

Quantum Circuits with Unbounded Fan-out

Peter Høyer^{1,*} and Robert Špalek^{2,**}

¹ Dept. of Comp. Sci., Univ. of Calgary, AB, Canada. hoyer@cpsc.ucalgary.ca

² Centrum voor Wiskunde en Informatica, Amsterdam, The Netherlands.

Robert.Spalek@cwi.nl

Abstract. We demonstrate that the unbounded fan-out gate is very powerful. Constant-depth polynomial-size quantum circuits with bounded fan-in and unbounded fan-out over a fixed basis (denoted by QNC_f^0) can approximate with polynomially small error the following gates: parity, $\text{mod}[q]$, And, Or, majority, $\text{threshold}[t]$, $\text{exact}[q]$, and counting. Classically, we need logarithmic depth even if we can use unbounded fan-in gates. If we allow arbitrary one-qubit gates instead of a fixed basis, then these circuits can also be made exact in log-star depth. Sorting, arithmetical operations, phase estimation, and the quantum Fourier transform can also be approximated in constant depth.

1 Introduction

In this paper, we study the power of shallow quantum circuits. Long quantum computations encounter various problems with decoherence, hence we want to speed them up as much as possible. We can exploit two types of parallelism:

1. Gates on different qubits can be applied at the same time.
2. Commuting gates can be applied on the same qubits at the same time.

The first possibility is straightforward. There are clues that also the second possibility might be physically feasible: ion-trap [3] and bulk-spin NMR [5]. If two quantum gates commute, so do their Hamiltonians and thus we can apply their joint operation by simply performing both evolutions at the same time.

We define an *unbounded fan-out gate* as a sequence of controlled-not gates sharing one control qubit. This gate is universal for all commuting gates: We show that the parallelisation method of [10,6] can apply general commuting gates in parallel using just the fan-out gate and one-qubit gates.

Classically, the main classes computed by polynomial-size, $(\log^k n)$ -depth circuits with unbounded fan-out are:

- NC^k : bounded fan-in gates,
- AC^k : unbounded fan-in gates,
- TC^k : unbounded threshold gates,
- $\text{AC}^k[q]$: unbounded fan-in and $\text{mod}[q]$ gates, and $\text{ACC}^k = \bigcup_q \text{AC}^k[q]$.

* Supported by the Alberta Ingenuity Fund and the Pacific Institute for the Mathematical Sciences.

** Work conducted in part while at Vrije Universiteit, Amsterdam. Partially supported by EU fifth framework project QAIP, IST-1999-11234 and RESQ, IST-2001-37559.

It is known that TC^k is strictly more powerful than ACC^k [11], and that $\text{AC}^k[q] \neq \text{AC}^k[q']$ for powers of distinct primes [14].

The main quantum circuit classes corresponding to the classical classes are QNC^k , QAC^k , QTC^k , and QACC^k . We use subscript ‘f’ to indicate circuits where we allow the fan-out gate (e.g. QNC_f^k). In contrast to the classical case, allowing $\text{mod}[q]$ gates with different moduli always leads to the same quantum classes: $\text{QACC}^k = \text{QAC}^k[q]$ for every q [6]. Furthermore, parity is equivalent to unbounded fan-out, hence $\text{QAC}_f^k = \text{QAC}^k[2] = \text{QACC}^k$.

In this paper, we show that even threshold gates can be approximated with fan-out and single qubit gates in constant depth. This implies that the bounded-error versions of the classes are equal: $\text{B-QNC}_f^k = \text{B-QAC}_f^k = \text{B-QTC}_f^k$.

We first construct a circuit for the exact $[q]$ gate (which outputs 1 if the input is of Hamming weight q , and 0 otherwise) and then use it for all other gates. The exact $[q]$ gate can be approximated in constant depth thanks to the parallelisation method. Furthermore, we show how to achieve exact computation at the cost of log-star depth.

Sorting and several arithmetical problems including addition and multiplication of n integers are computed by constant-depth threshold circuits [13], hence they are in B-QNC_f^0 . By optimising the methods of [4] to use the fan-out gate, we also put quantum phase estimation and the quantum Fourier transform in B-QNC_f^0 . By results of [12,4], polynomial-time bounded-error algorithms with oracle B-QNC_f^0 can factorise numbers and compute discrete logarithms. Thus, if B-QNC_f^0 can be simulated by a BPP machine, then factorisation can be done in polynomial time by bounded-error Turing machines.

2 Quantum Circuits with Unbounded Fan-out

Quantum circuits resemble classical reversible circuits. A quantum circuit is a sequence of quantum gates ordered into *layers*. The gates are consecutively applied in accordance with the order of the layers. Gates in one layer can be applied in parallel. The *depth* of a circuit is the number of layers and the *size* is the number of gates. A circuit can solve problems of a fixed size, so we define *families* of circuits containing one circuit for every input size. We consider only *uniform families*, whose description can be generated by a log-space Turing machine.

A *quantum gate* is a unitary operator applied on some subset of qubits. We usually use gates from a fixed *universal basis* (Hadamard gate, rotation by an irrational multiple of π , and the controlled-not gate) that can approximate any quantum gate with good precision [1]. The qubits are divided into 2 groups: *Input/output* qubits contain the description of the input at the beginning and they are measured in the computational basis at the end. *Ancilla qubits* are initialised to $|0\rangle$ at the beginning and the circuits usually clean them at the end, so that the output qubits are in a pure state and the ancillas could be reused.


Since unitary evolution is reversible, every operation can be undone. Running the computation backward is called *uncomputation* and is often used for cleaning ancilla qubits.

2.1 Definition of Quantum Gates

Quantum circuits cannot use a naive quantum fan-out gate mapping every superposition $|\phi\rangle|0\rangle \dots |0\rangle$ to $|\phi\rangle \dots |\phi\rangle$ due to the no-cloning theorem [16]. Such a gate is not linear, let alone unitary. Instead, our fan-out gate copies only classical bits and the effect on superpositions is determined by linearity. It acts as a controlled-not-not-...-not gate, i.e. it is an unbounded sequence of controlled-not gates sharing one control qubit. Parity is a natural counterpart of fan-out. It is an unbounded sequence of controlled-not gates sharing one target qubit.

Definition 1. *The fan-out gate maps $|x\rangle|y_1\rangle \dots |y_n\rangle \rightarrow |x\rangle|y_1 \oplus x\rangle \dots |y_n \oplus x\rangle$, where $x \oplus y = (x + y) \bmod 2$. The parity gate maps $|x_1\rangle \dots |x_n\rangle|y\rangle \rightarrow |x_1\rangle \dots |x_n\rangle|y \oplus (x_1 \oplus \dots \oplus x_n)\rangle$.*

Example 1. As used in [6], parity and fan-out can simulate each other in constant depth. Recall the Hadamard gate $H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$ and that $H^2 = I$. If a controlled-not gate is preceded and succeeded by Hadamard gates on both qubits, it just turns around. Since parity is a sequence of controlled-not gates, we can turn around all of them in parallel. The circuit is shown in the following figure:



In this paper, we investigate the circuit complexity of among others these gates:

Definition 2. *Let $x = x_1 \dots x_n$ and let $|x|$ denote the Hamming weight of x . The following $(n + 1)$ -qubit gates map $|x\rangle|y\rangle \rightarrow |x\rangle|y \oplus g(x)\rangle$, where $g(x) = 1$ iff*

$$\begin{array}{lll}
 |x| > 0: \text{ Or,} & |x| \geq \frac{n}{2}: \text{ majority,} & |x| = q: \text{ exact}[q], \\
 |x| = n: \text{ And (Toffoli),} & |x| \geq q: \text{ threshold}[q], & |x| \bmod q = 0: \text{ mod}[q].
 \end{array}$$

The counting gate is any gate that maps $|x\rangle|0^m\rangle \rightarrow |x\rangle| |x| \rangle$ for $m = \lceil \log(n+1) \rceil$.

2.2 Quantum Circuit Classes

Definition 3. $\text{QNC}_f(d(n))$ contains operators computed exactly (i.e. without error) by uniform families of quantum circuits with fan-out of depth $O(d(n))$, polynomial-size, and over a fixed basis. $\text{QNC}_f^k = \text{QNC}_f(\log^k n)$. R-QNC_f^k contains operators approximated with one-sided, and B-QNC_f^k with two-sided, polynomially small error.

Remark 1. Every s -qubit quantum gate can be decomposed into a sequence of one-qubit and controlled-not gates of length $O(s^3 4^s)$ [2]. Hence it does not matter whether we allow one-qubit or fixed-size gates in the basis. All our circuits below are over a fixed basis, unless explicitly mentioned otherwise. Some of our circuits need arbitrary one-qubit gates to be exact.

We do not distinguish between a quantum operator computed by a quantum circuit, a classical function induced by that operator by a measurement, and a language decided by that function. All of them are denoted by QNC_f^k .

3 Parallelisation Method

In this section, we describe a general parallelisation method for achieving very shallow circuits. Furthermore, we apply it on the rotation by Hamming weight and the rotation by value and show how to compute them in constant depth.


3.1 General Method

The unbounded fan-out gate is universal for commuting gates in the following sense: Using fan-out, gates can be applied on the same qubits at the same time whenever (1) they commute, and (2) we know the basis in which they all are diagonal, and (3) we can efficiently change into the basis. The method reduces the depth, however it costs more ancilla qubits.


Lemma 1. [8, Theorem 1.3.19] *For every set of pairwise commuting unitary gates, there exists an orthogonal basis in which all the gates are diagonal.*

Theorem 1. [10,6] *Let $\{U_i\}_{i=1}^n$ be pairwise commuting gates on k qubits. Gate U_i is controlled by $|x_i\rangle$. Let T be a gate changing the basis according to Lemma 1. There exists a quantum circuit with fan-out computing $U = \prod_{i=1}^n U_i^{x_i}$ having depth $\max_{i=1}^n \text{depth}(U_i) + 4 \cdot \text{depth}(T) + 2$, size $\sum_{i=1}^n \text{size}(U_i) + (2n+2) \cdot \text{size}(T) + 2$, and using $(n-1)k$ ancillas.*

Proof. Consider a circuit that applies all U_i sequentially. Put $TT^\dagger = I$ between U_i and U_{i+1} . Take $V_i = T^\dagger U_i T$ as new gates. They are diagonal in the computational basis, hence they just impose some phase shifts. The circuit follows:



Multiple phase shifts on entangled states multiply, so can be applied in parallel. We use fan-out gates twice: first to create n entangled copies of target qubits and then to destroy the entanglement. The final circuit with the desired parameters follows:



□


Example 2. As used in [6], it is simple to prove that $\text{mod}[q] \in \text{QNC}_f^0$: Each input qubit controls one increment modulo q on a counter initialised to 0. At the end, we obtain $|x| \text{ mod } q$. The modular increments commute and thus can be parallelised. Since q is fixed, changing the basis and the increment can both be done in constant depth.

3.2 Rotation by Hamming Weight and Value

In this paper, we often use a *rotation by Hamming weight* $R_z(\varphi|x|)$ and a *rotation by value* $R_z(\varphi x)$, where $R_z(\alpha)$ is one-qubit rotation around z -axis by angle α : $R_z(\alpha) = |0\rangle\langle 0| + e^{i\alpha}|1\rangle\langle 1|$. They both can be computed in constant depth.

First of all, it is convenient to use controlled one-qubit gates as basic elements. The following lemma shows that they can be simulated by one-qubit and controlled-not gates.

Lemma 2. [2, Lemma 5.1] *For every one-qubit gate U , there exist one-qubit gates A, B, C and rotation $P = R_z(\alpha)$ such that the controlled gate U is computed by the following constant-depth circuit:*




Remark 2. If a qubit controls more one-qubit gates, then we can still use this method. The controlled-not gate is just replaced by the fan-out gate and the rotations P are multiplied.

Lemma 3. *For every angle φ , there exist constant-depth, linear-size quantum circuits with fan-out computing $R_z(\varphi|x|)$ and $R_z(\varphi x)$ on input $x = x_{n-1} \dots x_1 x_0$.*

Proof. The left figure shows how to compute the rotation by Hamming weight: Each input qubit controls $R_z(\varphi)$ on the target qubit, hence the total angle is $\varphi|x|$. These controlled rotations are parallelised using the parallelisation method.

The right figure shows the rotation by value. It is similar to the rotation by Hamming weight, only the input qubit $|x_j\rangle$ controls $R_z(\varphi 2^j)$, hence the total angle is $\varphi \sum_{j=0}^{n-1} 2^j x_j = \varphi x$.



□

Remark 3. The construction uses rotations $R_z(\varphi)$ for arbitrary $\varphi \in \mathbb{R}$. However, we are only allowed to use a fixed set of one-qubit gates. It is easy to see that every rotation can be approximated with polynomially small error by $R_z(\theta q) = (R_z(\theta))^q$, where $\sin \theta = \frac{3}{5}$ and q is a polynomially large integer [1]. These q rotations commute, so can be applied in parallel and the depth is preserved.

4 Approximate Circuits

In this section, we present very shallow approximate circuits for all gates from Definition 2.

4.1 Quantum Fourier Transform

QFT is a very powerful tool used in several quantum algorithms, e.g. factorisation of integers [12]. In this section, we use it for a different purpose: parallelisation of increment gates.

Definition 4. *The quantum Fourier transform (QFT) performs the Fourier transform on the quantum amplitudes of the state, i.e. it maps*

$$F_n : |x\rangle \rightarrow |\psi_n^x\rangle = \frac{1}{2^{n/2}} \sum_{y=0}^{2^n-1} e^{2\pi i xy/2^n} |y\rangle. \tag{1}$$

Shor has shown in [12] how to compute QFT in quadratic depth, quadratic size, and without ancillas. The depth has further been improved to linear. Cleve and Watrous have shown in [4] that QFT can be approximated with error ε in depth $O(\log n + \log \log \frac{1}{\varepsilon})$ and size $O(n \log \frac{n}{\varepsilon})$. Furthermore, they have shown that logarithmic depth is necessary (in the model without fan-out).

4.2 Circuits of Double-Logarithmic Depth

Circuits in this sub-section are not optimal, however they give a good insight. They are based on counting the weight of the input, which is parallelised.

Definition 5. *The increment gate maps $\text{Incr}_n : |x\rangle \rightarrow |(x + 1) \bmod 2^n\rangle$.*

Lemma 4. *The increment gate is diagonal in the Fourier basis and its diagonal version is in QNC^0 .*

Proof. It is simple to prove the following equations:

1. $\text{Incr}_n = F_n^\dagger D_n F_n$ for diagonal $D_n = \sum_{x=0}^{2^n-1} e^{2\pi i x/2^n} |x\rangle\langle x|$,
2. $D_n = R_z(\pi) \otimes R_z(\pi/2) \otimes \dots \otimes R_z(\pi/2^{n-1})$. □

Remark 4. Classically, $\text{Incr} \in \text{NC}^1$. The circuits for QFT mentioned above also have logarithmic depth and they are only approximate. However, quantum circuits of this type can be parallelised, which we cannot do with classical circuits.

Furthermore, the addition of a fixed integer q is as hard as the increment: by Lemma 4, $\text{Incr}^q = F^\dagger D^q F$ and $(R_z(\varphi))^q = R_z(\varphi q)$, hence the diagonal version of the addition of q is also in QNC^0 .

Theorem 2. *Using fan-out, the counting gate can be approximated with error ε in depth $O(\log \log n + \log \log \frac{1}{\varepsilon})$ and size $O((n + \log \frac{1}{\varepsilon}) \log n)$.*

Proof. Compute the Hamming weight of the input: Each input qubit controls one increment on an m -qubit counter initialised to 0, where $m = \lceil \log(n + 1) \rceil$. The increments Incr_m are parallelised, so we apply the quantum Fourier transform F_m twice and the n constant-depth controlled D_m gates in parallel. □

Remark 5. Other gates are computed from the counting gate by standard methods: $\text{threshold}[t]$ can be computed as the most significant qubit of the counter if we align it to a power of 2 by adding fixed integer $2^m - t$.

4.3 Constant-Depth Circuits

Rotations by Hamming weight computed for many elementary angles in parallel can be used for approximating the Or and $\text{exact}[q]$ gates in constant depth.

Define one-qubit state $|\mu_\varphi^w\rangle = (H \cdot R_z(\varphi w) \cdot H) |0\rangle = \frac{1+e^{i\varphi w}}{2} |0\rangle + \frac{1-e^{i\varphi w}}{2} |1\rangle$. By Lemma 3, $|\mu_\varphi^{|x|}\rangle$ can be computed in constant-depth and linear-size.

Theorem 3. *$\text{Or} \in \text{R-QNC}_f^0$, i.e. it can be approximated with one-sided polynomially small error in constant-depth.*

Proof. Let $m = a \cdot n$, where a will be chosen later. For all $k \in \{0, 1, \dots, m - 1\}$, compute in parallel $|y_k\rangle = |\mu_{\varphi_k}^{|x|}\rangle$ for angle $\varphi_k = \frac{2\pi}{m}k$. If $|y_k\rangle$ is measured in the computational basis, the expected value is

$$E[Y_k] = \left| \frac{1 - e^{i\varphi_k|x|}}{2} \right|^2 = \left| e^{-i\varphi_k|x|} \cdot \frac{e^{i\varphi_k|x|} + e^{-i\varphi_k|x|} - 2}{4} \right| = \frac{1 - \cos(\varphi_k|x|)}{2}.$$

If all these m qubits $|y\rangle$ are measured, the expected Hamming weight is

$$E[|Y|] = E\left[\sum_{k=0}^{m-1} Y_k\right] = \frac{m}{2} - \frac{1}{2} \sum_{k=0}^{m-1} \cos\left(\frac{2\pi k}{m}|x|\right) = \begin{cases} 0 & \text{if } |x| = 0, \\ \frac{m}{2} & \text{if } |x| \neq 0. \end{cases}$$

The qubits $|y\rangle$ are actually not measured, but their Hamming weight $|y|$ controls another rotation on a new ancilla qubit $|z\rangle$. So compute $|z\rangle = |\mu_{\frac{2\pi}{m}|y}\rangle$.


Let Z be the outcome after $|z\rangle$ is measured. If $|y| = 0$, then $Z = 0$ with certainty. If $||y| - \frac{m}{2}| \leq \frac{m}{\sqrt{n}}$, then

$$P[Z = 0] = \left| \frac{1 + e^{i\frac{2\pi}{m}|y|}}{2} \right|^2 = \frac{1 + \cos\left(\frac{2\pi}{m}|y|\right)}{2} \leq \frac{1 - \cos\frac{2\pi}{\sqrt{n}}}{2} = O\left(\frac{1}{n}\right).$$

Assume that $|x| \neq 0$. Since $0 \leq Y_k \leq 1$, we can use Hoeffding’s Lemma 5 below and obtain $P\left[||Y| - \frac{m}{2}| \geq \varepsilon m\right] \leq \frac{1}{2\varepsilon^2 m}$. Fix $a = \log n$ and $\varepsilon = \frac{1}{\sqrt{n}}$. Now,

$$P\left[||y| - \frac{m}{2}| \geq \frac{m}{\sqrt{n}}\right] \leq \frac{1}{2^{a/n}} = \frac{1}{2^a} = \frac{1}{n}. \text{ Hence } P[Z = 0] = \begin{cases} 1 & \text{if } |x| = 0, \\ O\left(\frac{1}{n}\right) & \text{if } |x| \neq 0. \end{cases}$$

The circuit has constant depth and size $O(mn) = O(n^2 \log n)$. It is outlined in the following figure. The figure is slightly simplified: unimportant qubits and uncomputation of ancillas are omitted.



□

Lemma 5 (Hoeffding). [7] *If Y_1, \dots, Y_m are independent random variables bounded by $a_k \leq Y_k \leq b_k$, then, for all $\varepsilon > 0$,*

$$P[|S - E[S]| \geq \varepsilon m] \leq 2 \exp \frac{-2m^2\varepsilon^2}{\sum_{k=1}^m (b_k - a_k)^2}, \quad \text{where } S = \sum_{i=1}^m Y_i.$$

Remark 6. Since the outcome is a classical bit, we can save it and clean all ancillas by uncomputation. It remains to prove that the intermediate qubits $|y\rangle$ need not be measured, in order to be able to uncompute them.

We have only proved that the output qubit is a good approximation of the logical Or, if $|y\rangle$ is immediately measured. By the principle of deferred measurement, we can use controlled quantum operations and measure $|y\rangle$ at the end. However, the outcome is a classical bit hardly entangled with $|y\rangle$, hence it does not matter whether $|y\rangle$ is measured.


Remark 7. If we need smaller error $\frac{1}{n^c}$, we create c copies and compute exact Or of them by a binary tree of Or gates. The tree has depth $\log c = O(1)$. Using Theorem 6, the size can be reduced to $O(dn \log^{(d)} n)$ and the depth is $O(d)$.

Theorem 4. $\text{exact}[q] \in \text{R-QNC}_f^0$.

Proof. Slight modification of the circuit for Or: As outlined in the figure, by adding rotation $R_z(-\varphi q)$ to the rotation by Hamming weight in the first layer, we obtain $|\mu_\varphi^{|x|-q}\rangle$ instead of $|\mu_\varphi^{|x}|\rangle$. The second layer stays the same. If the output qubit $|z\rangle$ is measured, then

$$P[Z = 0] = \begin{cases} 1 & \text{if } |x| = q, \\ O(\frac{1}{n}) & \text{if } |x| \neq q. \end{cases}$$

We obtain an approximation of the exact[q] gate with one-sided polynomially small error. □



Remark 8. Other gates are computed from the exact[q] gate by standard methods: threshold[t] can be computed as the parity of exact[t], exact[t + 1], ..., exact[n]. The depth stays constant and the size is just n -times bigger, i.e. $O(n^3 \log n)$, hence threshold[t] $\in \text{B-QNC}_f^0$.

Using Theorem 7 and the technique of Theorem 2, threshold[t] can be computed in constant-depth and smaller size $O(n \log n)$.

4.4 Arithmetical Operations

The threshold gate is very powerful, so the fan-out gate is powerful too:

Theorem 5. *The following functions are in B-QNC_f^0 : addition and multiplication of n integers, division of two integers, and sorting of n integers.*

Proof. By [13], these functions are computed by constant-depth, polynomial-size threshold circuits. The depths are really small, from 2 to 5. A threshold circuit is built of weighted threshold gates. It is simple to prove that also the weighted threshold gate (with polynomially large integer weights) is in B-QNC_f^0 . □

5 Exact Circuits

In the previous section, we have shown how to approximate the exact[q] gate in constant depth. In this section, we show how to compute it exactly in log-star depth. The circuits in this section need arbitrary one-qubit gates instead of a fixed basis, otherwise they would not be exact.

Theorem 6. *Or on n qubits can be reduced exactly to Or on $m = \lceil \log(n + 1) \rceil$ qubits in constant-depth and size $O(n \log n)$.*

Proof. For $k \in \{1, 2, \dots, m\}$, compute in parallel $|y_k\rangle = |\mu_{\varphi_k}^{|x|}\rangle$ for angle $\varphi_k = \frac{2\pi}{2^k}$:

- If $|x| = 0$, then $|y_k\rangle = |0\rangle$ for each k .
- If $|x| \neq 0$, take unique decomposition $x = 2^a(2b + 1)$ where $a, b \in \mathbb{N}_0$. Then

$$\langle 1|y_{a+1}\rangle = \frac{1 - e^{i\varphi_{a+1}|x|}}{2} = \frac{1 - e^{i\pi(2b+1)}}{2} = \frac{1 - e^{i\pi}}{2} = 1.$$

It follows that $|x| = 0 \iff |y| = 0$. Hence the original problem is exactly reduced to a problem of logarithmic size. \square

Remark 9. If all input qubits are zero, then also all output qubits are zero. Otherwise the output qubits are in a general superposition such that the amplitude of the zero state is 0.

Corollary 1. $exact[q] \in \text{QAC}_f^0$.

Proof. Using the same method as in Theorem 4, also the $exact[q]$ gate can be reduced to Or, which is in QAC_f^0 . \square

Corollary 2. $exact[q] \in \text{QNC}_f(\log^* n)$.

Proof. Repeat the exact reduction $\log^* n$ times until the input size ≤ 2 . Compute and save the outcome and clean ancillas by uncomputation. \square

6 Quantum Fourier Transform and Phase Estimation

6.1 Constant-Depth QFT

We show that the approximate circuit for QFT from [4] can be compressed to constant depth, if we use the fan-out gate. Recall Definition 4 of QFT.

Theorem 7. $QFT \in \text{B-QNC}_f^0$.

Proof. The operator $F_n : |x\rangle \rightarrow |\psi_n^x\rangle$ can be computed by composing:

1. Fourier state construction (QFS): $|x\rangle|0\rangle \dots |0\rangle \rightarrow |x\rangle|\psi_n^x\rangle|0\rangle \dots |0\rangle$
2. Copying Fourier state: $|x\rangle|\psi_n^x\rangle|0\rangle \dots |0\rangle \rightarrow |x\rangle|\psi_n^x\rangle \dots |\psi_n^x\rangle$
3. Uncomputing phase estimation (QFP): $|\psi_n^x\rangle \dots |\psi_n^x\rangle|x\rangle \rightarrow |\psi_n^x\rangle \dots |\psi_n^x\rangle|0\rangle$
4. Uncopying Fourier state: $|\psi_n^x\rangle \dots |\psi_n^x\rangle|0\rangle \rightarrow |\psi_n^x\rangle|0\rangle \dots |0\rangle$

The following lemmas show that each of these four operators is in B-QNC_f^0 . \square

Lemma 6. $QFS \in \text{QNC}_f^0$.

Proof. QFS maps $|x\rangle|0\rangle \rightarrow |x\rangle|\psi_n^x\rangle$. Define $|\rho_r\rangle = \frac{|0\rangle + e^{2\pi r i}|1\rangle}{\sqrt{2}}$. It is simple to prove that $|\psi_n^x\rangle = |\rho_{x/2^1}\rangle|\rho_{x/2^2}\rangle \dots |\rho_{x/2^n}\rangle$. The n qubits $|\rho_{x/2^k}\rangle$ can be computed from x in parallel. Computation of $|\rho_{x/2^k}\rangle = R_z\left(\frac{2\pi}{2^k}x\right)\frac{|0\rangle + |1\rangle}{\sqrt{2}}$ is done by the rotation by value (Lemma 3) in constant depth and linear size. \square

Definition 6. Let x_1, \dots, x_m be n -bit integers. The reversible addition-gate maps $\text{add}_n^m : |x_1\rangle \dots |x_m\rangle \rightarrow |x_1\rangle \dots |x_{m-1}\rangle|y\rangle$, where $y = (\sum_{i=1}^m x_i) \bmod 2^n$.

Lemma 7. $\text{add}_n^m \in \text{B-QNC}_f^0$.

Proof. By Theorem 5, $y = (\sum_{i=1}^m x_i) \bmod 2^n$ can be approximated in constant depth and polynomial size. The result is, however, stored into ancilla qubits. Uncompute $x_m = (y - \sum_{i=1}^{m-1} x_i) \bmod 2^n$ in the same way (subtraction is as hard as addition). \square

Lemma 8. Copying Fourier state is in B-QNC_f^0 .

Proof. Take the reversible addition-gate: $(\text{add}_n^2)|y\rangle|x\rangle = |y\rangle|(x+y) \bmod 2^n\rangle$. It is simple to prove that $(\text{add}_n^2)^{-1}|\psi_n^y\rangle|\psi_n^x\rangle = |\psi_n^{x+y}\rangle|\psi_n^x\rangle$. Hence $(\text{add}_n^2)^{-1}|\psi_n^0\rangle|\psi_n^x\rangle = |\psi_n^x\rangle|\psi_n^x\rangle$. The state $|\psi_n^0\rangle = H^{\otimes n}|0^n\rangle$ is easy to prepare in constant depth.

By the same arguments, $(\text{add}_n^m)^{-1}|\psi_n^0\rangle \dots |\psi_n^0\rangle|\psi_n^x\rangle = |\psi_n^x\rangle \dots |\psi_n^x\rangle|\psi_n^x\rangle$. \square

Lemma 9. $QFP \in \text{B-QNC}_f^0$.

Proof. QFP maps $|\psi_n^x\rangle \dots |\psi_n^x\rangle|0\rangle \rightarrow |\psi_n^x\rangle \dots |\psi_n^x\rangle|x\rangle$. By Cleve and Watrous [4, sub-section 3.3], we can compute x with probability $\geq 1 - \varepsilon$ from $O(\log \frac{n}{\varepsilon})$ copies of $|\psi_n^x\rangle$ in depth $O(\log n + \log \log \frac{1}{\varepsilon})$ and size $O(n \log \frac{n}{\varepsilon})$. Use $\varepsilon = \frac{1}{\text{poly}(n)}$. It is easy to see that their circuit can have constant depth, if we use fan-out, parity, And, Or, majority gate. All these gates are in B-QNC_f^0 . \square

6.2 Quantum Phase Estimation

The method of computing QFT can be also used for phase estimation:

Theorem 8. Given a gate $S_x : |y\rangle|\phi\rangle \rightarrow |y\rangle R_z\left(\frac{2\pi x}{2^n}y\right)|\phi\rangle$ for basis states $|y\rangle$, where $x \in \mathbb{Z}_{2^n}$ is unknown, we can determine x with probability $\geq 1 - \varepsilon$ in constant depth, size $O(n \log \frac{n}{\varepsilon})$, and using the S_x gate $O(n \log \frac{n}{\varepsilon})$ times.

Proof. Compute $|\rho_{x/2^k}\rangle = R_z\left(\frac{2\pi x}{2^k}\right)\frac{|0\rangle + |1\rangle}{\sqrt{2}} = R_z\left(\frac{2\pi x}{2^n}2^{n-k}\right)\frac{|0\rangle + |1\rangle}{\sqrt{2}}$, which is the result of one application of $S_x\left(|2^{n-k}\rangle\frac{|0\rangle + |1\rangle}{\sqrt{2}}\right)$. Apply QFP on $O(\log \frac{n}{\varepsilon})$ copies of $|\psi_n^x\rangle = |\rho_{x/2^1}\rangle|\rho_{x/2^2}\rangle \dots |\rho_{x/2^n}\rangle$. \square

7 Concluding Remarks

7.1 Relations of Quantum Circuit Classes

We have shown that $\text{B-QNC}_f^0 = \text{B-QAC}_f^0 = \text{B-QACC}^0 = \text{B-QTC}_f^0$ (Theorem 4). If we allow arbitrary one-qubit gates, then also $\text{QAC}_f^0 = \text{QTC}_f^0 \subseteq \text{QNC}_f(\log^* n)$ (Corollaries 1 and 2). Several open problems of [6] have thus been solved.

Only little is known about classes that do not include the fan-out gate. For example, we do not know whether $\text{TC}^0 \subseteq \text{QTC}^0$, we only know that $\text{TC}^0 \subseteq \text{QTC}_f^0$. It is simple to prove that parity is in TC^0 : take Or of $\text{exact}[1]$, $\text{exact}[3]$, $\text{exact}[5]$, \dots , and compute $\text{exact}[k]$ from $\text{threshold}[k]$ and $\text{threshold}[k+1]$. However, this method needs fan-out to copy the input bits.

7.2 Randomised versus Quantum Depth

We compare depths of randomised classical circuits and quantum circuits, both with bounded fan-in and unbounded parity and fan-out. Quantum upper bounds are proved in this paper. Classical lower bounds can be proved by Yao's principle and the polynomial method (with polynomials modulo 2).

Gate	Randomised	Quantum
Or and $\text{threshold}[t]$ exactly	$\Theta(\log n)$	$O(\log^* n)$
$\text{mod}[q]$ exactly	$\Theta(\log n)$	$\Theta(1)$
Or with error $\frac{1}{n}$	$\Theta(\log \log n)$	$\Theta(1)$
$\text{threshold}[t]$ with error $\frac{1}{n}$	$\Omega(\log \log n)$	$\Theta(1)$

7.3 Upper Bounds for B-QNC_f^0

Shor's original factoring algorithm uses modular exponentiation and the quantum Fourier transform followed by a polynomial-time deterministic algorithm. The modular exponentiation a^x can be replaced by multiplication of some subset of numbers $a, a^2, a^4, \dots, a^{2^{n-1}}$ [4]. Numbers a^{2^k} are precomputed classically.

Since both multiplication of n numbers (Theorem 5) and QFT (Theorem 7) are in B-QNC_f^0 , there is a polynomial-time bounded-error algorithm with oracle B-QNC_f^0 factoring numbers, i.e. factoring $\in \text{RP}[\text{B-QNC}_f^0]$. If $\text{B-QNC}_f^0 \subseteq \text{BPP}$, then factoring $\in \text{RP}[\text{BPP}] \subseteq \text{BPP}[\text{BPP}] = \text{BPP}$.

Acknowledgements. We would like to thank Harry Buhrman, Hartmut Klauck, and Hein Röhrig at CWI in Amsterdam, and Fred Green at Clark University in Worcester for plenty helpful discussions, and Ronald de Wolf at CWI for help with writing the paper. We are grateful to Schloss Dagstuhl, Germany, for providing an excellent environment, where part of this work was carried out.

References

1. L. M. Adleman, J. DeMarrais, and M. A. Huang. Quantum computability. *SIAM Journal on Computing*, 26(5):1524–1540, 1997.
2. A. Barenco, C. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, and H. Weinfurter. Elementary gates for quantum computation. *Physical Review A*, 52:3457–3467, 1995. quant-ph/9503016.
3. J. I. Cirac and P. Zoller. Quantum computations with cold trapped ions. *Phys. Rev. Lett.*, 74:4091–4094, 1995.
4. R. Cleve and J. Watrous. Fast parallel circuits for the quantum Fourier transform. In *Proc. of the 41st IEEE Symp. on Foundations of Computer Science*, pages 526–536, 2000.
5. N. Gershenfeld and I. Chuang. Bulk spin resonance quantum computation. *Science*, 275:350–356, 1997. <http://citeseer.nj.nec.com/gershenfeld97bulk.html>.
6. F. Green, S. Homer, C. Moore, and C. Pollett. Counting, fanout, and the complexity of quantum ACC. *Quantum Information and Computation*, 2(1):35–65, 2002. quant-ph/0106017.
7. W. Hoeffding. Probability inequalities for sums of bounded random variables. *J. Amer. Statist. Assoc.*, 58:13–30, 1963.
8. R. A. Horn and C. R. Johnson. *Matrix Analysis*. Cambridge University Press, 1985.
9. C. Moore. Quantum circuits: Fanout, parity, and counting. quant-ph/9903046, 1999.
10. C. Moore and M. Nilsson. Parallel quantum computation and quantum codes. *SIAM Journal on Computing*, 31(3):799–815, 2002. quant-ph/9808027.
11. A. A. Razborov. Lower bounds for the size of circuits of bounded depth with basis $\{\&, \oplus\}$. *Math. Notes Acad. Sci. USSR*, 41(4):333–338, 1987.
12. P. W. Shor. Algorithms for quantum computation: discrete logarithms and factoring. In *Proc. of the 35th Annual Symp. on FOCS*, pages 124–134, Los Alamitos, CA, 1994. IEEE Press. <http://citeseer.nj.nec.com/14533.html>.
13. K.-Y. Siu, J. Bruck, T. Kailath, and T. Hofmeister. Depth efficient neural networks for division and related problems. *IEEE Transactions on Information Theory*, 39(3):946–956, 1993.
14. R. Smolensky. Algebraic methods in the theory of lower bounds for Boolean circuit complexity. In *Proc. 19th Annual ACM Symposium on Theory of Computing*, pages 77–82, 1987.
15. R. Špalek. Quantum circuits with unbounded fan-out. Master’s thesis, Faculty of Sciences, Vrije Universiteit, Amsterdam, 2002. <http://www.ucw.cz/~robert/qncwf/>. Shorter version and improved results in quant-ph/0208043.
16. W. K. Wootters and W. H. Zurek. A single quantum cannot be clone. *Nature*, 299:802–803, 1982.