

Bounds on the Power of Constant-Depth Quantum Circuits

November 25, 2003

Abstract

We show that if a language is recognized within certain error bounds by constant-depth quantum circuits over a finite family of gates, then it is computable in (classical) polynomial time. In particular, our results imply

$$\mathbf{EQNC}^0 \subseteq \mathbf{P},$$

where \mathbf{EQNC}^0 is the constant-depth analogue of the class \mathbf{EQP} .

On the other hand, we adapt and extend ideas of DiVincenzo & Terhal [?] to show that, for any family \mathcal{F} of quantum gates including Hadamard and CNOT gates, computing the acceptance probabilities of depth-five circuits over \mathcal{F} is just as hard as computing these probabilities for arbitrary quantum circuits over \mathcal{F} . In particular, this implies that

$$\mathbf{NQNC}^0 = \mathbf{NQACC} = \mathbf{NQP} = \mathbf{coC=P},$$

where \mathbf{NQNC}^0 is the constant-depth analogue of the class \mathbf{NQP} . This essentially refutes a conjecture of Green et al. that $\mathbf{NQACC} \subseteq \mathbf{TC}^0$ [?].

1 Introduction

Quantum decoherence is a major obstacle to maintaining long quantum computations, hence people are interested in the implementation and power of shallow quantum circuits. Many methods were developed to reduce the depth of the quantum circuits. Using unbounded fan-out gates, Hoyer and Spalek [] managed to apply a sequence of commuting gates on the same qubits at the same time, and thus greatly reduced the depth of the circuits under various circumstances. As for the power of shallow quantum circuits, several hardness results of simulating constant-depth quantum circuits were given by Terhal and Divincenzo [?]. They showed that if one can classically efficiently simulate quantum circuits of depth at least four using two-qubit gates then $\mathbf{BQP} \subseteq \mathbf{AM}$. They also showed that the polynomial hierarchy collapses if there exists an efficient counting simulation exists for the above circuits. In the same line of research, Green et al. [?] defined quantum circuit classes \mathbf{QNC}^k , \mathbf{QAC}^k and \mathbf{QACC}^k . While they proved a number of results about \mathbf{QAC}^k and \mathbf{QACC}^k , our paper focuses on the class \mathbf{QNC}^k , which is a proper subclass of \mathbf{QAC}^k by allowing only gates with

bounded fan-in. As \mathbf{QNC}^k being a pretty small circuit class, it seems reasonable that the first set of non-trivial quantum circuits successfully built will be of class \mathbf{QNC}^0 .

In this paper we study the power of constant-depth circuits by investigating several language classes about \mathbf{QNC}^k , especially \mathbf{QNC}^0 . Except minor modification which we will mention later, we use basically the same definitions for the language classes as used in Green et al. . We show that some classes such as \mathbf{EQNC}^0 are small classes inside \mathbf{P} , while other language classes are suprisingly large. In particular, we show that $\mathbf{NQNC}^0 = \mathbf{NQACC} = \mathbf{NQP} = \text{coC=P}$.

2 Preliminaries

2.1 Gates and circuits


We first review some standard quantum (unitary) gates. Among the single-qubit gates, we have the Pauli gates X , Y , and Z , the Hadamard gate H , and the $\pi/8$ gate T , which are defined thus, for $b \in \{0, 1\}$:

$$\begin{aligned} X|b\rangle &= |¬b\rangle, \\ Y|b\rangle &= i(-1)^b|¬b\rangle, \\ Z|b\rangle &= (-1)^b|b\rangle, \\ H|b\rangle &= (|0\rangle + (-1)^b|1\rangle)/\sqrt{2}, \\ T|b\rangle &= e^{i\pi b/4}|b\rangle. \end{aligned}$$

For $n \geq 1$, the $(n + 1)$ -qubit *generalized Toffoli gate* T_n satisfies

$$T_n|x_1, \dots, x_n, b\rangle = |x_1, \dots, x_n, b \oplus \bigwedge_{i=1}^n x_i\rangle.$$

Here b is the *target qubit* and x_1, \dots, x_n are the *control qubits*. T_n is the quantum analogue of the Boolean AND-gate with fanin n . T_2 is known simply as the Toffoli gate. T_1 is also known as the controlled NOT (CNOT) gate and is depicted below. Here, $a, b \in \{0, 1\}$.



For $q > 1$, the $(n + 1)$ -qubit Mod_q -gate acts on a basis state $|x_1, \dots, x_n, b\rangle$ by flipping the target qubit b iff $x_1 + \dots + x_n \not\equiv 0 \pmod{q}$. The control qubits x_1, \dots, x_n are left alone. The Mod_2 gate is also known as the parity gate.

Our notion of quantum circuit is fairly standard (see, for example, [?]): a series of quantum gates, drawn from some specified set of unitary operators, acting on some specified number of qubits, labeled $1, \dots, q$. The first few qubits are considered *input* qubits, which are assumed to be in some basis state initially (i.e., classical input); the rest are ancillæ, each assumed to be in the $|0\rangle$ state initially. Thus the initial state of the qubits is $|x, 00 \dots 0\rangle$,

for some binary string x . Some arbitrary set of qubits are specified as *output* qubits, and these qubits are assumed to be measured in the computational basis at the final state. We assume that the sets of input and output qubits are part of the description of the circuit. The circuit *accepts* its input if all the output qubits are observed to be 0 in the final state. Otherwise the circuit rejects. We let $\Pr[C(x)]$ denote the probability that C accepts input x .

If C is any quantum circuit, it will be convenient for us to define $|C|$, the *size* of C , to be the number of output qubits plus the number of “contact points” between qubits and gates, so for example, a single-qubit gate counts one towards the size, while a two-qubit gate counts two, etc. C may be laid out by partitioning its gates into *layers* $1, \dots, d$, such that (i) gates in the same layer all act on pairwise disjoint sets of qubits, and (ii) all gates in layer i are applied before any gates in layer $i + 1$, for $1 \leq i < d$. The *depth* of C is then the smallest possible value of d . The *width* of C is the number of qubits in C .

2.2 Complexity classes using QNC circuits

The circuit class **QNC** was first suggested by Moore and Nilsson [] as the quantum analogue of the class **NC** of bounded fan-in Boolean circuits with polylogarithmic depth and polynomial size. We define the class **QNC^k** in the same fashion as definitions in Green, Homer, Moore, and Pollett [] with some minor modifications.

Definition 2.1 **QNC^k** is the class of quantum circuit families $\{C_n\}_{n \geq 0}$ for which there exists a polynomial p such that each C_n contains n input qubits and at most $p(n)$ many ancillæ. Each C_n has depth $O(\log^k n)$ and uses only single-qubit gates and CNOT gates. The single-qubit gates must be from a fixed finite set. We say the circuit family $\{C_n\}$ is uniform if there is a (classical) polynomial-time algorithm that outputs a description of C_n on input 0^n .

Next we define the language classes **NQNC^k** and **EQNC^k**. These are **QNC^k** analogues of the classes **NQP** and **EQP**, respectively.

Definition 2.2 (cf. []) Let $k \geq 0$ be an integer.

- **NQNC^k** is the class of languages L such that there is a uniform $\{C_n\} \in \mathbf{QNC}^k$ such that, for all x ,

$$x \in L \iff \Pr[C_{|x|}(x)] > 0.$$

- **EQNC^k** is the class of languages L such that there is a uniform $\{C_n\} \in \mathbf{QNC}^k$ such that, for all x , $\Pr[C_{|x|}(x)] \in \{0, 1\}$ and

$$x \in L \iff \Pr[C_{|x|}(x)] = 1.$$

Remark. Green, Homer, Moore and Pollett implicitly consider the output qubits of C_n to be all the qubits in C_n []. In our model we allow any subset of qubits to be the output qubits of C_n , and we do not restrict our circuits to be clean, i.e., the non-output qubits could end up in an arbitrary state, possibly entangled with the output qubits. The reason we define our circuits this way is based on the observation that, in their model, if a language L is in \mathbf{EQNC}^k (or $\mathbf{BQNC}_{\epsilon,\delta}^k$ for large enough δ), then L can contain no more than one string of each length.

Bounded-error \mathbf{QAC}^k classes were mentioned in [?], and one can certainly ask about similar classes for \mathbf{QNC} circuits. It is not obvious that there is one robust definition of \mathbf{BQNC}^0 —perhaps because it is not clear how to reduce error significantly by amplification in constant depth.¹ In the next definition, we will try to be as general as possible while still maintaining our assumption that $\vec{0}$ is the only accepting output.

Definition 2.3 *Let ϵ and δ be functions mapping (descriptions of) quantum circuits into real numbers such that, for all quantum circuits C , $0 < \epsilon(C) \leq \delta(C) \leq 1$. We write ϵ_C and δ_C to denote $\epsilon(C)$ and $\delta(C)$, respectively. $\mathbf{BQNC}_{\epsilon,\delta}^k$ is the class of languages L such that there is a uniform $\{C_n\} \in \mathbf{QNC}^k$ such that for any string x of length n ,*

$$\begin{aligned} x \in L &\implies \Pr[C_n(x)] \geq \delta_{C_n}, \\ x \notin L &\implies \Pr[C_n(x)] < \epsilon_{C_n}. \end{aligned}$$

An interesting special case is when $\epsilon_C = \delta_C = 1$, that is, the input is accepted iff the circuit accepts with probability 1, and there is no promise on the acceptance probability. One might expect that, by the symmetry of the definitions, this class $\mathbf{BQNC}_{1,1}^0$ is the same as \mathbf{NQNC}^0 , but it is almost certainly not, as we will see.

2.3 Other classes of constant-depth quantum circuits

Definition 2.4 *Let $k \geq 0$ and $q > 1$ be integers.*

- \mathbf{QAC}^k is the same as \mathbf{QNC}^k except that generalized Toffoli gates are allowed in the circuits.
- $\mathbf{QACC}(q)$ is the same as \mathbf{QNC}^0 except that generalized Mod_q gates are allowed in the circuits.
- $\mathbf{QACC} = \bigcup_{q>1} \mathbf{QACC}(q)$.

¹One can always reduce error classically by just running the circuit several times on the same input. In this case, the best definition of \mathbf{BQNC}^0 may be that the gap between the allowed accept and reject probabilities should be at least $1/\text{poly}$.

3 Main results

3.1 Simulating QNC⁰ circuits exactly is hard

Theorem 3.1 $\text{NQNC}^0 = \text{NQP} = \text{C}_{\neq} \text{P}$.


As a corollary, we essentially solve an open problem of Green et al. [?]. They conjectured that $\text{NQACC} \subseteq \text{TC}^0$.

Corollary 3.2 For any $k \geq 0$,


$$\text{NQNC}^0 = \text{NQNC}^k = \text{NQAC}^k = \text{NQACC} = \text{C}_{\neq} \text{P}.$$

Thus, $\text{NQACC} \not\subseteq \text{TC}^0$ unless $\text{C}_{\neq} \text{P} = \text{TC}^0$.

Let B be the two-qubit Bell gate, defined as



Also let



which produces the EPR state $(|00\rangle + |11\rangle)/\sqrt{2}$. We prove the following lemma, from which the theorem follows quickly.

Lemma 3.3 For any quantum circuit \mathcal{C} using gates drawn from any family \mathcal{F} , there is a depth-three quantum circuit \mathcal{C}' of size linear in $|\mathcal{C}|$ using gates drawn from $\mathcal{F} \cup \{B, B^\dagger\}$ such that for any input x of the appropriate length,

$$\Pr[\mathcal{C}'(x)] = 2^{-m} \Pr[\mathcal{C}(x)],$$

for some $m \leq 2|\mathcal{C}|$ depending only on \mathcal{C} . The middle layer of \mathcal{C}' contains each gate in \mathcal{C} exactly once and no others. The third layer contains only B^\dagger -gates, and the first layer contains only B -gates, which are used only to create EPR states.

Proof. Our construction is a simplified version of the main construction in DiVincenzo and Terhal [?], but ours is stronger in one crucial respect discussed below: it does not significantly increase the family of gates used. To construct \mathcal{C}' , we start with \mathcal{C} and simply insert, for each qubit q of \mathcal{C} , a simplified teleportation module (shown in Figure 1) between any two consecutive quantum gates of \mathcal{C} acting on q . No further gates involve the qubits r_1 and r_2 to the right of the B^\dagger -gate. This module, which lacks the usual corrective Pauli gates,




Figure 1: The nonadaptive teleportation module [?]. The state in qubit q is teleported correctly iff the qubits r_1 and r_2 are both observed to be 0.

is a nonadaptive version of the standard single-qubit teleportation circuit [?]. It faithfully teleports the state if and only if the observed output of the B^\dagger -gate on the right is 00. After inserting each teleportation circuit, the gates acting before and after it are now acting on different qubits. Further, it is important to note that any entanglement the qubit state has with other qubits is easily seen to be preserved in the teleported qubit. The input qubits of \mathcal{C}' are those of \mathcal{C} . The output qubits of \mathcal{C}' are of two kinds: output qubits corresponding to outputs of \mathcal{C} are the *original outputs*; the other outputs are the qubits (in pairs) coming from the added B^\dagger -gates. We'll call the measurement of each such pair a *Bell measurement*, even though it is really in the computational basis.

In addition to the gates in \mathcal{C} , \mathcal{C}' uses only B -gates to make the initial EPR pairs and B^\dagger -gates for the Bell measurements. A sample transformation is shown in Figure 2. \mathcal{C}' has depth three since it uses the first layer to make the initial EPR states and the third layer to rotate the Bell basis back to the computational basis. All the gates of \mathcal{C} appear on the second layer. From the above construction and the properties of the teleportation module, it is not hard to see that for all x of the appropriate length,

$$\begin{aligned} \Pr[\mathcal{C}(x)] &= \Pr[\text{all original outputs of } \mathcal{C}' \text{ are } 0 \mid \text{all qubit states are teleported correctly}] \\ &= \Pr[\text{all original outputs of } \mathcal{C} \text{ are } 0 \mid \text{all Bell measurement results are } 00] \\ &= \frac{\Pr[\mathcal{C}'(x)]}{\Pr[\text{all Bell measurement results are } 00]}, \end{aligned}$$

since the Bell measurements are among the output measurements of \mathcal{C}' . Let k be the number of B^\dagger -gates on layer 3. Clearly, $k \leq |\mathcal{C}|$, and it is well-known that each Bell measurement will give 00 with probability 1/4, independent of all other measurements. So the lemma follows by setting $m = 2k$. \square

Proof of Theorem 3.1. **NQP** [?] is defined as the class of languages recognized by quantum Turing machines (equivalently, uniform quantum circuit families over a finite set of gates) where the acceptance criterion is that the accepting state appear with nonzero probability. It is known [?, ?] that **NQP** is equal to the counting class $\mathbf{C}_{\neq} \mathbf{P}$ [], which contains **NP** and is hard for the polynomial hierarchy. Since **QNC**⁰ circuit families must also draw their gates from some finite set, we clearly have $\mathbf{NQNC}^0 \subseteq \mathbf{NQP}$. The reverse containment follows from our construction: an arbitrary circuit \mathcal{C} is transformed into a depth-three circuit \mathcal{C}'




Figure 2: A sample transformation from \mathcal{C} to \mathcal{C}' . The circuit \mathcal{C} on the left has five gates: S , T , U , V , and W , with subscripts added to mark which qubits each gate is applied to. The qubits in \mathcal{C}' are numbered corresponding to those in \mathcal{C} .

with the same gates as \mathcal{C} plus B and B^\dagger . Moreover, \mathcal{C}' accepts with nonzero probability iff \mathcal{C} does. Thus an **NQP** language L recognized by a uniform quantum circuit family over a finite set of quantum gates is also recognized by a uniform depth-three circuit family over a finite set of quantum gates, and so $L \in \mathbf{NQNC}^0$. \square

Using the gate teleportation apparatus of Gottesmann and Chuang [?], DiVincenzo and Terhal also construct a depth-three² quantum circuit \mathcal{C}' out of an arbitrary circuit \mathcal{C} (over CNOT and single-qubit gates) with a similar relationship of acceptance probabilities. However, they only teleport the CNOT gate, and their \mathcal{C}' may contain single-qubit gates formed by compositions of arbitrary numbers of single-qubit gates from \mathcal{C} . (Such gates may not even be approximable in constant depth by circuits over a fixed finite family of gates.) When their construction is applied to each circuit in a uniform family, the resulting circuits are thus not generally over a finite gate set, even if the original circuits were.

Our construction solves this problem by teleporting every qubit state in between all gates involving it. Besides B and B^\dagger , we only use the gates of the original circuit. We also are able to bypass the CNOT gate teleportation technique of [?], using instead basic single-qubit teleportation [?], which works with arbitrary gates.

3.2 Simulating \mathbf{QNC}^0 circuits approximately is easy

In this section we prove that $\mathbf{BQNC}_{\epsilon,\delta}^0 \subseteq \mathbf{P}$ for certain ϵ, δ . For convenience we will assume that all gates used in quantum circuits are either one- or two-qubit gates that have “reasonable” matrix elements—algebraic numbers, for instance. Our results can apply more broadly, but they will then require greater care to prove.

For a quantum circuit \mathcal{C} , we define a dependency graph over the set of its output qubits.

Definition 3.4 *Let \mathcal{C} be a quantum circuit and let p and q be qubits of \mathcal{C} . We say that q depends on p if there is a forward path in \mathcal{C} starting at q before the first layer, possibly passing through gates, and ending at q after the last layer. More formally, we can define dependence by induction on the depth of \mathcal{C} . For depth zero, q depends on p iff $q = p$. For depth $d > 0$, let \mathcal{C}' be the same as \mathcal{C} but missing the first layer. Then q depends on p (in \mathcal{C}) iff there is a qubit r such that q depends on r (in \mathcal{C}') and either $p = r$ or there is a gate on the first layer of \mathcal{C} that involves both p and r .*

Definition 3.5 *For \mathcal{C} a quantum circuit and q a qubit of \mathcal{C} , define*

$$D_q = \{p \mid q \text{ depends on } p\}.$$

If S is a set of qubits of \mathcal{C} , define $D_S = \bigcup_{q \in S} D_q$. Let the dependency graph of \mathcal{C} be the undirected graph with the output qubits of \mathcal{C} as vertices, and with an edge between two qubits q_1 and q_2 iff $D_{q_1} \cap D_{q_2} \neq \emptyset$.

If \mathcal{C} has depth d , then it is easy to see that the degree of its dependency graph is less than 2^{2d} . The following lemma is straightforward.

²They count the depth as four, but they include the final measurement as an additional layer whereas we do not.

Lemma 3.6 *Let \mathcal{C} be a quantum circuit and let S and T be sets of output qubits of \mathcal{C} . Fix an input x and bit vectors u and v with lengths equal to the sizes of S and T , respectively. Let $E_{S=u}$ (respectively $E_{T=v}$) be the event that the qubits in S (respectively T) are observed to be in the state u (respectively v) in the final state of \mathcal{C} on input x . If $D_S \cap D_T = \emptyset$, then $E_{S=u}$ and $E_{T=v}$ are independent.*

For an algebraic number a , we let $\|a\|$ be the size of some reasonable representation of a . The results in this section follow from the next theorem.

Theorem 3.7 *There is a deterministic decision algorithm A which takes as input*

1. *a quantum circuit \mathcal{C} with depth d and n input qubits,*
2. *a binary string x of length n , and*
3. *an algebraic number $t \in [0, 1]$,*

and behaves as follows: Let D be one plus the degree of the dependency graph of \mathcal{C} . A runs in time $\text{Poly}(|\mathcal{C}|, 2^{2^d}, \|t\|)$, and for any $0 \leq t < \frac{1}{D}$,

- *if $\Pr[\mathcal{C}(x)] \geq 1 - t$, then A accepts, and*
- *if $\Pr[\mathcal{C}(x)] < 1 - Dt$, then A rejects.*

Note that since $D \leq 2^{2^d}$, if $t < 2^{-2^d}$, then A will reject when $\Pr[\mathcal{C}(x)] < 1 - 2^{2^d}t$.

Proof of Theorem 3.7. On input (\mathcal{C}, x, t) as above,

1. A computes the dependency graph $G = (V, E)$ of \mathcal{C} and its degree, and sets D to be the degree plus one.
2. A finds a D -coloring $c : V \rightarrow \{1, \dots, D\}$ of G via a standard greedy algorithm.
3. For each output qubit $q \in V$, A computes P_q —the probability that 0 is measured on qubit q in the final state (given input x).
4. For each color $i \in \{1, \dots, D\}$, let $B_i = \{q \in V \mid c(q) = i\}$. A computes

$$P_{B_i} = \prod_{q \in B_i} P_q,$$

which by Lemma 3.6 is the probability that all qubits colored i are observed to be 0 in the final state.

5. If $P_{B_i} \geq 1 - t$ for all i , the A accepts; otherwise, A rejects.

We first show that A is correct. If $\Pr[\mathcal{C}(x)] \geq 1 - t$, then for each $i \in \{1, \dots, D\}$,

$$1 - t \leq \Pr[\mathcal{C}(x)] \leq P_{B_i},$$

and so A accepts. On the other hand, if $\Pr[\mathcal{C}(x)] < 1 - Dt$, then

$$Dt < 1 - \Pr[\mathcal{C}(x)] \leq \sum_{i=1}^D (1 - P_{B_i}),$$

so there must exist an i such that $t < 1 - P_{B_i}$, and thus A rejects.

To show that A runs in the given time, first we show that the measurement statistics of any output qubit can be calculated in time polynomial in 2^{2^d} . Pick an output qubit q . By looking at \mathcal{C} we can find D_q in time $\text{Poly}(|\mathcal{C}|)$. Since \mathcal{C} has depth d and uses gates on at most two qubits each, D_q had cardinality at most 2^d . Then we simply calculate the measurement statistics of output qubit q from the input state restricted to D_q , i.e., with the other qubits traced out. This can be done by computing the state layer by layer, starting with layer one, and at each layer tracing out qubits when they no longer can reach q . Because of the partial traces, the state will in general be a mixed state so we maintain it as a density operator. We are multiplying matrices of size at most $2^{2^d} \times 2^{2^d}$ at most $O(d)$ times. All this will take time polynomial in 2^{2^d} , provided we can show that the individual field operations on the matrix elements do not take too long.

Since there are finitely many gates to choose from, their (algebraic) matrix elements generate a field extension F of \mathbb{Q} with finite index r . We can thus store values in F as r -tuples of rational numbers, with the field operations of F taking polynomial time. Furthermore, one can show that for $a, b \in F$, $\|ab\| = O(\|a\| + \|b\|)$ and $\|\sum_i 1^n a_i\| = O(n \cdot \max_i \|a_i\|)$ for any $a_1, \dots, a_n \in F$. A bit of calculation then shows that the intermediate representations of numbers do not get too large.

The dependency graph and its coloring can clearly be computed in time $\text{Poly}(|\mathcal{C}|)$. The only things left are the computation of the P_{B_i} and their comparison with $1 - t$. For reasons similar to those above for matrix multiplication, this can be done in time $\text{Poly}(|\mathcal{C}|, 2^{2^d}, \|t\|)$.

□

Corollary 3.8 *Suppose ϵ and δ are polynomial-time computable, and for any quantum circuit C of depth d , $\delta_C \geq 1 - 2^{-2^d}(1 - \epsilon_C)$. Then*

$$\text{BQNC}_{\epsilon, \delta}^0 \subseteq \text{P}.$$

Proof. For each C of depth d in the circuit family and each input x , apply the algorithm A of Theorem 3.7 with $t = 1 - \delta_C = 2^{-2^d}(1 - \epsilon_C)$, noting that $D \leq 2^{2^d}$. □

Corollary 3.9 *For quantum circuit C , let $\delta_C = 1 - 2^{-(2^d+1)}$, where d is the depth of C . Then*

$$\text{BQNC}_{(1/2), \delta}^0 \subseteq \text{P}.$$

Proof. Apply algorithm A to each circuit, setting $t = 2^{-(2d+1)}$. □

Corollary 3.10 $\text{BQNC}_{1,1}^0 \subseteq \mathbf{P}$.

Proof. Apply algorithm A to each circuit, setting $t = 0$. □

Corollary 3.11 $\text{EQNC}^0 \subseteq \mathbf{P}$.

Proof. Clearly, $\text{EQNC}^0 \subseteq \text{BQNC}_{1,1}^0$. □

4 Conclusions and further research

log log depth
inclusion in smaller classes than \mathbf{P}
classical preprocessing and postprocessing
narrowing the gap

Acknowledgments