

Bounds on the Power of Constant-Depth Quantum Circuits*

S. Fenner[†] F. Green[‡] S. Homer[§] Y. Zhang[†]

March 7, 2005

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, for $0 < \epsilon \leq \delta \leq 1$, we define $\mathbf{BQNC}_{\epsilon, \delta}^0$ to be the class of languages recognized by constant depth, polynomial-size quantum circuits with acceptance probability either $< \epsilon$ (for rejection) or $\geq \delta$ (for acceptance). We show that

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

provided that $1 - \delta \leq 2^{-2d}(1 - \epsilon)$, where d is the circuit depth.

On the other hand, we adapt and extend ideas of Terhal & DiVincenzo [TD04] 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 analog of the class \mathbf{NQP} . This essentially refutes a conjecture of Green et al. that $\mathbf{NQACC} \subseteq \mathbf{TC}^0$ [GHMP02].

1 Introduction

This paper investigates \mathbf{QNC}^0 circuits, that is, families of quantum circuits with polynomial size and constant depth, using quantum gates of bounded width. Informally speaking, we show that

1. decision problems computed by \mathbf{QNC}^0 circuits within certain error bounds can be computed classically in polynomial time (Corollary 3.10), yet
2. computing probability amplitudes of \mathbf{QNC}^0 circuits exactly is as hard as computing arbitrary $\#\mathbf{P}$ functions, even when we restrict the circuits to depth three over a fixed finite set of quantum gates (Theorem 3.1).

The second result extends and improves work of Terhal & DiVincenzo [TD04]. Combined with recent results of Aaronson [Aar04] it shows that

$$\mathbf{postBQNC}^0 = \mathbf{postBQP} = \mathbf{PP} \quad (\text{Theorem 3.5}),$$

where $\mathbf{postBQNC}^0$ (respectively, $\mathbf{postBQP}$) is the class of languages computable by \mathbf{QNC}^0 (respectively, polynomial-size quantum) circuits with bounded error and postselection.

*This work was supported in part by the National Security Agency (NSA) and Advanced Research and Development Agency (ARDA) under Army Research Office (ARO) contract numbers DAAD 19-02-1-0058 (for S. Homer, and F. Green) and DAAD 19-02-1-0048 (for S. Fenner and Y. Zhang).

[†]Dept. of Computer Science and Eng., University of South Carolina, Columbia, SC 29208, {fenner|zhang29}@cse.sc.edu

[‡]Dept. of Mathematics and Computer Science, Clark University, Worcester, MA 01610, fgreen@black.clarku.edu

[§]Dept. of Computer Science, Boston University, Boston, MA 02215, homer@bu.edu

Much can be done with $O(\log n)$ -depth quantum circuits (\mathbf{QNC}^1 circuits). Moore & Nilsson showed that many circuits can be parallelized to log depth—in particular, those implementing stabilizer error-correcting codes [MN02]. Cleve & Watrous were able to approximate the Quantum Fourier Transform over modulus 2^n with $O(\log n)$ -depth circuits [CW00]. At first glance, \mathbf{QNC}^0 circuits appear extremely weak; one might expect that nothing can be computed with \mathbf{QNC}^0 circuits that cannot already be computed classically in polynomial time (even in \mathbf{NC}^0), since each output qubit can only be connected to a constant number of input qubits. This is certainly the case for decision problems if we restrict ourselves to observing a *single* output qubit, but surprisingly, this is still open in the more reasonable case where we observe several outputs at once, then apply some classical Boolean acceptance criterion on the results. The reason why these circuits are probably hard to simulate classically is that, although each individual output probability amplitude is easy to compute, it may be the case that different output qubits are correlated with each other, and the correlation graph has a high rate of expansion. (Terhal & DiVincenzo show that this is *not* the case for a depth two circuit, which can thus be simulated easily classically [TD04].)

To get use out of $o(\log n)$ -depth quantum circuits, people have augmented them with quantum gates of unbounded fan-in. There are a number of unbounded-width gate classes studied in the literature, most being defined in analogy with classical Boolean gates. The generalized Toffoli gate (see Section 2.1) is the quantum equivalent of the unbounded Boolean AND-gate. Likewise, there are quantum equivalents of Mod-gates and threshold gates. One particular quantum gate corresponds to something taken almost completely for granted in Boolean circuits—fan-out. A fan-out gate copies the (classical) value of a qubit to several other qubits at once.¹ Using these gates, one can define quantum versions of various classical circuit classes: the previously mentioned \mathbf{QNC}^k (Moore & Nilsson [MN02]), \mathbf{QAC}^k and \mathbf{QACC}^k (Moore [Moo99], Green et al. [GHMP02]), and \mathbf{QTC}^k are analogous to \mathbf{NC}^k , \mathbf{AC}^k , \mathbf{ACC} , and \mathbf{TC}^k , respectively. The case of particular interest is when $k = 0$. All these classes are allowed constant-width gates drawn from a finite family. The classes differ in the additional gates allowed. \mathbf{QNC} is the most restrictive class; all gates must have bounded width. \mathbf{QAC} circuits are allowed generalized Toffoli gates, and \mathbf{QACC} circuits are allowed Mod_q -gates, where q is kept constant in each circuit family. \mathbf{QTC} circuits are allowed quantum threshold gates. See Section 2.1 for detailed definitions of most of these classes.

Although small-depth quantum circuit classes are defined analogously to Boolean classes, their properties have turned out to be quite different from their classical versions. A simple observation of Moore [Moo99] shows that the n -qubit fan-out gate and the n -qubit parity (Mod_2) gate are equivalent up to constant depth, i.e., each can be simulated by a constant-depth circuit using the other. This is completely different from the classical case, where parity cannot be computed even with \mathbf{AC}^0 circuits where fan-out is unrestricted [Ajt83, FSS84]. Later, Green et al. showed that all quantum Mod_q -gates are *constant-depth equivalent* for $q > 1$, and are thus all equivalent to fan-out. Thus, for any $q > 1$,

$$\mathbf{QNC}_f^0 = \mathbf{QACC}^0(q) = \mathbf{QACC}^0.$$

(The f subscript means, “with fan-out.”) The classical analogs of these classes are provably different. In particular, classical Mod_p and Mod_q gates are not constant-depth equivalent if p and q are distinct primes, and neither can be simulated by \mathbf{AC}^0 circuits [Raz87, Smo87].

Building on ideas in [MN02], Høyer & Špalek used \mathbf{QNC}^0 circuits with unbounded fan-out gates to parallelize a sequence of commuting gates applied to the same qubits, and thus greatly reduced the depth of circuits for various purposes [HŠ03]. They showed that threshold gates can be approximated in constant depth this way, and they can be implemented exactly if Toffoli gates are also allowed. Thus $\mathbf{QTC}_f^0 = \mathbf{QACC}^0$ as well. Threshold gates, and hence fanout gates, are quite powerful; many important arithmetic operations can be computed in constant depth with threshold gates [SBKH93]. This implies that the quantum Fourier transform—the quantum part of Shor’s factoring algorithm—can be approximated in constant depth using fanout gates.

All these results rely for their practicality on unbounded-width quantum gates being available, especially fan-out or some (any) Mod gate. Unfortunately, making such a gate in the lab remains a daunting prospect; it is hard enough just to fabricate a reliable CNOT gate. Much more likely in the short term is that only

¹There is no violation of the No-Cloning Theorem here; only the classical value is copied.

one- and two-qubit gates will be available, which brings us back to the now more interesting question of \mathbf{QNC}^0 . How powerful is this class?

A handful of hardness results about simulating constant-depth quantum circuits with constant-width gates are given by Terhal & DiVincenzo [TD04]. They show that if one can classically efficiently simulate, via sampling, the acceptance probability of quantum circuits of depth at least three using one- and two-qubit gates, then $\mathbf{BQP} \subseteq \mathbf{AM}$. They also showed that the polynomial hierarchy collapses if one can efficiently compute the acceptance probability exactly for such circuits. (Actually, a much stronger result follows from their proof, namely, $\mathbf{P} = \mathbf{PP}$.) Their technique uses an idea of Gottesman & Chuang for teleporting CNOT gates [GC99] to transform an arbitrary quantum circuit with CNOT and single-qubit gates into a depth-three circuit whose acceptance probability is proportional to, though exponentially smaller than, the original circuit. Their results, however, only hold on the supposition that depth-three circuits with *arbitrary* single-qubit and CNOT gates are simulatable. We weaken their hypothesis by showing how to produce a depth-three circuit with essentially the same gates as the original circuit. In addition, we can get by with only simple qubit state teleportation [BBC⁺93]. Our results immediately show that the class \mathbf{NQNC}^0 (the constant-depth analog of \mathbf{NQP} , see below), is actually the same as \mathbf{NQP} , which is known to be as hard as the polynomial hierarchy [FGHP99]. We give this result in Section 3.1. It underscores yet another drastic difference between the quantum and classical case: while \mathbf{AC}^0 is well contained in \mathbf{P} , \mathbf{QNC}^0 circuits (even just depth-three) can have amazingly complex behavior. Our result is also tight; Terhal & DiVincenzo showed that the acceptance probabilities of depth-two circuits over one- and two-qubit gates are computable in polynomial time.

In Section 3.2, we give contrasting upper bounds for \mathbf{QNC}^0 -related language classes. We show that various bounded-error versions of \mathbf{QNC}^0 (defined below) are contained in \mathbf{P} . Particularly, $\mathbf{EQNC}^0 \subseteq \mathbf{P}$, where \mathbf{EQNC}^0 is the constant-depth analog of the class \mathbf{EQP} (see below). Our proof uses elementary probability theory, together with the fact that single output qubit measurement probabilities can be computed directly, and the fact that output qubits are “largely” independent of each other. In hindsight, it is not too surprising that $\mathbf{EQNC}^0 \subseteq \mathbf{P}$. \mathbf{EQNC}^0 sets a severe limitation on the behavior of the circuit: it must accept with certainty or reject with certainty. This containment is more surprising (to us) for the bounded-error \mathbf{QNC}^0 classes.

We give open questions and suggestions for further research in Section 4.

2 Preliminaries

2.1 Gates and circuits

We assume prior knowledge of basic concepts in computational complexity: polynomial time, \mathbf{P} , \mathbf{NP} , as well as the counting class $\#\mathbf{P}$ [Val79]. Information can be found, for example, in Papadimitriou [Pap94]. The class $\mathbf{C}_{\neq}\mathbf{P}$ ($\mathbf{coC}_{=}\mathbf{P}$) was defined by Wagner [Wag86]. One way of defining $\mathbf{C}_{\neq}\mathbf{P}$ is as follows: a language L is in $\mathbf{C}_{\neq}\mathbf{P}$ iff there are two $\#\mathbf{P}$ functions f and g such that, for all x , $x \in L \iff f(x) \neq g(x)$. $\mathbf{C}_{\neq}\mathbf{P}$ was shown to be hard for the polynomial hierarchy by Toda & Ogihara [TO92].

We will also assume some background in quantum computation and the quantum circuit model. See Nielsen and Chuang [NC00] for a good reference of basic concepts and notation.

Our notion of quantum circuits is fairly standard (again see, for example, [NC00]): a series of quantum gates, drawn from some specified set of unitary operators, acting on some specified number of qubits, labeled $1, \dots, m$. 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 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. For the purposes of computing decision problems, we will say that the circuit *accepts* its input if all the output qubits are observed to be 0 in the final state. Otherwise the circuit rejects. This acceptance criterion is simple, and it is essentially the one given in [GHMP02]. Although we do not study it here, one may consider other acceptance criteria,

for example, feeding the observed outputs into an arbitrary polynomial time classical computation. To our knowledge, the power of such a model has not been studied.

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 sum of the arities of all the gates occurring in the circuit. C may be laid out by partitioning its gates into *layers* $1, \dots, d$, such that (i) gates in the same layer 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 .

The standard quantum complexity classes (of languages) can be defined in terms of quantum circuit families. A quantum circuit family is a sequence $\{C_n\}_{n \geq 0}$ of quantum circuits, where each C_n has n inputs. We say that $\{C_n\}$ is *uniform* if there is a (classical) polynomial-time algorithm that outputs a description of C_n on input 0^n . The classes **BQP**, **EQP**, and **NQP** are defined using polynomial size quantum circuits with gates drawn from some fixed finite universal set of gates (see [BV97, BBBV97, ADH97]). They differ only in their acceptance criteria: for **BQP**, there is a promise that $\Pr[C(x)] \in [0, \frac{1}{3}] \cup [\frac{2}{3}, 1]$ and $C(x)$ accepts iff $\Pr[C(x)] \in [\frac{2}{3}, 1]$; **EQP** is the same as **BQP** except that the promise is that It is known that $\Pr[C(x)] \in \{0, 1\}$; for **NQP**, there is no promise, and $C(x)$ accepts iff $\Pr[C(x)] > 0$. It is well-known that $\mathbf{P} \subseteq \mathbf{EQP} \subseteq \mathbf{BQP}$. It was shown in [FGHP99, YY99] that $\mathbf{NQP} = \mathbf{C}_{\neq} \mathbf{P}$, and is thus hard for the polynomial hierarchy.

2.2 Complexity classes using QNC circuits

The circuit class **QNC** was first suggested by Moore and Nilsson [MN02] as the quantum analog of the class **NC** of bounded fan-in Boolean circuits with polylogarithmic depth and polynomial size. We define the class \mathbf{QNC}^k in the same fashion as definitions in Green, Homer, Moore, & Pollett [GHMP02] with some minor modifications.

Definition 2.1 ([MN02]). \mathbf{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 ancilla. 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.

Next we define the language classes \mathbf{NQNC}^k and \mathbf{EQNC}^k . These are \mathbf{QNC}^k analogs of the classes **NQP** and **EQP**, respectively.

Definition 2.2 ([GHMP02]). Let $k \geq 0$ be an integer.

- \mathbf{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 ,

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

- \mathbf{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, & Pollett implicitly consider the output qubits of C_n to be all the qubits in C_n [GHMP02]. 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 [GHMP02], and one can certainly ask about similar classes for \mathbf{QNC}^k 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.² 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

We briefly mention some other quantum circuit classes related to \mathbf{QNC}^0 . See Green et al. [GHMP02] for details.

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

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

3 Main results

3.1 Simulating \mathbf{QNC}^0 circuits exactly is hard

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

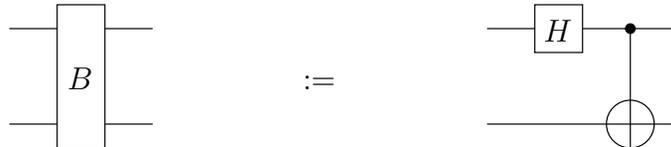
As a corollary, we essentially solve an open problem of Green et al. [GHMP02]. They conjectured that $\mathbf{NQACC} \subseteq \mathbf{TC}^0$, the class of constant-depth Boolean circuits with threshold gates.

Corollary 3.2. For any $k \geq 0$,

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

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

Let B be the two-qubit Bell gate, consisting of a Hadamard gate followed by a CNOT gate:



²One 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 $1/\text{poly}$.

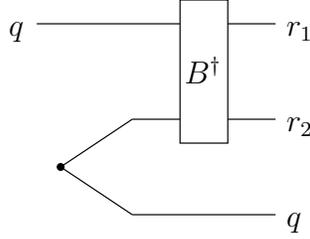
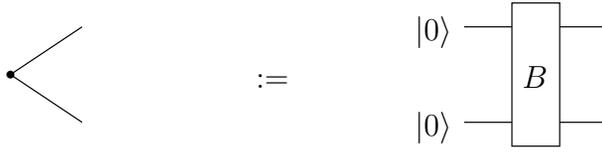


Figure 1: The nonadaptive teleportation module [TD04]. The state in qubit q is teleported correctly iff the qubits r_1 and r_2 are both observed to be 0.

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 Terhal & DiVincenzo [TD04], but ours is stronger in two respects, discussed below: it works for any family of gates allowed in the original circuit, and introduces no new gates except B and B^\dagger . 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, is a nonadaptive version of the standard single-qubit teleportation circuit [BBC⁺93]. 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 *check* qubits (in pairs) coming from the added B^\dagger -gates. We'll call the measurement of each such pair a *check measurement*.

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 check 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

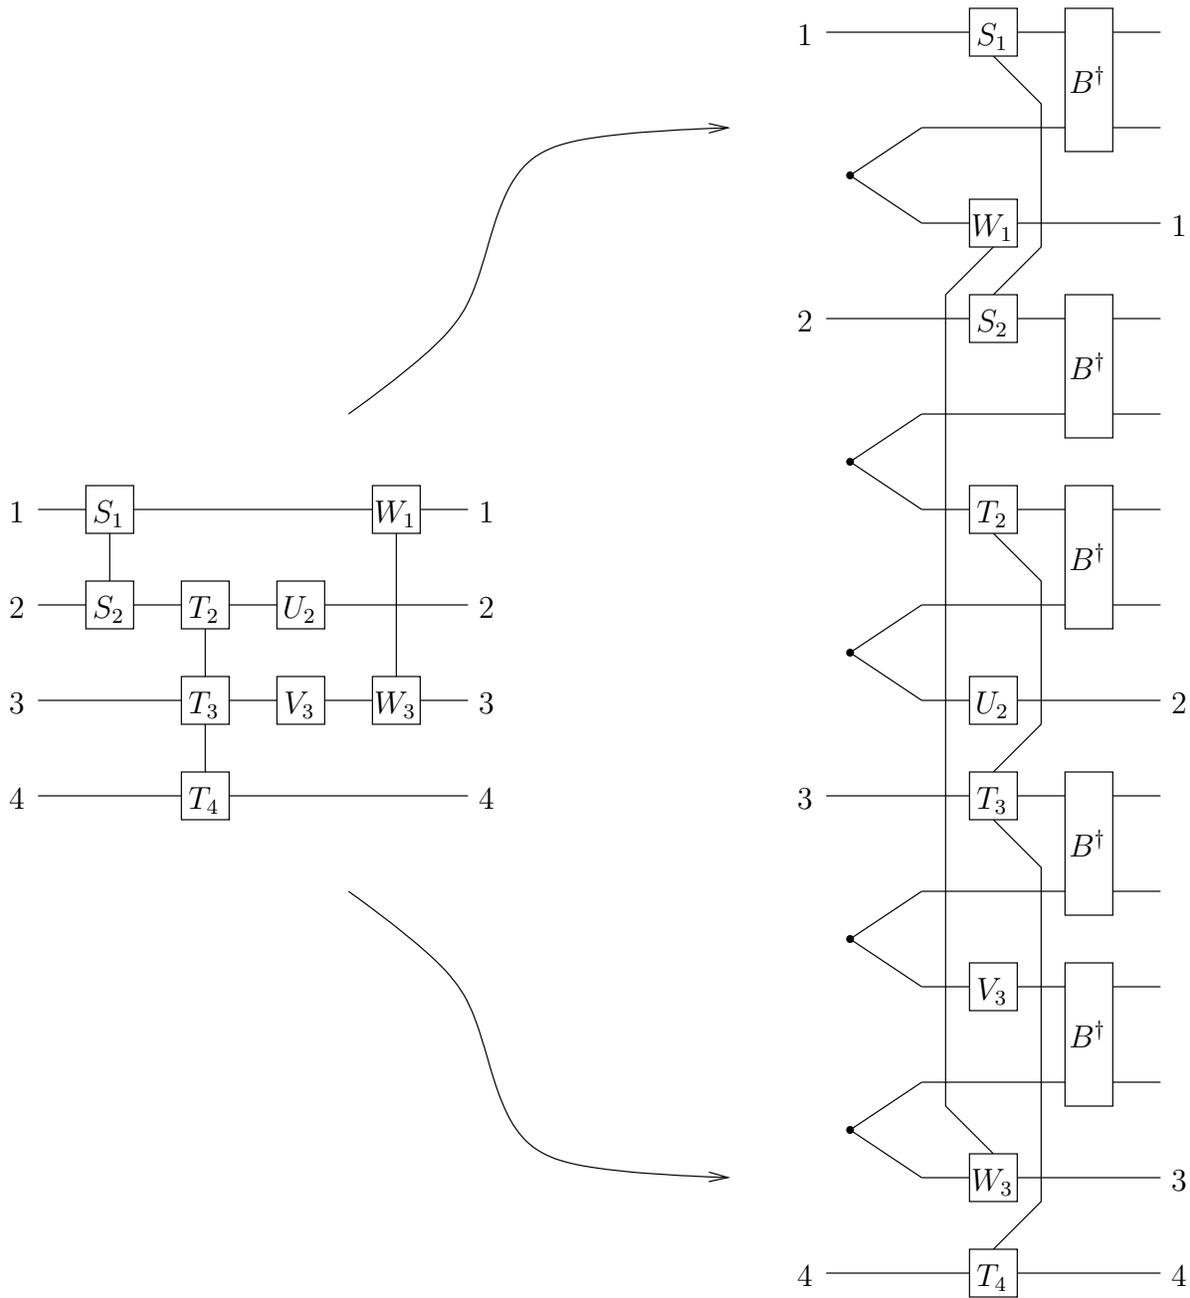


Figure 2: A sample transformation from \mathcal{C} to \mathcal{C}' . The circuit \mathcal{C} on the left has five gates: S , T , U , V , and W , with subscripts added to mark which qubits each gate is applied to. The qubits in \mathcal{C}' are numbered corresponding to those in \mathcal{C} .

teleportation module, it is not hard to see that for all x of the appropriate length,

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

since the check 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 check 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. As mentioned before, **NQP** [ADH97] 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 [FGHP99, YY99] that $\mathbf{NQP} = \mathbf{C}_{\neq} \mathbf{P}$, which contains **NP** and is hard for the polynomial hierarchy. Since \mathbf{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}' 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

Remarks. Using the gate teleportation apparatus of Gottesmann and Chuang [GC99], Terhal & DiVincenzo also construct a depth-three³ 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 [GC99], using instead basic single-qubit teleportation [BBC⁺93], which works with arbitrary gates.

Aaronson [Aar04] has recently considered the power of postselection in quantum circuits, defining a “new” class **postBQP** and showing that, actually, $\mathbf{postBQP} = \mathbf{PP}$, thereby giving a new quantum proof of a celebrated result of Beigel, Reingold, & Spielman [BRS91] stating that **PP** is closed under intersection. The class **postBQP** is defined using uniform families of polynomial-size quantum circuits \mathcal{C} with two distinguished qubits: a regular output qubit q and a postselection qubit p , subject to the promise that in the final state, p is observed to be 1 with positive probability. $\Pr[\mathcal{C}(x)]$ is then defined as the conditional probability that q is observed to be 1, given that p is observed to be 1. The acceptance criterion on $\Pr[\mathcal{C}(x)]$ is the same as with **BQP**.

Clearly, the definition of **postBQP** remains unchanged if we swap the roles of 0 and 1 in the measurement of p , and if we allow multiple postselection qubits and condition $\Pr[\mathcal{C}(x)]$ on their all being 0. Thus the following definition gives a reasonable constant-depth version of **postBQP**:

Definition 3.4. A language L is in $\mathbf{postBQNC}^0$ if there is a uniform family $\{C_n\}_{n \geq 0}$ of constant-depth, polynomial-size quantum circuits (over some fixed finite universal set of gates) with output qubit q and postselection qubits p_1, \dots, p_k such that, for all inputs x of length n , $\Pr[p_1, \dots, p_k \text{ of } C_n \text{ are all } 0] > 0$, and the quantity

$$\Pr[C_n(x)] = \Pr[q \text{ is } 1 \mid p_1, \dots, p_k \text{ are all } 0]$$

is in $[0, \frac{1}{3}] \cup [\frac{2}{3}, 1]$, and $x \in L$ iff $\Pr[C_n(x)] \geq 1/2$.

³They count the depth as four, but they include the final measurement as an additional layer whereas we do not.

The construction in the proof of Lemma 3.3 immediately yields the following:

Theorem 3.5. $\text{postBQNC}^0 = \text{postBQP}$, and hence $\text{postBQNC}^0 = \text{PP}$.

Proof. Given a **postBQP** circuit C with output q and postselection qubit p , we first place a final X gate on p (flipping its bit value), then we convert the modified C into a depth-three circuit C' as in the proof of Lemma 3.3, and we include all the qubits involved in check measurements as postselection qubits (in addition to p). C' then produces the same conditional probability for the value of q that C does. That $\text{postBQP} = \text{PP}$ follows from [Aar04]. \square

3.2 Simulating QNC^0 circuits approximately is easy

In this section we prove that $\text{BQNC}_{\epsilon, \delta}^0 \subseteq \mathbf{P}$ 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 3.6. 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 p 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 3.7. 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.

Lemma 3.8. 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 3.9. 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 $\text{Poly}(|\mathcal{C}|, 2^{2d}, \|t\|)$, and

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

Proof of Theorem 3.9. 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 3.8 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^{2d} . 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 $2^{2d} \times 2^{2d}$ at most $O(d)$ times. All this will take time polynomial in 2^{2d} , 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 $\text{Poly}(|\mathcal{C}|, 2^{2d}, \|t\|)$. \square

Corollary 3.10. *Suppose ϵ and δ are polynomial-time computable, and for any quantum circuit C of depth d , $\delta_C = 1 - 2^{-2d}(1 - \epsilon_C)$. Then*

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

Proof. For each C of depth d in the circuit family and each input x , apply the algorithm A of Theorem 3.9 with $t = 1 - \delta_C = 2^{-2d}(1 - \epsilon_C)$, noting that $D \leq 2^{2d}$. \square

The following few corollaries are instances of Corollary 3.10.

Corollary 3.11. For quantum circuit C , let $\delta_C = 1 - 2^{-(2d+1)}$, where d is the depth of C . Then

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

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

Corollary 3.12. $\mathbf{BQNC}_{1,1}^0 \subseteq \mathbf{P}$.

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

Corollary 3.13. $\mathbf{EQNC}^0 \subseteq \mathbf{P}$.

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

Corollaries 3.12 and 3.13 can actually be proven more directly. We simply compute, for each output, its probability of being 0. We accept iff all probabilities are 1.

We observe here that by a simple proof using our techniques, one can show that a \mathbf{QNC}^0 circuit cannot implement the generalized Toffoli gate, because its target depends on nonconstantly many input qubits.

4 Conclusions, open questions, and further research

Our upper bound results in Section 3.2 can be improved in certain ways. For example, the containment in \mathbf{P} is easily seen to apply to $(\log \log n + O(1))$ -depth circuits as well. Can we increase the depth further? Another improvement would be to put $\mathbf{BQNC}_{\epsilon,\delta}^0$ into classes smaller than \mathbf{P} . $\mathbf{LOGSPACE}$ seems manageable. \mathbf{NC}^1 ?

There are some general questions about whether we have the “right” definitions for these classes. For example, the accepting outcome is defined to be all outputs being 0. One can imagine more general accepting conditions, such as arbitrary classical polynomial-time postprocessing. If we allow this, then all our classes will obviously contain \mathbf{P} . If we allow arbitrary classical polynomial-time preprocessing, then all our classes will be closed under ptime m -reductions (Karp reductions).

Finally, there is the question of the probability gap in the definitions of \mathbf{BQNC}^0 . Ideally, we would like to narrow this gap (to $1/\text{poly}$, say) independent of the circuit depth, and still get containment in \mathbf{P} .

Acknowledgments

We would like to thank David DiVincenzo, Mark Heiligman, and Scott Aaronson for helpful conversations.

References

- [Aar04] Scott Aaronson. Quantum computing, postselection, and probabilistic polynomial-time, 2004, quant-ph/0412187. Manuscript.
- [ADH97] L. M. Adleman, J. DeMarras, and M.-D. A. Huang. Quantum computability. *SIAM Journal on Computing*, 26(5):1524–1540, 1997.
- [Ajt83] M. Ajtai. Σ_1^1 formulæ on finite structures. *Annals of Pure and Applied Logic*, 24:1–48, 1983.
- [BBBV97] C. H. Bennett, E. Bernstein, G. Brassard, and U. Vazirani. Strengths and weaknesses of quantum computation. *SIAM Journal on Computing*, 26(5):1510–1523, 1997, quant-ph/9701001.
- [BBC⁺93] C. H. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, and W. Wootters. Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels. *Physical Review Letters*, 70(13):1895–1899, 1993.

- [BRS91] R. Beigel, N. Reingold, and D. Spielman. PP is closed under intersection. In *Proceedings of the 23rd annual ACM Symposium on Theory of Computing*, pages 1–9, 1991.
- [BV97] E. Bernstein and U. Vazirani. Quantum complexity theory. *SIAM Journal on Computing*, 26(5):1411–1473, 1997.
- [CW00] R. Cleve and J. Watrous. Fast parallel circuits for the quantum Fourier transform. In *Proceedings of the 41st IEEE Symposium on Foundations of Computer Science*, pages 526–536, 2000, quant-ph/0006004.
- [FGHP99] S. Fenner, F. Green, S. Homer, and R. Prui. Determining acceptance possibility for a quantum computation is hard for the polynomial hierarchy. *Proceedings of the Royal Society London A*, 455:3953–3966, 1999, quant-ph/9812056.
- [FSS84] M. Furst, J. B. Saxe, and M. Sipser. Parity, circuits, and the polynomial time hierarchy. *Mathematical Systems Theory*, 17:13–27, 1984.
- [GC99] D. Gottesman and I. L. Chuang. Demonstrating the viability of universal quantum computation using teleportation and single-qubit operations. *Letters to Nature*, 402:390–393, 1999.
- [GHMP02] F. Green, S. Homer, C. Moore, and C. Pollett. Counting, fanout and the complexity of quantum ACC. *Quantum Information and Computation*, 2:35–65, 2002, quant-ph/0106017.
- [HŠ03] P. Høyer and R. Špalek. Quantum circuits with unbounded fan-out. In *Proceedings of the 20th Symposium on Theoretical Aspects of Computer Science*, volume 2607 of *Lecture Notes in Computer Science*, pages 234–246. Springer-Verlag, 2003.
- [MN02] C. Moore and M. Nilsson. Parallel quantum computation and quantum codes. *SIAM Journal on Computing*, 31(3):799–815, 2002.
- [Moo99] C. Moore. Quantum circuits: Fanout, parity, and counting, 1999, quant-ph/9903046. Manuscript.
- [NC00] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.
- [Pap94] C. H. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.
- [Raz87] 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.
- [SBKH93] 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.
- [Smo87] R. Smolensky. Algebraic methods in the theory of lower bounds for Boolean circuit complexity. In *Proceedings of the 19th ACM Symposium on the Theory of Computing*, pages 77–82, 1987.
- [TD04] B. M. Terhal and D. P. DiVincenzo. Adaptive quantum computation, constant depth quantum circuits and arthur-merlin games. *Quantum Information and Computation*, 4(2):134–145, 2004.
- [TO92] S. Toda and M. Ogiwara. Counting classes are at least as hard as the polynomial-time hierarchy. *SIAM Journal on Computing*, 21(2):316–328, 1992.
- [Val79] L. Valiant. The complexity of computing the permanent. *Theoretical Computer Science*, 8:189–201, 1979.
- [Wag86] K. Wagner. The complexity of combinatorial problems with succinct input representation. *Acta Informatica*, 23:325–356, 1986.
- [YY99] T. Yamakami and A. C.-C. Yao. $\text{NQP}_{\mathbb{C}} = \text{co-C}_{=}P$. *Information Processing Letters*, 71(2):63–69, 1999, quant-ph/9812032.