QUANTUM CIRCUITS FOR INCOMPLETELY SPECIFIED TWO-QUBIT OPERATORS

VIVEK V. SHENDE[a]

Department of Electrical Engineering and Computer Science, The University of Michigan
Ann Arbor, Michigan, 48109-2212, USA

IGOR L. MARKOV[b]

Department of Electrical Engineering and Computer Science, The University of Michigan
Ann Arbor, Michigan, 48109-2212, USA

While the question "how many CNOT gates are needed to simulate an arbitrary two-qubit operator" has been conclusively answered – three are necessary and sufficient – previous work on this topic assumes that one wants to simulate a given unitary operator up to global phase. However, in many practical cases additional degrees of freedom are allowed. For example, if the computation is to be followed by a given projective measurement, many dissimilar operators achieve the same output distributions on all input states. Alternatively, if it is known that the input state is $|0\rangle$, the action of the given operator on all orthogonal states is immaterial. In such cases, we say that the unitary operator is incompletely specified; in this work, we take up the practical challenge of satisfying a given specification with the smallest possible circuit. In particular, we identify cases in which such operators can be implemented using fewer quantum gates than are required for generic completely specified operators.

1 Introduction

Quantum circuits offer a common formalism to describe various quantum-mechanical effects and facilitate a unified framework for simulating such effects on a quantum computer [1]. The framework consists of two steps: (1) for a given unitary evolution, find a quantum circuit that computes it, (2) implement this circuit on a quantum computer. The first step is sometimes called quantum circuit synthesis [2], and is the focus of our work. Given that existing physical implementations are severely limited by the number of qubits, a considerable effort was made recently to synthesize small two-qubit circuits [3, 4, 5, 6, 7]. It has been shown that for such a circuit to implement a typical two-qubit operator, three CNOT gates are needed. However, this result was proven under the assumption that we know nothing about the circuit surrounding the given two-qubit operator. Thus, in the event that we have additional information, say the state of the input qubits or the basis in which the result of the computation is to be measured, the result no longer holds. In fact, we show that if the input state is $|0\rangle$, then three one-qubit gates and one CNOT suffice to simulate an arbitrary two-qubit operator. We also show that if a projective measurement in the computational basis follows the two-qubit operator, then it can be implemented by a circuit with two CNOTs.

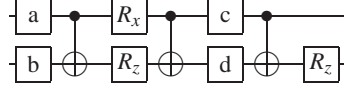[a] vshende@umich.edu
[b] imarkov@umich.edu

Fig. 1. A universal two-qubit circuit with three CNOT gates [5]. It contains seven one-qubit placeholders, which can be translated into 15 placeholders for one-parameter gates.

2   Background

The following family of "spin flip" or "$\sigma_y \otimes \sigma_y$" results are invaluable in the study of two-qubit operators. They are all related in some sense to the fact that a two-qubit pure state $|\phi\rangle$ is separable if and only if $\varepsilon(|\phi\rangle) := \langle\phi^*|\sigma_y^{\otimes 2}|\phi\rangle = 0$. For this reason, $|\varepsilon|^2$ is sometimes used to measure entanglement.

Facts about two-qubit operators.

1. **The Magic Basis [8, 9].** There exist matrices $E \in U(4)$ such that $E^\dagger SO(4)E = SU(2)^{\otimes 2}$. These are characterized by the property $EE^T = \sigma_y^{\otimes 2}$.

2. **The Makhlin Invariants [10]** Let $u, v \in SU(4)$. Then there exist $a, b, c, d \in SU(2)$ such that $(a \otimes b)u(c \otimes d) = v$ if and only if $u^T \sigma_y^{\otimes 2} u \sigma_y^{\otimes 2}$ and $v^T \sigma_y^{\otimes 2} v \sigma_y^{\otimes 2}$ have the same spectrum.

3. **The Canonical Decomposition [11, 12]** Any $u \in SU(4)$ can be written in the following form.

$$u = (a \otimes b)e^{i(I \otimes I + \theta_x \sigma_x \otimes \sigma_x + \theta_y \sigma_y \otimes \sigma_y + \theta_z \sigma_z \otimes \sigma_z)}(c \otimes d)$$

   Above, $a, b, c, d \in SU(2)$ and $\theta_x, \theta_y, \theta_z \in \mathbf{R}$.

These facts can be used to classify two-qubit pure states up to the action of local unitaries, as shown below, and this result is used later in our work. One can also classify mixed states, but the more general result is harder to state [10].

Proposition 1 *Let $|\phi\rangle$ and $|\psi\rangle$ be 2-qubit pure states. Then $|\phi\rangle$ and $|\psi\rangle$ can be interchanged by local unitary operators if and only if $|\varepsilon(|\phi\rangle)| = |\varepsilon(|\psi\rangle)|$.*

Proof. ($\Rightarrow$) Suppose first that ϕ and ψ are interconvertible by local unitaries, that is, there exist $a, b \in U(2)$ such that $(a \otimes b)|\phi\rangle = |\psi\rangle$. One can check that for m a 2×2 matrix, $m^T \sigma_y m = \sigma_y \det m$. Using this to simplify, we have $\langle\psi^*|\sigma_y^{\otimes 2}|\psi\rangle = \langle\phi^*|(a \otimes b)^T \sigma_y^{\otimes 2}(a \otimes b)|\phi\rangle = (\det a)(\det b)\langle\phi^*|\sigma_y|\phi\rangle$. The scalar vanishes upon taking absolute value. ($\Leftarrow$) Conversely, suppose $|\langle\phi^*|\sigma_y \otimes \sigma_y|\phi\rangle| = |\langle\psi^*|\sigma_y \otimes \sigma_y|\psi\rangle|$. By ignoring global phase, we may suppose that in fact $\langle\phi^*|\sigma_y \otimes \sigma_y|\phi\rangle = \langle\psi^*|\sigma_y \otimes \sigma_y|\psi\rangle$. Changing to the Magic Basis transforms the hypothesis into $\langle\phi|\phi\rangle = \langle\psi|\psi\rangle$, and the statement we want to prove into: there exists $p \in SO(4)$ such that $p|\psi\rangle = |\phi\rangle$. So, let $v \in \mathbf{C}^4$ be an arbitrary vector, and $v = v_r + iv_i$ be its decomposition into real and imaginary parts. Then we see that $v^T v = |v_r|^2 - |v_i|^2 + 2iv_r^T v_i$. Since we know v to be a unit vector, $v^T v$ encodes the magnitudes of the real and imaginary parts of v, and the angle between them. From this it is clear that two unit vectors v, w in $\mathbf{C}^4$ can be interchanged by an element of $SO(4)$ if and only if $v^T v = w^T w$, and we have proven our claim $\square$.

Another important consequence of the $\sigma_y \otimes \sigma_y$ theorems is the result that an arbitrary two-qubit operator can be implemented by a circuit containing three CNOTs and some one-qubit gates. It has been proven in various forms [3, 4, 5], of which we need the particular one described in Fig. 1.

It is also known that three CNOT gates are necessary to implement some two-qubit operators, such as a wire swap [3, 4, 5]. To prove this and other lower bound results, one considers *generic circuits*.

These are diagrams with placeholders for unspecified (variable) gates and may also contain specific (constant) gates. Each placeholder corresponds to some subset of possible gates. In this work, all placeholders are for one-qubit gates, and all constant gates are either CNOTs or one-qubit gates; we call such circuits *basic*. We label a placeholder for an unspecified element of $SU(2)$ with a lower-case roman letter, and placeholders for gates of the form $R_x(\alpha)$, $R_y(\beta)$, or $R_z(\gamma)$ by R_x, R_y, R_z respectively. Here, $R_n(\theta) = e^{i\sigma^n\theta/2}$. We refer to basic circuits whose only placeholders represent R_x, R_y, R_z gates as *elementary*. The motivation for restricting to elementary circuits is that each placeholder has one degree of freedom, which makes dimension counting easier and more precise. Moreover, nothing is lost by doing so since any $u \in SU(2)$ can be written in the form $R_k(\alpha)R_l(\beta)R_m(\gamma)$ for any $k, l, m \in \{x, y, z\}, k \neq l, l \neq m$.

We say that an *n*-qubit generic circuit is *universal* if, by specifying appropriate values for the placeholders, one can obtain a circuit simulating arbitrary $u \in U(2^n)$ up to global phase. The dimension of $U(2^n)$ is 4^n; subtracting one for global phase, we see that an elementary circuit on *n* qubits cannot be universal unless there are at least $4^n - 1$ placeholders. Our general strategy for showing that a given incompletely specified circuit is not universal is to convert it into an elementary circuit, eliminate as many placeholders as possible via circuit identities, and then count gates. For example, the following well-known identity is particularly instrumental: the R_x (respectively, R_z) gate can pass through the target (respectively, control) of a CNOT gate.

3 Preparation of Pure States

The three-CNOT lower bound applies when one must find a circuit to simulate a particular given two-qubit operator up to global phase. However, quantum-computational tasks arising in applications are often less completely specified, thus they can be performed by a greater variety of quantum circuits. One such task is state preparation. To prepare the *n*-qubit state $|\phi\rangle$ from $|0\rangle$, we can use any operator $u \in U(2^n)$ with $u|0\rangle = e^{i\theta}|\phi\rangle$. A poor choice of *u* ensures that *u* cannot be implemented with fewer than $O(4^n)$ gates. However, as the dimension of the space of pure states is $2^n - 1$, the lower bound by dimension counting techniques described in Section 2 only indicate that at least $\lceil(2^n - 3n - 1)/4\rceil$ CNOTs are necessary to prepare an arbitrary pure state.[c] We show below that this bound can be matched asymptotically by techniques based on the QR decomposition of matrices.

Proposition 2 *Preparing a generic n-qubit pure state from $|0\rangle$ requires $O(2^n)$ quantum gates.*

Proof. As shown in [13], an arbitrary *n*-qubit unitary operator can be simulated by a circuit containing approximately 8.7×4^n CNOT gates. Their technique is based on the QR decomposition and gives a circuit that builds up a unitary matrix column by column, with each of the 2^n columns built by a subcircuit containing $O(2^n)$ gates. For our present purposes, only the subcircuit responsible for the first column is needed. □.

Other decomposition algorithms find better circuits for arbitrary operators: the best currently known yields about $4^n/2$ CNOTs [14] and is a factor of two away from the lower bound of $\lceil(4^n - 3n - 1)/4\rceil$ given in [5]. However, as these algorithms do not build matrices column by column, they do not yield efficient techniques for state preparation. We note in passing that a significantly larger gap exists between the upper and lower bounds on the number of CNOT gates needed to prepare an arbitrary state, as compared to the corresponding bounds for the problem of simulating an arbitrary

---

[c] For more details on the use of these lower bounds methods, see Section 3 of [5].

unitary operator: in the first case, a factor of thirty, in the second, a factor of two.[d]

We now seek optimality results for the task of state preparation in the case of two qubits. As two-qubit states can be entangled, at least one use of a two-qubit gate is necessary to prepare any entangled state. To characterize two-qubit gates which are also sufficient for this purpose, we use some concepts from algebraic geometry, for whose explication the reader is referred to any introductory textbook, such as [17]. We also give an explicitly constructive proof of this result in the special case of the CNOT gate.

Proposition 3 *Let $G \in SU(4)$. Then an arbitrary pure state $|\psi\rangle$ can be prepared from $|0\rangle$ by a circuit containing one-qubit gates and a single gate G if and only if there exists a state $|\phi\rangle$ such that $\varepsilon(|\phi\rangle) = 0$ and $\varepsilon(G|\phi\rangle) = 1$.*

Proof.   ($\Leftarrow$) Note that $|\varepsilon(|0\rangle)| = 0$ and define $|B\rangle := (|00\rangle + |11\rangle)/\sqrt{2}$. Then $|\varepsilon(|B\rangle)| = 1$. Suppose there exist $a, b, c, d \in U(2)$ such that $(a \otimes b)G(c \otimes d)|0\rangle = |B\rangle$. Recalling from Proposition 1 that one-qubit operators preserve $|\varepsilon|$, we have $|\varepsilon((c \otimes d)|0\rangle)| = 0$, and $|\varepsilon(G(c \otimes d)|0\rangle)| = |\varepsilon(|B\rangle)| = 1$.

($\Rightarrow$) We note that by Proposition 1, it suffices to show that circuits of the form $G(c \otimes d)$ can prepare states with arbitrary $|\varepsilon|$ from $|0\rangle$. Again by Proposition 1, if $|\phi\rangle$ is the state given in the hypothesis, then there exist $a_1, b_1 \in U(2)$ such that $(a_1 \otimes b_1)|0\rangle = |\phi\rangle$. So, $\varepsilon(G(a_1 \otimes b_1)|0\rangle) = 1$. If we show that a state with $|\varepsilon| = 0$ can be prepared, it will follow by continuity of $|\varepsilon|$ that arbitrary states can be prepared as well.

It suffices to show that every two-qubit gate maps some $|\varepsilon| = 0$ state to another. For, if $|\phi\rangle$ is such a state for G, then we may choose $a_0, b_0 \in U(2)$ such that $|\phi\rangle = (a_0 \otimes b_0)|0\rangle$, and see that $|\varepsilon(G(a \otimes b)|0\rangle)| = 0$. Thus it suffices show that $\varepsilon(|\phi\rangle) = 0$ and $\varepsilon(G(|\phi\rangle)) = 0$ have common solutions for all G. Fix G, and fix a basis for the state space. Then, $\varepsilon(|\phi\rangle)$ and $\varepsilon(G|\phi\rangle)$ can be seen to be homogenous polynomials in the 4 coordinates (in fact, they are quadratic forms). In particular, the zeroes of these polynomials do not depend on global phase, so we may speak of their zeroes on the space $\mathbf{CP}^3$ of two-qubit pure states modulo global phase. It is a fact that any two (nonconstant) homogenous polynomials must have common zeroes here [17]. $\square$

As a single CNOT and a Hadamard gate can be used to prepare $(|00\rangle + |11\rangle)/\sqrt{2}$ from $|00\rangle$, the CNOT gate satisfies the hypothesis of the Proposition 3, and therefore a single CNOT suffices to prepare an arbitrary two-qubit pure state from $|0\rangle$. We now give a more explicit construction in this case.

Proposition 4 *A two-qubit pure state $|\phi\rangle$ can be prepared from $|0\rangle$ using the one CNOT gate and three one-qubit gates.*

Proof.   Let C_2^1 be the CNOT gate controlled on the higher qubit and acting on the lower. Let $c = u|0\rangle\langle 0| + v|1\rangle\langle 0| - \overline{v}|0\rangle\langle 1| + \overline{u}|1\rangle\langle 1|$ for some $u, v \in \mathbf{C}$; one can check that $c \in SU(2)$. Let $\phi_i = \langle i|\phi\rangle$. We explicitly compute

$$\varepsilon(C_2^1(I \otimes c)|\phi\rangle) = \phi_0\phi_2(u^2 - v^2) + \phi_1\phi_3(\overline{v}^2 - \overline{u}^2) - (\phi_0\phi_3 + \phi_1\phi_2)(u\overline{v} + v\overline{u})$$

Making the change of variables $z = u^2 - v^2$, $\lambda = (u\overline{v} + v\overline{u})$, we note that $\lambda \in \mathbf{R}$ and $|z|^2 + \lambda^2 = 1$; we want to solve $\phi_0\phi_2 z - \phi_1\phi_3\overline{z} = (\phi_0\phi_3 + \phi_1\phi_2)\lambda$ for z, λ. This is a linear system with two equations and three unknowns; thus we obtain z, λ up to a scalar multiple, and can choose the scalar so that $|z|^2 + \lambda = 1$.

---

[d]Since the first posting of this paper, several preprints have appeared to address this gap. In particular, it has been shown in [14, 15] that $2^{n+1} - 2n - 2$ CNOT gates suffice to prepare an arbitrary n-qubit state from $|0\rangle$. A different technique based on Grover's Search Algorithm also purports to do well in some special cases [16].

Let $|\eta\rangle = C_2^1(I \otimes c)|\phi\rangle$ and verify that $\varepsilon(|\eta\rangle) = 0$. Since $|\eta\rangle$ is separable, write it as $|s\rangle|t\rangle$. This allows one to define a and b so that $(a \otimes b)|0\rangle = |s\rangle|t\rangle$. Finally, we can write $(I \otimes c^\dagger)C_2^1(a \otimes b)|0\rangle = |\phi\rangle$ as desired □.

4 Measurement Don't-Cares

Fewer gates are required for state preparation because images of basis states other than $|0\rangle$ can be arbitrary (in other words, we are using additional information about the input). Similarly, we may be able to save gates if we know in advance how the circuit output will be used. In particular, we now suppose that we know the output is to be measured in some predetermined basis.

Suppose we intend to first simulate an operator u on a yet-unspecified input, then take a projective measurement with respect to some given orthogonal subspace decomposition $(\mathbf{C}^2)^{\otimes n} = \bigoplus E_i$, and we are interested only in having the measured state appear in a given subspace with the appropriate probability. In particular, if v is an operator that preserves each subspace E_i, then we do not care whether we implement u or vu. Conversely, if w is any operator which, upon any input, agrees with u after projective measurement with respect to the given subspace decomposition, then it is clear that wu^{-1} preserves each subspace E_i. If a given circuit simulates some such operator w up to phase, we say that this circuit simulates u up to the measurement don't care associated to the given subspace decomposition.

Mathematically speaking, the problem of state preparation is essentially a special case of a measurement don't care. To prepare the state $|\phi\rangle$ from $|0\rangle$, it is enough to have any operator whose matrix in the computational basis has first column $|\phi\rangle$. On the other hand, suppose we are interested in simulating some given operator u, then taking a projective measurement with respect to two orthogonal subspaces: one spanned by $|0\rangle$ and the other by the rest of the computational basis vectors. Then we may replace u with any operator v such that $\langle 0|u = \langle 0|v$; that is, the matrices of u and v must have the same first row in the computational basis. Thus the problem of state preparation amounts to specifying a single column of a matrix, whereas the aforementioned measurement don't care amounts to specifying a single row. Thus Propositions 2, 3, and 4 carry over to this context.

Proposition 5 *To simulate an arbitrary n-qubit operator up to a projective measurement onto two subspaces, one of which is one dimensional, at least* $\lceil (2^n - 3n - 1)/4 \rceil$ CNOT *gates are necessary, and* $O(2^n)$ CNOT *gates are sufficient. For n = 2, one* CNOT *is necessary and sufficient.*

Suppose now we have a subspace decomposition and an underspecified circuit S which we believe is universal up to the associated measurement don't care – that is, we believe that for any u, appropriate specification of parameters gives a circuit simulating an operator w such that there exists some operator v preserving the subspace decomposition with $vu = w$. Let T be an underspecified circuit that precisely captures the set of operators that fix the subspace decomposition. It is clear that S is universal up to the given measurement don't care if and only if the concatenated circuit ST is universal. Therefore, as we show below, one cannot claim asymptotic savings for this problem in general.

Proposition 6 *To simulate an arbitrary n-qubit operator up to a projective measurement in which each of the subspaces is one-dimensional, at least* $\lceil (4^n - 2^n - 3n)/4 \rceil$ CNOT *gates are required.*

Proof. First, note that an operator v can be right-multiplied by any diagonal operator δ (diagonal in the basis of the measurement) and that the group of diagonal matrices is 2^n-dimensional. Thus, $4^n - 2^n$ parameters must be accounted for. By the proof of Proposition 1 of [5], $(4^n - 2^n - 3n)/4$ CNOT gates are necessary to account for this many parameters □.

Given that the best known circuit synthesis technique for n-qubit circuits is still a factor of two away from the theoretical lower bound of $\lceil (4^n - 3n - 1)/4 \rceil$, it may be difficult to detect a savings of 2^n gates by analyzing specific circuits. Thus we turn to the two-qubit case, where all bounds are known, tight, and small — no more than three CNOT gates are required, and a savings of even one gate would be significant.

On two qubits, there are several different types of measurement possible. We classify them by the subspace dimensions, hence we have "$3 + 1$", "$2 + 2$", "$2 + 1 + 1$", and "$1 + 1 + 1 + 1$" measurements. In what follows, we generally require that each subspace is spanned by computational basis vectors. We refer collectively to the corresponding measurement don't-cares as CB-measurement don't-cares. Additionally, when dealing with $2 + 2$ measurements, we assume that one of the qubits is measured; that is, we do not consider the decomposition $\mathbf{C}^4 = \text{span}(|00\rangle, |11\rangle) \oplus \text{span}(|01\rangle, |10\rangle)$. Indeed, measuring a qubit is a common step in quantum algorithms and communication protocols.

We have already seen in Proposition 5 that one CNOT is necessary and sufficient in the $3 + 1$ case. We now show that at least two CNOTs are needed in the remaining cases

Proposition 7 *Let $\mathbf{C}^4 = \bigoplus E_i$ be a CB-subspace decomposition corresponding to the measurement don't-care M. Suppose no subspace is 3-dimensional, and further that the subspace decomposition is not $\mathbf{C}^4 = \{\text{span}(|00\rangle, |11\rangle) \oplus \text{span}(|01\rangle, |10\rangle).$[e] Then there exist two-qubit operators that cannot be simulated up to M by a circuit with only one CNOT.*

Proof. First, consider subspace decompositions in which neither $\text{span}(|00\rangle, |11\rangle)$, nor $\text{span}(|01\rangle, |10\rangle)$ occur. Remaining cases with 2+1+1 decompositions using one of those subspaces are considered separately below. Suppose an operator is universal up to a $1 + 1 + 1 + 1$ or $2 + 1 + 1$ CB-measurement don't-care satisfying the above condition. Then combining pairs of 1-dimensional subspaces into 2-dimensional subspaces, we see that the same circuit is universal up to a $2 + 2$ CB-measurement don't-care in which one of the qubits is measured. Suppose, without loss of generality, that it is the higher order qubit, hence that the subspaces are $\text{span}(|00\rangle, |01\rangle)$ and $\text{span}(|10\rangle, |11\rangle)$.

We now compose an arbitrary one-CNOT circuit with a circuit for operators preserving the relevant CB-subspaces, as outlined at the beginning of the section (see below-left). Conglomerating adjacent gates, we obtain the circuit below-right.
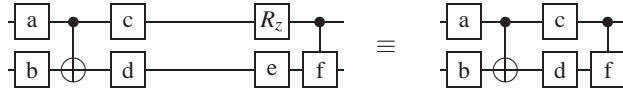

We now convert 3-dimensional place-holders to one-parameter gates, pass R_x and R_z through CNOT where desirable, and conglomerate adjacent gates again.
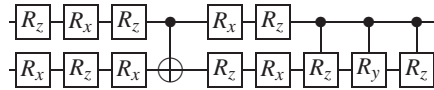
As this circuit has 13 one-parameter gates, the circuit we started with cannot be universal.

The $2 + 1 + 1$ cases where the 2-dimensional subspace is $\text{span}(|00\rangle, |11\rangle)$ or $\text{span}(|01\rangle, |10\rangle)$ can

---

[e] In the $2 + 2$ case where measurement is performed "across qubits", the key question is whether
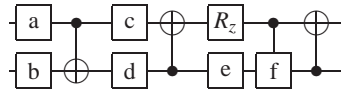
is universal. Unfortunately, we have neither been able to find circuit identities to reduce the number of one-parameter gates below 15, nor to show that this circuit is universal.

be dealt with similarly. We give the circuits preserving these subspace decompositions below; the left circuit corresponds to span($|01\rangle, |10\rangle$) and the right to span($|01\rangle, |10\rangle$).
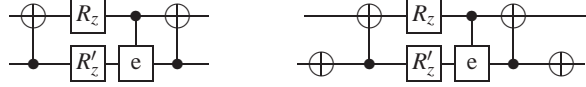

In both cases, the placeholder marked R'_z can be conglomerated with another placeholder, leaving a circuit with 14 one-parameter placeholders □.

It is a natural question whether one might do better with a different gate [7]. At least for the $1+1+1+1$ subspace decomposition, the answer is no.

Proposition 8 *Fix a two-qubit gate G. Some two-qubit operators cannot be simulated, up to the $1+1+1+1$ CB-measurement don't care, by a circuit limited to a single instance of G.*

**Proof.** We compose the circuit in question with a circuit for simulating a diagonal operator.
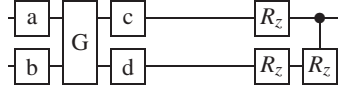
We now merge the R_z gates with the c and d placeholders; there remain 13 parameters — three each in the a, b, c, d placeholders and one in the controlled-R_z gate. This circuit fails to be universal □.

In a different direction, one may ask whether one can do better by measuring in a different basis.

Proposition 9 *Consider the $1+1+1+1$ measurement don't care, M, corresponding to a given fixed basis. Some two-qubit operators cannot be simulated up to M by a circuit with a single* CNOT.

Proof. We concatenate the circuit in question with a placeholder for a diagonal operator.
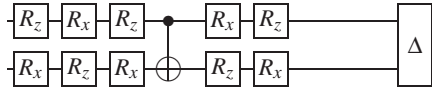
Counting parameters gives 13 (the placeholder for the diagonal operator counts for three.) Thus this circuit cannot be universal □.

Finally, we prove constructively that an arbitrary two-qubit operator can be implemented up to any CB-measurement don't care with a circuit containing two CNOT gates and various one-qubit gates.

Proposition 10 *The 2-qubit circuit below is universal up to any CB-measurement don't-care.*
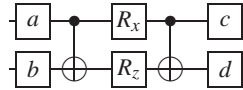
Proof. Consider a measurement with respect to any CB-subspace decomposition. The number and the probabilities of outcomes cannot change if we first measure along the $1+1+1+1$ subspace decomposition. Indeed, the number of outcomes is determined by the number of subspaces in the last measurement, and the probabilities of outcomes for a given pure state by squared norms of projections onto those subspaces. In a CB-subspace decomposition, the squared norm of a projection onto a 2- or 3-dimensional subspace equals, by the Pythagorean theorem, the sum of squared norms of projections onto the computational-basis vectors in that subspace. Therefore, a circuit which is universal up to a $1+1+1+1$ CB-measurement don't-care is universal up to any other CB-measurement don't-care, and it suffices to consider the $1+1+1+1$ case.

Recall that the circuit of Fig. 1 is universal. As adding a reversible constant gate (e.g., CNOT) to the end does not affect universality, the circuit below is universal as well.
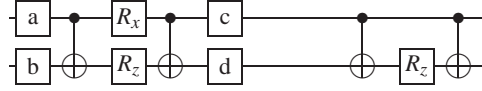
Observe that the right portion of this circuit simulates a diagonal operator, which preserves the subspaces spanned by the computational basis vectors. Thus, by the discussion earlier in the section, the left portion of this circuit is universal up to measurement in the computational basis □.

In applications such as Quantum Key Distribution, one may not know in advance which basis to measure in, but rather that one will choose at random between a given pair of bases for measurement. To save gates in this context, one could maintain two different circuits, one for each type of measurement. While it may seem counterintuitive that building two circuits would save on gates, note that the "circuit" here may consist of classical instructions to initiate a given laser pulse at a given time, thus we may maintain as many as we like in the memory of the classical computer we are using to control the quantum system. At issue is the execution time, which will be smaller when applying either of two smaller circuits (depending on the desired measurement) rather than a common, larger circuit followed by one of two measurements.

An alternative approach to saving gates in such a context is to try and find circuits which simulate the desired operator up to either of the possible measurements. The only fact we used about the computational basis in the proof of Proposition 10 was that operators expressible as $C_2^1(I \otimes R_z(\theta))C_2^1$ are diagonal in the computational basis. Such operators are also diagonal in any basis in which each vector lies in either span($|0\rangle, |3\rangle$) or span($|1\rangle, |2\rangle$). In particular, this includes bases of Bell states.

Proposition 11 *Two* CNOT*s suffice to simulate any two-qubit operator up to any measurement in a not necessarily predisclosed basis in which each vector lies in either span($|0\rangle, |3\rangle$) or span($|1\rangle, |2\rangle$).*

5   Conclusions

Algorithms and lower bounds for quantum circuit synthesis have significantly advanced in the last two years. In particular, several universal two-qubit circuits with optimal gate counts are available [3, 4, 5, 6, 7], and, in the general case of n-qubit circuits, asymptotically optimal gate counts can be realized by matrix-decomposition algorithms [13, 15].

In this context, we recall that quantum algorithms and cryptographic protocols often apply measurements, known in advance, after reversible quantum circuits. This allows a greater variety of circuits to be functionally equivalent, and we prove that useful information about measurement often facilitates finding smaller circuits. Taking into account a known input state also decreases circuit sizes. Both cases can be viewed as circuit synthesis for incompletely specified operators.

Our work has parallels in synthesis of classical irreversible logic circuits, where truth tables are sometimes underspecified, and the synthesis program must complete them so as to allow for smaller circuits. In other words, outputs produced for some input combinations can be arbitrary. Such unspecified behaviors of classical circuits are traditionally called "don't-cares". While covered in undergraduate circuits courses, they remain a worthy subject of research and appear in a variety of circumstances in practice. For example, if a given circuit operates on outputs of another circuit, the latter may not be able to produce certain combinations of bits. While this cannot happen with reversible quantum circuits, we may nonetheless know in advance that the input state will be $|0\rangle$. Indeed, it may happen that the purpose of the circuit all along was to prepare a given state form $|0\rangle$. To this end, we point out that an n-qubit state can be prepared using $O(2^n)$ gates — which is asymptotically optimal — whereas $O(4^n)$ gates are necessary to simulate a generic n-qubit unitary operator. We also show that at most

one maximally entangling gate is necessary and sufficient to prepare a 2-qubit state, and, in particular, that a single `CNOT` suffices. We have also shown that, if the final measurement is known to be in the computational basis, only two `CNOT` gates are necessary.

Acknowledgements

References

1. M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, England, 2000).
2. S. S. Bullock and I. L. Markov, "An Arbitrary Two-Qubit Quantum Computation In Twenty-Three Elementary Gates," *Phys. Rev. A* **68**, 012318 (2003).
3. G. Vidal and C. M. Dawson, "A universal quantum circuit for two-qubit transformations with three CNOT gates," *Phys. Rev. A* **69**, 032315 (2004).
4. F. Vatan and C. Williams, "Optimal Realization of an Arbitrary Two-Qubit Quantum Gate", *Phys. Rev. A* **69**, 010301 (2004).
5. V. V. Shende, I. L. Markov, and S. S. Bullock, "Minimal Universal Two-Qubit Controlled-NOT-based Circuits," *Phys. Rev. A* **69**, 062321 (2004).
6. V. V. Shende, S. S. Bullock, and I. L. Markov, "Recognizing Small-Circuit Structure in Two-Qubit Operators," *Phys. Rev. A* **70**, 012310 (2004).
7. J. Zhang, J. Vala, S. Sastry and K. B. Whaley, "Minimum construction of two-qubit quantum operations," *Phys. Rev. Lett.* **93**, 020502 (2004).
8. C. Bennett et al., "Mixed State Entanglement and Quantum Error Correction," *Phys. Rev. A* **54**, 3824 (1996).
9. S. Hill, K. Wootters, "Entanglement of a Pair of Quantum Bits," *Phys. Rev. Lett.* **78**, 5022 (1997).
10. Yu. Makhlin, "Nonlocal Properties of Two-qubit Gates and Mixed States and Optimization of Quantum Computations," *Quant. Info. Proc.* **1**, 243 (2002).
11. N. Khaneja, R. Brockett and S. J. Glaser, "Time Optimal Control In Spin Systems," *Phys. Rev. Lett.* **63**, 032308 (2001).
12. M. Lewenstein et al., "Characterization of Separable States and Entanglement Witnesses," *Phys. Rev. Lett.* **63**, 044304 (2001).
13. J. Vartiainen et al., "Efficient decomposition of quantum gates," *Phys. Rev. Lett.* **92**, 177902 (2004).
14. V. V. Shende, S. S. Bullock, and I. L. Markov, "A Practical Top-down Approach to Quantum Circuit Synthesis," to appear in *Proc. ACM/IEEE Asia and South Pacific Design Automation Conf.*, Shanghai, January 2005. `quant-ph/0406176`.
15. M. Mottonen et. al., "Transformation of Quantum States Using Uniformly Controlled Rotations," `quant-ph/0407010`.
16. A. N. Soklakov, R. Schack, "Efficient state preparation for a register of quantum bits," `quant-ph/0408045`.
17. I. R. Shafarevich, *Basic Algebraic Geometry 1*, Springer-Verlag, New York, 1994.