Power of Uninitialized Qubits in Shallow Quantum Circuits

Yasuhiro Takahashi

NTT Communication Science Laboratories, NTT Corporation, Japan
takahashi.yasuhiro@lab.ntt.co.jp

Seiichiro Tani

NTT Communication Science Laboratories, NTT Corporation, Japan
tani.seiichiro@lab.ntt.co.jp

──── **Abstract** ────

We study the computational power of shallow quantum circuits with $O(\log n)$ initialized and $n^{O(1)}$ uninitialized ancillary qubits, where n is the input length and the initial state of the uninitialized ancillary qubits is arbitrary. First, we show that such a circuit can compute any symmetric function on n bits that is classically computable in polynomial time. Then, we regard such a circuit as an oracle and show that a polynomial-time classical algorithm with the oracle can estimate the elements of any unitary matrix corresponding to a constant-depth quantum circuit on n qubits. Since it seems unlikely that these tasks can be done with only $O(\log n)$ initialized ancillary qubits, our results give evidences that adding uninitialized ancillary qubits increases the computational power of shallow quantum circuits with only $O(\log n)$ initialized ancillary qubits. Lastly, to understand the limitations of uninitialized ancillary qubits, we focus on near-logarithmic-depth quantum circuits with them and show the impossibility of computing the parity function on n bits.

1 Introduction

1.1 Background and Main Results

Much attention has been paid to the computational power of shallow (i.e., polylogarithmic-depth) quantum circuits [6, 16, 10, 9, 11, 8, 3, 22, 20, 4]. A major purpose of this line of research is to understand the differences between shallow quantum and classical circuits. In addition, it is strongly motivated by one of the most difficult problems concerning quantum circuit implementation: in current and near-future technologies, it would be very difficult to keep quantum coherence for a period of time long enough to apply many gates.

In discussing the computational power of shallow quantum circuits, polynomially many ancillary qubits initialized to, say, $|0\rangle$ are assumed to be available. The initialized ancillary qubits are particularly important for quantum circuits since many quantum operations require ancillary qubits to preserve their unitary property and store intermediate results. Another implementation problem arises here: it is difficult to prepare a large number of qubits that are simultaneously initialized to a certain state. Indeed, this problem has often been addressed

SYMPOSIUM
ON THEORETICAL
ASPECTS
OF COMPUTER
SCIENCE

in the literature [7, 14]. However, most papers concerning the problem assume a sufficiently long coherence time. In this paper, we address these two problems simultaneously.

A straightforward quantum computation model reflecting a short coherence time and a limited number of initialized ancillary qubits would be shallow quantum circuits with $O(\log n)$ initialized ancillary qubits, where n is the input length. However, their computational power seems quite low since each step of them can utilize only a small number of intermediate results. In fact, it is not even known whether such a circuit can compute the OR function on n bits, and it seems unlikely that it can. Therefore, it is highly desirable to find additional ancillary qubits satisfying the following conditions: they should be easier to prepare than initialized ancillary qubits and increase the computational power of shallow quantum circuits with only $O(\log n)$ initialized ancillary qubits. An interesting direction is to study qubits in the completely mixed state [13], but it would be better not to assume any particular initial state.

We consider polynomially many uninitialized qubits as additional ancillary qubits. More concretely, we study shallow quantum circuits with $O(\log n)$ initialized and $n^{O(1)}$ uninitialized ancillary qubits, where we assume that no intermediate measurements are allowed. The initial state of the uninitialized ancillary qubits is arbitrary and thus they are easier to prepare than initialized ancillary qubits, i.e., they satisfy the above first condition on additional ancillary qubits. But do they satisfy the second condition? Specifically, are shallow quantum circuits with $O(\log n)$ initialized and $n^{O(1)}$ uninitialized ancillary qubits more powerful than those without uninitialized ancillary qubits? Although uninitialized ancillary qubits are known to be useful for constructing a few efficient quantum circuits [1, 19], a complexity-theoretic analysis of quantum circuits with such ancillary qubits has not yet been done.

First, to give evidence of an affirmative answer to the question, we consider symmetric functions, which are Boolean functions whose output depends only on the number of ones in the input bits [12]. Let $\mathcal{S}_n$ be the class of symmetric functions on n bits that are classically computable in polynomial time. For example, $\mathcal{S}_n$ includes the OR function, for which it is not known whether there exists a shallow quantum circuit (consisting of one-qubit gates and CNOT gates) with only $O(\log n)$ initialized ancillary qubits, and it seems unlikely that it does. However, any function in $\mathcal{S}_n$ can be computed by adding uninitialized ancillary qubits:

▶ **Theorem 1.** *Any $f_n \in \mathcal{S}_n$ can be computed by an $O((\log n)^2)$-depth quantum circuit with n input qubits, one output qubit, and $O(\log n)$ initialized and $O(n(\log n)^2)$ uninitialized ancillary qubits such that it consists of the gates in the gate set $\mathcal{G}$, where $\mathcal{G}$ consists of a Hadamard gate, a phase-shift gate with angle $2\pi c/2^t$ for any integers $t \geq 1$ and c, and a CNOT gate.*

Theorem 1 gives evidence that shallow quantum circuits with $O(\log n)$ initialized and $n^{O(1)}$ uninitialized ancillary qubits are more powerful than those without uninitialized ancillary qubits in terms of computing symmetric functions. The proof of Theorem 1 immediately implies that the depth of the circuit can be decreased to $O(\log n)$ when the circuit is allowed to further include unbounded fan-out gates and unbounded Toffoli gates.

Then, to give further evidence of the computational advantage of using uninitialized ancillary qubits, we consider a classical algorithm with an oracle that can perform a shallow quantum circuit with them. When the oracle receives a bit string w, it performs the circuit with input qubits initialized to $|w\rangle$ and sends back the classical outcome of the measurement on the output qubit. Let $p(n)$ be a polynomial and C_n be a constant-depth quantum circuit on n qubits consisting of the gates in $\mathcal{G}$. The problem, denoted by $\mathrm{MAT}(p(n), C_n)$, is to compute a real number α_x such that $|\alpha_x - |\langle 0^n|C_n|x\rangle|^2| \leq 1/p(n)$ for any input $x \in \{0,1\}^n$, where C_n also denotes its matrix representation. It is not known whether the problem has a polynomial-time classical algorithm, and it seems unlikely that it does [17], even when we use

an oracle that can perform a shallow quantum circuit with only $O(\log n)$ initialized ancillary qubits. However, the problem can be solved by adding uninitialized ancillary qubits:

▶ **Theorem 2.** *For any polynomial $p(n)$ and a constant-depth quantum circuit C_n on n qubits consisting of the gates in $\mathcal{G}$, $\mathrm{MAT}(p(n), C_n)$ can be solved with probability exponentially (in n) close to 1 by a polynomial-time probabilistic classical algorithm with an oracle that can perform an $O(\log n)$-depth quantum circuit with $2n$ input qubits, one output qubit, and (no initialized and) n uninitialized ancillary qubits such that it consists of the gates in $\mathcal{G}$.*

As with Theorem 1, Theorem 2 gives evidence that shallow quantum circuits with $O(\log n)$ initialized and $n^{O(1)}$ uninitialized ancillary qubits are more powerful than those without uninitialized ancillary qubits. More concretely, by the proof of Theorem 2, this is evidence that there exists a probability distribution on $\{0, 1\}$ that can be generated with uninitialized ancillary qubits but cannot without them. This is because, otherwise, $\mathrm{MAT}(p(n), C_n)$ would be solved by using an oracle with only $O(\log n)$ initialized ancillary qubits. We give a brief comment on the number of input qubits in the circuit performed by the oracle. If the number is large, a classical algorithm can send 0^k for large k (besides another bit string) to the oracle and the circuit can use a part of the input qubits as a large number of initialized ancillary qubits. To avoid this, we restrict the number of input qubits to $2n$.

Lastly, to understand the limitations of uninitialized ancillary qubits, for an arbitrary constant $0 \le \delta < 1$, we focus on $O((\log n)^\delta)$-depth quantum circuits with them and consider the computability of the parity function on n bits. Since the depth is $o(\log n)$, it is easy to show that the parity function cannot be computed by any such circuit consisting of the gates in $\mathcal{G}$. This is also the case even when the circuit includes additional gates on a non-constant number of qubits:

▶ **Theorem 3.** *Let $0 \le \delta < 1$ be an arbitrary constant. Then, the parity function on n bits cannot be computed by any $O((\log n)^\delta)$-depth quantum circuit with n input qubits, one output qubit, and $O(\log n)$ initialized and $n^{O(1)}$ uninitialized ancillary qubits such that it consists of the gates in $\mathcal{G}$, unbounded fan-out gates on $(\log n)^{O(1)}$ qubits, and unbounded Toffoli gates.*

Theorem 3 means that $O((\log n)^\delta)$-depth quantum circuits with $O(\log n)$ initialized and $n^{O(1)}$ uninitialized ancillary qubits are *not* more powerful than those without uninitialized ancillary qubits in terms of computing the parity function, even when they include the two types of gates on a non-constant number of qubits. Moreover, Theorem 3 implies that the circuit in Theorem 1 is optimal in the following sense. As described in the paragraph following Theorem 1, the depth of the circuit becomes $O(\log n)$ when the circuit uses the gates in $\mathcal{G}$, unbounded fan-out gates, and unbounded Toffoli gates. As described in Section 1.3, the circuit is based on the computation of the number of ones in the input bits and thus can be regarded as a parity circuit. Thus, the circuit cannot be significantly improved simultaneously in terms of both the depth and the number of qubits on which unbounded fan-out gates act. This is because, otherwise, we would obtain a parity circuit that contradicts Theorem 3.

We comment on the states of uninitialized ancillary qubits in the above theorems. In the proofs of Theorems 1 and 2, we assume that the state of uninitialized ancillary qubits is an arbitrary computational basis (pure) state. These proofs can be simply extended for an arbitrary pure/mixed state by the linearity of quantum operations and the fact that a mixed state is a probabilistic mixture of pure states. Thus, Theorems 1 and 2 hold for an arbitrary pure/mixed state. We show Theorem 3 under the same assumption. Thus, Theorem 3 means that there does not exist an $O((\log n)^\delta)$-depth quantum circuit (with the property described in the theorem) that computes the parity function on n bits regardless of the initial state

of uninitialized ancillary qubits, where we assume that the initial state is restricted to an arbitrary computational basis (pure) state. In this statement, we can remove the restriction, i.e., we can assume that the initial state is an arbitrary pure/mixed state. This is because the resulting statement is weaker than the original one. In this sense, Theorem 3 holds for an arbitrary pure/mixed state.

1.2   Imposing the Quantum Catalytic Requirement

Buhrman et al. [5] defined a *classical* computation with a logarithmic-size clean space and a polynomial-size additional space, which they call a catalytic log-space computation. The initial state of the additional space is arbitrary, and they impose the catalytic requirement that its state has to be returned to the initial one at the end of the computation. They showed a surprising result: it appears that such a computation is more powerful than that without the additional space. The additional space seems like a catalyst in a chemical reaction.

The corresponding catalytic requirement in our quantum setting is that the state of uninitialized ancillary qubits has to be returned to the initial one at the end of computation. Since the circuit in Theorem 1 has no error, by the standard technique of uncomputation, it is easy to transform the circuit into the one that meets the quantum catalytic requirement without increasing the original asymptotic complexity. Thus, Theorem 1 means that uninitialized ancillary qubits seem like a catalyst as in the classical setting [5]. When shallow quantum circuits have an error, it is not easy to transform them into the ones that meet the quantum catalytic requirement and the analysis of such circuits is left for future work.

From a practical point of view, it is even better to decrease the number of uninitialized ancillary qubits we need to specially prepare in addition to decreasing the number of initialized ones. The quantum catalytic requirement allows us to do this in some cases. An example is when we use a shallow quantum circuit with uninitialized ancillary qubits in a quantum circuit for Shor's factoring algorithm [19]. The factoring circuit uses two registers and, during some operation, all qubits in one register are idle. Thus, when we use a shallow quantum circuit for the operation that meets the above requirement, we can regard the idle qubits as uninitialized ancillary qubits since the circuit returns their state to the initial one. The use of the circuit in this way requires that the computation has to be done with only qubits, which matches our quantum computation model. From a complexity-theoretic standpoint, it is also interesting to study a quantum computation model with an additional classical space [23].

1.3   Overview of Techniques

We construct two quantum circuits to obtain the circuit for $f_n \in \mathcal{S}_n$ in Theorem 1. The first one is an $O((\log n)^2)$-depth OR reduction circuit with $O(n(\log n)^2)$ uninitialized ancillary qubits, which reduces the computation of the OR function on n bits to that on $m = O(\log n)$ bits. Its first part is a modification of the original OR reduction circuit [11] and yields a state whose phase depends on the uninitialized ancillary qubits but has a convenient form to eliminate the dependency. We apply similar circuits repeatedly to add an appropriate phase to that of the state, which eliminates any dependency on the uninitialized ancillary qubits. The second circuit is an $O(m^2)$-depth one for g_m with $O(m2^m)$ uninitialized ancillary qubits. Here, g_m is a Boolean function on m bits satisfying that $g_m(s) = f_n(x)$ for any $x \in \{0,1\}^n$, where $s \in \{0,1\}^m$ is the binary representation of the number of ones in x. The circuit is based on the Fourier expansion of g_m [12] and the above method for eliminating any dependency on the uninitialized ancillary qubits. For any input $x \in \{0,1\}^n$, we first compute s using the OR reduction circuit and then compute $g_m(s) = f_n(x)$ using the circuit for g_m.

The algorithm in Theorem 2 is based on a polynomial-time probabilistic classical algorithm for $\mathrm{MAT}(p(n), C_n)$ with an oracle [17], where the oracle can perform a commuting quantum circuit for the Hadamard test [15]. Although initialized ancillary qubits can be used to parallelize the Hadamard test [22], it has not been known whether uninitialized ancillary qubits are useful for this purpose. We show that they can be used like initialized ancillary qubits in parallelizing the Hadamard test. We replace the commuting quantum circuit with a new circuit with our parallelizing techniques using uninitialized ancillary qubits in the algorithm for $\mathrm{MAT}(p(n), C_n)$, which yields the desired algorithm.

We show Theorem 3 by extending the proof of Bera [3]. Our proof is different from the previous one in that it deals with ancillary qubits and unbounded fan-out gates. The key to Theorem 3 is to show that, for any quantum circuit C_n with $O(\log n)$ initialized and $n^{O(1)}$ uninitialized ancillary qubits such that it may include unbounded Toffoli gates, there exists an initial state of the uninitialized ancillary qubits such that C_n with the initial state is well approximated by $\tilde{C}_n$ with the same initial state. Here, $\tilde{C}_n$ is the circuit obtained from C_n by removing unbounded Toffoli gates on a large number of qubits. Thus, if C_n is a small-depth quantum circuit for the parity function, then $\tilde{C}_n$ computes the same function with high probability. This is impossible since $\tilde{C}_n$ does not have any gate on a large number of qubits and thus its output does not depend on all input qubits.

2   Preliminaries

A quantum circuit consists of elementary gates, each of which is in the gate set $\mathcal{G}$, where $\mathcal{G}$ consists of a Hadamard gate H, a phase-shift gate $Z(\theta)$ with angle θ, and a CNOT gate. Here, $H = |+\rangle\langle 0| + |-\rangle\langle 1|$ and $Z(\theta) = |0\rangle\langle 0| + e^{i\theta}|1\rangle\langle 1|$, where $|\pm\rangle = (|0\rangle \pm |1\rangle)/\sqrt{2}$ and $\theta = 2\pi c/2^t$ for any integers $t \geq 1$ and c. We write $Z(\pi)$ and $HZ(\pi)H$ as Z and X, respectively. In some cases, we use a fan-out gate and a Toffoli gate as elementary gates. Let $k \geq 1$ be an integer. A fan-out gate on $k + 1$ qubits implements the operation defined as $|y\rangle \bigotimes_{j=1}^{k} |x_j\rangle \mapsto |y\rangle \bigotimes_{j=1}^{k} |x_j \oplus y\rangle$ for any $y, x_j \in \{0, 1\}$, where $\oplus$ denotes addition modulo 2. The first input qubit is called the control qubit. A k-controlled Toffoli gate implements the operation on $k + 1$ qubits defined as $\left(\bigotimes_{j=1}^{k} |x_j\rangle\right) |y\rangle \mapsto \left(\bigotimes_{j=1}^{k} |x_j\rangle\right) |y \oplus \bigwedge_{j=1}^{k} x_j\rangle$, where $\bigwedge$ denotes the logical AND. The first k input qubits are called the control qubits and the last input qubit is called the target qubit. These gates with $k = 1$ are CNOT gates. When it is permitted to apply a fan-out gate and a Toffoli gate on a non-constant number of qubits, they are called an unbounded fan-out gate and an unbounded Toffoli gate, respectively.

To simplify the descriptions of quantum circuits, we use a k-controlled $Z(\theta)$ gate for any θ described above, which will be decomposed into elementary gates. The gate implements the operation on $k + 1$ qubits defined as $\bigotimes_{j=1}^{k+1} |x_j\rangle \mapsto e^{i\theta \bigwedge_{j=1}^{k+1} x_j} \bigotimes_{j=1}^{k+1} |x_j\rangle$ for any $x_j \in \{0, 1\}$. We can choose an arbitrary qubit as the target qubit and the other qubits are called the control qubits. The inverse of the gate is the k-controlled $Z(-\theta)$ gate. When it is permitted to apply the gate on a non-constant number of qubits, it is called an unbounded $Z(\theta)$ gate.

The complexity measures of a quantum circuit are its size and depth. The size of a quantum circuit is the total size of all elementary gates in the circuit, where the size of an elementary gate is the number of qubits on which the gate acts. To define the depth, we regard the circuit as a set of layers $1, \ldots, d$ consisting of elementary gates, where gates in the same layer act on pairwise disjoint sets of qubits and any gate in layer j is applied before any gate in layer $j + 1$. The depth of the circuit is the smallest possible value of d [9].

We deal with a uniform family of polynomial-size quantum circuits $\{C_n\}_{n \geq 1}$, where no intermediate measurements are allowed. The uniformity means that the function $1^n \mapsto \overline{C_n}$ is classically computable in polynomial time, where $\overline{C_n}$ is the classical description of C_n. Each

C_n has n input qubits and can have one output qubit and $n^{O(1)}$ ancillary qubits that are divided into two groups: $p = O(\log n)$ qubits and the remaining q qubits. We assume that, for any $x \in \{0,1\}^n$ and $y \in \{0,1\}$, we can initialize the input qubits and output qubit to $|x\rangle$ and $|y\rangle$, respectively. We can also initialize the p ancillary qubits to $|0\rangle$, which we call initialized ancillary qubits, but we cannot initialize the q ancillary qubits and do not know their initial state. They are called uninitialized ancillary qubits. When C_n has the output qubit, a measurement in the Z basis is performed on it at the end of the computation. The classical outcome of the measurement, which is 0 or 1, is called the output of C_n. A symbol denoting a quantum circuit also denotes its matrix representation in the computational basis.

A Boolean function f_n on n bits is a mapping $f_n : \{0,1\}^n \rightarrow \{0,1\}$. We define its computability by a quantum circuit with uninitialized ancillary qubits as follows:

▶ **Definition 4.** Let f_n be a Boolean function on n bits and C_n be a quantum circuit with n input qubits, one output qubit, and p initialized and q uninitialized ancillary qubits. The circuit C_n computes f_n if, for any $x \in \{0,1\}^n$ and $y \in \{0,1\}$, when the input qubits and output qubit are initialized to $|x\rangle$ and $|y\rangle$, respectively, the output of C_n is $y \oplus f_n(x)$ with probability 1, regardless of the initial state of the q uninitialized ancillary qubits.

A Boolean function is called symmetric if its output depends only on the number of ones in the input bits [12]. Let $\mathcal{S}_n$ be the class of symmetric functions on n bits that are classically computable in polynomial time. For example, $\mathcal{S}_n$ includes the parity function PA_n and the OR function OR_n. Here, for any $x = x_1 \cdots x_n \in \{0,1\}^n$, $\mathrm{PA}_n(x) = 1$ if $|x|$ is odd and 0 otherwise, where $|x| = \sum_{j=1}^n x_j$. Moreover, $\mathrm{OR}_n(x) = 1$ if $|x| \geq 1$ and 0 otherwise.

We define the function associated with $f_n \in \mathcal{S}_n$ as follows:

▶ **Definition 5.** Let $f_n \in \mathcal{S}_n$. The function associated with f_n is the Boolean function g_m on $m = \lceil \log(n+1) \rceil$ bits defined as follows: For any $s = s_1 \cdots s_m \in \{0,1\}^m$, $g_m(s) = f_n(1^l 0^{n-l})$ if $l \leq n$ and 0 otherwise, where $l = \sum_{k=1}^m s_k 2^{k-1}$.

The function g_m is classically computable in time $n^{O(1)}$ and, for any $x \in \{0,1\}^n$, if $s = s_1 \cdots s_m$ is the binary representation of $|x|$, i.e., $|x| = \sum_{k=1}^m s_k 2^{k-1}$, then $g_m(s) = f_n(x)$.

We explain the idea of the original OR reduction quantum circuit [11]. The circuit has n input qubits and $O(n \log n)$ initialized ancillary qubits. Let $|x\rangle$ be an input state for any $x \in \{0,1\}^n$. The circuit transforms the state of m initialized ancillary qubits into the state $\bigotimes_{k=1}^m |\varphi_k\rangle$, where $m = \lceil \log(n+1) \rceil$ and $|\varphi_k\rangle = (|+\rangle + e^{\frac{2\pi i}{2^k}|x|}|-\rangle)/\sqrt{2}$ for any $1 \leq k \leq m$. If $|x| = 0$, then $|\varphi_k\rangle = |0\rangle$ for any $1 \leq k \leq m$ and thus the output state is $|0^m\rangle$. If $|x| \geq 1$, then $|\varphi_k\rangle = |1\rangle$ for some $1 \leq k \leq m$ and thus the output state is orthogonal to $|0^m\rangle$. This means that the circuit reduces the computation of OR_n to that of OR_m. The output state can be used to compute the binary representation $s_1 \cdots s_m$ of $|x|$. In fact, it is easy to show that the state $\bigotimes_{k=1}^m |s_k\rangle$ can be obtained by applying $\mathrm{QFT}_{2^m}^\dagger$ to the state $\bigotimes_{k=1}^m H|\varphi_k\rangle$, where $\mathrm{QFT}_{2^m}^\dagger$ is the inverse of the quantum Fourier transform modulo 2^m.

3 Shallow Quantum Circuits for Symmetric Functions

Let $f_n \in \mathcal{S}_n$. We compute f_n on input $x \in \{0,1\}^n$ using the following algorithm:
1. Compute the binary representation $s \in \{0,1\}^m$ of $|x|$, where $m = \lceil \log(n+1) \rceil$.
2. Compute $g_m(s) = f_n(x)$, where g_m is the function associated with f_n.
To implement Step 1, we construct an OR reduction circuit Q_n with uninitialized ancillary qubits. As described above, we can obtain s using Q_n (with a layer of H gates) and the standard $O(m)$-depth quantum circuit for $\mathrm{QFT}_{2^m}^\dagger$ with no ancillary qubits [18]. To implement Step 2, we construct a quantum circuit R_m for g_m with uninitialized ancillary qubits.
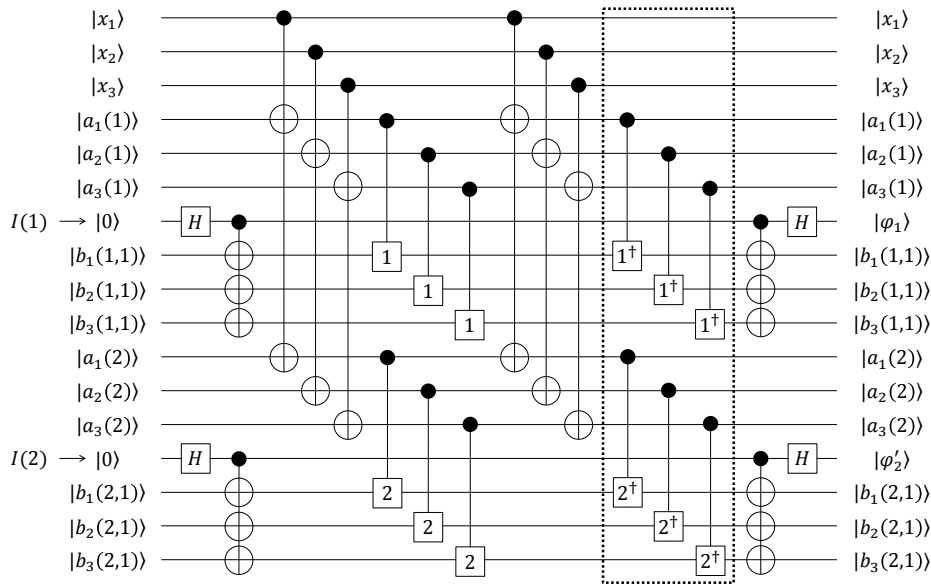
Figure 1 The first stage of our OR reduction circuit with input $x = x_1 x_2 x_3 \in \{0,1\}^3$. The gate next to the H gate is a fan-out gate on four qubits, where the top qubit is the control qubit. For any integer $t \geq 1$, the gates t and $t^\dagger$ represent a $Z(2\pi/2^t)$ gate and its inverse, i.e., a $Z(-2\pi/2^t)$ gate, respectively. The dashed box represents the gates added to the original OR reduction circuit.

3.1 OR Reduction Circuit with Uninitialized Ancillary Qubits

The circuit Q_n is an $O((\log n)^2)$-depth OR reduction circuit with n input qubits and $O(\log n)$ initialized and $O(n(\log n)^2)$ uninitialized ancillary qubits. To explain our idea for constructing Q_n, we consider the case where $n = 3$ (and thus $m = 2$). The first stage of Q_n is depicted in Fig. 1, where the initial state of the uninitialized ancillary qubits is represented by the (unknown) values $a_j(k), b_j(k,l) \in \{0,1\}$. This circuit is obtained by adding the gates in the dashed box to the original OR reduction circuit. We want to transform the initial states of the initialized ancillary qubits $I(1)$ and $I(2)$ into the states $|\varphi_1\rangle$ and $|\varphi_2\rangle$, respectively. If we do not apply the added gates, the output state of $I(k)$ is $(|+\rangle + e^{\frac{2\pi i}{2^k}\alpha(k,1)}|-\rangle)/\sqrt{2}$, where $\alpha(k,1) = \sum_{j=1}^{3}(-1)^{b_j(k,1)}(x_j \oplus a_j(k))$ and $k = 1,2$. The phase of this state depends on the initial state of the uninitialized ancillary qubits and we eliminate the dependency.

The point is that the added gates allow us to obtain an output state of $I(k)$ whose phase has a convenient form to eliminate the dependency. More concretely, by applying them, the output state of $I(k)$ is $(|+\rangle + e^{\frac{2\pi i}{2^k}\gamma(k,1)}|-\rangle)/\sqrt{2}$, where $\gamma(k,1) = |x| - 2\sum_{j=1}^{3}x_j(a_j(k) \oplus b_j(k,1))$. Since $e^{\frac{2\pi i}{2}\gamma(1,1)} = e^{\frac{2\pi i}{2}|x|}$, the output state of $I(1)$ is equal to $|\varphi_1\rangle$ as desired. The dependency is eliminated since the terms in $\gamma(1,1)$ other than $|x|$ yield only an angle of a multiple of 2π.

Unfortunately, the output state of $I(2)$, which is represented as $|\varphi_2'\rangle$ in Fig. 1, is not equal to $|\varphi_2\rangle$ in general since the phase $e^{\frac{2\pi i}{2^2}\gamma(2,1)}$ depends on the initial states of the uninitialized ancillary qubits, where $\gamma(2,1) = |x| - 2\sum_{j=1}^{3}x_j(a_j(2) \oplus b_j(2,1))$. To eliminate the dependency, we consider the second stage where we add an angle $\frac{2\pi}{2^2}\delta(2,2)$ to the original angle $\frac{2\pi}{2^2}\gamma(2,1)$ using three new uninitialized ancillary qubits (not depicted in Fig. 1). Here, their initial state is $|b_1(2,2)\rangle|b_2(2,2)\rangle|b_3(2,2)\rangle$ for any (unknown) $b_j(2,2) \in \{0,1\}$ and $\delta(2,2) = |x| - \gamma(2,1) - 2^2\sum_{j=1}^{3}x_j(a_j(2) \oplus b_j(2,1))(a_j(2) \oplus b_j(2,2))$. The value $\delta(2,2)$ has a form similar to $\gamma(2,1)$ and thus we can implement the second stage using a quantum circuit similar to the one in Fig. 1. Since $e^{\frac{2\pi i}{2^2}(\gamma(2,1)+\delta(2,2))} = e^{\frac{2\pi i}{2^2}|x|}$, we obtain $|\varphi_2\rangle$ as desired.

We generalize the idea. Let $x = x_1 \cdots x_n \in \{0,1\}^n$ be an input. We prepare n input qubits $X_1, \ldots, X_n$ and m initialized ancillary qubits $I(1), \ldots, I(m)$, where X_j is initialized to $|x_j\rangle$. We also prepare $nm(m+3)/2$ uninitialized ancillary qubits, which are divided into two groups, A and B. Group A consists of mn qubits, which are divided into m groups $A(1), \ldots, A(m)$. Each $A(k)$ consists of n qubits $A_1(k), \ldots, A_n(k)$, where the initial state of $A_j(k)$ is $|a_j(k)\rangle$ for any (unknown) $a_j(k) \in \{0,1\}$. Group B consists of $nm(m+1)/2$ qubits, which are divided into m groups $B(1), \ldots, B(m)$. Each $B(k)$ consists of kn qubits, which are divided into k groups $B(k,1), \ldots, B(k,k)$. Each $B(k,l)$ consists of n qubits $B_1(k,l), \ldots, B_n(k,l)$, where the initial state of $B_j(k,l)$ is $|b_j(k,l)\rangle$ for any (unknown) $b_j(k,l) \in \{0,1\}$. The circuit Q_n consists of m stages. For any $1 \le s \le m$, Stage s is defined as follows:

1. Apply a H gate to $I(k)$ for every $s \le k \le m$ in parallel.
2. Apply a fan-out gate on $n+1$ qubits to $B_1(k,s), \ldots, B_n(k,s)$, and $I(k)$ for every $s \le k \le m$ in parallel, where $I(k)$ is the control qubit.
3. If $s \ge 2$, then apply a fan-out gate on s qubits to $B_j(k,1), \ldots, B_j(k,s-1)$, and $A_j(k)$ for every $s \le k \le m$ and $1 \le j \le n$ in parallel, where $A_j(k)$ is the control qubit.
4. Apply a fan-out gate on $m-s+2$ qubits to $A_j(s), A_j(s+1), \ldots, A_j(m)$, and X_j for every $1 \le j \le n$ in parallel, where X_j is the control qubit.
5. Apply an s-controlled $Z(2\pi/2^{k-s+1})$ gate to $B_j(k,s)$ and the following qubits for every $s \le k \le m$ and $1 \le j \le n$ in parallel: $A_j(k)$ if $s = 1$ and $B_j(k,1), \ldots, B_j(k,s-1)$, and $A_j(k)$ otherwise.
6. Apply the gates in Step 4.
7. Apply the inverse of the gates in Step 5.
8. Apply the gates in Step 3, Step 2, and Step 1 (in this order).

The circuit Q_n outputs the desired state and has the desired complexity as follows. The proofs can be found in the full version of the paper [21].

▶ **Lemma 6.** *Let $x = x_1 \cdots x_n \in \{0,1\}^n$ be an input. For any $1 \le k \le m$ and $1 \le s \le k$, the state of $I(k)$ after Stage s is the state $(|+\rangle + e^{\frac{2\pi i}{2^k}\gamma(k,s)}|-\rangle)/\sqrt{2}$, where $\gamma(k,s) = |x| - 2^s \sum_{j=1}^n x_j \bigwedge_{l=1}^s (a_j(k) \oplus b_j(k,l))$. Moreover, the state of any qubit other than the initialized ancillary qubits is the same as its initial one. The state of $I(k)$ after Stage k is the state $|\varphi_k\rangle$.*

▶ **Lemma 7.** *The circuit Q_n uses $O(\log n)$ initialized and $O(n(\log n)^2)$ uninitialized ancillary qubits, and its depth is $O((\log n)^2)$, when the elementary gate set is $\mathcal{G}$.*

3.2 Circuit for the Function Associated with a Symmetric Function

We construct an $O(m^2)$-depth quantum circuit R_m for g_m with $m = \lceil \log(n+1) \rceil$ input qubits, one output qubit, and $O(m2^m)$ uninitialized ancillary qubits, where g_m is the function associated with $f_n \in \mathcal{S}_n$. The circuit uses (a slight modification of) the Fourier expansion of g_m [12]: For any $s = s_1 \cdots s_m \in \{0,1\}^m$, $g_m(s) = g_m(0^m) + \frac{2}{2^m} \sum_t c_t \bigoplus_{k=1}^m t_k s_k$, where $c_t = \sum_u g_m(u)(2\bigoplus_{k=1}^m u_k t_k - 1)$, $t = t_1 \cdots t_m$ ranges over $\{0,1\}^m \setminus \{0^m\}$, and $u = u_1 \cdots u_m$ ranges over $\{0,1\}^m$. Since $m = O(\log n)$ and g_m is classically computable in time $n^{O(1)}$, the number of c_t's with $t \in \{0,1\}^m \setminus \{0^m\}$ is $n^{O(1)}$ and the function $t \mapsto c_t$ is also classically computable in time $n^{O(1)}$. This implies the uniformity of our circuit family for f_n.

The circuit R_m with input $s = s_1 \cdots s_m \in \{0,1\}^m$ is based on the following algorithm:
1. Compute the parity value $\bigoplus_{k=1}^m t_k s_k$ for every $t \in \{0,1\}^m \setminus \{0^m\}$ in parallel.
2. Prepare $(|+\rangle + e^{\pi i g_m(s)}|-\rangle)/\sqrt{2} = |g_m(s)\rangle$ using the above representation of g_m.

Since we do not have any initialized ancillary qubit, in Step 1, we can only have the parity values on uninitialized ancillary qubits, i.e., $a_t \oplus \bigoplus_{k=1}^{m} t_k s_k$ for every $t \in \{0,1\}^m \setminus \{0^m\}$, where the initial state of the uninitialized ancillary qubits is represented by the (unknown) values $a_t \in \{0,1\}$. Thus, in Step 2, we have to use such values to prepare $(|+\rangle + e^{\pi i g_m(s)} |-\rangle)/\sqrt{2} = X^{g_m(0^m)}(|+\rangle + e^{\frac{2\pi i}{2^m} \sum_t c_t \bigoplus_{k=1}^{m} t_k s_k} |-\rangle)/\sqrt{2}$, which does not depend on a_t. The point is that this situation is essentially the same as the one where $|\varphi_m\rangle$ is prepared by Q_n as described in Section 3.1, i.e., where we can only have the values $a_j(m) \oplus x_j$ for every $1 \le j \le n$ and we have to use them to prepare $|\varphi_m\rangle = (|+\rangle + e^{\frac{2\pi i}{2^m} |x|} |-\rangle)/\sqrt{2}$, which does not depend on $a_j(m)$. Thus, roughly speaking, we can construct R_m in a similar way to a part of Q_n.

A slight difference between these situations is that, in Q_n, it is very easy to prepare the values $a_j(m) \oplus x_j$ from the input bits x_j, but, in R_m, we need to consider a quantum circuit for computing the parity values $a_t \oplus \bigoplus_{k=1}^{m} t_k s_k$ from the input bits s_k, i.e., for the operation on $2^m + m - 1$ qubits defined as $|s\rangle \bigotimes_t |a_t\rangle \mapsto |s\rangle \bigotimes_t |a_t \oplus \bigoplus_{k=1}^{m} t_k s_k\rangle$ for any $s \in \{0,1\}^m$ and $a_t \in \{0,1\}$. If we have $m2^{m-1}$ *initialized* ancillary qubits, it is easy to construct an $O(m)$-depth quantum circuit for the operation using the following algorithm:

1. Prepare 2^{m-1} copies of s_k on the ancillary qubits for every $1 \le k \le m$ in parallel.
2. Compute the parity value $a_t \oplus \bigoplus_{k=1}^{m} t_k s_k$ for every $t \in \{0,1\}^m \setminus \{0^m\}$ in parallel.

To implement Step 1, we apply fan-out gates on $2^{m-1} + 1$ qubits, each of which can be decomposed into an $O(m)$-depth quantum circuit. Since it is easy to construct an $O(\log m)$-depth quantum circuit for PA_m using a binary tree structure, we can implement Step 2 using a parallel application of such circuits. If we replace the initialized ancillary qubits with uninitialized ones, the circuit does not work. However, applying the circuit again yields the desired values. In fact, the first circuit outputs $a_t \oplus \bigoplus_{k=1}^{m} t_k s_k \oplus d$ for some $d \in \{0,1\}$ that is computed from the (unknown) values in $\{0,1\}$ representing the initial state of the uninitialized ancillary qubits, and the second one outputs $a_t \oplus \bigoplus_{k=1}^{m} t_k s_k \oplus d \oplus d = a_t \oplus \bigoplus_{k=1}^{m} t_k s_k$ as desired. Using this circuit, we construct R_m and show the following lemma. The details can be found in the full version of the paper [21].

▶ **Lemma 8.** *The circuit R_m computes g_m. It uses no initialized ancillary qubit and $O(m2^m)$ uninitialized ancillary qubits, and its depth is $O(m^2)$, when the elementary gate set is $\mathcal{G}$.*

Combining R_m with Q_n immediately implies Theorem 1:

Proof of Theorem 1. By Lemmas 6, 7, and 8, we can use Q_n and R_m to implement the algorithm for $f_n \in \mathcal{S}_n$ described at the beginning of Section 3 and the whole circuit has the desired complexity.   ◀

4   Classical Algorithms with Access to Shallow Quantum Circuits

Let $p(n)$ be a polynomial and C_n be a constant-depth quantum circuit on n qubits consisting of the gates in $\mathcal{G}$. The problem $\mathrm{MAT}(p(n), C_n)$ is to compute a real number α_x such that $|\alpha_x - |\langle 0^n | C_n | x\rangle|^2| \le 1/p(n)$ for any input $x \in \{0,1\}^n$. For any $x, w \in \{0,1\}^n$, we define $F_n(x,w) = \langle x | C_n^\dagger (\bigotimes_{j=1}^{n} Z_j^{w_j}) C_n | x\rangle$, where $w = w_1 \cdots w_n$ and Z_j is Z applied to the j-th qubit of C_n. As shown in [17], $\mathrm{MAT}(p(n), C_n)$ can be solved with probability exponentially (in n) close to 1 if there exists a probabilistic algorithm A_{F_n} such that, for any $x, w \in \{0,1\}^n$, the probability that $|A_{F_n}(x,w) - F_n(x,w)| \le 0.5/p(n)$ is exponentially close to 1. In fact, due to the Chernoff-Hoeffding bound, the algorithm for $\mathrm{MAT}(p(n), C_n)$ on input $x \in \{0,1\}^n$ is described with some $K = n^{O(1)}$ as follows: Choose $w(j) \in \{0,1\}^n$ uniformly at random and compute $A_{F_n}(x, w(j))$ for every $1 \le j \le K$, and output $(1/K) \sum_{j=1}^{K} A_{F_n}(x, w(j))$.

The probabilistic algorithm A_{F_n} in [17] can be considered as a repetition of a commuting quantum circuit D_{2n} for the Hadamard test with $2n$ input qubits and one output qubit. For any $x, w \in \{0, 1\}^n$, the output of D_{2n} with the input qubits initialized to $|x\rangle|w\rangle$ and output qubit initialized to $|0\rangle$ is 0 with probability $(1 + F_n(x, w))/2$. Thus, when the outputs 0 and 1 are regarded as 1 and -1, respectively, due to the Chernoff-Hoeffding bound, A_{F_n} is described with some $L = n^{O(1)}$ as follows, where the input is the pair of x and w: Perform D_{2n} with the input qubits initialized to $|x\rangle|w\rangle$ and output qubit initialized to $|0\rangle$, and obtain its output $z_j(x, w) \in \{1, -1\}$ for every $1 \leq j \leq L$. After that, output $(1/L) \sum_{j=1}^{L} z_j(x, w)$.

We construct a parallelized version of the Hadamard test, denoted by E_{2n}, by using uninitialized ancillary qubits. Although the standard Hadamard test is a sequential application of controlled gates with the same control qubit, roughly speaking, E_{2n} first prepares the copies of the state of the control qubit on uninitialized ancillary qubits and then applies the gates in parallel by using the copies. To be precise, let $x = x_1 \cdots x_n, w = w_1 \cdots w_n \in \{0, 1\}^n$. We prepare $2n$ input qubits $X_1, \ldots, X_n, W_1, \ldots, W_n$, one output qubit Y, and n uninitialized ancillary qubits $G(1), \ldots, G(n)$, where X_j, W_j, and Y are initialized to $|x_j\rangle, |w_j\rangle$, and $|0\rangle$, respectively. The initial state of the uninitialized ancillary qubits is arbitrary. The circuit E_{2n} is defined as follows:

1. Apply a H gate to Y.
2. Apply a fan-out gate on $n+1$ qubits to $G(1), \ldots, G(n)$, and Y, where Y is the control qubit.
3. Apply C_n to $X_1, \ldots, X_{n-1}$, and X_n.
4. Apply a 2-controlled Z gate to $G(j), X_j$, and W_j for every $1 \leq j \leq n$ in parallel.
5. Apply $C_n^\dagger$ to $X_1, \ldots, X_{n-1}$, and X_n.
6. Apply the gates in Step 2 and Step 1 (in this order).

Each fan-out gate can be decomposed into an $O(\log n)$-depth quantum circuit. Moreover, a 2-controlled Z gate can be decomposed into a constant number of the gates in $\mathcal{G}$ [1, 18]. Thus, E_{2n} is an $O(\log n)$-depth circuit consisting of the gates in $\mathcal{G}$. It has the desired output probability distribution. The proof can be found in the full version of the paper [21].

▶ **Lemma 9.** *For any $x, w \in \{0, 1\}^n$, the output of E_{2n} with the input qubits initialized to $|x\rangle|w\rangle$ and output qubit initialized to $|0\rangle$ is 0 with probability $(1 + F_n(x, w))/2$.*

This lemma immediately implies Theorem 2:

Proof of Theorem 2. We replace D_{2n} in the above-mentioned algorithm for $\text{MAT}(p(n), C_n)$ with E_{2n}. By Lemma 9, the output probability distribution of E_{2n} is the same as that of D_{2n}. Thus, as with the original algorithm, the resulting algorithm solves $\text{MAT}(p(n), C_n)$.   ◀

5   Limitations of Uninitialized Ancillary Qubits

5.1   Our Idea for Proving Theorem 3

For any integer $s \geq 1$, an s-controlled Toffoli gate is decomposed into an s-controlled Z gate sandwiched between two H gates [8]. Thus, to prove Theorem 3, it suffices to consider an unbounded Z gate in place of an unbounded Toffoli gate. We assume on the contrary that there exists a depth-d quantum circuit C_n for PA_n with n input qubits, one output qubit, $p = O(\log n)$ initialized ancillary qubits, and $q = n^{O(1)}$ uninitialized ancillary qubits such that it consists of the gates in $\mathcal{G}$, unbounded fan-out gates on $(\log n)^{O(1)}$ qubits, and unbounded Z gates, where $d = O((\log n)^\delta)$ for some constant $0 \leq \delta < 1$. When all unbounded Z gates in C_n act on a small number of qubits, such as $O(\log n)$ qubits, since d is sufficiently small, the proof of Bera [3] implies that there exists an input qubit of C_n such that the output of
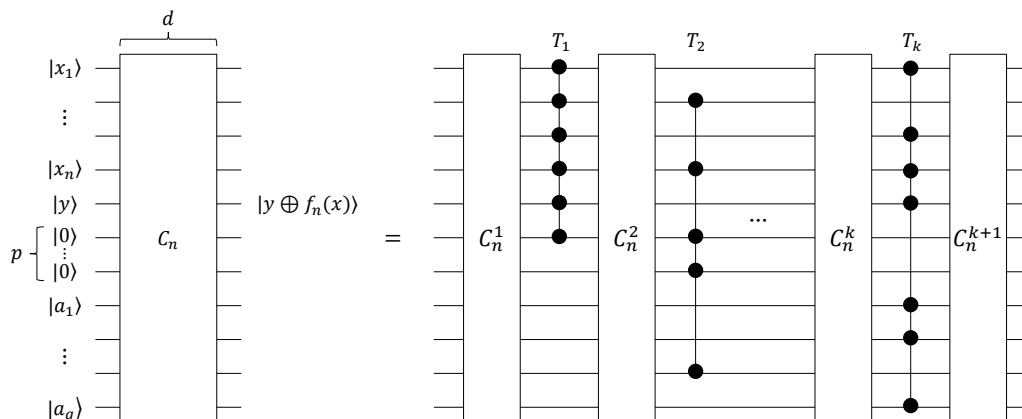
Figure 2 Circuit C_n for f_n and its decomposition. The initial states of the input qubits, output qubit, and uninitialized ancillary qubits are $|x_1\rangle \cdots |x_n\rangle$, $|y\rangle$, and $|a_1\rangle \cdots |a_q\rangle$, respectively, for any $x = x_1 \cdots x_n \in \{0, 1\}^n$, $y \in \{0, 1\}$, and $a_1 \cdots a_q \in \{0, 1\}^q$. Gates $T_1, \ldots, T_k$ are unbounded Z gates.

C_n does not depend on the input qubit. Thus, C_n cannot compute PA_n since the output of PA_n changes if any one of the n input bits changes. This contradicts the assumption.

The remaining case is when there exists an unbounded Z gate on a large number of qubits. Let $\tilde{C}_n$ be the circuit obtained from C_n by removing all such gates. Bera [3] showed that, when C_n does not have any ancillary qubit, it is well approximated by $\tilde{C}_n$ in the sense that, when the state of the input qubits is a computational basis state chosen uniformly at random, the output of C_n coincides with that of $\tilde{C}_n$ with high probability. Since C_n computes PA_n, $\tilde{C}_n$ computes PA_n with high probability. Thus, we obtain a contradiction as in the above case since all gates in $\tilde{C}_n$ act on a small number of qubits. To apply this idea to our setting, we show that C_n with p initialized ancillary qubits and q uninitialized ancillary qubits in state $|a\rangle$ for some $a \in \{0, 1\}^q$ is well approximated (in the sense described above) by $\tilde{C}_n$ with the same state. The former circuit computes PA_n since C_n with an arbitrary initial state of the uninitialized ancillary qubits computes PA_n. Thus, the latter circuit computes PA_n with high probability, and we obtain a contradiction as in the above simple case.

5.2 Analysis of a General Circuit and Its Application

We analyze a general depth-d quantum circuit C_n with n input qubits, one output qubit, and p initialized and q uninitialized ancillary qubits such that it consists of the gates in $\mathcal{G}$, unbounded fan-out gates, and unbounded Z gates. Its key property is described as follows:

▶ **Lemma 10** ([3, 2]). *Let C_n be a depth-d quantum circuit with n input qubits and one output qubit (possibly with ancillary qubits). If all gates in C_n act on at most w qubits, then the output of C_n can depend only on the states of at most w^d input qubits.*

Let $t \geq 2$ be an integer and $\mathcal{G}_t$ be the set of all unbounded Z gates in C_n that act on more than or equal to t qubits. We consider the case where $\mathcal{G}_t \neq \emptyset$ and assume that $\mathcal{G}_t = \{T_1, \ldots, T_k\}$ for some $k \geq 1$, where, for any $1 \leq l \leq k$, if T_l is in layer L of C_n, then T_{l+1} is in layer $L' \geq L$. We decompose C_n into the gates in $\mathcal{G}_t$ and the other parts as depicted in Fig. 2, where C_n computes a Boolean function f_n on n bits and C_n^j is a quantum circuit consisting of gates that are not in $\mathcal{G}_t$ for any $1 \leq j \leq k + 1$. Such a decomposition is not unique in general, but the point is to fix a decomposition. For any $1 \leq l \leq k$, we define a quantum circuit V_l as follows: $V_1 = C_n^1$ and $V_l = C_n^l T_{l-1} V_{l-1}$ for any $2 \leq l \leq k$. We also

define $\Delta_l(x, y, b) = ||T_l V_l|x \circ y \circ b\rangle - V_l|x \circ y \circ b\rangle||$ and $\Delta(x, y, b) = ||C_n|x \circ y \circ b\rangle - \tilde{C}_n|x \circ y \circ b\rangle||$ for any $x \in \{0,1\}^n$, $y \in \{0,1\}$, and $b \in \{0,1\}^{p+q}$. Here, the symbol "$\circ$" represents the concatenation of bit strings, $|||v\rangle|| = \sqrt{\langle v|v \rangle}$ for any vector $|v\rangle$, and $\tilde{C}_n = C_n^{k+1} C_n^k \cdots C_n^2 C_n^1$. Let U_n be a random variable uniformly distributed over $\{0,1\}^n$. We evaluate the probability $\Pr[\Delta(U_n, y, b) < \varepsilon]$ as follows. The proof can be found in the full version of the paper [21].

▶ **Lemma 11.** $\Pr[\Delta(U_n, y, b) < \varepsilon] \geq 1 - (k^2/\varepsilon^2) \sum_{l=1}^{k} \mathrm{E}[\Delta_l(U_n, y, b)^2]$ *for any* $\varepsilon > 0$, $y \in \{0,1\}$, *and* $b \in \{0,1\}^{p+q}$.

To evaluate the value $\sum_{l=1}^{k} \mathrm{E}[\Delta_l(U_n, y, b)^2]$, let t_l be the number of qubits on which T_l acts, $u_l = n + p + q + 1 - t_l$, and $t_{\min} = \min\{t_l | 1 \leq l \leq k\}$. We define $V_l|x \circ y \circ b\rangle = \sum_{i \in \{0,1\}^{t_l}} \sum_{j \in \{0,1\}^{u_l}} g_{xoyob}^{(l)}(i \circ j)|i \circ j\rangle$ for any $x \in \{0,1\}^n$, $y \in \{0,1\}$, and $b \in \{0,1\}^{p+q}$, where $g_{xoyob}^{(l)}(i \circ j)$ is a complex number. The qubits represented by $i \in \{0,1\}^{t_l}$ correspond to the qubits on which T_l acts. Of course, for any $1 \leq l \leq k$, T_l does not always act on the first t_l qubits in C_n. We therefore need to apply some permutation of all qubits; however, since such a permutation does not affect Lemma 13, which is the key to Theorem 3, we omit it.

We evaluate the above value as follows. The proof can be found in the full version [21].

▶ **Lemma 12.** $\sum_{l=1}^{k} \mathrm{E}[\Delta_l(U_n, y, b)^2] \leq k 2^{p+q+3}/2^{t_{\min}}$ *for any* $y \in \{0,1\}$ *and* $b \in \{0,1\}^{p+q}$. *Moreover, there exists some* $a \in \{0,1\}^q$ *such that* $\sum_{l=1}^{k} \mathrm{E}[\Delta_l(U_n, 0, 0^p \circ a)^2] \leq k 2^{p+3}/2^{t_{\min}}$.

Lemmas 11 and 12 immediately imply the following evaluation:

▶ **Lemma 13.** *There exists some* $a \in \{0,1\}^q$ *such that* $\Pr[\Delta(U_n, 0, 0^p \circ a) < \varepsilon] \geq 1 - k^3 2^{p+3}/(\varepsilon^2 2^{t_{\min}})$ *for any* $\varepsilon > 0$.

Lemmas 10 and 13 imply Theorem 3 as follows:

Proof of Theorem 3. We assume on the contrary that there exists a quantum circuit C_n for PA$_n$ described in Section 5.1. Since $p = O(\log n)$, there exists a constant $c > 0$ such that $p \leq c \log n$ when n is sufficiently large. We define $t = (c + 4) \log(n + p + q + 1)$ and consider $\mathcal{G}_t$ described above. When $\mathcal{G}_t = \emptyset$, all gates in C_n act on at most $w = (\log n)^{O(1)}$ qubits. By Lemma 10, the output of C_n can depend only on the states of at most $w^d = o(n)$ input qubits. Thus, there exists an input qubit of C_n such that the output of C_n does not depend on the input qubit. This yields a contradiction as described in Section 5.1.

We consider the remaining case where $\mathcal{G}_t \neq \emptyset$. In this case, we apply the above analysis of a general circuit. It holds that $p \leq c \log n$, $k \leq (n + p + q + 1)d/t_{\min}$, and $t_{\min} \geq (c + 4) \log(n + p + q + 1)$. Thus, by Lemma 13 with $\varepsilon = 0.1$,

$$\Pr[\Delta(U_n, 0, 0^p \circ a) < 0.1] \geq 1 - \left(\frac{d}{(c+4)\log(n+p+q+1)} \right)^3 \frac{800 n^c}{(n+p+q+1)^{c+1}}$$

for some $a \in \{0,1\}^q$. Let us express this value on the right-hand side by $1 - \gamma$. Thus, there exists a set $S \subseteq \{0,1\}^n$ such that S has at least $2^n(1-\gamma)$ elements and, for any $x \in S$, $\Delta(x, 0, 0^p \circ a) < 0.1$. Since γ goes to 0 as n goes to infinity, $2^n(1-\gamma) > 2^{n-1}$ when n is sufficiently large. A simple calculation shows that, for any $x \in \{0,1\}^n$ satisfying $\Delta(x, 0, 0^p \circ a) < 0.1$, the output of $\tilde{C}_n|x \circ 0 \circ 0^p \circ a\rangle$ coincides with that of $C_n|x \circ 0 \circ 0^p \circ a\rangle$ with probability of at least $1 - 0.1^2 = 0.99$ [3, 2]. When the initial state of the uninitialized ancillary qubits is $|a\rangle$, C_n computes PA$_n$. Thus, for any $x \in S$, the output of $\tilde{C}_n|x \circ 0 \circ 0^p \circ a\rangle$ is PA$_n(x)$ with probability of at least 0.99. This contradicts the fact obtained by the following argument. Since all gates in $\tilde{C}_n$ act on at most $(\log n)^{O(1)}$ qubits, as described for the case where $\mathcal{G}_t = \emptyset$, by Lemma 10, there exists an input qubit of $\tilde{C}_n$ such that the output of $\tilde{C}_n$ does not depend on the input qubit. This implies that, for at most 2^{n-1} elements $x \in \{0,1\}^n$, the output of $\tilde{C}_n|x \circ 0 \circ 0^p \circ a\rangle$ is PA$_n(x)$ with probability greater than 0.5. ◀

References

1 A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, and H. Weinfurter. Elementary gates for quantum computation. *Physical Review A*, 52(5):3457–3467, 1995.

2 D. Bera. *Quantum circuits: power and limitations*. PhD thesis, Boston University, 2010.

3 D. Bera. A lower bound method for quantum circuits. *Information Processing Letters*, 111(15):723–726, 2011.

4 S. Bravyi, D. Gosset, and R. König. Quantum advantage with shallow circuits, 2017. arXiv:1704.00690.

5 H. Buhrman, R. Cleve, M. Koucký, B. Loff, and F. Speelman. Computing with a full memory: catalytic space. In *Proceedings of the 46th ACM Symposium on Theory of Computing (STOC)*, pages 857–866, 2014.

6 R. Cleve and J. Watrous. Fast parallel circuits for the quantum Fourier transform. In *Proceedings of the 41st IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 526–536, 2000.

7 D. P. DiVincenzo. The physical implementation of quantum computation. *Fortschritte der Physik*, 48(9–11):771–783, 2000.

8 M. Fang, S. Fenner, F. Green, S. Homer, and Y. Zhang. Quantum lower bounds for fanout. *Quantum Information and Computation*, 6(1):46–57, 2006.

9 S. Fenner, F. Green, S. Homer, and Y. Zhang. Bounds on the power of constant-depth quantum circuits. In *Proceedings of Fundamentals of Computation Theory (FCT)*, volume 3623 of *Lecture Notes in Computer Science*, pages 44–55, 2005.

10 F. Green, S. Homer, C. Moore, and C. Pollett. Counting, fanout, and the complexity of quantum ACC. *Quantum Information and Computation*, 2(1):35–65, 2002.

11 P. Høyer and R. Špalek. Quantum fan-out is powerful. *Theory of Computing*, 1(5):81–103, 2005.

12 S. Jukna. *Boolean Function Complexity: Advances and Frontiers*. Springer, 2012.

13 E. Knill and R. Laflamme. Power of one bit of quantum information. *Physical Review Letters*, 81(25):5672–5675, 1998.

14 T. D. Ladd, F. Jelezko, R. Laflamme, Y. Nakamura, C. Monroe, and J. L. O'Brien. Quantum computing. *Nature*, 464:45–53, 2010.

15 G. De las Cuevas, W. Dür, M. van den Nest, and M. A. Martin-Delgado. Quantum algorithms for classical lattice models. *New Journal of Physics*, 13(093021), 2011.

16 C. Moore and M. Nilsson. Parallel quantum computation and quantum codes. *SIAM Journal on Computing*, 31(3):799–815, 2001.

17 X. Ni and M. van den Nest. Commuting quantum circuits: efficient classical simulations versus hardness results. *Quantum Information and Computation*, 13(1&2):54–72, 2013.

18 M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.

19 Y. Takahashi and N. Kunihiro. A quantum circuit for Shor's factoring algorithm using $2n+2$ qubits. *Quantum Information and Computation*, 6(2):184–192, 2006.

20 Y. Takahashi and S. Tani. Collapse of the hierarchy of constant-depth exact quantum circuits. *Computational Complexity*, 25(4):849–881, 2016.

21 Y. Takahashi and S. Tani. Power of uninitialized qubits in shallow quantum circuits, 2017. arXiv:1608.07020v3.

22 Y. Takahashi, T. Yamazaki, and K. Tanaka. Hardness of classically simulating quantum circuits with unbounded Toffoli and fan-out gates. *Quantum Information and Computation*, 14(13&14):1149–1164, 2014.

23 J. Watrous. On the complexity of simulating space-bounded quantum computations. *Computational Complexity*, 12(1–2):48–84, 2003.