Minimal universal two-qubit controlled-NOT-based circuits

Vivek V. Shende,[1,*] Igor L. Markov,[2,†] and Stephen S. Bullock[3,‡]
[1]*Department of Mathematics, The University of Michigan, Ann Arbor, Michigan 48109-1109, USA*
[2]*Department of Electrical Engineering and Computer Science, The University of Michigan, Ann Arbor, Michigan 48109-2122, USA*
[3]*Mathematical and Computational Sciences Division, National Institute of Standards and Technology, Gaithersburg, Maryland 20899, USA*
(Received 23 December 2003; revised manuscript received 10 February 2004; published 30 June 2004)

We give quantum circuits that simulate an arbitrary two-qubit unitary operator up to a global phase. For several quantum gate libraries we prove that gate counts are optimal in the worst and average cases. Our lower and upper bounds compare favorably to previously published results. Temporary storage is not used because it tends to be expensive in physical implementations. For each gate library, the best gate counts can be achieved by a single universal circuit. To compute the gate parameters in universal circuits, we use only closed-form algebraic expressions, and in particular do not rely on matrix exponentials. Our algorithm has been coded in C++.

DOI: 10.1103/PhysRevA.69.062321 PACS number(s): 03.67.Lx, 03.65.Fd, 03.65.Ud

I. 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 been studied only recently. Universal quantum circuits with six, four, and three (controlled-NOT) (CNOT) gates have been found that can simulate an arbitrary two-qubit operator up to a 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 previously [14,15] and more recently in [16].

Our work improves or broadens each of the above circuit constructions and lower bounds, as summarized in Table I. We rely on the Makhlin invariants [9] and simplify them for mathematical and computational convenience—our version facilitates circuit synthesis algorithms. 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].

The two lines in Table I give gate counts for circuits consisting of elementary and basic gates, respectively. Both types were introduced in [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, 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 II.

Indeed, all lower bounds in Table I and the n-qubit CNOT bound above rely on these circuit identities. Additionally, temporary decompositions into elementary gates may help in optimizing pulse sequences in physical implementations.

The remainder of this paper is structured as follows. Section II discusses gate libraries and circuit topologies. Section III derives the lower bounds of Table I. Section IV classifies two-qubit operators up to local unitaries. Section V develops some technical lemmas, and Sec. VI constructs small circuits that match upper bounds in Table I. Subtle complications caused by the lack of the R_y gate are discussed in the Appendix and in Sec. VII.

II. GATE LIBRARIES AND CIRCUIT TOPOLOGIES

We recall that the Bloch sphere isomorphism [1] 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 $U=e^{i\Phi}R_z(\theta)R_y(\phi)R_z(\psi)$ arises [1]. Of course, the choice of y, z is arbitrary; one may take any pair of orthogonal vectors in place of $\vec{y}$, $\vec{z}$.

*Electronic address: vshende@umich.edu
†Electronic address: imarkov@umich.edu
‡Electronic address: stephen.bullock@nist.gov

TABLE I. Constructive upper bounds on gate counts for generic circuits using several gate libraries. Each bound given for controlled-not (CNOT) gates is compatible with the corresponding overall bound. These bounds are tighter than those from [4,5] in all relevant cases.

Gate libraries	Lower and upper bounds CNOT	Overall	CNOT	Overall
{CNOT, any two or three of $\{R_x, R_y, R_z\}$}	3	18	3	18
{CNOT, arbitrary one-qubit gates}	3	9	3	10

Lemma 1. Let $\vec{n}$, $\vec{m} \in \mathbb{R}^3$, $\vec{n} \perp \vec{m}$, and $U \in \mathrm{SU}(2)$. 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 a phase. In the following, we always take m, n out of x, y, z.

We denote by C_a^b the controlled-NOT gate with control on the ath qubit and target on the bth. We recall that R_z gates commute past CNOTs on the control line and R_x gates commute past CNOTs on the target. Finally, for mathematical convenience, we multiply the CNOT gate by a global phase ξ such that $\xi^4 = -1$, to represent it as an element of SU (4).

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 contain the 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 1 indicates that if such a library includes two one-parameter subgroups of SU(2) (rotations around orthogonal axes) then the library is universal. In the literature, it is common to make assertions such as $\dim [\mathrm{SU}(2^n)] = 4^n - 1$. 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 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 $\mathrm{SU}(2^k)$, i.e., normalized. For an n-qubit circuit topology $\mathcal{T}$ we define $Q(\mathcal{T}) \subset \mathrm{SU}(2^n)$ to be the set of all operators that can be simulated, up to a global phase, by circuits conforming to $\mathcal{T}$. We say that $\mathcal{T}$ is universal if and only if $Q(\mathcal{T}) = \mathrm{SU}(2^n)$. 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 subscript x, y, or z.

We also allow for explicit relations between placeholders. For example, circuits conforming to the one-qubit circuit topology $aba^\dagger$ must contain three one-qubit gates and the first and last must be inverse to each other.

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 grow either, since $R_n(\psi) = R_n(0) R_n(\psi)$. We may similarly conglomerate arbitrary one-qubit gate placeholders, pass R_z (re-

TABLE II. Circuit identities used in our work. V^j represents an arbitrary one-quit operator on wire j.

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
$C_k^j R_k^j(\theta) = R_x^j(\theta) C_k^j, C_k^j S_x^j = S_x^j C_k^j$	Moving R_x, S_x via CNOT target
$C_k^j R_z^k(\theta) = R_z^k(\theta) C_k^j, C_k^j S_z^k = S_z^k C_k^j$	Moving R_z, S_z via CNOT control
$\sigma_x^k C_j^k = C_j^k \sigma_x^j \sigma_x^k$	Moving σ_x via CNOT control
$C_j^k \sigma_z^j = \sigma_z^j \sigma_z^k C_j^k$	Moving σ_z via CNOT target
$C_j^k \omega^{j,k} = \omega^{j,k} C_k^j$	Moving CNOT via SWAP
$V^j \omega^{j,k} = \omega^{j,k} V^k$	Moving a one-qubit gate via SWAP
$R_n(\theta) R_n(\phi) = R_n(\theta + \phi)$	Merging R_n gates
$\vec{n} \perp \vec{m} \Rightarrow S_n R_m(\theta) = R_{n \times m}(\theta) S_n$	Changing axis of rotation

spectively, R_x) placeholders through the control (respectively, target) of CNOT gates, decompose arbitrary one-qubit gate placeholders into $R_nR_mR_n$ placeholders for $n \perp m$, etc.

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

Lemma 2. Fix a gate library consisting of finitely many 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 $\ell < 4^n-1$ one-parameter placeholders. Observe that matrix multiplication and tensor products 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 the placeholders. Accounting for the global phase, $Q(\mathcal{T}) = \cup_{\xi^{2n}=1} \text{Image}(\xi f)$. Sard's theorem [[19], p. 39] demands that Image (ξf) 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. ■

III. LOWER BOUNDS

Lemma 2 implies that for any given elementary gate library one can find n-qubit operators requiring at least 4^n-1 one-qubit gates. We use this fact to obtain a lower bound for the number of CNOT gates required.

Proposition 1. 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 $\{R_z, R_x, \text{CNOT}\}$ 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 $3n+4k \geqslant 4^n-1$.

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 parametrize 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 formulas for the same circuits from 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 1. Fix an elementary-gate library. Then almost all two-qubit operators cannot be simulated without at least three CNOT gates and 15 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 Sec. VI.

Proposition 2. 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 1 implies that at least three CNOT gates are necessary in general; at least five one-qubit placeholders are required for dimensional 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 five one-qubit gates can be replaced by 15 one-parameter gates in such a way that the closest parametrized gates arising from the adjacent one-qubit gates can be combined. Thus, if five one-qubit placeholders and three CNOTs suffice, then so do 14 one-parameter placeholders and three CNOTs, which contradicts dimension-based lower bounds. ■

IV. INVARIANTS OF TWO-QUBIT OPERATORS

To study two-qubit operators that differ only by pre- or post-composing with one-qubit operators, we use the terminology of *cosets*, common in abstract algebra [20]. Let G be the group of operators that can be simulated entirely by one-qubit operations. That is, $G = SU(2)^{\otimes n} = \{a_1 \otimes a_2 \otimes \ldots \otimes a_n : a_i \in SU(2)\}$. Then two operators u, v are said to be in the same left coset of SU(4) modulo G (written as $uG = vG$) if and only if u differs from v only by precomposing with one-qubit operators; that is, if $u = vg$ for some $g \in G$. Similarly, we say that u and v are in the same right coset ($Gu = Gv$) if they differ only by postcomposition ($u = hv$ for some $h \in G$), and we say that u and v are in the same double coset ($u = GvG$) if they differ by possibly both pre- and postcomposition ($u = hvg$ for some $g, h \in G$). In the literature, the double cosets are often referred to as *local equivalence classes* [4].

Polynomial invariants classifying the double cosets have been proposed by Makhlin [9]. In what follows, we present equivalent invariants which generalize to n qubits and are more straight-forward to compute. Moreover, the proofs given here detail an explicit constructive procedure to find a, b, c, d such that $(a \otimes b)u(c \otimes d) = v$, once it has been determined by computing invariants that u, v are in the same double coset.

Definition 1. We define γ_n on $2^n \times 2^n$ matrices by the formula $u \mapsto u\boldsymbol{\sigma}_y^{\otimes} u^{\mathrm{T}} \boldsymbol{\sigma}_y^{\otimes n}$ When n is arbitrary or clear from the context, we write γ for γ_n.

Proposition 3. γ has the following properties: (1) $\gamma(I) = I$. (2) $\gamma(ab) = a\gamma(b)\gamma(a^T)^T a^{-1}$. (3) $\gamma(a \otimes b) = \gamma(a) \otimes \gamma(b)$. (4) $g \in M_{2\times 2}^{\otimes n} \Rightarrow \gamma(g) = \det(g) \cdot I$. (5) γ is constant on the left cosets $u \cdot SU(2)^{\otimes n}$. (6) $\chi[\gamma]$ is constant on double cosets $SU(2)^{\otimes n} \cdot u \cdot SU(2)^{\otimes n}$.

Proof. Properties 1, 2, and 3 are immediate from the definition. Property 4 can be checked explicitly for $n=1$, and then the general case follows from (3). For property (5), note first that $g \in \mathrm{SU}(2)^{\otimes n} \Rightarrow \gamma(g)=I$ by (4). Then expressing $\gamma(ag)$ and $\gamma(a \times I)$ using (1) and (2), we see they are equal. For property 6, we use (2), (4), and (5) to see that $g,h \in SU(2)^{\otimes n} \Rightarrow \gamma(gah)=g^{-1}\gamma(ah)g=g^{-1}\gamma(a)g$ and thus $\chi[\gamma(gah)]=\chi[\gamma(a)]$. Incidentally, (6) is closely related to [[16], Theorem I.3]. ■

While γ is constant. on left cosets and $\chi[\gamma]$ on double cosets, these invariants do not in general suffice to classify cosets. Roughly, a parameter space for double cosets would need dimension $\dim[\mathrm{SU}(2^n)]-2\dim[\mathrm{SU}(2)^{\otimes n}]=4^n-6n-1$, whereas the space of possible $\chi[\gamma]$ has dimension 2^n-1 (because the 2^n roots of $\chi[\gamma]$ must all have unit length and unit product). The first dimension is much larger except for $n=1,2$. In the case $n=1$, there is only one left coset (and only one double coset), so our invariants trivially suffice. For $n=2$, these numbers come out exactly equal, and γ and $\chi[\gamma]$ serve to classify, respectively, the left and double cosets.

Proposition 4. For $u,v \in \mathrm{SU}(4)$, $G=\mathrm{SU}(2)\otimes\mathrm{SU}(2)$ (1) $u \in G \Leftrightarrow \gamma(u)=I$; (2) $uG=vG \Leftrightarrow \gamma(u)=\gamma(v)$; (3) $GuG=GvG \Leftrightarrow \chi[\gamma(u)]=\chi[\gamma(v)]$.

Proof. Recall that $E\in \mathrm{U}(4)$ can be found such that $E\ \mathrm{SO}(4)E^\dagger=G$; such matrices are characterized by the property that $EE^T=-\sigma_y\otimes\sigma y$. This and related issues have been exhaustively dealt with in several papers [10–13,16], where it is shown that E can be chosen as

$$\frac{1}{\sqrt{2}}\begin{pmatrix}1 & i & 0 & 0\\ 0 & 0 & i & 1\\ 0 & 0 & i & -1\\ 1 & -i & 0 & 0\end{pmatrix}.$$

Observe that the properties $\gamma(u)=I$, $\gamma(u)=\gamma(v)$, $\chi[\gamma(u)]=\chi[\gamma(v)]$ are not changed by replacing γ with $E^\dagger\gamma E$. Then using the fact that $-\sigma_y\otimes\sigma_y=EE^T=(EE^T)^\dagger$ compute

$$E^\dagger\gamma(g)E=E^\dagger gEE^Tg^TE^{t\dagger}E^\dagger E=(E^\dagger gE)(E^\dagger gE)^T.$$

Therefore it suffices to prove the proposition after making the following substitutions: $g\mapsto u=E^\dagger gE$, $G\mapsto \mathrm{SO}(4)$, $\gamma(g)\mapsto uu^T$. Now property 1 is immediate and property 2 follows from $uu^T=vv^T \Leftrightarrow v^\dagger u=(v^\dagger u)^{t\dagger}\Leftrightarrow v^\dagger u\in \mathrm{SO}(4)$.

To prove property 3, note that 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 $a\in\mathrm{SO}(4)$ such that $auu^Ta^\dagger$ is diagonal. By reordering (and negating) the colunmns of a, we can reorder the diagonal elements of $auu^Ta^\dagger$ as desired. Thus if $\chi[uu^T]=\chi[vv^T]$, we can find $a,b\in\mathrm{SO}(4)$ such that $auu^Ta^T=bvv^Tb^T$ by diagonalizing both; then $(v^\dagger b^Tau)(v^\dagger b^Tau)^T=I$. Let $c=v^\dagger b^Tau\in\mathrm{SO}(4)$. We have $a^Tbvc=u$, as desired. ■

The proof above gives an algorithm for computing a,b,c,d for given two-qubit u and v so that $(a\otimes b)u(c\otimes d)=v$. Also, u may be chosen as a relative phasing of Bell states.

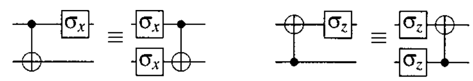

FIG. 1. Circuit identities to move σ_x,σ_z past CNOT. The σ_x identity is standard in the theory of classical reversible circuits, where σ_x is just the NOT gate, and amounts to the statement that $(1\oplus a)\oplus(1\oplus b)=(a\oplus b)$. The σ_z identity can be obtained from it by conjugating by $H\otimes H$.

V. TECHNICAL LEMMAS

We present two parametrizations of the space of double cosets described in Sec. IV. These will be used in the constructions of universal two-qubit circuit topologies to follow.

We will use the following general technique 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 $\sigma_y^{\otimes 2}g^T\sigma_y^{\otimes 2}$. As will be shown below, 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 $\gamma(u)$.

Proposition 5. For any $u\in\mathrm{SU}(4)$, one can find α,β,δ such that $\chi[\gamma(u)]=\chi[\gamma(C_1^2(I\otimes R_y(\alpha))C_2^1(R_z(\delta)\otimes R_y(\beta))C_1^2)]$.

Proof. Let $v=C_1^2(I\otimes R_y(\alpha))C_2^1(R_z(\delta)\otimes R_y(\beta))C_1^2$. As v is given explicitly by a circuit, we use the technique described above to determine the following circuit for $\gamma(v)$.

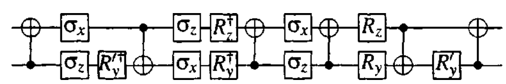

Here, $R_y'=R_y(\alpha)$, $R_y=R_y(\beta)$, and $R_z=R_z(\delta)$. We now use the circuit identities in Fig. 1 and $\sigma_iR_j(\theta)=R_j(-\theta)\sigma_i$ to push all the σ_i gates to the left of the circuit, where they cancel up to an irrelevant global phase of −1. All gates in the wake of their passing become inverted, and we obtain the following circuit.

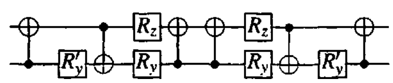

For invertible matrices, $\chi[AB]=\chi[A^{-1}(AB)A]=\chi[BA]$. In view of the fact that we are ultimately interested only in $\chi[\gamma(V)]$ we may move gates from the left of the circuit to the right. Conglomerating R_y' gates and canceling paired CNOT gates, in this way we obtain

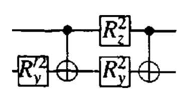

We have shown that $\chi[\gamma(v)]=\chi[C_2^1(R_z(\delta)\otimes R_y(\beta))C_2^1(I\otimes R_y(\alpha))]$. Again, since $\chi[B]=\chi[A^{-1}BA]$, we conjugate by $I\otimes S_x$. This fixes the CNOT gate and replaces R_y gates with R_z:

$$\chi[\gamma(v)]=\chi[C_2^1(R_z(\delta)\otimes R_z(\beta))C_2^1(I\otimes R_z(\alpha))].$$

Finally, 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 $\alpha=(x+y)/2$, $\beta=(x+z)/2$, and $\delta=(y+z)/2$ for

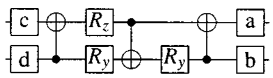

FIG. 2. A universal two-qubit circuit with three CNOT gates. It requires 10 basic gates [3] or 18 gates from {CNOT, R_y,R_z}.

e^{ix}, e^{iy}, and e^{iz} any three eigenvalues of $\gamma(U)$. ■

Proposition 6. For any $u \in \mathrm{SU}(4)$, one can find θ,ϕ,ψ such that $\chi[\gamma(uC_2^1(I\otimes R_z(\psi))C_2^1)]=\chi[\gamma(C_2^1(R_x(\theta)\otimes R_z(\phi))C_2^1)]$.

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, this is $\mathrm{tr}[\gamma(u^T)^T\gamma(\Delta)]$. Explicit computation as in the previous proposition gives $\gamma(\Delta)=\Delta^2$, and one obtains $\mathrm{tr}[\gamma(u\Delta)]=(t_1+t_4)e^{-i\psi}+(t_2+t_3)e^{i\psi}$, where t_1,t_2,t_3,t_4 are the diagonal entries of $\gamma(u^T)^T$. We may ensure that this number is real by requiring $\tan(\psi)=\mathrm{Im}(t_1+t_2+t_3+t_4)/\mathrm{Re}(t_1+t_2-t_3-t_4)$.

Now consider $m\in\mathrm{SU}(N)$, $\chi[m]=\Sigma\dot{a}_iX^i=\Pi(X-r_i)$, where the r_i form the spectrum of m. Since $m\in\mathrm{SU}(N)$, we must have $\Pi r_i=1=\Pi\overline{r_i}$. Therefore, $\chi[m]=\chi[m]\Pi\overline{r_i}=\Pi(\overline{r_i}X-1)$. Expanding the equality $\Pi(X-r_i)=\Pi(\overline{r_i}X-1)$ gives $\overline{a_i}=a_{N-i}$. In particular, for $N=4$, $a_2\in\mathbb{R}$, and $\mathrm{tr}(m)=a_3=\overline{a_1}$. Since $a_4=a_0=1$, $\chi[m]$ has all real coefficients if and only if $\mathrm{tr}[m]\in\mathbb{R}$. In this case, the roots of $\chi[m]$ must come in conjugate pairs: $\chi[m]=(X-e^{ir})(X-e^{-ir})(X-e^{is})(X-e^{-is})$. On the other hand, for $w=C_2^1(R_x((r+s)/2)\otimes R_z((r-s)/2))C_2^1$, one can verify that $\chi[\gamma(w)]$ takes this form.

Taking $m=\gamma(UC_2^1(I\otimes R_z(\psi))C_2^1)$, with ψ as determined above, we obtain $\theta=(r+s)/2$, $\phi=(r-s)/2$. ■

VI. MINIMAL TWO-QUBIT CIRCUITS

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

In view of Lemma 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 contain only 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].

Theorem 1. Fifteen $\{R_y,R_z\}$ gates and three CNOTs suffice to simulate an arbitrary two-qubit operator.

Proof. Choose α,β,δ as in Proposition 5. Then by Proposition 4 one can find $a,b,c,d\in\mathrm{SU}(2)$ 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 Fig. 2 is universal. ■

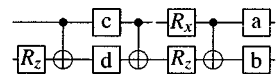

FIG. 3. Another universal two-qubit circuit with three CNOT gates. It requires 10 basic gates [3] or 18 gates from {CNOT, R_x,R_z}.

Theorem 2. Fifteen $\{R_x,R_y\}$ gates and three CNOTs suffice to simulate an arbitrary two-qubit operator.

Proof. Conjugation by $H^{\otimes n}$ fixes $\mathrm{SU}(2^n)$ and R_y. It also flips CNOT gates $(H^{\otimes 2}C_1^2H^{\otimes 2}=C_2^1)$ 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 the Appendix that no such circuit can be universal and articulate the implications of this distinction in Sec. VII. Nonetheless, we demonstrate here a universal two-qubit circuit topology with gates from the {R_x, R_z, CNQT} gate library that contains 15 one-qubit gates and three CNOT gates.

Theorem 3. Fifteen {R_x, R_z} gates and three CNOTs suffice to simulate an arbitrary two-qubit operator.

Proof. Let U' be the desired operator; set $U=U'C_2^1$. Choose θ,ϕ,ψ for U' as in Proposition 6. By Proposition 4, one can find $a,b,c,d\in\mathrm{SU}(2)$ such that

$$U(I\otimes R_z(\psi))C_2^1=(a\otimes b)C_2^1(R_z(\theta)\otimes R_x(\phi))C_2^1(c\otimes d)$$

Solving for U gives the overall circuit topology in Fig. 3. ■

Unlike the circuit of Theorem 1, the circuit in Fig. 3 can be adapted to both other gate libraries. We can replace c by $S_z(S_z^\dagger c)$ and a by $(aS_z)S_z^\dagger$ and 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 Theorem 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 6, 5, and 4, and Theorems 1, 2, and 3 as algorithms. By reordering eigenvalues in the proof of Proposition 4, one may typically produce several different circuits. Similar degrees of freedom are discussed in [5].

To complete Table I, count *basic* gates in Figs. 2 or 3.

VII. CONCLUSIONS

Two-qubit circuit synthesis is relevant to ongoing physics experiments and can be used in peephole optimization of larger circuits, where small subcircuits 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 constructively synthesize small circuits for arbitrary two-qubit operators with respect to several gate libraries. Most of our lower and upper bounds on worst-case gate counts are tight, and rely on circuit identities summarized in Table II. We also prove that n-qubit circuits require $\left\lceil\frac{1}{4}(4^n-3n-1)\right\rceil$ CNOT gates in the worst case.

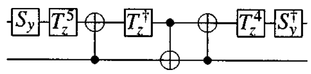

FIG. 4. The result of our algorithm applied to the two-qubit quantum Fourier transform. The circuit contains three one-qubit gates and three CNOTs, but the one-qubit gates are broken up into elementary gates for specificity. Here, $T_z = R_z(\pi/4)$ is the T gate defined in [1] up to a global phase.

While our techniques do not guarantee optimal circuits for non-worst-case operators, they perform well in practice: one run of our algorithm produced the circuit shown in Fig. 4 for the two-qubit quantum Fourier transform. We show elsewhere that this circuit has a minimal basic-gate count.

A somewhat surprising result of our work is the apparent asymmetry between R_x, R_y, and R_z gates. While one would expect any circuit topology for CNOT, R_z, and R_y to carry over to other elementary-gate libraries, we prove a negative result for the library CNOT, R_z, and R_x; namely, using R_y gates appears essential for the minimal universal circuit topology shown in Fig. 2, which exhibits the maximal possible number of one-qubit gates that are not between any two CNOT gates.

The asymmetry between elementary one-qubit gates directly impacts peephole optimization of n-qubit circuits, where decompositions like that in Fig. 2 are preferable to that in Fig. 3. For example, consider a three-qubit circuit consisting of two two-qubit blocks on lines (1) one and. two, (2) two and three. If both blocks are decomposed as in Fig. 2, 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 Fig. 3 is used.

ACKNOWLEDGMENTS

This work is funded by the DARPA QuIST program and an NSF grant.

APPENDIX

We now illustrate the counterintuitive difference between (i) the *CXZ* library and (ii) libraries *CYZ* and *CXY*. That is, universal circuit topologies with certain properties exist only for the *CYZ* and *CXY* libraries.

The proof of Theorem 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. Fix an elementary-gate library. There exist unitary operators $U \in \mathrm{SU}(4)$ 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 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 1, there exist operators that cannot be simulated without 15 one-parameter gates; the remaining three must go in the middle of the circuit. ■

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 and 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.

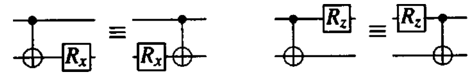

When one CNOT gate occurs immediately after another in a circuit, we say that they are *adjacent*. When such pairs of CNOTs share control lines, they cancel out, and otherwise may still lead to reductions as discussed below. We will be interested in circuits that do not allow such simplifications. To 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 if and only if 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.

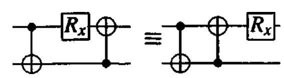

On the other hand, two CNOT gates of the same orientation can be effectively separated by a single R_x or R_z gate, as shown below. Up to swapping wires, these are the only ways to effectively separate two CNOTs with a single R_x or R_z.

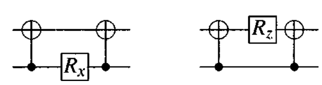

Proposition 8. At least four gates from $\{R_x, R_z\}$ are necessary to effectively separate four or more CNOT gates.

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. ■

Denote by ω^{ij} the SWAP gate that exchanges the ith and jth qubits. It can be simulated using CNOTs as $C_i^j C_j^i C_i^j = \omega^{ij} = C_j^i C_i^j C_j^i$. 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., $C_1^2 C_1^2$ or $C_2^1 C_2^1$), then they cancel out and can be removed from the circuit. Otherwise, use the identity $C_1^2 C_2^1 = C_2^1 \omega^{12}$ or $C_2^1 C_1^2 = C_1^2 \omega^{12}$ and push the SWAP to the end of the circuit. We apply this technique at the level of circuit topologies and observe that $Q(\mathcal{T}\omega^{12})$ is measure zero (or universal) if and only if $Q(\mathcal{T})$ is. By the above discussion, we can always reduce to an effectively separated circuit before checking these properties.

Proposition 9. Almost all unitary operators $U \in \mathrm{SU}(4)$ 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 simulate only. 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 8 and the discussion immediately following it, we need only consider circuit topologies with no more than three CNOTs. On the other hand, we know from Proposition 2 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 CNOTs partition a minimal two-qubit circuit into 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 CNOTs. To this end, we consider two pairs of CNOTs (the central CNOT is in both pairs), and distinguish these three cases: (1) both pairs of CNOTs consist of gates of the same orientation, (2) both consist of gates of opposite orientations, or (3) one pair has gates of 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 CNOTs 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.

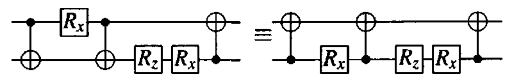

Finally, consider the case in which all three CNOT gates have the same orientation. Each pair of consecutive CNOTs must have at least one R_x or R_z between them, to be effectively separated. Thus one of the pairs has a single R_x or R_z between its members, and the other has two one-qubit gates. We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

Suppose that the one-qubit gates separating the 2-pair of CNOTs occur on different lines. If either one-qubit gate can commute past the CNOTs of the 2-pair, then it can move to the edge of the circuit; in this case Proposition 7 implies that the circuit topology we are looking at can simulate only a measure-zero subset of SU(4) (one can show that two R_x, R_z gates cannot effectively separate three CNOTs). Otherwise, we use the identity $C_2^1\big(R_x(\alpha)\otimes R_z(\beta)\big)C_2^1 = C_1^2\big(R_z(\beta)\otimes R_x(\alpha)\big)C_1^2$ to flip the 2-pair, and thus the 1-pair now have opposite orientations. As there is only one one-qubit gate between them, this pair is not effectively separated. For example,

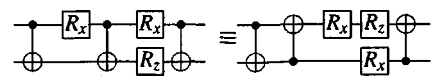

We are left with the possibility that all the CNOT gates have the same orientation and that the 2-pair's one-qubit gates appear on the same line. Both R_z, R_x must occur, or else we could combine them and apply Proposition 7 to show that such a circuit topology can simulate only a measure-zero subset of $\mathrm{SU}(2^n)$. 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 again implies that the circuit topology simulates only a measure-zero subset of $\mathrm{SU}(2^n)$. 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 that originally separated the 1-pair must be on the same line as the commuting gate. It follows that all gates are on the same line. Up to conjugating by $H\otimes H$, swapping wires, and inverting the circuit, this leaves exactly one possibility:

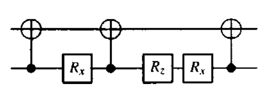

Finally, we add the four one-qubit gates on the sides, decompose each into $R_xR_zR_x$ via Lemma 1, and observe that an R_x gate can commute across the top and be absorbed on the other side. This leaves 14 one-parameter gates, and by Lemma 2, such a circuit topology simulates only a measure-zero subset of SU(4). ∎

[1] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, England, 2000).
[2] D. P. DiVincenzo, Phys. Rev. A **51**, 1015 (1995).
[3] A. Barenco *et al.*, Phys. Rev. A **52**, 3457 (1995).
[4] J. Zhang *et al.*, Phys. Rev. Lett. **91**, 027903 (2003).
[5] S. S. Bullock and I. L. Markov, Phys. Rev. A **68**, 012318 (2003).
[6] F. Vatan and C. Williams, Phys. Rev. A **69**, 032315 (2004).
[7] G. Vidal and C. M. Dawson, Phys. Rev. A **69**, 010301 (2004).
[8] J. Zhang *et al.*, e-print quant-ph/0308167.
[9] Yu. Makhlin, Quantum Inf. Process. **1**, 243 (2002).
[10] C. H. Bennett *et al.*, Phys. Rev. A **54**, 3824, (1996).
[11] S. Hill and W. K. Wootters, Phys. Rev. Lett. **78**, 5022, (1997).
[12] N. Khaneja, R. Brockett, and S. J. Glaser, Phys. Rev. A **63**, 032308 (2001).
[13] M. Lewenstein *et al.*, Phys. Rev. A **63**, 044304 (2001).
[14] E. M. Rains, e-print quant-ph/9704042.
[15] M. Grassl *et al.*, Phys. Rev. A **58**, 1833 (1998).
[16] S. S. Bullock and G. K. Brennen, J. Math. Phys. **45**, 2447 (2004).
[17] E. Knill, e-print quant-ph/9508006.
[18] D. Wineland *et al.*, J. Res. Natl. Inst. Stand. Technol. **103**, 259 (1998).
[19] V. Guillemin and A. Pollack, *Differential Topology* (Prentice-Hall, Englewoods Cliffs, NJ, 1974).
[20] M. Artin, *Algebra* (Prentice-Hall, Englewood Cliffs, NJ, 1991).