

Another Way to Perform the Quantum Fourier Transform in Linear Parallel Time

Cristopher Moore

Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, New Mexico 87501
moore@santafe.edu

Abstract. We exhibit a quantum circuit that performs the Quantum Fourier Transform on n qubits in $\mathcal{O}(n)$ depth. Thus, a parallel quantum computer can carry out the QFT in linear time. Griffiths and Niu have already shown this, so this paper is little more than an exercise in quantum circuit design; but perhaps it illustrates a worthwhile idea. We also speculate as to whether the QFT might be in the class **QNC**¹ of problems solvable in logarithmic parallel time.

Shor's factoring algorithm [6] suggests that quantum computers can do things in polynomial time that classical computers cannot. However, since decoherence due to storage errors is a function of time, we should also ask to what extent we can parallelize quantum algorithms; if we can do many quantum operations at once, rather than serially, we can solve larger problems before our computer decoheres.

Consider a quantum circuit operating on a set of qubits, containing one-qubit gates (2×2 unitary matrices) and the two-qubit controlled-not gate; these are universal for quantum computation [1, 4]. We can define the *depth* of this circuit as the number of layers, where each layer consists of gates operating on mutually disjoint sets of qubits; that is, each qubit interacts with at most one other qubit at a time. (In a model of quantum computation where one qubit can simultaneously interact with several others, we could allow gates operating on the same qubit in the same level, as long as these gates all mutually commute.)

The heart of Shor's algorithm is the *Quantum Fourier Transform*. If we represent n -digit numbers $|a\rangle$ with n qubits, the QFT maps $|a\rangle$ to

$$2^{-n/2} \sum_{b=0}^{2^n-1} e^{2\pi i ab/2^n} |b\rangle$$

In this paper, we exhibit a circuit with depth $\mathcal{O}(n)$ for performing the QFT. Griffiths and Niu have already done this, in fact in a more natural way [3]. However, perhaps the reader will enjoy a new construction using slightly different ideas.

The standard quantum algorithm for the QFT takes $n(n-1)/2$ gates [2, 6]. One way to construct it is to reshuffle the rows of the matrix by putting the

digits of the input in reverse order. Then for $n = 3$, for instance, we have

$$\begin{pmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 & 1 & -1 & 1 & -1 \\ 1 & i & -1 & -i & 1 & i & -1 & -i \\ 1 & -i & -1 & i & 1 & -i & -1 & i \\ 1 & e^{\pi i/4} & i & e^{3\pi i/4} & -1 & e^{5\pi i/4} & -i & e^{7\pi i/4} \\ 1 & e^{5\pi i/4} & i & e^{7\pi i/4} & -1 & e^{\pi i/4} & -i & e^{3\pi i/4} \\ 1 & e^{3\pi i/4} & -i & e^{\pi i/4} & -1 & e^{7\pi i/4} & i & e^{5\pi i/4} \\ 1 & e^{7\pi i/4} & -i & e^{5\pi i/4} & -1 & e^{3\pi i/4} & i & e^{\pi i/4} \end{pmatrix}$$

where we are suppressing a factor of $2^{-3/2}$.

If we call this $F(3)$, we immediately notice that its upper-left and upper-right quadrants are

$$\begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 \\ 1 & i & -1 & -i \\ 1 & -i & -1 & i \end{pmatrix}$$

which is simply $F(2)$. The lower-left and lower-right quadrants of $F(3)$ are $F(2)$ and $-F(2)$, with a series of phase shifts applied to the columns; this can be expressed by multiplying on the right by the matrix

$$\begin{pmatrix} 1 & & & \\ & e^{\pi i/4} & & \\ & & i & \\ & & & e^{3\pi i/4} \end{pmatrix}$$

which we will call M . In general, we can write

$$\begin{aligned} F(n+1) &= \frac{1}{\sqrt{2}} \begin{pmatrix} F & F \\ FM & -FM \end{pmatrix} \\ &= \begin{pmatrix} F & \\ & F \end{pmatrix} \cdot \begin{pmatrix} 1 & \\ & M \end{pmatrix} \cdot \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \end{aligned}$$

We recognize this as the circuit for $F(n)$ applied to the n least significant qubits, followed by a gate where the most significant qubit controls whether or not to apply the phase shifts M , followed by the *Hadamard operator*

$$H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$$

applied to the most significant qubit.

Finally, note that M is simply a tensor product of independent one-qubit operations

$$M = \begin{pmatrix} 1 & \\ & i \end{pmatrix} \otimes \begin{pmatrix} 1 & \\ & e^{\pi i/4} \end{pmatrix} \otimes \begin{pmatrix} 1 & \\ & e^{\pi i/8} \end{pmatrix} \otimes \cdots$$

Then the controlled- M gate becomes a series of controlled phase-shift gates

$$\begin{pmatrix} 1 & \\ & M \end{pmatrix} = \begin{pmatrix} 1 & & \\ & 1 & \\ & & 1 \end{pmatrix} \otimes \begin{pmatrix} 1 & & \\ & 1 & \\ & & e^{\pi i/4} \end{pmatrix} \otimes \dots$$

These gates are symmetric, in that the “controlled” and “controlling” qubits are interchangeable. Putting all this together gives us the recursive construction shown in Figure 1.

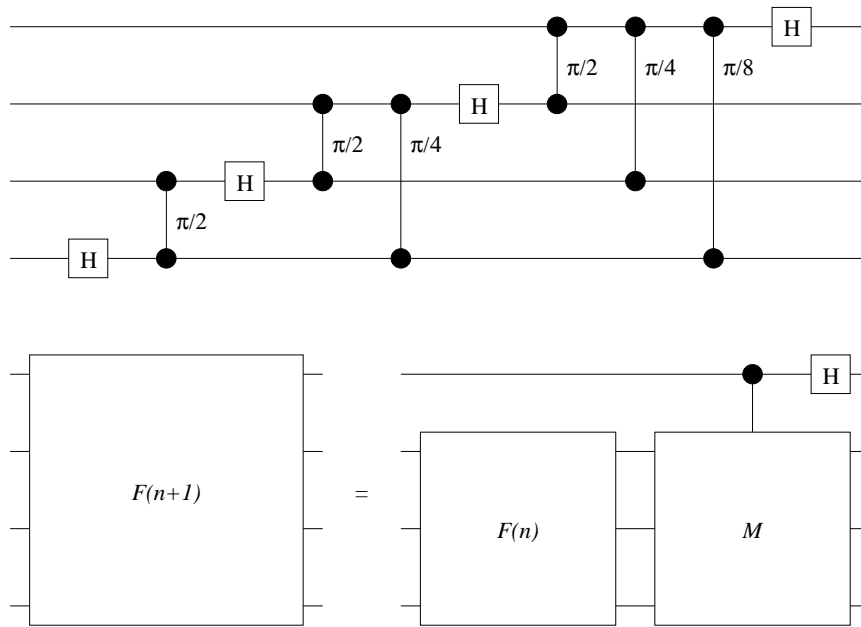


Fig. 1. The standard algorithm for the QFT, with $\mathcal{O}(n^2)$ gates applied serially. The controlled phase-shift gates are symmetric, and are shown as bonds between two qubits. Here $n = 4$.

To what extent can this circuit be parallelized? Even though all the phase shift gates within a given pair of H 's commute with each other, we can't perform them simultaneously unless we can couple one qubit to multiple qubits at the same time, and they don't commute with the H preceding them. Thus, it would appear that all $\mathcal{O}(n^2)$ gates have to be applied in series.

However, we can turn this circuit into one where most of the gates commute, so that many can be performed simultaneously, in the following way. Note that H is its own inverse. Conjugating a phase shift gate with H gives

$$H \begin{pmatrix} 1 & \\ & e^{i\theta} \end{pmatrix} H = \frac{1}{2} \begin{pmatrix} 1 + e^{i\theta} & 1 - e^{i\theta} \\ 1 - e^{i\theta} & 1 + e^{i\theta} \end{pmatrix}$$

Call this matrix R_θ . Then if we pass the H operators through the phase shifts to the right, we get the circuit shown in Figure 2, where the controlled phase-shift gates have been replaced by controlled- R_θ gates.




Fig. 2. The circuit of figure 1, after the H operators have been moved to the left, leaving controlled- R_θ gates behind.

Now note that two controlled- R_θ gates commute in every case except when the ‘control’ of one is the ‘controlled’ qubit of the other. Formally, if R_{ij} is a controlled- R_θ gate with qubit i controlling qubit j , then R_{ij} and R_{kl} commute unless $j = k$ or $i = l$. We can perform commuting gates simultaneously, as long as we respect the ordering between pairs of this kind. Adding the constraint that each qubit only interact with one other in each layer gives the circuit of Figure 3; it has depth $2n - 2$, linear in n .




Fig. 3. The circuit of figure 2, after the controlled- R_θ gates have been collected into groups that can be performed simultaneously. This gives a circuit with $\mathcal{O}(n)$ layers. Here $n = 6$.

It is easy to show that $2n - 2$ is the minimal depth for this set of gates. We have one gate R_{ij} for every pair $i < j$, and R_{ij} must be performed after R_{jk} . Therefore, two gates R_{ij} and R_{kl} cannot be in the same layer if $i < j < k < l$, since R_{jk} has to precede R_{ij} but follow R_{kl} .

This means that the $n - 1$ gates R_{ij} where $j = i + 1$ must all be in separate layers; since each qubit can only interact with one gate per layer, the $n - 2$ gates R_{ij} where $j = i + 2$ also need their own layers. Adding these to a final layer of H 's gives depth $2n - 2$.

Of course, this does not mean that a different set of gates couldn't solve the QFT more efficiently. It would be especially nice if the QFT could be accomplished by a quantum circuit with depth $\mathcal{O}(\log n)$. This would put it in **QNC**¹, the quantum analog of the class **NC**¹ of problems solvable in logarithmic time by a parallel computer [5]. We would also add the requirement that only a polynomial number of 'ancilla' qubits be used, corresponding to a polynomial number of processors.

How would this be done? Each qubit controls and receives phase shifts on and from $\mathcal{O}(n)$ other qubits. We can easily 'fan out' $\mathcal{O}(n)$ copies of each controlling qubit with a reversible circuit of depth $\mathcal{O}(\log n)$ consisting of controlled-not gates. Classically, we could 'fan in' n phase shifts on a given qubit in depth $\mathcal{O}(\log n)$ by composing them in pairs.

However, it does not seem to be so easy to combine quantum gates in this way. We need some representation of phases so that they can be added in pairs with a linear, unitary operator. Perhaps a clever reader can find such a representation.

In conclusion, we have shown how, in one case, a quantum circuit can be parallelized by re-writing its gates, and lumping them into mutually commuting groups that can be performed simultaneously.

Acknowledgements. I would like to thank the organizers of the First International Conference on Unconventional Models of Computation in Auckland, New Zealand; Seth Lloyd, Tom Knight, David DiVincenzo, and Artur Ekert for helpful conversations; and Fat Mikey and Lady Fox for cuddles.

References

1. A. Barenco, C. Bennett, R. Cleve, D.P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J.A. Smolin and H. Weifurter, "Elementary gates for quantum computation." quant-ph/9503016, *Phys. Rev. A* **52** (1995) 3457-3467.
2. D. Coppersmith, "An approximate Fourier transform useful in quantum factoring." IBM Research Report RC 19642.
3. R.B. Griffiths and C.-S. Niu, "Semiclassical Fourier transform for quantum computation." quant-ph/9511007, *Phys. Rev. Lett.* **76** (1996) 3228-3231.
4. S. Lloyd, "Almost any quantum logic gate is universal." *Phys. Rev. Lett.* **75** (1995) 346-349.
5. C.H. Papadimitriou, *Computational Complexity*. Addison-Wesley, 1994.
6. P.W. Shor, "Algorithms for quantum computation: discrete logarithms and factoring." quant-ph/9508027, in *Proc. 35th Symp. on Foundations of Computer Science* (1994) 124-134.