Optimal Realizations of Controlled Unitary Gates

Guang Song and Andreas Klappenecker
Department of Computer Science
Texas A&M University
College Station, TX 77843-3112
{gsong,klappi}@cs.tamu.edu

Abstract

The controlled-not gate and the single qubit gates are considered elementary gates in quantum computing. It is natural to ask how many such elementary gates are needed to implement more elaborate gates or circuits. Recall that a controlled-U gate can be realized with two controlled-not gates and four single qubit gates. We prove that this implementation is optimal if and only if the matrix U satisfies the conditions $\operatorname{tr} U \neq 0$, $\operatorname{tr}(UX) \neq 0$, *and* $\det U \neq 1$. *We also derive optimal implementations in the remaining non-generic cases.*

1 Introduction

It was shown in the seminal paper [1] that any unitary $2^n \times 2^n$ matrix M can be realized on a quantum computer with n quantum bits by a finite sequence of controlled-not and single qubit gates. We will refer to controlled-not and single qubit gates as elementary gates. It is natural to ask how many elementary gates are necessary and sufficient to realize a given unitary matrix M. Answering such questions is a notoriously difficult task.

It was shown in [1] that a controlled unitary operation can be realized with at most six elementary gates, that is, given a unitary 2×2 matrix U, there exist unitary matrices A, B, C, and E such that

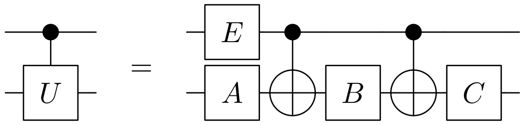

(1)

Our main result shows that this implementation is optimal:

Theorem A *Suppose that* U *is a unitary* 2×2 *matrix satisfying* $\det U \neq 1$, $\operatorname{tr} U \neq 0$, *and* $\operatorname{tr} UX \neq 0$. *Then six elementary gates are necessary and sufficient to implement a controlled-U gate.*

Notations. We use the following abbreviations throughout this paper:

$$H = \frac{1}{\sqrt{2}}\left(\begin{array}{rr} 1 & 1 \\ 1 & -1 \end{array}\right), \quad X = \left(\begin{array}{cc} 0 & 1 \\ 1 & 0 \end{array}\right), \quad Z = \left(\begin{array}{rr} 1 & 0 \\ 0 & -1 \end{array}\right).$$

We denote by $\mathbf{C}$ the field of complex numbers, and by $\mathbf{R}$ the field of real numbers. We will say that a unitary matrix U is generic if and only if the conditions $\det U \neq 1$, $\operatorname{tr} U \neq 0$, and $\operatorname{tr} UX \neq 0$ are satisfied. Notice that the inverse $U^\dagger$ of a generic unitary matrix U is again generic.

Theorem A gives a sharp lower bound on the number of elementary gates that are needed to implement a generic controlled-U operation. The non-generic case is discussed in Theorem B in Section 3.

2 Proof of Theorem A

We will show that any implementation of a generic controlled-U operation requires at least six elementary gates. We classify the possible implementation in terms of the number of controlled-not operations used. We will use entanglement properties to rule out various potential implementations. The following simple fact will turn out to be particularly helpful:

Lemma 1 *Let $|\psi\rangle$, $|\phi\rangle$ be nonzero elements of $\mathbf{C}^2$. The input $|\psi\rangle \otimes |\phi\rangle$ to a controlled-U gate will produce an entangled output state if and only if $|\phi\rangle$ is not an eigenvector of U and $|\psi\rangle = a|0\rangle + b|1\rangle$ with $a, b \neq 0$.*

Proof. The input $|\psi\rangle \otimes |\phi\rangle = (a|0\rangle + b|1\rangle) \otimes |\phi\rangle$ to the controlled-U gate produces the result

$$|r_{out}\rangle = a|0\rangle \otimes |\phi\rangle + b|1\rangle \otimes U|\phi\rangle$$

Denote by $|\phi^\perp\rangle$ a nonzero vector in $\mathbf{C}^2$ satisfying $\langle\phi^\perp|\phi\rangle = 0$. Consequently, $U|\phi\rangle = c|\phi\rangle + d|\phi^\perp\rangle$ with $c, d \in \mathbf{C}$, and $d \neq 0$. Therefore, the output state $|r_{out}\rangle$ can be expressed in the form

$$|r_{out}\rangle = a\,|0\rangle \otimes |\phi\rangle + bc\,|1\rangle \otimes |\phi\rangle + bd\,|1\rangle \otimes |\phi^\perp\rangle \tag{2}$$

with $a, b, d \neq 0$. Seeking a contradiction, we assume that $|r_{out}\rangle$ is not an entangled state. This would mean that there exist complex coefficients $\alpha, \beta, \gamma, \delta$ such that

$$\begin{aligned} |r_{out}\rangle &= (\alpha|0\rangle + \beta|1\rangle) \otimes (\gamma|\phi\rangle + \delta|\phi^\perp\rangle) \\ &= \alpha\gamma|0\rangle \otimes |\phi\rangle + \alpha\delta|0\rangle \otimes |\phi^\perp\rangle + \beta\gamma|1\rangle \otimes |\phi\rangle + \beta\delta|1\rangle \otimes |\phi^\perp\rangle. \end{aligned}$$

Comparing coefficients with (2) shows that $\alpha\delta = 0$, hence α or δ has to be zero. Either choice leads to a contradiction.

On the other hand, if $|\psi\rangle$ is a multiple of $|0\rangle$ or of $|1\rangle$ or if $|\phi\rangle$ is an eigenvector of U, then it follows from the definitions that the output of the controlled-U gate will not be entangled. □

Corollary 2 *Assume that $|\phi\rangle$ is an eigenvector of U with eigenvalue λ_ϕ, and a state $|\psi\rangle \in \mathbf{C}^2$. If we input $|\psi\rangle \otimes |\phi\rangle$ to the controlled-U gate, then the output is of the form* $\mathrm{diag}(1, \lambda_\phi)|\psi\rangle \otimes |\phi\rangle$. *In particular, the output is not entangled.*

Another simple consequence of this lemma is that a controlled-U gate is able to produce an entangled output state from some input state $|\psi\rangle \otimes |\phi\rangle$, as long as U is not a multiple of the identity matrix. Any matrix U with $\mathrm{tr}(UX) \neq 0$ satifies this condition. In particular, we need at least one controlled-not gate to implement a controlled-U gate with $\mathrm{tr}(UX) \neq 0$.

One Controlled-Not Gate. We consider now possible implementations of a generic controlled-U gate with only one controlled-not gate and some single qubit gates. Recall that it is possible to switch the control and the target qubit of a controlled-not gate by conjugation with Hadamard matrices H:

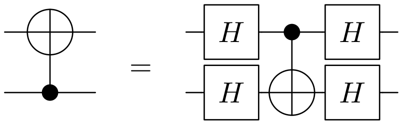

(3)

Therefore, we can assume without loss of generality that the controlled-U gate is expressed in the following form:

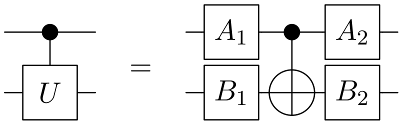

(4)

Lemma 3 *The unitary matrices A_1, A_2 used in the single qubit gates in (4) have to be both diagonal or both antidiagonal.*

Proof. Suppose that we input $|i\rangle \otimes |\alpha\rangle$, where $i = 0, 1$ and $|\alpha\rangle$ is some arbitrary vector in $\mathbf{C}^2$ such that $B_1|\alpha\rangle$ is not an eigenvector of X. Lemma 1 shows that the output of the controlled-U gate is not entangled when provided with such an input state, since the most significant qubit is not in superposition. Notice that the circuit on the right hand side in (4) will produce an entangled output unless A_1 is diagonal or antidiagonal. It follows that A_1 is of the desired form. The same argument applied to the inverse circuits proves that A_2 has to be diagonal or antidiagonal. It is clear that A_1 and A_2 are either both diagonal or both antidiagonal, because $|00\rangle$ has to be an eigenstate of the circuit (4). □

We can assume that A_1 and A_2 are diagonal. Indeed, if the A_i's are antidiagonal, then we can replace the controlled-not gate in (4) by

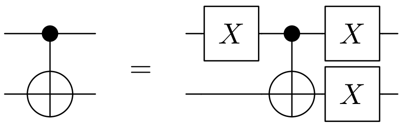

(5)

Here we used the fact that XA_1 and A_2X will be both diagonal, when A_1, A_2 are antidiagonal.

Lemma 4 *If* $\operatorname{tr} U \neq 0$, *then the circuit (4) cannot implement a controlled-U operation.*

Proof. The preceding discussion shows that A_1 and A_2 are both diagonal or antidiagonal, and we may assume that A_1 and A_2 are diagonal. Thus, there exist real numbers ϑ_0 and ϑ_1 such that $A_2A_1|0\rangle = e^{i\vartheta_0}|0\rangle$ and $A_2A_1|1\rangle = e^{i\vartheta_1}|1\rangle$. It follows that $B_2B_1 = e^{-i\vartheta_0}I$ and $B_2XB_1 = e^{-i\vartheta_1}U$. Consequently, $B_1^\dagger XB_1 = e^{i(\vartheta_0-\vartheta_1)}U$, which implies $\operatorname{tr} U = 0$. Therefore, it is in general not possible to implement a controlled-U operation with one controlled-not gate and several single qubit gates. □

Two Controlled-Not Gates. Assume that we have now two controlled-not gates and several single qubit gates at our disposal. This allows to express the controlled-U gate in the form

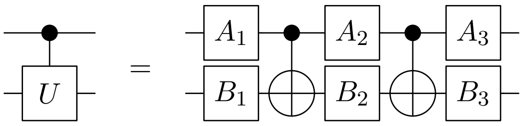

(6)

In fact, any implementation of a controlled-U gate with two controlled-not gates can be reduced to this form. Indeed, it is possible to swap the control and target qubits of a controlled-not gate by conjugation with Hadamard gates, as we have seen in our discussion of the previous case. We will see what kind of properties have to be satisfied by the matrices A_i and B_i. The following Lemmas will prepare us to prove Proposition 13.

We say that a unitary 2×2 matrix V is sparse if and only if V is diagonal or antidiagonal.

Lemma 5 *If the matrix A_1 in (6) is sparse, then A_2, A_3 are sparse as well.*

Proof. Suppose that A_1 is diagonal. Choose a state $|0\rangle \otimes B_1^\dagger B_2^\dagger |\psi\rangle$ as input, where $|\psi\rangle$ is not an eigenstate of X. Notice that a controlled-U operation leaves the input $|0\rangle \otimes |\psi\rangle$ invariant. Consider now the evolution of the input state through the circuit on the right hand side of (6). Because A_1 is diagonal, the resulting state after the first controlled-not operation is $\alpha|0\rangle \otimes B_2^\dagger|\psi\rangle$, where α is a scalar phase factor. Applying the single qubit operations A_2 and B_2 yields $\alpha A_2|0\rangle \otimes |\psi\rangle$. Lemma 1 shows that the output after the second controlled-not operation will be entangled, unless A_2 is sparse. Therefore, A_2 has to be sparse. Thus, the state after the second controlled-not is of the form $\beta|i\rangle \otimes |\psi'\rangle$, where β is some phase factor, $i = 0, 1$, and $|\psi'\rangle$ is some element of $\mathbf{C}^2$. Hence A_3 has to map $|i\rangle$ to $|0\rangle$, up to a phase factor, i.e., A_3 has to be sparse.

If A_1 is antidiagonal, then we can use the identity (5) to replace the antidiagonal matrix A_1 by the diagonal matrix XA_1, which allows us to conclude that A_2X, hence A_2, and A_3 are sparse. □

Lemma 6 *Suppose that U is not a multiple of the identity matrix. If A_1 in the circuit (6) is not sparse, then A_2, A_3 are not sparse either.*

Proof. Assume that the input state is of the form $A_1^\dagger|0\rangle \otimes |\psi\rangle$, where $|\psi\rangle$ is not an eigenvector of U. Since A_1 is not sparse, $A_1^\dagger|0\rangle = a|0\rangle + b|1\rangle$ with $a, b \neq 0$. Therefore, the input $A_1^\dagger|0\rangle \otimes |\psi\rangle$ to the controlled-U operation will yield an entangled output state, according to Lemma 1.

On the other hand, consider the right hand side of (6). The input state $A_1^\dagger|0\rangle \otimes |\psi\rangle$ produces after the first controlled-not gate a state of the form $|0\rangle \otimes B_1|\psi\rangle$. The input to the second controlled-not gate is then $A_2|0\rangle \otimes B_2 B_1|\psi\rangle$. Lemma 1 shows that the output of the second controlled-not operation cannot be entangled, unless A_2 is not sparse. Therefore, A_2 is not sparse. However, A_3 cannot be sparse either, because this would imply that A_2 and A_1 are sparse, as can be seen by applying Lemma 5 to the inverse circuit. □

Lemma 7 *Let U be a unitary 2×2 matrix. Assume that A_1, A_2, A_3 in (6) are sparse. If $\operatorname{tr} U \neq 0$, then $B_2 \neq H$. If $\operatorname{tr}(UX) \neq 0$, then none of the matrices B_1, B_2, B_3 can be equal to an identity matrix, and B_1, B_3 cannot both be equal to H.*

Proof. Comparing the result of the inputs $|0\rangle \otimes |\psi\rangle$ and $|1\rangle \otimes |\psi\rangle$ on the left and right hand side of (6) yields

$$e^{i\theta_0} I = B_3 X^k B_2 X^\ell B_1, \tag{7}$$

$$e^{i\theta_1} U = B_3 X^{1-k} B_2 X^{1-\ell} B_1, \tag{8}$$

for some $k, \ell \in \{0,1\}$. Notice that equation (7) implies $B_2 = e^{i\theta_0} X^k B_3^\dagger B_1^\dagger X^\ell$. Substituting B_2 in (8) yields

$$U = e^{i(\theta_0 - \theta_1)} B_3 X B_3^\dagger B_1^\dagger X B_1. \tag{9}$$

Step 1. We show that $B_i \neq I$ for $i = 1, 2, 3$:

i) Suppose that $B_1 = I$. Equation (9) implies $\operatorname{tr}(UX) = 0$, contradicting our assumptions.

ii) Suppose that $B_2 = I$. Equations (7) and (8) then imply $U = e^{i(\theta_0 - \theta_1)} I$, thus $\operatorname{tr}(UX) = 0$, which contradicts our assumptions.

iii) Suppose that $B_3 = I$. Equation (9) implies $\operatorname{tr}(XU) = 0$, whence $\operatorname{tr}(UX) = 0$. This contradicts our assumptions.

Step 2. We show that $B_2 \neq H$. Seeking a contradiction, we suppose that $B_2 = H$. From (7) and (8), $U = e^{i(\theta_0 - \theta_1)} B_3 X^{1-k} Z X^k B_3^\dagger$, which implies $\operatorname{tr}(U) = 0$. Contradiction.

Step 3. The case $B_1 = B_3 = H$ immediately leads to a contradiction, because (9) would imply $U = e^{i(\theta_0 - \theta_1)} I$, and thus $\operatorname{tr}(UX) = 0$. □

Lemma 8 *If $\det U \neq 1$, then at least one of the matrices A_i in (6) is not equal to the identity matrix.*

Proof. Seeking a contradiction, we assume that A_1, A_2, A_3 are identity matrices. It follows at once that $B_3 B_2 B_1 = e^{i\phi} I$, and $B_3 X B_2 X B_1 = e^{i\phi} U$, hence $\det U = 1$. It follows that one of the matrices A_i has to differ from the identity matrix. □

Lemma 9 *If the matrices A_1, A_2, A_3 in (6) are not sparse, then $B_2 \neq H$ and $B_1^\dagger |\omega_0\rangle$ and $B_1^\dagger |\omega_1\rangle$ are eigenvectors of U, where*

$$|\omega_0\rangle = (|0\rangle + |1\rangle)/\sqrt{2} \quad \textit{and} \quad |\omega_1\rangle = (|0\rangle - |1\rangle)/\sqrt{2}$$

are eigenvectors of X.

Proof. Consider the input $|i\rangle \otimes B_1^\dagger|\omega_k\rangle$, with $i, k = 0, 1$. Note that the controlled-U operation does not produce an entangled output state when provided with such an input. On the other hand, consider the evolution of these states in the circuit (6). The first two single qubit operations yield the state $A_1|i\rangle \otimes |\omega_k\rangle$. The controlled-not operation produces the state $Z^k A_1|i\rangle \otimes |\omega_k\rangle$, where we have used the fact that $|\omega_k\rangle$ is an eigenvector of X with eigenvalue $(-1)^k$. The result of the next two single qubit operations is then $A_2 Z^k A_1|i\rangle \otimes B_2|\omega_k\rangle$. Notice that the matrix $A_2 Z^k A_1$ cannot be sparse, because this would imply that A_3 is sparse. In other words, the input to the second controlled-not gate is a state of the form $(a_i|0\rangle + b_i|1\rangle) \otimes B_2|\omega_k\rangle$ with $a_i, b_i \neq 0$. Since the circuit implements a controlled-U operation, this gate should not produce an entangled output state. Therefore, $B_2|\omega_k\rangle$ has to be an eigenvector of X. However, this means that the input $|\psi\rangle \otimes B_1^\dagger|\omega_k\rangle$ to (6) does not get entangled for arbitrary states $|\psi\rangle$. Consequently, $B_1^\dagger|\omega_k\rangle$ has to be an eigenvector of U by Lemma 1.

The previous discussion showed that $B_2|\omega_k\rangle$, $k = 0, 1$, has to be an eigenvector of X, i.e., is mapped to a multiple of $|\omega_\ell\rangle$ for $\ell = 0, 1$. In particular, B_2 cannot be the Hadamard matrix H. □

Lemma 10 *If the matrices A_1, A_2, A_3 in the circuit (6) are all nonsparse, then either $A_3A_2A_1$ or $A_3ZA_2A_1$ is a diagonal matrix. In particular, it is not possible that $A_i = A_{i+1} = H$ for $i = 1, 2$.*

Proof. Recall that $B_1^\dagger|\omega_0\rangle$ is an eigenstate of U, where $|\omega_0\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$, as is shown in Lemma 9. The circuit on the right hand side of (6) maps the input $|0\rangle \otimes B_1^\dagger|\omega_0\rangle$ to $A_3 Z^\ell A_2 A_1|0\rangle \otimes B_3 B_2|\omega_0\rangle$, where $\ell = 0$ if B_2 maps $|\omega_0\rangle$ to a multiple of itself, and $\ell = 1$ otherwise. Comparing this state with the supposed output state $|0\rangle \otimes B_1^\dagger|\omega_0\rangle$ shows that $A_3 Z^\ell A_2 A_1|0\rangle$ coincides, up to a phase factor, with $|0\rangle$. Hence $A_3 Z^\ell A_2 A_1$ has to be a diagonal unitary matrix. The second statement is obvious. □

Lemma 11 *Let U be a unitary 2×2 matrix with $\mathrm{tr}(UX) \neq 0$. If A_1, A_2, A_3 in (6) are not sparse, then $B_1, B_2, B_3 \neq H$, and at least one of the matrices B_1, B_2, B_3 differs from the identity matrix.*

Proof. Let $|\omega_0\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$ and $|\omega_1\rangle = (|0\rangle - |1\rangle)/\sqrt{2}$. Lemma 9 showed that $B_1^\dagger|\omega_0\rangle$ and $B_1^\dagger|\omega_1\rangle$ are eigenvectors of U, that is, $UB_1^\dagger|\omega_k\rangle = \alpha_k B_1^\dagger|\omega_k\rangle$ for $k = 0, 1$. Hence $B_1 U B_1^\dagger = H\,\mathrm{diag}(\alpha_0, \alpha_1)H$. The choice $B_1 = H$ would force U to be diagonal, which would contradict $\mathrm{tr}(UX) \neq 0$. The same argument applied to the inverse circuit proves $B_3 \neq H$. We already know that $B_2 \neq H$ by Lemma 9.

Seeking a contradiction, we assume that $B_1 = B_2 = B_3 = I$. A potential implementation of the controlled-U gate is given by:

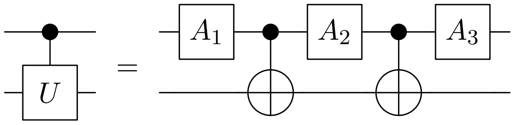

(10)

The choice $B_1 = I$ implies $U = B_1 U B_1^\dagger = H \operatorname{diag}(\alpha_0, \alpha_1) H$, according to our discussion above. This special form of U shows that $|\omega_0\rangle$ and $|\omega_1\rangle$ are eigenvectors of U corresponding to the eigenvalues α_0 and α_1. If we input a state $|\psi\rangle \otimes |\omega_0\rangle$, then Corollary 2 shows that $A_3 A_2 A_1 = \operatorname{diag}(1, \alpha_0)$. Similarly, we obtain $A_3 Z A_2 Z A_1 = \operatorname{diag}(1, \alpha_1)$, considering input states of the form $|\psi\rangle \otimes |\omega_1\rangle$.

The determinants of $A_3 A_2 A_1$ and $A_3 Z A_2 Z A_1$ are the same, hence α_0 has to coincide with α_1. However, this implies that U is diagonal, because $U = H \operatorname{diag}(\alpha_0, \alpha_1) H = \operatorname{diag}(\alpha_0, \alpha_0)$, whence $\operatorname{tr}(UX) = 0$. Therefore, at least one of the matrices B_i has to differ from the identity matrix. □

Lemma 12 *If* $\det U \neq 1$, *then the circuit (11) cannot implement a controlled-U gate.*

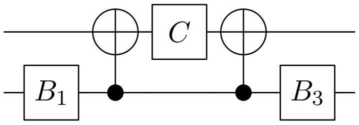

(11)

Proof. Transforming (11) into the form (6) yields $A_1 = H$, $A_2 = HCH$, and $A_3 = H$. Lemma 10 shows that $A_3 Z^\ell A_2 A_1 = \operatorname{diag}(\alpha_0, \alpha_1)$, with $\ell = 0$. Therefore, $C = \operatorname{diag}(\alpha_0, \alpha_1)$. A diagonal matrix C satisfies

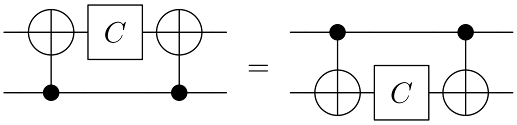

It follows that the circuit (11) can be written in the form (6) with $A_i = I$ for $i = 1, 2, 3$. This contradicts Lemma 8. □

Proposition 13 *Suppose that* U *is a unitary* 2×2 *matrix satisfying* $\operatorname{tr} U \neq 0$ *and* $\operatorname{tr}(UX) \neq 0$. *Any implementation of a controlled-U gate with two controlled-not gates and some single qubit gates needs at least a total of six gates provided that* $\det U \neq 1$, *and at least a total of five gates otherwise.*

Proof. Suppose we are given a fixed control-U gate. The implementations of a controlled-U gate with two controlled not gates and single qubit gates can be classified according to the positions of the target qubits of the two controlled-not gates. We will show that any of the four implementation types will require at least six elementary gates.

Case 1. Suppose that the target bit of both controlled-not operations is the least significant bit, as shown in (6).

Suppose that A_1 is sparse. We know from Lemma 7 that none of the matrices B_i, $i = 1, 2, 3$, can be an identity matrix, whence we have a total of five or more gates. If $\det U \neq 1$, then Lemma 8 shows that at least one of the matrices A_i is not the identity matrix, giving an additional gate.

Suppose that A_1 is not sparse. Then A_2 and A_3 are not sparse either, by Lemma 6. We know from Lemma 11 that at least one of the matrices B_1, B_2, B_3 is not an identity matrix, whence we have a total of at least six gates.

Case 2. Suppose that the first controlled-not gate acts on the most significant bit and the second controlled-not gate acts on the least significant bit. So the circuit is of the form:

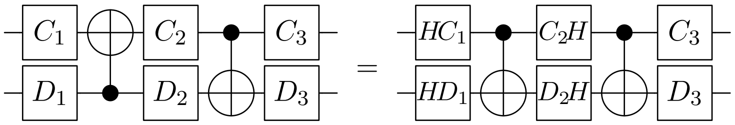

(12)

We use the circuit on the right hand to show that the circuit on the left hand side cannot have less than six elementary gates.

Assume that HC_1 is sparse. Then C_2H has to be sparse as well, hence C_1 and C_2 cannot be identity matrices. Lemma 7 shows that $D_3 \neq I$ and $D_2H \neq H$, hence $D_2 \neq I$. Thus we have at least six elementary gates.

Assume that HC_1 is not sparse. Then C_3 cannot be sparse. Either C_1 or C_2 has to differ from the identity matrix, because Lemma 10 shows that HC_1 and C_2H cannot both be equal to H. Lemma 11 shows that $HD_1 \neq H$ and $D_2H \neq H$, hence $D_1, D_2 \neq I$. Thus we have at least six elementary gates.

Case 3. Suppose that the first controlled-not gate acts on the least significant bit and the second controlled-not gate acts on the most significant bit. The inverse circuit cannot implement a controlled-$U^\dagger$ operation with less than six elementary gates, because it is of the form discussed in Case 2.

Case 4. Finally, suppose that the target qubit of both controlled-not gates is the most significant qubit. Thus, the circuit is of the form

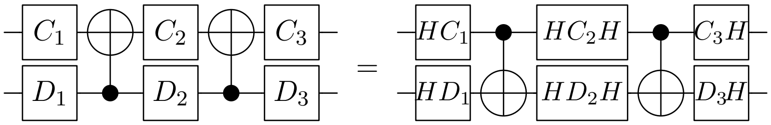

(13)

Assume that HC_1 is sparse, then C_3H is sparse as well, thus $C_1, C_3 \neq I$. Either D_1 or D_3 differs from the identity, because Lemma 7 shows that HD_1 and D_3H cannot both be equal to H. Futhermore, Lemma 7 shows that $HD_2H \neq I$, hence $D_2 \neq I$. Therefore, we have at least six gates.

Assume that HC_1 is not sparse. Lemma 11 shows that $HD_1 \neq H$ and $D_3H \neq H$ holds. Therefore $D_1, D_3 \neq I$. Lemma 6 shows that HC_2H cannot be sparse, hence $C_2 \neq I$. Therefore, we have at least five gates. If $\det U \neq 1$, then D_1, C_2, D_3 cannot be the only nontrivial single qubit gates, as Lemma 12 shows, proving that we have at least six elementary gates in that case. □

More than Two Controlled-Not Gates. Although it is undesirable to use more than two controlled-not gates, we need to show (for mathematical completeness) that implementations of a controlled-U gate with three or more controlled-not gates cannot reduce the total number of elementary gates. Fortunately, it turns out that the proof of this case is much simpler, because an implementation with five or fewer gates can then have at most two single qubit gates. Thus, the circuit can be expressed in the following form:

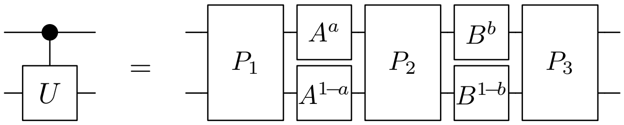

(14)

where $a, b \in \{0, 1\}$ determine the target bit of the single qubit gates A and B, respectively. The P_i implement permutations of the basis vectors realized by controlled-not operations.

We collect some general observations about implementations of controlled-U gates with at most two single qubit gates in Lemma 14–16. It is then shown in Lemma 17–19 that an implementation of a controlled-U operation with three controlled-not gates and at most two single qubit gates cannot exist, when $\operatorname{tr} U \neq 0$ and $\operatorname{tr}(UX) \neq 0$. The remaining cases are simple consequences of Lemma 14.

Lemma 14 *Let U be a unitary 2×2 matrix. Suppose that there exists an implementation of the controlled-U gate with two single qubit gates A and B, and some controlled-not gates. Then A is sparse if and only if B is sparse.*

Proof. The input $|00\rangle$ remains unchanged by a controlled-U operation. If A is sparse, then the state after P_2 is of the form $\alpha|b_1 b_0\rangle$, with $b_1, b_0 = 0, 1$. This state has to be mapped by B to a state of the form $|c_1 c_0\rangle$, with $c_1, c_0 = 0, 1$. Thus, B has to be sparse. The same argument applied to the inverse circuit shows that if B is sparse, then A has to be sparse as well. $\Box$

Lemma 15 *Let U be a unitary 2×2 matrix. Suppose that there exists an implementation of the controlled-U gate with two single qubit gates A and B, and some controlled-not gates. If $\operatorname{tr} U \neq 0$ and $\operatorname{tr}(UX) \neq 0$, then A and B cannot be sparse.*

Proof. If A and B are sparse, then the circuit (14) implements a monomial matrix. This would imply that U is sparse, contradicting either $\operatorname{tr} U \neq 0$ or $\operatorname{tr}(UX) \neq 0$. $\Box$

Denote by $c(P_i)$ the number of controlled-not gates used to realize the permutation P_i in (14).

Lemma 16 *Let U be a unitary 2×2 matrix. If $\operatorname{tr} U \neq 0$ and $\operatorname{tr}(UX) \neq 0$, then $c(P_2) > 0$ in the circuit (14).*

Proof. If $c(P_2) = 0$, then (14) implies

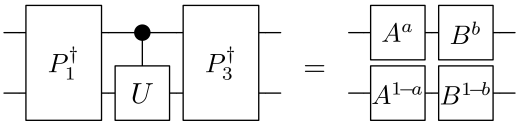

(15)

The state input $|00\rangle$ will not be changed by the circuit on the left hand side, because $P_1^\dagger$ and $P_3^\dagger$ are merely sequences of controlled-not gates. On the other hand, if $a \neq b$, then the circuit on the right hand side would map $|00\rangle$ to a superposition of base states, because A and B are not sparse. Thus, $a = b$, and BA has to be a diagonal matrix. However, this would imply that (14) is realizing a monomial matrix, which contradicts $\operatorname{tr} U \neq 0$ or $\operatorname{tr}(UX) \neq 0$. Therefore, $c(P_2)$ cannot be zero. $\Box$

Three Controlled-Not Gates. We assume now that three controlled-not gates are used in the circuit (14), that is, $c(P_1) + c(P_2) + c(P_3) = 3$ and $0 \leq c(P_i) \leq 3$. We may assume without loss of generality that $c(P_1) \geq c(P_3)$, for otherwise we can consider the inverse circuits. Consequently, $c(P_3)$ is either 0 or 1. In Lemma 18 we consider the case $c(P_3) = 1$, and in Lemma 19 the case $c(P_3) = 0$.

Lemma 17 *Let U be a unitary 2×2 matrix. If $\operatorname{tr} U \neq 0$ and $\operatorname{tr}(UX) \neq 0$, then $c(P_2) < 3$ in the circuit (14) with three controlled-not gates.*

Proof. Seeking a contradiction, suppose that $c(P_2) = 3$. The three controlled-not operations in P_2 must have alternating target qubits, because otherwise it would be possible to reduce the number of controlled-not gates. Therefore, the circuit (14) can be rewritten in the form:

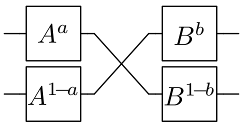

(16)

Clearly, this circuit is not able to realize a controlled-U operation, because it cannot entangle any unentangled input state. □

Lemma 18 *Let U be a unitary matrix with $\operatorname{tr} U \neq 0$ and $\operatorname{tr}(UX) \neq 0$. If $c(P_3) = 1$, then the circuit (14) with three controlled-not gates cannot implement a controlled-U operation.*

Proof. We have $c(P_1) \geq c(P_3)$ and $c(P_2) > 0$, hence $c(P_i) = 1$ for $i = 1, 2, 3$. The target bit of the controlled-not gate in P_1 (or P_3) cannot be the least significant bit, because this would imply that there exists an implementation of a controlled-XU (or a controlled-UX) gate with two controlled-not gates and at most four elementary gates. Proposition 13 shows that this is not possible. Therefore, the circuit has to be of the form:

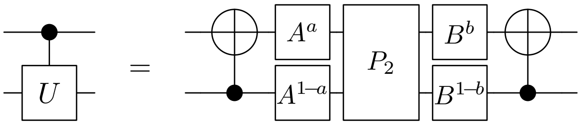

(17)

Regardless of the target bit of the controlled-not gate in P_2, we get

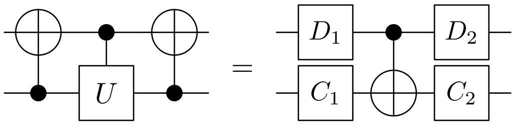

for some unitary matrices C_1, C_2, D_1, D_2. If we input $|10\rangle$, then the circuit on the left hand side produces an entangled output state. This means D_1 has to be nonsparse, and $C_1|0\rangle$ cannot be an eigenstate of X. The input $|00\rangle$ will not produce an entangled state in the circuit on the left hand side, but produces an entangled state on the right hand side. □

Lemma 19 *Suppose that U is a unitary 2×2 matrix satisfying $\operatorname{tr} U \neq 0$ and $\operatorname{tr} UX \neq 0$. If $c(P_3) = 0$ then the circuit (14) with three controlled-not gates cannot implement a controlled-U gate.*

Proof. Since $c(P_3) = 0$ and $c(P_2) = 1, 2$, we have $c(P_1) = 2, 1$, respectively. Therefore, the circuit is of the form

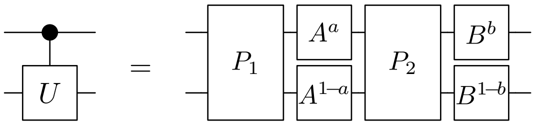

(18)

The permutation P_1 is of the form

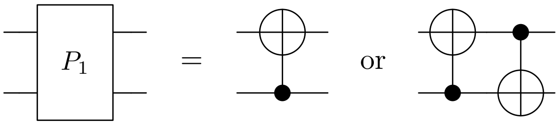

(19)

because otherwise it would be possible to realize a controlled-XU operation with less than five operations, which contradicts Proposition 13. Notice that

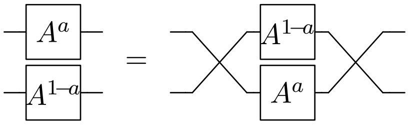

We will take advantage of this identity to derive the desired contradiction. Realizing that

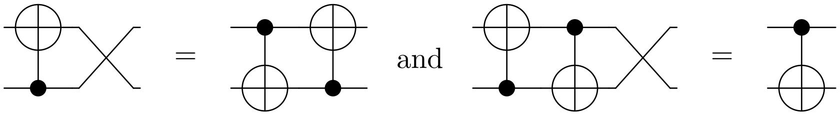

we find that the circuit (18) can be re-written in the form

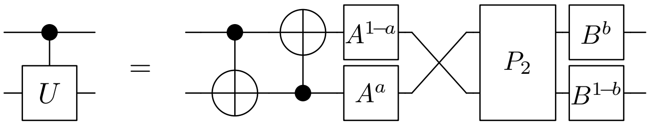

(20)

or

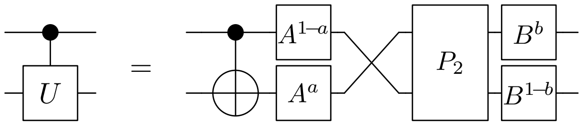

(21)

depending on the form of P_1 shown in (19), respectively. The circuits (20) and (21) can both be simplified to contain at most five elementary gates, by reducing the combination of the swap operation with P_2 to merely 1 and 2 controlled-not gates, respectively. This would imply that the controlled-XU gate can be realized with at most four elementary gates, contradicting Proposition 13. □

Proposition 20 *Let U be a unitary matrix with $\operatorname{tr} U \neq 0$ and $\operatorname{tr}(UX) \neq 0$. It is impossible to implement a controlled-U gate with two or fewer single qubit gates and three controlled-not gates.*

Proof. This follows from Lemma 18–19. □

Four or more Controlled-Not Gates. If we have more than three controlled-not gates, then an implementation with less than six elementary gates is not possible, because of Lemma 15.

Summary. We have shown that a generic controlled-U gate cannot be implemented with less than six elementary gates. In fact, we could rule out implementations based on a single controlled-not gate. Implementations with two controlled-not gates are possible, but at least four single qubit gates are necessary. The previous discussion showed that this gate count cannot be improved by implementations based on three or more controlled-not gates. This concludes the proof of Theorem A. □

3 Further Ramifications

Let $m(U)$ denote the minimal number of elementary gates that are needed to implement a controlled-U gate. We know from [1] that $m(U) \leq 6$. We have shown in the previous section that $m(U) = 6$ provided that U is generic. We will show next that $m(U) \leq 5$ when U is not generic.

Theorem B *Let U be a unitary 2×2 matrix. Let ϕ and ϕ_0 be real numbers in the range $0 < \phi < 2\pi$ and $0 \leq \phi_0 < 2\pi$.*
a) If $U = I$, then $m(U) = 0$.
b) If $U = e^{i\phi} I$, then $m(U) = 1$.
c) If $U = X$, then $m(U) = 1$.
d) If $U = e^{i\phi} X$, then $m(U) = 2$.
e) If $U = e^{i\phi_0} Z$, then $m(U) = 3$.
f) If $\operatorname{tr} U = 0$, $\det U = -1$, $U \neq \pm X$, then $m(U) = 3$.
g) If $\operatorname{tr} U = 0$, $\det U \neq -1$, $U \neq e^{i\phi_0} X$, $U \neq e^{i\phi_0} Z$, then $m(U) = 4$.
h) If $\operatorname{tr} UX = 0$, $\operatorname{tr} U \neq 0$, $\det U = 1$, $U \neq \pm I$, then $m(U) = 4$.
i) If $\operatorname{tr} UX = 0$, $\operatorname{tr} U \neq 0$, $\det U \neq 1$, $U \neq e^{i\phi_0} I$, then $m(U) = 5$.
j) If $\det U = 1$, $\operatorname{tr} U \neq 0$, $\operatorname{tr} UX \neq 0$, then $m(U) = 5$.

Theorem B captures all non-generic cases. The upper bounds on the number of gates are straightforward to see with the help of Table 1.

Case	Form	Circuit
if $\operatorname{tr} U = 0$	$U = e^{i\phi} P X P^\dagger$	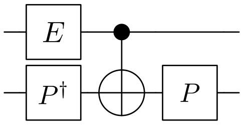 (C1)
else if $\operatorname{tr}(UX) = 0$	$U = e^{i\phi} P X P^\dagger X$	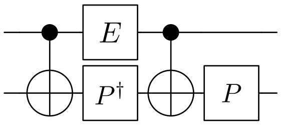 (C2)
else if $\det U = 1$	$U = CXBXA$	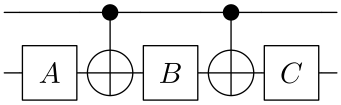 (C3)
else	$U = e^{i\phi} CXBXA$	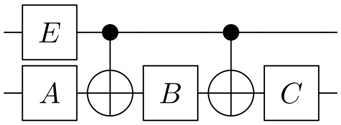 (C4)

Table 1: Quantum circuits for the implementation of controlled-U gates.

We formally prove tight upper bounds on $m(U)$ in the following simple Lemma.

Lemma 21 *The number $m(U)$ of elementary gates given in the statement of Theorem B are sufficient to realize the corresponding controlled-U gates.*

Proof. The cases *a)*–*c)* of Theorem B are obvious.

Case d) If $U = e^{i\phi}X$, then the circuit (C1) in Table 1 with $P = I$ and $E = \mathrm{diag}(1, e^{i\phi})$ implements a controlled-U gate, hence $m(U) \leq 2$.

Case e) If $U = e^{i\phi_0}Z$, then

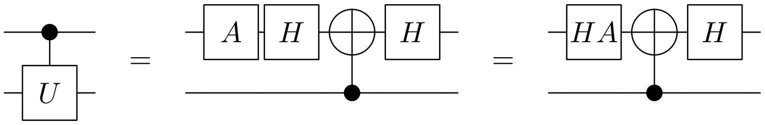

with $A = \mathrm{diag}(1, e^{i\phi_0})$, hence $m(U) \leq 3$.

Cases f, g) If $\mathrm{tr}\, U = 0$, then U is of the form $U = e^{i\phi}PXP^\dagger$. Therefore, circuit (C1) of Table 1 with $E = \mathrm{diag}(1, e^{i\phi})$ shows that $m(U) \leq 4$, which proves the upper bound of *g*). If in addition $\det U = -1$, then $E = I$ in (C1), whence $m(U) \leq 3$.

Cases h, i) If $\mathrm{tr}(UX) = 0$, then the matrix U is of the form $U = e^{i\phi}PXP^\dagger X$ for some unitary matrix P, and $\phi \in \mathbf{R}$. Let $E = \mathrm{diag}(1, e^{i\phi})$. The circuit (C2) in Table 1 shows that $m(U) \leq 5$, proving the upper bound of *i*). If in addition $\det U = 1$, then necessarily $E = I$, hence $m(U) \leq 4$, proving the upper bound of *h*).

Case *j*). If $\det U = 1$, then it is possible to realize the controlled-U gate in the form (C3), as was shown in [1], hence $m(U) \leq 5$. □

It remains to prove the lower bounds on $m(U)$. The following lemma allows to prove the cases a)–f):

Lemma 22 *Suppose that U is a unitary 2×2 matrix. If the controlled-U gate can be implemented with one single qubit gate A and some controlled-not gates, then U has to be of the form $U = e^{i\phi}I$ or $U = e^{i\phi}X$ for some $\phi \in \mathbf{R}$. Furthermore, A has to be diagonal and can be assumed to be of the form $A = \mathrm{diag}(1, e^{i\phi})$.*

Proof. The controlled-U gate is realized by a circuit of the form

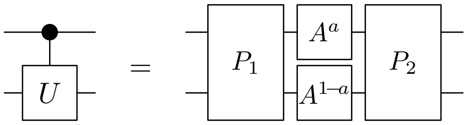

where P_1 and P_2 are permutations realized by controlled-not operations, and the target bit of A is selected by $a \in \{0, 1\}$. The state $|00\rangle$ remains invariant under the action of a controlled-U gate. Since $P_1|00\rangle = |00\rangle = P_2^\dagger|00\rangle$, it follows that $(A^a \otimes A^{1-a})|00\rangle = |00\rangle$, whence A has to be a diagonal matrix. By multiplying with an irrelevant global phase factor, we can assume that A is of the form $A = \operatorname{diag}(1, e^{i\phi})$. Notice that the phase of exactly two out of the four computational base states are changed to $e^{i\phi}$ by A, hence U has to be of the stated form. □

Lemma 23 *Parts a)–f) of Theorem B hold.*

Proof. It is clear that no gate is needed to implement a controlled-identity gate, whence *a*) holds. In *b*), only one phase gate is needed to affect the phase change, hence *b*) is true. At least one controlled-not gate is needed in the cases *c)–f)*, because U has two different eigenvalues a and $-a$. We have $m(U) \geq 2$ in cases *d)–f)*, because $U \neq X$. If $U = e^{i\phi}X$, then $m(U) = 2$ by Lemma 21, which proves *d*).

We have $U = e^{i\phi_0}Z$ in case *e)*. This gate can affect a phase change, hence at least one single qubit gate is needed. Lemma 22 shows that circuits with one single qubit gate and controlled-not gate cannot implement $U = e^{i\phi_0}Z$. Therefore, another single qubit gate is needed, that is, $m(e^{i\phi_0}Z) \geq 3$, whence $m(e^{i\phi_0}Z) = 3$ by Lemma 21.

In case f), U is a unitary matrix with $\operatorname{tr} U = 0$, $\det U = -1$, and $U \neq \pm X$. We know that $m(U) \geq 2$. Two controlled-not gates cannot implement such a gate because of the determinant condition. Lemma 22 shows that a single qubit gate and a controlled-not gate cannot implement U. Therefore, $m(U) \geq 3$, hence $m(U) = 3$ by Lemma 21. □

The remaining cases need a little bit more work. The next lemma gives some partial information about circuit with two single qubit gates and some controlled-not gates. Lemma 14 showed that A and B are either both sparse or both non-sparse. The sparse case is covered by the following lemma:

Lemma 24 *Suppose that U is a unitary 2×2 matrix. If the controlled-U gate can be implemented with two sparse single-qubit gates A and B, and some controlled-not gates, then the matrix U has to be of the form $U = e^{i\phi}I$, $U = e^{i\phi}X$, $U = \operatorname{diag}(e^{i\phi}, e^{-i\phi})$, or $U = \operatorname{antidiag}(e^{i\phi}, e^{-i\phi})$, where $\phi \in \mathbf{R}$.*

Proof. Since the matrices A and B are sparse, U has to be sparse as well, that is, $U = \operatorname{diag}(e^{i\phi_1}, e^{i\phi_2})$ or $U = \operatorname{antidiag}(e^{i\phi_1}, e^{i\phi_2})$.

Suppose that A or B is of the form $e^{i\phi}I$, $\phi \in \mathbf{R}$. Such a gate affects only a global phase change, and thus may be deleted. It follows from Lemma 22 that U is of the desired form.

Suppose that A and B are not multiples of the identity matrix. We may multiply A and B with global phase factors without changing the functionality of the circuit. Therefore, we can assume that A and B are of the form $A = X^a \operatorname{diag}(1, e^{i\phi_A})$ and $B = \operatorname{diag}(1, e^{i\phi_B})X^b$ for some $a, b \in \{0, 1\}$, and $0 < \phi_A, \phi_B < 2\pi$. Consider the four computational base states $\mathcal{B} = \{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$ as inputs to our circuit. The circuit realizes a monomial matrix, because A and B are sparse. Therefore, ignoring order, the output of these four input states is given by a set of four states $\{\alpha|00\rangle, \beta|01\rangle, \gamma|10\rangle, \delta|11\rangle\}$ with phase factors $\alpha, \beta, \gamma, \delta$. Since the circuit realizes a controlled-U operation with sparse U, the multiset of these four phase factors should be of the form $P_U = \{1, 1, e^{i\phi_1}, e^{i\phi_2}\}$.

Notice that A and B each affect a phase change in exactly two of the four computational base states. During the evolution of a input state from $\mathcal{B}$, the state might be multiplied by phases $e^{i\phi_A}$ and $e^{i\phi_B}$. If we record the combinatorial possibilities, then the multiset of phase factors $\{\alpha, \beta, \gamma, \delta\}$ can be of the form:

a) $\{e^{i\phi_A}, e^{i\phi_A}, e^{i\phi_B}, e^{i\phi_B}\}$, provided A and B affect disjoint inputs,

b) $\{1, e^{i\phi_A}, e^{i\phi_B}, e^{i(\phi_A+\phi_B)}\}$, provided a single input is affected by A and B,

c) $\{1, 1, e^{i(\phi_A+\phi_B)}, e^{i(\phi_A+\phi_B)}\}$, provided A, B affect the same two inputs.

We compare these multisets with P_U to derive some constraints about U. It is clear that case a) cannot occur, because $e^{i\phi_A}, e^{i\phi_B} \neq 1$. In case b), we necessarily have $\phi_A = -\phi_B$, therefore U is of the form $U = \operatorname{diag}(e^{i\phi_A}, e^{-i\phi_A})$ or $U = \operatorname{antidiag}(e^{i\phi_A}, e^{-i\phi_A})$. In case c), it follows that U is of the form $U = e^{i\phi}I$ or $U = e^{i\phi}X$ with $\phi = \phi_A + \phi_B$. □

Remark 25 *Suppose that a unitary matrix U is of the form $U = e^{i\phi}X$, $U = e^{i\phi}I$, $U = \operatorname{diag}(e^{i\phi}, e^{-i\phi})$, or $U = \operatorname{antidiag}(e^{i\phi}, e^{-i\phi})$. Notice that*
i) if $\operatorname{tr}(U) = 0$, then $U = e^{i\phi}X$, $U = \pm iZ$, or $U = \operatorname{antidiag}(e^{i\phi}, e^{-i\phi})$,
ii) if $\operatorname{tr}(U) \neq 0$, then $U = e^{i\phi}I$, or $U = \operatorname{diag}(e^{i\phi}, e^{-i\phi})$.

The following lemma proves case *g)* of Theorem B:

Lemma 26 *Suppose that U is a unitary 2×2 matrix. If $\operatorname{tr}(U) = 0$, $\det(U) \neq -1$, $U \neq e^{i\phi}X$, and $U \neq e^{i\phi}Z$, then $m(U) = 4$.*

Proof. Seeking a contradiction, we assume that $m(U) \leq 3$. The condition $\operatorname{tr} U = 0$ implies that $U \neq e^{i\phi}I$, i.e., at least one controlled-not gate is needed in the implementation.

Suppose that at least two controlled-not gates are used in the implementation. This means that at most one single qubit gate can be used. Lemma 22 shows that U would have to be of the form $U = e^{i\phi}X$, contradicting the assumptions. Therefore, the potential implementation of U must have one controlled-not gate.

Suppose now that one controlled-not gate and at most two single qubit gates A and B are used in the implementation. The matrices A and B are either both sparse or both not sparse by Lemma 14.

Case 1. Assume that A and B are sparse. Then U has to be sparse. According to Lemma 24 and Remark 25 the matrix U would have to be of the form $U = e^{i\phi}X$, $U = \pm iZ$ or $U = \operatorname{antidiag}(e^{i\phi}, e^{-i\phi})$. None of these matrices satisfies the assumptions of the lemma, contradiction.

Case 2. Suppose now that A and B are not sparse. If the controlled-not gate has the same target bit as the controlled-U gate, then the circuit is of the form (4). Lemma 3 shows that A_1 and A_2 are sparse. It follows that the circuit has to be of the form

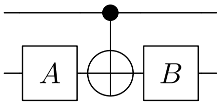

This implies $BA = e^{i\phi}I$ and $BXA = e^{i\phi}U$, whence $\det U = -1$, contradicting the assumptions.

Assume now that target bit of the controlled-not gate is the most significant bit. One easily sees that the two single qubit gates have to act on the most significant bit as well, i.e., the circuit is of the form

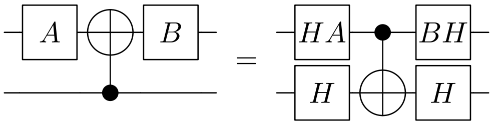

It follows from Lemma 3 that HA and BH are sparse, whence A and B are both not sparse. Notice that $|0\rangle$, $|1\rangle$ are eigenvectors of U, say with eigenvalues α_0, α_1, respectively. Corollary 2 shows that $BA = \operatorname{diag}(1, \alpha_0)$, and $BXA = \operatorname{diag}(1, \alpha_1)$. Comparing determinants shows that $\alpha_2 = -\alpha_1$. This implies that U is of the form $U = \operatorname{diag}(\alpha_1, \alpha_2) = \alpha_1 Z$, contradicting the assumptions.

Therefore, we can conclude that $m(U) \geq 4$. We obtain $m(U) = 4$ with Lemma 21. □

We proceed with the proof of case h) of Theorem B.

Lemma 27 *Let U be a unitary 2×2 matrix. If $\operatorname{tr}(UX) = 0$, $\operatorname{tr} U \neq 0$, $\det U = 1$, $U \neq \pm I$, then $m(U) = 4$.*

Proof. Seeking a contradiction, we assume that $m(U) \leq 3$. Lemma 4 shows that more than one controlled-not gate has to be used in the implementation of the controlled-U gate. We cannot have an implementation with three controlled-not gates, because the matrix corresponding to this circuit would have determinant -1. In the remaining case, one single qubit gate and two controlled-not gates are used for the implementation. Lemma 22 shows that a solution U with $\det U = 1$ would have to be of the form $U = \pm I$ or $U = -X$, all of which contradict the assumptions. We can conclude that $m(U) \geq 4$, hence $m(U) = 4$ by Lemma 21. □

The next lemma covers case *i)* of Theorem B.

Lemma 28 *Suppose that U is a unitary 2×2 matrix satisfying $\operatorname{tr}(U) \neq 0$, $\det U \neq 1$, $U \neq e^{i\phi} I$, and $\operatorname{tr}(UX) = 0$. Then five elementary gates are necessary and sufficient to implement such a controlled-U gate.*

Proof. Seeking a contradiction, we assume that $m(U) \leq 4$.

Case 1. Suppose that the implementation uses at least *three* controlled-not gates. According to Lemma 22, U would have to be of the form $U = e^{i\phi} I$ or $U = e^{i\phi} X$, which contradicts the assumptions $U \neq e^{i\phi}$ and $\operatorname{tr}(U) \neq 0$.

Case 2. Suppose that only one controlled-not gate is used in the implementation of the controlled-U gate. It follows that $\operatorname{tr} U = 0$ by Lemma 4. This contradicts the assumption $\operatorname{tr} U \neq 0$.

Case 3. Suppose that two controlled-not gates and at most two single qubit gates A and B are used in the implementation of the controlled-U gate. We know from Lemma 14 that A and B are either both sparse or both not sparse.

Case 3.1. Suppose that A and B are both sparse. It follows from Lemma 24 and Remark 25 that U has to be of the form $U = e^{i\phi} I$ or $U = \operatorname{diag}(e^{i\phi}, e^{-i\phi})$, which contradicts the assumptions $U \neq e^{i\phi} I$ and $\det U \neq 1$.
Case 3.2. Suppose that A and B are both not sparse. We distinguish four different cases, depending on the target bit of the two controlled-not gates.

Case $\downarrow\downarrow$. Suppose that the target bit of both controlled-not gates is the least significant bit. It follows from Lemma 6 and Lemma 7 that the single qubit gates have to act on the least significant bit as well. Assume without loss of generality that the control-U operation is implemented by a circuit of the form

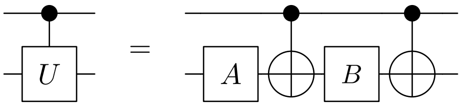

Comparing the result of the input $|0\rangle\otimes|\psi\rangle$ and $|1\rangle\otimes|\psi\rangle$ shows that $e^{i\theta}I = BA$ and $e^{i\theta}U = XBXA$. Comparing determinants yields $\det U = 1$, which contradicts our assumptions.

Case $\uparrow\uparrow$. Suppose that the most significant bit is the target bit of both controlled-not gates. There must be a single qubit gate, say B, on the most significant bit between the two controlled-not gates, because of Lemma 6. The other single quantum bit gate has to be on the most significant bit as well, in order to map $|00\rangle$ to $|00\rangle$.

Therefore, we may assume without loss of generality that the circuit is of the form

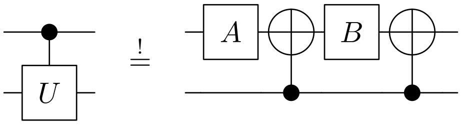

Notice that $|0\rangle$, $|1\rangle$ are eigenvectors of U, say with eigenvalues α_0, α_1, respectively. Corollary 2 shows that $BA = \text{diag}(1, \alpha_0)$, and $XBXA = \text{diag}(1, \alpha_1)$. Comparing determinants shows that $\alpha_0 = \alpha_1$. This implies that U is of the form $U = \text{diag}(\alpha_0, \alpha_1) = \alpha_0 I$, contradicting the assumptions.

Cases $\uparrow\downarrow$ and $\downarrow\uparrow$. Finally, consider the case that the two controlled-not gates have different target bits. We may assume that the first controlled-not gate has the most significant bit as target bit. If this is not the case, the we simply consider the inverse circuit. The circuit is of the general form

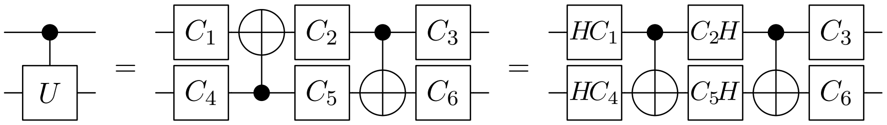

where two of the matrices C_i are given by A and B, and the remaining four are identity matrices. We use the circuit on the right hand side to derive a contradiction.

Lemma 7 and Lemma 9 show that $C_5H \neq H$, hence $C_5 \neq I$. Consequently, at least one of the matrices C_1 or C_2 has to be the identity matrix. Therefore, HC_1 or C_2H has to be nonsparse, whence C_3 is nonsparse by Lemma 6. It follows that $C_1 = C_2 = C_4 = C_6 = I$. However, we know from Lemma 10 that HC_1 and C_2H cannot both be equal to H, thus it is impossible that $C_1 = C_2 = I$, contradiction.

Therefore, we can conclude that it is impossible to implement a controlled-U gate with $m(U) \leq 4$ operations. It follows from Lemma 21 that $m(U) = 5$, which concludes the proof. □

It remains to show case *j)* of Theorem B:

Lemma 29 *Let U be a unitary 2×2 matrix. If $\det U = 1$, $\operatorname{tr} U \neq 0$, and $\operatorname{tr}(UX) \neq 0$, then $m(U) = 5$.*

Proof. It is not possible to implement such a controlled-U gate with only one controlled-not gate, cf. Lemma 4. If two controlled-not gates are used in the implementation, then Proposition 13 shows that three additional single qubit gates are necessary.

Assume that $m(U) \leq 4$ elementary gates are enough. If three or more controlled-not gates are used in the implementation, then at most one single qubit gate can be used. Lemma 22 shows that U would have to be of the form $U = e^{i\phi} I$ or $U = e^{i\phi} X$, contradicting $\operatorname{tr}(UX) \neq 0$ and $\operatorname{tr} U \neq 0$.

Therefore, $m(U) \geq 5$. It was shown in [1] that $m(U) \leq 5$ when $\det U = 1$, which proves the claim. □

4 Conclusions

We have derived the minimal number of elementary gates that are necessary in any implementation of a controlled unitary gate. It would be interesting to know tight lower bounds for other fundamental constructions of quantum circuits. In particular, it would be nice to know the minimal number of elementary gates that are needed to realize doubly controlled-U gates, such as the Toffoli gate.

Acknowledgments. We thank Martin Rötteler for numerous comments that helped to improve this paper.

Reference

[1] Adriano Barenco, Charles H. Bennett, Richard Cleve, David P. DiVincenzo, Norman Margolus, Peter Shor, Tycho Sleator, John A. Smolin, and Harald Weinfurter. Elementary gates for quantum computation. *Phys. Rev. A*, 52:3457–3467, 1995.