

# Counting, Fanout, and the Complexity of Quantum ACC

Frederic Green

Department of Mathematics and Computer Science  
Clark University, Worcester, MA 01610  
fgreen@black.clarku.edu

Steven Homer \*

Computer Science Department  
Boston University, Boston, MA 02215  
homer@cs.bu.edu

Cristopher Moore †

Computer Science Department  
University of New Mexico, Albuquerque NM 87131  
and the Santa Fe Institute  
moore@cs.unm.edu

Christopher Pollett

Department of Mathematics  
University of California, Los Angeles, CA  
cpollett@math.ucla.edu

## Abstract

*We propose definitions of  $\text{QAC}^0$ , the quantum analog of the classical class  $\text{AC}^0$  of constant-depth circuits with AND and OR gates of arbitrary fan-in, and  $\text{QACC}[q]$ , the analog of the class  $\text{ACC}[q]$  where  $\text{Mod}_q$  gates are also allowed. We prove that parity or fanout allows us to construct quantum  $\text{MOD}_q$  gates in constant depth for any  $q$ , so  $\text{QACC}[2] = \text{QACC}$ . More generally, we show that for any  $q, p > 1$ ,  $\text{MOD}_q$  is equivalent to  $\text{MOD}_p$  (up to constant depth). This implies that  $\text{QAC}^0$  with unbounded fanout gates, denoted  $\text{QAC}_{\text{wf}}^0$ , is the same as  $\text{QACC}[q]$  and  $\text{QACC}$  for all  $q$ . Since  $\text{ACC}[p] \neq \text{ACC}[q]$  whenever  $p$  and  $q$  are distinct primes,  $\text{QACC}[q]$  is strictly more powerful than its classical counterpart, as is  $\text{QAC}^0$  when fanout is allowed. This adds to the growing list of quantum complexity classes which are provably more powerful than their classical counterparts.*

---

\*Supported in part by the NSF under grant NSF-CCR-9988310

†Supported in part by the NSF under grant NSF-PHY-0071139

We also develop techniques for proving upper bounds for  $\text{QACC}^0$  in terms of related language classes. We define classes of languages  $\text{EQACC}$ ,  $\text{NQACC}$  and  $\text{BQACC}_{\mathbf{Q}}$ . We define a notion of log-planar  $\text{QACC}$  operators and show the appropriately restricted versions of  $\text{EQACC}$  and  $\text{NQACC}$  are contained in  $\text{P/poly}$ . We also define a notion of log-gate restricted  $\text{QACC}$  operators and show the appropriately restricted versions of  $\text{EQACC}$  and  $\text{NQACC}$  are contained in  $\text{TC}^0$ .

## 1 Introduction

Advances in quantum computation in the last decade have been among the most notable in theoretical computer science. This is due to the surprising improvements in the efficiency of solving several fundamental combinatorial problems using quantum mechanical methods in place of their classical counterparts. These advances led to considerable efforts in finding new efficient quantum algorithms for classical problems and in developing a complexity theory of quantum computation.

While most of the original results in quantum computation were developed using quantum Turing machines, they can also be formulated in terms of quantum circuits, which yield a more natural model of quantum computation. For example, Shor [27] has shown that quantum circuits can factor integers more efficiently than any known classical algorithm for factoring. And quantum circuits have been shown (see Yao [33]) to provide a universal model for quantum computation.

The theory of circuit complexity has long been an important branch of theoretical computer science. Shallow circuits correspond to parallel algorithms that can be performed in small amounts of time on a massively parallel computer with constant communication delays, and so circuit complexity can be thought of as a study of how to solve problems in parallel. In addition, some low-lying circuit classes have beautiful algebraic characterizations, e.g. [4, 5, 17].

In [18, 19], Moore and Nilsson suggested a definition of  $\text{QNC}$ , the quantum analog of the class  $\text{NC}$  of problems solvable by circuits with polylogarithmic depth and polynomial size [21]. Here, we will study quantum versions of some additional circuit classes. Recall the following definitions:

1.  $\text{NC}^k$  consists of problems solvable by families of circuits of AND, OR, and NOT gates with depth  $\mathcal{O}(\log^k n)$  and size polynomial in  $n$ , where  $n$  is the size of the input, and where the AND and OR gates have just two inputs each.
2.  $\text{AC}^k$  is like  $\text{NC}^k$ , but where we allow AND and OR gates with unbounded fan-in, i.e. arbitrary numbers of inputs, in each layer of the circuit.
3.  $\text{ACC}^k[q]$  is like  $\text{AC}^k$ , but where we also allow  $\text{Mod}_q$  gates with unbounded fan-in, where  $\text{Mod}_q(x_1, \dots, x_n)$  outputs 1 iff the sum of the inputs is not a multiple of  $q$ .
4.  $\text{ACC}^k = \cup_q \text{ACC}^k[q]$ .
5.  $\text{NC} = \cup_k \text{NC}^k = \cup_k \text{AC}^k = \cup_k \text{ACC}^k$ .

Then we have

$$\text{AC}^0 \subset \text{ACC}^0[2] \subset \text{ACC}^0 \subseteq \text{NC}^{(1)} \subseteq \dots \subseteq \text{NC}$$

In fact, these first two inclusions are known to be proper [2, 13, 22, 30]. Neither MAJORITY nor PARITY are in  $AC^0$ , while the latter is trivially in  $ACC^0[2]$ . In addition,  $ACC^0[p]$  and  $ACC^0[q]$  are known to be incomparable whenever  $p$  and  $q$  are distinct primes. Thus these classes give us some of the few strict inclusions known in computational complexity theory. However, for all anyone knows,  $ACC^0[6]$  could contain PP, NP, and the entire polynomial hierarchy!

Quantum analogs of  $AC^0$  and ACC are defined and studied here. One central class that we examine is a quantum analog of  $AC^0$  that we denote  $QAC_{wf}^0$ .  $QAC_{wf}^0$  is the class of families of operators which can be built out of products of constantly many layers consisting of polynomial-sized tensor products of one-qubit gates (analogous to NOT's), Toffoli gates (analogous to AND's and OR's) and fan-out gates. The subscript “ $wf$ ” in the notation denotes “with fan-out.” The idea of fan-out in the quantum setting is subtle, as is made clear in Section 3 of this paper. The sub-class of  $QAC_{wf}^0$  that does not include fan-out gates is denoted simply  $QAC^0$ . An analog of  $ACC[q]$  (i.e., ACC circuit families only allowing  $Mod_q$  gates) is  $QACC[q]$ , defined similarly to  $QAC_{wf}^0$ , but replacing the fan-out gates with quantum  $Mod_q$  gates (which we denote as  $MOD_q$ ). The class QACC is  $\cup_q QACC[q]$ .

In this paper, we prove a number of results about QAC and QACC, and address some definitional difficulties. We show that an ability to form a “cat state” with  $n$  qubits, or fan out a qubit into  $n$  copies in constant depth, is equivalent to being able to construct an  $n$ -ary parity gate in constant depth. We discuss how best to compare these circuit classes to classical ones.

We prove the surprising result that, for any integer  $q > 1$ ,  $QAC_{wf}^0 = QACC[q] = QACC$ . This is in sharp contrast to the classical result of Smolensky [30] that says  $ACC^0[q] \neq ACC^0[p]$  for any pair of distinct primes  $q, p$ , which implies that for any prime  $p$ ,  $AC^0 \subset ACC^0[p] \subset ACC$ . This result shows that parity gates are as powerful as any other mod gates in QACC, and more generally, that any  $MOD_q$  gate is as good as any other, up to polynomial size and constant depth. Thus we conclude that  $QAC_{wf}^0$ , or, for any  $q$ ,  $QACC[q]$ , is strictly more powerful than  $ACC[q]$  and  $AC^0$ .

We also develop methods for proving upper bounds for QACC. The definition of QACC immediately leads to a problem in this regard: QACC is a class of operators that only have a natural interpretation quantum mechanically. In order to clarify the relationship with classical computation we assign properties to QACC circuits based on measurements we can perform on them. In particular, we define several natural languages classes related to QACC. These language classes arise from considering a quantum circuit family in the class and specifying a condition on the expectation of observing a particular state after applying a circuit from the family to an input state. The condition might simply be that the expectation is non-zero, or that it is bounded away from zero by some constant, or that it is exactly equal to some constant. We call the language classes obtained by these conditions on the expectation NQACC, BQACC and EQACC, respectively. For example, the class NQACC corresponds to the case where  $x$  is in the language if the expectation of the observed state after applying the QACC operator is non-zero. This is analogous to the definition of the class NQP as defined in Adleman et al. [1] and discussed in Fenner et al. [12]. In this way we obtain natural classes of languages which correspond to those defined classically by families of small depth circuits. In these terms, for example, we can more succinctly and precisely express the statement “QACC[q] is strictly more powerful than ACC[q]” by writing  $ACC[q] \subset EQACC[q]$ .

We desire upper bounds showing that these language classes are contained in classically defined circuit classes, thus delimiting the power of these quantum computations. In particular, we

believe that the languages arising in this way from our definitions are contained within  $TC^0$ , those problems computed by constant-depth threshold circuits. We have been unable to verify this, and in fact the only classical upper bound for these language classes that we know of is the very powerful counting class  $coC=P$  (see [12]). We do give some evidence for this proposed  $TC^0$  upper bound here and further provide some techniques which may prove useful in solving this problem.

Our methods result in upper bounds for restricted QACC circuits. Roughly speaking, we show that QACC is no more powerful than  $P/poly$  provided that a layer of “wire-crossings” in the QACC operator can be written as log many compositions of Kronecker products of controlled-not gates. We call this class  $QACC_{pl}^{log}$ , where the “pl” is for this planarity condition. We show if one further restricts attention to the case where the number of multi-line gates (gates whose input is more than 1 qubit) is log-bounded then the circuits are no more powerful than  $TC^0$ . We call this class  $QACC_{gates}^{log}$ . These results hold for arbitrary complex amplitudes in the QACC circuits.

In terms of our language classes, we show that  $NQACC_{gates}^{log}$  is in  $TC^0$  and  $NQACC_{pl}^{log}$  is in  $P/poly$ . Although the proof uses some of the techniques developed by Fenner, Green, Homer and Pruim [12] and by Yamakami and Yao [31] to show that  $NQP_C = coC=P$ , the small depth circuit case presents technical challenges not present in their setting. In particular, given a QACC operator built out of layers  $M_1, \dots, M_t$  and an input state  $|x, 0^{p(n)}\rangle$ , we must show that a  $TC^0$  circuit can keep track of the amplitudes of each possible resulting state as each layer is applied. After all layers have been applied, the  $TC^0$  circuit then needs to be able to check that the amplitude of one possible state is non-zero. Unfortunately, there could be exponentially many states with non-zero amplitudes after applying a layer. To handle this problem we introduce the idea of a “tensor-graph,” a new way to represent a collection of states. We can extract from these graphs (via  $TC^0$  or  $P/poly$  computations) whether the amplitude of any particular vector is non-zero.

The exponential growth in the number of states is one of the primary obstacles to proving that all of  $NQACC$  is in  $TC^0$  (or even  $P/poly$ ), and thus the tensor graph formalism represents a significant step towards such an upper bound. The reason the bounds apply only in the restricted cases is that although tensor graphs can represent any QACC operator, in the case of operators with layers that might do arbitrary permutations, the top-down approach we use to compute a desired amplitude from the graph no longer seems to work. We feel that it is likely that the amplitude of any vector in a tensor graph can be written as a polynomial product of a polynomial sum in some extension algebra of the ones we work with in this paper, in which case it is quite likely it can be evaluated in  $TC^0$ .

Another important obstacle to obtaining a  $TC^0$  upper bound is that one needs to be able to add and multiply a polynomial number of complex amplitudes that may appear in a QACC computation. We solve this problem. It reduces to adding and multiplying polynomially many elements of a certain transcendental extension of the rational numbers. We show that in fact  $TC^0$  is closed under iterated addition and multiplication of such numbers (Lemma 5.1 below). This result is of independent interest, and our application of tensor-graphs and these closure properties of  $TC^0$  may prove useful in further investigations of small-depth quantum circuits.

We now discuss the organization of the rest of this paper. Section 2 contains definitions for the quantum operator classes we will be considering as well as other background definitions. Section 3 shows the constant-depth quantum circuit equivalence of fan-out and parity gates. Section 4 establishes for arbitrary  $p$  and  $q$  the constant-depth quantum equivalence of  $Mod_p$  and  $Mod_q$ . Section 5 contains our upper bound results. Finally, the last section has a conclusion and some

open problems.

Preliminary versions of these results appeared in [16] and [15].

## 2 Preliminaries

In this section we define the gates used as building blocks for our quantum circuits. Classes of operators built out of these gates are then defined. We define language classes that can be determined by these operators and give a couple of definitions from algebra. Lastly, some closure properties of  $\text{TC}^0$  are described.

**Definition 2.1** *We define various quantum gates as follows:*

- *By a one-qubit gate we mean an operator from the group  $U(2)$ .*
- *Let  $U = \begin{pmatrix} u_{00} & u_{01} \\ u_{10} & u_{11} \end{pmatrix} \in U(2)$ .  $\wedge_m(U)$  is defined as:  $\wedge_0(U) = U$  and for  $m > 0$ ,  $\wedge_m(U)$  is*


$$\wedge_m(U)(|\vec{x}, y\rangle) = \begin{cases} u_{y0}|\vec{x}, 0\rangle + u_{y1}|\vec{x}, 1\rangle & \text{if } \wedge_{k=1}^m x_k = 1 \\ |\vec{x}, y\rangle & \text{otherwise} \end{cases}$$

- *Let  $X = \sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ . A Toffoli gate is a  $\wedge_m(X)$  gate for some  $m \geq 0$ . A controlled-not gate is a  $\wedge_1(X)$  gate.*
- *The Hadamard gate is the one-qubit gate  $H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$ .*
- *An (m-)spaced controlled-not gate is an operator that maps  $|y_1, \dots, y_m, x\rangle$  to  $|x \oplus y_1, y_2, \dots, y_m, x\rangle$  or  $|x, y_1, \dots, y_m\rangle$  to  $|x, y_1, \dots, y_{m-1}, y_m \oplus x\rangle$*
- *An (m-ary) fan out gate  $F$  is an operator that maps  $|y_1, \dots, y_m, x\rangle$  to  $|x \oplus y_1, \dots, x \oplus y_m, x\rangle$ .*
- *The classical Boolean  $\text{Mod}_q$ -function on  $n$  bits is defined so that  $\text{Mod}_q(x_1, \dots, x_n) = 1$  iff  $\sum_{i=1}^n x_i \not\equiv 0 \pmod{q}$ . We also define  $\text{Mod}_{q,r}(x_1, \dots, x_n)$  to output 1 iff  $\sum_{i=1}^n x_i \equiv r \pmod{q}$ . A quantum  $\text{MOD}_q$  gate is an operator that maps  $|y_1, \dots, y_m, x\rangle$  to  $|y_1, \dots, y_m, x \oplus \text{Mod}_q(y_1, \dots, y_m)\rangle$ . A quantum  $\text{MOD}_{q,r}$  gate maps  $|y_1, \dots, y_m, x\rangle$  to  $|y_1, \dots, y_m, x \oplus \text{Mod}_{q,r}(y_1, \dots, y_m)\rangle$ . We write  $\neg\text{MOD}_q$  for  $\text{MOD}_{q,0}$ . A parity gate is a  $\text{MOD}_2$  gate.*

Note that, since negation is built into the output (via the exclusive OR), it is easy to simulate negations using  $\text{MOD}_{q,r}$  gates (unlike the classical case). For example, by setting  $b = 1$ , we can compute  $\neg\text{Mod}_{q,r}$ . More generally, using one work bit, it is possible to simulate “ $\neg\text{MOD}_{q,r}$ ,” defined so that,

$$|x_1, \dots, x_n, b\rangle \mapsto |x_1, \dots, x_n, b \oplus (\neg\text{Mod}_{q,r}(x_1, \dots, x_n))\rangle$$

using just  $\text{MOD}_{q,r}$  and a controlled-not gate. Thus  $\text{MOD}_{q,r}$  and  $\neg\text{MOD}_{q,r}$  are equivalent up to constant depth. Finally, observe that  $\text{MOD}_{q,r}^{-1} = \text{MOD}_{q,r}$ .



**Figure 1.** Our notation for  $n$ -ary Toffoli, controlled- $U$ , and  $\text{MOD}_q$  gates, fanout gates, symmetric phase shift gates, and the Hadamard gate. On the top right, we show a useful identity between the controlled-not, the controlled  $\pi$ -shift, and the Hadamard gate. On the bottom right we show a controlled- $U$  gate with one of its inputs negated by conjugation with  $X$ .

We will use the notation in Figure 1 for our various gates.

As discussed in further detail in section 3 below, the no-cloning theorem of quantum mechanics makes it difficult to directly fan out qubits in constant depth (although constant fan-out in constant depth is no problem, since we can make multiple copies of the inputs). Thus it is necessary to define the operator  $F$  as in the above definition. Also, in the literature it is frequently the case that one says a given operator  $M$  on  $|y_1, \dots, y_m\rangle$  can be written as a tensor product of certain gates  $M_j$ . What is meant is that there is an permutation operator  $\Pi$  ( a map from  $|y_1, \dots, y_m\rangle$  to  $|y_{\pi(1)}, \dots, y_{\pi(m)}\rangle$ ) for some permutation  $\pi$ ) such that

$$M|y_1, \dots, y_m\rangle = \Pi \otimes_j^n M_j \Pi^{-1} |y_1, \dots, y_m\rangle$$

where the  $M_j$ 's are our base gates, i.e. those gates for which no inherent ordering on the  $y_i$  is assumed *a priori*, and  $\otimes$  is the Kronecker product, which flattens a tensor product into a matrix with blocks indexed in a particular way. Since it is important to keep track of such details in our upper bounds proofs, we will always use Kronecker products of the form  $\otimes_j^n M_j$  without unspoken permutations. Nevertheless, being able to do permutation operators (not conjugation by a permutation) intuitively allows our circuits to simulate classical wire crossings. To handle permutations, we allow our circuits to have controlled-not layers. A *controlled-not layer* is a gate which performs, in one step, controlled-not's between an arbitrary collection of disjoint pairs of lines in its domain. That is, it performs  $\Pi \otimes_j^n \wedge_1(X) \Pi^{-1}$  for some permutation operator  $\Pi$ . It is easy to see [18] that any permutation can be written as a product of a constant number of controlled-not layers. We say a controlled-not layer is *log-depth* if it can be written as the composition of log many matrices each of which is the Kronecker product of identities and spaced controlled-not gates.

$M^{\otimes n}$  is the  $n$ -fold Kronecker product of  $M$  with itself.

**Definition 2.2**

$\text{QAC}^k$  is the class of families  $\{F_n\}$ , where  $F_n$  is in  $U(2^{n+p(n)})$ ,  $p$  a polynomial, and each  $F_n$  is writable as a product of  $\mathcal{O}(\log^k n)$  layers, where a layer is a Kronecker product of one-qubit gates

and Toffoli gates or is a controlled-not layer. Also for all  $n$  the number of distinct types of one qubit gates used must be fixed.

$\text{QACC}^k[q]$  is the same as  $\text{QAC}^k$  except we also allow  $\text{MOD}_q$  gates.  $\text{QACC}^k = \cup_q \text{QACC}^k[q]$ .

$\text{QAC}_{wf}^k$  is the same as  $\text{QAC}^k$  but we also allow fan-out gates.

$\text{QACC}$  is defined as  $\text{QACC}^0$  and  $\text{QACC}[q]$  is defined as  $\text{QACC}^0[q]$ .  $\text{QACC}_{pl}^{\log}$  is  $\text{QACC}$  restricted to log-depth controlled not layers.  $\text{QACC}_{\text{gates}}^{\log}$  is  $\text{QACC}$  restricted so that the total number of multi-line gates in all layers is log-bounded.

If  $\mathcal{C}$  is one of the above classes and  $K \subseteq \mathbf{C}$ , then  $\mathcal{C}_K$  are the families in  $\mathcal{C}$  with coefficients restricted to  $K$ .

Let  $\{F_n\}$  and  $\{G_n\}$ ,  $G_n, F_n \in U(2^n)$  be families of operators. We say  $\{F_n\}$  is  $\text{QAC}^0$  reducible to  $\{G_n\}$  if there is a family  $\{R_n\}$ ,  $R_n \in U(2^{n+p(n)})$  of  $\text{QAC}^0$  operators augmented with operators from  $\{G_n\}$  such that for all  $n$ ,  $\mathbf{x}, \mathbf{y} \in \{0, 1\}^n$ , there is a setting of  $z_1, \dots, z_{p(n)} \in \{0, 1\}$  for which  $\langle \mathbf{y} | F_n | \mathbf{x} \rangle = \langle \mathbf{y}, \mathbf{z} | R_n | \mathbf{x}, \mathbf{z} \rangle$ . Operator families are  $\text{QAC}^0$  equivalent if they are  $\text{QAC}^0$  reducible to each other. If  $\mathcal{C}_1$  and  $\mathcal{C}_2$  are families of  $\text{QAC}^0$  equivalent operators, we write  $\mathcal{C}_1 = \mathcal{C}_2$ .

We refer to the  $z_i$ 's above as ‘‘work bits’’ (also called ‘‘ancillae’’ in [18]). Note that in proving  $\text{QAC}^0$  equivalence, the work bits must be returned to their original values in a computation so that they are disentangled from the rest of the circuit, and can be re-used by subsequent layers.

It follows for any  $\{F_n\} \in \text{QAC}^0$  that  $F_n$  is writable as a product of finite number of layers. In an earlier paper, Moore [16] places no restriction on the number of distinct types of one-qubit gates used in a given family of operators. Here we restrict these so that the number of distinct amplitudes which appear in matrices in a layer is fixed with respect to  $n$ . This restriction arises implicitly in the quantum Turing machine case of the upper bounds proofs in Fenner, et al. [12] and Yamakami and Yao [31]. Also, it seems fairly natural since in the classical case one builds circuits using a fixed number of distinct gate types. Our classes here are, thus, more ‘‘uniform’’ than those defined earlier [16]. We now define language classes based on our classes of operator families.

**Definition 2.3** Let  $\mathcal{C}$  be a class of families of  $U(2^{n+p(n)})$  operators where  $p$  is a polynomial and  $n = |x|$ .

1.  $E\mathcal{C}$  is the class of languages  $L$  such that for some  $\{F_n\} \in \mathcal{C}$  and  $\{|\vec{z}_n\rangle\} = \{|\langle z_{n,1}, \dots, z_{n,n+p(n)} \rangle|\}$  a family of states,  $m := |\langle \vec{z}_n | F_n | x, 0^{p(n)} \rangle|^2$  is 1 or 0 and  $x \in L$  iff  $m = 1$ .
2.  $N\mathcal{C}$  is the class of languages  $L$  such that for some  $\{F_n\} \in \mathcal{C}$  and  $\{|\vec{z}_n\rangle\}$  a family of states,  $x \in L$  iff  $|\langle \vec{z}_n | F_n | x, 0^{p(n)} \rangle|^2 > 0$ .
3.  $B\mathcal{C}$  is the class of languages  $L$  where for  $\{F_n\} \in \mathcal{C}$  and  $\{|\vec{z}\rangle\}$ ,  $x \in L$  if  $|\langle \vec{z}_n | F_n | x, 0^{p(n)} \rangle|^2 > 3/4$  and  $x \notin L$  if  $|\langle \vec{z}_n | F_n | x, 0^{p(n)} \rangle|^2 < 1/4$ .

It follows  $E\mathcal{C} \subseteq N\mathcal{C}$  and  $E\mathcal{C} \subseteq B\mathcal{C}$ . We frequently will omit the ‘.’ when writing a class, so  $E\text{QACC}$  is written as  $\text{EQACC}$ . Let  $|\Psi\rangle := |F_n | x, 0^{p(n)}\rangle$ . Notice that  $|\langle \vec{z}_n | F_n | x, 0^{p(n)} \rangle|^2 = \langle \Psi | P_{|\vec{z}_n\rangle} | \Psi \rangle$ , where  $P_{|\vec{z}_n\rangle}$  is the projection matrix onto  $|\vec{z}_n\rangle$ . We could allow in our definitions measurements of up to polynomially many such projection observables and not affect our results below. However,

this would shift the burden of the computation in some sense away from the QACC operator and instead onto preparation of the observable.

Next are some variations on familiar definitions from algebra.

**Definition 2.4** *Let  $k > 0$ . A subset  $\{\beta_i\}_{1 \leq i \leq k}$  of  $\mathbf{C}$  is linearly independent if  $\sum_{i=1}^k a_i \beta_i \neq 0$  for any  $(a_1, \dots, a_k) \in \mathbf{Q}^k - \{\vec{0}^k\}$ . A set  $\{\beta_i\}_{1 \leq i \leq k}$  is algebraically independent if the only  $p \in \mathbf{Q}[x_1, \dots, x_k]$  with  $p(\beta_1, \dots, \beta_k) = 0$  is the zero polynomial.*

We now briefly mention some closure properties of  $\text{TC}^0$  computable functions that are useful in proving  $\text{NQACC}_{\text{gates}}^{\log} \subseteq \text{TC}^0$ . For proofs of the statements in the next lemma see [28, 29, 10].

**Lemma 2.5** (1)  $\text{TC}^0$  functions are closed under composition. (2) The following are  $\text{TC}^0$  computable:  $x + y$ ,  $x \div y := x - y$  if  $x - y > 0$  and 0 otherwise,  $|x| := \lceil \log_2(x + 1) \rceil$ ,  $x \cdot y$ ,  $\lfloor x/y \rfloor$ ,  $2^{\min(i, p(|x|))}$ , and  $\text{cond}(x, y, z) := y$  if  $x > 0$  and  $z$  otherwise. (3) If  $f(i, x)$  is  $\text{TC}^0$  computable then  $\sum_{k=0}^{p(|x|)} f(k, x)$ ,  $\prod_{k=0}^{p(|x|)} f(k, x)$ ,  $\forall i \leq p(|x|) (f(i, x) = 0)$ ,  $\exists i \leq p(|x|) (f(i, x) = 0)$ , and  $\mu_{i \leq p(|x|)} (f(i, x) = 0) :=$  the least  $i$  such that  $f(i, x) = 0$  and  $i \leq p(|x|)$  or  $p(x) + 1$  otherwise, are  $\text{TC}^0$  computable.


We drop the min from the  $2^{\min(i, p(|x|))}$  when it is obvious a suitably large  $p(|x|)$  can be found. We define  $\text{max}(x, y) := \text{cond}(1 \div (y \div x), x, y)$  and define

$$\text{max}_{i \leq p(|x|)} (f(i)) := \mu_{i \leq p(|x|)} (\forall j \leq p(|x|) (f(j) \div f(i) = 0))$$

Using the above functions we describe a way to do sequence coding in  $\text{TC}^0$ . Let  $\beta_{|t|}(x, w) := \lfloor (w - \lfloor w/2^{(x+1)|t|} \rfloor \cdot 2^{(x+1)|t|}) / 2^{x|t|} \rfloor$ . The function  $\beta_{|t|}$  is useful for block coding. Roughly,  $\beta_{|t|}$  first gets rid of the bits after the  $(x+1)|t|$ th bit then chops off the low order  $x|t|$  bits. Let  $B = 2^{\lfloor \max(x, y) \rfloor}$ , so that  $B$  is longer than either  $x$  or  $y$ . Hence, we code pairs as  $\langle x, y \rangle := (B + y) \cdot 2B + B + x$ , and projections as  $(w)_1 := \beta_{\lfloor \frac{1}{2}|w| \rfloor - 1}(0, \beta_{\lfloor \frac{1}{2}|w| \rfloor}(0, w))$  and  $(w)_2 := \beta_{\lfloor \frac{1}{2}|w| \rfloor - 1}(0, \beta_{\lfloor \frac{1}{2}|w| \rfloor}(1, w))$ . We can encode a polynomial,  $\text{TC}^0$  computable sequence of numbers  $\langle f(1), \dots, f(k) \rangle$  as the pair  $\langle \sum_i^k (f(i) 2^{i \cdot m}), m \rangle$  where  $m := \lfloor f(\text{max}_i(f(i))) \rfloor + 1$ . We then define the function which projects out the  $i$ th member of a sequence as  $\beta(i, w) := \beta_{(w)_2}(i, w)$ .

We can code integers using the positive natural numbers by letting the negative integers be the odd natural numbers and the positive integers be the even natural numbers.  $\text{TC}^0$  can use the  $\text{TC}^0$  circuits for natural numbers to compute both the polynomial sum and polynomial product of a sequence of  $\text{TC}^0$  definable integers. It can also compute the rounded quotient of two such integers. For instance, to do a polynomial sum of integers, compute the natural number which is the sum of the positive numbers in the sum using  $\text{cond}$  and our natural number iterated addition circuit. Then compute the natural number which is the sum of the negative numbers in the sum. Use the subtraction circuit to subtract the smaller from the larger number and multiply by two. One is then added if the number should be negative. For products, we compute the product of the natural numbers which results by dividing each integer code by two and rounding down. We multiply the result by two. We then sum the number of terms in our product which were negative integers. If this number is odd we add one to the product we just calculated. Finally, division can be computed using the Taylor expansion of  $1/x$ .





**Figure 2. Two ways to make a cat state on  $n$  qubits. The circuit on the left uses only two-qubit gates and has depth  $\log n$ . On the right, we define a “fanout gate” that simultaneously performs  $n$  controlled-nots from one input qubit.**

### 3 Fanout, Cat States, and Parity

To make a shallow parallel circuit, it is often important to *fan out* one of the inputs into multiple copies. One of the differences between classical circuits and quantum ones as we have defined them here is that in classical circuits, we usually assume that we get arbitrary fanout for free, simply by splitting a wire into as many copies as we like. This is difficult in quantum circuits, since making an unentangled copy requires non-unitary, and in fact non-linear, processes:

$$(\alpha|0\rangle + \beta|1\rangle) \otimes (\alpha|0\rangle + \beta|1\rangle) = \alpha^2|00\rangle + \alpha\beta(|01\rangle + |10\rangle) + \beta^2|11\rangle$$

has coefficients quadratic in  $\alpha$  and  $\beta$ , so it cannot be derived from  $\alpha|0\rangle + \beta|1\rangle$  using any linear operator, let alone a unitary one. This is one form of the so-called “no cloning” theorem.

However, the controlled-not gate can be used to copy a qubit onto a work bit in the pure state  $|0\rangle$  by making a non-destructive measurement:


$$(\alpha|0\rangle + \beta|1\rangle) \otimes |0\rangle \rightarrow \alpha|00\rangle + \beta|11\rangle$$

Note that the final state is not a tensor product of two independent qubits, since the two qubits are completely entangled. This means that whatever we do to one copy, we do to the other. Except when the states are purely Boolean, we have to treat this kind of “fanout” more gingerly than we would in the classical case.

By making  $n$  copies of a qubit in this sense, we can make a “cat state”  $\alpha|000 \cdots 0\rangle + \beta|111 \cdots 1\rangle$ . Such states are useful in making quantum computation fault-tolerant (e.g. [9, 26]). We can do this in  $\log n$  depth with controlled-not gates, as shown on the left-hand side of Figure 2. When preceded by a Hadamard gate on the top qubit, this circuit will map an initial state  $|0000\rangle$  onto a cat state  $\frac{1}{\sqrt{2}}(|0000\rangle + |1111\rangle)$ . However, we will also consider circuits which can do this in a single layer, with a “fanout gate” that simultaneously copies a qubit onto  $n$  target qubits. This is simply the product of  $n$  controlled-not gates, as shown on the right-hand side of Figure 2.

We now show that in quantum circuits, we can do fanout in constant depth if and only if we can construct a parity gate in constant depth.

**Proposition 3.1** *In any class of quantum circuits that includes Hadamard and controlled-not gates, the following are equivalent:*



**Figure 3.** The parity and fanout gates are conjugates of each other by a layer of Hadamard gates.

1. It is possible to map  $\alpha|0\rangle + \beta|1\rangle$  and  $n - 1$  work bits in the state  $|0\rangle$  onto an  $n$ -qubit cat state  $\alpha|000\cdots 0\rangle + \beta|111\cdots 1\rangle$  in constant depth.
2. The  $n$ -ary fanout gate on the right-hand side of Figure 2 can be implemented in constant depth with at most  $n - 1$  additional work bits.
3. An  $n$ -ary parity or  $\text{MOD}_2$  gate as defined above can be implemented in constant depth with at most  $n - 1$  additional work bits.

**Proof.** First, note that (1) is *a priori* weaker than (2), since (1) only requires that an operator map  $|100\cdots 0\rangle$  to  $|111\cdots 1\rangle$  and  $|000\cdots 0\rangle$  to itself. In fact, the two circuits shown in Figure 2 both do this, even though they differ on other initial states.


To prove  $(2 \Leftrightarrow 3)$ , we simply need to notice that the parity gate is a fanout gate going the other way conjugated by a layer of Hadamard gates, since parity is simply a product of controlled-nots with the same target qubit, and conjugating with  $H$  reverses the direction of a controlled-not. This is shown in Figure 3. Clearly the number of work bits used to perform either gate will be the same. (We prove this equivalence in greater detail and generality in Proposition 4.2 below.)

To prove  $(1 \Rightarrow 3)$ , we use a slightly more elaborate circuit shown in Figure 4. Here we use the identity shown in Figure 1 to convert the parity gate into a product of controlled  $\pi$ -shifts. Since these are diagonal, they can be parallelized as in [18] by copying the target qubit onto  $n - 1$  work bits, and applying each one to a different copy. While we have drawn the circuit with two fanout gates, any gate that satisfies the conditions in (1), and its inverse after the  $\pi$ -shifts, will do.

Finally,  $(2 \Rightarrow 1)$  is obvious. □

This brings up an interesting issue. It is not clear that a  $\text{QAC}^0$  operator as we have defined  $\text{QAC}^0$  here can “simulate” any  $\text{AC}^0$  circuit, since we are not allowing arbitrary fanout in each layer. An alternate definition, which we might call  $\text{QAC}$  with fanout or  $\text{QAC}_{\text{wf}}^0$ , would allow us to perform controlled- $U$  gates or Toffoli gates in the same layer whenever they have different target qubits, even if their input qubits overlap. This seems reasonable, since these gates commute. Since we can fan out to  $n$  copies in  $\log n$  layers as in Figure 2, we have  $\text{QAC}^k \subseteq \text{QAC}_{\text{wf}}^k \subseteq \text{QAC}^{(k+1)}$ . We can define  $\text{QACC}_{\text{wf}}^k$  in the same way, and Proposition 3.1 implies that  $\text{QAC}_{\text{wf}}^k = \text{QACC}_{\text{wf}}^k[2] = \text{QACC}^k[2]$ .

It is partly a matter of taste whether  $\text{QAC}^0$  or  $\text{QAC}_{\text{wf}}^0$  is a better analog of  $\text{AC}^0$ . However, fanout does seem possible in several proposed technologies for quantum computing. In an ion trap computer [8], vibrational modes can couple with all the atoms simultaneously, so we could apply a controlled-not from one atom to the “bus qubit” and then from the bus to the other  $n$  atoms. In bulk-spin NMR [14], we can activate the couplings from one atom to  $n$  others, and perform  $n$  controlled  $\pi$ -shifts simultaneously, which is equivalent to fanout with the target qubits



**Figure 4.** The parity gate can also be written as a product of controlled  $\pi$ -shifts, with the target qubit conjugated by  $H$ . Since these are diagonal, we can parallelize them using any gate that can make a cat state .

conjugated with the Hadamard gate. Thus allowing fanout may in fact be the most reasonable model of constant-depth quantum circuits.

## 4 Constant Depth Equivalence of $\text{MOD}_p$ and $\text{MOD}_q$ Gates

As stated in the Introduction, in the classical case  $\text{Mod}_p$  and  $\text{Mod}_q$  gates are not easy to build from each other whenever  $p$  and  $q$  are relatively prime. In fact, to do it in constant depth requires a circuit of exponential size [30]. In this section, we will show this is not true in the quantum case. Specifically, we show that any  $\text{MOD}_q$  gate can be built in constant depth from any  $\text{MOD}_p$  gate, for any two numbers  $p$  and  $q$ . We start by showing that any  $\text{MOD}_q$  gate can be built from parity gates in constant depth.

**Proposition 4.1** *In any circuit class containing  $n$ -ary parity gates and one-qubit gates, we can construct an  $n$ -ary  $\text{MOD}_q$  gate, with  $\mathcal{O}(n \log q)$  work bits, in depth depending only on  $q$ .*

**Proof.** Let  $k = \lceil \log_2 q \rceil$ , and let  $M$  be a Boolean matrix on  $k$  qubits where the zero state has period  $q$ . For instance, if we write  $|x\rangle$  as shorthand for  $|x_{k-1} \cdots x_1 x_0\rangle$  where  $x_i$  is the  $2^i$  digit of  $x$ 's binary expansion and  $0 \leq x < 2^k$ , we can define  $M$  so that it permutes the  $|x\rangle$  as follows:

$$M|x\rangle = \begin{cases} |(x+1) \bmod q\rangle & \text{if } x < q \\ |x\rangle & \text{if } x \geq q \end{cases}$$

Then if we start with  $k$  work bits in the state  $|0\rangle$  and apply a controlled- $M$  gate to them from each input, the state will differ from  $|0\rangle$  on at least one qubit if and only if the number of true inputs is not a multiple of  $q$ . (Note that this controlled- $M$  gate applies to  $k$  target qubits at once in an entangled way.) We can then apply an  $n$ -ary OR of these  $k$  qubits to the target qubit, i.e. a Toffoli gate with its inputs conjugated with  $X$  and its target qubit negated before or after the gate. We end by applying the inverse series of controlled- $M^\dagger$  gates to return the  $k$  work bits to  $|0\rangle$ .

Now we use Proposition 4 of [18] to parallelize this set of controlled- $M$  gates. We can convert them to diagonal gates by conjugating the  $k$  qubits with a unitary operator  $T$ , where  $T^\dagger D T = M$

and  $D$  is diagonal. If we have a parity gate, we can fan out the  $k$  work bits to  $n$  copies each using Proposition 3.1. We can then simultaneously apply the  $n$  controlled- $D$  gates from each input to the corresponding copy, and then uncopy them back.

This is shown in Figure 5. For  $q = 3$ , for instance,  $M = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \\ & & & 1 \end{pmatrix}$ ,  $T = \frac{1}{\sqrt{3}} \begin{pmatrix} & & & \sqrt{3} \\ 1 & & & \\ e^{4\pi i/3} & e^{2\pi i/3} & 1 & \\ e^{2\pi i/3} & e^{4\pi i/3} & 1 & \end{pmatrix}$ ,

and  $D = \begin{pmatrix} 1 & & & \\ & 1 & & \\ & & e^{2\pi i/3} & \\ & & & e^{4\pi i/3} \end{pmatrix}$ .

The operators  $T$ ,  $T^\dagger$ , and the controlled- $D$  gate can be carried out in some finite depth by controlled-nots and one-qubit gates by the results of [3]. The total depth of our  $\text{MOD}_q$  gate is a function of these and so of  $q$ , but not of  $n$ . Finally, the number of work bits used is  $(n - 1)k = \mathcal{O}(n \log q)$  as promised.  $\square$

To look more closely at the depth as a function of  $q$ , we note that using the methods of Reck et al. [23] and Barenco et al. [3], any operator on  $k$  qubits can be performed with  $\mathcal{O}(k^3 4^k)$  two-qubit gates. Since  $k = \lceil \log_2 q \rceil$ , this means that the depths of  $T$ ,  $T^\dagger$  and the controlled- $D$  gates are at most  $\mathcal{O}(q^2 \log^3 q)$ .

Since we can construct  $\text{MOD}_q$  gates in constant depth, we have  $\text{QACC}^k[q] \subset \text{QACC}^k[2]$  for all  $q$ , so  $\text{QACC}^k = \text{QACC}^k[2]$ . By Proposition 3.1, these are both also equal to  $\text{QACC}_{\text{wf}}^k$ . In particular, we have

$$\text{QACC}_{\text{wf}}^0 = \text{QACC}[2] = \text{QACC}$$

while classically both equalities are strict inclusions. Note that allowing fanout immediately gives  $\text{QACC}_{\text{wf}}[q] = \text{QACC}[2] = \text{QACC}$  for any  $q$ , but we will show below that including fanout explicitly is not necessary.

We now show the converse  $\text{QACC}[2] \subseteq \text{QACC}[q]$  for any  $q$ , i.e. parity can be built from  $\text{MOD}_q$  for any  $q$ . This shows the constant depth equivalence of  $\text{MOD}_q$  gates for all  $q$ .

Let  $q \in \mathbf{N}$ ,  $q \geq 2$  be fixed for the remainder of this section. Consider quantum states labeled by digits in  $D = \{0, \dots, q - 1\}$ . By analogy with ‘‘qubit,’’ we refer to a state of the form,

$$\sum_{k=0}^{q-1} c_k |k\rangle$$

with  $\sum_k |c_k|^2 = 1$  as a ‘‘qudigit.’’

We define three important operations on qudigits. The  $n$ -ary *modular addition* operator  $M_q$  acts as follows:



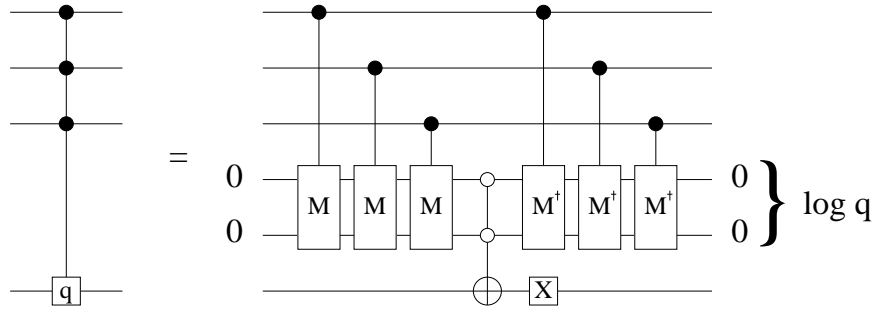
$$M_q |x_1, \dots, x_n, b\rangle = |x_1, \dots, x_n, (b + x_1 + \dots + x_n) \bmod q\rangle$$

We use the same graphical notation for  $M_q$  as we do for a  $\text{MOD}_q$  gate, but interpreting the lines as qudigits, as illustrated in Figure 6.


Since  $M_q$  merely permutes the states, it is clear that it is unitary. Similarly, the  $n$ -ary unitary *base  $q$  fanout* operator  $F_q$  acts as,

$$F_q |x_1, \dots, x_n, b\rangle = |(x_1 + b) \bmod q, \dots, (x_n + b) \bmod q, b\rangle$$

We write  $F$  for  $F_2$ , since it is the ‘‘standard’’ fan-out gate introduced in Definition 2.1. Note that  $M_q^{-1} = M_q^{q-1}$  and  $F_q^{-1} = F_q^{q-1}$ .



**Figure 5. Building a  $\text{Mod}_q$  gate.** We choose a matrix  $M$  on  $k = \lceil \log_2 q \rceil$  qubits such that  $M^q = 1$ , apply controlled- $M$  gates from the  $n$  inputs to  $k$  work bits, apply an OR from these  $k$  qubits to the target qubit, and reverse the process to return the work bits to  $|0\rangle$ . To parallelize this, we can diagonalize  $M$  by writing it as  $T^\dagger DT$ , fan the  $k$  qubits out into  $n$  copies each using Proposition 3.1, and apply controlled- $D$  gates simultaneously from each input to a set of copies. The total depth depends on  $q$  but not on  $n$ .



**Figure 6. An  $M_q$  gate.**

Finally, the Quantum Fourier Transform  $H_q$  (which generalizes the Hadamard transform  $H$  on qubits) acts on a single qudigit as,

$$H_q|a\rangle = \frac{1}{\sqrt{q}} \sum_{b=0}^{q-1} \zeta^{ab} |b\rangle$$

where  $\zeta = e^{\frac{2\pi i}{q}}$  is a primitive complex  $q^{\text{th}}$  root of unity. It is easy to see that  $H_q$  is unitary, via the fact that  $\sum_{\ell=0}^{q-1} \zeta^{a\ell} = 0$  iff  $a \not\equiv 0 \pmod{q}$ .

The first observation is that, analogous to parity and fanout for Boolean inputs, the operators  $M_q$  and  $F_q$  are “conjugates” in the following sense. This is a generalization of the equivalence of assertions (2) and (3) of Proposition 3.1.

**Proposition 4.2**  $M_q = (H_q^{\otimes(n+1)})^{-1} F_q^{-1} H_q^{\otimes(n+1)}$ .

**Proof.** We apply the operators  $H_q^{\otimes(n+1)}$ ,  $F_q^{-1}$ , and  $(H_q^{\otimes(n+1)})^{-1}$  in that order to the state  $|x_1, \dots, x_n, b\rangle$ , and check that the result has the same effect as  $M_q$ .

The operator  $H_q^{\otimes(n+1)}$  simply applies  $H_q$  to each of the  $n + 1$  qudigits of  $|x_1, \dots, x_n, b\rangle$ , which yields,

$$\frac{1}{q^{\frac{(n+1)}{2}}} \sum_{\mathbf{y} \in D^n} \sum_{a=0}^{q-1} \zeta^{\mathbf{x} \cdot \mathbf{y} + ab} |y_1, \dots, y_n, a\rangle,$$

where  $\mathbf{y}$  is a compact notation for  $y_1, \dots, y_n$ , and  $\mathbf{x} \cdot \mathbf{y}$  denotes  $\sum_{i=1}^n x_i y_i$ . Then applying  $F_q^{-1}$  to the above state yields,

$$\frac{1}{q^{\frac{(n+1)}{2}}} \sum_{\mathbf{y} \in D^n} \sum_{a=0}^{q-1} \zeta^{\mathbf{x} \cdot \mathbf{y} + ab} |(y_1 - a) \bmod q, \dots, (y_n - a) \bmod q, a\rangle.$$

By a change of variable, the above can be re-written as,

$$\frac{1}{q^{\frac{(n+1)}{2}}} \sum_{\mathbf{y} \in D^n} \sum_{a=0}^{q-1} \zeta^{\sum_{i=1}^n x_i (y_i + a) + ab} |y_1, \dots, y_n, a\rangle$$

Finally, applying  $(H_q^{\otimes(n+1)})^{-1}$  to the above undoes the Fourier transform and puts the coefficient of  $a$  in the exponent into the last slot of the state.

The result is,

$$(H_q^{\otimes(n+1)})^{-1} F_q^{-1} H_q^{\otimes(n+1)} |x_1, \dots, x_n, b\rangle = |x_1, \dots, x_n, (b + x_1 + \dots + x_n) \bmod q\rangle,$$

which is exactly what  $M_q$  would yield.  $\square$

We now describe how the operators  $M_q$ ,  $F_q$  and  $H_q$  can be modified to operate on registers consisting of qubits rather than qudits. Firstly, we encode each digit using  $\lceil \log q \rceil$  bits. Thus, for example, when  $q = 3$ , the basis states  $|0\rangle$ ,  $|1\rangle$  and  $|2\rangle$  are represented by the two-qubit registers  $|00\rangle$ ,  $|01\rangle$  and  $|10\rangle$ , respectively. Note that there remains one state (in the example,  $|11\rangle$ ) which does not correspond to any of the qudits. In general, there will be  $2^{\lceil \log q \rceil} - q$  such “non-qudigit” states.  $M_q$ ,  $F_q$  and  $H_q$  can now be defined to act on qubit registers, as follows. Consider a state  $|x\rangle$  where  $x$  is a number represented as  $m$  bits (i.e., an  $m$ -qubit register). If  $m < \lceil \log q \rceil$ , then  $H_q$  leaves  $|x\rangle$  unaffected. If  $0 \leq x \leq q - 1$  (where here we are identifying  $x$  with the number it represents), then  $H_q$  acts exactly as one expects, namely,  $H_q|x\rangle = (1/\sqrt{q}) \sum_{y=0}^{q-1} \zeta^{xy} |y\rangle$ . If  $x \geq q$ , again  $H_q$  leaves  $|x\rangle$  unchanged. Since the resulting transformation is a direct sum of unit matrices and matrices of the form of  $H_q$  as it was originally set down, the result is a unitary transformation.  $M_q$  and  $F_q$  can be defined to operate similarly on  $m$ -qubit registers for any  $m$ : Break up the  $m$  bits into blocks of  $\lceil \log q \rceil$  bits. If  $m$  is not divisible by  $\lceil \log q \rceil$ , then  $M_q$  and  $F_q$  do not affect the “remainder” block that contains fewer than  $\lceil \log q \rceil$  bits. Likewise, in a quantum register  $|x_1, \dots, x_n\rangle$  where each of the  $x_i$ 's (with the possible exception of  $x_n$ ) are  $\lceil \log q \rceil$ -bit numbers,  $M_q$  and  $F_q$  operate on the blocks of bits  $x_1, \dots, x_n$  exactly as expected, except that there is no affect on the “non-qudigit” blocks (in which  $x_i \geq q$ ), or on the (possibly) one remainder block for which  $|x_n| < \lceil \log q \rceil$ . Since  $M_q$  and  $F_q$  operate exactly as they did originally on blocks representing qudits, and like unity for non-qudigit or remainder blocks, it is clear that they remain unitary.

Henceforth,  $M_q$ ,  $F_q$ , and  $H_q$  should be understood to act on qubit registers as described above. Nevertheless, it will usually be convenient to think of them as acting on qudigit registers consisting of  $\lceil \log q \rceil$  qubits in each.

**Lemma 4.3**  $F_q$  and  $M_q$  are  $\text{QAC}^0$ -equivalent.


**Proof.** By Barenco et al. [3], any fixed dimension unitary matrix can be computed in fixed depth using one-qubit gates and controlled nots. Hence  $H_q$  can be computed in  $\text{QAC}^0$ , as can  $H_q^{\otimes(n+1)}$ . The result now follows immediately from Proposition 4.2.  $\square$

**Lemma 4.4**  $\text{MOD}_q$  and  $M_q$  are  $\text{QAC}^0$ -equivalent.

**Proof.** First note that  $\neg\text{MOD}_q$  and  $\text{MOD}_{q,r}$  are equivalent, since a  $\text{MOD}_{q,r}$  gate can be simulated by a  $\neg\text{MOD}_q$  gate with  $q - r$  extra inputs set to the constant 1. Since  $\neg\text{MOD}_q$  and  $\text{MOD}_q$  gates are equivalent, we can freely use  $\text{MOD}_{q,r}$  gates in place of  $\text{MOD}_q$  gates and vice versa.

It is easy to see that, given an  $M_q$  gate, we can simulate a  $\text{MOD}_q$  gate. Applying  $M_q$  to  $n + 1$  digits (represented as bits, but each digit only taking on the values 0 or 1) transforms,

$$|x_1, \dots, x_n, 0\rangle \mapsto |x_1, \dots, x_n, (\sum_i x_i) \bmod q\rangle.$$



**Figure 7.** A  $\widehat{\text{MOD}}_{3,r}$  circuit for  $r = 0$ . In the figure,  $\text{mod}(x)$  denotes  $\text{Mod}_{3,r}(x_1, \dots, x_n)$ . The notation on the right will be used as a shorthand for this circuit.

Now send the bits of the last block ( $\sum_i x_i \bmod q$ ) to an  $n$ -ary OR gate with control bit  $b$  (see the proof of Proposition 4.1). The resulting output is exactly  $b \oplus \text{Mod}_q(x_1, \dots, x_n)$ . The bits in the last block can be erased by reversing the  $M_q$  gate. This leaves only  $x_1, \dots, x_n$ ,  $\mathcal{O}(n)$  work bits, and the output  $b \oplus \text{Mod}_q(x_1, \dots, x_n)$ .

The converse (simulating  $M_q$  given  $\text{MOD}_q$ ) requires some more work. The first step is to show that  $\text{MOD}_{q,0}$  can also determine if a sum of *digits* is divisible by  $q$ . Let  $x_1, \dots, x_n \in D$  be a set of digits represented as  $\lceil \log q \rceil$  bits each. For each  $i$ , let  $x_i^{(k)}$  ( $0 \leq k \leq \lceil \log q \rceil - 1$ ) denote the bits of  $x_i$ . Since the numerical value of  $x_i$  is  $\sum_{k=0}^{\lceil \log q \rceil - 1} x_i^{(k)} 2^k$ , it follows that

$$\sum_{i=1}^n x_i = \sum_{k=0}^{\lceil \log q \rceil - 1} \sum_{i=1}^n x_i^{(k)} 2^k.$$

The idea is to express this last sum in terms of a set of Boolean inputs that are fed into a  $\text{MOD}_{q,0}$  gate. To account for the factors  $2^k$ , each  $x_i^{(k)}$  is fanned out  $2^k$  times before plugging it into the  $\text{MOD}_{q,0}$  gate. Since  $k < \lceil \log q \rceil$ , this requires only constant depth and  $\mathcal{O}(n)$  work bits (which of course are set back to 0 in the end by reversing the fanout). Thus, just using  $\text{MOD}_{q,0}$  and constant fanout, we can determine if  $\sum_{i=1}^n x_i \equiv 0 \pmod q$ . More generally, we can determine if  $\sum_{i=1}^n x_i \equiv r \pmod q$  using just a  $\text{MOD}_{q,r}$  gate and constant fanout. Let  $\widehat{\text{MOD}}_{q,r}(x_1, \dots, x_n)$  denote the resulting circuit, that determines if a sum of digits is congruent to  $r \pmod q$ . The construction of  $\widehat{\text{MOD}}_{q,r}(x_1, \dots, x_n)$  is illustrated in Figure 7 for the case of  $q = 3$ .

We can get the bits in the value of the sum  $\sum_{i=1}^n x_i \bmod q$  using  $\widehat{\text{MOD}}_{q,r}$  circuits. This is done, essentially, by implementing the relation  $x \bmod q = \sum_{r=0}^{q-1} r \cdot \text{Mod}_{q,r}(x)$ . For each  $r$ ,  $0 \leq r \leq q-1$ , we compute  $\text{Mod}_{q,r}(x_1, \dots, x_n)$  (where now the  $x_i$ 's are digits). This can be done by applying the  $\widehat{\text{MOD}}_{q,r}$  circuits in series (for each  $r$ ) to the same inputs, introducing a 0 work bit for each application, as illustrated in Figure 8.






Figure 8. Applying  $\widehat{\text{MOD}}_{q,r}$  circuits in series.

Let  $r_k$  denote the  $k^{\text{th}}$  bit of  $r$ . For each  $r$  and for each  $k$ , we take the AND of the output of the  $\widehat{\text{MOD}}_{q,r}$  with  $r_k$  (again by applying the AND's in series, which is still constant depth, but introduces  $q$  extra work inputs). Let  $a_{k,r}$  denote the output of one of these AND's. For each  $k$ , we OR together all the  $a_{k,r}$ 's, that is, compute  $\bigvee_{r=0}^{q-1} a_{k,r}$ , again introducing a constant number of work bits. Since only one of the  $r$ 's will give a non-zero output from  $\widehat{\text{MOD}}_{q,r}$ , this collection of OR gates outputs exactly the bits in the value of  $\sum_{i=1}^n x_i \bmod q$ . Call the resulting circuit  $C$ , and the sum it outputs  $S$ .

Finally, to simulate  $M_q$ , we need to include the input digit  $b \in D$ . To do this, we apply a unitary transformation  $T$  to  $|S, b\rangle$  that transforms it to  $|S, (b + S) \bmod q\rangle$ . By Barenco, et al. [3] (as in the proof of Lemma 4.3),  $T$  can be computed in fixed depth using one-qubit gates and controlled NOT gates. Now using  $S$  and all the other work inputs, we reverse the computation of the circuit  $C$ , thus clearing the work inputs. This is illustrated in figure 9.




Figure 9. Combining circuits to compute  $M_q$ .

The result is an output consisting of  $x_1, \dots, x_n$ ,  $\mathcal{O}(n)$  work bits, and  $(b + \sum_{i=1}^n x_i) \bmod q$ , which is the output of an  $M_q$  gate.  $\square$

It is clear that we can fan out digits, and therefore bits, using an  $F_q$  gate (setting  $x_i = 0$  for  $1 \leq i \leq n$  fans out  $n$  copies of  $b$ ). It is slightly less obvious (but still straightforward) that, given an  $F_q$  gate, we can fully simulate an  $F$  gate.

**Lemma 4.5** *For any  $q > 2$ ,  $F$  and  $F_q$  are  $\text{QAC}^0$ -equivalent.*

**Proof.** By the preceding lemmas,  $F_q$  and  $\text{MOD}_q$  are  $\text{QAC}^0$ -equivalent. By Proposition 4.1,  $\text{MOD}_q$  is  $\text{QAC}^0$ -reducible to  $F$ . Hence  $F_q$  is  $\text{QAC}^0$ -reducible to  $F$ .

Conversely, arrange each block of  $\lceil \log q \rceil$  input bits to an  $F_q$  gate as follows. For the control-bit block (which contains the bit we want to fan out), set all but the last bit to zero, and call the last bit  $b$ . Set all bits in the  $i^{\text{th}}$  input-bit block to 0. Now the  $i^{\text{th}}$  output of the  $F_q$  circuit is  $b$ , represented as  $\lceil \log q \rceil$  bits with only one possibly nonzero bit. Send this last output bit  $b$  and the input bit  $x_i$  to a controlled-NOT gate. The outputs of that gate are  $b$  and  $b \oplus x_i$ . Now apply  $F_q^{-1}$  to the bits that were the outputs of the  $F_q$  gate (which are all left unchanged by the controlled-not's). This returns all the  $b$ 's to 0 except for the control bit which is always unchanged. The outputs of the controlled-not's give the desired  $b \oplus x_i$ . Thus the resulting circuit simulates  $F$  with  $\mathcal{O}(n)$  work bits.  $\square$

**Theorem 4.6** *For any  $q \in \mathbf{N}$ ,  $q \neq 1$ ,  $\text{QACC} = \text{QACC}[q]$ .*

**Proof.** By the preceding lemmas, fanout of bits is equivalent to the  $\text{MOD}_q$  function. Thus we can do fanout, and hence  $\text{MOD}_2$ , if we can do  $\text{MOD}_q$ . By the result of Proposition 4.1, we can do  $\text{MOD}_q$  if we can do fanout in constant depth. Hence  $\text{QACC} = \text{QACC}[2] \subseteq \text{QACC}[q]$ .  $\square$

To compare these results with classical circuits requires a little care. For any Boolean function  $\phi$  with  $n$  inputs and  $m$  outputs, we can define a reversible version  $\phi'$  on  $n + m$  bits where  $\phi'(x, y) = (x, y \oplus \phi(x))$  keeps the input  $x$  and XORs the output  $\phi(x)$  with  $y$ . Then if  $\phi$  has a circuit with depth  $d$  and width  $w$ , it is easy to construct a reversible circuit for  $\phi'$  of depth  $2d - 1$  where  $wd$  work bits start and end in the zero state. We do this by assigning a work bit to each gate in the original circuit, and replacing each gate with a reversible one that XORs that work bit with the output. Then we can erase the work bits by moving backward through the layers of the circuit.

Then if we adopt the convention that a Boolean function with  $n$  inputs and  $m$  outputs is in a quantum circuit class if its reversible version is, we clearly have, for any  $k$ ,  $\text{AC}^k \subseteq \text{QAC}_{\text{wf}}^k$  and  $\text{ACC}^k \subseteq \text{QACC}^k$ . Thus we have

$$\text{AC}^0 \subset \text{ACC}^0[2] \subset \text{ACC}^0 \subseteq \text{QAC}_{\text{wf}}^0 = \text{QACC}[q] = \text{QACC}$$

showing that  $\text{QAC}_{\text{wf}}^0$  and  $\text{QACC}[2]$  are more powerful than  $\text{AC}^0$  and  $\text{ACC}^0[2]$  respectively.

Interestingly, if  $\text{QAC}^0$  as we first defined it cannot do fanout, i.e. if  $\text{QAC}^0 \subset \text{QAC}_{\text{wf}}^0$ , then in a sense it fails to include  $\text{AC}^0$ , since the fanout function from  $\{0, 1\}$  to  $\{0, 1\}^n$  is trivially in  $\text{AC}^0$ . However, it is not clear whether it fails to include any  $\text{AC}^0$  functions with a one-bit output. On the other hand, if  $\text{QAC}^0$  can do fanout, it can also do parity and is greater than  $\text{AC}^0$ , so either way  $\text{AC}^0$  and  $\text{QAC}^0$  are different. We are indebted to Pascal Tesson for pointing this out.

## 5 Upper Bounds

In this section, we prove the upper bounds results  $\text{NQACC}_{\text{gates}}^{\log} \subseteq \text{TC}^0$ ,  $\text{BQACC}_{\text{Q,gates}}^{\log} \subseteq \text{TC}^0$ ,  $\text{NQACC}_{\text{pl}}^{\log} \subseteq \text{P/poly}$ , and  $\text{BQACC}_{\text{Q,pl}}^{\log} \subseteq \text{P/poly}$ .

Suppose  $\{F_n\}$  and  $\{z_n\}$  determine a language  $L$  in **NQACC**. Let  $F_n$  be the product of the layers  $U_1, \dots, U_t$  and  $E$  be the distinct entries of the matrices used in the  $U_j$ 's. By our definition of **QACC**, the size of  $E$  is fixed with respect to  $n$ . We need a canonical way to write sums and products of elements in  $E$  to be able to check  $|\langle \vec{z} | U_1 \cdots U_t | x, 0^{p(n)} \rangle|^2 > 0$  with a  $\text{TC}^0$  function. To do this let  $A = \{\alpha_i\}_{1 \leq i \leq m}$  be a maximal algebraically independent subset of  $E$ . Let  $F = \mathbf{Q}(A)$  and let  $B = \{\beta_i\}_{0 \leq i < d}$  be a basis for the field  $G$  generated by the elements in  $(E - A) \cup \{1\}$  over  $F$ . Since the size of the bases of  $F$  and  $G$  are less than the cardinality of  $E$  the size of these bases is also fixed with respect to  $n$ .

As any sum or product of elements in  $E$  is in  $G$ , it suffices to come up with a canonical form for elements in  $G$ . Our representation is based on Yamakami and Yao [31]. Let  $\alpha \in G$ . Since  $B$  is a basis,  $\alpha = \sum_{j=0}^{d-1} \lambda_j \beta_j$  for some  $\lambda_j \in F$ . We encode an  $\alpha$  as a  $d$ -tuple (we iterate the pairing function from the preliminaries to make  $d$ -tuples)  $\langle \lceil \lambda_0 \rceil, \dots, \lceil \lambda_{d-1} \rceil \rangle$  where  $\lceil \lambda_j \rceil$  encodes  $\lambda_j$ . As the elements of  $A$  are algebraically independent, each  $\lambda_j = s_j / u_j$  where  $s_j$  and  $u_j$  are of the form

$$\sum_{\vec{k}_j, |\vec{k}_j| \leq e} a_{\vec{k}_j} \left( \prod_{i=1}^m \alpha_i^{k_{ij}} \right).$$

Here  $\vec{k}_j = (k_{1j}, \dots, k_{mj}) \in \mathbf{Z}^m$ ,  $|\vec{k}_j|$  is  $\sum_i k_{ij}$ ,  $a_{\vec{k}_j} \in \mathbf{Z}$ , and  $e \in \mathbf{N}$ . In particular, any product  $\beta_m \cdot \beta_l = \sum_{j=0}^{d-1} \lambda_j \beta_j$  with  $\lambda_j = s_j / u_j$  and  $s_j$  and  $u_j$  in this form. We take a common denominator  $u$  for elements of  $E \cup \{\beta_m \cdot \beta_l\}$  and not just  $E$  since the  $\lambda_j$ 's associated with the  $\beta_m \cdot \beta_l$  might have additional factors in their denominators not in  $E$ . Also fix an  $e$  large enough to bound the  $|\vec{k}_j|$ 's which might appear in any element of  $E$  or a product  $\beta_m \cdot \beta_l$ . This  $e$  will be constant with respect to  $n$ . In multiplying  $t$  layers of **QACC** circuit against an input, the entries in the result will be polynomial sums and products of elements in  $E \cup \{\beta_m \cdot \beta_l\}$ , so we can bound  $|\vec{k}_j|$  for  $\vec{k}_j$ 's which appear in the  $\lambda_j$ 's of such an entry by  $e \cdot p(n)$ . To complete our representation of  $\alpha \in G$  we encode  $\lambda_j$  as the sequence  $\langle r, \langle \langle a_{\vec{k}_j}, k_{1j}, \dots, k_{mj} \rangle \rangle \rangle$  where  $r$  is the power to which  $u$  is raised and  $\langle \langle a_{\vec{k}_j}, k_{1j}, \dots, k_{mj} \rangle \rangle$  is the sequence of  $\langle a_{\vec{k}_j}, k_{1j}, \dots, k_{mj} \rangle$ 's that appear in  $s_j$ . By our discussion, the encoding of an  $\alpha$  that appears as an entry in the output after applying a **QACC** operator to the input is of polynomial length and so can be manipulated in  $\text{TC}^0$ .

We have need of the following lemma:

**Lemma 5.1** *Let  $p$  be a polynomial. (1) Let  $f(i, x) \in \text{TC}^0$  output encodings of  $a_{i,x} \in \mathbf{Z}[A]$ . Then  $\mathbf{Z}[A]$  encodings of  $\sum_{i=1}^{p(|x|)} a_{i,x}$  and  $\prod_{i=1}^{p(|x|)} a_{i,x}$  are  $\text{TC}^0$  computable. (2) Let  $f(i, x) \in \text{TC}^0$  output encodings of  $a_{i,x} \in G$ . Then  $G$  encodings of  $\sum_{i=1}^{p(|x|)} a_{i,x}$  and  $\prod_{i=1}^{p(|x|)} a_{i,x}$  are  $\text{TC}^0$  computable.*

**Proof.** We will abuse notation in this proof and identify the encoding  $f(i, x)$  with its value  $a_{i,x}$ . So  $\sum_i f(i, x)$  and  $\prod_i f(i, x)$  will mean the encoding of  $\sum_i a_{i,x}$  and  $\prod_i a_{i,x}$  respectively.

(1) To do sums, the first thing we do is form the list  $L1 = \langle f(0, x), \dots, f(p(|x|), x) \rangle$ . Then we create a flattened list  $L2$  from this with elements which are the  $\langle a_{\vec{k}_j}, k_{1j}, \dots, k_{mj} \rangle$ 's from the  $f(i, x)$ 's.  $L1$  is in  $\text{TC}^0$  using our definition of sequence from the preliminaries, and closure under sums and  $\max_i$  to find the length of the longest  $f(i, x)$ . To flatten  $L1$  we use  $\max_i$  to find the length  $d$  of the longest  $f(i, x)$  for  $i \leq p(|x|)$ . Then using  $\max$  twice we can find the length of the longest  $\langle a_{\vec{k}_j}, k_{1j}, \dots, k_{mj} \rangle$ . This will be the second coordinate in the pair used to define sequence

$L2$ . We then do a sum of size  $d \cdot p(|x|)$  over the subentries of  $L1$  to get the first coordinate of the pair used to define  $L2$ . Given  $L2$ , we make a list  $L3$  of the distinct  $\vec{k}_j$ 's that appear as  $\langle a_{\vec{k}_j}, k_{1j}, \dots, k_{mj} \rangle$  in some  $f(i, x)$  for some  $i \leq p(|x|)$ . This list can be made from  $L2$  using sums, *cond* and  $\mu$ . We sum over the  $t \leq \text{length}(L2)$  and check if there is some  $t' < t$  such that the  $t'$ th element of  $L2$  has same  $\vec{k}_j$  as  $t$  and if not add the  $t$ th elements  $\vec{k}_j$  times 2 raised to the appropriate power. We know what power by computing the sum of the number of smaller  $t'$  that passed this test. Using *cond* and closure under sums we can compute in  $\text{TC}^0$  a function which takes a list like  $L2$  and a  $\vec{k}_j$  and returns the sum of all the  $a_{\vec{k}_j}$ 's in this list. So using this function and the lists  $L2$  and  $L3$  we can compute the desired encoding.

For products, since the  $\alpha_i$ 's of  $A$  are algebraically independent,  $\mathbf{Z}[A]$  is isomorphic to the polynomial ring  $\mathbf{Z}[y_1, \dots, y_m]$  under the natural map which takes  $\alpha_j$  to  $y_j$ . We view our encodings  $f(i, x)$  as  $m$ -variate polynomials in  $\mathbf{Z}[y_1, \dots, y_m]$ . We describe for any  $p'$  a circuit that works for any  $\text{TC}^0$  computable  $f(i, x)$  such that  $\prod_i f(i, x)$  is of degree less than  $p'$  viewed as an  $m$ -variate polynomial. In  $\text{TC}^0$  we define  $g(i, x)$  to consist of the sequence of polynomially many integer values which result from evaluating the polynomial encoded by  $f(i, x)$  at the points  $(i_1, \dots, i_m) \in \mathbf{N}^m$  where  $0 \leq i_s$  and  $\sum_s i_s \leq p'$ . To compute  $f(i, x)$  at a point involves computing a polynomial sum of a polynomial product of integers, and so will be in  $\text{TC}^0$ . Using closure under polynomial integer products we compute  $k(j, x) := \prod_i \beta(j, g(i, x))$  where  $\beta$  is the sequence projection function from the preliminaries. Our choice of points is what is called by Chung and Yao [7] the  $p'$ -th order principal lattice of the  $m$ -simplex given by the origin and the points  $p'$  from the origin in each coordinate axis. By Theorems 1 and 4 of that paper (proved earlier by a harder argument in Nicolaidis [20]) the multivariate Lagrange Interpolant of degree  $p'$  through the points  $k(j, x)$  is unique. This interpolant is of the form  $P(y_1, \dots, y_m) = \sum_j p_j(y_1, \dots, y_m)k(j, x)$  where the  $p_j$ 's are polynomials which do not depend on the function  $f$ . An explicit formula for these  $p_j$ 's is given in Corollary 2 of Chung and Yao [7] as a polynomial product of linear factors. Since these polynomials are all of degree less than  $p'$ , they have only polynomial in  $p'$  many coefficients and in PTIME these coefficients can be computed by iteratively multiplying the linear factors together. We can then hard code these  $p_j$ 's (since they don't depend on  $f$ ) into our circuit and with these  $p_j$ 's,  $k(j, x)$ , and closure under sums we can compute the polynomial of the desired product in  $\text{TC}^0$ .

(2) We do sums first. Assume  $f(i, x) := \sum_{j=0}^{d-1} \lambda_{ij} \beta_j$ . One immediate problem is that the  $\lambda_{ij}$  and  $\lambda_{i'j}$  might use different  $u^r$ 's for their denominators. Since  $\text{TC}^0$  is closed under poly-sized maximum, it can find the maximum value  $r_0$  to which  $u$  is raised. Then it can define a function  $g(i, x) = \sum_{j=0}^{d-1} \gamma_{ij} \beta_j$  which encodes the same element of  $G$  as  $f(i, x)$  but where the denominators of the  $\gamma_{ij}$ 's are now  $u^{r_0}$ . If  $\lambda_j$  was  $s_j/u^r$  we need to compute the encoding  $s_j \cdot u^{r_0-r}/u^{r_0}$ . This is straightforward from (1). Now

$$\sum_{i=1}^{p(|x|)} f(i, x) = \sum_{i=1}^{p(|x|)} g(i, x) = \sum_{j=0}^{d-1} [(\sum_{i=1}^{p(|x|)} s_{ij})/u^{r_0}] \beta_j,$$

where  $s_{ij}$ 's are the numerators of the  $\gamma_{ij}$ 's in  $g(i, x)$ . From part (1) we can compute the encoding  $e_j$  of  $(\sum_{i=1}^{p(|x|)} s_{ij})$  in  $\text{TC}^0$ . So the desired answer  $\langle \langle r_0, e_0 \rangle, \dots, \langle r_0, e_{d-1} \rangle \rangle$  is in  $\text{TC}^0$ .


For products  $\prod_{i=1}^{p(|x|)} f(i, x)$ , we play the same trick as the in the  $\mathbf{Z}[A]$  product case. We view our encodings of elements of  $G$  as  $d$ -variate polynomials in  $F(y_0, \dots, y_{d-1})$  under the map  $\beta_k$  goes

to  $y_k$ . (Note that this map is not necessarily an isomorphism.) We then create a function  $g(i, x)$  which consists of the sequence of values obtained by evaluating  $f(i, x)$  at polynomially many points in a lattice as in the first part of this lemma. Evaluating  $f(i, x)$  at a point can easily be done using the first part of this lemma. We then use part (1) of this lemma to compute the products  $k(j, x) = \beta(j, g(i, x))$ . We then get the interpolant  $P(y_0, \dots, y_{d-1}) = \sum_j p_j(y_0, \dots, y_m)k(j, x)$ . We non-uniformly obtain the encoding of  $p_j(\beta_0, \dots, \beta_{d-1})$  expressed as an element of  $G$ . i.e., in the form  $\sum_{w=0}^{d-1} \lambda_{jw} \beta_w$ . Thus, the product  $\prod_{i=1}^{p(|x|)} f(i, x)$  is

$$\sum_{w=0}^{d-1} \left( \sum_j \lambda_{jw} k(j, w) \right) \beta_w$$

The encoding of the products is the d-tuple given by  $\langle \sum_j \lambda_{j0} k(j, 0), \dots, \sum_j \lambda_{jd-1} k(j, d-1) \rangle$ . Each of its components is a polynomial sum of a product of two things in  $F$  and can be computed using the first part of the lemma.  $\square$

For  $\{F_n\} \in \text{QAC}_{wf}^0 = \text{QACC}$ , the vectors that  $F_n$  act on are elements of a  $2^{n+p(n)}$  dimensional space  $\mathcal{E}_{1,n+p(n)}$  which is a tensor product of the 2-dimensional spaces  $\mathcal{E}_1, \dots, \mathcal{E}_{n+p(n)}$ , which in turn are each spanned by  $|0\rangle, |1\rangle$ . We write  $\mathcal{E}_{j,k}$  for the subspace  $\otimes_{i=j}^k \mathcal{E}_i$  of  $\mathcal{E}_{1,n+p(n)}$ . We now define a succinct way to represent a set of vectors in  $\mathcal{E}_{1,n+p(n)}$  which is useful in our argument below. A *tensor graph* is a directed acyclic graph with one source node of indegree zero, one terminal node of outdegree zero, and two kinds of edges: horizontal edges, which are unlabeled, and vertical edges, which are labeled with a pair of amplitudes and a product of *colors* and *anticolors* (which are defined below). We require that all paths from the source to the terminal traverse the same number of vertical edges and that no vertex can have vertical edge indegree greater than one or outdegree greater than one. The *height* of a node in a tensor graph is the number of vertical edges traversed to get to it on any path from the source; the *height* of an edge is the height of its end node. The *width* of a tensor graph is maximum number of nodes of the same height. As an example of a tensor graph where our color product is the number 1, consider the following figure:



The rough idea of tensor graphs is that paths through the graph correspond to collections of vectors in  $\mathcal{E}_{1,n}$ . For this particular figure the left path from the source node (s) to the terminal node (t) corresponds to the vectors given by

$$|1\rangle \otimes \left( \frac{1}{\sqrt{2}}|0\rangle + \frac{1}{\sqrt{2}}|1\rangle \right) \otimes \frac{1}{2}|0\rangle$$

and the right hand path corresponds to

$$|0\rangle \otimes \left( \frac{1}{\sqrt{2}}|0\rangle + \frac{-1}{\sqrt{2}}|1\rangle \right) \otimes \frac{1}{2}|0\rangle.$$


A  $\mathcal{E}_{j,k}$ -term in a tensor graph is a maximal induced tensor subgraph between a node of height  $j - 1$  and a node of height  $k$ . If the horizontal indegree of the node at height  $j - 1$  is zero and the horizontal outdegree of the node at height  $k$  is zero then we say the term is *good*. For the graph we considered above there are two good  $\mathcal{E}_{1,2}$ -terms and two good  $\mathcal{E}_{2,3}$ -terms but only one  $\mathcal{E}_{1,3}$ -term corresponding to the whole figure.

“Colors” are used to handle controlled-not layers. A color  $c$  and its anticolor  $\tilde{c}$  are defined to obey the following multiplicative properties:  $c \cdot c = \tilde{c} \cdot \tilde{c} = 1$  and  $c \cdot \tilde{c} = 0$ . Given a color  $b$  and a product of colors  $c$  not involving  $b$  or its anticolor we require  $b \cdot c = c \cdot b$  and  $\tilde{b} \cdot c = c \cdot \tilde{b}$ . If  $a$  is a product of colors and anticolors not involving the color  $b$  or  $\tilde{b}$  and  $c$  is another product of colors we have  $a(bc) = (ab)c$ . We consider formal sums of products of complex numbers times colors. We require complex numbers to commute with colors and require colors and anticolors to distribute, i.e., if  $a, b, c$  are colors or anticolors then  $a \cdot (b+c) = a \cdot b + a \cdot c$  and  $(b+c) \cdot a = b \cdot a + c \cdot a$ . Finally, we require addition to work so that the above structure satisfies the axioms of an  $\mathbf{C}$ -algebra. Given a tensor graph  $G$  denote this  $\mathbf{C}$ -algebra by  $\mathcal{A}_G$ . Since

$$(a \cdot a) \cdot \tilde{a} = \tilde{a} \neq 0 = a \cdot (a \cdot \tilde{a})$$

this algebra is not associative. However, in the sums we will consider, the terms will never have more than two positions where a color or its anticolor can occur, so the products we will consider are associative.

Using our earlier encoding for the elements of  $\mathbf{C}$  which could appear in a QACC computation, it is straightforward to use sequence coding to get a  $\text{TC}^0$  encodings of the relevant elements of  $\mathcal{A}_G$ . As an example of how colors affect amplitudes, consider the following picture:



The amplitude of  $|1, 0, 0\rangle$  in the left hand dotted path is  $b \cdot \frac{-1}{\sqrt{2}} \cdot 1 \cdot \frac{-1}{\sqrt{2}} \cdot b \cdot 1 = 1/2$  using commutativity and  $b^2 = 1$ . Its amplitude in the right hand dotted path would be zero because of the last vertical edge. However, vectors such as  $|0, 0, 1\rangle$  would have nonzero amplitude in the right hand dotted path. Nevertheless, the amplitude of any vector  $|x\rangle$  in any path other than the dotted ones from  $s$  to  $t$  will be 0 as  $b \cdot \tilde{b} = 0$ . More formally, we define the amplitude of an  $|x\rangle$  in a vertical edge as

equal to the left amplitude times the color product in the edge if  $|x\rangle$  is  $|0\rangle$  and equal to the right amplitude times the color product in the edge if  $|x\rangle$  is  $|1\rangle$ . The amplitude of a vector  $|x_1, \dots, x_j\rangle$  in a path in a tensor graph is the product over  $k$  from 1 to  $j$  of the amplitude of the vectors  $|x_k\rangle$  in the vertical edge of height  $k$ . The amplitude of a vector  $|x_j, \dots, x_k\rangle$  in an  $\mathcal{E}_{j,k}$ -term is the sum of its amplitude in its paths. The amplitude of a vector  $|x_1, \dots, x_{p(n)}\rangle$  in a tensor graph  $G$  is defined to be the sum of its amplitudes in  $G$ 's  $\mathcal{E}_{1,p(n)}$ -terms.

As we will be interested in families of tensor graphs  $\{G_n\}$ , corresponding to our circuit families we want to look at those families with a certain degree of uniformity. We say a family of tensor graphs  $\{G_n\}$  is *color consistent* if: (1) the number of colors for edges of the same height is bounded by a constant  $k$  with respect to  $n$ , (2) the number of heights in which a given color/anticolor can appear is exactly two (colors and their anticolors must appear on the same heights), (3) each color product at the same height is of the form  $\prod_{i=0}^k l_i$  where  $l_i$  must be either a color  $c_i$  or  $\bar{c}_i$  (it follows there are  $2^k$  possible color products for edges at a given height). We say that a color/anticolor is *active* at a given height if the height is at or after the first height at which the color/anticolor occurs and is below the height of its second occurrence. The family is further said to be *log-color depth* if the number of active colors/anticolors of a given height is log-bounded.

**Theorem 5.1** *Let  $\{F_n\}$  be a family of QACC operators and let  $\{\langle z_n|\}$  a family of observables. (1) There is a color-consistent family of tensor graphs of width  $2^{2^t}$  and polynomial size representing the output amplitudes of  $U_1 \cdots U_t |z_n\rangle$  where  $U_i$  are the layers of  $F_n$ . (2) If  $\{F_n\}$  is in  $\text{QACC}_{pt}^{\log}$  then the family of tensor graphs will be of log-color depth. (3) If  $\{F_n\}$  is in  $\text{QACC}_{gates}^{\log}$  then the number of paths from the source to the terminal node is polynomially bounded.*

**Proof.** The proof is by induction on  $t$ . In the base case,  $t = 0$ , we do not multiply any layers, and we can easily represent this as a tensor graph of width 1. Assume for  $j < t$  that  $U_j \cdots U_1 |\vec{x}, 0^{p(n)}\rangle$  can be written as color consistent tensor graph of width  $2^{2^t}$  and polynomial size. There are two cases to consider: In the first case the layer is a tensor product of matrices  $M_1 \otimes \cdots \otimes M_\nu$  where the  $M_k$ 's are Toffoli gates, one qubit gates, or fan-out gates (since  $\text{QACC}_{wf}^0 = \text{QACC}$ ); in the second case the layer is a controlled-not layer.

For the first case we “multiply”  $U_t$  against our current graph by “multiplying” each  $M_j$  in parallel against the terms in our sum corresponding to  $M_j$ 's domain, say  $\mathcal{E}_{j',k'}$ . If  $M_j = \begin{pmatrix} u_{00} & u_{01} \\ u_{10} & u_{11} \end{pmatrix}$  with domain  $\mathcal{E}_{j'}$  is a one-qubit gate, then we multiply the two amplitudes in each vertical edge of height  $j'$  in our tensor graph by  $M_j$ . This does not effect the width, size, or number of paths through the graph. If  $M_j$  is a Toffoli gate, then for each good term  $S$  in  $\mathcal{E}_{j',k'}$  in our tensor graph we add one new term to the resulting graph. This term is added by adding a horizontal edge going out from the source node of  $S$  followed by the new  $\mathcal{E}_{j',k'}$ -term followed by a horizontal edge into the terminal node of  $S$ . The new term is obtained from  $S$  by setting to 0 the left hand amplitudes of all edges in  $S$  of height between  $j'$  and  $k' - 1$  and then if  $\alpha, \gamma$  is the amplitude of an edge of height  $k'$  in the new term we change it to  $\gamma - \alpha, \alpha - \gamma$ . This new term adjusts the amplitude for the case of a  $|1\rangle^{\otimes(k'-j'-1)}$  vector in  $\mathcal{E}_{j',k'-1}$  tensored with either a  $|0\rangle$  or  $|1\rangle$ . This operation increases the width of the new tensor graph by the width of the good  $\mathcal{E}_{j',k'}$ -term for each good  $\mathcal{E}_{j',k'}$ -term in the graph. Since the original graph has width  $2^{2^{(t-1)}}$  there are at most this many starting and ending vertices for such terms. So there at most  $(2^{2^{(t-1)}})^2$  such terms. Each of these

terms has width at most  $2^{2^{(t-1)}}$ . Thus, the new width is at most

$$2^{2^{(t-1)}} + (2^{2^{(t-1)}})^2 \cdot 2^{2^{(t-1)}} < 2^{2^{2t}}.$$

Notice this action adds one new path through the  $\mathcal{E}_{j',k'}$  part of the graph for every existing one.

Now suppose  $M_j$  is a fan-out gate, let  $S$  be a good  $\mathcal{E}_{j',k'}$ -term in our tensor graph and let  $e$  be any vertical edge in  $S$  in  $\mathcal{E}_{k'}$ . Suppose  $e$  has amplitude  $\alpha$  for  $|0\rangle$  and amplitude  $\gamma$  for  $|1\rangle$ . In the new graph we change the amplitude of  $e$  to  $\alpha, 0$ . We then add a horizontal edge out of the source node of  $S$  followed by a new  $\mathcal{E}_{j',k'}$ -term followed by a horizontal edge into the terminal node of  $S$ . The new term is obtained from  $S$  by changing the amplitude for edges in  $\mathcal{E}_{k'}$  with amplitudes  $\alpha, \gamma$  in  $S$  to  $0, \gamma$ . The amplitudes of the non- $\mathcal{E}_{k'}$  edges in this term are the reverse of the corresponding edge in  $S$ , i.e., if the edge in  $S$  had amplitude  $\delta, \zeta$  then the new term edge would have amplitude  $\zeta, \delta$ . The same argument as in the Toffoli case shows the new width is bounded by  $2^{2^{2t}}$  and that this action adds one new path through the  $\mathcal{E}_{j',k'}$  part of the graph for every existing one.

For the case of a controlled-not layer, suppose we have a controlled-not going from line  $i$  onto line  $j$ . Let  $c, \bar{c}$  be a new color, anti-color pair not yet appearing in the graph. Let  $e_i$  be a vertical edge of height  $i$  in the graph and let  $C_i, \alpha_i, \gamma_i$  be respectively its color product and two amplitudes. Similarly, let  $e_j$  be a vertical edge of height  $j$  in the graph and  $C_j, \alpha_j, \gamma_j$  be its color product and two amplitudes. In the new graph we multiply  $c$  times the color product of  $e_i$  and  $e_j$  and change the amplitude of  $e_i$  to  $\alpha_i, 0$ . We then add a horizontal edge going out from the starting node of  $e_i$ , followed by a vertical edge with values  $C_i \cdot \bar{c}, 0, \gamma_i$  followed by a horizontal edge into the terminal node of  $e_i$ . In turn, we add a horizontal edge going out of the starting node of  $e_j$ , followed by a vertical edge with values  $C_j \cdot \bar{c}, \gamma_j, \alpha_j$  followed by a horizontal edge into the terminal node of  $e_j$ . We handle all other controlled gates in this layer in a similar fashion (recall they must go to disjoint lines). We add at most a new vertex of a given height for every existing vertex of a given height. So the total width is at most doubled by this operation and  $2 \cdot 2^{2^{(t-1)}} < 2^{2^{2t}}$ . In the  $\text{QACC}_{pl}^{\log}$  case, simulating a layer which is a Kronecker product of spaced controlled-not gates and identity matrices, notice we would at most add one to the color depth at any place. So if a controlled-not layer is a composition of  $\mathcal{O}(\log)$  many such layers it will increase the color depth by  $\mathcal{O}(\log)$ . In the  $\text{QACC}_{gates}^{\log}$  case, notice that simulating a single controlled-not we add one new path for each existing path through the graph at each of the two heights affected. This gives three new paths on the whole subspace for each old one.

Since we have handled the two possible layer cases and the changes we needed to make only increase the resulting tensor graph polynomially, we thus have established the induction step and (1) and (2) of the theorem. For (3), observe for each multi-line gate we handle in adding a layer we at most quadruple the number of paths through the subspace where that gate applies. Since there are at most logarithmically many such gates, the number of paths through the graph increases polynomially.  $\square$

**Theorem 5.2** *Let  $\{G_n\}$  be a family of constant width color-consistent tensor graphs of vectors in  $\mathcal{E}_{1,p(n)}$ . Assume the coefficients of amplitudes in the  $\{G_n\}$  can be encoded in  $\text{TC}^0$  using our encoding scheme described earlier and that  $\{G_n\}$  has log-color depth. Then the amplitude of any basis vector of  $\mathcal{E}_{1,p(n)}$  in  $G_n$  is P/poly computable. If the number of paths through the graph from the source to the terminal node is polynomially bounded then the amplitude of any basis vector is  $\text{TC}^0$  computable.*



**Proof.** Let  $G_n$  be a particular graph in the family and let  $|\vec{x}_n\rangle$  be the vector whose amplitude we want to compute. Assume that all graphs in our family have fewer than  $k$  colors in any color product and have a width bounded by  $w$ . We will proceed from the source to the terminal node one height at a time to compute the amplitude. Since the width is  $w$  the number of  $\mathcal{E}_1$ -terms is at most  $w$  and each of these must have width at most  $w$ . Let  $\alpha_{1,1}, \dots, \alpha_{1,w}$  (some of which may be zero) denote the amplitudes in  $\mathcal{A}_{G_n}$  of  $|x_{n,1}\rangle$  in each of these terms. The  $\alpha_{1,i}$  are each sums of at most  $w$  amplitudes times the color products of at most  $k$  colors and anticolors, so the encoding of these  $w$  amplitudes is  $\text{TC}^0$  computable. Because of the restriction on the width of  $G_n$  there are at most  $w$  many  $\mathcal{E}_{1,j}$ -terms,  $w^2$  many  $\mathcal{E}_{j,j+1}$ -terms, and  $w$  many  $\mathcal{E}_{1,j+1}$ -terms. Fixing some ordering on the nodes of height  $j$  and  $j+1$  let  $\gamma_{j,i,k}$  be the amplitude of  $|x_{n,j+1}\rangle$  in the  $\mathcal{E}_{j,j+1}$ -term with source the  $i$ th node of height  $j$  and with terminal node the  $k$ th node of height  $j+1$ . The amplitude is zero if there is no such  $\mathcal{E}_{j,j+1}$ -term. Then the amplitudes  $\alpha_{j+1,1}, \dots, \alpha_{j+1,w}$  of the  $\mathcal{E}_{1,j+1}$ -terms can be computed from the amplitudes  $\alpha_{j,1}, \dots, \alpha_{j,w}$  of the  $\mathcal{E}_{1,j}$ -terms using the formula

$$\alpha_{j+1,k} = \sum_{i=1}^w \alpha_{j,i} \cdot \gamma_{j,i,k}.$$

Thus  $\alpha_{j+1,k}$  can be computed from the  $\alpha_{j,i}$  using a polynomial sized circuit to do these adds and multiplies. Similarly, each  $\alpha_{j,k}$  can be computed by polynomial sized circuits from the  $\alpha_{j-1,k}$ 's and so on. Since we have log-color depth the number of terms consisting of elements in our field times color products in a  $\alpha_{j,k}$  will be polynomial. So the size of the  $\alpha_{j,k}$ 's  $j \leq p(n)$ ,  $k \leq w$  will be polynomial in the input  $\vec{x}_n$ . So the size of the circuits for each  $\alpha_{j,k}$  where  $j \leq p(n)$  and  $k \leq w$  will be polynomial size. There is only one  $\mathcal{E}_{1,p(n)}$ -term in  $G_n$  and its amplitude is that of  $|\vec{x}_n\rangle$ , so this shows it has polynomial sized circuits. For the  $\text{TC}^0$  result, if the number of paths is polynomially bounded, then the amplitude can be written as the polynomial sum of the amplitudes in each path. The amplitude in a path can in turn be calculated as a polynomial product of the amplitudes times the colors on the vertical edges in the path. Our condition on every color appearing at exactly two heights guarantees the color product along the whole path will be 1 or 0, and will be zero iff we get a color and its anticolor on the path. This is straightforward to check in  $\text{TC}^0$ , so this sum of products can thus be computed in  $\text{TC}^0$  using Lemma 5.1.  $\square$

### Corollary 5.3

- (1)  $\text{EQACC}_{pl}^{\log} \subseteq \text{NQACC}_{pl}^{\log} \subseteq \text{P/poly}$ , and  $\text{BQACC}_{\mathbb{Q},pl}^{\log} \subseteq \text{P/poly}$ .
- (2)  $\text{EQACC}_{\text{gates}}^{\log} \subseteq \text{NQACC}_{\text{gates}}^{\log} \subseteq \text{TC}^0$ , and  $\text{BQACC}_{\mathbb{Q},\text{gates}}^{\log} \subseteq \text{TC}^0$ .

**Proof.** Given a family  $\{F_n\}$  of  $\text{QACC}_{pl}^{\log}$  operators and a family  $\{|\vec{z}_n|\}$  of states we can use Theorem 5.1 to get a family  $\{G_n\}$  of log color depth, color-consistent tensor graphs representing the amplitudes of  $F_n^{-1}|\vec{z}_n\rangle$ . Note  $\{F_n^{-1}\}$  is also a family of  $\text{QACC}_{pl}^{\log}$  operators since Toffoli and fan-out gates are their own inverses, the inverse of any one qubit gate is also a one qubit gate (albeit usually a different one), and finally a controlled-not layer is its own inverse. Theorem 5.2 shows there is a P/poly circuit computing the amplitude of any vector  $|\vec{x}_n\rangle$  in this graph. This amounts to calculating

$$\langle \vec{x}_n | F_n^{-1} | \vec{z}_n \rangle = \langle \vec{z}_n | F_n | \vec{x}_n \rangle.$$

If this is nonzero, then  $|\langle \vec{z}_n | F_n | \vec{x}_n \rangle|^2 > 0$ , and we know  $\vec{x}$  is in the language. In the  $\text{BQACC}_{\mathbf{Q}}$  case everything is a rational so P/poly can explicitly compute the magnitude of the amplitude and check if it is greater than  $3/4$ . The  $\text{TC}^0$  result follows similarly from the  $\text{TC}^0$  part of Theorem 5.1.

□

Finally, we note that some of the inclusions in the previous corollary can be strengthened if we assume that the circuit families are polynomial-time uniform and their coefficients polynomial-time computable. In particular p-uniform  $\text{NQACC}_{pl}^{\log}$  is contained in P and p-uniform  $\text{NQACC}_{gates}^{\log}$  is contained in p-uniform  $\text{TC}^0$ .

## 6 Discussion and Open Problems

A number of open questions are suggested by our work.

- Is  $\text{QACC}^0 = \text{QACC}_{wf}^0$ ? That is, can the fanout gate be constructed in constant depth when each qubit can only act as an input to one gate in each layer?
- Is  $\text{QACC}_{wf}^0 = \text{QTC}^0$ ? That is, can the techniques used here be extended to construct quantum threshold gates in constant depth?
- Is all of  $\text{NQACC}$  in  $\text{TC}^0$  or even P/poly? We conjecture that  $\text{NQACC}$  is in  $\text{TC}^0$ . As mentioned in the introduction, we have developed techniques that remove some of the important obstacles to proving this.
- Are there any natural problems in  $\text{NQACC}$  that are not known to be in  $\text{ACC}$ ?
- What exactly is the complexity of the languages in  $\text{EQACC}$ ,  $\text{NQACC}$  and  $\text{BQACC}_{\mathbf{Q}}$ ? We entertain two extreme possibilities. Recall that the class  $\text{ACC}$  can be computed by quasipolynomial size depth 3 threshold circuits [32]. It would be quite remarkable if  $\text{EQACC}$  could also be simulated in that manner. However, it is far from clear if any of the techniques used in the simulations of  $\text{ACC}$  (the Valiant-Vazirani lemma, composition of low-degree polynomials, modulus amplification via the Toda polynomials, etc.), which seem to be inherently irreversible, can be applied in the quantum setting. At the other extreme, it would be equally remarkable if  $\text{NQACC}$  and  $\text{NQTC}^0$  (or  $\text{BQACC}_{\mathbf{Q}}$  and  $\text{NQTC}^0$ ) coincide. Unfortunately, an optimal characterization of  $\text{QACC}$  language classes anywhere between those two extremes would probably require new (and probably difficult) proof techniques.
- How hard are the fixed levels of  $\text{QACC}$ ? While lower bounds for  $\text{QACC}$  itself seem impossible at present, it might be fruitful to study the limitations of small depth  $\text{QACC}$  circuits (depth 2, for example).

**Acknowledgments:** We thank Bill Gasarch for helpful comments and suggestions.

## References

- [1] L. Adleman, J. DeMarrais, and M. Huang. Quantum computability. *SIAM J. Computing* **26** (1997) 1524–1540.
- [2] M. Ajtai. “ $\Sigma_1^1$  formulae on finite structures.” *Ann. Pure Appl. Logic* **24** (1983) 1–48.
- [3] A. Barenco, C. Bennett, R. Cleve, D.P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J.A. Smolin, and H. Weifurter, “Elementary gates for quantum computation.” *Phys. Rev. A* **52**, pages 3457–3467, 1995.
- [4] D. A. Barrington, “Bounded-width polynomial-size branching programs recognize exactly those languages in  $\text{NC}^1$ .” *J. Comput. System Sci.* **38** (1989) 150–164.
- [5] M. Beaudry, P. McKenzie, P. Péladeau, and D. Thérien, “Circuits with monoidal gates.” *Proc. Symposium on Theoretical Aspects of Computer Science* (1993) 555–565.
- [6] C.H. Bennett, “Time/space trade-offs for reversible computation.” *SIAM J. Computing* **4(18)** (1989) 766–776.
- [7] K. C. Chung and T. H. Yao, “On Lattices Admitting Unique Lagrange Interpolations.” *SIAM Journal of Numerical Analysis* **14** (1977) 735–743.
- [8] J.I. Cirac and P. Zoller, “Quantum computers with cold trapped ions.” *Phys. Rev. Lett.* **74** (1995) 4091–4094.
- [9] D.P. DiVincenzo and P.W. Shor, “Fault-tolerant error correction with efficient quantum codes.” quant-ph/9605031.
- [10] P. Clote, “On polynomial Size Frege Proofs of Certain Combinatorial Principles.” In P. Clote and J. Krajicek, editors, *Arithmetic, Proof Theory, and Computational Complexity*, Oxford (1993) 164–184.
- [11] L. Fortnow and J. Rogers, “Complexity Limitations on Quantum Computation.” *Proc. 13th Conference on Computational Complexity* (1998) 202–209.
- [12] S. Fenner, F. Green, S. Homer, and R. Pruiam, “Quantum NP is hard for PH.” *Royal Society of London A* (1999) 455, pp 3953 - 3966.
- [13] M. Furst, J.B. Saxe, and M. Sipser, “Parity, circuits, and the polynomial-time hierarchy.” *Math. Syst. Theory* **17** (1984) 13–27.
- [14] N. Gershenfeld and I. Chuang, “Bulk spin resonance quantum computation.” *Science* **275** (1997) 350.
- [15] F. Green, S. Homer, and C. Pollett, “On the Complexity of Quantum ACC,” in *Fifteenth Annual Conference on Computational Complexity Theory*, IEEE Computer Society Press, (2000) 250-262.

- [16] C. Moore, “Quantum Circuits: Fanout, Parity, and Counting.” quant-ph/9903046.
- [17] C. Moore, D. Thérien, F. Lemieux, J. Berman, and A. Drisko, “Circuits and Expressions with Non-Associative Gates.” *Journal of Computer and System Sciences* **60** (2000) 368–394.
- [18] C. Moore and M. Nilsson, “Parallel quantum computation and quantum codes.” quant-ph/9808027, and to appear in *SIAM J. Computing*.
- [19] C. Moore and M. Nilsson, “Some notes on parallel quantum computation.” quant-ph/9804034.
- [20] R. A. Nicolaides, “On a class of finite elements generated by Lagrange interpolation.” *SIAM Journal of Numerical Analysis* **9** (1972) 177–199.
- [21] C.H. Papadimitriou, *Computational Complexity*. Addison-Wesley, 1994.
- [22] A.A. Razborov, “Lower bounds for the size of circuits of bounded depth with basis  $\{\&, \oplus\}$ .” *Math. Notes Acad. Sci. USSR* **41(4)** (1987) 333–338.
- [23] M. Reck, A. Zeilinger, H.J. Bernstein, and P. Bertani, “Experimental realization of any discrete unitary operator.” *Phys. Rev. Lett.* **73** (1994) 58-61.
- [24] M. Saks, “Randomization and derandomization in space-bounded computation.” *Proc. 11th IEE Conference on Computational Complexity* (1996) 128–149.
- [25] P.W. Shor, “Algorithms for quantum computation: discrete logarithms and factoring.” *Proc. 35th IEEE Symposium on Foundations of Computer Science* (1994) 124–134.
- [26] P.W. Shor, “Fault-tolerant quantum computation.” quant-ph/9605011.
- [27] P. W. Shor, “Polynomial-time algorithms for prime number factorization and discrete logarithms on a quantum computer.” *SIAM J. Computing* **26** (1997) 1484–1509.
- [28] K.-Y. Siu and V Rowchowdhury, “On optimal depth threshold circuits for multiplication and related problems.” *SIAM J. Discrete Math.* **7** (1994) 284–292.
- [29] K.-Y. Siu and J Bruck, “On the power of threshold circuits with small weights.” *SIAM J. Discrete Math.* **4** (1991) 423–435.
- [30] R. Smolensky, “Algebraic methods in the theory of lower bounds for Boolean circuit complexity.” *Proc. 19th Annual ACM Symposium on Theory of Computing* (1987) 77–82.
- [31] T. Yamakami and A.C. Yao, “ $NQP_C = co-C=P$ .” To appear in *Information Processing Letters*.
- [32] A. C.-C. Yao, “On ACC and threshold circuits.” *Proc. 31st IEEE Symposium on Foundations of Computer Science* (1990) 619–627.
- [33] A. C.-C. Yao, “Quantum circuit complexity.” *Proc. 34th IEEE Symposium on Foundations of Computer Science* (1993) 352–361.