A particular useful application of Definition 2.7 is the decomposition of the direct sum of two matrices as a sum of two Kronecker product terms:

$$A \oplus B = E_1 \otimes A + E_2 \otimes B. \tag{2.4}$$

3 Matrix Decomposition for Fast Fourier Transform

The discrete Fourier transform maps a series of N complex numbers in another series of N complex numbers as defined below.

Definition 3.1 (DFT [20]). *The discrete Fourier transform of a vector $x = [x_0, \dots, x_{N-1}]^T \in \mathbb{C}^N$ is defined as the vector $y = [y_0, \dots, y_{N-1}]^T \in \mathbb{C}^N$ with*

$$y_k = \frac{1}{\sqrt{N}} \sum_{j=0}^{N-1} \omega_N^{kj} x_j, \quad k = 0, 1, \dots, N-1, \tag{3.1}$$

where $\omega_N := e^{-\frac{2\pi i}{N}}$ is an N th root of unity.

Similarly, the inverse discrete Fourier transform (IDFT) is given by (3.1) where $\omega_N := e^{\frac{2\pi i}{N}}$ is now the principal N th root of unity. The expression (3.1) can be written as the multiplication of a DFT matrix F_N to be defined below with the vector x .

3.1 The DFT matrix

Definition 3.1, or variants thereof with a scalar factor different from $\frac{1}{\sqrt{N}}$ ¹, are most commonly used throughout the literature. As we focus on the matrix representation and decompositions of that matrix representation, the following equivalent characterization of the DFT is more useful. A similar approach is followed in [15, 18, 8, 20]. Van Loan [20] even states: “I am convinced that life as we know it would be considerably different if, from the 1965 Cooley–Tukey paper onwards, the FFT community had made systematic and heavy use of matrix-vector notation!”.

¹ The scalar factors of the DFT and the inverse DFT need to multiply to $\frac{1}{N}$. The choice $\frac{1}{\sqrt{N}}$ leads to unitary transformations.

Definition 3.2 (DFT matrix). *The DFT matrix is defined as the unitary matrix*

$$F_N := \frac{1}{\sqrt{N}} \begin{bmatrix} \omega_N^0 & \omega_N^0 & \omega_N^0 & \cdots & \omega_N^0 \\ \omega_N^0 & \omega_N^1 & \omega_N^2 & \cdots & \omega_N^{N-1} \\ \omega_N^0 & \omega_N^2 & \omega_N^4 & \cdots & \omega_N^{2(N-1)} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \omega_N^0 & \omega_N^{N-1} & \omega_N^{2(N-1)} & \cdots & \omega_N^{(N-1)(N-1)} \end{bmatrix} \in \mathbb{C}^{N \times N}, \quad (3.2)$$

where $\omega_N := e^{-\frac{2\pi i}{N}}$.

Using Definition 3.2, the DFT and IDFT become matrix-vector multiplications $y = F_N x$ and $y = F_N^* x$, respectively. As an example, the first four DFT matrices are given by:

$$F_1 = [1], \quad F_2 = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}, \quad F_3 = \frac{1}{\sqrt{3}} \begin{bmatrix} 1 & 1 & 1 \\ 1 & \omega_3 & \omega_3^2 \\ 1 & \omega_3^2 & \omega_3 \end{bmatrix}, \quad F_4 = \frac{1}{2} \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & -i & -1 & i \\ 1 & -1 & 1 & -1 \\ 1 & i & -1 & -i \end{bmatrix},$$

where we used that ω_N simplifies for $N = 1, 2$, and 4, and the equality $\omega_N^m = \omega_N^{(m \bmod N)}$. The 2×2 DFT matrix F_2 is also called the *Hadamard* matrix and denoted as H .

In order to simplify notation, we will make use of the exponent notation for the DFT matrix as introduced by Pease [15]:

$$\tilde{F}_N = \begin{bmatrix} 0 & 0 & 0 & \cdots & 0 \\ 0 & 1 & 2 & \cdots & N-1 \\ 0 & 2 & 4 & \cdots & 2(N-1) \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & N-1 & 2(N-1) & \cdots & (N-1)(N-1) \end{bmatrix},$$

where only the exponents of ω_N are listed in the matrix \tilde{F}_N . For the remainder of the paper, all exponent matrices will be indicated by a tilde. Note that by using exponent notation, an multiplication of ω_N factors becomes an addition of its exponents.

3.2 Radix-2 decomposition of the discrete Fourier transform matrix

A straightforward implementation of the (inverse) DFT as a matrix-vector product requires $\mathcal{O}(N^2)$ operations. The radix-2 FFT algorithm, popularized by Cooley and Tukey [3], computes the DFT for vectors of length $N = 2^n$ in only $\mathcal{O}(N \log_2 N) = \mathcal{O}(2^n n)$ operations. This speedup is achieved by a divide-and-conquer approach which relates a permuted DFT matrix of dimension N to a permuted DFT matrix of size $N/2$. Let us illustrate this matrix decomposition for the DFT matrix of dimension 8.

Example 1. *The exponent DFT matrix of dimension 8 is given by:*

$$\tilde{F}_8 = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 2 & 3 & 4 & 5 & 6 & 7 \\ 0 & 2 & 4 & 6 & 8 & 10 & 12 & 14 \\ 0 & 3 & 6 & 9 & 12 & 15 & 18 & 21 \\ 0 & 4 & 8 & 12 & 16 & 20 & 24 & 28 \\ 0 & 5 & 10 & 15 & 20 & 25 & 30 & 35 \\ 0 & 6 & 12 & 18 & 24 & 30 & 36 & 42 \\ 0 & 7 & 14 & 21 & 28 & 35 & 42 & 49 \end{bmatrix} \stackrel{\text{mod } 8}{\equiv} \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 2 & 3 & 4 & 5 & 6 & 7 \\ 0 & 2 & 4 & 6 & 0 & 2 & 4 & 6 \\ 0 & 3 & 6 & 1 & 4 & 7 & 2 & 5 \\ 0 & 4 & 0 & 4 & 0 & 4 & 0 & 4 \\ 0 & 5 & 2 & 7 & 4 & 1 & 6 & 3 \\ 0 & 6 & 4 & 2 & 0 & 6 & 4 & 2 \\ 0 & 7 & 6 & 5 & 4 & 3 & 2 & 1 \end{bmatrix} \begin{matrix} 000 \\ 001 \\ 010 \\ 011 \\ 100 \\ 101 \\ 110 \\ 111 \end{matrix},$$

where in the matrix on the right the exponents are given modulo 8. The binary representation of the row indices are shown on the right side of the matrix. Permuting the rows such that the binary representation of the row

indices reverse in order results in the matrix \tilde{F}'_8 :

$$\tilde{F}'_8 = \left[\begin{array}{cccc|cccc} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 4 & 0 & 4 & 0 & 4 & 0 & 4 \\ 0 & 2 & 4 & 6 & 0 & 2 & 4 & 6 \\ 0 & 6 & 4 & 2 & 0 & 6 & 4 & 2 \\ \hline 0 & 1 & 2 & 3 & 4 & 5 & 6 & 7 \\ 0 & 5 & 2 & 7 & 4 & 1 & 6 & 3 \\ 0 & 3 & 6 & 1 & 4 & 7 & 2 & 5 \\ 0 & 7 & 6 & 5 & 4 & 3 & 2 & 1 \end{array} \right] \begin{array}{l} 000 \\ 100 \\ 010 \\ 110 \\ 001 \\ 101 \\ 011 \\ 111 \end{array} = \left[\begin{array}{c|c} \tilde{F}'_4 & \tilde{F}'_4 \\ \hline \tilde{F}'_4 + \tilde{\Omega}_4 & \tilde{F}'_4 + \tilde{\Omega}_4 + 4 \end{array} \right], \text{ with } \tilde{\Omega}_4 = \begin{bmatrix} 0 & 1 & 2 & 3 \\ 0 & 1 & 2 & 3 \\ 0 & 1 & 2 & 3 \\ 0 & 1 & 2 & 3 \end{bmatrix}.$$

The left side of the above expression shows a block partitioning of \tilde{F}'_8 in terms of the bit-reversed exponent DFT matrix of half the dimension, \tilde{F}'_4 . Here we again applied the modulo 8 equivalence. This partitioning can be written in terms of the permuted DFT matrices as:

$$F'_8 = \frac{1}{\sqrt{2}} \left[\begin{array}{c|c} F'_4 & F'_4 \\ \hline F'_4 \Omega_4 & -F'_4 \Omega_4 \end{array} \right], \quad \text{with } \Omega_4 = \text{diag}(0, 1, 2, 3),$$

where the minus sign in the (2,2)-block comes from $\omega_8^4 = -1$.

The generalization of the permutation and partition to a DFT matrix of dimension 2^n for $n > 2$ requires the definition of Ω_{2^n} , which we give below.

Definition 3.3. Define Ω_n as the following $2^n \times 2^n$ diagonal matrix:

$$\Omega_n = \Omega_{2^n} := \begin{bmatrix} \omega_{2^{n+1}}^0 & & & \\ & \omega_{2^{n+1}}^1 & & \\ & & \ddots & \\ & & & \omega_{2^{n+1}}^{2^n-1} \end{bmatrix},$$

where $\omega_{2^{n+1}} := e^{\frac{-2\pi i}{2^{n+1}}}$.

The block partitioning of the 4×4 DFT matrix that we demonstrated in Example 1 can be generalized for any $2^n \times 2^n$ DFT matrix as stated in the following theorem.

Theorem 3.4 (See [18, 15]). Let $\mathbf{F}'_n = \mathbf{P}_n \mathbf{F}_n$ be the $2^n \times 2^n$ bit-reversed DFT matrix, where \mathbf{P}_n is the bit reversal permutation matrix acting on n bits. Then \mathbf{F}'_n admits the following factorization:

$$\begin{aligned} \mathbf{F}'_n &= \frac{1}{\sqrt{2}} \left[\begin{array}{cc} \mathbf{F}'_{n-1} & \mathbf{F}'_{n-1} \\ \mathbf{F}'_{n-1} \Omega_{n-1} & -\mathbf{F}'_{n-1} \Omega_{n-1} \end{array} \right], \\ &= \left[\begin{array}{c} \mathbf{F}'_{n-1} \\ \mathbf{F}'_{n-1} \end{array} \right] \left[\begin{array}{cc} \mathbf{I}_{n-1} & \\ & \Omega_{n-1} \end{array} \right] \left(\frac{1}{\sqrt{2}} \left[\begin{array}{cc} \mathbf{I}_{n-1} & \mathbf{I}_{n-1} \\ \mathbf{I}_{n-1} & -\mathbf{I}_{n-1} \end{array} \right] \right), \\ &= (\mathbf{I}_2 \otimes \mathbf{F}'_{n-1})(\mathbf{I}_{n-1} \oplus \Omega_{n-1})(H \otimes \mathbf{I}_{n-1}), \end{aligned} \quad (3.3)$$

where Ω_{n-1} is given by Definition 3.3.

The subdivision of the permuted DFT matrix in a 2×2 block matrix is also called a *radix-2* splitting in the literature. The bit reversal permutation \mathbf{P}_n satisfies the following identity

$$\mathbf{P}_n(v_1 \otimes \cdots \otimes v_n) = v_n \otimes \cdots \otimes v_1. \quad (3.4)$$

where $v_k \in \mathbb{C}^2$ for $k = 1, \dots, n$. The matrix \mathbf{P}_n is involutory, because $\mathbf{P}_n^2 = \mathbf{I}_n$, i.e., reversing the bits twice gives the original ordering. It is also unitary because it is a permutation matrix.

Theorem 3.4 can be applied repeatedly to obtain the complete radix-2 factorization of the DFT matrix. This is formalized in the following theorem.

Theorem 3.5. The $2^n \times 2^n$ DFT matrix \mathbf{F}_n can be factored as:

$$\mathbf{F}_n = \mathbf{P}_n \mathbf{F}'_n = \mathbf{P}_n \mathbf{A}_n^{(0)} \mathbf{A}_n^{(1)} \cdots \mathbf{A}_n^{(n-1)}, \quad (3.5)$$

where, for $k = 0, 1, \dots, n-1$,

$$\mathbf{A}_n^{(k)} = \mathbf{I}_{n-k-1} \otimes \mathbf{B}_{k+1}, \quad \text{and} \quad \mathbf{B}_{k+1} = \frac{1}{\sqrt{2}} \left[\begin{array}{cc} \mathbf{I}_k & \mathbf{I}_k \\ \Omega_k & -\Omega_k \end{array} \right] = (\mathbf{I}_k \oplus \Omega_k)(H \otimes \mathbf{I}_k), \quad (3.6)$$

with Ω_k given by Definition 3.3.

Proof. We give the proof for \mathbf{F}'_n by induction based on Theorem 3.4. The case $n = 1$ can be directly verified:

$$\mathbf{F}'_1 = \mathbf{A}'_1 = \mathbf{1} \otimes \mathbf{B}_1 = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}.$$

Assume $\mathbf{F}'_n = \mathbf{A}'_n \mathbf{A}'_{n-1} \cdots \mathbf{A}'_1$ as prescribed by Theorem 3.5. For the bit-reversed DFT matrix of size 2^{n+1} we have by Theorem 3.4,

$$\begin{aligned} \mathbf{F}'_{n+1} &= (I_2 \otimes \mathbf{F}'_n)(\mathbf{I}_n \oplus \mathbf{\Omega}_n)(H \otimes \mathbf{I}_n), \\ &= \left(I_2 \otimes \left[\mathbf{A}'_n \mathbf{A}'_{n-1} \cdots \mathbf{A}'_1 \right] \right) (\mathbf{I}_n \oplus \mathbf{\Omega}_n)(H \otimes \mathbf{I}_n), \\ &= \left(I_2 \otimes \mathbf{A}'_n \right) \left(I_2 \otimes \mathbf{A}'_{n-1} \right) \cdots \left(I_2 \otimes \mathbf{A}'_1 \right) \mathbf{A}'_{n+1}, \\ &= \mathbf{A}'_{n+1} \mathbf{A}'_{n+1} \cdots \mathbf{A}'_{n+1} \mathbf{A}'_{n+1}, \end{aligned}$$

where we used $I_2 \otimes \mathbf{A}'_k = \mathbf{A}'_{k+1}$ for $k = 0, 1, \dots, n-1$, and $\mathbf{A}'_{n+1} = \mathbf{B}_{n+1}$. \square

The key observation here is that each matrix \mathbf{A}'_k in (3.6) only has two nonzero elements on every row. Consequently, the matrix-vector product $\mathbf{A}'_k v$ can be computed in only $\mathcal{O}(2^k)$ operations, resulting in an overall $\mathcal{O}(2^{2n})$ or $\mathcal{O}(N \log N)$ computational complexity for the Fourier transform of a vector of size $N = 2^n$. This is why the FFT is called *fast*.

4 Matrix Decomposition for Quantum Fourier Transform

As we can see from the previous section, the reduction in complexity from $\mathcal{O}(N^2)$ to $\mathcal{O}(N \log N)$ in the FFT algorithm essentially results from the Kronecker product structure that appears in the matrix factors of \mathbf{F}'_n in (3.3). It is the $\mathbf{I}_k \oplus \mathbf{\Omega}_k$ factor that retains the N factor in the complexity of the FFT algorithm because the multiplication of that diagonal matrix with a vector has to be performed in $\mathcal{O}(N)$ operations.

It is conceivable that the complexity of the computation can potentially be reduced further, at least in some special cases, if we can somehow rewrite this diagonal matrix as a Kronecker product of 2×2 matrices. It turns out that we can almost do that. In this section, we show that the diagonal matrix $\mathbf{I}_k \oplus \mathbf{\Omega}_k$ can be written as a product of k simpler matrices, each of which is the sum of two Kronecker products of 2×2 matrices. This factorization yields a decomposition of the DFT matrix that enables the DFT to be performed efficiently on a quantum computer. Exactly how that is achieved will be discussed in the next section. Here we will simply refer to this decomposition as the decomposition used by the quantum Fourier transform (QFT).

We begin by showing that the diagonal matrix $\mathbf{\Omega}_n$ can be written as a Kronecker product of 2×2 unitary matrices of the following form.

Definition 4.1. *Define*

$$R_n := \begin{bmatrix} \omega_{2^n}^0 & \\ & \omega_{2^n}^1 \end{bmatrix} = \begin{bmatrix} 1 & \\ & \omega_{2^n} \end{bmatrix}, \quad (4.1)$$

where $\omega_{2^n} := e^{\frac{-2\pi i}{2^n}}$.

The matrix R_n satisfies the identity stated in the following lemma.

Lemma 4.2. *Let R_n be defined by Definition 4.1. Then*

$$R_n^{2^j} = R_{n-j},$$

for $j \in \mathbb{N}$.

Proof. We start from the following identity

$$\omega_{2^n}^{2^j} = \left(e^{\frac{-2\pi i}{2^n}} \right)^{2^j} = e^{\frac{-2\pi i}{2^{n-j}}} = \omega_{2^{n-j}}.$$

Hence,

$$(R_n)^{2^j} = \begin{bmatrix} 1 & \\ & \omega_{2^n}^{2^j} \end{bmatrix} = \begin{bmatrix} 1 & \\ & \omega_{2^{n-j}} \end{bmatrix} = R_{n-j}. \quad \square$$

We can use this result to decompose $\mathbf{\Omega}_n$ as a Kronecker product of n R_i matrices.

Lemma 4.3. *Let Ω_n be defined by Definition 3.3. Then*

$$\Omega_n = R_2 \otimes R_3 \otimes \cdots \otimes R_n \otimes R_{n+1}, \quad (4.2)$$

where R_n is defined by Definition 4.1.

Proof. From Lemma 2.5, we can use the binary representation of j to rewrite the j th diagonal element of Ω_n as

$$\Omega_n(j, j) = \omega_{2^{n+1}}^j = \omega_{2^{n+1}}^{j_1 2^{n-1}} \omega_{2^{n+1}}^{j_2 2^{n-2}} \cdots \omega_{2^{n+1}}^{j_{n-1} 2^1} \omega_{2^{n+1}}^{j_n 2^0}, \quad (4.3)$$

for $j \in \mathbb{N} : 0 \leq j \leq 2^n - 1$. We can rewrite (4.3) as

$$\Omega_n(j, j) = R_{n+1}^{2^{n-1}}(j_1, j_1) R_{n+1}^{2^{n-2}}(j_2, j_2) \cdots R_{n+1}^{2^1}(j_{n-1}, j_{n-1}) R_{n+1}^{2^0}(j_n, j_n). \quad (4.4)$$

It follow from Lemma 4.2 that (4.4) simplifies to

$$\Omega_n(j, j) = R_2(j_1, j_1) R_3(j_2, j_2) \cdots R_n(j_{n-1}, j_{n-1}) R_{n+1}(j_n, j_n), \quad (4.5)$$

such that (4.2) follows from Lemma 2.6, which extends trivially to diagonal matrices. \square

Our objective is to decompose $\mathbf{I}_k \oplus \Omega_k$ into a product of k matrices, each of which can be written as the sum of Kronecker products of 2×2 matrices. The follow theorem and its proof shows how this can be done.

Theorem 4.4. *The diagonal operators $\mathbf{I}_k \oplus \Omega_k \in \mathbb{C}^{2^{k+1} \times 2^{k+1}}$ from Theorem 3.5 admit the decomposition*

$$\mathbf{I}_k \oplus \Omega_k = \prod_{i=1}^k \left[E_1 \otimes \mathbf{I}_{i-1} \otimes I_2 \otimes \mathbf{I}_{k-i} + E_2 \otimes \mathbf{I}_{i-1} \otimes R_{i+1} \otimes \mathbf{I}_{k-i} \right], \quad (4.6)$$

where R_i is defined in Definition 4.1.

Proof. We start by splitting the diagonal matrix $\mathbf{I}_k \oplus \Omega_k$ into 2 terms using (2.4), which yields

$$\mathbf{I}_k \oplus \Omega_k = E_1 \otimes \mathbf{I}_k + E_2 \otimes \Omega_k. \quad (4.7)$$

Next, we rewrite these terms in a redundant form based on the mixed product identity of Kronecker products (2.2). Using the property $E_1 = E_1^2$, we get for the first term

$$\begin{aligned} E_1 \otimes \mathbf{I}_k &= E_1 \otimes I_2 \otimes \cdots \otimes I_2, \\ &= (E_1 \otimes I_2 \otimes I_2 \otimes \cdots \otimes I_2)(E_1 \otimes I_2 \otimes I_2 \otimes \cdots \otimes I_2) \cdots (E_1 \otimes I_2 \otimes I_2 \otimes \cdots \otimes I_2), \\ &= \prod_{i=1}^k E_1 \otimes \mathbf{I}_{i-1} \otimes I_2 \otimes \mathbf{I}_{k-i}, \end{aligned}$$

and using Lemma 4.3 and the property $E_2 = E_2^2$, the second term results in

$$\begin{aligned} E_2 \otimes \Omega_k &= E_2 \otimes R_2 \otimes \cdots \otimes R_{k+1}, \\ &= (E_2 \otimes R_2 \otimes I_2 \otimes \cdots \otimes I_2)(E_2 \otimes I_2 \otimes R_3 \otimes \cdots \otimes I_2) \cdots (E_2 \otimes I_2 \otimes I_2 \otimes \cdots \otimes R_{k+1}), \\ &= \prod_{i=1}^k E_2 \otimes \mathbf{I}_{i-1} \otimes R_{i+1} \otimes \mathbf{I}_{k-i}. \end{aligned}$$

Hence, the sum of Kronecker products (4.7) is equal to the product of sums, i.e.,

$$\begin{aligned} \mathbf{I}_k \oplus \Omega_k &= \prod_{i=1}^k E_1 \otimes \mathbf{I}_{i-1} \otimes I_2 \otimes \mathbf{I}_{k-i} + \prod_{i=1}^k E_2 \otimes \mathbf{I}_{i-1} \otimes R_{i+1} \otimes \mathbf{I}_{k-i}, \\ &= \prod_{i=1}^k \left[E_1 \otimes \mathbf{I}_{i-1} \otimes I_2 \otimes \mathbf{I}_{k-i} + E_2 \otimes \mathbf{I}_{i-1} \otimes R_{i+1} \otimes \mathbf{I}_{k-i} \right], \end{aligned}$$

because the mixed terms in the latter expression cancel out due to the property $E_1 E_2 = 0$. \square

The order of the terms in the Kronecker product expression (4.6) can be changed because of the specific form of the R_{i+1} matrices.

Lemma 4.5. *The Kronecker product matrices in (4.6) satisfy*

$$E_1 \otimes \mathbf{I}_{i-1} \otimes I_2 \otimes \mathbf{I}_{k-i} + E_2 \otimes \mathbf{I}_{i-1} \otimes R_{i+1} \otimes \mathbf{I}_{k-i} \quad (4.8)$$

$$= I_2 \otimes \mathbf{I}_{i-1} \otimes E_1 \otimes \mathbf{I}_{k-i} + R_{i+1} \otimes \mathbf{I}_{i-1} \otimes E_2 \otimes \mathbf{I}_{k-i}, \quad (4.9)$$

for $i = 1, \dots, k$.

Proof. Splitting I_2 and R_{i+1} in their E_1 and E_2 components, i.e.,

$$\begin{aligned} I_2 &= E_1 + E_2, \\ R_{i+1} &= E_1 + \omega_{2^{i+1}} E_2, \end{aligned}$$

yields

$$\begin{aligned} E_1 \otimes \mathbf{I}_{i-1} \otimes I_2 \otimes \mathbf{I}_{k-i} + E_2 \otimes \mathbf{I}_{i-1} \otimes R_{i+1} \otimes \mathbf{I}_{k-i} &= E_1 \otimes \mathbf{I}_{i-1} \otimes E_1 \otimes \mathbf{I}_{k-i} + E_1 \otimes \mathbf{I}_{i-1} \otimes E_2 \otimes \mathbf{I}_{k-i} + \\ &\quad E_2 \otimes \mathbf{I}_{i-1} \otimes E_1 \otimes \mathbf{I}_{k-i} + E_2 \otimes \mathbf{I}_{i-1} \otimes \omega_{2^{i+1}} E_2 \otimes \mathbf{I}_{k-i}, \\ &= E_1 \otimes \mathbf{I}_{i-1} \otimes E_1 \otimes \mathbf{I}_{k-i} + E_1 \otimes \mathbf{I}_{i-1} \otimes E_2 \otimes \mathbf{I}_{k-i} + \\ &\quad E_2 \otimes \mathbf{I}_{i-1} \otimes E_1 \otimes \mathbf{I}_{k-i} + \omega_{2^{i+1}} E_2 \otimes \mathbf{I}_{i-1} \otimes E_2 \otimes \mathbf{I}_{k-i}, \\ &= (E_1 + E_2) \otimes \mathbf{I}_{i-1} \otimes E_1 \otimes \mathbf{I}_{k-i} + \\ &\quad (E_1 + \omega_{2^{i+1}} E_2) \otimes \mathbf{I}_{i-1} \otimes E_2 \otimes \mathbf{I}_{k-i}, \\ &= I_2 \otimes \mathbf{I}_{i-1} \otimes E_1 \otimes \mathbf{I}_{k-i} + R_{i+1} \otimes \mathbf{I}_{i-1} \otimes E_2 \otimes \mathbf{I}_{k-i}, \end{aligned}$$

where we subsequently used the scalar shift and distributivity properties (2.1), and combined the terms with E_1 , respectively E_2 , in the $(i+1)$ st position. \square

Observe that (4.6) splits $\mathbf{I}_k \oplus \mathbf{\Omega}_k$ into a product of k terms where each term has the form (4.8). Lemma 4.5 allows us to modify each term in the product (4.6) to an alternative representation (4.9). In what follows (section 5.3), we will see that every single term in the product corresponds to a specific instance of an elementary operation on a quantum computer and that the combination of Theorem 4.4 and lemma 4.5 leads to a new class of algorithms for implementing $\mathbf{I}_k \oplus \mathbf{\Omega}_k$.

5 Quantum Fourier transform on a quantum computer

We now discuss how the decomposition shown in section 4 enables the discrete Fourier transform to be computed efficiently on a quantum computer.

5.1 Quantum Fourier transform of rank-1 tensors on a classical computer

The Kronecker product form that appears in (4.6) allows for a very efficient application to a rank-1 tensor of the form

$$\mathbf{x}_n = x_1 \otimes x_2 \otimes \dots \otimes x_n, \quad (5.1)$$

where $x_i \in \mathbb{C}^2$ are vectors of length 2. The multiplication of each factor in (4.6) with (5.1) takes $\mathcal{O}(n)$ operations, whereas multiplying the diagonal matrix $\mathbf{I}_n \oplus \mathbf{\Omega}_n$ with a vector of length 2^n takes $\mathcal{O}(2^n)$ operations. Therefore, it is conceivable that the decomposition of the DFT matrix derived in the previous section may yield a more efficient algorithm for performing a discrete Fourier transform of a rank-1 tensor even on a classical computer. In this section, we show that this is not the case. However, in the next section, we will discuss a key feature of a quantum computer that allows to take full advantage of the Kronecker product structure of the decomposition given by Theorem 4.4, and yields an algorithm that only requires $\mathcal{O}(n^2)$ operations.

Recall from Theorem 3.4 that \mathbf{F}'_n can be decomposed as $\mathbf{F}'_n = (I_2 \otimes \mathbf{F}'_{n-1})(\mathbf{I}_{n-1} \oplus \mathbf{\Omega}_{n-1})(H \otimes \mathbf{I}_{n-1})$, where the factor $\mathbf{I}_{n-1} \oplus \mathbf{\Omega}_{n-1}$ can be further decomposed according to (4.6). The multiplication of the rightmost factor of \mathbf{F}'_n with (5.1) yields

$$(H \otimes \mathbf{I}_{n-1})\mathbf{x}_n = (Hx_1) \otimes x_2 \otimes \dots \otimes x_n.$$

Note that this multiplication only requires a constant number of operations to produce Hx_1 , which is independent of the problem size. Furthermore, the result remains a rank-1 tensor. As long as we keep the result in this rank-1 tensor product form, no additional computation is required.

Next, we consider multiplying one factor of $\mathbf{I}_{n-1} \oplus \mathbf{\Omega}_{n-1}$ derived in Theorem 4.4 with a rank-1 tensor of the form (5.1). This amounts to a product of the matrix (4.8), where $k = n - 1$, with the vector \mathbf{x}_n , yielding

$$\begin{aligned} \mathbf{y}_n &= \left[E_1 \otimes \mathbf{I}_{i-1} \otimes I_2 \otimes \mathbf{I}_{n-i-1} + E_2 \otimes \mathbf{I}_{i-1} \otimes R_{i+1} \otimes \mathbf{I}_{n-i-1} \right] \mathbf{x}_n \\ &= (E_1 x_1) \otimes x_2 \otimes \cdots \otimes x_n + (E_2 x_1) \otimes x_2 \otimes \cdots \otimes (R_{i+1} x_{i+1}) \otimes \cdots \otimes x_n. \end{aligned} \quad (5.2)$$

Thus, this multiplication also requires a constant number of operations. However, the rank of the product is generally increased to 2, unless one of the components of x_1 is zero.

If \mathbf{y}_n were to remain a rank-1 tensor after successive multiplications of all factors of $\mathbf{I}_{n-1} \oplus \mathbf{\Omega}_{n-1}$, the complexity of multiplying the last two factors of \mathbf{F}'_n with \mathbf{x}_n would be $\mathcal{O}(n)$. Consequently, by applying this estimate recursively to \mathbf{F}'_k , for $k = n - 1, n - 2, \dots, 1$, the overall complexity of $\mathbf{F}'_n x$ calculation would be $\mathcal{O}(1 + 2 + \cdots + n) = \mathcal{O}(n^2)$, even on a classical computer.

Unfortunately, as we can already see from (5.2), that successive applications of factors of the form $E_1 \otimes \mathbf{I}_{i-1} \otimes I_2 \otimes \mathbf{I}_{n-i} + E_2 \otimes \mathbf{I}_{i-1} \otimes R_{i+1} \otimes \mathbf{I}_{n-i}$ for $i = 1, 2, \dots, n$ tend to increase the rank of the tensor. Nonetheless, as we saw in the proof of Theorem 4.4, the product formulation (4.6) is a redundant representation where most terms cancel out. If instead we decompose the diagonal matrix as

$$\mathbf{I}_{n-1} \oplus \mathbf{\Omega}_{n-1} = E_1 \otimes I_2 \otimes \cdots \otimes I_2 + E_2 \otimes R_2 \otimes \cdots \otimes R_n, \quad (5.3)$$

based on (2.4) and Lemma 4.3, then we get that the product of the diagonal matrix with a rank-1 tensor is given by

$$(\mathbf{I}_{n-1} \oplus \mathbf{\Omega}_{n-1}) \mathbf{x}_n = (E_1 x_1) \otimes x_2 \otimes \cdots \otimes x_n + (E_2 x_1) \otimes (R_2 x_2) \otimes \cdots \otimes (R_n x_n). \quad (5.4)$$

So we see that in general the rank increases to 2 after multiplication with the diagonal matrix, just like after multiplication with a single term of the product representation (5.2).

The redundant product representation (4.6) is more expensive to evaluate classically than (5.4), but will allow an efficient implementation on a quantum computer.

Although, in a classical algorithm, we can attempt to reduce the rank of each product before the multiplication with the next diagonal matrix is initiated, which is the approach taken in [7, 16] for a different tensor representation, it is, in general, not clear how much reduction can be achieved. Even if a rank reduction can be achieved, it will not come free. Consequently, the complexity of computing the discrete Fourier transform on a classical computer through a sequence of tensor products is higher than $\mathcal{O}(n^2)$ even if no explicit sum is performed on the linear combination of rank-1 tensors, which requires a tremendous amount of storage for even a modest sized n , e.g., $n = 100$.

5.2 Qubits and quantum efficiency

However, rank increase is not an issue on a quantum computer. On a quantum computer, a normalized vector of the corresponding tensor \mathbf{x}_n is kept as a quantum state, which can be viewed as a linear combination of a set of tensor product basis states, i.e.,

$$\mathbf{x}_n = \sum_{j_1, \dots, j_n=1}^2 \alpha_{j_1 j_2 \dots j_n} \cdot e_{j_1} \otimes e_{j_2} \otimes \cdots \otimes e_{j_n}. \quad (5.5)$$

The 2^n coefficients $\alpha_{j_1 j_2 \dots j_n}$ of the expansion can be encoded and updated in n quantum bits known as *qubits*. The basis vectors $e_{j_1} \otimes e_{j_2} \otimes \cdots \otimes e_{j_n}$ are the canonical basis of \mathbb{C}^{2^n} . For $n = 1$, a normalized vector $\mathbf{x}_1 \in \mathbb{C}^2$ can simply be written as $\mathbf{x}_1 = \alpha e_1 + \beta e_2$ with α and β satisfying $|\alpha|^2 + |\beta|^2 = 1$. It follows from the superposition principle of quantum mechanics that a single *qubit* can keep both the α and β coefficients simultaneously. For $n > 1$, the 2^n coefficients satisfy the normalization condition $\sum_{j_1, \dots, j_n=1}^2 |\alpha_{j_1 j_2 \dots j_n}|^2 = 1$ and they can be stored in the state of just n qubits. Observe that an n -qubit state (5.5) can be in a complete superposition, i.e., having all coefficients $\alpha_{j_1 j_2 \dots j_n}$ different from zero, and still allow a data sparse representation up to some finite precision in $\mathcal{O}(n)$ space in memory of a classical computer if it is a rank-1 tensor (5.1).

An additional feature of a quantum computer is that it can store a tensor \mathbf{x}_n which has tensor rank $r > 1$ in only n qubits, such a state is also known as an entangled state in quantum mechanics. The minimal representation of an entangled state,

$$\mathbf{x}_n = \sum_{i=1}^r \alpha_i \cdot x_1^{(i)} \otimes x_2^{(i)} \otimes \cdots \otimes x_n^{(i)}, \quad (5.6)$$

where $x_k^{(i)} \in \mathbb{C}^2$, $\|x_k^{(i)}\|_2 = 1$ for $i = 1, 2, \dots, r$, $k = 1, 2, \dots, n$, $\sum_{i=1}^r |\alpha_i|^2 = 1$, and r is an integer rank, is a canonical polyadic decomposition of the n -way tensor \mathbf{x}_n , see Definition 2.3. Notice that if (5.6) is not a *minimal*

representation for \mathbf{x}_n , then the actual tensor rank can be strictly smaller than r . Storing a tensor (5.6) up to finite precision requires $\mathcal{O}(nr)$ classical memory. For $n = 2$, every tensor \mathbf{x}_2 can be represented as a sum of two terms,

$$\mathbf{x}_2 = \alpha_1 \cdot x_1^{(1)} \otimes x_2^{(1)} + \alpha_2 \cdot x_1^{(2)} \otimes x_2^{(2)}, \quad (5.7)$$

which is a singular value decomposition of a 2×2 matrix.

How a quantum computer does this is beyond the scope of this paper, and not relevant for our discussion. It is also worth pointing out that, although a quantum computer can store a quantum state (5.6), it cannot be accessed in the same way a register or memory is accessed on a classical computer. We will discuss the implication of this feature in terms of how QFT can be used later.

The multiplication of a matrix operator with a vector \mathbf{x}_n of tensor rank r (5.6) can be carried out by first multiplying the matrix with the tensor factors $x_1^{(i)} \otimes x_2^{(i)} \otimes \dots \otimes x_n^{(i)}$, $i = 1, \dots, r$. When the matrix operation can be written in terms of Kronecker products of 2×2 matrices, the multiplication does not require a vector to be explicitly formed. All we have to do is to multiply each 2×2 matrix in a Kronecker product with an appropriate component of each of tensor factors $x_j^{(i)}$. For example, applying the matrix in (4.8), with $k = n - 1$, to a tensor factor in (5.6) yields

$$(E_1 x_1^{(i)}) \otimes x_2^{(i)} \otimes \dots \otimes x_j^{(i)} \otimes x_{j+1}^{(i)} \otimes x_{j+2}^{(i)} \dots \otimes x_n^{(i)} + (E_2 x_1^{(i)}) \otimes x_2^{(i)} \otimes \dots \otimes x_j^{(i)} \otimes (R_{j+1} x_{j+1}^{(i)}) \otimes x_{j+2}^{(i)} \otimes \dots \otimes x_n^{(i)}, \quad (5.8)$$

similar to the rank-1 case in (5.2). As the tensor rank representation (5.6) has a minimal number of terms by definition, the multiplication of the matrix (4.9) with \mathbf{x}_n requires classically at least r of the evaluations shown above for $i = 1, \dots, r$. In essence this amounts to computing $R_{j+1} x_{j+1}^{(i)}$ for every tensor factor $x_1^{(i)}$. These changes are managed by a quantum computer instantaneously at a cost that is independent of r , even as the matrix operator would increase the rank r . No additional computation is required to collect or recompute the factors.

Because each factor in the matrix decomposition derived in Theorem 4.4 contains a sum of Kronecker products of 2×2 identity matrices with at most one non-identity matrix. The only operations we need to account for are the multiplications of R_i 's and H with tensor factors $x_j^{(i)}$ and the total cost for updating the state tensor (5.6) is the same for $r = 1$ as it is for a full rank tensor. This particular feature of a quantum computer results in quantum efficiency. It also means that the complexity of multiplying (4.6) with a vector of size 2^n can be evaluated in terms of the number of R_i 's in all of the factors, regardless how large the rank of the vector is when viewed as a n -dimensional tensor. Because the total number of R_i 's is $\mathcal{O}(n^2)$, the QFT can achieve $\mathcal{O}(n^2)$ complexity overall.

5.3 Quantum gates and quantum circuits

Writing down and analyzing the QFT in terms of products of matrices in the form of (4.6) can be cumbersome. In the quantum computing literature, there is a convenient graphical way to depict the unitary transformations on a single or an n -qubit system. The building blocks of the transformation are 2×2 unitary matrices such as the matrix H in (3.6) or R_{i+1} in (4.6). Each one of them is drawn as a square box labelled by a letter and referred to as a quantum gate. A single quantum gate is applied to a single qubit, which is drawn as a line to the left of the gate. The result $\phi = H\psi$ is also drawn as a line to the right of the gate. In quantum mechanics, ψ and ϕ are often written in Dirac's ket notation as shown in the diagram below

$$|\psi\rangle \text{ --- } \boxed{H} \text{ --- } |\phi\rangle.$$

Successive applications of a sequence of m 2×2 unitaries U_1, U_2, \dots, U_m to a single qubit quantum state ψ , i.e., $\phi = U_m \dots U_2 U_1 \psi$, can be drawn as

$$|\psi\rangle \text{ --- } \boxed{U_1} \text{ --- } \boxed{U_2} \text{ --- } \dots \text{ --- } \boxed{U_m} \text{ --- } |\phi\rangle.$$

Such a diagram is often referred to as a quantum circuit [5]. Notice the first unitary to be applied, U_1 , is the leftmost quantum gate in the quantum circuit.

The application of a unitary in the form of $U \otimes I_2$ to a two qubit state \mathbf{x}_2 can be drawn as


$$\begin{array}{c} q_1 \text{ --- } \boxed{U} \text{ ---} \\ q_2 \text{ ---} \end{array}$$

The q_1 and q_2 symbols are used to label the qubits. If the input to this circuit is a rank-1 tensor $\mathbf{x}_2 = x_1 \otimes x_2$, the output of the circuit is simply $Ux_1 \otimes x_2$, and a classical simulation of the circuit has the same cost as its executing on a quantum computer. Notice that the 2×2 identity matrix does not require a quantum gate. In general, if \mathbf{x}_2 is a two qubit state of maximal rank 2, cfr. (5.7), the circuit performs the computation

$$\mathbf{y}_2 = \alpha_1 \cdot (Ux_1^{(1)}) \otimes x_2^{(1)} + \alpha_2 \cdot (Ux_1^{(2)}) \otimes x_2^{(2)},$$

which requires two products with U to compute classically. It follows that if an n -qubit state has rank $\mathcal{O}(2^n)$, the quantum efficiency of quantum computers can lead to exponential speedups in comparison to classical computers.

This type of diagram can be easily generalized for unitary transformations applied to state vectors encoded by several qubits. In the context of the QFT, a particularly interesting unitary transformation is the one represented by the matrices of Lemma 4.5. As we saw in (5.8), the action on the $(j+1)$ st tensor factor $x_{j+1}^{(i)}$ depends on the coefficients of the first tensor factor, $E_1 x_1^{(i)}$ and $E_2 x_1^{(i)}$. If the first tensor factor is equal to e_1 , then the identity matrix is applied to $x_{j+1}^{(i)}$. If $x_1^{(i)}$ is equal to e_2 , then the $(j+1)$ st factor changes to $R_{j+1} x_{j+1}^{(i)}$. Consequently, this type of transformation is called a *controlled* unitary with the first qubit as the controlling qubit and the $(j+1)$ st qubit as the target qubit. In general, $x_1^{(i)}$ can be in a superposition of e_1 and e_2 and the result is the linear combination (5.8) which can increase the tensor rank as we discussed in Section 5.1. The diagrammatic notation for this controlled unitary is drawn on the left in the diagram below. Note that a solid circle is drawn on the line representing the controlling qubit. It is connected to the qubit to be transformed via a vertical line. Lemma 4.5 shows the equivalence between the controlled unitaries on the left and the right for controlled- R gates.




Another useful controlled unitary is the *controlled*-NOT or CNOT unitary. As a 2-qubit operator with the first qubit being the control and the second qubit being the target, it can be written as

$$X_c = E_1 \otimes I_2 + E_2 \otimes X,$$

where X is the NOT operator

$$X = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix},$$


that maps e_1 to e_2 and vice versa. The conventional diagram for a CNOT gate in quantum computing is



One particular use of the CNOT gate is to implement the SWAP operator defined as

$$\text{SWAP}(\psi_1 \otimes \psi_2) = \psi_2 \otimes \psi_1.$$

This is sometimes represented by




in a quantum circuit. It can be easily verified that a SWAP gate can be decomposed as the product of three CNOTS with alternating controlling qubits. As a result, the CNOT implementation of a SWAP gate can be drawn as



Controlled unitaries can also be defined to have multiple target qubits. The matrix $E_1 \otimes \mathbf{I}_k + E_2 \otimes \mathbf{U}_k$ is a controlled- \mathbf{U}_k gate with the first qubit as a control and qubits $2, \dots, k+1$ as target. It follows from (3.4) that a bit reversal operator \mathbf{P}_n can be written in terms of a sequence of SWAP operations. Therefore, \mathbf{P}_n can be implemented using sequence of CNOT gates.

In general, a unitary transformation \mathbf{U}_n applied to an n -qubit state ψ_n can be drawn as



An alternative and simplified way to draw this transformation is

$$|\psi_n\rangle \text{ --- } \boxed{\mathbf{U}_n} \text{ --- } |\phi_n\rangle,$$

where the wire with a ‘/’ through indicates a wire representing n qubits.

5.4 Quantum circuit for the quantum Fourier transform

We now show how the QFT matrix decomposition derived in Theorem 3.5 and Theorem 4.4 can be expressed succinctly by using the quantum circuit diagrams introduced in section 5.3. We start from Theorem 3.5 and write down the circuit representation of the decomposition of the DFT matrix $\mathbf{F}_n \in \mathbb{C}^{2^n \times 2^n}$ as follows

$$|\psi_n\rangle \text{ --- } \boxed{\mathbf{F}_n} \text{ --- } = \begin{array}{c} q_1 \\ q_2 \\ \vdots \\ q_{n-1} \\ q_n \end{array} \begin{array}{c} \boxed{\mathbf{B}_n} \\ \boxed{\mathbf{B}_{n-1}} \\ \vdots \\ \boxed{\mathbf{B}_2} \\ \boxed{\mathbf{B}_1} \end{array} \begin{array}{c} \cdots \\ \cdots \\ \vdots \\ \cdots \\ \cdots \end{array} \begin{array}{c} \boxed{\mathbf{P}_n} \end{array} \quad (5.9)$$

Note that the gate associated with \mathbf{B}_{k+1} is applied to the last $k+1$ qubits.

Each multi-qubit gate block of this quantum circuit can be decomposed into a sequence of single or two-qubit gates. For example, the bit-reversal permutation matrix \mathbf{P}_n can first be written in terms of $\lfloor n/2 \rfloor$ SWAP gates as follows

$$|\psi_n\rangle \text{ --- } \boxed{\mathbf{P}_n} \text{ --- } = \begin{array}{c} q_1 \\ q_2 \\ \vdots \\ q_{n-1} \\ q_n \end{array} \begin{array}{c} \cdots \\ \cdots \\ \vdots \\ \cdots \\ \cdots \end{array} \quad (5.10)$$

Since each SWAP can be implemented with 3 CNOT gates, \mathbf{P}_n requires $\lfloor 3n/2 \rfloor$ CNOT gates.

It follows from (3.6) that the \mathbf{B}_{k+1} gates can be decomposed further as

$$|\psi_{k+1}\rangle \text{ --- } \boxed{\mathbf{B}_{k+1}} \text{ --- } = \begin{array}{c} q_1 \\ q_2 \\ \vdots \\ q_k \\ q_{k+1} \end{array} \begin{array}{c} \boxed{H} \\ \vdots \\ \boxed{\mathbf{I}_k \oplus \mathbf{\Omega}_k} \end{array} \quad (5.11)$$

Furthermore, the decomposition given in Theorem 4.4 can be used to rewrite the $\mathbf{I}_k \oplus \mathbf{\Omega}_k$ block in terms of controlled operations involving two qubits, as shown in the rightmost circuit below


$$\begin{array}{c} q_1 \\ q_2 \\ \vdots \\ q_k \\ q_{k+1} \end{array} \begin{array}{c} \boxed{\mathbf{I}_k \oplus \mathbf{\Omega}_k} \end{array} = \begin{array}{c} q_1 \\ q_2 \\ \vdots \\ q_k \\ q_{k+1} \end{array} \begin{array}{c} \bullet \\ \boxed{\mathbf{\Omega}_k} \end{array} = \begin{array}{c} q_1 \\ q_2 \\ \vdots \\ q_k \\ q_{k+1} \end{array} \begin{array}{c} \bullet \\ \boxed{R_2} \\ \vdots \\ \boxed{R_k} \\ \boxed{R_{k+1}} \end{array} = \begin{array}{c} q_1 \\ q_2 \\ \vdots \\ q_k \\ q_{k+1} \end{array} \begin{array}{c} \bullet \\ \bullet \\ \vdots \\ \bullet \\ \bullet \end{array} \begin{array}{c} \cdots \\ \cdots \\ \vdots \\ \cdots \\ \cdots \end{array} \begin{array}{c} \boxed{R_2} \\ \vdots \\ \boxed{R_k} \\ \cdots \\ \boxed{R_{k+1}} \end{array} \quad (5.12)$$

The circuit on the left is a dense $k+1$ qubit diagonal gate $\mathbf{I}_k \oplus \mathbf{\Omega}_k$, which we can rewrite as a controlled- $\mathbf{\Omega}_k$ with the first qubit as control and the next k qubits as target based on (4.7). From Lemma 4.3 we have that $\mathbf{\Omega}_k$ can be decomposed as a Kronecker product of R matrices, which gives the third circuit above, where the control qubit controls the whole group $R_2 \otimes \cdots \otimes R_{k+1}$. This circuit can evidently be expanded as a product of separate controlled- R gates to obtain the rightmost circuit, which gives a compact visual proof of Theorem 4.4 only using

circuit diagrams. The decomposition of $\mathbf{I}_k \oplus \mathbf{\Omega}_k$ in k gates is exceptionally efficient, because the implementation of an n -qubit diagonal operator requires $\mathcal{O}(2^n)$ gates in general [21, 1].

All the controlled- R gates in this decomposition are diagonal matrices and commute so the order in which they are applied does not change the outcome. Furthermore, we have from Lemma 4.5 that the roles of the control and target qubits can be swapped without changing the result. This leads us to a whole class of equivalent quantum circuits for $\mathbf{I}_k \oplus \mathbf{\Omega}_k$.

Finally combining (5.9)–(5.12) gives the following complete QFT quantum circuit



In this QFT circuit, we swapped the control and target qubits compared to Theorem 4.4 for all controlled- R gates. This particular choice gives the same QFT circuit that is discussed in [14, section 5.1].

The gate count of for each type of gate used in a QFT is summarized in Table 1. The total number of elementary gates required to implement the QFT is $\mathcal{O}(n^2)$, which defines the complexity of the QFT algorithm. It represents an exponential improvement over the $\mathcal{O}(2^{2n})$ complexity of the FFT algorithm.

Table 1: Gate count for the QFT circuit on n qubits.

Matrix	Count	Gate type
\mathbf{P}_n	$\lfloor 3n/2 \rfloor$	CNOT
$\mathbf{A}_n^{(0)} \cdots \mathbf{A}_n^{(n-1)}$	n	Hadamard
	$n(n-1)/2$	controlled- R

6 Generalization to Radix- d Quantum Fourier transform

The QFT decomposition presented in section 4 can be easily generalized for DFT matrices of dimension $d^n \times d^n$ for any integer $d > 2$. The generalization relies on a radix- d FFT decomposition and the base- d representation of an integer j

$$j = [j_1 j_2 \cdots j_{n-1} j_n] = j_1 \cdot d^{n-1} + j_2 \cdot d^{n-2} + \cdots + j_{n-1} \cdot d^1 + j_n \cdot d^0,$$

where $j_i = \{0, 1, \dots, d-1\}$ for $i = 1, \dots, n$. Hence, Lemmas 2.5 and 2.6 generalize in a straightforward manner to base- d .

6.1 Radix- d decomposition of the discrete Fourier transform matrix

In this section, we use $\mathbf{A}_n = A_{d^n}$ to denote matrices of dimension d^n .

Definition 6.1. Define $\mathbf{\Omega}_n \in \mathbb{C}^{d^n \times d^n}$, $R_n \in \mathbb{C}^{d \times d}$ as the following matrices:

$$\mathbf{\Omega}_n = \mathbf{\Omega}_{d^n} := \begin{bmatrix} \omega_{d^{n+1}}^0 & & & \\ & \omega_{d^{n+1}}^1 & & \\ & & \ddots & \\ & & & \omega_{d^{n+1}}^{d^n-1} \end{bmatrix}, \quad R_n = \begin{bmatrix} \omega_{d^n}^0 & & & \\ & \omega_{d^n}^1 & & \\ & & \ddots & \\ & & & \omega_{d^n}^{d-1} \end{bmatrix},$$

where $\omega_{d^i} = e^{\frac{-2\pi i}{d^i}}$.

If the DFT matrix is of dimension $d^n \times d^n$, then it can be factored in $d \times d$ block partitioning, also known as the radix- d factorization of the DFT matrix [18, 15]. This is summarized in the following two theorems which naturally extend the results of the radix-2 framework of section 3.2.

Theorem 6.2 (See [18]). *Let $\mathbf{F}'_n = \mathbf{P}_n \mathbf{F}_n$ be the $d^n \times d^n$ base- d reversed DFT matrix. Then \mathbf{F}'_n admits the following factorization:*

$$\mathbf{F}'_n = (I_d \otimes \mathbf{F}'_{n-1})(\mathbf{I}_{n-1} \oplus \mathbf{\Omega}_{n-1} \oplus \cdots \oplus \mathbf{\Omega}_{n-1}^{d-1})(F_d \otimes \mathbf{I}_{n-1}), \quad (6.1)$$

where $\mathbf{\Omega}_{n-1}$ is given by Definition 6.1 and F_d is the $d \times d$ DFT matrix defined in Definition 3.2.

Equation (6.1) essentially decomposes \mathbf{F}'_n into the Kronecker product of the $d^{n-1} \times d^{n-1}$ DFT matrix \mathbf{F}'_{n-1} and a $d \times d$ block matrix with block size $d^{n-1} \times d^{n-1}$. The full radix- d factorization of the DFT matrix follows from Theorem 6.2 via the same simple induction argument used in the proof of Theorem 3.5.

Theorem 6.3. *The $d^n \times d^n$ DFT matrix \mathbf{F}_n can be factored as:*

$$\mathbf{F}_n = \mathbf{P}_n \mathbf{F}'_n = \mathbf{P}_n \mathbf{A}_n^{(0)} \cdots \mathbf{A}_n^{(n-1)}, \quad (6.2)$$

where, for $k = 0, \dots, n-1$,

$$\mathbf{A}_n^{(k)} = \mathbf{I}_{n-k-1} \otimes \mathbf{B}_{k+1}, \quad \text{and} \quad \mathbf{B}_{k+1} = (\mathbf{I}_k \oplus \mathbf{\Omega}_k \oplus \cdots \oplus \mathbf{\Omega}_k^{d-1})(\mathbf{F}_1 \otimes \mathbf{I}_k). \quad (6.3)$$

with $\mathbf{\Omega}_k$ defined in Definition 6.1 and $\mathbf{F}_1 = F_d$ the $d \times d$ DFT matrix given by Definition 3.2.

Definition 6.4. *The base- d reversal permutation matrix $\mathbf{P}_n \in \mathbb{C}^{d^n \times d^n}$ is defined as the permutation matrix which satisfies:*

$$\mathbf{P}_n(v_1 \otimes \cdots \otimes v_n) = v_n \otimes \cdots \otimes v_1,$$

where $v_k \in \mathbb{C}^d$, for $k = 1, \dots, n$.

Definition 6.4 shows that the permutation matrix can be implemented in $\lfloor n/2 \rfloor$ SWAPs of d dimensional vectors in a Kronecker product of n such vectors.

6.2 Decomposition of the diagonal matrix

In this section we decompose the diagonal matrix $\mathbf{I}_k \oplus \mathbf{\Omega}_k \oplus \cdots \oplus \mathbf{\Omega}_k^{d-1}$ into k controlled- R unitaries, with R defined in Definition 6.1, following an analogous strategy to section 4. We start again with a decomposition of $\mathbf{\Omega}_n$ as a Kronecker product of $d \times d$ R_i matrices.

Lemma 6.5. *Let R_n be defined by Definition 6.1. Then*

$$R_n^{d^j} = R_{n-j},$$

for $j \in \mathbb{N}$.

Proof. The proof proceeds analogously to Lemma 4.2. □

Lemma 6.6. *Let $\mathbf{\Omega}_n$ and R_n be defined by Definition 6.1. Then*

$$\mathbf{\Omega}_n = R_2 \otimes R_3 \otimes \cdots \otimes R_n \otimes R_{n+1}.$$

Proof. The proof proceeds analogously to Lemma 4.3, replacing the binary representation with base- d representation and using Lemma 6.5. □

The generalization of Theorem 4.4 to the decomposition of the matrix $\mathbf{I}_k \oplus \mathbf{\Omega}_k \oplus \cdots \oplus \mathbf{\Omega}_k^{d-1}$ can be stated as follows.

Theorem 6.7. *The diagonal operators $\mathbf{I}_k \oplus \mathbf{\Omega}_k \oplus \cdots \oplus \mathbf{\Omega}_k^{d-1} \in \mathbb{C}^{d^{k+1} \times d^{k+1}}$ from Theorem 6.3 admit the decomposition:*

$$\mathbf{I}_k \oplus \mathbf{\Omega}_k \oplus \cdots \oplus \mathbf{\Omega}_k^{d-1} = \prod_{i=1}^k \sum_{\ell=1}^d E_\ell \otimes \mathbf{I}_{i-1} \otimes R_{i+1}^{\ell-1} \otimes \mathbf{I}_{k-i} \quad (6.4)$$

with $E_\ell := e_\ell e_\ell^\top \in \mathbb{C}^{d \times d}$ and R_{i+1} defined in Definition 6.1.

Proof. We start by splitting the direct sum of d diagonal matrices as a sum of Kronecker products, which yields

$$\mathbf{I}_k \oplus \mathbf{\Omega}_k \oplus \cdots \oplus \mathbf{\Omega}_k^{d-1} = E_1 \otimes \mathbf{I}_k + E_2 \otimes \mathbf{\Omega}_k + \cdots + E_d \otimes \mathbf{\Omega}_k^{d-1}, \quad (6.5)$$

where we used a generalization of (2.4). Next, we rewrite the ℓ th term on the right-hand side in a redundant form based on the mixed product identity of Kronecker product (2.2). Using Lemma 6.6 and the property $E_\ell = E_\ell^2$, we get

$$\begin{aligned} E_\ell \otimes \Omega_k^{\ell-1} &= E_\ell \otimes R_2^{\ell-1} \otimes \cdots \otimes R_{k+1}^{\ell-1}, \\ &= (E_\ell \otimes R_2^{\ell-1} \otimes I_2 \otimes \cdots \otimes I_2)(E_\ell \otimes I_2 \otimes R_3^{\ell-1} \otimes \cdots \otimes I_2) \cdots (E_\ell \otimes I_2 \otimes I_2 \otimes \cdots \otimes R_{k+1}^{\ell-1}), \\ &= \prod_{i=1}^k E_\ell \otimes \mathbf{I}_{i-1} \otimes R_{i+1}^{\ell-1} \otimes \mathbf{I}_{k-i}. \end{aligned}$$

Hence, the sum of Kronecker products (6.5) is equal to the product of sums, i.e.,

$$\mathbf{I}_k \oplus \Omega_k \oplus \cdots \oplus \Omega_k^{d-1} = \sum_{\ell=1}^d \prod_{i=1}^k E_\ell \otimes \mathbf{I}_{i-1} \otimes R_{i+1}^{\ell-1} \otimes \mathbf{I}_{k-i} = \prod_{i=1}^k \sum_{\ell=1}^d E_\ell \otimes \mathbf{I}_{i-1} \otimes R_{i+1}^{\ell-1} \otimes \mathbf{I}_{k-i}$$

because the mixed terms in the latter expression cancel out due to the property $E_i E_j = 0$ for $i \neq j$. \square

The position of the R_{i+1} and E_ℓ matrices in the Kronecker product expression can again be swapped.

Lemma 6.8. *The Kronecker product matrices in (6.4) satisfy*

$$\sum_{\ell=1}^d E_\ell \otimes \mathbf{I}_{i-1} \otimes R_{i+1}^{\ell-1} \otimes \mathbf{I}_{k-i} = \sum_{\ell=1}^d R_{i+1}^{\ell-1} \otimes \mathbf{I}_{i-1} \otimes E_\ell \otimes \mathbf{I}_{k-i}, \quad (6.6)$$

for $i = 1, \dots, k$.

Proof. Splitting R_{i+1} in its E_1, E_2, \dots, E_d components, i.e.,

$$R_{i+1} = E_1 + \omega_{d^{i+1}} E_2 + \cdots + \omega_{d^{i+1}}^{d-1} E_d,$$

yields

$$\begin{aligned} &\sum_{\ell=1}^d E_\ell \otimes \mathbf{I}_{i-1} \otimes R_{i+1}^{\ell-1} \otimes \mathbf{I}_{k-i} \\ &= E_1 \otimes \mathbf{I}_{i-1} \otimes (E_1 + E_2 + \cdots + E_d) \otimes \mathbf{I}_{k-i} + \\ &\quad E_2 \otimes \mathbf{I}_{i-1} \otimes (E_1 + \omega_{d^{i+1}} E_2 + \cdots + \omega_{d^{i+1}}^{d-1} E_d) \otimes \mathbf{I}_{k-i} + \\ &\quad \cdots + \\ &\quad E_d \otimes \mathbf{I}_{i-1} \otimes (E_1 + \omega_{d^{i+1}}^{d-1} E_2 + \cdots + \omega_{d^{i+1}}^{(d-1)^2} E_d) \otimes \mathbf{I}_{k-i}, \\ &= E_1 \otimes \mathbf{I}_{i-1} \otimes E_1 \otimes \mathbf{I}_{k-i} + E_2 \otimes \mathbf{I}_{i-1} \otimes E_1 \otimes \mathbf{I}_{k-i} + \cdots + E_d \otimes \mathbf{I}_{i-1} \otimes E_1 \otimes \mathbf{I}_{k-i} + \\ &\quad E_1 \otimes \mathbf{I}_{i-1} \otimes E_2 \otimes \mathbf{I}_{k-i} + \omega_{d^{i+1}} E_2 \otimes \mathbf{I}_{i-1} \otimes E_2 \otimes \mathbf{I}_{k-i} + \cdots + \omega_{d^{i+1}}^{d-1} E_d \otimes \mathbf{I}_{i-1} \otimes E_2 \otimes \mathbf{I}_{k-i} + \\ &\quad \cdots + \\ &\quad E_1 \otimes \mathbf{I}_{i-1} \otimes E_d \otimes \mathbf{I}_{k-i} + \omega_{d^{i+1}}^{d-1} E_2 \otimes \mathbf{I}_{i-1} \otimes E_d \otimes \mathbf{I}_{k-i} + \cdots + \omega_{d^{i+1}}^{(d-1)^2} E_d \otimes \mathbf{I}_{i-1} \otimes E_d \otimes \mathbf{I}_{k-i}, \\ &= (E_1 + E_2 + \cdots + E_d) \otimes \mathbf{I}_{i-1} \otimes E_1 \otimes \mathbf{I}_{k-i} + \\ &\quad (E_1 + \omega_{d^{i+1}} E_2 + \cdots + \omega_{d^{i+1}}^{d-1} E_d) \otimes \mathbf{I}_{i-1} \otimes E_2 \otimes \mathbf{I}_{k-i} + \\ &\quad \cdots + \\ &\quad (E_1 + \omega_{d^{i+1}}^{d-1} E_2 + \cdots + \omega_{d^{i+1}}^{(d-1)^2} E_d) \otimes \mathbf{I}_{i-1} \otimes E_{d-1} \otimes \mathbf{I}_{k-i}, \\ &= \sum_{\ell=1}^d R_{i+1}^{\ell-1} \otimes \mathbf{I}_{i-1} \otimes E_\ell \otimes \mathbf{I}_{k-i}, \end{aligned}$$

where we subsequently used the distributivity and scalar shift properties (2.1) of the Kronecker product, and combined the terms with the same matrix, E_1, E_2 up to E_d , in the $(i+1)$ st position. \square

6.3 Quantum circuit for Radix-d QFT


The generalization of a qubit for encoding a d -dimensional vector is a *qudit*. When $d = 3$, it is called a *qutrit*. The implementation of the QFT on a quantum computer working with qudits is studied in [2, 13, 6]. In the current section, we show that our derivation of the QFT on a qubit system extends to qudit systems in a straightforward manner, and leads to a class of equivalent QFT circuits on a qudit system.

Single qudit gates are unitary $d \times d$ matrices and the state space of multi-qudit systems is $(\mathbb{C}^d)^{\otimes n}$. A controlled- U operation acting on two qudits is denoted with the same circuit symbol as for the qudit case, but the action on the target qudit is now dependent on the d different states of the control qudit:

$$\begin{array}{c} \text{---} \\ | \\ \text{---} \\ \boxed{U} \\ | \\ \text{---} \end{array} \leftrightarrow (E_1 \otimes I_d) + (E_2 \otimes U) + \dots + (E_d \otimes U^{d-1}) =: U_c,$$


with $E_k := e_k e_k^\top \in \mathbb{C}^{d \times d}$, for $k = 1, \dots, d$. This definition is in agreement with the definition of a controlled-NOT operation on a d -level quantum system, also called a SUM gate in [9].

Theorem 6.3 gives the following quantum circuit for the Fourier transform of dimension $d^n \times d^n$ acting on n qudits:



The structure of this circuit is exactly the same as in the qubit case. The difference is that the Hadamard gate H is replaced by the *Fourier gate* F_d which applies the d -dimensional DFT matrix to a single qudit, and that the diagonal operators now consist of a direct sum of d diagonal matrices.

Just like before, these diagonal operators can be synthesized in controlled- R_j operators acting on two qudits by Theorem 6.7 and Lemma 6.8:



The representation is only one of the many possible implementations of the diagonal matrix under the freedom allowed by the commutativity and swap property (Lemma 6.8) of controlled- R gates.

We conclude that the gate count for a quantum Fourier transform on an n qudit d -level system is still $\mathcal{O}(n^2)$, but now for a QFT of a state vector of dimension d^n . This results in a $\log_2 d$ reduction in gates for a qudit system compared to a qubit system that has a state space of the same dimension at the cost of more complex controlled operations.

7 Conclusion

In this paper, we showed how the quantum Fourier transform algorithm can be derived as a decomposition of the discrete Fourier matrix. The derivation starts from the radix-2 decomposition of the DFT matrix that yields the FFT algorithm, and only makes use of Kronecker products to further decompose a diagonal matrix into a product of simpler unitary matrices, each of which can be written as the sum of two Kronecker products of 2×2 matrices. This alternative approach to the derivation of the quantum Fourier transform in [14, 4] requires little knowledge of quantum computing.

We showed in section 5.1 that the Kronecker structure of the QFT decomposition does not lead to an immediate reduction in complexity compared to that of the FFT algorithm on a classical computer, even when the QFT is applied to a rank-1 tensor. We explained why the complexity of the QFT on a quantum computer can be evaluated by counting the number of 2×2 matrices produced by the QFT matrix decomposition, which is $\mathcal{O}((\log N)^2)$ for

a transformation of a vector of size $N = 2^n$. We made the connection between the matrix decomposition of the QFT and the quantum circuit representation widely used in the quantum computing literature. We pointed out that the QFT decomposition of the DFT matrix and the corresponding quantum circuit is not unique. The non-uniqueness of the decomposition can potentially provide some flexibility in the quantum circuit topology in terms of the placement of controlled- R gates. We also generalized the radix-2 decomposition to a radix- d decomposition which requires the QFT to be implemented on a quantum computer equipped with qudits.

Although the QFT can be performed efficiently on a quantum computer, the result of the transform is not easily accessible. This is a key difference between the QFT and FFT, and between a quantum and classical algorithm in general. Because of this, one cannot use the QFT in the same way the FFT is used. For example, it is not straightforward to perform a fast convolution of two vectors [12], which is the most widely used application of the FFT, on a quantum computer using QFTs. On the other hand, the QFT is used as a building block for several quantum algorithms such as in phase estimation [11, 17]. However, how the QFT is used in these algorithms is beyond the scope of this paper.

Acknowledgments

This work was supported by the Laboratory Directed Research and Development Program of Lawrence Berkeley National Laboratory under U.S. Department of Energy Contract No. DE-AC02-05CH11231.

References

- [1] S. S. Bullock and I. L. Markov. Asymptotically optimal circuits for arbitrary n-qubit diagonal computations. *Quantum Inf. Comput.*, 4(1):27–47, 2004.
- [2] Y. Cao, S.-G. Peng, C. Zheng, and G.-L. Long. Quantum Fourier transform and phase estimation in qudit system. *Commun. Theor. Phys. (Beijing)*, 55(5):790–794, 2011.
- [3] J. W. Cooley and J. W. Tukey. An algorithm for the machine calculation of complex Fourier series. *Math. Comp.*, 19:297–301, 1965.
- [4] D. Coppersmith. An approximate Fourier transform useful in quantum factoring. *IBM Research Report RC 19642*, 1994.
- [5] D. Deutsch. Quantum Computational Networks. *Proc. R. Soc. Lond. Ser. A Math. Phys. Eng. Sci.*, 425(1868):73–90, 1989.
- [6] S. Dogra, A. Dorai, and K. Dorai. Implementation of the quantum Fourier transform on a hybrid qubit-qudit NMR quantum emulator. *Int. J. Quantum Inf.*, 13(07):1550059, 2015.
- [7] S. Dolgov, B. Khoromskij, and D. Savostyanov. Superfast Fourier transform using QTT approximation. *J. Fourier Anal. Appl.*, 18(5):915–953, 2012.
- [8] M. Drubin. Kronecker product factorization of the FFT matrix. *IEEE Trans. Comput.*, C-20(5):590–593, 1971.
- [9] D. Gottesman. Fault-tolerant quantum computation with higher-dimensional systems. *Chaos Solitons Fractals*, 10(10):1749–1758, 1999.
- [10] F. L. Hitchcock. The expression of a tensor or a polyadic as a sum of products. *Journal of Mathematics and Physics*, 6(1-4):164–189, 1927.
- [11] A. Y. Kitaev. Quantum measurements and the Abelian stabilizer problem, 1995.
- [12] C. Lomont. Quantum convolution and quantum correlation algorithms are physically impossible. *arXiv: quant-ph/0309070*, 2003.
- [13] A. Muthukrishnan and C. R. Stroud. Quantum fast Fourier transform using multilevel atoms. *J. Modern Opt.*, 49(13):2115–2127, 2002.
- [14] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, Cambridge, 2010.
- [15] M. C. Pease. An adaptation of the fast Fourier transform for parallel processing. *J. ACM*, 15(2):252–264, 1968.
- [16] D. Savostyanov. QTT-rank-one vectors with QTT-rank-one and full-rank Fourier images. *Linear Algebra Appl.*, 436(9):3215–3224, 2012.

- [17] P. W. Shor. Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer. *SIAM J. Comput.*, 26(5):1484–1509, 1997.
- [18] H. Sloate. Matrix representations for sorting and the fast Fourier transform. *IEEE Trans. Circuits Syst.*, 21(1):109–116, 1974.
- [19] F. Sullivan and J. Dongarra. The Top 10 Algorithms. *Comput. Sci. Eng.*, 2(1):22–23, 2000.
- [20] C. Van Loan. *Computational Frameworks for the Fast Fourier Transform*. Society for Industrial and Applied Mathematics, 1992.
- [21] J. M. Welch. *On the Synthesis of Quantum Circuits for Diagonal Operators in Quantum Computation*. PhD thesis, Harvard University, Graduate School of Arts & Sciences, 2015.