Quantum Information and Computation, Vol. 9, No. 5&6 (2009) 0461–0486

On THE CNOT-COST OF TOFFOLI GATES

VIVEK V. SHENDE[a]
Princeton University, Princeton, NJ 08544

IGOR L. MARKOV[b]
Department of EECS, University of Michigan, Ann Arbor, MI 48109

Received March 22, 2008
Revised January 2, 2009

The three-input TOFFOLI gate is the workhorse of circuit synthesis for classical logic operations on quantum data, e.g., reversible arithmetic circuits. In physical implementations, however, TOFFOLI gates are decomposed into six CNOT gates and several one-qubit gates. Though this decomposition has been known for at least 10 years, we provide here the first demonstration of its CNOT-optimality. We study three-qubit circuits which contain less than six CNOT gates and implement a block-diagonal operator, then show that they implicitly describe the cosine-sine decomposition of a related operator. Leveraging the canonical nature of such decompositions to limit one-qubit gates appearing in respective circuits, we prove that the n-qubit analogue of the TOFFOLI requires at least $2n$ CNOT gates. Additionally, our results offer a complete classification of three-qubit diagonal operators by their CNOT-cost, which holds even if ancilla qubits are available.

Keywords:
Communicated by: H-K Lo & R Laflamme

1 Introduction

The three-qubit TOFFOLI gate appears in key quantum logic circuits, such as those for modular exponentiation. However, in physical implementations it must be decomposed into one- and two-qubit gates. Figure 1 reproduces the textbook circuit from [14] with six CNOT gates, as well as Hadamard (H), $T = \exp(i\pi\sigma_z/8)$ and $T^\dagger$ gates.

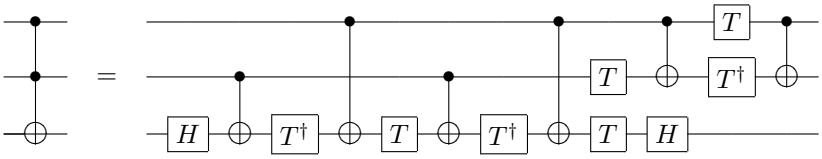

Fig. 1 Decomposing the TOFFOLI gate into one-qubit and six CNOT gates.

The pursuit of efficient circuits for standard gates has a long and rich history. DiVincenzo and Smolin found numerical evidence [4] that five two-qubit gates are necessary and sufficient to implement the TOFFOLI. Margolus showed that a phase-modified TOFFOLI gate admits a

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

three-CNOT implementation [6, 5], whose optimality was eventually demonstrated by Song and Klappenecker [20]. Unfortunately, this MARGOLUS gate can replace TOFFOLI only in rare cases. The detailed case analysis used in the optimality proof from [20] does not extend easily to circuits with four or five CNOTs. The omnibus Barenco et al. paper offers circuits for many standard gates, including an eight-CNOT circuit for the TOFFOLI [1, Corollary 6.2], as well as a six-CNOT circuit for the controlled-controlled-σ_z, which differs from the TOFFOLI only by one-qubit operators [1, Section 7]. Problem 4.4b of the textbook by Nielsen and Chuang asks whether the circuit of Figure 1 could be improved. The problem was marked as unsolved, and we report the following progress.

Theorem 1 *A circuit consisting of* CNOT *gates and one-qubit gates which implements the* n*-qubit* TOFFOLI *gate without ancillae requires at least* $2n$ CNOT *gates. For* $n = 3$*, this bound holds even when ancillae are permitted, and is achieved by the circuit of Figure 1.*

Our main tool is the Cartan decomposition in its "KAK" form, which provides a Lie-theoretic generalization of the singular-value decomposition [8]. Several special cases have previously proven useful for the synthesis and analysis of quantum circuits, notably the two-qubit *magic decomposition* [10, 11, 24, 23, 22, 16, 17], *the cosine-sine decomposition* [7, 2, 13, 18], and *the demultiplexing decomposition* [18]. The canonical nature of the two-qubit "magic" decomposition was used previously to perform CNOT-counting for two-qubit operators [16]. The magic decomposition is a two-qubit phenomenon,[c] but the cosine-sine and demultiplexing decompositions hold for n-qubit operators are similarly canonical. Moreover, the components of these decompositions are *multiplexors* [18] — block-diagonal operators that commute with many common circuit elements. Commutation properties facilitate circuit restructuring that can dramatically reduce the number of circuit topologies to be considered in proofs. These results and observations allow us to perform CNOT-counting using the Cartan decomposition in a divide-and-conquer manner.

In the remaining part of this paper, we first review basic properties of quantum gates in Section 2 and make several elementary simplifications to reduce the complexity of the subsequent case analysis. In particular, we pass from the CNOT and TOFFOLI gates to the symmetric, diagonal CZ and CCZ gates, and recall circuit decompositions which yield operators commuting with Z and CZ gates. We also define qubit-local CZ-costs, and observe that the total CZ-cost can be lower-bounded by half the sum of the local CZ counts for each qubit. Though weak, this bound suffices for our purposes and we can compute it in simple cases. Further technique is developped in Section 3, where we compute matrix entries to derive constraints on gates from circuit equations. This approach was employed by Song and Klappenecker in the two-qubit case, and we generalize several of their results to n-qubit circuits.

Section 4 is the heart of the present work, in which we prove our result on the CNOT-cost of the TOFFOLI gate. It starts by motivating and outlining the methods involved, previews key intermediate results, and proves that the CNOT-cost of the TOFFOLI is 6, based on these results. In Section 4.2, we use the canonicality of the cosine-sine decomposition derive circuit constraints. Section 4.1, motivated by [17], employs the canonicality of the demultiplexing decomposition, captured by a spectral invariant, to lower-bound CZ gates required in circuit

[c] While the Cartan decomposition $\mathrm{SU}(n) = \mathrm{SO}(n) \cdot [\text{diagonals}] \cdot \mathrm{SO}(n)$ is general, the utility of the magic decomposition arises from the isomorphism $\mathrm{SU}(2) \times \mathrm{SU}(2) \simeq \mathrm{SO}(4)$ being represented as an inner automorphism of $\mathrm{SU}(4)$. Such coincidental isomorphisms are few and confined to low dimensions.

implementations of operators. The results apply, *mutatis mutandis*, to CNOT-based implementations as well. Finally, in Section 4.3, we deduce as corollaries that the three-qubit PERES gate requires exactly 5 CNOTs and the n-qubit TOFFOLI gate requires at least $2n$. In Section 5, we extend our techniques to all three-qubit diagonal operators, completely classifying them according to CZ-cost. Generalizations to circuits with ancillae are obtained in Section 6. Concluding discussion can be found in Section 7.

2 Preliminaries

We review notation and properties of useful quantum gates, then characterize operators that commute with Pauli-Z gates on multiple qubits. We then review circuit decompositions from [3, 13, 18]. Finally, we introduce terminology appropriate for quantifying gate costs of unitary operators in terms of the CNOT and CZ and state elementary but useful observations about these costs.

2.1 Notation and properties of standard quantum gates

We write $\mathtt{X}, \mathtt{Y}, \mathtt{Z}$ for the Pauli operators, and $\mathtt{CX}, \mathtt{CCX}$ for $\mathtt{CNOT}, \mathtt{TOFFOLI}$. Rotation gates $\exp(i\mathtt{Z}\theta)$ are denoted by $R_z(\theta)$, and we analogously use R_x, R_y.[d] We work throughout on some fixed number of qubits N. For a one-qubit gate g and a qubit q, we denote by $g^{(q)}$ the N-qubit operator implemented by applying the gate g on qubit q. Similarly, $\mathtt{C}^{(i)}\mathtt{X}^{(j)}$ is the operator implemented by a controlled-X with the control on qubit i and target on qubit j. The controlled-Z being symmetric with respect to exchanging qubits, we do not distinguish control from target in the notation $\mathtt{CZ}^{(i,j)}$. We similarly denote the operator of a controlled-controlled-Z on qubits i, j, k by $\mathtt{CCZ}^{(i,j,k)}$. *In choosing qubit labels, we follow throughout the convention that the high-to-low significance order of qubits is the same as the lexicographic order of their labels.*

We follow the standard but sometimes confusing convention that *typeset operators act on vectors from the left*, but *circuit diagrams process inputs from the right.* Consistently with the established notation for the CNOT gate, we denote the X gate by "$\oplus$" in circuit diagrams. We denote the Z gate by a "$\bullet$" symbol, which does not lead to ambiguity in the matching notation for CZ because CZ is symmetric. Thus the following diagram expresses the identity $\mathtt{CZ}^{(\ell,m)}\mathtt{X}^{(\ell)} = \mathtt{Z}^{(m)}\mathtt{X}^{(\ell)}\mathtt{CZ}^{(\ell,m)}$ and rearranges gates in quantum circuits, like de Morgan's law does in digital logic.

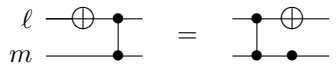

(1)

Another standard identity relates the X, Z, and one-qubit HADAMARD (H) gates: $\mathtt{HXH} = \mathtt{Z}$. By case analysis on control qubits, one obtains the further identities $\mathtt{H}^{(i)}\mathtt{C}^{(j)}X^{(i)}\mathtt{H}^{(i)} = \mathtt{CZ}^{(i,j)}$ and $\mathtt{H}^{(i)}\mathtt{CC}^{(j,k)}\mathtt{X}^{(i)}\mathtt{H}^{(i)} = \mathtt{CCZ}^{(i,j,k)}$. Despite this equivalence, we prefer the X family of gates for some applications and the Z family for others, as summarized in Table 1.

Circuits consisting entirely of one-qubit gates and CZ (respectively CNOT) gates will be called CZ-circuits (respectively CNOT-circuits). Using the above identities, CZ-circuits and CNOT-circuits can be interchanged at the cost of adding one-qubit H gates. It will also be

[d] We omit the factor of $\pm 1/2$ used by other authors.

	CNOT and TOFFOLI	CZ and CCZ
Advantages	With one-qubit gates added, either CNOT or CZ would be universal	
	Implement addition and multiplication	Symmetric
	Universal for reversible computation	Fewer circuit topologies
	Block-diagonal	Diagonal
	With 1-qubit diagonals, implement any diagonal	—
	Commute with X on target	Commute with Z on target
Other properties	Change direction after two H-conjugations	
	One can map back and forth by H-conjugation on target	
Applications	**Circuit synthesis**	**Circuit analysis**

Table 1 Relative advantages of standard controlled gates.

convenient to consider $\mathtt{CZ}^{(\ell)}$-circuits, which by definition are arbitrary circuits where all multi-qubit gates touching qubit ℓ are CZ. While these are not a subclass of CZ-circuits, a $\mathtt{CZ}^{(\ell)}$-circuit can be converted into a CZ-circuit without any changes affecting qubit ℓ.

2.2 Operators commuting with Z

An operator is diagonal if and only if it commutes with $\mathtt{Z}^{(\ell)}$ for all qubits ℓ. Similarly, Q commutes with Z on the highest order qubit if and only if

$$Q = \begin{bmatrix} Q_0 & 0 \\ 0 & Q_1 \end{bmatrix}$$

For ℓ any qubit, Q commutes with $\mathtt{Z}^{(\ell)}$ if and only if the (s,t)-th entry of the matrix of Q whenever is zero when the binary expansions of s and t differ at digit ℓ. It is more convenient to write this in the following way:

Observation 2 *Let Q be a unitary operator and ℓ be a qubit. The following are equivalent.*

1. *Q commutes with $\mathtt{Z}^{(\ell)}$*
2. $\langle 0|^{(\ell)} Q |1\rangle^{(\ell)} = 0$
3. $\langle 1|^{(\ell)} Q |0\rangle^{(\ell)} = 0$
4. *Q decomposes as a $Q = |0\rangle\langle 0| \otimes Q_0 + |1\rangle\langle 1| \otimes Q_1$, where the projectors $|i\rangle\langle i|$ operate on qubit ℓ and the unitary Q_i operate on the qubits other than ℓ.*
5. *Q decomposes as a product of a positively ℓ-controlled operator a negatively ℓ-controlled operator: $Q = \mathtt{C}^{(\ell)}(Q_1) \cdot \overline{\mathtt{C}}^{(\ell)}(Q_0)$.*

The Q_0, Q_1 of 4 and 5 are the same.

Whenever we have an operator commuting with $\mathtt{Z}^{(\ell)}$ we will employ subscripts as in Observation 2 to denote its diagonal blocks. More generally,

Notation. Fix qubits $\ell_1, \ldots, \ell_k$, and let Q commute with $\mathtt{Z}^{(\ell_i)}$ for all i. For any bitstring $j_1 \ldots j_k$ we write $Q_{j_1 \ldots j_k}$ for $\langle j_1 \ldots j_k|^{(\ell_1 \ldots \ell_k)} Q |j_1 \ldots j_k\rangle^{(\ell_1 \ldots \ell_k)}$. Note that $Q_{j_1 \ldots j_k}$ acts on k fewer qubits than Q.

For example, the two-qubit operator $U = |0\rangle\langle 0|^{(i)} \otimes Z^{(j)} + |1\rangle\langle 1|^{(i)} \otimes X^{(j)}$ commutes only with Z^i and hence we would write $U_0 = Z$ and $U_1 = X$. Below, we write the matrix of U in two bases; on the left, i is the most significant qubit, but on the right, j is.

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix} \quad \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & -1 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix}$$

$$i > j \qquad\qquad j > i$$

Observe that only the left matrix is block-diagonal.

In circuit diagrams, we indicate that a given operator commutes with $\mathtt{Z}^{(\ell)}$ by using a "control-on-box", as below. The backslash on the bottom line indicates that it represents an arbitrary number of qubits (a multi-qubit bus).

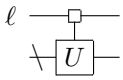

Unlike the notation for the positively and negatively controlled-U gates, the U above is not to be understood as acting only on the lower line. By way of illustration, we translate Observation 2, item 5, into a diagram:

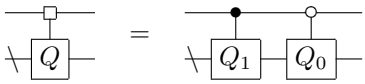

Observation 3 *Let Q, R be two gates such that for every qubit ℓ, either one of them does not affect ℓ, or both of them commute with $\mathtt{Z}^{(\ell)}$. Then $QR = RQ$. In picture:*

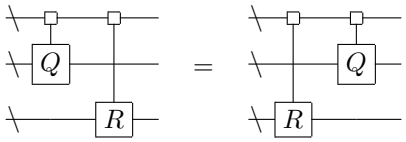

We now recall the *multiplexed rotation* gates [13, 18], which generalize the R_x, R_y, R_z gates. Let Δ be a diagonal Hermitian matrix acting on the qubits $\ell_1, \ldots, \ell_k$, and fix another qubit $m \neq \ell_i$. We define the operator $R_z^{(m)}(\Delta)$ on the qubits $\ell_1, \ldots, \ell_k, m$ by the conditions (1) that it commute with $\mathtt{Z}^{(\ell_i)}$ for all i, and (2) for any bitstring $j_1 \ldots j_k$, we have $R_z^{(m)}(\Delta)_{j_1 \ldots j_k} = R_z(\Delta_{\ell_1 \ldots \ell_k})$. Explicitly, $R_z^{(m)}(\Delta) = \exp(i\mathtt{Z}^{(m)}\Delta^{(\ell_1 \ldots \ell_k)})$. Multiplexed R_x, R_y gates are defined similarly. Since such operators commute with $\mathtt{Z}^{(\ell_i)}$, we depict them in circuit diagrams with the appropriate control-on-boxes.

It is natural to ask when an operator commuting with various $\mathtt{Z}$ gates can be implemented in a $\mathtt{CZ}$-circuit containing only gates commuting with the same $\mathtt{Z}$ gates. The answer is given in terms of the *partial determinant*.

Definition 1 *Fix qubits $\ell_1 \ldots \ell_k$, and let U be an operator commuting with $\mathtt{Z}^{(\ell_1)}, \ldots, \mathtt{Z}^{(\ell_k)}$. We define its* partial determinant $\det_{\ell_1 \ldots \ell_k}(U)$*, to be the diagonal operator acting on the qubits $\ell_1, \ldots, \ell_k$, whose $j_1 \ldots j_k$'th diagonal entry is* $\det(U_{j_1 \ldots j_k})$*. In short, it is given by the formula* $(\det_{\ell_1 \ldots \ell_k}(U))_{j_1 \ldots j_k} = \det(U_{j_1 \ldots j_k})$.

When computing partial determinants of a single gate or subcircuit acting on m qubits, we first tensor respective operators with $I_{2^{N-m}}$ to form operators acting on all N qubits (which may affect the determinants). When applied to such "full" operators, the partial determinant mapping is a group homomorphism.

For example, consider a CCZ gate in a three qubit circuit. Then for any two qubits i, j, we have $\det_{i,j}(\mathtt{CCZ}) = \mathtt{CZ}^{i,j}$. However, in a four qubit circuit, we $\det_{i,j}(\mathtt{CCZ}) = I$.

Proposition 4 *Fix qubits $\ell_1 \dots \ell_k$ among $N > k$ qubits. A unitary U commuting with $\mathtt{Z}^{(\ell_1)}, \dots, \mathtt{Z}^{(\ell_k)}$ can be implemented by a* CZ*-circuit in which only diagonal gates operate on qubits ℓ_i if and only if $\det_{\ell_1 \dots \ell_k}(U)$ is separable (can be implemented by one-qubit gates).*

Proof: ($\Rightarrow$). To establish the result in the forward direction, it suffices to show the separability of $\det_{\ell_1 \dots \ell_k}(U)$ for a generating set of operators. By definition, such a generating set is provided by CZs, one-qubit diagonals on the ℓ_i, and gates not affecting any of the ℓ_i.

Note first that any diagonal gate D acting on qubits $\ell_1, \dots, \ell_k$ has partial determinant given by $\det_{\ell_1 \dots \ell_k}(D) = D^{2^{N-k}}$, understood as an operator on qubits $\ell_1 \dots \ell_k$. In particular, if D were separable, then so is $\det_{\ell_1 \dots \ell_k}(D)$. If $D = \mathtt{CZ}^{(\ell_i, \ell_j)}$, then from $\mathtt{CZ}^2 = I$ and $N > k$ we deduce $\det_{\ell_1 \dots \ell_k}(\mathtt{CZ}^{(\ell_i, \ell_j)}) = I$. The remaining gates we need to consider are:
(i) any gate not affecting qubits ℓ_i implements $U = Q^{(1..N) \setminus (\ell_1 \dots \ell_k)}$ for some Q.
In this case $U_{j_1 \dots j_k} = Q$, and furthermore $\det_{\ell_1 \dots \ell_k}(U) = \det(Q) I$.
(ii) CZ gates connecting qubits $\ell_i, m \notin \{\ell_1, \dots, \ell_k\}$. We compute $\det_{\ell_1 \dots \ell_k}(\mathtt{CZ}^{(\ell_i, m)}) = (\mathtt{Z}^{(\ell_i)})^{2^{N-k-1}}$.

($\Leftarrow$). This part of the result is not used in the rest of the paper, and we therefore defer the proof to the Appendix. □

2.3 Cartan decompositions in quantum logic

This section recalls two important operator decompositions (*cosine-sine* and *demultiplexing*) and casts them as circuit decompositions. Readers willing to accept their use in our proofs may skip to Section 2.4.

Observe that an operator can be implemented with a single one-qubit gate if and only if it commutes with the Pauli operators Z and X on all other qubits. Thus to produce a CNOT-circuit for a given operator U, one may use the following algorithmic framework.

1. Decompose U into a circuit in which each non-CNOT gate, $V, W, \dots$, commutes with X and Z on more qubits that U does.

2. Apply the algorithm recursively to $V, W, \dots$ until one-qubit gates are reached.

As Z is self-adjoint, the requirement that U commutes with $\mathtt{Z}^{(i)}$ can be rephrased as the condition that U is fixed under the involution $U \mapsto \mathtt{Z}^{(i)} U \mathtt{Z}^{(i)}$. Given such an involution, a fundamental Lie-theoretic result produces an operator decomposition [8]. Here we recite the result for completeness, but do not require the reader to understand all terminology.

The Cartan Decomposition. Let G be a reductive Lie group, and $\iota : G \to G$ an involution. Let $K = \{g : \iota(g) = g\}$ and A be maximal over subgroups contained in $\{g : \iota(g) = g^{-1}\}$. Then K is reductive, A is abelian, and $G = KAK$.

In order to restate decompositions of unitary operators as circuit decompositions, we employ the notation of *set-valued* quantum gates [18]. Completely unlabelled gates (as in Equation 4) denote the set of all gates satisfying all control-on-box commutativity conditions imposed by the diagram, and gates labelled R_x, R_y, R_z denote the appropriate set of (possibly multiplexed) rotations. An equivalence of circuits with set-valued gates means that if we pick an element from each set on one side, there is a way to choose elements on the other so that the two circuits compute the same operator. The backslashed wires which usually indicate multiple qubits may also carry *zero* qubits.

The involution $\phi_Z : U \mapsto \mathtt{Z}^{(\ell)} U \mathtt{Z}^{(\ell)}$ corresponds to the *cosine-sine decomposition.*[e]

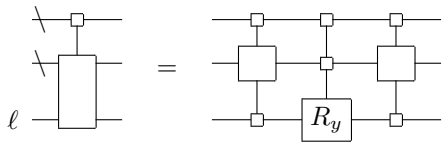

(2)

The involution $\phi_Y : U \mapsto \mathtt{Y}^{(\ell)} U \mathtt{Y}^{(\ell)}$ yields the *demultiplexing decomposition* [18].

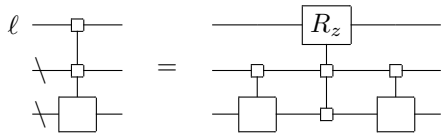

(3)

The map ϕ_Y restricts to the subgroup of diagonal operators. This group being abelian, the K and A factors commute, leaving the following decomposition of diagonal operators.

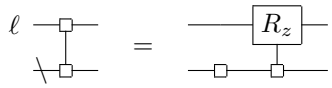

(4)

The involution ϕ_Y further restricts to the subgroup of multiplexed $\mathtt{Z}$ rotations, which we can demultiplex again. The K and A factors again commute; the A factor is computed by the last 3 gates in the circuit below.

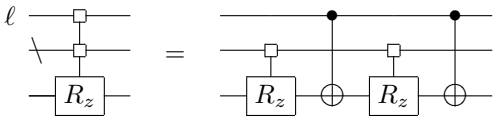

(5)

To establish the existence of these decompositions, it remains to verify in each case that the purported K and A satisfy the appropriate properties with respect to the relevant involution. This can be checked after passing to the Lie algebra where it is easy. Alternatively, explicit constructions of the cosine-sine and demultiplexing decompositions are given in [15] and [18], respectively.

To decompose general n-qubit operators, Equation 2 can be applied iteratively until all remaining gates are either multiplexed R_y gates or diagonal. The R_y gates can be replaced by R_z gates at the cost of introducing some one-qubit operators; the R_z and other diagonal gates can be decomposed as described above; for details and optimizations see [13]. Smaller

[e] The terminology comes from the numerical linear algebra literature; see [15] and references therein.

circuits are obtained by another algorithm, which alternates cosine-sine decompositions with demultiplexing decompositions; for details and optimizations, see [18].

When circuit decompositions are applied recursively, some gates can be reduced by local circuit transformations. For example, when iteratively demultiplexing multiplexed R_z gates, some CNOTs may be cancelled as shown below.

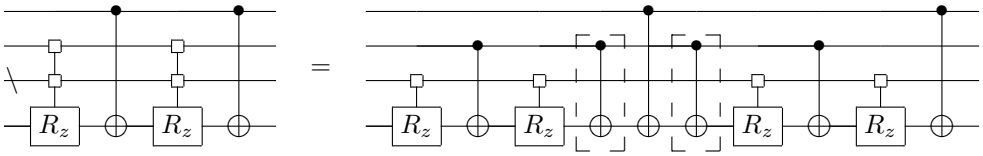

This technique produces a circuit with 2^n CNOT gates for an n-ply multiplexed R_z gate. Using Equation 4, we obtain a circuit with $2^n - 2$ CNOT gates for an arbitrary n-qubit diagonal operator [3]. Applying this result to CCZ gate leads to the circuit in Figure 1.

2.4 Basic facts about CZ*-counting*

The CZ-cost $|U|_{\mathtt{CZ}}$ of an N-qubit operator U is the minimum number of CZs which appear in any N-qubit CZ-circuit for U; we define the CNOT-cost analogously. The identity $\mathtt{H}^{(i)}\mathtt{C}^{(j)}X^{(i)}\mathtt{H}^{(i)} = \mathtt{CZ}^{(i,j)}$ ensures that $|U|_{\mathtt{CZ}} = |U|_{\mathtt{CNOT}}$. The further identity $\mathtt{H}^{(i)}\mathtt{CC}^{(j,k)}\mathtt{X}^{(i)}\mathtt{H}^{(i)} = \mathtt{CCZ}^{(i,j,k)}$ yields:

Observation 5 $|\mathtt{CCZ}|_{\mathtt{CZ}} = |\mathtt{CCX}|_{\mathtt{CNOT}} \leq 6$.

By way of illustration, the following modification of the circuit in Figure 1 implements the CCZ in terms of CZs.

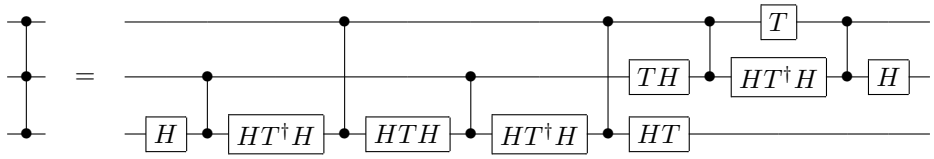

(6)

It shall prove more convenient to compute $|\mathtt{CCZ}|_{\mathtt{CZ}}$ rather than $|\mathtt{CCZ}|_{\mathtt{CNOT}}$. To do so, we are going to study the number of CZs which must touch a given qubit in any CZ-circuit for a given operator. More precisely, the $\mathtt{CZ}^{(\ell)}$-cost $|U|_{\mathtt{CZ};\ell}$ is the minimum number of CZ gates incident on ℓ in any $\mathtt{CZ}^{(\ell)}$-circuit for U. These cost functions are related through the following estimate.[f]

Observation 6 *For any operator P,*

$$|P|_{\mathtt{CZ}} \geq \frac{1}{2}\sum_j |P|_{\mathtt{CZ};j}$$

Proof: Each CZ gate touches two qubits. □

As the costs $|\mathtt{CCZ}|_{\mathtt{CZ};j}$ are the same for $j = 1, 2, 3$ (by symmetry),

$$|\mathtt{CCZ}|_{\mathtt{CZ}} \geq \frac{3}{2}|\mathtt{CCZ}|_{\mathtt{CZ};j} \tag{7}$$

[f] This bound is very weak in general. It can be shown using the results of [18] that $|U|_{\mathtt{CZ};\ell} < 6N$. Thus a direct application of the techniques developped here cannot yield a bound better than $|U|_{\mathtt{CZ}} \geq N(6N-1)$. On the other hand, dimension-counting shows that a generic N-qubit operator U requires on the order of 4^N CZ gates [9].

We emphasize that the number of qubits, N, is an unspecified parameter in both $|\cdot|_{\mathtt{CZ}}$ and $|\cdot|_{\mathtt{CZ};\ell}$. In the presence of ancillae, we define $|U|^a_{\mathtt{CZ}} := \min_t |U \otimes I_2^{\otimes t}|_{\mathtt{CZ}}$. Obviously $|U|^a_{\mathtt{CZ}} \leq |U|_{\mathtt{CZ}}$. While $|U|^a_{\mathtt{CZ}} = |U|_{\mathtt{CZ}}$ seems unlikely to always hold, we are not aware of any counterexamples. Indeed, we will show in Section 6 that this equality holds for all two-qubit operators and all three-qubit diagonal operators.

3 Deriving gate constraints from circuit equations

The circuit decompositions of Section 2.3 are essentially unique, and from one can derive various constraints on which gates may appear in certain circuit equations. We will pursue this route in Section 4.2. However, the simplest cases are easier to treat from the more elementary point of view adopted by Song and Klappenecker in their classification of two-qubit controlled-U operators by CNOT-cost [19]. Considering the operator computed by a candidate circuit, they first focus on matrix elements which vanish if the operator is a controlled-U. In order to produce such zero elements, the gates in the candidate circuit must satisfy certain constraints. Below we derive a series of more general results for n-qubit circuits. One-qubit gates which become diagonal when multiplied by X occur frequently; we refer to them as anti-diagonal.

Lemma 1 *The following equation imposes at least one of the following constraints.*

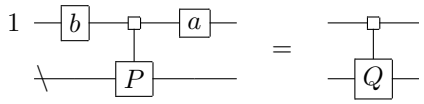

1. *a, b are both diagonal or both anti-diagonal.*
2. *P takes the form $d \otimes P_0$ for some one-qubit diagonal d.*

Proof: $0 = \langle 0|^{(1)} aPb\,|1\rangle^{(1)} = \langle 0|\,a\,|0\rangle\,\langle 0|\,b\,|1\rangle\,P_0 + \langle 0|\,a\,|1\rangle\,\langle 1|\,b\,|1\rangle\,P_1$. As the coefficients do not vanish, P_0 and P_1 are linearly dependent. It follows that $P = d \otimes P_0$ for some one-qubit diagonal d. □

Corollary 1 *If $a^{(i)}\mathtt{CZ}^{(i,j)}b^{(i)}$ commutes with $\mathtt{Z}^{(i)}$, then a, b are both diagonal or anti-diagonal.*

Corollary 2 *In the situation of Lemma 1, there exist one-qubit operators a', b' which are either diagonal or anti-diagonal, such that $a'^{(1)}Pb'^{(1)} = Q$.*

Proof: In Case 1, $a = a'$ and $b = b'$. In Case 2, write $P = d \otimes P_0$. Take $a' = adbd^{-1}$ and $b' = I$; then $a'^{(1)}Pb'^{(1)} = a^{(1)}Pb^{(1)}$. As $a'^{(1)} = QP^\dagger$ commutes with $\mathtt{Z}^{(1)}$, it is diagonal. □

We turn now to circuits with two CZ gates.

Lemma 2 *Suppose the following equation holds.*

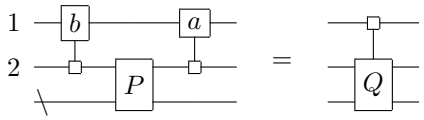

Then (I) $a_i b_j$ is diagonal for all i, j or (II) one of P, $\mathtt{X}^{(2)}P$ commutes with $\mathtt{Z}^{(2)}$.

Proof: We compute:

$$0 = \langle 0|^{(1)} \langle i|^{(2)} aPb\,|1\rangle^{(1)} |j\rangle^{(2)} = \langle 0|^{(1)} a_i b_j\,|1\rangle^{(1)} \langle i|^{(2)} P\,|j\rangle^{(2)}$$

Either $\langle i|^{(2)} P\,|j\rangle^{(2)} = 0$ for some i, j, or $\langle 0|\,a_i b_j\,|1\rangle$ vanishes for all i, j. □

Corollary 3 *Suppose the following equation holds.*

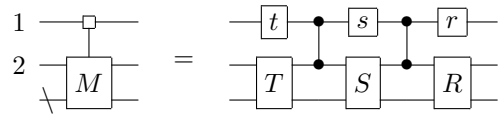

Then either (I) an even number of r, s, t are anti-diagonal, and the remainder diagonal, or (II) S or $S\mathtt{X}^{(2)}$ commutes with $\mathtt{Z}^{(2)}$.

Proof: In order to apply Lemma 2, We move R and T to the other side.

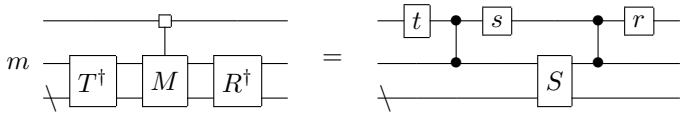

The cases here will correspond to the cases of Lemma 2. Case II is preserved verbatim. For Case I, the "$a_i b_j$" which must be diagonal are $rst, rs\mathtt{Z}t, r\mathtt{Z}st, r\mathtt{Z}s\mathtt{Z}t$. Since $(rst)^\dagger rs\mathtt{Z}t = t\mathtt{Z}t^\dagger$ is diagonal, we deduce that either t or $t\mathtt{X}$ is diagonal. Likewise, $r\mathtt{Z}st(rst)^\dagger = r\mathtt{Z}r^\dagger$ is diagonal, so either r or $r\mathtt{X}$ is diagonal. Finally, rst is diagonal, so from what we know about r, t, either s or $s\mathtt{X}$ is diagonal, and the number of r, s, t which are not diagonal is even. □

The following reformulation will be useful later.

Corollary 4 *Suppose Q commutes with $\mathtt{Z}^{(\ell)}$ and let $\mathcal{C}$ be a $\mathtt{CZ}^{(\ell)}$-circuit computing Q in which exactly two* CZ*s are incident on ℓ, say $\mathtt{CZ}^{(\ell,m)}$ and $\mathtt{CZ}^{(\ell,n)}$. Then all non-diagonal one-qubit gates may be eliminated from qubit ℓ at the cost of possibly (i) replacing $\mathtt{CZ}^{(\ell,n)}$ with $\mathtt{CZ}^{(\ell,m)}$ and (ii) adding one-qubit gates on qubits m, n.*

Proof: By hypothesis, $\mathcal{C}$ takes the form

$$Q = [r \otimes R]\mathtt{CZ}^{(\ell,m)}[s \otimes S]\mathtt{CZ}^{(\ell,n)}[t \otimes T]$$

where r, s, t are subcircuits of one-qubit operators acting on ℓ, and R, S, T are subcircuits containing no gates acting on ℓ. We immediately replace r, s, t by the one-qubit operators they compute. Moreover, if $m \neq n$, then replace S and T by $S \cdot \mathtt{SWAP}^{(m,n)}$ and $\mathtt{SWAP}^{(m,n)} \cdot T$, where SWAP is the gate which exchanges qubits. The swaps will be restored and canceled at the end of the proof. We are in the situation of Lemma 2.

Case I. We are done, with the exception that the r, s, t may be anti-diagonal rather than diagonal. In this case, Equation 1 allows the extraneous Xs to be pushed through and cancelled at the cost of introducing Z gates on qubit m. The diagonal gates remaining on qubit ℓ may be commuted through the CZs and conglomerated into one. Finally, the possible swap introduced between the S, T terms may be cancelled.

Case II. Using Equation 1 and replacing s by $s\mathtt{Z}$ if necessary, we commute S past one of the CZs. We now have:

$$Q = [r \otimes R]\mathtt{CZ}^{(\ell,m)}s^{(\ell)}\mathtt{CZ}^{(\ell,m)}[t \otimes ST]$$

Rearranging the equation,

$$[I \otimes R^\dagger]Q[I \otimes T^\dagger S^\dagger] = r^{(\ell)}\mathtt{CZ}^{(\ell,m)}s^{(\ell)}\mathtt{CZ}^{(\ell,m)}t^{(\ell)} \tag{8}$$

Let V be the value of either side of the equation above. Then from the LHS we see that V commutes with $Z^{(\ell)}$, and from the RHS we see that V is a two-qubit operator commuting with $Z^{(m)}$. Thus V is a two-qubit diagonal, and admits the following decomposition.

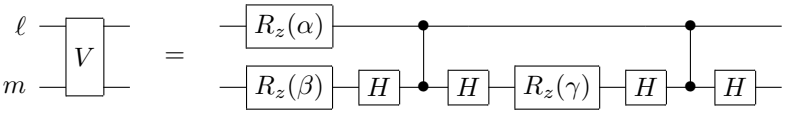

Substituting this decomposition for the RHS of Equation 8 and restoring the R, S, T gates completes the proof. □

4 The CNOT-cost of the TOFFOLI gate

So far we have reduced CNOT-counting for the TOFFOLI gate to CZ-counting for the CCZ gate, with the latter two being diagonal and symmetric. Having derived the inequality $3|\texttt{CCZ}|_{\texttt{CZ};\ell}/2 \leq |\texttt{CCZ}|_{\texttt{CZ}}$, we seek to determine the qubit-local costs $|\texttt{CCZ}|_{\texttt{CZ};\ell}$.

The idea is to find an equivalence relation $\sim_\ell$ such that (i) $U \sim_\ell V \implies |U|_{\texttt{CZ};\ell} = |V|_{\texttt{CZ};\ell}$ and (ii) the equivalence classes of $\sim_\ell$ are easy to characterize.

Definition 2 *For P, Q commuting with $\texttt{Z}^{(\ell)}$, we write $P \sim_\ell Q$ if there exist a, b, A, B satisfying the following equation.*

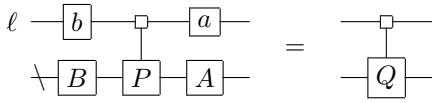

(9)

The fact that $|\cdot|_{\texttt{CZ};\ell}$ is constant on equivalence classes is obvious; the ability to characterize the equivalence classes comes from a comparison between Equation 9 and the demultiplexing decomposition of Equation 3. We construct invariants of the equivalence classes in Theorem 8. The reductions of Section 4.2 provide circuit forms on which the invariants are easy to compute; as a consequence, we arrive at a complete characterization of U such that $|U|_{\texttt{CZ};\ell} = 0, 1, 2$ in Theorem 9. The CCZ gate falls into none of these classes, and thus $|\texttt{CCZ}|_{\texttt{CZ};\ell} \geq 3$, and hence $|\texttt{CCZ}|_{\texttt{CZ}} \geq 5$. Unfortunately, qubit-local CZ-counting can take us no further: one can show by construction that in fact $|\texttt{CCZ}|_{\texttt{CZ};\ell} = 3$.

We now consider a hypothetical five-CZ circuit for the CCZ and seek a contradiction, using a divide-and-conquer strategy. There are many possible arrangements of the CZs, and we do not deal with them case by case. Nonetheless, we fix one here for clarity.

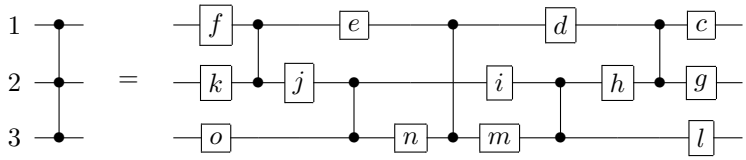

(10)

We define a, b, P, Q as follows.

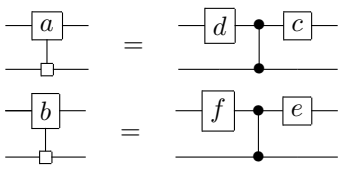

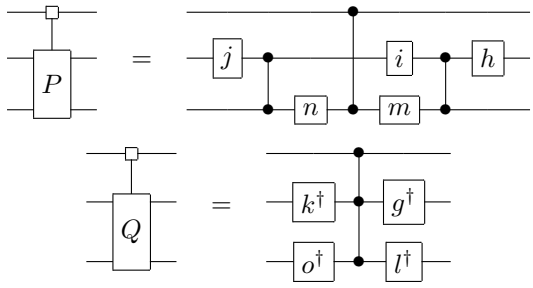

Our circuit decomposition now takes the following form.

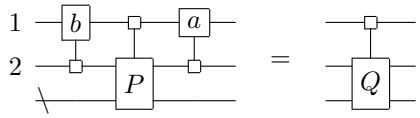

(11)

Up to some two-qubit diagonal fudge factors, this equation says that the cosine-sine decomposition of $b^\dagger \otimes I$ is $Q^\dagger[a \otimes I]P$. In Section 4.2, we translate the well-known canonicality of this Cartan decomposition into constraints on the components a, b, P and Q. The formulae of Theorem 9 further strengthen these constraints in the $|\cdot|_{\mathtt{CZ};\ell} = 3$ case. Specifically, we show in Theorem 10 that if $|U|_{\mathtt{CZ};\ell} = 3$ and $\mathcal{C}$ computes U using the minimum required three CZ gates incident on ℓ, then all one-qubit gates on ℓ are diagonal or anti-diagonal. The anti-diagonal gates can be made diagonal at the cost of introducing Z gates elsewhere in the circuit.

This is the last result needed to determine the CZ-cost of the CCZ. From $|\mathtt{CCZ}|_{\mathtt{CZ};\ell} \geq 3$, we see that in any five-CZ circuit for the CCZ, two of the qubits, m, n touch exactly three CZ gates and the remaining one touches four. By Theorem 10, we can assume all one-qubit operators on m, n are diagonal. Proposition 4 would then require $\det_{m,n} \mathtt{CCZ} = \mathtt{CZ}^{(m,n)}$ to be separable, which it is not.

Theorem 7 $|\mathtt{CCZ}|_{\mathtt{CZ}} = 6$.

We show in Section 6 that the use of ancillae can not lower the CZ-cost of the CCZ.

4.1 CZ *counting via the demultiplexing decomposition*

We now turn to the study of qubit-local CZ-cost. To apply $P \sim_\ell Q \implies |P|_{\mathtt{CZ};\ell} = |Q|_{\mathtt{CZ};\ell}$, we first seek to determine when $P \sim_\ell Q$. This will be done under the assumption that P and Q both commute with $\mathtt{Z}^{(\ell)}$.

Definition 3 *Let U commute with $\mathtt{Z}^{(\ell)}$. Then the ℓ-mux-spectrum $\Im^{(\ell)}(U)$ is the multi-set of eigenvalues, taken with multiplicity, of $U_1^\dagger U_0$. Two multi-sets S, T are said to be congruent, $S \cong T$, if there exists a nonzero scalar λ such that either $\lambda S = T$ or $\lambda S = T^\dagger$.*

We note that before taking the ℓ-mux-spectrum of U, it is necessary to fix the number of qubits on which U acts : $\Im^{(\ell)}(U \otimes I)$ contains $\dim I$ copies of $\Im^{(\ell)}(U)$.

Theorem 8 *Suppose P, Q commute with $\mathtt{Z}^{(\ell)}$. Then $P \sim_\ell Q \iff \Im^{(\ell)}(P) \cong \Im^{(\ell)}(Q)$.*

Proof: ($\Rightarrow$). As $P \sim_\ell Q$, there are gates a, b, A, B such that

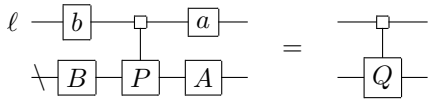

By Corollary 1, we may assume that either a, b or $a\mathtt{X}, b\mathtt{X}$ are diagonal. In the first case, $Q_0 = a_0 b_0 A P_0 B$ and $Q_1 = a_1 b_1 A P_1 B$. Thus $Q_1^\dagger Q_0 = (a_1 b_1)^\dagger a_0 b_0 B^\dagger P_1^\dagger P_0 B$, which has the same eigenvalues as $(a_1 b_1)^\dagger a_0 b_0 P_1^\dagger P_0$. Thus $\Im^{(\ell)}(P) \cong \Im^{(\ell)}(Q)$.

Otherwise, $a' = a\mathtt{X}$ and $b' = \mathtt{X}b$ are diagonal. Now $Q_0^\dagger Q_1 = (a_1' b_1')^\dagger a_0' b_0' B^\dagger P_0^\dagger P_1 B$, which has the same eigenvalues as $(a_1' b_1')^\dagger a_0' b_0' P_0^\dagger P_1$, whose eigenvalues in turn are the complex conjugates of those of $a_1' b_1' (a_0' b_0')^\dagger P_1^\dagger P_0$; again $\Im^{(\ell)}(P) \cong \Im^{(\ell)}(Q)$.

($\Leftarrow$). By supposition, the $\Im^{(\ell)}(P) \cong \Im^{(\ell)}(Q)$ We note $\Im^{(\ell)}(\mathtt{X}^{(\ell)} P \mathtt{X}^{(\ell)}) = \Im(P)^\dagger$ and $\Im((R_z^{(\ell)}(\lambda) P) = e^{2i\lambda}\Im(P)$. Therefore we can readily find an operator $P' \sim_\ell P$ such that the ℓ-mux-spectrum of P is identical, rather than merely congruent, to that of Q. It remains to show that $P' \sim_\ell Q$.

By the demultiplexing decomposition (Equation 3) there exist unitary operators M_P, N_P and a real diagonal matrix δ_P, all of which operate on the qubits other than ℓ, such that $P' = [I \otimes M_P] R_z^{(\ell)}(\delta_P)[I \otimes N_P]$. Likewise we decompose $Q = [I \otimes M_Q] R_z^{(\ell)}(\delta_Q)[I \otimes N_Q]$. If we let $\Delta_P = \exp(i\delta_P)$ and $\Delta_Q = \exp(i\delta_Q)$, then the ℓ-mux-spectra of P' and Q are respectively the entries of Δ_P^2 and Δ_Q^2. Since $\Im^{(\ell)}(P) = \Im^{(\ell)}(Q)$, there must exist a permutation matrix π acting on the qubits other than ℓ such that $\pi \Delta_P^2 \pi^\dagger = \Delta_Q^2$. Rearranging, we have $\Delta_Q^\dagger \pi \Delta_P = \Delta_Q \pi \Delta_P^\dagger$. Writing K for this term, $[I \otimes M_Q K M_P^\dagger] P' [I \otimes N_P^\dagger \pi^\dagger N_Q] = Q$. Thus $P' \sim_\ell Q$. □

We now apply Theorem 8 to prove the following result relating $\Im^{(\ell)}(P)$ and $|P|_{\mathtt{CZ};\ell}$. We emphasize that the number of qubits on which P acts is an unspecified parameter in both of these functions.

Theorem 9 *Let P commute with $\mathtt{Z}^{(\ell)}$.*

- $|P|_{\mathtt{CZ};\ell} = 0$ *iff* $\Im^{(\ell)}(P) \cong \{1, 1, \ldots\}$.
- $|P|_{\mathtt{CZ};\ell} = 1$ *iff* $\Im^{(\ell)}(P) \cong \{1, -1, 1, -1, \ldots\}$
- $|P|_{\mathtt{CZ};\ell} \leq 2$ *iff* $\Im^{(\ell)}(P)$ *is congruent to some multi-set S of unit norm complex numbers which come in conjugate pairs.*

Proof: The first and second statements follow immediately from Theorem 8 and the calculations $\Im^{(\ell)}(I) = \{1, 1, \ldots\}$ and $\Im^{(\ell)}(\mathtt{CZ}^{(\ell,m)}) = \{1, -1, 1, -1, \ldots\}$. To perform the relevant calculation for the third statement, we will use Corollary 4.

Let ℓ be the most significant qubit. For δ a diagonal real operator acting on all qubits but ℓ, define $\Phi(\delta)$ by

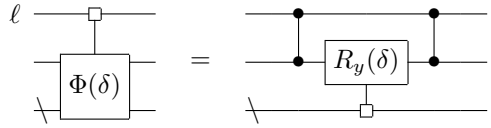

By construction, $|\Phi(\delta)|_{\mathtt{CZ};\ell} \leq 2$. We compute $\Im^{(\ell)}(\Phi(\delta)) = \{e^{2i\delta_0}, e^{-2i\delta_0}, e^{2i\delta_1}, e^{-2i\delta_1}, \ldots, \}$.

($\Leftarrow$) Write the entries of S as $e^{i\phi} \cdot \{e^{i\theta_0}, e^{-i\theta_0}, e^{i\theta_1}, e^{-i\theta_1}, \ldots\}$, and let θ be the real diagonal operator acting on all qubits but ℓ whose diagonal entries are $\theta_0, \theta_1, \ldots$. By construction, $\Im^{(\ell)}(\Phi(\theta/2)) = S$, and $S \cong \Im^{(\ell)}(Q)$ by hypothesis. By Theorem 8, $\Phi(\theta/2) \sim_\ell Q$ are ℓ-equivalent. It follows that $|Q|_{\mathtt{CZ};\ell} = |\Phi(\theta/2)|_{\mathtt{CZ};\ell} \leq 2$.

($\Rightarrow$) By hypothesis $|Q|_{\mathtt{CZ};\ell} \leq 2$. If in fact $|Q|_{\mathtt{CZ};\ell} = 0, 1$, note by the first two statements of the Theorem, which have been proven, the ℓ-mux-spectrum of Q has the desired property.

Thus we assume $|Q|_{\mathtt{CZ};\ell} = 2$. Let $\mathcal{C}$ be a circuit in which this minimal CZ count is achieved. By Corollary 4, we can find an equivalent circuit $\mathcal{C}'$ of the following form.

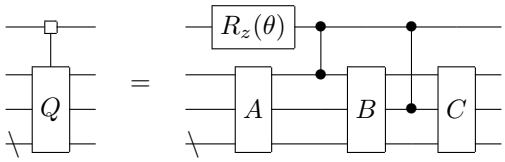

We have drawn the CZs with different lower contacts, but of course they might be the same. Actually, we prefer the latter case, and ensure it by incorporating swaps into B, C if necessary. We take a cosine-sine decomposition (see Equation 2) of B

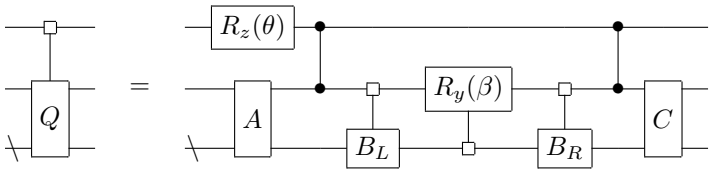

Note that the B_L and B_R gates commute with the CZs. Thus $Q \sim_\ell \Phi(\beta)$. By Theorem 8, the $\Im^{(\ell)}(Q) \cong \Im^{(\ell)}(\Phi(\beta))$. But we have already seen that $\Im^{(\ell)}(\Phi(\cdot))$ always consists of conjugate pairs of unit-norm complex numbers. □

4.2 Circuit constraints from the cosine-sine decomposition

This section is devoted to the study of Equation 11. We take cosine-sine decompositions of a, b. Below, A_l, A_r, B_l, B_r are two-qubit diagonal operators, and α, β are 2×2 real diagonal matrices of angular parameters.

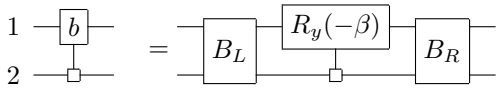

(12)

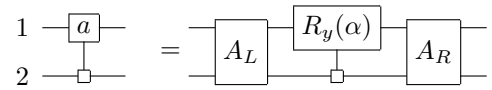

(13)

Define $\tilde{P} = A_L P B_R$ and $\tilde{Q} = A_R^\dagger Q B_L^\dagger$ to obtain:

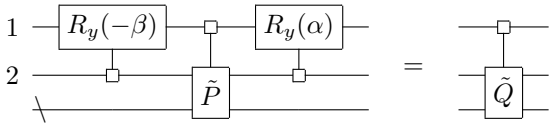

(14)

We recall the standard argument used to measure the uniqueness of the KAK decomposition [8]. Throughout this discussion, we will write simply $R_y(\alpha)$ for $R_y^{(1)}(\alpha^{(2)})$, and similarly for $R_y(\beta)$. Rearrange the equation to obtain $\tilde{Q}^\dagger R_y(\alpha)\tilde{P} = R_y(\beta)$. Transforming the equation by $k \mapsto \mathtt{Z}^{(1)} k^\dagger \mathtt{Z}^{(1)}$, we get $\tilde{P}^\dagger R_y(\alpha)\tilde{Q} = R_y(\beta)$. Multiplying these equations yields $\tilde{P}^\dagger R_y(2\alpha)\tilde{P} = R_y(2\beta)$. Thus $R_y(2\alpha)$ and $R_y(2\beta)$ have the same eigenvalues. One can check that in fact they are conjugate under an element of the group W generated by $\mathtt{X}^{(2)}$ and $\mathtt{CZ}^{(1,2)}$; note that these operators commute with $\mathtt{Z}^{(1)}$. That is, there exists $w \in W$ such that $wR_y(2\alpha)w^\dagger = R_y(2\beta)$. Now let $t = wR_y(\alpha)w^\dagger R_y(-\beta)$. We have both $t = R_y(\xi)$ for some

2×2 real diagonal matrix ξ acting on qubit 2, and $t^2 = I$; it follows that $t \in \{\pm I, \pm \mathtt{Z}^{(2)}\}$. Defining $\bar{P} = \tilde{P} \cdot [tw \otimes I]$ and $\bar{Q} = \tilde{Q} \cdot [w \otimes I]$ reduces our equation to the following.

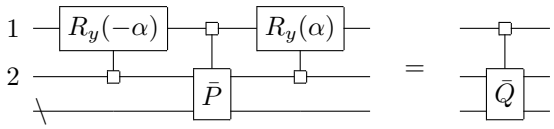

(15)

By an argument similar to that given for $\tilde{P}$ and $\tilde{Q}$, the operators $\bar{P}$ and $\bar{Q}$ both commute with $R_y(2\alpha)$. Conjugation by $R_y(\alpha)$ is an involution on the set of operators commuting with $R_y(2\alpha)$; Equation 15 says that P and Q are interchanged by this involution. In fact, this involution always has a simpler description:

Lemma 3 *Equation 15 also holds for some $\tilde{\alpha}$ for which $\tilde{\alpha}_i$ is an integer or half-integer multiple of π. Half-integers occur if and only if $2\alpha_i$ is an odd integer multiple of π.*

Proof: Decompose $2\alpha_i = \phi_i + \psi_i \pmod{2\pi}$ where $\phi_i \in (-\pi, \pi)$, where $\psi_i = 0$ unless $\phi_i = 0$, and $\psi_i \in \{0, \pi\}$ in any event. Then any operator which commutes with $R_y(2\alpha)$ also commutes with $R_y(\phi/2)$. Thus, on operators commuting with $R_y(2\alpha)$, conjugation by $R_y(\alpha)$ is the same as conjugation by $R_y(\alpha - \phi/2) = R_y(\alpha - \phi/2 - \psi/2)R_y(\psi/2)$. But $2(\alpha - \phi/2 - \psi/2) = 0 \pmod{2\pi}$. □

We also record the constraints imposed on possible $\bar{P}, \bar{Q}$ by the value of $\theta = 2\alpha$.

Lemma 4 *Fix distinct qubits ℓ, m. Let U be a unitary operator commuting with $\mathtt{Z}^{(\ell)}$, and let θ be a two-by-two real diagonal matrix of angular parameters which is understood to operate on m. Then U commutes with $R_y^{(\ell)}(\theta)$ if and only if one of the following holds:*

1. $\cos(\theta)$ *is scalar, and either*
   - *(a)* $\sin(\theta) = 0$.
   - *(b)* $\sin(\theta)$ *is a nonzero scalar and* $U_0 = U_1$.
   - *(c)* $\mathtt{Z}\sin(\theta)$ *is a nonzero scalar and* $U_0 = \mathtt{Z}^{(m)} U_1 \mathtt{Z}^{(m)}$.
2. $\cos(\theta)$ *is not scalar,* U *commutes with* $\mathtt{Z}^{(m)}$*, and either*
   - *(a)* $\sin(\theta_0) = 0$ *and* $\sin(\theta_1) = 0$.
   - *(b)* $\sin(\theta_0) = 0$ *and* $\sin(\theta_1) \neq 0$ *and* $U_{01} = U_{11}$.
   - *(c)* $\sin(\theta_0) \neq 0$ *and* $\sin(\theta_1) = 0$ *and* $U_{00} = U_{10}$.
   - *(d)* $\sin(\theta_0) \neq 0$ *and* $\sin(\theta_1) \neq 0$ *and* $U_0 = U_1$.

Proof: The ($\Leftarrow$) direction is trivial. For ($\Rightarrow$), suppose $[R_y^{(\ell)}(\theta^{(m)}), U] = 0$ and expand using the expression $R_y^{(\ell)}(\theta^{(m)}) = \exp(i\mathtt{Y}^{(\ell)}\theta^{(m)}) = \cos(\theta)^{(m)} + iY^{(\ell)}\sin(\theta)^{(m)}$ in order to observe that U_0 and U_1 both commute with $\cos(\theta)^{(m)}$, and $U_0 \sin(\theta)^{(m)} = \sin(\theta)^{(m)} U_1$. Now repeatedly apply the fact that two-by-two matrices which commute with a two-by-two diagonal matrix with distinct entries are themselves diagonal. □

Finally, we translate these results back to the original operators P, Q.

Lemma 5 *In the situation of Equation 11, at least one of the following must hold.*

1. *Either a, b are diagonal or $a\mathtt{X}^{(1)}, b\mathtt{X}^{(1)}$ are diagonal.*

2. *There exists a two-qubit operator U and two-qubit diagonals D, D' such that*

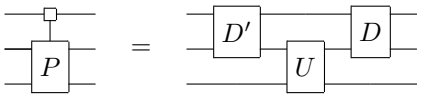

Similarly, there exists a two-qubit operator V and two-qubit diagonals C, C' such that

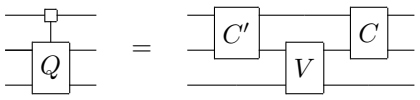

3. *Either P or $P\mathtt{X}^{(2)}$ commute with $\mathtt{Z}^{(2)}$. There exist replacements a', b' for a, b which are in the subgroup generated by two-qubit diagonal operators on qubits 1 and 2, $\mathtt{C}^{(2)}\mathtt{X}^{(1)}$, and $\mathtt{X}^{(1)}$, such that Equation 11 continues to hold.*

Proof: This amounts to unwinding the above discussion in light of Lemma 4. Case I comes from Case 1.a of the Lemma; the X appears because of the 2 in $\theta = 2\alpha$. Case II comes from Cases 1.b and 1.c. The first claim in Case III is just Case 2 of the Lemma; the possible X here comes from the w factor in $\bar{P} = \tilde{P}tw$ from the discussion above. The second claim follows from Lemma 3. □

While we cannot completely characterize operators with $|\cdot|_{\mathtt{CZ};\ell} = 3$, we can characterize $\mathtt{CZ}^{(\ell)}$-minimal circuits which compute them.

Theorem 10 *Fix a qubit ℓ, and suppose M commutes with $\mathtt{Z}^{(\ell)}$. Suppose $|M|_{\mathtt{CZ};j} = 3$, and let $\mathcal{C}$ be a $\mathtt{CZ}^{(j)}$-circuit exhibiting this bound. Then all one-qubit gates of $\mathcal{C}$ on ℓ are diagonal or anti-diagonal.*

Proof: Consider $M, \mathcal{C}$ satisfying the hypothesis. Without loss of generality, $\ell = 1$ and $\mathcal{C}$ takes the form

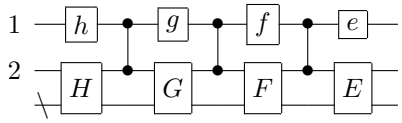

The CZs may have originally had different terminals, but we can incorporate swaps into E, F, G, H to suppress this behavior. This affects neither the hypothesis nor the conclusion.

(*) Define P by

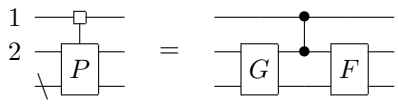

If $P\mathtt{X}^{(2)}$ commutes with $\mathtt{Z}^{(2)}$, then return to (*) and replace G by $G\mathtt{X}^{(2)}$, H by $\mathtt{X}^{(2)}H$, and h by $Z^{(1)}h$. This does not affect the conclusion, and by Equation 1, the resulting circuit still computes M. We have ensured that if one of $P, P\mathtt{X}^{(2)}$ commutes with $\mathtt{Z}^{(2)}$, then it is P.

Define a, b, Q by

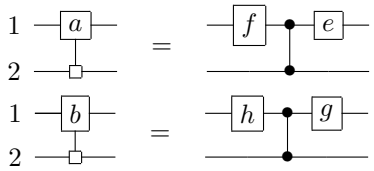

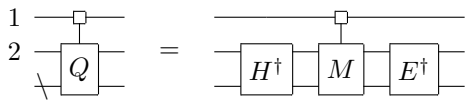

Note $|Q|_{\mathtt{CZ};1} = |M|_{\mathtt{CZ};1}$. We also have $Q = [a \otimes I]P[b \otimes I]$, hence are in the situation of Equation 11. Lemma 5 allows us to reduce to the following cases.

Case I. a, b are diagonal, or $a\mathtt{X}^{(1)}, b\mathtt{X}^{(1)}$ are diagonal. In either case, Corollary 1 applied to the circuits defining a, b shows that e, f, g, h are each diagonal or anti-diagonal.

Case II. Q takes the form

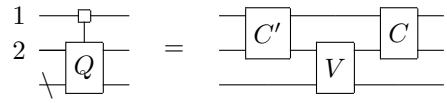

The cosine-sine decomposition (see Equation 2) of V along qubit 2 determines unitary operators R, S and a real diagonal operator δ such that:

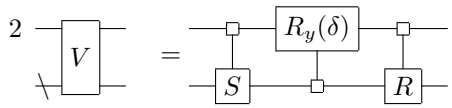

(16)

We substitute, commute the S, T outwards past C, C', and decompose the diagonals C, C'.

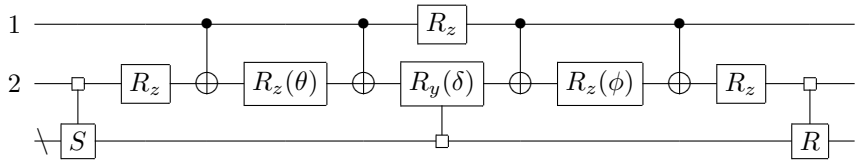

Evidently $\Im^{(1)}(Q)$ depends only on θ, δ, ϕ. We calculate that, up to a global scalar multiple, $\Im^{(1)}(Q)$ consists of the roots of the following quadratics in T:

$$T^2 - 2T(\cos(2\theta + 2\phi)\cos(\delta_i)^2 + \cos(2\theta - 2\phi)\sin(\delta_i)^2) + 1$$

The equations being real, each has complex conjugate roots. By Theorem 9, $|M|_{\mathtt{CZ};1} = |Q|_{\mathtt{CZ};1} = 2$, contrary to hypothesis.

Case III. We have already ensured that P, rather than $P\mathtt{X}^{(2)}$, commutes with $\mathtt{Z}^{(2)}$. We replace a, b by the a', b' of Lemma 5. We demultiplex P (see Equation 3) to obtain a decomposition of the following form, where D is diagonal.

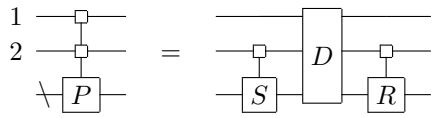

The operators S, R commute past a', b' to the edges of the circuit, and thus do not affect the $\mathtt{CZ}$-cost of Q. That is, $|Q|_{\mathtt{CZ};\ell} = |[a' \otimes I]D[b' \otimes I]|_{\mathtt{CZ};\ell}$.

By construction, $|P|_{\mathtt{CZ};\ell} = |D|_{\mathtt{CZ};\ell} = 1$. If $D = |0\rangle\langle 0|^{(\ell)} \otimes D_0 + |1\rangle\langle 1|^{(\ell)} \otimes D_1$, Theorem 9 asserts the entries of $D_0^\dagger D_1$ are $e^{i\theta}\{1, -1, 1, -1, \ldots\}$. Thus D can be written as

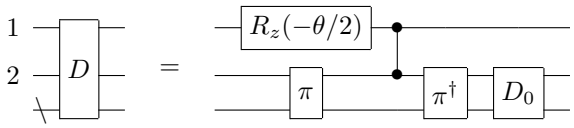

for some permutation π. We set $N := \mathtt{X}^{(1)}([a' \otimes I]D[b' \otimes I])^\dagger \mathtt{X}^{(1)}[a' \otimes I]D[b' \otimes I]$, so that $\Im^{(1)}([a' \otimes I]D[b' \otimes I])$ is given by the entries of $\langle 0|^{(1)} N |0\rangle^{(1)}$. Evidently D_0 commutes past a' and cancels with $D_0^\dagger$. Applying Equation 1 to eliminate X gates, the following circuit computes N.

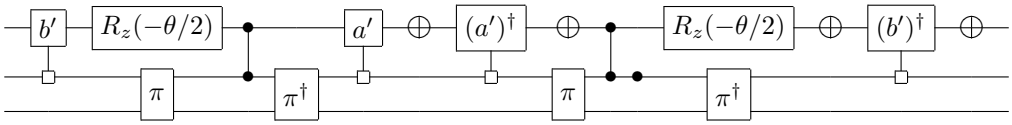

The condition on a' implies that $(a')^\dagger \mathtt{X}^{(1)} a' \mathtt{X}^{(1)}$ is diagonal. It follows that the subcircuit sandwiched between the two CZs computes a diagonal operator, and so the CZs cancel. Then the π, $\pi^\dagger$ pair on the left cancel. The $\pi^\dagger \mathtt{Z}^{(m)} \pi$ term on the right commutes past the $(b')^\dagger$. What remains is a circuit of the form

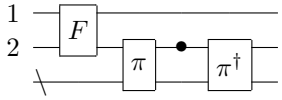

By construction, N commutes with both $\mathtt{Z}^{(1)}$ and $\mathtt{Z}^{(2)}$. It follows that F is diagonal. Then $f = \langle 0|^{(1)} F |0\rangle^{(1)}$ is some one-qubit diagonal acting on m. We have $\langle 0|^{(1)} N |0\rangle^{(1)} = \pi^\dagger \mathtt{Z}^{(2)} \pi f^{(2)}$. Denote by f_0, f_1 the entries of f. Then the entries of $\langle 0|^{(1)} N |0\rangle^{(1)}$ are $f_0, f_1, -f_0, -f_1$, and moreover f_0 will occur with the same multiplicity as $-f_1$; likewise $-f_0$ will occur with the same multiplicity as f_1. We see that $\sqrt{-f_0/f_1}\Im^{(1)}([a' \otimes I]D[b' \otimes I])$ come in conjugate pairs. By Theorem 9, $|[a' \otimes I]D[b' \otimes I]|_{CZ;1} \leq 2$. But now $|M|_{CZ;1} = |Q|_{CZ;1} = |[a' \otimes I]D[b' \otimes I]|_{CZ;1}$, contrary to hypothesis. □

4.3 Corollaries

The PERES gate implements a three-qubit transformation from classical reversible logic $\mathtt{PERES}^{(\ell;m;n)} = \mathtt{C}^{(\ell)}\mathtt{X}^{(m)} \cdot \mathtt{CC}^{(\ell,m)} X^{(n)}$. As shown in [12], it can be a useful alternative to the TOFFOLI gate in reversible circuits.

Corollary 5 $|\mathtt{PERES}|_{\mathtt{CZ}} = 5$.

Proof: As is clear from its definition, the PERES gate can be implemented by the circuit of Figure 1, save the rightmost CNOT. Thus, $|\mathtt{PERES}|_{\mathtt{CZ}} \leq 5$. On the other hand, it also follows from the definition that any circuit for the PERES can, with the addition of a single CNOT, become a circuit for the TOFFOLI. Thus $|\mathtt{PERES}|_{\mathtt{CZ}} \geq |\mathtt{TOFFOLI}|_{\mathtt{CZ}} - 1 = 5$, and all inequalities are equalities. □

In a different direction, we consider below multiply-controlled Z gates:

Corollary 6 $|(n-1)$*-controlled-*$\mathtt{Z}|_{\mathtt{CZ}} \geq 2n$ *for any* $n \geq 3$.

Proof: We proceed by induction on n. Suppose the Corollary is false; choose minimal falsifying n, and a falsifying circuit $\mathcal{C}$. By Theorem 7, $n > 2$. As before, at least three CZ gates are incident to each qubit, and counting shows that at least one, say ℓ touches exactly three. As before, we can assume that all one-qubit operators which appear on ℓ are diagonal. Form the circuit $\mathcal{C}' = \langle 1|^{(\ell)} \mathcal{C} |1\rangle^{(\ell)}$ by replacing every gate g of $\mathcal{C}$ with $g' = \langle 1|^{(\ell)} g |1\rangle^{(\ell)}$. This has no effect on gates which do not touch ℓ; it turns one-qubit gates on ℓ into scalars, and replaces $\mathtt{CZ}^{(\ell,s)}$ with $\mathtt{Z}^{(s)}$. At any rate, $\mathcal{C}'$ is a CZ-circuit on $(n-1)$ qubits which computes the $(n-2)$-controlled-Z. We deduce by induction that it contains at least $2(n-1)$ CZ gates. Adding the (at least) three CZs incident to ℓ, there are at least $2n+1$ total CZs in $\mathcal{C}$. □

5 Three-qubit diagonal operators

We give here a complete classification of three-qubit diagonal operators by their CZ-cost. Throughout this section, we assume no ancillae are available and label our qubits 1, 2, 3, from most significant to least significant. We abbreviate $\langle i|^{(1)} \langle j|^{(2)} \langle k|^{(3)} D |i\rangle^{(1)} |j\rangle^{(2)} |k\rangle^{(3)}$ by D_{ijk}. We also write $\Delta(\eta)$ for the one-qubit gate given by $|0\rangle\langle 0| + |1\rangle\langle 1|\,\eta$. Define

$$\lambda_1(D) = \frac{D_{011}D_{000}}{D_{001}D_{010}}, \qquad \lambda_2(D) = \frac{D_{101}D_{000}}{D_{100}D_{001}}, \qquad \lambda_3(D) = \frac{D_{110}D_{000}}{D_{100}D_{010}}, \qquad \xi(D) = \frac{D_{111}D_{000}^2}{D_{100}D_{010}D_{001}}$$

Then any three-qubit diagonal D admits the expansion

$$D = D_{000}\cdot\Delta\left(\frac{D_{100}}{D_{000}}\right)^{(1)}\cdot\Delta\left(\frac{D_{010}}{D_{000}}\right)^{(2)}\cdot\Delta\left(\frac{D_{001}}{D_{000}}\right)^{(3)}\cdot\mathrm{diag}(1,1,1,\lambda_1(D),1,\lambda_2(D),\lambda_3(D),\xi(D))$$

The $\lambda_i(D)$ are multiplicative, $\lambda_i(DD') = \lambda_i(D)\lambda_i(D')$, and likewise for ξ. We denote by $S(D)$ the ordered quadruple $(\lambda_1(D),\lambda_2(D),\lambda_3(D),\xi(D))$.

Observation 11 *For D, D' three-qubit diagonal operators, $S(D) = S(D')$ iff $S(D^\dagger D') = (1,1,1,1)$ iff $D^\dagger D'$ is a tensor product of one-qubit diagonal operators. It follows that $S(D) = S(D') \implies |D|_{\mathtt{CZ};i} = |D'|_{\mathtt{CZ};i}$.*

Observation 12 *$\Im^{(i)}(D) = \{1, \lambda_j(D)^\dagger, \lambda_k(D)^\dagger, \xi(D)^\dagger\lambda_i(D)\}$ where $\{i,j,k\} = \{1,2,3\}$.*

Lemma 6 *A three-qubit diagonal D can be implemented in a three-qubit CZ-circuit with:*

- *0 CZs on touching qubit 1 iff $S(D) = (\xi, 1, 1; \xi)$.*
- *1 CZ touching qubit 1 iff $S(D) = (\xi, -1, -1; \xi), (-\xi, 1, -1; \xi), (-xi, -1, 1\ \xi)$.*
- *2 CZs touching qubit 1 iff $S(D) = (a, b, c; abc), (a, b, c; ab/c), (a, b, c; ac/b)$.*

Proof: This is just a translation of Theorem 9 using Observation 12, involving a straightforward but tedious calculation which we omit. □

The two possibilities $S(D) = (a, b, c; abc), (a, b, c; ab/c)$ are quite different, and the following result helps distinguish between them.

Lemma 7 *Let D be a three-qubit diagonal operator and u be a one-qubit gate. Suppose $|Du^{(3)}\mathtt{CZ}^{(1,3)}|_{\mathtt{CZ};1} = 1$ or $|\mathtt{CZ}^{(1,3)}u^{(3)}D|_{\mathtt{CZ};1} = 1$. Then $\lambda_1(D)\lambda_2(D) = \lambda_3(D)\xi(D)$.*

Proof: The conclusion being stable under $D \to D^\dagger$, we assume $|Du^{(3)}\mathtt{CZ}^{(1,3)}|_{\mathtt{CZ};1} = 1$. Decompose $u^\dagger = e^{i\theta}R_z(\alpha)R_y(\beta)R_z(\gamma)$. Then $\Im^{(\ell)}(A)$ is given by the roots of the polynomials

$$x^2 - \cos(2\beta)(1 - \lambda_2(D))x - \lambda_2(D)$$

$$x^2 - \cos(2\beta)(\lambda_3(D) - \xi/(D)\lambda_1(D))x - \lambda_3(D)\xi(D)/\lambda_1(D)$$

For these to have roots either $\{p, p, -p, -p\}$ or $\{p, p, p, p\}$, the two equations must have the same constant terms – either both p^2 or both $-p^2$. □

We turn to computing CZ-costs. These being invariant under relabelling of qubits, we write $s(D)$ for $(\lambda_1(D),\lambda_2(D),\lambda_3(D);\xi(D))$, where we ignore the order of the λ_i.

Observation 13 *Given two three-qubit diagonals D, D', $s(D) = s(D')$ if and only if there exist one-qubit diagonals d, d', d'' and a wire permutation ω such that $D' = (d\otimes d'\otimes d'')\cdot\omega D\omega^\dagger$. Thus $s(D) = s(D') \implies |D|_{\mathtt{CZ}} = |D'|_{\mathtt{CZ}}$.*

Theorem 14 *Let D be a three-qubit diagonal operator. Then there exists a* CZ*-circuit for D containing*

- *0* CZ*s iff $s(D) = (1, 1, 1; 1)$.*
- *1* CZ *iff $s(D) = (1, 1, -1; -1)$.*
- *2* CZ*s iff $s(D) = (1, 1, \xi; \xi), (1, -1, -1; 1)$.*
- *3* CZ*s iff $s(D) = (1, 1, \xi; \xi), (\xi, -1, -1; \xi), (-\xi, 1, -1; \xi)$.*
- *4* CZ*s iff $s(D) = (a, b, c; ab/c)$.*
- *5* CZ*s iff $s(D) = (a, b, c; ab/c), (a, b, c; abc)$.*
- *6* CZ*s always.*

Proof: We assume without loss of generality that D takes the form $\mathrm{diag}(1, 1, 1, \lambda_1, 1, \lambda_2, \lambda_3, \xi)$. We number the qubits 1,2,3 from most to least significant.

($\Leftarrow$). We can assume that in fact $S(D)$ takes the form given. Our constructions will use the CX, which may be replaced by the CZ at the cost of inserting HADAMARD gates.

Case 0. $S(D) = (1, 1, 1; 1) \implies D = I$.

Case 1. $S(D) = (1, 1, -1; -1) \implies D = \mathtt{CZ}^{(1,2)}$.

Case 2a. $S(D) = (\xi, 1, 1; \xi)$. Fix $\eta = \sqrt{\xi}$;

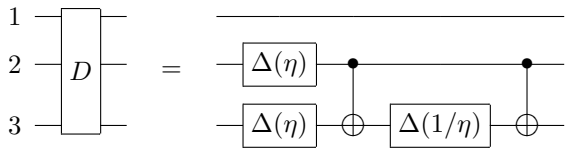

Case 2b. $S(D) = (1, -1, -1; 1) \implies D = \mathtt{CZ}^{(1,3)}\mathtt{CZ}^{(1,2)}$.

Case 3a. $S(D) = (\xi, 1, 1; \xi)$. By Case 2a, the CZ can be implemented in a circuit containing 2 CZs. It follows that any operator that can be implemented with $n > 0$ CZs can be implemented with $n+1$. Thus since D can be implemented with 2 CZs, it can be implemented with 3.

Case 3b. $S(D) = (\xi, -1, -1; \xi)$. Fix $\eta = \sqrt{\xi}$;

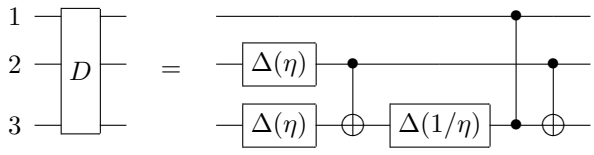

Case 3c. $S(D) = (-\xi, 1, -1; \xi)$. Fix $\eta = \sqrt{-\xi}$.

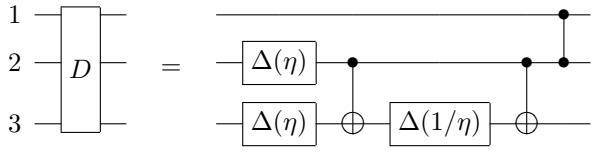

Case 4. $S(D) = (a, b, c; ab/c)$. Fix square roots α, β, γ for a, b, c;

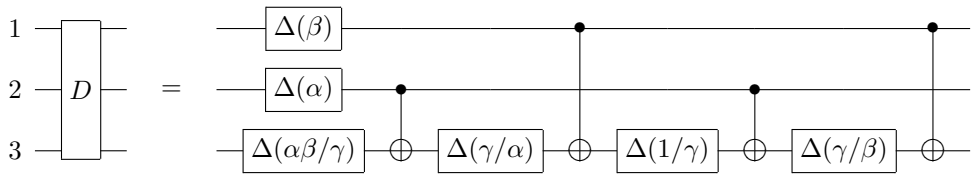

Case 5a. $S(D) = (a, b, c; ab/c)$. As D can be implemented with 4 CZs, it can be implemented with 5.

Case 5b. $S(D) = (a, b, c; abc)$. Fix square roots α, β, γ for a, b, c;

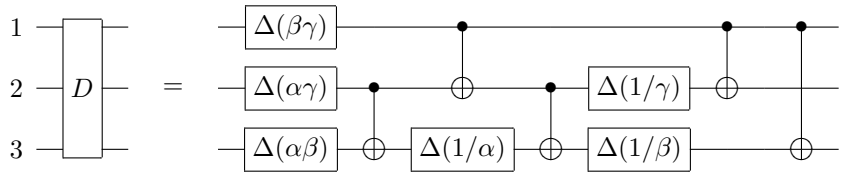

Case 6. More generally, any n-qubit diagonal operator has CZ-cost bounded by $2^n - 2$. See [3] or Section 2.3.

($\Rightarrow$).

Case 0. D must be locally equivalent to I, hence $s(D) = (1, 1, 1; 1)$.

Case 1. D must be locally equivalent to some CZ, hence $s(D) = (1, 1, -1; -1)$.

Case 2 Suppose there exists a minimal implementation of D in which both CZ gates connect the same two qubits. Then D is locally equivalent to a two-qubit diagonal; in which case one can compute $s(D) = (\xi, 1, 1; \xi)$.

Otherwise, there is a minimal implementation of D in which the two CZ gates are $\mathtt{CZ}^{(i,j)}$, $\mathtt{CZ}^{(j,k)}$. By Corollary 4, we may pass to an implementation with only diagonal one-qubit gates along j; by Corollary 2, we may pass to an implementation with only diagonal one-qubit gates along i, k as well. But then D is locally equivalent to $\mathtt{CZ}^{(i,j)}\mathtt{CZ}^{(j,k)}$ and we may compute $s(D) = (1, -1, -1; 1)$.

Case 3. It suffices to show that $|D|_{\mathtt{CZ};j} \leq 1$ for some j. For, if $|D|_{\mathtt{CZ};j} = 0$, then D is a two-qubit diagonal, with $s(D) = (\xi, 1, 1; \xi)$, and if $|D|_{\mathtt{CZ};j} = 1$, then by Lemma 6, $s(D) = (-\xi, 1, -1; \xi)$ or $(\xi, -1, -1; \xi)$.

Consider an implementation of D containing three CZs. We have $|D|_{\mathtt{CZ};\ell} \leq 1$ for some ℓ unless the CZs are distributed so that each qubit touches exactly two. Let j be a qubit touching the middle CZ. By Corollary 4, we can assume the circuit contains only diagonal gates on qubit j; it follows by inspection that $D \sim_j \mathtt{CZ}^{(i,j)}\mathtt{CZ}^{(j,k)}$. But we have already determined that $|\mathtt{CZ}^{(i,j)}\mathtt{CZ}^{(j,k)}|_{\mathtt{CZ};j} = 1$.

Case 4. Consider an implementation of D containing four CZs. If any qubit touches fewer than two CZs, we reduce to the previous case and observe that the desired condition on s holds. Thus suppose each qubit touches at least two CZs. Then there are only two possibilities for the number of CZs touched by each qubit: $(2, 2, 4)$ and $(2, 3, 3)$.

For the configuration $(2, 2, 4)$, say qubits ℓ, m touch two CZs and qubit n touches four. Note that no CZs connect ℓ, m. Thus we may assume by Corollary 4 all one-qubit gates on ℓ, m are diagonal. By Proposition 4, $\det_{\ell,m} D$ is separable; this says precisely that $\lambda_\ell(D)\lambda_m(D) = \lambda_n(D)\xi(D)$.

For the configuration $(2,3,3)$, say qubit 1 touches two CZs and qubits 2,3 touch three. Then there are two CZs connecting qubits 2 and 3, one connecting qubits 1 and 3 and one connecting qubits 1 and 2. By Corollary 4, we ensure that all one-qubit gates on qubit 1 are diagonal. If the CZs connecting qubits 2 and 3 are outermost, $D \sim_\ell \mathtt{CZ}^{(1,2)}\mathtt{CZ}^{(1,3)}$, hence can be implemented with three CZ s by Case 3. Otherwise, one of the CZs incident on qubit 1 is outermost; without loss of generality let it be $\mathtt{CZ}^{(1,3)}$. Then we have an equation of the form $D = u^{(3)}\mathtt{CZ}^{(1,3)}A$ where by construction A commutes with $\mathtt{Z}^{(1)}$ and $|A|_{\mathtt{CZ};1} = 1$. Lemma 7 yields the desired result.

Case 5. It suffices by Lemma 6 to show that $|D|_{\mathtt{CZ};\ell} \leq 2$ for some ℓ. Suppose not; then in any five-CZ implementation for D, each qubit must touch three CZs. It follows that two of the qubits, say ℓ, m touch exactly three CZs, and the remaining qubit touches four. By Theorem 10, all one-qubit gates on ℓ, m are diagonal or anti-diagonal. Enough applications of Equation 1 will ensure that all one-qubit gates on ℓ, m are in fact diagonal. Move the CZ which connects ℓ, m to the edge of the circuit. This yields $D = \mathtt{CZ}^{(\ell,m)}A$, where $|A|_{\mathtt{CZ};\ell} \leq 2$. By Lemma 6, it follows that $|D|_{\mathtt{CZ};\ell} \leq 2$ as well. □

6 Circuits with ancillae

The proofs of Theorems 7 and 14 assume that only three qubits were present, and use this assumption when enumerating possible circuit configurations with a given total number of CZ gates. This dependency can be eliminated. Indeed, these cases involved so few CZs that one could eliminate configurations with ancillae by performing explicit checks.

More significant is the use of Proposition 4 and the characterization by Theorem 9 of $|D|_{\mathtt{CZ};\ell} \leq 2$. Both of these statements are true for any fixed N, but suffer when N is allowed to vary. For example if only $N = 3$ qubits are available, then $\det_{1,2} \mathtt{CCZ}^{(1,2,3)} = \mathtt{CZ}^{(1,2)}$, so by Proposition 4, the CCZ cannot be implemented in any three-qubit CZ-circuit in which all gates commute with $\mathtt{Z}^{(1)}, \mathtt{Z}^{(2)}$. But if $N = 4$ qubits are present, $\det_{1,2}(\mathtt{CCZ}^{(1,2,3)}) = I^{(1,2)}$, so $\mathtt{CCZ}^{(1,2,3)} \otimes I^{(4)}$ *can* be implemented in a four-qubit CZ-circuit in which all one-qubit gates commute with $\mathtt{Z}^{(1)}$ and $\mathtt{Z}^{(2)}$.

Similarly, for $N = 3$ qubits, we have $\Im^{(\ell)}(\mathtt{CCZ}) = \{1, 1, 1, -1\}$ and thus by Theorem 9 $|\mathtt{CCZ}|_{\mathtt{CZ};\ell} \geq 3$. However, for $N = 4$ qubits, $\Im^{\ell}(\mathtt{CCZ}^{(\ell,m,n)}) = \{1, 1, 1, -1, 1, 1, 1, -1\}$, so now Theorem 9 implies that $|\mathtt{CCZ}^{(1,2,3)} \otimes I^{(4)}|_{\mathtt{CZ};1} = 2$. Indeed:

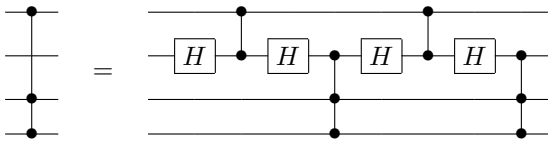

On the other hand, the properties $\Im^{(\ell)}(U) \cong \{1, 1, \ldots\}$ and $\Im^{(\ell)}(U) \cong \{1, -1, 1, -1 \ldots\}$ are stable under adding ancillae. By Theorem 9, so are the properties $|U|_{\mathtt{CZ};\ell} = 0$ and $|U|_{\mathtt{CZ};\ell} = 1$. Since only these properties are used in the proof of Lemma 7, it too holds even in the presence of ancillae. This leads to an extension of the CZ-cost classification of three-qubit diagonals to the case where ancilla qubits are permitted.

Lemma 8 *Let A be a unitary operator; let $\mathcal{C}$ be qubit minimal among* CZ*-circuits computing A, possibly with the use of ancillae, using only $|A|^a_{\mathtt{CZ}}$* CZ *gates. Then every ancilla in $\mathcal{C}$ touches at least three* CZ *gates.*

Proof: Fix an ancilla qubit ℓ. If no CZ gates touch ℓ, then it may be removed. If one (respectively two) CZ touches ℓ, then by Corollary 2 (respectively Corollary 4), then there is a circuit with no more CZs in which the only one-qubit gates on a are diagonal.

Now form the circuit $\langle 0|^{(\ell)} \mathcal{C} |0\rangle^{(\ell)}$ as in the proof of Corollary 6. This circuit computes the operator A using one fewer ancilla, fewer CZs than $\mathcal{C}$. □

Corollary 7 *For any two-qubit operator V, $|V|^a_{\mathtt{CZ}} = |V|_{\mathtt{CZ}}$.*

Proof: If no ancillae are needed to minimize CZ-count, then the result holds. Otherwise, each ancilla used in a qubit-minimal CZ-minimal implementation must touch at least three CZgates. Thus $|\cdot|_{\mathtt{CZ}} \geq |\cdot|^a_{\mathtt{CZ}} \geq 3$. However it is known [23, 22, 16] that two-qubit operators have $|\cdot|_{\mathtt{CZ}} \leq 3$. Thus all the inequalities are equalities. □

Proposition 15 *For any three-qubit diagonal operator, D, $|D|^a_{\mathtt{CZ}} = |D|_{\mathtt{CZ}}$.*

Proof: Suppose $|D|^a_{\mathtt{CZ}} < |D|_{\mathtt{CZ}}$. By Lemma 8, a qubit-minimal circuit for D achieving the bound for $|D|^a_{\mathtt{CZ}}$ contains at least three CZ gates incident on each ancilla. By assumption at least one ancilla is used, so $|D|_{\mathtt{CZ}} > |D|^a_{\mathtt{CZ}} \geq 3$. It follows from Theorem 14 and Lemma 6 that $|D|_{\mathtt{CZ};\ell} > 1$ for the three qubits $\ell = 1, 2, 3$. By Theorem 9, this property is stable under addition of ancilla. Thus a qubit-minimal circuit for D achieving the bound for $|D|^a_{\mathtt{CZ}}$ contains at least 3 CZs incident to each ancilla, and at least 2 CZs incident to each non-ancilla qubit. If k ancillae are used, then we have $|D|^a_{\mathtt{CZ}} \geq (3k+6)/2$. From Theorem 14 and the supposition we have $|D|^a_{\mathtt{CZ}} < |D|_{\mathtt{CZ}} = 6$; it follows that $k = 1$, that $|D|^a_{\mathtt{CZ}} = 5$, and that $|D|_{\mathtt{CZ}} = 6$.

In any four-qubit, five-CZ circuit for D, we must have two of the non-ancilla, say x_1, x_2 touching two CZs, and both the remaining non-ancilla z and the ancilla a touching three. By Corollary 4, we can assume that the only one-qubit operators appearing on x_1, x_2 are diagonal. We may also assume that the graph where vertices are qubits and edges are CZ gates is connected; otherwise D could be split into the tensor product of a two-qubit and a one-qubit diagonal, and hence would have $|D| \leq 2$. Then there are only three possibilities regarding which wires are connected by CZs.

I	(x_1, x_2)	(x_1, z)	(x_2, a)	(z, a)	(z, a)
II	(x_1, z)	(x_1, z)	(x_2, a)	(x_2, a)	(z, a)
III	(x_1, z)	(x_1, a)	(x_2, z)	(x_2, a)	(z, a)

We will show that any circuit with those CZ gates can be transformed so that (i) one of the outermost CZgates does not touch the ancilla, and (ii) one of the x-qubits on which this CZ gate acts has the property that all one-qubit gates acting on it are diagonal. As this x-qubit only touched 2 CZ gates to begin with, it follows from Lemma 7 that $s(D)$ takes the form $(a, b, c; ab/c)$. By Theorem 14, $|D|_{\mathtt{CZ}} = 4$, which is a contradiction.

We return to checking (i) and (ii). Eliminate non-diagonal one-qubit gates on x_i using Corollary 4. In Case (I), the (x_1, x_2) CZ can therefore only be prevented from moving by the (x_1, a). This can be on only one side, so the (x_1, x_2) can be moved outwards to the other. Similarly, in Case (II), an (x, z) can only be blocked by (z, a) and the other (x, z). In this case, the second (x, z) is blocked on only one side and can be moved to the edge. In Case (III), we use Corollary 4 to clear both the x_1 and x_2 qubits of non-diagonal gates; the possible additional one-qubit gates will only fall on the z and a qubits. Now the (x_1, z) can only be blocked by the (x_2, z) and the (z, a), and also the (x_2, z) can only be blocked by (z, a) and (x_1, z). Thus one of (x_1, z) and (x_2, z) can be made outermost. □

Corollary 8 $|\mathtt{CCZ}|^a_{\mathtt{CZ}} = |\mathtt{TOFFOLI}|^a_{\mathtt{CZ}} = 6$ *and* $|PERES|^a_{\mathtt{CZ}} = 5$.

7 Conclusion

While our work is primarily focused on quantum circuit implementations, the TOFFOLI gate originally arose as a universal gate for classical reversible logic [21]. In contrast, the NOT and CNOT gates are *not* universal for reversible logic: their action on bit-strings is affine-linear over F_2, and thus the same is true for any operator computed by any circuit containing only these gates.

Augmenting CNOT gates with single-qubit rotations to express the TOFFOLI gate provides the lacking non-linearity. Thus the number of one-qubit gates (excluding inverters) needed to express the TOFFOLI, or more generally any reversible computation, can be thought of as a measure of its non-linearity. In this inverted cost model (also relevant to some quantum implementation technologies) the following question remains open: *how many one-qubit gates are needed to implement the* TOFFOLI*?* Furthermore, *are there circuits that simultaneously minimize the number of* CNOT *and one-qubit gates ?*

In a different direction, recall our results showing that diagonality and block-diagonality of an operator impose strong constraints on small circuits that compute this operator. We believe other conditions may act in a similar way. In particular, we ask *what can be said about minimal quantum circuits for operators computable by classical reversible circuits, i.e., operators expressed by 0-1 matrices?* Very little is known even for three-qubit operators. In particular, the CNOT-cost of the controlled-swap (Fredkin gate) remains unresolved.

Closest to our present work, the exact CNOT-cost of the n-qubit analogue of the TOFFOLI gate remains unknown. We have shown that $2n$ CNOTs are necessary if ancillae are not permitted, but already for $n = 4$ we only know that $8 \leq |\mathtt{CCCZ}|_{\mathtt{CZ}} \leq 14$, where the upper bound is provided by a generic decomposition of diagonal operators [3]. Existing constructions of the n-qubit TOFFOLI gate require a quadratic number of CNOT gates without the use of ancillae. With one ancilla, such constructions require linearly many CNOTs, but the leading coefficient is in double-digits [1, 12].

Finally, we hope that our proof can be simplified and our techniques generalized. In particular, we have relied on repeated comparisons of various Cartan decompositions to each other. A careful study of the proof will reveal the simultaneous use of six Cartan decompositions — those corresponding to conjugation by X and Z on each of three wires. Keeping track of these decompositions in a more systematic manner may simplify the proof, while using additional decompositions may lead to new results. A related challenge is gauging the power of the qubit-by-qubit gate counting we have used. It follows from the results of [18] that $|U|_{\mathtt{CZ};\ell} < 6(n-1)$ for U an n-qubit operator, and hence no technique relying solely on this process can achieve better than a quadratic lower bound. On the other hand, we have only been able to characterize cases when $|U|_{\mathtt{CZ};\ell} > 2$, and thus have achieved only linear lower bounds.

Acknowledgements

We thank Mikko Mottonen, Jun Zhang, K. Birgitta Whaley, and Yaoyun Shi for helpful discussions. This work was sponsored in part by the Air Force Research Laboratory under Agreement No. FA8750-05-1-0282.

References

1. A. Barenco, C. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, and H. Weinfurter, Elementary gates for quantum computation, *PRA* **52**, 3457 (1995).
2. S. S. Bullock, A note on the Khaneja-Glaser decomposition, *QIC* **4**.5, 396-400 (2004).
3. S. S. Bullock, I. L. Markov, Asymptotically Optimal Circuits for Arbitrary n-qubit Diagonal Computations, *QIC* **4**.1, 027-047 (2004).
4. D. P. DiVincenzo and J. A. Smolin, Results on two-bit gate design for quantum computers, *Proc. of the Workshop on Physics and Computation* (1994).
5. D.P. DiVincenzo, Quantum gates and circuits. *Proc. R. Soc. Lond. A*, 454:261276, (1998).
6. N. Margolus, Simple quantum gates, Unpublished manuscript (circa 1994).
7. N. Khaneja and S. J. Glaser, Cartan decomposition of SU(n) and control of spin systems, *Chem. Physics* **267**, 11-23 (2001).
8. A. W. Knapp, *Lie Groups Beyond an Introduction*, Progress in Mathematics, vol. 140, *Birkhäuser*, 1996.
9. E. Knill, Approximation by quantum circuits, LANL report LAUR-95-2225.
10. M. Lewenstein, B. Kraus, J. I. Cirac, and P. Horodecki, Optimization of entanglement witnesses, *PRA* **62**, 052310 (2000).
11. Yu. G. Makhlin, Nonlocal properties of two-qubit gates and mixed states and optimizations of quantum computations, *QIP* **1**, 243-252 (2002).
12. D. Maslov and G. W. Dueck. *IEE Electronics Letters* **39**.25, 1790-1791 (2003).
13. M. Möttönen, J. J. Vartiainen, V. Bergholm, and M. M. Salomaa, Quantum circuits for general multiqubit gates, *PRL* **93**, 130502 (2004).
14. M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, Cambridge University Press (2000).
15. C. C. Paige and M. Wei, History and generality of the CS decomposition, *Linear Algebra and Applications* **208**, 303 (1994).
16. V. V. Shende, I. L. Markov, and S. S. Bullock, CNOT-optimal circuits for generic two-qubit operators, *PRA* **69**, 062321 (2004).
17. V. V. Shende, S. S. Bullock, and I. L. Markov, Recognizing small-circuit structure in two-qubit operators. *PRA* **70**, 012310 (2004).
18. V. V. Shende, S. S. Bullock, and I. L. Markov, Synthesis of quantum logic circuits, *IEEE. Trans. on CAD* **25**, 1000 (2006).
19. G. Song and A. Klappenecker, Optimal realizations of controlled unitary gates, *QIC* **3**, 139-155 (2003).
20. G. Song and A. Klappenecker, The simplified Toffoli gate implementation by Margolus is optimal, *QIC* **4**, 361-372 (2004).
21. T. Toffoli, Reversible Computing, *MIT Technical Report* MIT/LCS/TM-151 (1980).
22. F. Vatan and C. Williams, Optimal quantum circuits for general two-qubit gates, *PRA* **69**, 032315 (2004).
23. G. Vidal and C. M. Dawson, A universal quantum circuit for two-qubit transformations with three CNOT gates, *PRA* **69**, 010301 (2004).
24. J. Zhang, J. Vala, K. B. Whaley, and S. Sastry, A geometric theory of non-local two-qubit operations, *PRA* **67**, 042313 (2003).

Appendix: Proof of Proposition 4

Below we restate Proposition 4 and complete its proof.

Proposition 5 *Fix qubits $\ell_1 \dots \ell_k$ among $N > k$ qubits. A unitary U commuting with $\mathtt{Z}^{(\ell_1)}, \dots, \mathtt{Z}^{(\ell_k)}$ can be implemented by a* CZ*-circuit in which only diagonal gates operate on qubits ℓ_i if and only if $\det_{\ell_1 \dots \ell_k}(U)$ is separable (can be implemented by one-qubit gates).*

Proof: ($\Rightarrow$). It suffices to show the separability of $\det_{\ell_1 \ldots \ell_k}(U)$ for a small generating set of operators. Direct calculation confirms this for (i) CZ gates, (ii) diagonal one-qubit gates on the ℓ_i, and (iii) any gate not affecting qubits ℓ_i.

($\Leftarrow$). By hypothesis, $\det_{\ell_1 \ldots \ell_k}(U)$, and hence $\mathcal{D} = \det_{\ell_1 \ldots \ell_k}(U)^{-2^{k-N}}$, can be implemented using only one-qubit diagonal gates. It remains to implement $\tilde{U} = U/\mathcal{D}$, which satisfies the normalization $\tilde{U}_{j_1 \ldots j_k} \in \mathrm{iSU}(2^{N-k})$. We will construct a circuit for $\tilde{U}$ by multiplexing circuits for $\tilde{U}_{j_1 \ldots j_k}$. Let $\mathcal{C}$ be a $(N-k)$-qubit circuit containing only CZs and one-qubit R_x, R_y, R_z gates such that any operator in $\mathrm{iSU}(2^{N-k})$ can be implemented by making the appropriate choice of parameter for the R_x, R_y, R_z gates. Such universal circuits exist [1]; see Section 2.3 for modern constructions. Choose specifications $\mathcal{C}_{j_1 \ldots j_k}$ implementing the $\tilde{U}_{j_i \ldots j_k}$; let the s-th rotation gate in $\mathcal{C}_{j_1 \ldots j_k}$ be given by $R_{d(s)}(\theta_{j_1 \ldots j_k}(s))^{(q(s))}$, where $q(s)$ is a qubit, $\theta_{j_1 \ldots j_k}(s)$ is an angle, and $d(s) = x, y, z$. Define $\Theta(s)$ to be the real diagonal operator on qubits $\ell_i \ldots \ell_k$ such that $\Theta(s)_{j_i \ldots j_k} = \theta_{j_1 \ldots j_k}(s)$. Form the N-qubit circuit $\tilde{\mathcal{C}}$ by replacing the s-th rotation gate of $\mathcal{C}$ by the multiplexed rotation $R_{d(s)}(\Theta(s))^{(q(s))}$; then $\tilde{\mathcal{C}}$ implements $\tilde{U}$. Implement $R_{d(s)}(\Theta(s))^{(q(s))}$ by a CZ-circuit containing no one-qubit operator on any qubit save $q(s)$, which is not one of the ℓ_i (see [13] or Section 2.3). □

Corollary 9 *N-qubit operators which commute with* Z *on k qubits can be implemented using on the order of $2^k 4^{N-k}$ one-qubit and* CZ *gates.*[g]

Proof: This follows from the construction in the proof of Proposition 4 and the known estimates in the cases $k = 0, N-1$ [13] and $k = N$ [3]. □

[g]Dimension-counting following [9] shows that roughly this many are necessary for almost all such operators.