

Universal Quantum Circuits

Debajyoti Bera¹, Stephen Fenner², Frederic Green³, and Steve Homer¹

¹ Boston University, Department of Computer Science, Boston, MA 02134. *

² University of South Carolina, Department of Computer Science and Engineering,
Columbia, SC 29208. **

³ Clark University, Department of Mathematics and Computer Science, Worcester,
MA 01610. ***

Abstract. We define and construct efficient depth-universal and almost-size-universal quantum circuits. Such circuits can be viewed as general-purpose simulators for central classes of quantum circuits and can be used to capture the computational power of the circuit class being simulated. For depth we construct universal circuits whose depth is the same order as the circuits being simulated. For size, there is a log factor blow-up in the universal circuits constructed here. We prove that this construction is nearly optimal.

1 Introduction

Like resource-bounded universal Turing machines, efficiently constructed universal circuits capture the hardness of languages computed by circuits in a given circuit class. As a result, the study of the existence and complexity of universal circuits for quantum circuit classes provides insight into the computational strength of such circuits, as well as their limits.

There is both a theoretical and a practical aspect to this study. The existence of a universal circuit family for a complexity class defined by resource bounds (depth, size, gate width, etc.) provides an upper bound on the resources needed to compute any circuit in that class. It also opens up possibilities for proving lower bounds on the hard languages in the class, as such bounds would follow from a lower bound proof for the language computed by a universal circuit family for the circuit class.

More precisely, the specific, efficient construction of a universal circuit for a class of circuits yields, for a fixed input size, a single circuit which can be used to carry out the computation of every circuit (with that same input size) in that family, basically a chip or processor for that class of circuits. The more efficient the construction of the universal circuit, the smaller the processor for that class.

* `{(dbera|homer)}@cs.bu.edu`. Partially supported by the National Security Agency (NSA) and Advanced Research and Development Agency (ARDA) under Army Research Office (ARO) contract number DAAD 19-02-1-0058.

** `fenner@cse.sc.edu`. Partially supported by NSF grant CCF-05-15269.

*** `fgreen@black.clarku.edu`. Partially supported by the NSA and ARDA under ARO contract number DAAD 19-02-1-0058.

Furthermore, the universal circuit is in a sense a compiler for all possible computations of all circuits in this family. It can be used to efficiently program all possible computations capable of being carried out by circuits in this circuit class, and in doing so automatically acts as a general purpose simulator and with as little loss of efficiency as is possible.

In the case of quantum circuits there are particular issues relating to the requirements that computations must be clean and reversible which come into play, and to an extent complicate the classical methods. Still much of our motivation for this work originates with classical results due to Cook, Valiant, and others [CH85, Val76]. Cook and Hoover considered depth universality and described a depth-universal uniform circuit family for circuits of depth $\Omega(\log n)$. Valiant studied size universality and showed how to construct universal circuits of size $O(s \log s)$ to simulate any circuit of size s . (See Section 1.1.)

Definition 1 (Universal Quantum Circuits). Fix $n > 0$ and let \mathcal{C} be a collection of quantum circuits on n qubits. A quantum circuit U on $n+m$ qubits is *universal for \mathcal{C}* if, for every circuit $C \in \mathcal{C}$, there is a string $x \in \{0, 1\}^m$ (the *encoding*) such that for all strings $y \in \{0, 1\}^n$ (the *data*),

$$U(|y\rangle \otimes |x\rangle) = C|y\rangle \otimes |x\rangle.$$

The circuit collections we are interested in are usually defined by bounding various parameters such as the size (number of gates), depth (number of layers of gates acting simultaneously on disjoint sets of qubits), or palette of allowed gates (e.g., Hadamard, $\pi/8$, CNOT).

As in the classical case, we also want our universal circuits to be *efficient* in various ways. For one, we restrict them to using the same gate family as the circuits they simulate. We may also want to restrict their size or the number m of qubits they use for the encoding. We are particularly concerned with the depth of universal circuits.

Definition 2 (Depth-Universal Quantum Circuits). Fix a family \mathcal{F} of unitary quantum gates. A family of quantum circuits $\{U_{n,d}\}_{n,d>0}$ is *depth-universal over \mathcal{F}* if

1. $U_{n,d}$ is universal for n -qubit circuits with depth $\leq d$ using gates from \mathcal{F} ,
2. $U_{n,d}$ only uses gates drawn from \mathcal{F} ,
3. $U_{n,d}$ has depth $O(d)$, and
4. the number of encoding qubits of $U_{n,d}$ is polynomial in n and d .

Depth-universal circuits are desirable because they can simulate any circuit within a constant slow-down factor. Thus they are as time-efficient as possible.

Our first result, presented in Section 3, shows that depth-universal quantum circuits exist for the gate families $\mathcal{F} = \{H, T\} \cup \{F_n \mid n \geq 1\}$ and $\mathcal{F}' = \{H, T\} \cup \{F_n \mid n \geq 1\} \cup \{\wedge_n(X) \mid n \geq 1\}$, where H and T are the Hadamard and $\pi/8$ gates, respectively, and F_n and $\wedge_n(X)$ are the $(n+1)$ -qubit fanout and $(n+1)$ -qubit Toffoli gates, respectively (see Section 2).

Theorem 3. *Depth-universal quantum circuits exist over \mathcal{F} and over \mathcal{F}' . Such circuits use $O(n^2d)$ qubits and can be built log-space uniformly in n and d .*

Note that the results for the two circuit families are independent, because it is not known whether n -qubit Toffoli gates can be implemented exactly in constant depth using single-qubit gates and fanout gates, although they can be approximated this way [HS05].

It would be nice to find depth-universal circuits over families of bounded-width gates⁴ such as $\{H, T, \text{CNOT}\}$. Depth-universal circuits with bounded-width gates, if they exist, must have depth $\Omega(\log n)$ and thus can only depth-efficiently simulate circuits with depth $\Omega(\log n)$. This can be easily seen as follows: Suppose all you wanted was a universal circuit U for depth-1 circuits on n qubits that use CNOT gates *only*. Since any pair of the n qubits could potentially be connected with a CNOT gate, that pair must be connected somehow (indirectly perhaps) within the circuit U . Thus any data input qubit can potentially affect any of the other $n-1$ data output qubits. Since U only has constant-width gates, the number of qubits affected by any given data input increases by only a constant factor per layer, and so U must have $\Omega(\log n)$ layers.

One can therefore only hope to find depth-universal circuits for circuits of depth $\Omega(\log n)$ over bounded-width gates. Although such circuits exist in the classical case (see below), we are unable to construct them in the quantum case (see Section 6).

1.1 Other relevant work

The study of quantum circuit complexity was originated by Yao [Yao]. The basic definitions and first results in this research area can be found in Nielsen and Chuang [NC00]. Most of the research on universal quantum circuit classes deals with finding small, natural, universal sets of gates which can be used in quantum circuits to efficiently simulate any quantum computation. Our problem and point of view here is quite different. We have the goal of constructing, for a natural class C of quantum circuits, a single family of quantum circuits which can efficiently simulate all circuits on the class C . In this paper we consider classes C which have significant resource bounds (small or even constant depth, or fixed size) and ask that the corresponding universal circuits family to have similar depth or size bounds.

Cook and Hoover [CH85] considered the problem of constructing general-purpose classical (Boolean) circuits using gates with fanin two. They asked whether, given n, c, d , there is a circuit U of size $c^{O(1)}$ and depth $O(d)$ that can simulate any n -input circuit of size c and depth d . Cook and Hoover constructed a depth-universal circuit for depth $\Omega(\log n)$ and polynomial size, but which takes as input a nonstandard encoding of the circuit, and they also presented a circuit with depth $O(\log n \log \log n)$ to convert the standard encoding of the circuit to the required encoding.

⁴ The width of a gate is the number of qubits it acts upon.

Valiant looked at a similar problem—trying to minimize the size of the universal circuit [Val76]. He considered classical circuits built from fanin 2 gates (but with unbounded fanout) and embedded the circuit in a larger universal graph. Using switches at key vertices of the universal graph, any graph (circuit) can be embedded in it. He managed to create universal graphs for different types of circuits and showed how to construct a $O(c \log c)$ -size and $O(c)$ -depth universal circuit. He also showed that his constructions have size within a constant multiplicative factor of the information theoretic lower bound.

For quantum circuits, Nielsen and Chuang (in [NC97]) considered the problem of building generic universal circuits, or *programmable universal gate arrays* as they call them. Their universal circuits work on two quantum registers, a data register and a program register. They do not consider any size or depth bound on the circuits and show that simulating every possible unitary operation requires completely orthogonal programs in the program register. Since there are infinitely many possible unitary operations, any universal circuit would require an infinite number of qubits in the program register. This shows that it is not possible to have a generic universal circuit which works for all circuits of a certain input length. However they showed that it is possible to construct an extremely weak type of probabilistic universal circuit with size linear in the number of inputs to the simulated circuit.

Sousa and Ramos considered a similar problem of creating a universal quantum circuit to simulate any quantum gate [SR07]. They construct a basic building block which can be used to implement any single-qubit or CNOT gate on n qubits by switching certain gates on and off. They showed how to combine several of these building blocks to implement any n -qubit quantum gate.

1.2 Outline of the paper

For the rest of the paper, we will use U to denote the universal circuit and C to denote the circuit being simulated. We define the quantum gates we will use in Section 2. The construction of depth-universal circuits is in Section 3. We briefly describe the construction of almost-size-universal quantum circuits in Section 4. We mention a couple of miscellaneous results in Section 5.

2 Preliminaries

We assume the standard notions of quantum states, quantum circuits, and quantum gates described in [NC00], in particular, H (Hadamard), T ($\pi/8$), $S = T^2$ (phase), and CNOT (controlled NOT). We will also need some additional gates, which we now motivate.

The depth-universal circuits we construct require the ability to feed the output of a single gate to many other gates. While this operation, commonly known as fanout, is common in classical circuits, copying an arbitrary quantum state unitarily is not possible in quantum circuits due to the no-cloning

theorem [NC00]. It turns out that we can construct our circuits using a classical notion of fanout operation, defined as the *fanout gate* $F_n : |c, t_1, \dots, t_n\rangle \mapsto |c, c \oplus t_1, \dots, c \oplus t_n\rangle$ for any of the standard basis states $|c\rangle$ (the control) and $|t_1\rangle, \dots, |t_n\rangle$ (the targets) and extended linearly to other states⁵ [FFGHZ06]. F_n can be constructed in depth $\lg n$ using CNOT gates. We need to use unbounded fanout gates to achieve full depth universality. We also use the *unbounded Toffoli gate* $\wedge_n(X) : |c_1, \dots, c_n, t\rangle \mapsto |c_1, \dots, c_n, t \oplus \bigwedge_{i=1}^n c_i\rangle$. We reserve the term “Toffoli gate” to refer to the (standard) Toffoli gate $\wedge_2(X)$, which is defined on three qubits.

In addition to the fanout gate, our construction requires us to use controlled versions of the gates used in the simulated circuit. For most of the commonly used basis sets of gates (e.g., Toffoli gate, Hadamard gate, and phase gate S), the gates themselves are sufficient to construct their controlled versions (e.g., a controlled Hadamard gate can be constructed using a Toffoli gate and Hadamard and phase gates). Depth or size universality requires that the controlled versions of the gates should be constructible using the gates themselves within proper depth or size, as required.

Definition 4 (Closed under controlled operation). A set of quantum gates $G = \{G_1, \dots\}$ is said to be *closed under controlled operation* if for each $G_i \in G$, the controlled version of the gate $C\text{-}G_i|c\rangle|t\rangle \longrightarrow |c\rangle G_i^c|t\rangle$ can be implemented in constant depth and size using the gates in G . Here, $|c\rangle$ is a single qubit and G_i could be a single or a multi-qubit gate.

Note that $\text{CNOT} = F_1$, and given H , T , and CNOT we can implement the Toffoli gate via a standard constant-size circuit [NC00]. We can implement the phase gate S as T^2 , and since $T^8 = I$, we can implement $S^\dagger = T^6$ and $T^\dagger = T^7$.

A *generalized Z gate*, which we will hereafter refer to simply as a *Z gate*, is an extension of the single-qubit Pauli Z gate ($|x\rangle \mapsto (-1)^x|x\rangle$) to multiple qubits:

$$|x_1, \dots, x_n\rangle \xrightarrow{Z} (-1)^{x_1 x_2 \dots x_n} |x_1, \dots, x_n\rangle.$$

A Z gate can be constructed easily (in constant depth and size) from a single unbounded Toffoli gate (and vice versa) by conjugating the target qubit of the unbounded Toffoli gate with H gates (i.e., placing H on both sides of the Toffoli gate on its target qubit).

Similarly, a *Z-fanout gate* Z_n applies the single-qubit Z gate to each of n target qubits if the control qubit is set:

$$|c, t_1, \dots, t_n\rangle \xrightarrow{Z_n} (-1)^{c \cdot (t_1 + \dots + t_n)} |c, t_1, \dots, t_n\rangle.$$

A Z_n gate can be constructed from a single F_n gate and vice versa in constant depth (although not constant size) by conjugating each target with H gates. So, in our depth-universal circuit construction, we can use these either or both of these types of gates. Similarly for unbounded Toffoli versus Z gates. Z gates

⁵ This does not contradict the no-cloning theorem as only classical states are copied.

and Z -fanout gates are important because they only change the phase, leaving the values of the qubits intact (they are represented by diagonal matrices in the computational basis). This allows us to use a trick due to Høyer and Špalek [HS05] and run all possible gates for a layer in parallel.

3 Depth-universal quantum circuits

In this section, we prove Theorem 3, i.e., that depth-universal circuits exist for each of the gate families

$$\mathcal{F} = \{H, T\} \cup \{F_n \mid n \geq 1\},$$

$$\mathcal{F}' = \{H, T\} \cup \{F_n \mid n \geq 1\} \cup \{\wedge_n(X) \mid n \geq 1\}.$$

We first give the proof for \mathcal{F} then show how to modify it for \mathcal{F}' .

The depth-universal circuit U we construct simulates the input circuit C layer by layer, where a layer consists of the collection of all its gates at a fixed depth. C is encoded in a slightly altered form, however. First, all the fanout gates in C are replaced with Z -fanout gates on the same qubits with H gates conjugating the targets. At worst, this may roughly double the depth of C (adjacent H gates cancel). Each layer of the resulting circuit is then separated into three adjacent layers: the first having only the H gates of the original layer, the second only the T gates, and the third only the Z -fanout gates. U then simulates each layer of the modified C by a constant number of its own layers. We describe next how these layers are constructed.

Simulating single-qubit gates. The circuit to simulate an n -qubit layer of single-qubit gates of type G , say, consists of a layer of controlled- G gates where the control qubits are fed from the encoding and the target qubits are the data qubits. Figure 1 shows a layer of G gates, where $G \in \{H, T\}$, controlled using H , S , T , CNOT, and Toffoli gates. To simulate G gates on qubits i_1, \dots, i_k , say, set c_{i_1}, \dots, c_{i_k} to 1 and the rest of the c -qubits to 0.




Fig. 1. Simulating a layer of single-qubit G gates with controlled G gates. The ancilla in the implementation of the controlled T gate is assumed part of the encoding. The ancilla is reset to 0 at the end and hence can be reused for implementing all T layers.

Simulating Z-fanout gates. The circuit to simulate a Z-fanout layer is shown in Figure 2. The top n qubits are the original data qubits. The rest are ancilla




Fig. 2. Simulating a layer of Z-fanout gates.

qubits. All the qubits are arranged in n blocks B_1, \dots, B_n of n qubits per block. The qubits in block B_i are labeled b_{i1}, \dots, b_{in} .

Each A_i subcircuit looks like Figure 3. The qubits c_{i1}, \dots, c_{in} are encoding qubits. The large gate between the two columns of Toffoli gates is a Z-fanout gate with its control on the i th ancilla (corresponding to b_{ii} and c_{ii}) and targets on all the other ancillæ.

Here is the state evolution from $|\mathbf{d}\rangle = |d_1 \dots d_n\rangle$, suppressing the c_{ij} qubits and ancillæ internal to the A_i subcircuits in the ket labels. Note that after the first layer of fanouts, each qubit b_{ij} carries the value d_j .

$$\begin{aligned}
 |\mathbf{d}, \mathbf{0}, \dots, \mathbf{0}\rangle &\mapsto |\mathbf{d}, \mathbf{d}, \dots, \mathbf{d}\rangle \\
 &\mapsto (-1)^{\sum_i d_i c_{ii} (\sum_{j \neq i} d_j c_{ij})} |\mathbf{d}, \mathbf{d}, \dots, \mathbf{d}\rangle \\
 &\mapsto (-1)^{\sum_i d_i c_{ii} (\sum_{j \neq i} d_j c_{ij})} |\mathbf{d}, \mathbf{0}, \dots, \mathbf{0}\rangle.
 \end{aligned}$$

To simulate some Z-fanout gate G of C whose control is on the i th qubit, say, we do this in block B_i by setting c_{ii} to 1 and setting c_{ij} to 1 for every j where the j th qubit is a target of G . All the other c -qubits in B_i are set to 0. We can do this in separate blocks for multiple Z-fanout gates on the same layer, because no two gates can share the same control qubit. Any c -qubits in unused blocks are set to 0.




Fig. 3. Subcircuit A_i in the simulation of Z -fanout gates.




Fig. 4. Subcircuit A_i for a layer of Z gates.

Simulating unbounded Toffoli gates. We can modify the construction above to accommodate unbounded Toffoli gates (gate family \mathcal{F}'), or equivalently Z gates, by breaking each layer of C into four adjacent layers, the first three being as before, and the fourth containing only Z gates. The top-level circuit to simulate a layer of Z gates looks just as before (Figure 2), except now each A_i subcircuit looks a bit different and is shown in Figure 4, where the central gate is a Z gate connecting the ancillæ.

As before, the qubits c_{i1}, \dots, c_{in} are encoding qubits. The XZX gates on c_{ii} multiply the overall phase by $(-1)^{c_{ii}}$. When the Z gate of A_i is applied, its j th contact point is in the state $\overline{b_{ij}c_{ij}}$. Note that $\overline{b_{ij}c_{ij}} = b_{ij}$ if $c_{ij} = 1$ and 1 otherwise. The Z gate then multiplies the overall phase by $(-1)^{\prod_j (\overline{b_{ij}c_{ij}})} = (-1)^{\prod_{j:c_{ij}=1} b_{ij}}$. The state thus evolves as given below:

$$\begin{aligned}
 |\mathbf{d}, \mathbf{0}, \dots, \mathbf{0}\rangle &\mapsto |\mathbf{d}, \mathbf{d}, \dots, \mathbf{d}\rangle \\
 &\mapsto (-1)^{\sum_i \overline{c_{ii}} + \prod_{j:c_{ij}=1} b_{ij}} |\mathbf{d}, \mathbf{d}, \dots, \mathbf{d}\rangle \\
 &\mapsto (-1)^{\sum_i \overline{c_{ii}} + \prod_{j:c_{ij}=1} b_{ij}} |\mathbf{d}, \mathbf{0}, \dots, \mathbf{0}\rangle
 \end{aligned}$$

To simulate some Z gate G of C whose first qubit is i , say, we do this in block B_i by setting c_{ii} to 1 and setting c_{ij} to 1 for every j where the j th qubit is part of G . All the other c -qubits in B_i are set to 0. As before, we can do this in separate blocks for multiple gates on the same layer, because no two gates can share the same first qubit. Any c -qubits in unused blocks are set to 0, and it is easy to check that this makes the block have no net effect.

4 Size-universal quantum circuits

Similar to a depth-universal circuit, a *size-universal circuit* is a universal circuit with the same order of the number of gates as the circuit it is simulating. Formally,

Definition 5. A family $\{U_{n,c}\}$ of universal circuits for n -qubit circuits of size $\leq c$ is *size-universal* if $\text{SIZE}(U_{n,c}) = O(c)$.

A simple counting argument shows that it is not possible to obtain a completely size-universal circuit for fanin-2 circuits. Consider all circuits with c fanin-2 gates where one input of each gate is the first qubit. There are $(n-1)^c$ possible circuits. Then consider similar circuits where there is no gate with input as the first qubit and continue recursively. Thus the number of possible fanin-2 circuits is $\Omega((n-1)^{c+1})$. Since all the encoding bits have to be connected to some of the fanin-2 gates in the universal circuit, it must have $\Omega(c \log n)$ gates.

We use Valiant's idea of universal graphs [Val76] to construct a universal family of fanin-2 circuits that are very close to the aforementioned lower bound. As before, we would like to simulate C by using the same set of gates used in C . Our construction works for any circuit using unbounded Toffoli gates and any set of single-qubit and 2-qubit gates closed under the controlled operation.

First we will define a universal directed acyclic graph with n special vertices (called *poles*) in which we can embed any circuit with n gates (considering the inputs also as gates). The embedding will map the wires in the circuit to paths in the graph.

Definition 6 (Edge-embedding [Val76]). An *edge-embedding* ρ of $G = (V, E)$ into $G' = (V', E')$ maps V one-to-one to V' and maps each edge $(i, j) \in E$ to a directed path $\rho(i) \rightsquigarrow \rho(j)$ in G' such that distinct edges are mapped to edge-disjoint paths.

The graph of any circuit of size n can be represented as a directed acyclic graph with vertices $\{1, \dots, n\}$ such that there is no edge from j to i for $i < j$ and each vertex has fanin and fanout 2. Let $\Gamma_2(n)$ be the set of all such graphs.

Definition 7 (Edge-universal graph [Val76]). A graph G' is *edge-universal* for $\Gamma_2(n)$ if it has distinct poles p_1, \dots, p_n such that any graph $G \in \Gamma_2(n)$ can be edge-embedded into G' where each vertex $i \in G$ is mapped to vertex $\rho(i) = p_i \in G'$.

Then, Valiant shows how to construct a universal graph.

Theorem 8 ([Val76]). *There is a constant k such that for all n there exists an acyclic graph G' that is edge-universal for $\Gamma_2(n)$, and G' has $kn \lg n$ vertices, each vertex having fanin and fanout 2.*

It is fairly easy to construct a universal circuit using the universal graph. In fact, the universal circuit for circuits with n inputs and c gates will be any edge-universal graph for $\Gamma_2(n+c)$.

Consider any such edge-universal graph G' . Then G' has $c' = k(n+c) \log(n+c)$ vertices for some k . These c' vertices include fixed poles $p_1, \dots, p_n, p_{n+1}, \dots, p_{n+c}$ and non-pole vertices. Create a quantum circuit C' with c' gates (including the inputs and outputs) where G' describes how the gates connect to each other. For each of the vertices p_1, \dots, p_n of G' , remove their incoming edges and replace the vertices by the input as shown in Figure 5. Replace each of the vertices p_{n+1}, \dots, p_{n+c} with a subcircuit that applies any of the single- or 2-qubit gates on the inputs, where the gate to apply is controlled by the encoding. E.g., Figure 7 shows the gates at a pole vertex in a universal circuit simulating CNOT and H gates. For a non-pole vertex, replace it with a subcircuit that swaps the incoming and outgoing wires (i.e., first input is connected to second output and second input is connected to first output) or directly connects them (i.e., first input is connected to first output and similarly for the second input). Again, the subcircuit is controlled by the encoding which controls whether to swap or directly connect (see Figure 6). The edge disjointness property guarantees that wires in the embedded circuit are mapped to paths in C' which can share a vertex but cannot share any edge.




Fig. 5. The gate for a pole vertex p_i is mapped to input x_i .




Fig. 6. The gates at a non-pole vertex v . The encoding bit c_v specifies if first output qubit should be mapped to first input or second input qubit and similarly for second output qubit.




Fig. 7. Example of the gates at a pole vertex v simulating a circuit with CNOT and H gates. The encoding bits c_v^g specify which kind of gate is at vertex v , and the c_v^d specify which qubit the gate acts on (for H gate) or which is the control qubit (for CNOT gate).

To simulate any fanin-2 circuit C with c gates acting on n qubits, construct the edge-universal graph G' for $\Gamma_2(n+c)$. Embed the graph of C into G' such that the input nodes of C are mapped to the poles p_1, \dots, p_n in G' . Now for each

gate of the circuit, consider the pole to which it was mapped. Set a bit in the encoding to denote the type of the gate at that pole. For the non-pole vertices, set a bit in the encoding to specify whether the two input values should be swapped or mapped directly to the two output values. The size of the encoding is $(n+c)(\log |\text{gates}|+1)+(|I_2(n+c)|-(n+c))$ which is $O(c \log c)$ for polynomial-size circuits. This construction gives us a universal circuit with a logarithmic blow-up in size.

Theorem 9. *There is a constant k and a family of universal circuits $U_{n,c}$ that can simulate every circuit with c gates acting on n qubits such that $\text{SIZE}(U_{n,c}) = k(n+c) \log(n+c)$.*

We can use a similar idea for circuits with unbounded fanin. First we decompose the unbounded fanin gates using bounded fanin gates (fanin 2 in this case). This is doable for most of the common unbounded fanin gates. For example, an unbounded Toffoli gate of size f can be constructed using $\Theta(f)$ successive Toffoli gates of size 3, which can in turn be implemented using Hadamard, phase, $\pi/8$ and CNOT gates [NC00]. So any circuit of size c consisting of Hadamard, $\pi/8$ and unbounded Toffoli gates can be transformed into an equivalent circuit with size at most $O(cn)$ consisting of these single-qubit gates and CNOT gates. The rest of the construction follows as before.

Corollary 1. *There is a family of universal circuits $U_{n,c}$ that can simulate quantum circuits of size c on n qubits and consisting of Hadamard, $\pi/8$, and unbounded Toffoli gates such that $\text{SIZE}(U_{n,c}) = O(nc \log(nc))$.*

5 Other results

Circuit encoding. We have been mostly concerned with the actual simulation of a quantum circuit C by the universal circuit U . It is possible, however, to hide some complexity of the simulation in U 's description of C itself. Usually, the description of a classical circuit describes the underlying graph of the circuit and specifies the gates at each vertex. We can similarly describe a quantum circuit by its graph structure. The description is extremely compact with size proportional to the size of the circuit. However, we use a description that is more natural for quantum circuits and especially suitable for simulation. The description stores the grid structure of the circuit; the rows of the grid correspond to the qubits, and the columns correspond to the different layers of the circuit. This description is not unique for any given circuit and its size is $O(nd)$, where n is the number of qubits and d is the depth of the circuit. A graph-based description can be easily converted to this grid-based description in polynomial time.

Depth-universal classical circuits. The techniques of Section 3 can be easily adapted to build depth-universal circuits for a variety of classical (Boolean) circuit classes with unbounded gates, e.g., AC, ACC, and TC circuits. The key reason is that these big gates are all “self-similar” in the sense that fixing some of the inputs can yield a smaller gate of the same type. We will present these results in the full paper.

6 Open Problems

A number of natural, interesting open problems remain.

Fanout gates are used in our construction of a depth-universal circuit family. Is the fanout gate necessary in our construction? We believe it is. In fact, we do not know how to simulate depth- d circuits over $\{H, T, \text{CNOT}\}$ universally in depth $O(d)$ without using fanout gates, even assuming that the circuits being simulated have depth $\Omega(\log n)$. The shallowest universal circuits with bounded-width gates we know of have a $\lg n$ blow-up factor in the depth, just by replacing the fanout gates with log-depth circuits of CNOT gates.

Our results apply to circuits with very specific gate sets. How much can these gate sets be generalized? Are similar results possible for any countable set of gates containing Hadamard, unbounded Toffoli, and fanout gates?

We showed how to construct a universal circuit with a logarithmic blow-up in size. The construction is within a constant factor of the minimum possible size for polynomial-size, bounded-fanin circuits. However for constant-size circuits, we believe the lower bound can be tightened to match the proven upper bound. For unbounded-fanin circuits, we construct a universal circuit with size $O(nc \log nc)$ which is significantly larger than the bounded fanin lower bound of $\Omega(c \log n)$. We think that a better lower bound is possible for the unbounded-fanin case.

Acknowledgments

We thank Michele Mosca and Debbie Leung for insightful discussions. The second author is grateful to Richard Cleve and IQC (Waterloo) and to Harry Buhrman and CWI (Amsterdam) for their hospitality.

References

- [CH85] Stephen A. Cook and H. James Hoover. A depth-universal circuit. *SIAM Journal of Computing*, 14(4):833–839, 1985.
- [FFGHZ06] M. Fang, S. Fenner, F. Green, S. Homer and Y. Zhang. Quantum lower bounds for fanout. *Quantum Information and Computation*, 6(1):046–057, 2006.
- [HS05] P. Høyer and R. Špalek. Quantum circuits with unbounded fan-out. *Theory of Computing*, 1:81–103, 2005.
- [NC97] M. A. Nielsen and I. L. Chuang. Programmable quantum gate arrays. arxiv:quant-ph/9703032v1, 1997.
- [NC00] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.
- [SR07] P. B. M. Sousa and R. V. Ramos. Universal quantum circuit for n -qubit quantum gate: A programmable quantum gate. *Quantum Information and Computation*, 7(3):228–242, 2007.
- [Val76] Leslie G. Valiant. Universal circuits (preliminary report). In *Proceedings of the 8th ACM Symposium on the Theory of Computing*, 196–203, 1976.
- [Yao] A. C.-C. Yao. Quantum circuit complexity. In *Proceedings of the 34th IEEE Symposium on Foundations of Computer Science*, 352–361, 1993.