Query and Depth Upper Bounds for Quantum Unitaries via Grover Search

Gregory Rosenthal

Institute for Quantum Computing, University of Waterloo

We prove that any n-qubit unitary can be implemented (i) approximately in time $\tilde{O}(2^{n/2})$ with query access to an appropriate classical oracle, and also (ii) exactly by a circuit of depth $\tilde{O}(2^{n/2})$ with one- and two-qubit gates and $2^{O(n)}$ ancillae. The proofs involve similar reductions to Grover search. The proof of (ii) also involves a linear-depth construction of arbitrary quantum states using one- and two-qubit gates (in fact, this can be improved to constant depth with the addition of fanout and generalized Toffoli gates) which may be of independent interest. We also prove a matching $\Omega(2^{n/2})$ lower bound for (i) and (ii) for a certain class of implementations.

1 Introduction

This paper addresses two seemingly disparate questions in quantum circuit complexity via a common proof technique. The first of these questions is as follows:

Question 1.1 (The unitary synthesis problem [1, 2])**.** Is there a polynomial-time quantum algorithm A such that for every unitary U, there exists a classical oracle f such that A^f approximately implements U?

By A^f we mean A with query access to the Boolean function f. Note that Question 1.1 is concerned with the *overall* runtime rather than just the number of queries. Question 1.1 was posed by Aaronson and Kuperberg [2] and named the "unitary synthesis problem" by Aaronson [1]. An affirmative answer would imply that to obtain a small quantum circuit for a unitary U, it suffices to give an efficient algorithm for computing the oracle f, which is interesting because more is known about Boolean function complexity than is known about quantum circuit complexity. Aaronson [3] and Lombardi, Ma and Wright [4] discuss this motivation in the context of certain physically motivated unitaries and quantum cryptography respectively.

We also consider the following question:

Question 1.2. Given n, what is the minimum circuit depth required to exactly implement a worst-case n-qubit unitary using one- and two-qubit gates?

The depth of a circuit is the number of layers of gates in it. Circuit depth corresponds to parallel computation time, and quantum circuits of higher depth are believed to be more difficult to physically implement. An upper bound for the unitary synthesis problem implies a similar depth upper bound for *approximately* implementing worst-case unitaries, because any n-bit function can be computed by a circuit of depth $O(n)$ (e.g. a CNF or DNF), but otherwise Questions 1.1 and 1.2 are not obviously related.

We prove $\tilde{O}(2^{n/2})$ upper bounds for both Questions 1.1 and 1.2, which improve on the previously best known upper bounds by a constant factor in the exponent. This is discussed in more detail in Sections 1.1 and 1.2 respectively. Both upper bounds are proved using a reduction from the task of implementing a unitary U to that of implementing what we call a *U-column-constructor*, or U-CC for short:

Gregory Rosenthal: grosenth@uwaterloo.ca

Definition 1.3 (U-CC)**.** Given an n-qubit unitary U, call a unitary A acting on $m \geq 2n$ qubits a *U-column-constructor (U-CC)* if $A|x, 0^{m-n}\rangle = |x\rangle \otimes U|x\rangle \otimes |0^{m-2n}\rangle$ for all $x \in \{0,1\}^n$.

Informally, controlled on an input string $x \in \{0,1\}^n$, a U-CC constructs the corresponding output state $U|x\rangle$ of U in a separate register;[1] if this separate register is not initialized to the all-zeros state then a U-CC's behavior is unspecified subject to unitarity. Using a zero-error variant of Grover search we prove the following, where by C^A we mean C with A and $A^\dagger$ oracles:

Theorem 1.4. *There is a uniform family $(C_{n,m})_{n,m}$ for $m \geq 2n$ of quantum circuits, each making $O\big(2^{n/2}\big)$ queries to an m-qubit quantum oracle, such that for all n-qubit unitaries U and all m-qubit U-CCs A it holds that $C^A_{n,m}$ implements U.*

To see why Theorem 1.4 is nontrivial, suppose that we wish to apply a unitary U on the input state $\sum_{x\in\{0,1\}^n} \alpha_x|x\rangle$. A natural first step is to query a U-CC to obtain the state $\sum_{x\in\{0,1\}^n} \alpha_x|x\rangle \otimes U|x\rangle$. But now to obtain $U\sum_{x\in\{0,1\}^n} \alpha_x|x\rangle$, it is necessary to uncompute $|x\rangle$ in superposition controlled on $U|x\rangle$.

The analogous "state synthesis problem" for constructing quantum states using a classical oracle has polynomial-time solutions [1, 5], and in this paper we prove that any n-qubit state can be constructed in $O(n)$ depth using $\tilde{O}(2^n)$ ancillae. These upper bounds for constructing states generalize to implementing U-CCs, and plugging these implementations of U-CCs into the algorithm from Theorem 1.4 yields $\tilde{O}\big(2^{n/2}\big)$ upper bounds for Questions 1.1 and 1.2.

Finally, we prove a matching $\Omega\big(2^{n/2}\big)$ query lower bound for Theorem 1.4 when U is Haar random and the U-CC is defined appropriately, implying that new techniques are needed to make further progress on Questions 1.1 and 1.2. This is discussed further in Section 1.3.

Below we state our results informally; more precise statements will be given in subsequent sections.

1.1 Upper bound for the unitary synthesis problem

We prove the following:

Theorem 1.5. *There is an $\tilde{O}\big(2^{n/2}\big)$-time quantum algorithm A such that for every n-qubit unitary U, there exists a classical oracle f such that A^f approximately implements U to within exponentially small error.*

Theorem 1.5 is the first nontrivial upper or lower bound for the unitary synthesis problem. For comparison, every n-qubit unitary can be implemented using $\tilde{O}\big(2^{2n}\big)$ one- and two-qubit gates [6, Section 4.5], and hence by the Solovay-Kitaev theorem [7] can be approximated to within exponentially small error using $\tilde{O}\big(2^{2n}\big)$ gates from a finite gate set. Thus the trivial algorithm in which the oracle encodes the description of a circuit for approximating U runs in time $\tilde{O}\big(2^{2n}\big)$. Irani, Natarajan, Nirkhe, Rao and Yuen [8, Section 7.2] also observed that *with postselection* there is a polynomial-time solution to the unitary synthesis problem, using the Choi–Jamiołkowski isomorphism and quantum teleportation. Subsequently to our work, Lombardi et al. [4] proved that there is no one-query solution to the unitary synthesis problem using $o(2^n)$ qubits.

1.2 Circuit depth upper bounds for states and unitaries

For brevity we assign the following name to the "standard" quantum circuit model:

Definition 1.6 (QNC circuits)**.** A QNC circuit is a quantum circuit consisting of one- and two-qubit gates.

We prove the following:

Theorem 1.7. *Every n-qubit unitary can be implemented by a* QNC *circuit of depth $\tilde{O}\big(2^{n/2}\big)$ with $\tilde{O}\big(2^{2n}\big)$ ancillae.*

[1] Here and throughout this paper, we allow ancillae that start in the all-zeros state and are required to end in the all-zeros state.

Sun, Tian, Yang, Yuan and Zhang [9] proved that every n-qubit unitary can be implemented by a QNC circuit of depth $\tilde{O}(2^n)$ with $O(2^n)$ ancillae, compared to which the circuit from Theorem 1.7 has lower depth but more ancillae. More generally, Sun et al. [9] proved that for $m \leq 2^n$, any n-qubit unitary can be implemented by a QNC circuit of size (i.e. number of gates) $O(4^n)$ and depth $\tilde{O}(4^n/m)$ with m ancillae. In followup work, Yuan and Zhang [10] generalized our proof of Theorem 1.7 to show that for $2^n \leq m \leq 4^n$, any n-qubit unitary can be implemented by a QNC circuit of depth $\tilde{O}\left(2^{3n/2}m^{-1/2}\right)$ with m ancillae; when $m = 4^n$ this matches Theorem 1.7 up to $\mathrm{poly}(n)$ factors.

The size of a circuit is trivially at most its depth times number of qubits acted on, so the circuit from Theorem 1.7 has size $\tilde{O}\left(2^{2.5n}\right)$. This raises the following question:

Question 1.8. Can every n-qubit unitary be implemented by a *single* QNC circuit that is *both* of size $\tilde{O}(4^n)$ and depth $\tilde{O}\left(2^{n/2}\right)$?

Our proof of Theorem 1.7 uses a low-depth construction of quantum states which may be of independent interest, and of which we first state the following corollary:

Corollary 1.9. *Every n-qubit state can be constructed by a QNC circuit of depth $O(n)$ with $\tilde{O}(2^n)$ ancillae.*

Sun et al. [9] and Zhang, Li and Yuan [11] independently proved Corollary 1.9, respectively shortly before and shortly after we did, and with just $O(2^n)$ ancillae and $O(2^n)$ size. Yuan and Zhang [10] proved in followup work that every n-qubit state can be constructed by a QNC circuit of size $O(2^n)$ and depth $O\left(n + \frac{2^n}{n+m}\right)$ using $m \geq 0$ ancillae, and that these size and depth upper bounds are tight for all n, m. This improves on a slightly weaker tradeoff of Sun et al. [9], and the proof of Yuan and Zhang's [10] upper bound cites ideas from our proof of Corollary 1.9.

Corollary 1.9 follows from a *constant*-depth construction of quantum states over a larger gate set. The following class was defined by Green, Homer, Moore and Pollett [12]:

Definition 1.10 ($\mathsf{QAC}^0_\mathsf{f}$ [12])**.** A QAC_f circuit is a quantum circuit consisting of arbitrary one-qubit gates, as well as *generalized Toffoli gates* of arbitrary arity defined by

$$|b, x\rangle \mapsto \left| b \oplus \prod_{j=1}^{n} x_j, x \right\rangle \quad \text{for} \quad b \in \{0,1\}, x = (x_1, \ldots, x_n) \in \{0,1\}^n,$$

and *fanout gates* of arbitrary arity defined by

$$|b, x\rangle \mapsto |b, x \oplus b^n\rangle \quad \text{for} \quad b \in \{0,1\}, x \in \{0,1\}^n.$$

A $\mathsf{QAC}^0_\mathsf{f}$ circuit is a constant-depth QAC_f circuit.

Analogously to in classical circuit complexity, one motivation for studying restricted quantum circuit classes such as $\mathsf{QAC}^0_\mathsf{f}$ circuits is that they seem potentially easier to prove lower bounds against than general QNC circuits. $\mathsf{QAC}^0_\mathsf{f}$ circuits can trivially simulate AC^0 circuits [12] (i.e. constant-depth Boolean circuits with NOT gates and unbounded-fanin AND and OR gates), and in fact are even more powerful than their classical counterparts, because polynomial-size $\mathsf{QAC}^0_\mathsf{f}$ circuits can also compute the majority function [13, 14] whereas AC^0 circuits require exponential size to do so [15, 16]. We prove the following:

Theorem 1.11. *Every n-qubit state can be constructed by a $\mathsf{QAC}^0_\mathsf{f}$ circuit with $\tilde{O}(2^n)$ ancillae.*

We will see that QNC circuits of logarithmic depth can efficiently simulate $\mathsf{QAC}^0_\mathsf{f}$ circuits, so Theorem 1.11 implies Corollary 1.9.

It is well known that every function from n bits to one bit can be computed by a DeMorgan circuit of depth $O(n)$ and size $O(2^n/n)$, and that for most functions this upper bound is tight [16–18].[2] The above results can be seen as progress toward analogous statements about the quantum circuit complexity of constructing quantum states and implementing unitary transformations.

[2]The $\Omega(n)$ depth lower bound follows from the $\Omega(2^n/n)$ size lower bound, because any Boolean circuit of depth d has size less than 2^d.

1.3 Lower bound for implementing U given a U-CC

We prove the following, where by C^A we mean C with A and $A^\dagger$ oracles:

Theorem 1.12. *For all sequences of quantum circuits $(C_n)_n$ making $o\left(2^{n/2}\right)$ queries to a $2n$-qubit quantum oracle, with probability $1 - o(1)$ over a Haar random n-qubit unitary U, there exists a $2n$-qubit U-CC A such that C_n^A is (in some sense) almost maximally far from implementing U.*

Theorem 1.12 matches the upper bound from Theorem 1.4. Since any U-CC tensored with the identity is also a U-CC, we may replace "$2n$-qubit" with "m-qubit" in Theorem 1.12 for any $m \geq 2n$. However, some restrictions on A are still necessary for a lower bound such as Theorem 1.12 to hold, at least if we allow A to act on more than $2n$ qubits. For example, the unitary A defined by

$$\forall x, y \in \{0,1\}^n, b \in \{0,1\} : A|x, y, b\rangle = \begin{cases} |x\rangle \otimes U|x \oplus y\rangle \otimes |0\rangle & \text{if } b = 0 \\ U|x\rangle \otimes |y\rangle \otimes |1\rangle & \text{if } b = 1 \end{cases}$$

is a U-CC, and can trivially be used to implement U when applied with $b = 1$.

It is well known that unstructured search on a list of length N requires $\Omega(\sqrt{N})$ quantum queries [6], but this does not immediately imply that Theorem 1.4 is tight, since there also exist algorithms that do not simulate unstructured search. As we will explain more precisely when we prove Theorem 1.12, it also takes $\Omega(\sqrt{N})$ quantum queries to compute $\sigma^{-1}(1)$ given query access to a permutation σ of $\{1, \ldots, N\}$ [19, 20], and this almost immediately implies an $\Omega\left(2^{n/2}\right)$ lower bound for Theorem 1.4 when U is a permutation matrix and A is defined appropriately. However this example is unsatisfying if our ultimate goal is to prove lower bounds for Questions 1.1 and 1.2, since n-qubit permutation matrices can be efficiently synthesized with a classical oracle and also implemented by a QNC circuit of depth $O(n)$.[3] In contrast, if *any* family of unitaries is hard to implement in the sense of Questions 1.1 and 1.2, then Haar random unitaries are also hard to implement for the following reason:

Observation 1.13. *Any fixed unitary U can be written as $U = UR \cdot R^\dagger$ pointwise where R is Haar random, so since UR and $R^\dagger$ are also Haar random, the task of implementing U reduces to that of successively implementing two (dependent) Haar random unitaries.*

This reduction, along with the U-CCs from our proofs of Theorems 1.5 and 1.7, shows that an improved upper bound for Theorem 1.4 in the case where U is Haar random would imply improved upper bounds for Questions 1.1 and 1.2 in the general case. (Provided that in this hypothetical improvement to Theorem 1.4, the complexity of the non-query operations is not too large.) However, Theorem 1.12 rules out this approach.

Theorem 1.12 does, however, leave open the possibility of obtaining a tighter upper bound for Questions 1.1 and 1.2 by reducing to some *other* "generalized CC":

Question 1.14. Is there a sequence $(C_n)_n$ of quantum circuits, each making $o\left(2^{n/2}\right)$ queries to a $(p(n) + q(n))$-qubit quantum oracle where $p(n), q(n) = \text{poly}(n)$, such that for all n-qubit unitaries U there exists a family of $q(n)$-qubit states $\Psi = (|\psi_x\rangle)_{x \in \{0,1\}^{p(n)}}$ such that for all "Ψ-CCs" A (i.e. $A|x, 0 \ldots 0\rangle = |x\rangle|\psi_x\rangle$) it holds that C_n^A implements U?

If $\Psi = (|f(x)\rangle)_{x \in \{0,1\}^{p(n)}}$ for some function $f : \{0,1\}^{p(n)} \to \{0,1\}$ then a Ψ-CC computes f, so Question 1.14 is essentially a rephrasing of the unitary synthesis problem with a more modest runtime requirement. However, this rephrasing suggests an approach to proving lower bounds for the unitary synthesis problem by considering increasingly general classes of state families Ψ.

At a high level, we prove Theorem 1.12 by using Observation 1.13 to reduce to the previously mentioned lower bound for the case where U is a permutation matrix.

[3] On input x, first compute $\sigma(x)$ (by querying the oracle, or by simulating an appropriate Boolean circuit, depending on the model of computation), and then uncompute x given $\sigma(x)$ by running a similar procedure in reverse.

1.4 Organization and preliminaries

In Section 2 we prove upper and lower bounds for implementing U given query access to a U-CC and its inverse. In Section 3 we prove an upper bound for the unitary synthesis problem. Finally in Section 4 we prove circuit depth upper bounds for constructing states and implementing unitaries.

We denote the n-qubit identity transformation by I_n, or I when n is implicit. We write $\|\cdot\|$ to denote the 2-norm of a vector or the operator norm of a matrix.

2 Bounds for implementing U given a U-CC

In Sections 2.1 and 2.2 we prove upper and lower bounds respectively on the complexity of implementing a unitary U, given query access to a U-CC and its inverse. First we define the quantum query model more precisely. By a *quantum circuit making k queries to an n-qubit quantum oracle*, we mean a circuit of the form $C = C_k Q_k C_{k-1} Q_{k-1} \cdots C_0$ where each C_j is a unitary and each Q_j is a placeholder for either a "forward" or "backward" query. For an n-qubit unitary A, by C^A we mean the unitary defined by substituting A and $A^\dagger$ respectively for the forward and backward queries in C. Claims about the quantum circuit complexity of C are in reference to the circuit $C_k C_{k-1} \cdots C_0$ defined by removing the queries from C. Let $C^\dagger = C_0^\dagger Q_1^\dagger C_1^\dagger Q_2^\dagger \cdots C_k^\dagger$, where the "conjugate transpose" of the forward query symbol is the backward query symbol and vice versa, and note that $\left(C^\dagger\right)^A = \left(C^A\right)^\dagger$.

2.1 Upper bound

We will use a variant of Grover search that finds the marked string with certainty rather than just with high probability:

Lemma 2.1. *There is a uniform sequence of* QAC_f *circuits $(G_n)_n$—each of depth $O\left(2^{n/2}\right)$, making $O\left(2^{n/2}\right)$ queries, and acting on $n+1$ qubits—such that for all $x \in \{0,1\}^n$ it holds that $G_n^{I-2|x,1\rangle\langle x,1|}|0^{n+1}\rangle = |x,0\rangle$.*

Imre and Balázs [21] survey several proofs of the query upper bound from Lemma 2.1 in detail, and below we include an alternate proof of Lemma 2.1 due to Wiebe [22].

Proof. Let $x \in \{0,1\}^n$ denote the marked string that we wish to find, and let

$$t = \left\lceil \frac{\pi}{4} 2^{n/2} \right\rceil, \qquad \theta = \frac{\pi/2}{2t+1}, \qquad p = 2^n \sin^2\theta.$$

Since $p \le 2^n\theta^2 \le 2^n\left(\frac{\pi/2}{2t}\right)^2 \le 1$ we can define states

$$|\psi_0\rangle = |+^n\rangle \otimes \left(\sqrt{1-p}|0\rangle + \sqrt{p}|1\rangle\right), \qquad |\psi_t\rangle = \left((2|\psi_0\rangle\langle\psi_0| - I)(I - 2|x,1\rangle\langle x,1|)\right)^t |\psi_0\rangle$$

where $|+\rangle = \frac{|0\rangle+|1\rangle}{\sqrt{2}}$. Since $\langle x,1|\psi_0\rangle = 2^{-n/2}\sqrt{p} = \sin\theta$, we may write $|\psi_0\rangle = \cos\theta|\alpha\rangle + \sin\theta|\beta\rangle$ where $|\beta\rangle = |x,1\rangle$ and $|\alpha\rangle$ is a superposition of standard basis states besides $|x,1\rangle$. In this basis,

$$\begin{aligned}|\psi_t\rangle &= \left(\begin{pmatrix} 2\cos^2\theta - 1 & 2\cos\theta\sin\theta \\ 2\cos\theta\sin\theta & 2\sin^2\theta - 1 \end{pmatrix}\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}\right)^t \begin{pmatrix} \cos\theta \\ \sin\theta \end{pmatrix} = \begin{pmatrix} \cos 2\theta & -\sin 2\theta \\ \sin 2\theta & \cos 2\theta \end{pmatrix}^t \begin{pmatrix} \cos\theta \\ \sin\theta \end{pmatrix} \\ &= \begin{pmatrix} \cos((2t+1)\theta) \\ \sin((2t+1)\theta) \end{pmatrix} = \begin{pmatrix} \cos(\pi/2) \\ \sin(\pi/2) \end{pmatrix} = \begin{pmatrix} 0 \\ 1 \end{pmatrix},\end{aligned}$$

i.e. $|\psi_t\rangle = |x,1\rangle$. □

Now we reduce the task of implementing a unitary U to that of implementing a U-CC:

Theorem 2.2 (formal version of Theorem 1.4)**.** *There is a uniform family of* QAC_f *circuits $(C_{n,m})_{n,m}$ for $m \ge 2n$—each of depth $O\left(2^{n/2}\right)$, making $O\left(2^{n/2}\right)$ queries to an m-qubit quantum oracle, and acting on $O(m)$ qubits—such that for all n-qubit unitaries U and all m-qubit U-CCs A it holds that $C_{n,m}^A(I_n \otimes |0\ldots0\rangle) = U \otimes |0\ldots0\rangle$.*

Proof. By linearity we may assume that the input is a string $x \in \{0,1\}^n$; our goal is to output $U|x\rangle \otimes |0\dots0\rangle$. First apply A, yielding $|x\rangle \otimes U|x\rangle \otimes |0\dots0\rangle$. The challenge now is to uncompute $|x\rangle$.

Let

$$C = (A \otimes I_1)\big(I_n \otimes \big(I_{m-n+1} - 2|0^{m-n},1\rangle\langle 0^{m-n},1|\big)\big)\big(A^\dagger \otimes I_1\big),$$

and observe that C can be implemented by a $\mathsf{QAC}^0_\mathsf{f}$ circuit making two queries. By the definition of A we have that

$$\begin{aligned} C &= (A \otimes I_1)\left(I_{m+1} - 2\sum_{y\in\{0,1\}^n} |y\rangle\langle y| \otimes |0^{m-n}\rangle\langle 0^{m-n}| \otimes |1\rangle\langle 1|\right)\big(A^\dagger \otimes I_1\big) \\ &= I_{m+1} - 2\sum_{y\in\{0,1\}^n} |y\rangle\langle y| \otimes U|y\rangle\langle y|U^\dagger \otimes |0^{m-2n}\rangle\langle 0^{m-2n}| \otimes |1\rangle\langle 1| \\ &= I_{m+1} - 2\sum_{y\in\{0,1\}^n} |y,1\rangle\langle y,1| \otimes U|y\rangle\langle y|U^\dagger \otimes |0^{m-2n}\rangle\langle 0^{m-2n}|, \end{aligned}$$

where in the last line we reorder the qubits for future convenience. Therefore

$$C\big(I_{n+1} \otimes U|x\rangle \otimes |0^{m-2n}\rangle\big) = (I_{n+1} - 2|x,1\rangle\langle x,1|) \otimes U|x\rangle \otimes |0^{m-2n}\rangle,$$

so using our copy of $U|x\rangle$, the circuit C can implement the reflection $I_{n+1} - 2|x,1\rangle\langle x,1|$ in a disjoint register without disturbing the copy of $U|x\rangle$.

We could therefore simulate the circuit G_n from Lemma 2.1, with queries to $I - 2|x,1\rangle\langle x,1|$ answered in this manner, to construct a copy of $|x\rangle$. Instead perform this simulation *in reverse*, to *uncompute* the existing copy of $|x\rangle$ while preserving the copy of $U|x\rangle$. Finally swap $U|x\rangle$ into the appropriate register. □

2.2 Lower bound

For linear transformations L, M from n qubits to m qubits where $n \le m$ let $\langle L, M\rangle = 2^{-n}\operatorname{tr}\big(L^\dagger M\big)$, i.e. $\langle\cdot,\cdot\rangle$ is the Frobenius inner product normalized such that $\langle A, A\rangle = 1$ for all isometries A.

Theorem 2.3 (formal version of Theorem 1.12)**.** *For all sequences of quantum circuits $(C_n)_n$ making $o\big(2^{n/2}\big)$ queries to a $2n$-qubit quantum oracle, with probability $1-o(1)$ over a Haar random n-qubit unitary U, there exists a $2n$-qubit U-CC A such that for all states $|\psi\rangle$,*

$$\big|\big\langle C_n^A(I_n \otimes |0\dots0\rangle), U \otimes |\psi\rangle\big\rangle\big| \le o(1).$$

Proof. For a permutation σ of $\{0,1\}^n$ let A_σ be the unitary defined by $A_\sigma|x,y\rangle = |x, y \oplus \sigma(x)\rangle$ for all $x, y \in \{0,1\}^n$. Nayak [20, Corollary 1.2] proved that any quantum circuit making $o\big(2^{n/2}\big)$ queries to A_σ outputs $\sigma^{-1}(0^n)$ with probability less than $1/2$, where the probability is over a uniform random permutation σ of $\{0,1\}^n$ as well as the randomness of the output measurement. (We remark that Ambainis [19] previously proved a similar result using different techniques.)

Let $\varepsilon, \delta > 0$ be universal constants, and assume for the sake of contradiction that there exists a quantum circuit C making $o\big(2^{n/2}\big)$ queries to a $2n$-qubit quantum oracle, such that with probability at least ε over a Haar random n-qubit unitary U, for all $2n$-qubit U-CCs A, there exists a state $|\psi\rangle$ such that $\big|\big\langle C^A(I_n \otimes |0\dots0\rangle), U \otimes |\psi\rangle\big\rangle\big| \ge \delta$. We will prove that there exists a quantum circuit making $o\big(2^{n/2}\big)$ queries to A_σ that outputs $\sigma^{-1}(0^n)$ with probability $\Omega(1)$, where the probability is over a uniform random permutation σ of $\{0,1\}^n$ as well as the randomness of the output measurement. By executing this circuit constantly many times until it outputs $\sigma^{-1}(0^n)$, we can boost the success probability to be greater than $1/2$ which contradicts Nayak's result. Therefore no such circuit C exists.

Write $C = C_s Q_s C_{s-1} Q_{s-1} \cdots C_0$ for $s = o\big(2^{n/2}\big)$, where each C_i is a unitary and each Q_i is a placeholder for either a forward or backward query. For an n-qubit unitary R, define a quantum circuit C_R by replacing each forward query Q_i in C with $(I_n \otimes R)Q_i$, and replacing each backward query Q_i in C with $Q_i\big(I_n \otimes R^\dagger\big)$. For a permutation σ of $\{0,1\}^n$ let P_σ denote the corresponding permutation matrix on n qubits, i.e. $P_\sigma|x\rangle = |\sigma(x)\rangle$ for all $x \in \{0,1\}^n$. Clearly for all R, σ it holds

that $C_R^{A_\sigma} = C^{(I_n \otimes R)A_\sigma}$, and that $(I_n \otimes R)A_\sigma$ is a $2n$-qubit RP_σ-CC. If σ is fixed and R is Haar random, then RP_σ is also Haar random, so for some state $|\psi\rangle$ depending on R and σ,

$$\Pr_R\left(\left|\left\langle C_R^{A_\sigma}(I_n \otimes |0\dots0\rangle), RP_\sigma \otimes |\psi\rangle\right\rangle\right| \geq \delta\right) \geq \varepsilon.$$

Call a fixed unitary R "good with respect to σ" if $\left|\left\langle C_R^{A_\sigma}(I_n \otimes |0\dots0\rangle), RP_\sigma \otimes |\psi\rangle\right\rangle\right| \geq \delta$ for some $|\psi\rangle$. Also let $D_R = C_R^\dagger(R \otimes |\psi\rangle)$. (For intuition, if R is good with respect to σ then $C_R^{A_\sigma}$ approximately implements RP_σ, and so D_R approximately implements $P_{\sigma^{-1}}$.) If R is good with respect to σ then

$$\begin{aligned}
\delta &\leq \left|2^{-n}\operatorname{tr}\left((I_n \otimes \langle 0\dots0|)\left(C_R^{A_\sigma}\right)^\dagger(RP_\sigma \otimes |\psi\rangle)\right)\right| && \text{(definition of } \langle\cdot,\cdot\rangle\text{)}\\
&= \left|2^{-n}\sum_{x\in\{0,1\}^n}\langle x, 0\dots0|\left(C_R^{A_\sigma}\right)^\dagger(RP_\sigma|x\rangle \otimes |\psi\rangle)\right| && \text{(definition of trace)}\\
&\leq 2^{-n}\sum_{x\in\{0,1\}^n}\left|\langle x, 0\dots0|D_R^{A_\sigma}|\sigma(x)\rangle\right| && \text{(triangle ineq., definitions of } D_R, P_\sigma\text{)}\\
&\leq 2^{-n}\sum_{x\in\{0,1\}^n}\left\|\left(\langle\sigma^{-1}(x)| \otimes I\right)D_R^{A_\sigma}|x\rangle\right\| && \text{(Cauchy-Schwarz, } x \leftarrow \sigma^{-1}(x)\text{).}
\end{aligned}$$

For $x \in \{0,1\}^n$ let $p_{\sigma,R,x}$ be the probability that if we run $D_R^{A_\sigma}$ on input x and measure the first n qubits of the output state, then the result is $\sigma^{-1}(x)$. Then we can phrase the above inequality as $\delta \leq 2^{-n}\sum_{x\in\{0,1\}^n}\sqrt{p_{\sigma,R,x}}$, and by Cauchy-Schwarz it follows that $\delta^2 \leq 2^{-n}\sum_{x\in\{0,1\}^n} p_{\sigma,R,x}$.

Therefore for every fixed permutation σ of $\{0,1\}^n$, for Haar random R and uniform random $x \in \{0,1\}^n$, it holds that

$$\varepsilon \leq \Pr_R(R \text{ is good w.r.t. } \sigma) \leq \Pr_R\left(\delta^2 \leq \mathop{\mathbb{E}}_x[p_{\sigma,R,x}]\right) \leq \delta^{-2}\mathop{\mathbb{E}}_{R,x}[p_{\sigma,R,x}]$$

where the last step is by Markov's inequality. If we also take σ to be uniform random then $\mathbb{E}_{\sigma,R,x}[p_{\sigma,R,x}] \geq \varepsilon\delta^2$, so there exist *fixed* values of R and x such that $\mathbb{E}_\sigma[p_{\sigma,R,x}] \geq \varepsilon\delta^2$. Thus there exists a quantum circuit (specifically $D_R^{A_\sigma}|x\rangle$ for these fixed values of R and x) making $o\left(2^{n/2}\right)$ queries to A_σ that outputs $\sigma^{-1}(x)$ with probability at least $\varepsilon\delta^2$, where the probability is over a uniform random permutation σ of $\{0,1\}^n$ as well as the randomness of the output measurement. By symmetry such a circuit exists with $x = 0^n$ as desired. □

3 Upper bound for the unitary synthesis problem

We prove the following:

Theorem 3.1 (formal version of Theorem 1.5)**.** *Let $\varepsilon(n) = \exp(-\mathrm{poly}(n))$. Then there is a uniform sequence of QAC_f circuits $(C_n)_n$—each of depth $O\left(2^{n/2}\right)$, making $O\left(2^{n/2}\right)$ queries, and with $\mathrm{poly}(n)$ ancillae—such that for all n-qubit unitaries U there exists a classical oracle f such that $\left\|C_n^f(I_n \otimes |0\dots0\rangle) - U \otimes |0\dots0\rangle\right\| \leq \varepsilon(n)$.*

Although Theorem 3.1 is stated in terms of QAC_f circuits for convenience, a similar statement for QNC circuits follows easily using Lemma A.1. A similar statement also holds for diamond norm error: For an isometry A define a channel $\mathcal{C}_A$ by $\mathcal{C}_A(\rho) = A\rho A^\dagger$. For all states $|\psi\rangle$ and $|\phi\rangle$ it holds that $\frac{1}{2}\||\psi\rangle\langle\psi| - |\phi\rangle\langle\phi|\|_1 \leq \||\psi\rangle - |\phi\rangle\|$ [5, Eq. 2.3], so for all isometries A and B it holds that

$$\begin{aligned}
\frac{1}{2}\|\mathcal{C}_A - \mathcal{C}_B\|_\diamond &= \frac{1}{2}\max_{|\psi\rangle}\|((\mathcal{C}_A - \mathcal{C}_B) \otimes I)(|\psi\rangle\langle\psi|)\|_1\\
&= \frac{1}{2}\max_{|\psi\rangle}\left\|(A \otimes I)|\psi\rangle\langle\psi|\left(A^\dagger \otimes I\right) - (B \otimes I)|\psi\rangle\langle\psi|\left(B^\dagger \otimes I\right)\right\|_1\\
&\leq \max_{|\psi\rangle}\|(A \otimes I)|\psi\rangle - (B \otimes I)|\psi\rangle\|
\end{aligned}$$

$$= \|(A - B) \otimes I\|$$
$$= \|A - B\|. \tag{1}$$

In particular, if we define a channel $\mathcal{C}_n^f$ by tracing out the ancillae after applying $C_n^f(I_n \otimes |0 \dots 0\rangle)$, then

$$\frac{1}{2}\left\|\mathcal{C}_n^f - \mathcal{C}_U\right\|_\diamond \leq \frac{1}{2}\left\|\mathcal{C}_{C_n^f(I_n \otimes |0\dots 0\rangle)} - \mathcal{C}_{U \otimes |0\dots 0\rangle}\right\|_\diamond \leq \varepsilon(n),$$

where the first inequality holds because tracing out qubits at the end cannot increase the diamond norm distance between two channels, and the second inequality is by Eq. (1) and Theorem 3.1.

Queries to a classical oracle (i.e. a Boolean function) can be modeled in either of two standard ways. In the first, a function $f : \{0,1\}^n \to \{0,1\}^m$ is encoded as the oracle U_f defined by $U_f|x, y\rangle = |x, y \oplus f(x)\rangle$. In the second, which is only applicable when $m = 1$, the function f is instead encoded as the oracle V_f defined by $V_f|x\rangle = (-1)^{f(x)}|x\rangle$. These models are equivalent, because $V_f = (I_n \otimes \langle -|)U_f(I_n \otimes |-\rangle)$ where $|-\rangle = \frac{|0\rangle - |1\rangle}{\sqrt{2}}$, and if $g(x, y) = \bigoplus_{j=1}^m f(x)_j y_j$ (where the subscript j indicates the j-th bit of an m-bit string) then $U_f = (I_n \otimes H^{\otimes m})V_g(I_n \otimes H^{\otimes m})$ where H denotes the Hadamard gate. We write C^f to abbreviate C^{U_f} or C^{V_f}, as defined in the beginning of Section 2; since U_f and V_f are Hermitian we do not need to distinguish between forward and backward queries.

Our proof uses the following result:[4]

Theorem 3.2 (Rosenthal [5])**.** *Let $\varepsilon(n) = \exp(-\text{poly}(n))$. Then there is a uniform sequence of $\text{poly}(n)$-qubit $\mathsf{QAC}^0_\mathsf{f}$ circuits $(C_n)_n$, each making four queries, such that for all n-qubit states $|\psi\rangle$ there exists a classical oracle f such that $\left\|C_n^f|0\dots 0\rangle - |\psi\rangle|0\dots 0\rangle\right\| \leq \varepsilon(n)$.*

Our proof also uses the following lemma, which says that if the task of implementing an isometry J reduces to that of implementing an isometry A, then the task of *approximately* implementing J reduces to that of *approximately* implementing A:

Lemma 3.3. *Let C be an $(m+a)$-qubit quantum circuit making k queries to an n-qubit quantum oracle, and let J be an isometry from m qubits to $m+a$ qubits. Assume there exists a subspace $S \subseteq \left(\mathbb{C}^2\right)^{\otimes n}$ and an isometry $A : S \to \left(\mathbb{C}^2\right)^{\otimes n}$ such that for all n-qubit unitaries U consistent with A it holds that $C^U(I_m \otimes |0^a\rangle) = J$. Then for all isometries $B : S \to \left(\mathbb{C}^2\right)^{\otimes n}$ and all n-qubit unitaries V consistent with B, it holds that $\left\|C^V(I_m \otimes |0^a\rangle) - J\right\| \leq \sqrt{2} \cdot k\|A - B\|$.*

First we prove Theorem 3.1 assuming Lemma 3.3, and then we prove Lemma 3.3.

Proof of Theorem 3.1. Theorem 3.2 generalizes from constructing states to implementing $\text{poly}(n)$-qubit U-CCs (for an n-qubit unitary U), because if f_x is the oracle associated with constructing $U|x\rangle$ in Theorem 3.2, then the function $(x, y) \mapsto f_x(y)$ can simulate queries to f_x controlled on x. Therefore there is a uniform sequence of $\text{poly}(n)$-qubit $\mathsf{QAC}^0_\mathsf{f}$ circuits $(A_n)_n$, each making four queries, such that for all n-qubit unitaries U there exists a classical oracle f such that

$$\max_{x \in \{0,1\}^n} \left\|A_n^f|x, 0^{n+m}\rangle - |x\rangle \otimes U|x\rangle \otimes |0^m\rangle\right\| \leq \varepsilon(n)/\left(c2^n \cdot \sqrt{2}\right)$$

for $m = \text{poly}(n)$. Here c is a constant such that the circuit in Theorem 2.2 makes at most $c2^{n/2}$ queries. Since the operator norm of a matrix is at most the Frobenius norm, it follows that

$$\left\|A_n^f\left(I_n \otimes |0^{n+m}\rangle\right) - \sum_{x \in \{0,1\}^n} |x\rangle\langle x| \otimes U|x\rangle \otimes |0^m\rangle\right\| \leq \sqrt{\sum_{x \in \{0,1\}^n} \left\|A_n^f|x, 0^{n+m}\rangle - |x\rangle \otimes U|x\rangle \otimes |0^m\rangle\right\|^2}$$
$$\leq \varepsilon(n)/\left(c2^{n/2} \cdot \sqrt{2}\right).$$

The result then follows by Theorem 2.2 and Lemma 3.3, the latter applied with isometries $J = U \otimes |0\dots 0\rangle$ and A a U-CC and $B = A_n^f(I_n \otimes |0^{n+m}\rangle)$. □

[4] An earlier version of our proof used a similar result of Aaronson [1, Proposition 3.3.5] that required $O(n)$ queries, resulting in a multiplicative $O(n)$ blowup in the query complexity in Theorem 3.1.

3.1 Proof of Lemma 3.3

We use the fact that

$$\|U_m U_{m-1} \cdots U_1 - V_m V_{m-1} \cdots V_1\| \leq \sum_{j=1}^{m} \|U_j - V_j\| \tag{2}$$

for all unitaries U_j, V_j by the triangle inequality [6, Eq. 4.69]. (The reason that Lemma 3.3 does not trivially follow from Eq. (2), and without the $\sqrt{2}$ factor, is that some of the states acted on by applications of V in $C^V(I_m \otimes |0^a\rangle)$ might be far from S.)

For a unitary $U \in \mathbb{C}^{n\times n}$ and a subspace $S \subseteq \mathbb{C}^n$, we write $U_{|S}$ to denote the isometry from S to $\mathbb{C}^n$ defined by restricting the domain of U to S. The following inequality is trivially tight to within a $\sqrt{2}$ factor:

Claim 3.4. *For all unitaries $V \in \mathbb{C}^{n\times n}$, subspaces $S \subseteq \mathbb{C}^n$, and isometries $W : S \to \mathbb{C}^n$, there exists a unitary $U \in \mathbb{C}^{n\times n}$ such that $U_{|S} = W$ and $\|U - V\| \leq \sqrt{2}\big\|W - V_{|S}\big\|$.*

Lemma 3.3 follows immediately from Eq. (2) and Claim 3.4, along with the fact that $\big\|U^\dagger - V^\dagger\big\| = \|U - V\|$ (to handle backward queries).

Proof of Claim 3.4. Let $P\Sigma Q^\dagger$ be a singular value decomposition of $W^\dagger V_{|S}$, and write $\Sigma = \mathrm{diag}(\sigma_1, \ldots, \sigma_k)$ where $k = \dim(S)$ and $\sigma_1 \geq \cdots \geq \sigma_k \geq 0$. By definition $P, Q : \mathbb{C}^k \to S$ are (bijective) isometries such that $P\Sigma Q^\dagger = W^\dagger V_{|S}$. Then

$$\begin{aligned}\big\|W - V_{|S}\big\| &\geq \big\|(W - V_{|S})Q|k\rangle\big\| \geq \sqrt{2 - 2\big|\langle k|Q^\dagger W^\dagger V_{|S} Q|k\rangle\big|} = \sqrt{2 - 2\sigma_k |\langle k|Q^\dagger P|k\rangle|} \\ &\geq \sqrt{2 - 2\sigma_k},\end{aligned}$$

where the last inequality is by Cauchy-Schwarz.

Define isometries $A, B \in \mathbb{C}^{n\times k}$ by $A = WP$ and $B = V_{|S}Q$, and note that $A^\dagger B = P^\dagger W^\dagger V_{|S} Q = \Sigma$. Let

$$\overline{A} = \mathrm{trnc}\Big[(B - A\Sigma)\big(I - \Sigma^2\big)^{-1/2}\Big], \qquad \overline{B} = \mathrm{trnc}\Big[(B\Sigma - A)\big(I - \Sigma^2\big)^{-1/2}\Big],$$

where the "truncation" operator trnc[·] removes the j'th column of a matrix for all j such that $\sigma_j = 1$ (thus avoiding division by zero). It is straightforward to verify that $\overline{A}, \overline{B}$ are isometries satisfying $A^\dagger \overline{A} = B^\dagger \overline{B} = 0$ and $\overline{A}^\dagger \overline{B} = \mathrm{trnc}[\Sigma]$, so

$$\big\|\overline{A} - \overline{B}\big\| = \max_{\||\psi\rangle\|=1} \big\|\overline{A}|\psi\rangle - \overline{B}|\psi\rangle\big\| = \max_{\||\psi\rangle\|=1} \sqrt{2 - 2\langle\psi|\mathrm{trnc}[\Sigma]|\psi\rangle} = \sqrt{2 - 2\sigma_k} \leq \big\|W - V_{|S}\big\|.$$

Let $S' \subseteq \mathbb{C}^n$ be the subspace that V maps to the image of $\overline{B}$. The subspaces S and S' are orthogonal, because $B^\dagger \overline{B} = 0$ and V maps S to the image of B, so we can write $\mathbb{C}^n = S \oplus S' \oplus S''$ for some subspace S''. Let $R : S' \to \mathbb{C}^{|\{j:\sigma_j \neq 1\}|}$ be the (bijective) isometry such that $V_{|S'} = \overline{B}R$, and define a unitary $U \in \mathbb{C}^{n\times n}$ by

$$U_{|S} = W, \qquad U_{|S'} = \overline{A}R, \qquad U_{|S''} = V_{|S''}.$$

To see that U is in fact unitary, note that $W^\dagger \cdot \overline{A}R = PA^\dagger \overline{A}R = 0$, and that the image of $U_{|S\oplus S'}$ (i.e. the span of A and $\overline{A}$) equals the image of $V_{|S\oplus S'}$ (i.e. the span of B and $\overline{B}$). The latter condition holds because for all columns j, if $\sigma_j \neq 1$ then $\mathrm{span}(A|j\rangle, \overline{A}|j\rangle) = \mathrm{span}(A|j\rangle, B|j\rangle) = \mathrm{span}(B|j\rangle, \overline{B}|j\rangle)$, and if $\sigma_j = 1$ then $A|j\rangle = B|j\rangle$ (recalling that $A^\dagger B = \Sigma$) and $\overline{A}|j\rangle, \overline{B}|j\rangle$ have not been defined (recalling the definition of trnc[·]).

Let $|\psi\rangle \in \mathbb{C}^n$ be a unit vector such that $\|U - V\| = \|(U - V)|\psi\rangle\|$, and for a subspace $T \subseteq \mathbb{C}^n$ let $|\psi_T\rangle$ be the projection of $|\psi\rangle$ onto T. Then by the triangle inequality and Cauchy-Schwarz,

$$\begin{aligned}\|U - V\| = \|(U - V)|\psi\rangle\| &\leq \sum_{T\in\{S,S',S''\}} \big\|(U - V)_{|T}|\psi_T\rangle\big\| \leq \big\|W - V_{|S}\big\| \cdot \||\psi_S\rangle\| + \big\|\overline{A} - \overline{B}\big\| \cdot \||\psi_{S'}\rangle\| \\ &\leq \big\|W - V_{|S}\big\|(\||\psi_S\rangle\| + \||\psi_{S'}\rangle\|) \leq \sqrt{2}\big\|W - V_{|S}\big\| \cdot \||\psi_{S\oplus S'}\rangle\| \leq \sqrt{2}\big\|W - V_{|S}\big\|.\end{aligned}$$

□

4 Circuit depth upper bounds for states and unitaries

In Sections 4.1 and 4.2 respectively we prove circuit depth upper bounds for constructing arbitrary states and implementing arbitrary unitaries.

4.1 States

We prove the following:

Theorem 4.1 (formal version of Theorem 1.11)**.** *For all n-qubit states $|\psi\rangle$ there exists an $\tilde{O}(2^n)$-qubit $\mathsf{QAC}^0_\mathsf{f}$ circuit C such that $C|0\dots0\rangle = |\psi\rangle|0\dots0\rangle$.*

Theorem 4.1 and Lemma A.1 imply the following:

Corollary 4.2 (formal version of Corollary 1.9)**.** *For all n-qubit states $|\psi\rangle$ there exists an $O(n)$-depth, $\tilde{O}(2^n)$-qubit QNC circuit C such that $C|0\dots0\rangle = |\psi\rangle|0\dots0\rangle$.*

A proof sketch of Theorem 4.1 is as follows. First consider the analogous problem of sampling a string s from a given distribution over $\{0,1\}^n$. One way to sample s is to first sample

$$b_x \sim \text{Bernoulli}(\Pr(s \text{ begins with } x1 \mid s \text{ begins with } x))$$

independently for all binary strings x of length less than n, and then output the string y defined by $y_i = b_{y_1 y_2 \cdots y_{i-1}}$ for i from 1 to n. Furthermore each bit of y can be computed by a DNF formula of size $\tilde{O}(2^n)$ as a function of $(b_x)_x$. Similarly we can construct a quantum state $\sum_{y\in\{0,1\}^n} \alpha_y|y\rangle$ using unentangled one-qubit states in place of $(b_x)_x$; this actually yields a state of the form $\sum_{y\in\{0,1\}^n} \alpha_y|y\rangle|\text{garbage}_y\rangle$, but it turns out that $|\text{garbage}_y\rangle$ can be efficiently uncomputed controlled on y.

Our proof will use the following notation. Let $\{0,1\}^{\le n}$ (resp. $\{0,1\}^{<n}$) denote the set of strings of length at most (resp. less than) n over $\{0,1\}$, including the empty string ϵ. For $x \in \{0,1\}^*$ let $x_k, x_{<k}, x_{\le k}$ respectively denote the k-th bit, first $k-1$ bits, and first k bits of x, and let $|x|$ denote the length of x. For $x, y \in \{0,1\}^*$ let xy denote the concatenation of x and y.

Proof. Let $|\psi\rangle = \sum_{x\in\{0,1\}^n} \alpha_x|x\rangle$ denote the n-qubit state to be constructed, and define "conditional amplitudes" β_x for $x \in \{0,1\}^{\le n}\backslash\{\epsilon\}$ as follows: Let $|\psi_\epsilon\rangle = |\psi\rangle$, and for $x \in \{0,1\}^{<n}$, given an $(n-|x|)$-qubit state $|\psi_x\rangle$, write

$$|\psi_x\rangle = \begin{cases} \beta_{x0}|0\rangle|\psi_{x0}\rangle + \beta_{x1}|1\rangle|\psi_{x1}\rangle & \text{if } |x| \le n-2, \\ \beta_{x0}|0\rangle + \beta_{x1}|1\rangle & \text{if } |x| = n-1 \end{cases}$$

for $(n-|x|-1)$-qubit states $|\psi_{x0}\rangle, |\psi_{x1}\rangle$ (if $|x| \le n-2$) and complex numbers β_{x0}, β_{x1} such that $|\beta_{x0}|^2 + |\beta_{x1}|^2 = 1$. Let

$$|\phi_x\rangle = \beta_{x0}|0\rangle + \beta_{x1}|1\rangle$$

for $x \in \{0,1\}^{<n}$, and observe that $\alpha_x = \prod_{i=1}^n \beta_{x_{\le i}}$ for all $x \in \{0,1\}^n$.

Let $f : \{0,1\}^{\{0,1\}^{<n}} \to \{0,1\}^n$ be the function defined by $f(x)_i = x_{f(x)_{<i}}$ for i from 1 to n. The function f is illustrated in Fig. 1 and can be computed as follows:

$$f(x)_j = \bigvee_{\substack{t\in\{0,1\}^j \\ t_j=1}} \bigwedge_{1\le i\le j} \mathbb{1}_{x_{t_{<i}}=t_i} \qquad \text{for } 1 \le j \le n.$$

(The conjunction indicates whether t equals the first j bits of $f(x)$, and the disjunction indicates whether the satisfying t is such that $t_j = 1$.) Each of the n AC^0 formulas has $\sum_{j=1}^n 2^{j-1}j \le \tilde{O}(2^n)$ leafs, so a $\mathsf{QAC}^0_\mathsf{f}$ circuit on $\tilde{O}(2^n)$ qubits can compute the unitary

$$U_f|x,a\rangle = |x, a \oplus f(x)\rangle \qquad \text{for } x \in \{0,1\}^{\{0,1\}^{<n}}, a \in \{0,1\}^n$$

by first fanning out the input bits to make as many copies as needed, then simulating the gates in the AC^0 formulas, and finally uncomputing the garbage.

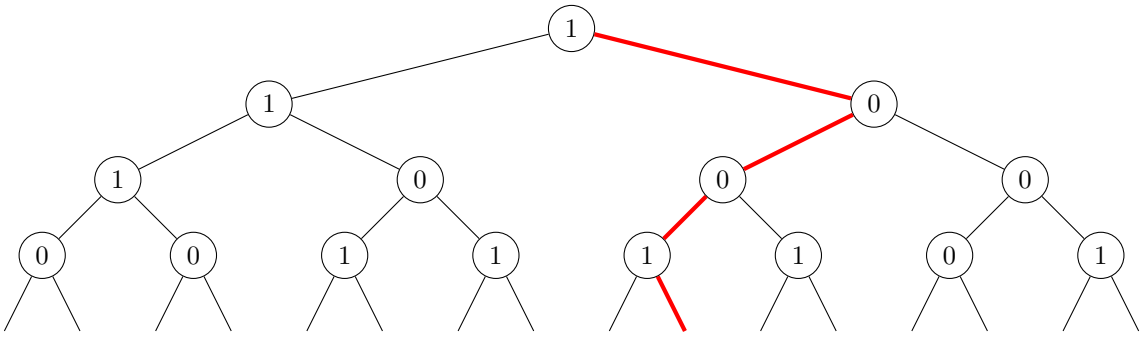

Figure 1: Nodes are labeled with the inputs to f. The highlighted path represents the output of f, and is defined by starting at the root and repeatedly walking to the left or right child depending on whether the current node is labeled 0 or 1.

Let $(\mathsf{R}_x)_{x\in\{0,1\}^{<n}}$ be one-qubit registers and let S be an n-qubit register. The first step toward constructing $|\psi\rangle$ is to construct the state

$$U_f\left(\bigotimes_{x\in\{0,1\}^{<n}} |\phi_x\rangle_{\mathsf{R}_x} \otimes |0^n\rangle_{\mathsf{S}}\right),$$

using a layer of one-qubit gates followed by the aforementioned circuit for U_f. Here, when computing U_f, the x-th input bit to f is in R_x for all x, and the output register of f is S. Observe that

$$U_f(I\otimes|0^n\rangle) = \sum_{x\in\{0,1\}^{\{0,1\}^{<n}}} |x\rangle\langle x|\otimes|f(x)\rangle = \sum_{t\in\{0,1\}^n}\left(\sum_{x\in f^{-1}(t)}|x\rangle\langle x|\right)\otimes|t\rangle = \sum_{t\in\{0,1\}^n}\left(\bigotimes_{i=1}^{n}|t_i\rangle\langle t_i|_{\mathsf{R}_{t_{<i}}}\right)\otimes|t\rangle_{\mathsf{S}},$$

where the t-th tensor product above implicitly acts as the identity on all R_x for which x does not equal $t_{<i}$ for any i. Therefore

$$U_f\left(\bigotimes_{x\in\{0,1\}^{<n}} |\phi_x\rangle_{\mathsf{R}_x}\otimes|0^n\rangle_{\mathsf{S}}\right) = \sum_{t\in\{0,1\}^n}\bigotimes_{x\in\{0,1\}^{<n}}\begin{cases}|t_i\rangle\langle t_i|\phi_{t_{<i}}\rangle_{\mathsf{R}_x} & \text{if } x=t_{<i} \text{ for some } i\\ |\phi_x\rangle_{\mathsf{R}_x} & \text{otherwise}\end{cases}\otimes|t\rangle_{\mathsf{S}}.$$

By the definition of $|\phi_{t_{<i}}\rangle$ it holds that $\langle t_i|\phi_{t_{<i}}\rangle = \beta_{t_{<i}t_i} = \beta_{t_{\le i}}$, so since $\alpha_t = \prod_{i=1}^n \beta_{t_{\le i}}$ for all $t\in\{0,1\}^n$ it follows that

$$U_f\left(\bigotimes_{x\in\{0,1\}^{<n}} |\phi_x\rangle_{\mathsf{R}_x}\otimes|0^n\rangle_{\mathsf{S}}\right) = \sum_{t\in\{0,1\}^n}\alpha_t\bigotimes_{x\in\{0,1\}^{<n}}\begin{cases}|t_i\rangle_{\mathsf{R}_x} & \text{if } x=t_{<i} \text{ for some } i\\ |\phi_x\rangle_{\mathsf{R}_x} & \text{otherwise}\end{cases}\otimes|t\rangle_{\mathsf{S}}.$$

All that remains to construct the state $|\psi\rangle = \sum_{t\in\{0,1\}^n}\alpha_t|t\rangle$ is to uncompute the above content of $(\mathsf{R}_x)_{x\in\{0,1\}^{<n}}$ controlled on the state $|t\rangle$ of S. To do so, first make $|\{0,1\}^{<n}|$ copies of t using fanout. Then for each $x\in\{0,1\}^{<n}$ in parallel, controlled on one of these copies of t, if $x=t_{<i}$ for some i then perform in R_x an operation that maps $|t_i\rangle$ to $|0\rangle$, and otherwise perform in R_x an operation that maps $|\phi_x\rangle$ to $|0\rangle$. Finally, uncompute the extra copies of t using fanout. □

4.2 Unitaries

First we establish some basic properties of QAC_f circuits:

Lemma 4.3. *There is a uniform family of $O(mn\log n)$-qubit $\mathsf{QAC}^0_\mathsf{f}$ circuits $(C_{n,m})_{n,m}$, where $C_{n,m}$ takes as input a $(\log n)$-qubit register K and m-qubit registers $\mathsf{A}_0,\dots,\mathsf{A}_{n-1},\mathsf{B}$ (and ancillae) and $C_{n,m}$ swaps A_k and B controlled on the classical state $|k\rangle_\mathsf{K}$.*

Proof. We can assume without loss of generality that $m = 1$, because then the general case follows by swapping the i-th qubits of A_k and B for all i in parallel. By linearity we may assume that the input is a standard basis state

$$|k\rangle_{\mathsf{K}}|x_0\rangle_{\mathsf{A}_0}\cdots|x_{n-1}\rangle_{\mathsf{A}_{n-1}}|y\rangle_{\mathsf{B}}.$$

For now assume that y is promised to be 0. First compute $x_k = \bigvee_{j=0}^{n-1}(\mathbb{1}_{j=k} \wedge x_j)$ in B, using that $\mathsf{QAC}^0_{\mathsf{f}}$ circuits can simulate AC^0 circuits; note that comparing j and k requires $O(\log n)$ qubits for any given value of j. Then controlled on the state $|x_k\rangle_{\mathsf{B}}$, make n copies of x_k with fanout, use them to XOR the bit $\mathbb{1}_{j=k} \wedge x_k$ into A_j for all $j < n$ in parallel, and finally uncompute the extra copies of x_k using fanout.

For the general case where y might not be 0, let C be a one-qubit register in the ancillae. First swap A_k and C as described above, then swap B and A_k as described above, and finally swap C and B. □

Lemma 4.4. *If C is an n-qubit, size-s, depth-d $\mathsf{QAC}_{\mathsf{f}}$ circuit then controlled-C can be implemented by an $O(n)$-qubit, size-$O(s)$, depth-$O(d)$ $\mathsf{QAC}_{\mathsf{f}}$ circuit.*

Proof. Controlled on a bit $b \in \{0,1\}$, each gate in a $\mathsf{QAC}_{\mathsf{f}}$ circuit can be implemented controlled on b as follows. A k-qubit generalized Toffoli gate controlled on b is equivalent to a $(k+1)$-qubit generalized Toffoli gate, fanning out a bit c controlled on b is equivalent to fanning out bc, and applying a one-qubit gate controlled on b is a two-qubit operation and can therefore be implemented with a constant number of one-qubit and CNOT gates [6, Section 4.5.2]. The result follows by making n copies of b, and using these copies to implement all gates in a given layer of C in parallel controlled on b, where the same ancillae are reused in simulations of successive layers of C. □

Now we prove that $O\big(2^{n/2}\big)$-depth $\mathsf{QAC}_{\mathsf{f}}$ circuits can implement any n-qubit unitary:

Theorem 4.5. *For all n-qubit unitaries U there exists an $O\big(2^{n/2}\big)$-depth, $\tilde{O}\big(2^{2n}\big)$-qubit $\mathsf{QAC}_{\mathsf{f}}$ circuit C such that $C(I_n \otimes |0\ldots0\rangle) = U \otimes |0\ldots0\rangle$.*

Proof. By Theorem 2.2 it suffices to implement a U-CC with an $\tilde{O}\big(2^{2n}\big)$-qubit $\mathsf{QAC}^0_{\mathsf{f}}$ circuit, and this can be achieved as follows. On input $x \in \{0,1\}^n$ to the U-CC, first make 2^n copies of x using fanout. Then for all $y \in \{0,1\}^n$ in parallel, in a register R_y use Theorem 4.1 and Lemma 4.4 to construct $U|y\rangle$ controlled on $x = y$, using a different copy of x for each string y. Then swap R_x into the output register using Lemma 4.3, and use fanout to uncompute the extra copies of x. □

Theorem 4.5 and Lemma A.1 imply the following:

Corollary 4.6 (formal version of Theorem 1.7)**.** *For all n-qubit unitaries U there exists an $\tilde{O}\big(2^{n/2}\big)$-depth, $\tilde{O}\big(2^{2n}\big)$-qubit QNC circuit C such that $C(I_n \otimes |0\ldots0\rangle) = U \otimes |0\ldots0\rangle$.*

Acknowledgments

Thanks to Scott Aaronson, Karen J. Morenz Korol, Fermi Ma, Adrian She, Nathan Wiebe, and Henry Yuen for helpful discussions. Part of this work was done while the author was visiting the Simons Institute for the Theory of Computing.

A QNC simulation of $\mathsf{QAC}_{\mathsf{f}}$ circuits

Lemma A.1. *For all n-qubit, depth-d $\mathsf{QAC}_{\mathsf{f}}$ circuits U, there exists an $O(n)$-qubit, depth-$O(d \log n)$, size-$O(dn)$ QNC circuit C such that $C(I_n \otimes |0\ldots0\rangle) = U \otimes |0\ldots0\rangle$.*

Proof. Green et al. [12] observed that the transformation $|b, 0^{n-1}\rangle \mapsto |b^n\rangle$ for $b \in \{0,1\}$ can be implemented by a size-$(n-1)$, depth-$\lceil \log n \rceil$ circuit consisting of CNOT gates with no ancillae. Therefore the fanout transformation $|b, x\rangle \mapsto |b, x \oplus b^n\rangle$ for $b \in \{0,1\}, x \in \{0,1\}^n$ can be implemented by first computing b^n as described above, then XORing b^n onto x, and finally uncomputing b^n.

Similarly an n-qubit generalized Toffoli gate can be cleanly simulated by a size-$O(n)$, depth-$O(\log n)$ QNC circuit with $O(n)$ ancillae. This follows by simulating a log-depth DeMorgan formula for the AND function (i.e. the circuit whose graph is a balanced binary tree of 2-bit AND gates), with one ancilla qubit allocated to store the value of each gate in the DeMorgan formula, and then uncomputing the garbage.

A general n-qubit, depth-1 QAC_f circuit can be written as $\bigotimes_j G_j$, where each G_j is a k_j-qubit gate such that $\sum_j k_j \leq n$, and if $k_j > 1$ then G_j is either a generalized Toffoli or fanout gate. It follows that $\bigotimes_j G_j$ can be cleanly simulated by a QNC circuit where the size and number of ancillae are $O\left(\sum_j k_j\right) \leq O(n)$ and the depth is $O(\max_j \log k_j) \leq O\left(\log \sum_j k_j\right) \leq O(\log n)$. The lemma follows by successively implementing each layer of a QAC_f circuit in this way, reusing the same ancillae to simulate each layer. □

References

[1] Scott Aaronson. "The complexity of quantum states and transformations: from quantum money to black holes". arXiv:1607.05256 (2016). 1, 2, 8

[2] Scott Aaronson and Greg Kuperberg. "Quantum versus classical proofs and advice". Theory Comput. **3**, 129–157 (2007). arXiv:quant-ph/0604056. 1

[3] Scott Aaronson. "Open problems related to quantum query complexity". ACM Trans. Quantum Comput. **2**, 1–9 (2021). arXiv:2109.06917. 1

[4] Alex Lombardi, Fermi Ma, and John Wright. "A one-query lower bound for unitary synthesis and breaking quantum cryptography". In STOC. Pages 979–990. (2024). arXiv:2310.08870. 1, 2

[5] Gregory Rosenthal. "Efficient quantum state synthesis with one query". In SODA. Pages 2508–2534. (2024). arXiv:2306.01723. 2, 7, 8

[6] Michael A. Nielsen and Isaac L. Chuang. "Quantum computation and quantum information: 10th anniversary edition". Cambridge University Press. (2010). 2, 4, 9, 12

[7] Christopher M. Dawson and Michael A. Nielsen. "The Solovay–Kitaev algorithm". Quantum Inf. Comput. **6**, 81–95 (2006). arXiv:quant-ph/0505030. 2

[8] Sandy Irani, Anand Natarajan, Chinmay Nirkhe, Sujit Rao, and Henry Yuen. "Quantum search-to-decision reductions and the state synthesis problem". In CCC. Volume 234, pages 5:1–5:19. (2022). arXiv:2111.02999. 2

[9] Xiaoming Sun, Guojing Tian, Shuai Yang, Pei Yuan, and Shengyu Zhang. "Asymptotically optimal circuit depth for quantum state preparation and general unitary synthesis". IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst. **42**, 3301–3314 (2023). arXiv:2108.06150. 3

[10] Pei Yuan and Shengyu Zhang. "Optimal (controlled) quantum state preparation and improved unitary synthesis by quantum circuits with any number of ancillary qubits". Quantum **7**, 956 (2023). arXiv:2202.11302. 3

[11] Xiao-Ming Zhang, Tongyang Li, and Xiao Yuan. "Quantum state preparation with optimal circuit depth: Implementations and applications". Physical Review Letters **129**, 230504 (2022). arXiv:2201.11495. 3

[12] Frederic Green, Steven Homer, Cristopher Moore, and Christopher Pollett. "Counting, fanout, and the complexity of quantum ACC". Quantum Inf. Comput. **2**, 35–65 (2002). arXiv:quant-ph/0106017. 3, 12

[13] Peter Høyer and Robert Špalek. "Quantum fan-out is powerful". Theory Comput. **1**, 81–103 (2005). 3

[14] Yasuhiro Takahashi and Seiichiro Tani. "Collapse of the hierarchy of constant-depth exact quantum circuits". Comput. Complexity **25**, 849–881 (2016). arXiv:1112.6063. 3

[15] Johan Håstad. "Almost optimal lower bounds for small depth circuits". In STOC. Pages 6–20. (1986). 3

[16] Stasys Jukna. "Boolean function complexity". Volume 27 of Algorithms and Combinatorics. Springer, Heidelberg. (2012). 3

[17] Oleg Lupanov. "On a method of circuit synthesis". Izvestia VUZ **1**, 120–140 (1958).

[18] Claude Shannon. "The synthesis of two-terminal switching circuits". Bell System Tech. J. **28**, 59–98 (1949). 3

[19] Andris Ambainis. "Quantum lower bounds by quantum arguments". J. Comput. System Sci. **64**, 750–767 (2002). arXiv:quant-ph/0002066. 4, 6
[20] Ashwin Nayak. "Inverting a permutation is as hard as unordered search". Theory Comput. **7**, 19–25 (2011). arXiv:1007.2899. 4, 6
[21] Sándor Imre and Ferenc Balázs. "Quantum computing and communications: an engineering approach". Chapter 7. John Wiley & Sons. (2005). 5
[22] Nathan Wiebe (2021). Personal communication. 5