Finding Small Two-Qubit Circuits

Vivek V. Shende[a], Igor L. Markov[a], and Stephen S. Bullock[b]

[a]The University of Michigan, Department of Electrical Engineering and Computer Science
[b] National Institute of Standards and Technology, I.T.L.-M.C.S.D.

ABSTRACT

An important result from the mid nineties shows that any unitary evolution may be realized as a sequence of controlled-not and one-qubit gates. This work surveys especially efficient circuits in this library, in the special case of evolutions on two-quantum bits. In particular, we show that to construct an arbitrary two-qubit state from $|00\rangle$, one CNOT gate suffices. To simulate an arbitrary two-qubit operator up to relative phases, two CNOTs suffice. To simulate an arbitrary two-qubit operator up to global phase, three CNOTs suffice. In each case, we construct an explicit circuit and prove optimality in the generic case. We also contribute a procedure to determine the minimal number of CNOT gates necessary to simulate a given two-qubit operator up to global phase. We use this procedure to discuss timing a given Hamiltonian to simulate the CNOT and to determine an optimal circuit for the two-qubit Quantum Fourier Transform. Our constructive proofs amount to circuit synthesis algorithms and have been coded in C++.

1. INTRODUCTION

Recent empirical work on quantum communication, cryptography and computation[1] resulted in a number of experimental systems that can implement two-qubit circuits. Thus, decomposing arbitrary two-qubit operators into fewer gates from a universal library may simplify such physical implementations. While the universality of various gate libraries has been established in the past,[2,3] the minimization of gate counts has only been studied recently. Universal quantum circuits with six, four and three CNOT gates have been found that can simulate an arbitrary two-qubit operator up to global phase.[4–7] It has also been shown that if the CNOT gate is the only two-qubit gate available, then three CNOT gates are required.[6–8] Many of these results rely on the Makhlin invariants[9] or the related *magic basis* and *canonical decomposition*.[10–13] Similar invariants have been investigated elsewhere.[14–16]

Our work improves or broadens each of the above circuit constructions and lower bounds. Moreover, we consider operators that can be simulated using fewer gates than are necessary in the worst case, which has not been done previously. We rely on the Makhlin invariants,[9] and simplify them for mathematical and computational convenience. We have coded the computation of specific gate parameters in several hundred lines of C++ and note that it involves only closed-form algebraic expressions in the matrix elements of the original operator (no matrix logarithms or exponents.) We articulate the degrees of freedom in our algorithm, and our program produces multiple circuits for the same operator. This may be useful with particular implementation technologies where certain gate sequences are more likely to experience errors. Additionally, this paper contributes a lower bound for the number of CNOT gates required to simulate an arbitrary n-qubit operator, which is tighter than the generic bound for arbitrary two-qubit operators.[3,17]

We discuss gate counts for circuits both in terms of elementary and of basic gates. Both types are common in the literature,[3] but basic gates better reflect gate costs in some physical implementations where all one-qubit gates are equally accessible. Yet, when working with ion traps, possible operations are better characterized by elementary gates; moreover, R_z gates are significantly easier to implement than R_x and R_y gates.[18] Our work uncovers another asymmetry, which is of theoretical nature and does not depend on the implementation technology — a subtle complication arises when only CNOT, R_x and R_z gates are available.

Our work shows that basic-gate circuits can be simplified by temporarily decomposing basic gates into elementary gates, so as to apply convenient circuit identities summarized in Table 1. Indeed, our lower bounds and the n-qubit CNOT bound above rely on these circuit identities. Additionally, temporary decompositions into elementary gates may help optimizing pulse sequences in physical implementations.

The remainder of this paper is structured as follows. Section 2 discusses gate libraries, circuit topologies, and lower bounds. Section 3 provides background on using invariants to classify two-qubit operators up to local unitaries. Section 4 classifies operators which can be simulated using only one CNOT, discusses preparing two-qubit states, and explains how to time a given Hamiltonian to simulate a CNOT without using Trotterization.[19] Section 5 classifies operators that can be simulated using two CNOTs, proves optimality of circuits for the wire swap and two-qubit Quantum Fourier Transform, and discusses simulation up to relative phase. Section 6 gives constructive upper bounds to match the lower bounds of Section 2. We summarize our results in Section 7. Necessary background on unentangled states is developed in Appendix A, and subtle complications caused by the lack of the R_y gate are discussed in Appendix B.

2. GATE LIBRARIES & LOWER BOUNDS

We denote the Pauli matrices by $\sigma_x = |1\rangle\langle 0| + |0\rangle\langle 1|$, $\sigma_y = i|1\rangle\langle 0| - i|0\rangle\langle 1|$, and $\sigma_z = |0\rangle\langle 0| - |1\rangle\langle 1|$. We recall the Bloch sphere isomorphism[1] which identifies identifies a unit vector $\vec{n} = (n_x, n_y, n_z)$ with $\sigma_n = n_x\sigma_x + n_y\sigma_y + n_z\sigma_z$. Under this identification, rotation by the angle θ around the vector $\vec{n}$ corresponds to the special unitary operator $R_n(\theta) = e^{-i\sigma_n\theta/2}$. It is from this identification that the decomposition of an arbitrary one-qubit gate $U = \mathrm{e}^{i\Phi}R_z(\theta)R_y(\phi)R_z(\psi)$ arises.[1] Of course, the choice of y, z is arbitrary; one may take any pair of orthogonal vectors in place of $\vec{y}, \vec{z}$.

LEMMA 2.1. *Let $\vec{n}, \vec{m} \in \mathbb{R}^3$, $\vec{n} \perp \vec{m}$, and $U \in SU(2)$. Then one can find θ, ϕ, and ψ such that $U = R_n(\theta)R_m(\phi)R_n(\psi)$.*

In the case of $\vec{n} \perp \vec{m}$, we have $\sigma_n R_m(\theta)\sigma_n = R_m(-\theta)$ and $R_n(\pi/2)R_m(\phi)R_n(-\pi/2) = R_p(\phi)$ for $\vec{p} = \vec{m} \times \vec{n}$. For convenience, we set $S_n = R_n(\pi/2)$; then S_z is the usual S gate, up to phase. In the sequel, we always take m, n out of x, y, z.

We denote by C_a^b the controlled-not (CNOT) gate with control on the a-th qubit and target on the b-th. We recall that R_z gates commute past CNOTs on the control line and R_x gates commute past CNOTs on the target.

In this work we distinguish two types of gate libraries for quantum operators that are universal in the exact sense (compare to approximate synthesis and the Solovay-Kitaev theorem.) The *basic-gate* library[3] contains the CNOT, and all one-qubit gates. *Elementary-gate* libraries also CNOT gate and one-qubit gates, but we additionally require that they contain only finitely many one-parameter subgroups of $SU(2)$. We call these *elementary-gate* libraries, and Lemma 2.1 indicates that if such a library includes two one-parameter subgroups of $SU(2)$ (rotations about around orthogonal axes) then the library is universal. In the literature, it is common to make assertions like: $\dim[SU(2^n)] = 4^n - 1$. Thus if a given gate library contains only gates from one-parameter families and fully-specified gates such as CNOT, at least $4^n - 1$ one-parameter gates are necessary.[3,17] Such dimension-counting arguments lower-bound the number of R_x, R_y, R_z gates required in the worst case.[3]

To formalize dimension-counting arguments, we introduce the concept of *circuit topologies* — underspecified circuits that may have *placeholders* instead of some gates, only with the gate type specified. Before studying a circuit topology, we must fix a gate library and thus restrict the types of fully-specified (constant) gates and placeholders. We say that a fully-specified circuit $\mathcal{C}$ conforms to a circuit topology $\mathcal{T}$ if $\mathcal{C}$ can be obtained from $\mathcal{T}$ by specifying values for the variable gates. All k-qubit gates are to be in $SU(2^k)$, i.e., normalized. For an n-qubit circuit topology $\mathcal{T}$, we define $Q(\mathcal{T}) \subset SU(2^n)$ to be the set of all operators that can be simulated, up to global phase, by circuits conforming to $\mathcal{T}$. We say that $\mathcal{T}$ is universal iff $Q(\mathcal{T}) = SU(2^n)$. In this work, constant gates are CNOTs, and placeholders represent either all one-qubit gates or a given one-parameter subgroup of $SU(2)$. We label one-qubit gate placeholders by $a, b, c, \ldots$, and one-parameter placeholders by R_* with subscripts x, y or z.

We now formalize the intuition that the dimension of $SU(2^n)$ should match the number of one parameter gates.

LEMMA 2.2. *Fix a gate library consisting of constant gates and finitely many one-parameter subgroups. Then almost all n-qubit operators cannot be simulated by a circuit with fewer than $4^n - 1$ gates from the one-parameter subgroups.*

Proof. Fix a circuit topology $\mathcal{T}$ with fewer than $\ell < 4^n - 1$ one-parameter placeholders. Observe that matrix multiplication and tensor product are infinitely differentiable mappings and let $f : \mathbb{R}^\ell \to SU(2^n)$ be the smooth function that evaluates the operator simulated by $\mathcal{T}$ for specific values of parameters in placeholders. Accounting for global phase, $Q(\mathcal{T}) = \bigcup_{\xi^{2^n}=1} \mathrm{Image}(\xi f)$. Sard's theorem[20] demands that $\mathrm{Image}(\xi f)$ be a measure-zero subset of $SU(2^n)$ for dimension reasons, and a finite union of measure-zero sets is measure-zero.

For a given library, there are only countably many circuit topologies. Each captures a measure-zero set of operators, and their union is also a measure-zero set. □

Circuit identities such as $R_n(\theta)R_n(\phi) = R_n(\theta + \phi)$ can be performed at the level of circuit topologies. This identity indicates that two R_n gates may always be combined into one R_n gate, hence anywhere we find two consecutive R_n placeholders in a circuit topology $\mathcal{T}$, we may replace them with a single one without shrinking $Q(\mathcal{T})$. Of course, $Q(\mathcal{T})$ does not

Circuit identities	Descriptions
$C_j^k C_j^k = 1$	CNOT-gate cancellation
$\omega^{j,k}\omega^{j,k} = 1$	SWAP-gate cancellation
$C_j^k C_k^j = \omega^{j,k} C_j^k$	CNOT-gate elimination
$C_j^k \omega^{j,k} = \omega^{j,k} C_k^j$	moving CNOT past SWAP
$V^j \omega^{j,k} = \omega^{j,k} V^k$	moving a 1-qubit gate via SWAP
$\sigma_x^k C_j^k = C_j^k \sigma_x^j \sigma_x^k$	moving σ_x past CNOT control
$C_j^k \sigma_z^j = \sigma_z^j \sigma_z^k C_j^k$	moving σ_z past CNOT target
$C_k^j V^j = V^j C_k^j$ for $V = R_x(\theta), S_x, \sigma_x$	moving x gates past CNOT target
$C_k^j V^k = V^k C_k^j$ for $V = R_z(\theta), S_z, \sigma_z$	moving z gates CNOT control
$R_n(\theta)R_n(\phi) = R_n(\theta+\phi)$	merging R_n gates.
$\vec{n} \perp \vec{m} \implies S_n R_m(\theta) = R_{n\times m}(\theta) S_n$	changing axis of rotation

Table 1. Circuit identities used in out work. Superscripts indicate the wires on which a gate acts.

grow, either, since $R_n(\psi) = R_n(0)R_n(\psi)$. We may similarly conglomerate arbitrary one-qubit gate placeholders, pass R_z (R_x) placeholders through the control (target) of CNOT gates, decompose arbitrary one-qubit gate placeholders into $R_nR_mR_n$ placeholders for $n \perp m$, etc. More circuit identities are listed in Table 1. We use these techniques and Lemma 2.2 to lower bound the number of CNOT gates required to simulate an arbitrary n-qubit operator.

PROPOSITION 2.3. *Fix any gate library containing only the* CNOT *and one-qubit gates. Then almost all n-qubit operators cannot be simulated by a circuit with fewer than* $\lceil \frac{1}{4}(4^n - 3n - 1)\rceil$ CNOT *gates.*

Proof. Enlarging the gate library cannot increase the minimum number of CNOTs in a universal circuit. Thus we may assume the library is the basic-gate library. We show that any n-qubit circuit topology $\mathcal{T}$ with k CNOT gates can always be replaced with an n-qubit circuit topology $\mathcal{T}'$ with gates from the $\{R_z$, R_x, CNOT$\}$ gate library such that $Q(\mathcal{T}) = Q(\mathcal{T}')$ and $\mathcal{T}'$ has k CNOTs and at most $3n+4k$ one-parameter gates. The proposition follows from $3n+4k \geq 4^n - 1$.

We begin by conglomerating neighboring one-qubit gates; this leaves at most $n+2k$ one-qubit gates in the circuit. Now observe that the following three circuit topologies parametrise the same sets of operators: $C_1^2(a\otimes b) = C_1^2(R_xR_zR_x \otimes R_zR_xR_z) = (R_x \otimes R_z)C_1^2(R_zR_x \otimes R_xR_z)$ We use this identity iteratively, starting at the left of the circuit topology. This ensures that each CNOT has exactly four one-parameter gates to its left. (Note that we apply gates in circuits left to right, but read formulae for the same circuits right to left.) The n one-qubit gates at the far right of the circuit can be decomposed into three one-parameter gates apiece. □

COROLLARY 2.4. *Fix an elementary-gate library. Then almost all two-qubit operators cannot be simulated without at least three* CNOT *gates and fifteen one-qubit gates.*

For elementary-gate libraries containing two out of the three subgroups R_x, R_y, R_z, we give explicit universal two-qubit circuit topologies matching this bound in Section 6.

PROPOSITION 2.5. *Using the basic-gate library, almost all two-qubit operators require at least three* CNOT *gates, and at least basic nine gates total.*

Proof. Proposition 2.3 implies that at least three CNOT gates are necessary in general; at least five one-qubit placeholders are required for dimension reasons. The resulting overall lower bound of eight basic gates can be improved further by observing that given any placement of five one-qubit gates around three CNOTs, one can find two one-qubit gates on the same wire, separated only by a CNOT. Using the $R_zR_xR_z$ or $R_xR_zR_x$ decomposition as necessary, the 5 one-qubit gates can be replaced by fifteen one-parameter gates in such a way that the closest parameterized gates arising from the adjacent one-qubit gates can be combined. Thus, if five one-qubit placeholders and three CNOTs suffice, then so do fourteen one-parameter placeholders and three CNOTs, which contradicts dimension-based lower bounds. □

3. BACKGROUND ON TWO-QUBIT ENTANGLEMENT

We show in Appendix A, Lemma 7.2, that any operator $u \in U(2^n)$ that maps unentangled states to unentangled states can be written as $g\omega$ for $g \in U(2)^{\otimes n} = \{\bigotimes_{i=0}^{n-1} a_i : a_i \in U(2)\}$ and ω a wire permutation. The unentangled two-qubit states $|\phi\rangle$ are characterized by the property $\phi^T \sigma_y^{\otimes 2} \phi = 2(\langle 0|\phi\rangle\langle 3|\phi\rangle - \langle 1|\phi\rangle\langle 2|\phi\rangle) = 0$. We use these facts to characterize $SU(2)^{\otimes 2}$.

LEMMA 3.1. *Let $g \in SU(4)$. Then $g \in SU(2)^{\otimes 2}$ iff $g^T \sigma_y^{\otimes 2} g = \sigma_y^{\otimes 2}$.*

Proof. For m a 2×2 matrix, $m^T \sigma_y m = \sigma_y \cdot \det m$. For $a, b \in SU(2)$, $(a\otimes b)^T \sigma_y^{\otimes 2}(a \otimes b) = \sigma_y^{\otimes 2} \cdot \det a \cdot \det b = \sigma_y^{\otimes 2}$. Conversely, suppose $u^T \sigma_y^{\otimes 2} u = \sigma_y^{\otimes 2}$. Then for any unentangled state ϕ, we have $0 = \phi^T \sigma_y^{\otimes 2} \phi = \phi^T u^T \sigma_y^{\otimes 2} u \phi = (u\phi)^T \sigma_y^{\otimes 2} (u\phi)$, hence u maps unentangled states to unentangled states. Thus $u = g\omega$ for some scalar λ, $g \in U(2)^{\otimes 2}$, and ω a wire permutation. We compute $\omega^T g^T \sigma_y^{\otimes 2} g \omega = \sigma_y^{\otimes 2} \cdot \det g$, hence $\det g = 1$. By assumption, $\det u = \det g\omega = \det \omega = 1$, but on two qubits, the only wire permutation with determinant 1 is the identity. □

The following useful facts derive from this observation.

PROPOSITION 3.2. *(The Magic Basis[10, 11]) Let $E = C_1^2(S_z \otimes S_x)$. Then $E^{-1} SU(2)^{\otimes 2} E = SO(4)$.*

Proof. We compute: $EE^T = \sigma_y^{\otimes 2}$. Therefore, $u\sigma_y^{\otimes 2} u^T = \sigma_y^{\otimes 2} \iff uEE^T u^T = EE^T \iff (E^{-1}uE)(E^{-1}uE)^T = I$, and hence $u \in SU(2)^{\otimes 2} \iff E^{-1}uE \in SO(4)$. □

We use $\chi[M]$ to denote the characteristic polynomial of M.

PROPOSITION 3.3. *(The Makhlin Invariants[9]) Let $u, v \in SU(4)$, and set $U = E^{-1}uE$, $V = E^{-1}vE$. Then there exist $a, b, c, d \in SU(2)$ such that $(a\otimes b)v(c \otimes d) = u$ iff $\chi[UU^T] = \chi[VV^T]$.*

Proof. In light of Proposition 3.2, it suffices to show that there exist matrices $x, y \in SO(4)$ such that $xVy = U$.

For P symmetric unitary, $P^{-1} = \overline{P}$, hence $[P + \overline{P}, P - \overline{P}] = 0$. It follows that the real and imaginary parts of P share an orthonormal basis of eigenvectors. As they are moreover real symmetric matrices, we know from the spectral theorem that their eigenvectors can be taken to be real. Thus one can find an $a \in SO(4)$ such that $aUU^T a^T$ is diagonal. By re-ordering (and negating) the columns of a, we can re-order the diagonal elements of $auu^T a^T$ as desired. Thus if $\chi[UU^T] = \chi[VV^T]$, we can find $a, b \in SO(4)$ such that $aUU^T a^T = bVV^T b^T$ by diagonalizing both; then $(V^\dagger b^T au)(V^\dagger b^T au)^T = I$. Let $c = V^\dagger b^T au \in SO(4)$. We have $a^T bVc = U$, as desired. □

We note that since techniques exist to compute the spectral decomposition, one can effectively compute a, b, c, d.

COROLLARY 3.4. *(The Canonical Decomposition[13]) Given any $u \in SU(4)$, there exist $a, b, c, d \in SU(2)$ and $\Delta \in SU(4)$ such that Δ is diagonal in the Magic Basis and $u = (a\otimes b)\Delta(c \otimes d)$.*

To ease computation of the Makhlin invariants, we observe that $(E^{-1}vE)(E^{-1}vE)^T = E^{-1}(v\sigma_y^{\otimes 2} v^T \sigma_y^{\otimes 2})E$.

DEFINITION 3.5. *We define γ on $2^n \times 2^n$ matrices by the formula $u \mapsto u\sigma_y^{\otimes n} u^T \sigma_y^{\otimes n}$.*

PROPOSITION 3.6. *Let $a, b \in SU(2^m), c \in SU(2^n)$, $g, h \in SU(2)^{\otimes n}$, $u, v \in SU(4)$, $G = SU(2)^{\otimes 2}$.*

1. $\gamma(I) = I$

2. $\gamma(ab) = a\gamma(b)\gamma(a^T)^T a^{-1}$

3. $\gamma(a \otimes c) = \gamma(a) \otimes \gamma(c)$

4. $\gamma(g) = I$

5. $\gamma(ag) = \gamma(a)$

6. $\chi[\gamma(gah)] = \chi[\gamma(a)]$.

7. $u \in G \iff \gamma(u) = I$

8. $uG = vG \iff \gamma(u) = \gamma(v)$

9. $GuG = GvG \iff \chi[\gamma(u)] = \chi[\gamma(v)]$

10. There exists $m, n, o, p \in SU(2)$ such that u is a scalar multiple of $(m \otimes n)v(o \otimes p)$ iff $\chi[\gamma(u)] = \chi[\pm\gamma(v)]$.

Proof. (1), (2), and (3) are immediate from the definition. (4) can be checked explicitly for $n = 1$, and then the general case follows from (3). For (5), note first that $g \in SU(2)^{\otimes n} \implies \gamma(g) = I$ by (4). Then expressing $\gamma(ag)$ and $\gamma(a \cdot I)$ using (1) and (2), we see they are equal. For (6), we use (2), (4), and (5) to see that $g, h \in SU(2)^{\otimes n} \implies \gamma(gah) = g^{-1}\gamma(ah)g = g^{-1}\gamma(a)g$ thus $\chi[\gamma(gah)] = \chi[\gamma(a)]$. (7) and (8) are immediate upon changing to the Magic Basis, and (9) follows from Proposition 3.3 as noted above. (10) follows from (9) and the fact that the only scalars in $SU(4)$ are $\pm 1, \pm i$. □

The following general technique can be used to compute $\gamma(u)$. First, determine a circuit, C, simulating the operator u. Given C, it is straightforward to obtain a circuit simulating $\sigma_y^{\otimes 2} u^T \sigma_y^{\otimes 2}$: reverse the order of gates in C, and replace a given gate g by $\sigma_y^{\otimes 2} g^T \sigma_y^{\otimes 2}$. If g is a one-qubit gate, then $\sigma_y^{\otimes 2} g^T \sigma_y^{\otimes 2} = g^\dagger$. For the `CNOT`, we note that $\sigma_y^{\otimes 2} C_1^2 \sigma_y^{\otimes 2} = C_1^2(\sigma_x \otimes \sigma_z)$ and similarly $\sigma_y^{\otimes 2} C_2^1 \sigma_y^{\otimes 2} = C_2^1(\sigma_z \otimes \sigma_x)$. Now, combine the circuits for u and $\sigma_y^{\otimes 2} u^T \sigma_y^{\otimes 2}$ to obtain a circuit simulating $\gamma(u)$.

4. CIRCUITS WITH ONE CNOT GATE

PROPOSITION 4.1. *Fix $u \in SU(4)$. Then u can be simulated (up to global phase)*

(i.) By a circuit with no `CNOT` *gates iff $\chi[\gamma(u)] = (x+1)^4$ or $(x-1)^4$.*

(ii.) By a circuit with one `CNOT` *gate iff $\chi[\gamma(u)] = (x-i)^2(x+i)^2$.*

Proof. By Proposition 3.6.(10), u can be simulated up to global phase by one-qubit operators iff $\chi[\gamma(u)] = \chi[\pm\gamma(I)] = (x\pm 1)^4$. Similarly, u can be simulated up to global phase by one-qubit operators and a single `CNOT` gate iff $\chi[\gamma(u)] = \chi[\pm\gamma(C_1^2\sqrt{i})]$ or $\chi[\gamma(u)] = \chi[\pm\gamma(C_2^1\sqrt{i})]$. The global phase $\sqrt{i}$ appears because 3.6.10 applies to elements of $SU(4)$, whereas $\det C_1^2 = \det C_2^1 = -1$. But $\chi[\pm\gamma(C_1^2\sqrt{i})] = \chi[\pm\gamma(C_2^1\sqrt{i})] = (x-i)^2(x+i)^2$. □

One-`CNOT`-circuits are the minimal solutions to the problem of simulating a matrix of which only one row (or column) is specified. This arises when preparing states: if you want to transform the state $|0\rangle$ into the state $|\psi\rangle$, it suffices to simulate any operator whose first column (in the computational basis) is $|\psi\rangle$.

PROPOSITION 4.2. *For any $u \in SU(4)$ and $i \in \{0,1,2,3\}$, there exists some $u'_i \in SU(4)$ such that u'_i can be simulated with one* `CNOT` *and $u'_i|i\rangle = v|i\rangle$. For generic u, one* `CNOT` *is required to simulate any such u'_i.*

Proof. We consider the case $|i\rangle = 0$, the others are similar.

We claim that the orbits of the action of $SU(2)^{\otimes 2}$ on the unit vectors in $\mathbb{C}^4$ are classified by the quadratic form $v \to \varepsilon(v) = \frac{1}{2}v^T\sigma_y^{\otimes 2}v$. By Proposition 3.2 it suffices to check that the orbits of $SO(4)$ on this set are classified by $v \to v^T v$. Letting $v = v_r + iv_i$, 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. Now it is obvious geometrically that if w is another unit vector such that $w^T w = v^T v$, then there is some transformation in $SO(4)$ taking v to w.

Now let $v = u|0\rangle$ be desired first column of the matrix. We begin by seeking an operator $c \in SU(2)$ such that $w = (c\otimes I)|0\rangle$ has the property that $\langle 0|w\rangle\langle 2|w\rangle - \langle 1|w\rangle\langle 3|w\rangle = \varepsilon(v)$. Clearly $w_1 = w_3 = 0$. By proper choice of c, we may take w_0, w_2 to be any complex numbers whose squared lengths sum to unity. It follows that we can force $w_0w_2 - w_1w_3$ to be any complex number of norm bounded by $1/2$. Since $v_0v_3 - v_1v_2$ must be some such number, we can in fact find c as desired. Now let $w' = C_2^1 w = w_0|0\rangle + w_1|1\rangle + w_3|2\rangle + w_2|3\rangle$; by construction w' has the property that $\varepsilon(w') = \varepsilon(v)$, hence there exist $a, b \in SU(2)$ such that $v = (a\otimes b)w' = (a\otimes b)C_2^1(c\otimes I)|0\rangle$ as desired.

Finally, we show that one `CNOT` is necessary in general. Any operator that can be simulated with no `CNOT`s is a Kronecker product; hence v is unentangled and thus by Lemma 7.1, $v_0v_3 = v_1v_2$, which is in general false. □

We point out the following practical application of Proposition 4.1, which somewhat simplifies an earlier account of Hamiltonian timing.[21] Given a Hamiltonian that can be timed to compute a `CNOT` modulo one-qubit gates, Proposition 4.1 can be used to finding the correct duration. Our example is a perturbation of $\sigma^x \otimes \sigma^x$: $H_{42} = (0.42)I\otimes\sigma^z + \sigma^x\otimes\sigma^x$ Recall that a `CNOT` can be constructed using one-qubit gates and some time-iterate of the Hamiltonian $\sigma^x\otimes\sigma^x$. However, to handle the noise term, existing techniques resort to Trotterization, which implements $exp(A+B)$ by separately turning on A and B for short periods of time. Below we find a simpler, direct implementation of `CNOT` from H_{42}. It is especially interesting in light of concerns about the scalability of Trotterization.[19]

We compute $\gamma(e^{iH_{42}t})$ for uniformly-spaced trial values of t and seek out those values at which the characteristic polynomial nears $\chi[\gamma(x)] = (x^2+1)^2 = x^4+2x^2+1$. Our implementation in `C++` finds $t_{\texttt{CNOT}} = 0.80587$ in twenty seconds on a common workstation. Hence, we can produce a `CNOT` from H_{42} and one-qubit gates without Trotterization. Explicitly, for a, b, c, d as given below, $C_2^1 = (a\otimes b)e^{iH_{42}t_{\texttt{CNOT}}}(c\otimes d)$ with numerical precision of 10^{-6}.

$$a = \frac{1}{2}\begin{pmatrix} 1-i & -1+i \\ 1+i & 1+i \end{pmatrix} \quad b = \begin{pmatrix} -0.21503-0.976607i & 0 \\ 0 & -0.21503+0.976607i \end{pmatrix}$$

$$c = 0.707107\begin{pmatrix} -1 & -1 \\ 1 & -1 \end{pmatrix} \quad d = \begin{pmatrix} 0.152049+0.690566i & 0.690566-0.152049i \\ -0.690566-0.152049i & 0.152049-0.690566i \end{pmatrix}$$

Further numerical experiments suggest that building a `CNOT` is possible whenever 0.42 is replaced by a weight w, $0 \le w \le 1$. However, we have no analytical proof of this. Numerical experiments also suggest the *impossibility* of timing the Hamiltonian $H_{XYZ} = \sigma^x\otimes\sigma^x + \sigma^y\otimes\sigma^y + \sigma^z\otimes\sigma^z$ so as to compute a `CNOT`. In other words, trying values of t in the range $-10 \le t \le 10$ as above produced no candidate durations.

5. CIRCUITS WITH TWO CNOT GATES

PROPOSITION 5.1. *Fix $u \in SU(4)$. The following are equivalent: (i.) u can be simulated up to global phase by a circuit with at most two* `CNOT` *gates (ii.) u can be simulated up to global phase by a circuit with exactly two* `CNOT` *gates (iii.) $\chi[\gamma(u)]$ has all real coefficients. (iv.) $tr[\gamma(u)] \in \mathbb{R}$.*

Proof. (iii $\iff$ iv). For $M \in SU(N)$, $\chi(M) = \prod(x-\lambda_i)$, we have $\prod \lambda_i = 1$. Thus $\chi(u) = \left(\prod \overline{\lambda_i}\right) \prod(x-\lambda_i) = \prod(\overline{\lambda_i}x - 1)$. It follows that the coefficient of x^k is the complex conjugate of the coefficient of x^{N-k}. In particular, for $N = 4$, the coefficient of x^2 is real and the coefficients of x^3, x are $\mathrm{tr}(M)$ and its conjugate. Since the constant term and the x^4 coefficient are 1, we see $\chi(M)$ has all real coefficients iff $\mathrm{tr}(M)$ is real.

(ii $\iff$ iii). Since $\chi[\gamma(C_1^2)] = \chi[\gamma(C_2^1)]$, it follows from Proposition 3.6.(9) that we may flip `CNOT`s by introducing one-qubit gates. Explicitly, $C_1^2 = (H \otimes H) C_2^1 (H \otimes H)$, where H is the Hadamard gate. Thus u can be simulated using two `CNOT` gates iff $(p \otimes q)u(r \otimes s) = C_1^2 (a \otimes b) C_1^2$. We decompose $a = R_x(\alpha)R_z(\delta)R_x(\beta)$ decomposition and $b = R_z(\theta)R_x(\phi)R_z(\psi)$, and pass R_x gates and R_z gates outward through the target and control of the `CNOT` gates. Thus we are left with $(p' \otimes q')u(r' \otimes s') = C_1^2[R_z(\delta) \otimes R_x(\phi)]C_1^2$. By Proposition 3.6, this occurs iff $\chi[\gamma(u)] = \chi[\gamma(C_1^2[R_z(\delta) \otimes R_x(\phi)]C_1^2)]$, and explicit computation yields this to be $(x + e^{i(\delta+\phi)})(x + e^{-i(\delta+\phi)})(x + e^{i(\delta-\phi)})(x + e^{-i(\delta-\phi)})$. The roots come in conjugate pairs, hence the polynomial has real coefficients; conversely if we know that $\chi[\gamma(u)]$ has real coefficients, then the roots come in conjugate pairs; thus $\chi[\gamma(u)]$ is as above for some δ, ϕ.

(i $\iff$ ii). Explicit computation shows that $\mathrm{tr}[\gamma(C_1^2)] = 0$ and $\mathrm{tr}[\gamma(I)] = 4$. It follows that circuits that require zero or one `CNOT` gates can also be simulated using exactly two `CNOT`s. □

COROLLARY 5.2. *The usual implementation of the wire swap ω as $C_1^2 C_2^1 C_1^2$ or $C_2^1 C_1^2 C_2^1$ is optimal.*

Proof. One checks directly that $\mathrm{tr}[\gamma(\omega)] \notin \mathbb{R}$, so ω requires at least three `CNOT` gates. □

COROLLARY 5.3. *Consider the circuit in Figure 1 that implements the two-qubit Quantum Fourier Transform, $\mathcal{F}$. This circuit achieves the smallest possible number (3) of* `CNOT`*s and the smallest possible number (3) of one-qubit gates.*

Proof. Since $\mathrm{tr}[\gamma(\mathcal{F})] \notin \mathbb{R}$, we see that $\mathcal{F}$ requires at least three `CNOT`s. We have a circuit containing three `CNOT`s and three one-qubit gates; suppose that in fact two one-qubit gates would suffice. Suppose one occurred at each end of the circuit; then the three `CNOT` gates in the middle would either cancel to a single `CNOT`, or form a wire swap. The first case is a contradiction since we know $\mathcal{F}$ requires three `CNOT`s, and to rule out the second, it suffices to compute that $\chi[\gamma(\mathcal{F})] \neq \chi[\pm\gamma(\omega)]$. So, at least one edge of the circuit has an exposed `CNOT`. But this implies that one of $\mathcal{F}C_1^2, \mathcal{F}C_2^1, C_1^2\mathcal{F}, C_2^1\mathcal{F}$ can be simulated with two `CNOT`s. Computing $\mathrm{tr}[\gamma]$ for each of these cases shows this to be impossible as well. □

We make more precise the relationship of generic operators to two-`CNOT` operators in the following proposition.

PROPOSITION 5.4. *For any $u \in SU(4)$, one can find $v \in SU(4)$ and an angle ψ such that v can be simulated by a circuit containing two* `CNOT` *gates and $uC_2^1(I \otimes R_z(\psi))C_2^1) = v$.*

Proof. We set $\Delta = C_2^1(I \otimes R_z(\psi))C_2^1$ and compute $\mathrm{tr}[\gamma(u\Delta)]$. By Proposition 3.6, this is $\mathrm{tr}[\gamma(u^T)^T \gamma(\Delta)]$. We compute $\gamma(\Delta) = \Delta^2$, and obtain $\mathrm{tr}[\gamma(u\Delta)] = (t_1 + t_4)e^{-i\psi} + (t_2 + t_3)e^{i\psi}$, where t_1, t_2, t_3, t_4 are the diagonal entries of $\gamma(u^T)^T$. To ensure this number is real, we require $\tan(\psi) = \mathrm{Im}(t_1 + t_2 + t_3 + t_4)/\mathrm{Re}(t_1 + t_2 - t_3 - t_4)$. By Proposition 5.1, $v = u\Delta$ can be simulated using two `CNOT`s. □

In particular, every operator in $SU(4)$ is the product of a diagonal operator and a two-`CNOT` operator. This allows the following application for two-`CNOT` circuits. Suppose we are interested in simulating the operator u *up to relative phase*. This situation occurs when we know in advance that immediately afterwards we will take a measurement with respect to the computational basis, as such measurements kill relative phase. In this event, we cannot distinguish Δu from u for Δ

Figure 1. A minimal circuit for the Quantum Fourier Transform on two qubits The circuit contains 3 one-qubit gates and 3 `CNOT`s, but the one-qubit gates are broken up into elementary gates for specificity. Here, $T_z = R_z(\pi/4)$ is the T gate[1] up to a global phase.

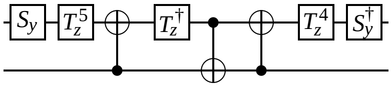

diagonal. Thus, we can replace u by the operator $v = C_2^1(I \otimes R_z(\psi))C_2^1$, which can be simulated using two CNOTs. We can also see that two CNOTs are in general necessary for simulation up to relative phase. If one CNOT were sufficient in general, every two-qubit operator could be simulated by a one-CNOT circuit, followed by a diagonal computation. But an arbitrary circuit of this form can be written as follows.

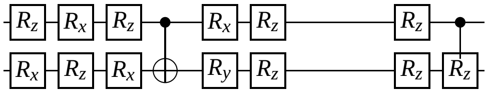

Even without cancelling redundant R_z gates, we know from Lemma 2.2 that this circuit cannot be universal.

PROPOSITION 5.5. *An arbitrary two-qubit operator can be simulated up to relative phase by a circuit containing two* CNOT*s. Two* CNOT*s are necessary for generic operators.*

6. CIRCUITS WITH THREE CNOT GATES

We now construct universal two-qubit circuit topologies. We consider three different gate libraries: each contains the CNOT, and two out of the three one-qubit gates $\{R_x, R_y, R_z\}$. We refer to these as the CXY, CYZ, and CXZ gate libraries.

In view of Lemma 2.1, one might think that there is no significant distinction between these cases. Indeed, conjugation by the Hadamard gate transforms will allow us to move easily between the CXY and CYZ gate libraries. However, we will see that the CXZ gate library is fundamentally different from the other two. Roughly, the reason is that R_x and R_z can be respectively moved past the target and control of the CNOT gate, while no such identity holds for the R_y gate. While the CXY and CYZ libraries each only contain one of $\{R_x, R_z\}$, the CXZ gate library contains both, and consequently has different characteristics. Nonetheless, gate counts will be the same in all cases. We begin with the CYZ case, which has been considered previously.[5]

PROPOSITION 6.1. *15 $\{R_y, R_z\}$ gates and 3* CNOT*s suffice to simulate an arbitrary two-qubit operator up to global phase.*

Proof. Let $v = C_1^2(I \otimes R_y(\alpha))C_2^1(R_z(\delta) \otimes R_y(\beta))C_1^2$. We compute $\gamma(v)$ in terms of α, β, δ. Using the technique described in Section 3, we build a circuit to simulate $\gamma(v)$.

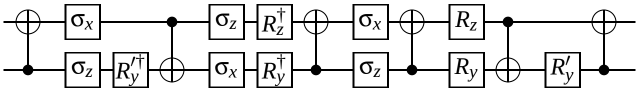

Here, $R_y' = R_y(\alpha)$, $R_y = R_y(\beta)$, and $R_z = R_z(\delta)$. We now use the circuit identities from the 3rd section of Table 1 and $\sigma_i R_j(\theta) = R_j(-\theta)\sigma_i$ to push all the σ_i gates to the left of the circuit, where they cancel (up to a global phase of -1). All one-qubit gates in the wake of their passing become inverted, and we obtain the following circuit.

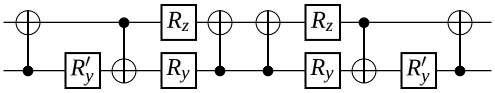

For invertible matrices, $\chi(AB) = \chi(A^{-1}(AB)A) = \chi(BA)$. In view of the fact that we are ultimately interested only in $\chi[\gamma(v)]$ we may move gates from the left of the circuit to the right. Thusly conglomerating R_y' gates and canceling paired CNOT gates, we obtain. $\chi[\gamma(v)] = \chi[C_2^1(R_z(\delta) \otimes R_y(\beta))C_2^1(I \otimes R_y(\alpha))]$. Since $\chi[B] = \chi[A^{-1}BA]$, we conjugate by $I \otimes S_x$. This fixes the CNOT gate and replace R_y gates with R_z: $\chi[\gamma(v)] = \chi[C_2^1(R_z(\delta) \otimes R_z(\beta))C_2^1(I \otimes R_z(\alpha))]$

Now suppose we had set out to simulate an operator u. We ensure that the entries of the diagonal matrix $C_2^1(R_z(\delta) \otimes R_z(\beta))C_2^1(I \otimes R_z(\alpha))$ match the spectrum of $\gamma(u)$ by specifying $\alpha = \frac{x+y}{2}$, $\beta = \frac{x+z}{2}$, and $\delta = \frac{y+z}{2}$ for e^{ix}, e^{iy}, e^{iz} any three eigenvalues of $\gamma(u)$. Then by Proposition 3.6.(9), one can find $a, b, c, d \in SU(2)$ such that $u = (a \otimes b)C_1^2(I \otimes R_y(\alpha))C_2^1(R_z(\delta) \otimes R_y(\beta))C_1^2(c \otimes d)$. Thus, the circuit topology depicted in Figure 2-left is universal. □

PROPOSITION 6.2. *15 $\{R_x, R_y\}$ gates and 3* CNOT*s suffice to simulate an arbitrary two-qubit operator up to global phase.*

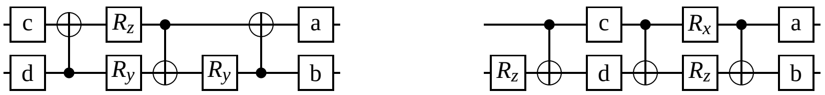

Figure 2. Universal two-qubit circuit with **three** CNOT gates, and a total of **10** basic gates.[3] The left circuit requires **18** gates from $\{\mathtt{CNOT}, R_y, R_z\}$, whereas the right circuit requires **18** gates from $\{\mathtt{CNOT}, R_x, R_z\}$.

Proof. Conjugation by $H^{\otimes n}$ fixes $SU(2^n)$ and R_y. It also flips CNOT gates ($H^{\otimes 2}C_1^2 H^{\otimes 2} = C_2^1$) and swaps R_x with R_z. □

Unfortunately, no such trick transforms CYZ into CXZ. Any such transformation would yield a universal two-qubit circuit topology in the CXZ library in which only three one-parameter gates occur in the middle. We show in Appendix B that no such circuit can be universal and articulate there the implications of this distinction. Nonetheless, we demonstrate here a universal two-qubit circuit topology with gates from the $\{R_x, R_z, \mathtt{CNOT}\}$ gate library that contains 15 one-qubit gates and 3 CNOT gates.

PROPOSITION 6.3. *15 $\{R_x, R_z\}$ gates and 3 CNOTs suffice to simulate an arbitrary two-qubit operator up to global phase.*

Proof. Let u' be the desired operator; set $u = u'C_2^1$. Choose v, ψ relative to u', as in Proposition 5.4. By the proof of Proposition 5.1, one can write $v = (a \otimes b)C_2^1(R_z(\theta) \otimes R_x(\phi))C_2^1(c \otimes d)$ for some $a, b, c, d \in SU(2)$. Setting $u(I \otimes R_z(\psi))C_2^1 = v$ and solving for u gives the overall circuit topology in Figure 2-right. □

Unlike the circuit of Figure 2-left, the circuit in Figure 2-right can be adapted to both other gate libraries. We can replace c by $S_z(S_z^\dagger c)$ and a by $(aS_z)S_z^\dagger$, then use the $S_z, S_z^\dagger$ gates to change the R_x gate into an R_z. A similar trick using S_x can change the bottom R_z gates into R_y; this yields a circuit in the CYZ gate library. As in Proposition 6.2, conjugating by $H \otimes H$ yields a circuit in the CXY gate library.

Given an arbitrary two-qubit operator, individual gates in universal circuits can be computed by interpreting proofs of Propositions 3.3, 3.6, 5.1, 6.1, 6.2, 6.3. By re-ordering eigenvalues in the proof of Proposition 3.3, one may typically produce several different circuits. Similar degrees of freedom have been discussed elsewhere.[5]

7. CONCLUSIONS

Two-qubit circuit synthesis is relevant to on-going physics experiments and can be used in peephole optimization of larger circuits, where small sub-circuits are identified and simplified one at a time. This is particularly relevant to quantum communication, where protocols often transmit one qubit at a time and use encoding/decoding circuits on three qubits.

We have given constructive techniques to find small circuits for two-qubit operators. Specifically, we have shown that one CNOT is necessary and sufficient to prepare an arbitrary two-qubit state from $|0\rangle$, two CNOTs are necessary and sufficient for simulation up to relative phase, and three necessary and sufficient for simulation up to global phase. We have also given formulae to determine how many CNOT gates are required to simulate a given two qubit operator. Lower bounds use the circuit identities summarized in Table 1. In particular, Section 2 showss that n-qubit circuits require $\lceil \frac{1}{4}(4^n - 3n - 1) \rceil$ CNOT gates in the worst case.

Acknowledgments and disclaimers. This work is funded by the DARPA QuIST program and an NSF grant. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing official policies or endorsements of employers and funding agencies.

Appendix A

In this Appendix, we develop some basic results about unentangled states and operators that preserve them. Below, we use & and | to indicate “bitwise OR” and “bitwise AND”. (For example, $0011\&0101 = 0001$, and $0011|0101 = 0111$.)

LEMMA 7.1. *The state vector $|\psi\rangle = \sum_i \psi_i |i\rangle$ is unentangled iff $\psi_i\psi_j - \psi_k\psi_l = 0$ whenever $i\&j = k\&l$ and $i|j = k|l$.*

Proof. This is a special case of the Segre construction.[22] We sketch a proof, for completeness.

($\Rightarrow$) Assume $|\psi\rangle$ is unentangled. We note that to prove our claim, it suffices to show that $\psi_i\psi_j$ can be described solely in terms of $i|j$ and $i\&j$. So write $|\psi\rangle = (a_{n-1}|0\rangle + b_{n-1}|1\rangle)(a_{n-2}|0\rangle + b_{n-2}|1\rangle)\ldots(a_1|0\rangle + b_1)$. Expanding, we see that $\psi_i\psi_j$ is a product of the a_i and b_j with two elements of subscript k appearing for each k. Specifically, if $i\&j$ has a 1 as its k-th binary digit, then b_k appears twice; if $i|j$ has a 0 as its k-th binary digit, then a_k appears twice, and otherwise a_k and b_k both appear.

($\Leftarrow$) Conversely, suppose $|\psi\rangle$ satisfies the condition given. If $|\psi\rangle = 0$, then it is clearly unentangled; if not then one of the ψ_i is nonzero. Without loss of generality, suppose it is ψ_0. We leave it to the reader to verify that in this case, $|\psi\rangle = \psi_0(|0\rangle + \frac{\psi_{2^{n-1}}}{\psi_0}|1\rangle)(|0\rangle + \frac{\psi_{2^{n-2}}}{\psi_0}|1\rangle)\ldots(|0\rangle + \frac{\psi_1}{\psi_0}|1\rangle)$ and hence is unentangled. □

LEMMA 7.2. *Let $u \in U(2^n)$ map unentangled states to unentangled states. Then there exist $g, g' \in U(2)^{\otimes n}$ and ω a permutation of wires such that $u = g\omega = \omega g'$.*

Proof. For $0 \leq i < 2^n$, define $X_i = \sigma_{i_{n-1}} \otimes \sigma_{i_{n-2}} \otimes \ldots \otimes \sigma_{i_0}$, where i_k is the k-th binary digit of i and as usual $\sigma_0 = I$, $\sigma_1 = \sigma_x$. Then $X_i|0\rangle = |i\rangle$. We need the following three subclaims.

(1.) *If $|\psi(\lambda)\rangle = \lambda|i\rangle + |\psi\rangle$ for some given $0 \leq i < 2^n$ and all scalars λ, then $|\psi\rangle = \eta|i'\rangle$ for some scalar η and $0 \leq i' < 2^n$ differing from i by at most one binary digit.* Applying X_i to both sides, we may assume $i = 0$. Now let k be any integer with more than one 1 in its binary expansion, then k can be written as $l|m$ with $l\&m = 0$. Then, from above, $\psi(\lambda)_0\psi(\lambda)_k = \psi(\lambda)_l\psi(\lambda)_m$. Expanding, this says that $(\psi_0 + \lambda)\psi_k = \psi_l\psi_m$. This can only hold for all λ if both sides are zero; hence $\psi_k = 0$. Similarly, for $2^r \neq 2^s$ any integers which each have one 1 in their binary expansion, note $2^r\&2^s = 0$ and hence $(\psi_0 + \lambda)\psi_{2^r|2^s} = \psi_{2^r}\psi_{2^s}$, and at least one of ψ_{2^r}, ψ_{2^s} must be zero.

(2.) *Let $u \in U(2^n)$ map unentangled states to unentangled states, and fix $|i_0\rangle$. Then, up to scalar multiplication, u permutes the $|i\rangle$ which differ from i_0 in one binary digit.* Replacing u with X_iuX_i, we may assume $i_0 = 0$. Now let 2^t be an integer with exactly one 1 in its binary expansion. Then $\lambda|0\rangle + |2^t\rangle$ is unentangled for any t, hence so is $u(\lambda|0\rangle + |2^t\rangle$ by assumption. Expanding out, we have $\lambda|0\rangle + u|2^t\rangle$ unentangled for all λ, hence by (1), $u|2^t\rangle = \eta_t|2^{t'}\rangle$ for some scalar η_t and some t'.

(3.) *Let $u \in U(2^n)$ map unentangled states to unentangled states, fix $|i_0\rangle$, and fix $|i\rangle$ for all i differing from i_0 by exactly one binary digit. Then u fixes $|j\rangle$ for j differing from i_0 by two binary digits.* Replacing u by $X_{i_0}uX_{i_0}$, we assume $i_0 = 0$. So consider $u|2^s + 2^t\rangle$, $s \neq t$. By (2), this must be a scalar multiple of $|i\rangle$ for i differing by one binary digit from 2^s and by one binary digit from 2^t. There are two such numbers, $|0\rangle$ and $|2^s + 2^t\rangle$; but since $u \in U(2^n)$ and $u|0\rangle = |0\rangle$, we cannot have $u|2^s + 2^t\rangle = |0\rangle$. Finally, by Lemma 7.1, the state $|\psi\rangle = |0\rangle + |2^s\rangle + |2^t\rangle + |2^s + 2^t\rangle$ is unentangled, whereas the state $u|\psi\rangle = |0\rangle + |2^s\rangle + |2^t\rangle + \eta|2^s + 2^t\rangle$ is entangled unless $\eta = 1$.

Finally, we prove the lemma. Since $|0\rangle$ is unentangled, so is $u|0\rangle$, hence there exists $k \in U(2)^{\otimes n}$ such that $ku|0\rangle = |0\rangle$. By (2.), $ku|2^t\rangle = \eta_t|2^s\rangle$; observe that the scalars must have norm 1 and let $h = (|0\rangle\langle 0| + \eta_{n-1}|1\rangle\langle 1|)(|0\rangle\langle 0| + \eta_{n-2}|1\rangle\langle 1|)\ldots(|0\rangle\langle 0| + \eta_0|1\rangle\langle 1|)$. Then hku permutes the $|2^s\rangle$, and we may choose an appropriate wire permutation ω so that ωhku fixes the $|2^s\rangle$. Finally, we induct: suppose all the $|i\rangle$, where $0 \leq i < 2^n$ has r or fewer 1s in its binary expansion, are fixed by ωhku. Then every $|i_0\rangle$ with i_0 having $r-1$ digits 1 in its binary expansion satisfies the condition of (3), and every j with $r+1$ digits 1 in its binary expansion differs from some i_0 by two binary digits.

Thus ωhku is the identity, and $u = \omega^{-1}(hk)^{-1}$. Clearly $(hk)^{-1} \in U(2)^{\otimes n}$ and ω^{-1} is a wire permutation. Finally, we observe that for any wire permutation ω and any $g \in U(2)^{\otimes n}$, $\omega g = \omega g \omega^{-1}\omega$, and $\omega g \omega^{-1} \in U(2)^{\otimes n}$. □

Appendix B

In this Appendix, we show that no 18-gate universal two-qubit circuit from the CXZ gate library can have 12 one-qubit gates at the edges, as does the circuit in Figure 2-left. This result has consequences for peephole optimization of n-qubit circuits, where decompositions like that in Figure 2-left are preferrable over that in Figure 2-right. For example, consider a three-qubit circuit consisting of two two-qubit blocks on lines (i) one and two, (ii) two and three. If both blocks are decomposed as in Figure 2-left, then the b gate from the first block and the c gate from the second block merge into one gate on line two. However, no such reduction would happen if the decomposition from Figure 2-right is used.

We now formalize the above claims. The proof of Proposition 6.1 contains a universal generic circuit with three CNOT gates and 15 R_y or R_z gates with the property that all but three of the one-qubit gates appear either before the first or after the last CNOT gate. This is minimal.

PROPOSITION 7.3. *Fix an elementary-gate library. There exist unitary operators $U \in SU(4)$ that cannot be simulated by any two-qubit circuit in which all but two of the one-qubit gates appear either before the first or after the last* CNOT *gate.*

Proof. There are four places where the one-parameter gates can appear: at the left or right of the first or second line. If more than three gates appear in one such place, conglomerate them into a single one-qubit gate, and decompose the result into three one-parameter gates via Lemma 2.1. By this method, any two-qubit circuit can be transformed into an equivalent circuit with at most 12 one-parameter gates on its sides. By Corollary 2.4, there exist operators that cannot be simulated without 15 one-parameter gates; the remaining three must go in the middle of the circuit. $\square$

We have seen that for the CYZ and the CXY gate libraries, this lower bound is tight. We will show that this is not the case for the CXZ gate library. Before beginning the proof, we make several observations about the CXZ gate library.

Note that conjugating a circuit identity by $H \otimes H$ exchanges R_x and R_z gates, and flips CNOTs. Two other ways to produce new identities from old are: swapping wires, and inverting the circuit – reversing the order of gates & replacing each with its inverse. For example, one may obtain one of the commutativity rules below from the other by conjugating by $H \otimes H$ and then swapping wires.

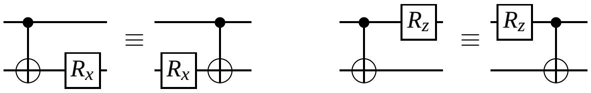

When one CNOT gate occurs immediately after another in a circuit, we say that they are *adjacent*. When such pairs of CNOTs share control lines, they cancel out, and otherwise may still lead to reductions as discussed below. We will be interested in circuits which do not allow such simplifications. To this end, recall that R_x gates commute past the target of a CNOT, and R_z gates commute past the control. Moreover, we have the following circuit identity: $C_2^1(R_x(\alpha) \otimes R_z(\beta))C_2^1 = C_1^2(R_z(\beta) \otimes R_x(\alpha))C_1^2$. We say that a given collection of one-qubit gates *effectively separates* a chain of CNOTs iff there is no way of applying the aforementioned transformation rules to force two CNOT gates to be adjacent. For example, there is no way to effectively separate two CNOTs of opposite orientation by a single R_x or R_z gate. This is illustrated below.

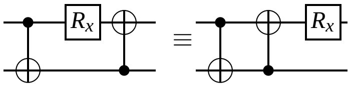

On the other hand, two CNOT gates of the same orientation can be effectively separated by a single R_x or R_z gate, as shown below. Up to swapping wires, these are the only ways to effectively separate two CNOTs with a single R_x or R_z.

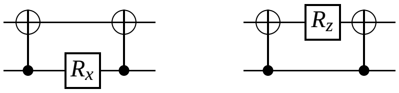

PROPOSITION 7.4. *At least four gates from $\{R_x, R_z\}$ are necessary to effectively separate four or more* CNOT *gates.*

Proof. Clearly it suffices to check this in the case of exactly four CNOTs. If three R_x, R_z gates sufficed, then one would have to go between each pair of CNOT gates. Suppose all the CNOT gates have the same orientation, say with control on the bottom wire. Then the first pair must look like one of the pairs above. In either case, we may use the identity $C_2^1(R_x(\alpha) \otimes R_z(\beta))C_2^1 = C_1^2(R_z(\beta) \otimes R_x(\alpha))C_1^2$ to flip these CNOT gates, thus ensuring that there is a consecutive pair of CNOT gates with opposite orientations. As remarked above, there is no way to effectively separate these using the single one-qubit gate allotted them. $\square$

Denote by ω^{ij} the SWAP gate which exchanges the i-th and j-th qubits. It can be simulated using CNOTs as $C_i^jC_j^iC_i^j = \omega^{ij} = C_j^iC_i^jC_j^i$. SWAP gates can be pushed through an elementary-gate circuit without introducing new gates. So, consider a two-qubit circuit in which adjacent CNOT gates appear. If they have the same orientation (e.g., $C_1^2C_1^2$ or $C_2^1C_2^1$), then they cancel out and can be removed from the circuit. Otherwise, use the identity $C_1^2C_2^1 = C_2^1\omega^{12}$ or $C_2^1C_1^2 = C_1^2\omega^{12}$ and push the SWAP to the end of the circuit. We apply this technique at the level of circuit topologies and observe that since $Q(\mathcal{T}\omega^{12})$ is measure-zero (or universal) iff $Q(\mathcal{T})$ is. By the above discussion, we can always reduce to an effectively separated circuit before checking these properties.

PROPOSITION 7.5. *Almost all unitary operators $U \in SU(4)$ cannot be simulated by any two-qubit circuit with CXZ gates in which all but three of the R_x, R_z gates appear either before the first or after the last* `CNOT`.

Proof. We show that any circuit topology of the form above can only simulate a measure-zero subset of $SU(4)$; the result then follows from the fact that a countable union of measure-zero sets is measure-zero.

The assumption amounts to the fact that only three gates are available to effectively separate the `CNOT` gates. By Proposition 7.4 and the discussion immediately following it, we need only consider circuit topologies with no more than three `CNOT`s. On the other hand, we know from Proposition 2.5 that any two-qubit circuit topology with fewer than three `CNOT` gates can simulate only a measure-zero subset of $SU(4)$. Thus it suffices to consider circuit topologies with exactly three `CNOT` gates. Moreover, we can require that they be effectively separated, since otherwise we could reduce to a two-`CNOT` circuit.

Three `CNOT`s partition a minimal two-qubit circuit in four regions. We are particularly interested in the two regions limited by `CNOTs` on both sides because single-qubit gates in those regions must effectively separate the `CNOT`s. To this end, we consider two pairs of `CNOT`s (the central `CNOT` is in both pairs), and distinguish these three cases: (1) both pairs of `CNOT`s consist of gates of the same orientation, (2) both consist of gates of opposite orientations, or (3) one pair has gates of the opposite orientations and the other pair has gates of the same orientation. In the second case, the `CNOT` gates cannot be effectively separated, since each pair of gates with opposite orientations requires two one-parameter gates to be effectively separated, and only three R_x, R_z gates are available. In the third case, two `CNOT`s with opposite orientations must be separated by two one-parameter gates, leaving only one R_x or R_z to separate the pair with the same orientation. Thus, the pair with the same orientation may be flipped, reducing to Case 1, as shown below.

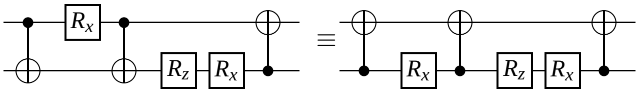

Finally, consider the case in which all three `CNOT` gates have the same orientation. Each pair of consecutive `CNOT`s must have at least one R_x or R_z between them, to be effectively separated. Thus one of the pairs has a single R_x or R_z between its members, and the other has two one-qubit gates. We refer to these as the 1-pair and the 2-pair, respectively.

Suppose that the one-qubit gates separating the 2-pair of `CNOT`s occur on different lines. If either one-qubit can commute past the `CNOT`s of the 2-pair, then it can move to the edge of the circuit; in this case Proposition 7.3 implies that the circuit topology we are looking at can only simulate a measure-zero subset of $SU(4)$ (one can show that two R_x, R_z gates cannot effectively separate three `CNOT`s.) Otherwise, we use the identity $C_2^1(R_x(\alpha) \otimes R_z(\beta))C_2^1 = C_1^2(R_z(\beta) \otimes R_x(\alpha))C_1^2$ to flip the 2-pair, and thus 1-pair now have opposite orientations. As there is only one one-qubit gate between them, this pair is not effectively separated. For example:

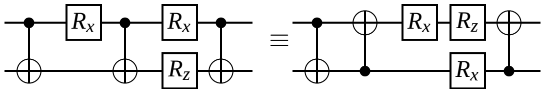

We are left with the possibility that all the `CNOT` gates have the same orientation and that the 2-pair's one-qubit gates appear on the same line. Both R_z, R_x must occur, or else we could combine them and apply Proposition 7.3 to show that such a circuit topology can only simulate a measure-zero subset of $SU(2^n)$. Now, if R_xR_z appears between two `CNOT` gates of the same orientation, then either the R_x or the R_z can commute past one of them. If the outermost gate can commute, Proposition 7.3 again implies that the circuit topology simulates only a measure-zero subset of $SU(2^n)$. Thus the inner gate can commute with the 1-pair. We have now interchanged the roles of the 1-pair and the 2-pair, thus by the previous paragraph, the gate which originally separated the 1-pair must be on the same line as the commuting gate. Consequently, all one-qubit gates are on the same line. The following circuit is the only possibility, up to inverting, swapping wires, and conjugating by $H \otimes H$.

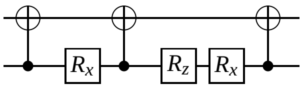

Finally, we add the four one-qubit gates on the sides, decompose each into $R_xR_zR_x$ via Lemma 2.1, and observe that an R_x gate can commute across the top and be absorbed on the other side. This leaves 14 one-parameter gates, and by Lemma 2.2, such a circuit topology simulates only a measure-zero subset of $SU(4)$. □

REFERENCES

1. M. Nielsen and I. Chuang, *Quantum Computation and Quantum Information*, Cambridge Univ. Press, 2000.
2. D. P. DiVincenzo, "Two-bit gates are universal for quantum computation," *PRA* **51**, p. 1015, 1995, cond-mat/94.
3. A. Barenco et al., "Elementary Gates For Quantum Computation," *PRA* **52**, p. 3457, 1995, quant-ph/9503016.
4. J. Zhang et al., "Exact two-qubit universal quantum circuit," *PRL* **91**, p. 027903, 2003, quant-ph/0212109.
5. S. S. Bullock and I. L. Markov, "An Elementary Two-Qubit Quantum Computation In Twenty-Three Elementary Gates," *PRA* **68**, p. 012318, 2003, quant-ph/0211002.
6. F. Vatan and C. Williams, "Optimal Realization of an Arbitrary Two-Qubit Quantum Gate", quant-ph/0308006.
7. G.Vidal and C. M. Dawson, "A Universal Quantum Circuit For Two-qubit Transformations With Three CNOT gates," quant-ph/0307177.
8. J. Zhang et al., "Optimal quantum circuit synthesis from Controlled-U gates", quant-ph/0308167.
9. Yu. Makhlin, "Nonlocal Properties of Two-qubit Gates and Mixed States and Optimization of Quantum Computations," *Quant. Info. Proc.* **1**, p. 243, 2002, quant-ph/0002045.
10. C. Bennett et al., "Mixed State Entanglement and Quantum Error Correction," *PRA* **54**, p. 3824, 1996, quant-ph/9604024.
11. S. Hill, K. Wooters, "Entanglement of a Pair of Quantum Bits," *PRL* **78**, p. 5022, 1997, quant-ph/9703041.
12. N. Khaneja, R. Brockett and S. J. Glaser, "Time Optimal Control In Spin Systems," 2001 *PRA* **63**, 032308-032320, 2001, quant-ph/0006114.
13. M. Lewenstein et al., "Characterization of Separable States and Entanglement Witnesses," *PRA* **63**, p. 044304, 2001, quant-ph/0011050.
14. E. Rains, "Polynomial Invariants of Quantum Codes," quant-ph/9704042.
15. M. Grassl et al., "Computing Local Invariants of Qubit Systems," *PRA* **58**, p. 1833, 1998, quant-ph/9712040.
16. S. Bullock and G. Brennen, "Canonical Decompositions of *n*-qubit Quantum Computations and Concurrence," 2003, quant-ph/0309104.
17. E. Knill, "Approximation by Quantum Circuits," 1995, quant-ph/9508006.
18. D. Wineland et. al., "Experimental Issues in Coherent Quantum Manipulation of Trapped Atomic Ions," *NIST Journal of Research* **103**, p. 259, 1998.
19. A. M. Childs et. al., "Lower bounds on the complexity of simulating quantum gates," *PRA* **68**, p. 052311, 2003.
20. V. Guillemin and A. Pollack, *Differential Topology*, Prentice-Hall, Inc., Englewoods Cliffs New Jersey, 1974.
21. J. Zhang et al., "A geometric theory of non-local two-qubit operations," *PRA* **67**, 042313, 2003, quant-ph/0209120.
22. I. Shafarevich, *Basic Algebraic Geometry 1*, Springer-Verlag, New York, 1994.