Advantages of using relative-phase Toffoli gates with an application to multiple control Toffoli optimization

Dmitri Maslov*
National Science Foundation, Arlington, Virginia 22230, USA
(Received 17 August 2015; revised manuscript received 29 October 2015; published 10 February 2016)

Various implementations of the Toffoli gate up to a relative phase have been known for years. The advantage over the regular Toffoli gate is its smaller circuit size. However, its use has been often limited to a demonstration of quantum control in designs such as those where the Toffoli gate is being applied last or otherwise for some specific reasons the relative phase does not matter. It was commonly believed that the relative-phase deviations would prevent the relative-phase Toffoli gates from being very helpful in practical large-scale designs. In this paper, we report three circuit identities that provide the means for replacing certain configurations of the multiple control Toffoli gates with their simpler relative-phase implementations, up to a selectable unitary on certain qubits, and without changing the overall functionality. We illustrate the advantage via applying those identities to the optimization of the known circuits implementing multiple control Toffoli gates, and report the reductions in the controlled-NOT count, T count, as well as the number of ancillae used. We suggest that a further study of the relative-phase Toffoli implementations and their use may yield other optimizations.

DOI: 10.1103/PhysRevA.93.022311

I. INTRODUCTION

Multiple control Toffoli gates are the staple of quantum arithmetic and reversible circuits. They are employed widely within quantum algorithms, including in reversible transformations, such as arithmetic circuits and all sorts of Boolean operations over quantum registers, as well as subroutines within other specialized quantum transforms. Unfortunately, multiple control Toffoli gates are not simple operations, and must be implemented using a certain library of elementary gates—physically attainable transformations for physical-level implementations, and fault-tolerant gates on the logical level. As of the time of this writing, most advanced and developed trapped ions [1] and superconducting [2] quantum information processing approaches allow computations over at most a few dozen qubits using at most a few dozen two-qubit gates. The smallest of the multiple control Toffoli gates, the three-qubit Toffoli gate, requires six controlled-NOT (CNOT) gates as a physical-level circuit over controlling apparatus allowing the application of the CNOT and arbitrary single qubit gates, and seven T gates, as a logical fault-tolerant circuit over the Clifford $+ T$ library without ancillae. The known implementations of larger multiple control Toffoli gates come at a substantially higher cost. This makes the multiple control Toffoli gates be expensive computing primitives. As such, the ability to replace them with their simpler counterparts that nevertheless can guarantee the overall functional integrity, as well as their optimization (multiple control Toffoli gates are implemented using smaller size multiple control Toffoli gates [3,4]), are important in practice. Ultimately, the difficulty of implementing Toffoli gates may even be a deciding factor in the ability to run an experiment of a desired size. Indeed, consider a scenario where only a fixed number of certain elementary gates can be applied. Imagine the goal is to run a discrete logarithm type computation [4]. Since circuits implementing such an algorithm are dominated by reversible arithmetic operations, which in turn rely on the Toffoli gates, it is conceivable that optimizing Toffoli implementations would yield a resource count that is possible to execute for a desired size computation. Multiple control Toffoli gates are, of course, important beyond just the discrete logarithm type algorithms.

The goal of this paper is to provide a framework for replacing multiple control Toffoli gates with their simpler relative-phase implementations. The advantage is illustrated through an optimization of the implementations of the multiple control Toffoli gates. The reported optimization is viewed as a motivating example rather than a complete and finished study. An in-depth look at the implementations of the relative-phase multiple control Toffoli gates and their use in the optimization of arbitrary quantum circuits may likely yield more results.

To draw a classical analogy, relative-phase Toffoli gates may turn out to play a role analogous to the classical NAND gates: while classical (quantum/reversible) circuits are designed using a convenient for a human set of operations (multiple control Toffoli gates), a compiler may decompose those into NAND gates (relative-phase multiple control Toffoli gates) before they are mapped into lowest-level transistors (elementary quantum gates).

II. DEFINITIONS

In this paper, we will work with pure n-qubit quantum states $\sum_{i=0}^{2^n-1} \alpha_i |i\rangle$ and quantum transformations described by the $2^n \times 2^n$ unitary matrices U. Recall that a square matrix U is called unitary if its inverse equals to its conjugate transpose, $U^{-1} = U^\dagger$. While the property of unitarity defines evolutions that are possible to attain physically, it does not prescribe which ones may be implemented directly. To assist with the presentation of the material, we will discretize the family of transformations that may be obtained physically, and call them elementary quantum gates. This does not limit the applicability of the results—indeed, discrete circuits may be thought of as certain versions of continuous Hamiltonians, but are otherwise easier to work with. In particular, in this

*dmitri.maslov@gmail.com; Present address: Joint Center for Quantum Information and Computer Science, University of Maryland, College Park, MD 20742, USA.

work we will rely on the following elementary gates: Pauli-X, $X = \text{NOT} = \left(\begin{smallmatrix}0 & 1\\ 1 & 0\end{smallmatrix}\right)$, Pauli-$Z$, $Z = \left(\begin{smallmatrix}1 & 0\\ 0 & -1\end{smallmatrix}\right)$, and its roots Phase, $P = \sqrt{Z} = \left(\begin{smallmatrix}1 & 0\\ 0 & i\end{smallmatrix}\right)$, $T = \sqrt[4]{Z} = \left(\begin{smallmatrix}1 & 0\\ 0 & \frac{1+i}{\sqrt{2}}\end{smallmatrix}\right)$, and Pauli-$Y$, $Y = \left(\begin{smallmatrix}0 & -i\\ i & 0\end{smallmatrix}\right)$. A fourth root of Y will be mentioned in some constructions, in the form of $R_Y(\pi/4)$, that is equivalent to the fourth root of Y up to a global phase. Recall that $R_Y(\theta) = \left(\begin{smallmatrix}\cos\frac{\theta}{2} & -\sin\frac{\theta}{2}\\ \sin\frac{\theta}{2} & \cos\frac{\theta}{2}\end{smallmatrix}\right)$. Finally, for completeness we will need the Hadamard gate, $H = \frac{1}{\sqrt{2}}\left(\begin{smallmatrix}1 & 1\\ 1 & -1\end{smallmatrix}\right)$, and the two-qubit CNOT gate that we introduce via the mapping of kets, rather than the 4×4 matrix, as $\text{CNOT}(a,b) : |a,b\rangle \mapsto |a,b \oplus a\rangle$, and everywhere else by linearity, due to the simplicity of such a definition.

Quantum circuits are defined as the strings of quantum gates, or otherwise products of matrices that correspond to the individual gates. For multiple qubit circuit computations via matrices, a proper Kronecker product needs to be taken to compute matrix products. For example, a two-qubit operation corresponding to the Hadamard gate on the first qubit is given by the matrix $H \otimes \text{Id}$, where Id is the identity applied to the second qubit. Recall that the product of matrices is taken in reverse order with respect to the order of gates in the corresponding circuit. Following the standard notations, circuits or unitaries composed of quantum gates or matrices X, Y, Z, P, H, and CNOT are called Clifford. These unitaries play an important role in quantum error correction, but are not complete (moreover, simulable classically with a polynomial size effort) for quantum computation. As such, for completeness, a circuit library needs to contain a non-Clifford gate, such as the T gate. The addition of any non-Clifford gate to the Clifford circuits furthermore turns out to result in the computational universality [4].

The above is meant to be a quick reminder of some basic facts and an introduction of the notations used in this paper. For an in-depth review we refer the reader to [4].

For convenience, we furthermore use the following notations: For a set of variables or qubits $X = \{x_1,x_2,...,x_n\}$, $|X|$ equals n, being the number of individual qubits in this set, and the conjugation (Boolean AND) of variables, $x_1\&x_2\&...\&x_n$ is denoted as simply x. When the number of variables in the set X is zero, we assign x the value of 1. When the set of variables X consists of a single element $\{x\}$, the conjugation of the variables within the set, as well as the name of the variable, coincide; this does not, however, cause any issues.

We next define the multiple control Toffoli gates.

Definition 1. A *multiple control Toffoli gate* over a set of n qubits with the set $X = \{x_1,x_2,...,x_{n-1}\}$ being the controls, and qubit y being the target, $\text{TOF}^n(X; y)$, is defined as the matrix,

$$\text{diag}\left\{1,1,....,1,\begin{pmatrix}0 & 1\\ 1 & 0\end{pmatrix}\right\}.$$

We will sometimes omit the superscript and write $\text{TOF}(X; y)$ when the controls and the target are explicitly specified and the size of the multiple control Toffoli gate can thus be restored. Similarly, we may omit the specification of the qubits the gate operates on and write TOF^n when we are only concerned with the size of the gate. Finally, we may write TOF when the goal is to specify the kind of gate being the Toffoli and distinguish it from other kinds of gates. Observe, that when $|X| = 0$ the above definition reports the Pauli-X (NOT) gate, for $|X| = 1$ the definition introduces the CNOT gate, when $|X| = 2$, it reduces to the usual Toffoli gate TOF^3, and for larger sets X, the multiply controlled Toffoli gates—Toffoli-4, Toffoli-5, etc.

An alternate definition of the multiple control Toffoli gate may cast it in the form of the mapping of kets, as follows, $\text{TOF}^n(X; y) : |X,y\rangle \mapsto |X,y \oplus x\rangle$. In some cases, the mapping of kets may be easier to operate with than the corresponding unitary matrix.

In our constructions, relative-phase implementations of quantum unitary transformations play a major role. For the purpose of this work, we define relative-phase implementations as follows.

Definition 2. A *relative-phase* version of a quantum n-qubit unitary operation $U = \{u_{i,j}\}|_{i,j=0..2^n-1}$ is any n-qubit unitary $V = \{v_{i,j}\}|_{i,j=0..2^n-1}$ such that $|v_{i,j}| = |u_{i,j}|$ for all i and j.

In other words, a relative-phase version or otherwise implementation of a unitary U is a unitary V such that the elements of the two matrices differ by $e^{i\pi\phi}$, where $\phi \in \mathbb{R}$, and ϕ may be different for different matrix elements. Observe that $e^{i\pi\phi}0 = 0$, therefore relative-phase versions of unitaries have zeros everywhere the original unitary does.

To illustrate, a relative-phase multiple control Toffoli gate over the set of controls $X = \{x_1,x_2,...,x_{n-1}\}$ with the target y, $\text{RTOF}(X; y)$, can be written as follows,

$$\text{diag}\left\{z_0,z_1,....,z_{2^n-3},\begin{pmatrix}0 & z_{2^n-2}\\ z_{2^n-1} & 0\end{pmatrix}\right\},$$

where z_i are arbitrary length-1 complex numbers. Prefix "R" is used to distinguish the relative-phase version from the multiple control Toffoli gate itself. Observe that when all $z_i = 1$, the respective relative-phase Toffoli gate $\text{RTOF}(X; y)$ becomes the multiple control Toffoli gate $\text{TOF}(X; y)$, and when all z_i take the same but fixed value z, the respective relative-phase Toffoli gate $\text{RTOF}(X; y)$ implements the multiple control Toffoli gate $\text{TOF}(X; y)$ up to an undetectable global phase z.

A relative-phase multiple control Toffoli gate RTOF^n may be thought of as a product of the multiple control Toffoli gate TOF^n and an n-qubit diagonal unitary D^n. Indeed, for a diagonal unitary $D^n := \text{diag}\{z_0,z_1,....,z_{2^n-1}\}$ circuit TOF^nD^n implements a generic relative-phase multiple control Toffoli gate $\text{R}_1\text{TOF}^n = \text{diag}\{z_0,z_1,....,z_{2^n-3},\left(\begin{smallmatrix}0 & z_{2^n-2}\\ z_{2^n-1} & 0\end{smallmatrix}\right)\}$, whereas circuit D^nTOF^n implements a generic relative-phase multiple control Toffoli gate $\text{R}_2\text{TOF}^n = \text{diag}\{z_0,z_1,....,z_{2^n-3},\left(\begin{smallmatrix}0 & z_{2^n-1}\\ z_{2^n-2} & 0\end{smallmatrix}\right)\}$. Observe how both gates are relative-phase multiple control Toffoli gates, but different in the last two nonzero elements, that are being permuted. We will exploit this property in the circuit diagrams. In particular, of the two possible decompositions of the relative-phase multiple control Toffoli gate into a product of the multiple control Toffoli and a diagonal unitary, we will select TOF^nD^n to be the canonic one, and draw the respective relative-phase multiple control Toffoli gate with same controls as the diagonal gate D^n and a distorted target, such as illustrated in Fig. 1(c). The helpful intuition behind this pictorial representation is as follows: A Toffoli gate $\text{TOF}(X; y)$

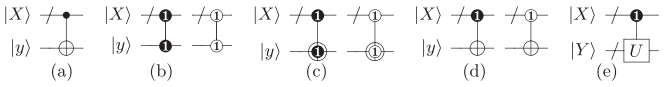

FIG. 1. (a) A multiple control Toffoli gate TOF(X; y). (b) A diagonal gate $D_1(X;y)$ and its inverse. Observe how different diagonal gates can be visually distinguished by the number within the control, and a diagonal gate and its inverse are related by the different color of the control. (c) A relative-phase multiple control Toffoli gate $\mathrm{R_1TOF}(X;y)$ and its inverse $\mathrm{R_1^{-1}TOF}(X;y)$. (d) A type-$y$ special form $\mathrm{S^yR_1TOF}(X;y)$, and its inverse. (e) A controlled-unitary U implemented up to some relative phase. The $-\!\!/-$ symbol denotes a multiqubit register.

may be combined with a diagonal gate $D(Z)$, $Z \in \{X,y\}$, following it to obtain a relative-phase Toffoli gate, or a Toffoli gate TOF(X; y) may be combined with a diagonal gate $D(Z)$, $Z \in \{X,y\}$, preceding it to obtain the inverse of a relative-phase Toffoli gate; conversely, each relative-phase Toffoli gate or its inverse may be broken down into a suitable pair of the multiple control Toffoli gate and the diagonal gate.

An important property of the relative-phase multiple control Toffoli gates is that every one of those is an inverse of some other relative-phase multiple control Toffoli gate. Indeed, for $\mathrm{R_1TOF}^n = \mathrm{diag}\{z_0,z_1,....,z_{2^n-3},\left(\begin{smallmatrix}0 & z_{2^n-2}\\ z_{2^n-1} & 0\end{smallmatrix}\right)\}$ and $\mathrm{R_2TOF}^n = \mathrm{diag}\{w_0,w_1,....,w_{2^n-3},\left(\begin{smallmatrix}0 & w_{2^n-2}\\ w_{2^n-1} & 0\end{smallmatrix}\right)\}$ $\mathrm{R_1TOF}^n = \mathrm{R_2^{-1}TOF}^n$ when $w_i = z_i^{-1}$ for $i = 0...2^n-3$, $w_{2^n-2} = z_{2^n-1}^{-1}$, and $w_{2^n-1} = z_{2^n-2}^{-1}$.

We next define special form relative-phase multiple control Toffoli gates, that are important in some of the constructions that follow.

Definition 3. For a set $X = \{x_1,x_2,...,x_n\}$, and its subset $X' = \{x_{i_1},x_{i_2},...,x_{i_k}\}$ a *type-X' special form relative-phase multiple control Toffoli gate*, $\mathrm{S^{X'}RTOF}(x_1,x_2,...,x_{n-1};x_n)$, is defined as the matrix,

$$\mathrm{diag}\left\{z_0,z_1,...,z_{2^n-3},\begin{pmatrix}0 & z_{2^n-2}\\ z_{2^n-1} & 0\end{pmatrix}\right\},$$

where every pair of complex numbers z_s and z_t are equal whenever the binary expansions of s and t are different only in the digits $i_1,i_2,...,i_{k-1}$, and i_k.

To illustrate, a type-$\{x_1\}$ $\mathrm{S^{x_1}RTOF}(x_1,x_2,...,x_{n-1};x_n)$ is given by the matrix,

$$\mathrm{diag}\left\{z_0,z_1,...,z_{2^{n-1}-1},\right.$$
$$\left.z_0,z_1,...,z_{2^{n-1}-3},\begin{pmatrix}0 & z_{2^{n-1}-2}\\ z_{2^{n-1}-1} & 0\end{pmatrix}\right\}.$$

The type-$\{x_1\}$ special form relative-phase Toffoli gate $\mathrm{S^{x_1}RTOF}$ has half the number of the degrees of freedom compared to the equal size unrestricted relative-phase Toffoli gate RTOF. In practice, this suggests that it should be easier to find an efficient circuit implementing a relative-phase Toffoli gate than it is to find one of the same size for a type-$\{x_1\}$ special form relative-phase Toffoli gate. To give another example, a type-X $\mathrm{S^XRTOF}(x_1,x_2,...,x_{n-1};x_n)$ is the most restrictive of the kind. It is equal to the respective Toffoli gate up to a global phase, and thereby does not give much freedom in implementing by a circuit over the $\mathrm{TOF}(x_1,x_2,...,x_{n-1};x_n)$. This means that in the practical constructions, and whenever possible, we will try to use a type-X' special form relative-phase multiple control Toffoli gate with the smallest size set X'.

An alternate and equivalent definition of a type-X' special form relative-phase Toffoli gate is via a transformation given by the circuit $\mathrm{TOF}(x_1,x_2,...,x_{n-1};x_n)D(X \setminus X')$. It furthermore serves as a basis for how we draw SRTOF gates in the circuit diagrams. Compared to the multiple control Toffoli gate, every control or target in the set $X \setminus X'$ of $\mathrm{S^{X'}RTOF}$ appears distorted by the dingbat originating from the respective $D(X \setminus X')$, and every control or target in the set X' appears undistorted [see Fig. 1(d)].

Beyond having fewer degrees of freedom compared to an unrestricted relative-phase Toffoli gate, there is one more important difference between the special form relative-phase Toffoli gates and the relative-phase Toffoli gates: The inverse of a type-X' special form relative-phase Toffoli gate is not always a type-X' special form relative-phase Toffoli gate.

The use of the subscripts in formulas and different numbers within dingbats in the quantum circuit diagrams allows one to distinguish different versions of the relative phase and special form relative-phase multiple control Toffoli gates. For instance, notations $\mathrm{R_1TOF}$ and $\mathrm{R_2TOF}$ indicate that both gates are some relative-phase Toffoli gates, but they are not necessarily related. In contrast, an $\mathrm{R_1^{-1}TOF}$ is the inverse of the $\mathrm{R_1TOF}$. Recall that a circuit implementing the inverse operation may be constructed by conjugating the gates in the circuit implementing the given unitary and inverting their order. Observe further that any two TOF gates of the same size are represented by identical matrices; this is not always true for some two RTOF or a pair of SRTOF, therefore the ability to distinguish different versions of the relative-phase implementations is important, as these could be different gates.

We will draw quantum gates and circuits using standard notations, including the relative-phase gates per diagrams found in Fig. 1, with time propagating from left to right. Some useful circuit identities clarifying and summarizing the above discussions are shown next.

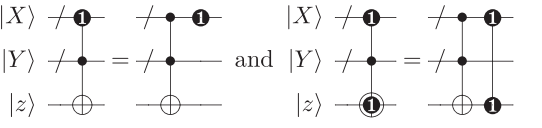

show canonic decomposition of $\mathrm{S^{Y,z}R_1TOF}$ and $\mathrm{S^YR_1TOF}$ into the product of the multiple control Toffoli gate TOF and the diagonal gate D_1; reading right-to-left, these rules show how to combine a suitable pair of the multiple control Toffoli gates and the diagonal gate into a (special form) relative-phase Toffoli gate. When $Y = \emptyset$, the second circuit illustrates the $\mathrm{R_1TOF}$ gate.

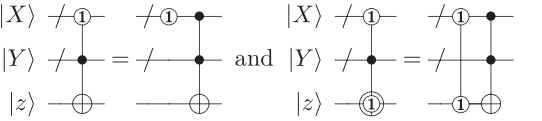

show canonic decomposition of $\mathrm{S^{Y,z}R_1^{-1}TOF}$ and $\mathrm{S^YR_1^{-1}TOF}$ into the product of the diagonal gate and the multiple control Toffoli gate. Indeed, looking at the second of the two

identities,

$$\begin{aligned} &\mathrm{S}^{y}\mathrm{R}^{-1}\mathrm{TOF}(X,Y;z)\\ &= (\mathrm{TOF}(X,Y;z)D_1(X,z))^{-1}\\ &= D_1^{-1}(X,z)\mathrm{TOF}(X,Y;z), \end{aligned}$$

being the circuit pictured on the right-hand side.

$\forall\textcircled{1}\ \exists\textcircled{2}$:

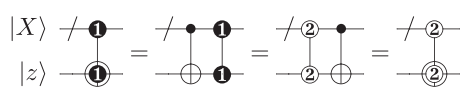

in other words, every $\mathrm{R}_1\mathrm{TOF}$ is also an $\mathrm{R}_2^{-1}\mathrm{TOF}$ under the proper choice of relative phases.

In general, for any reversible gate $R(X)$ its relative-phase version could be thought of as a product $R(X)D(X)$, for a proper diagonal unitary $D(X)$. This suggests a possible route in which the work reported in this paper may be extended.

III. MAIN RESULT

Our main result is summarized in the next three propositions. We apply it to obtain multiple corollaries, and to optimize multiple control Toffoli gates in the section that follows. The proofs of these three propositions rely on the three circuit identities concluding the previous section, as well as the following notion: The controlled-U implemented up to a relative phase, $\mathrm{RCU}(V,W;X)$, commutes with the controlled V implemented up to a relative phase, $\mathrm{RCV}(V,Y;Z)$, where the qubit sets V,W,X,Y, and Z are disjoint. This rule also applies to show that any two nonintersecting unitaries commute. We assume the reader's familiarity with the above commutation rule, and do not explicitly prove it here.

Proposition 1. The conjugation of the controlled unitary U over the qubit set Z implemented up to a possible relative phase, $R_1CU(Y,a;Z)$, by a pair of multiple control Toffoli gates $\mathrm{TOF}(X;a)$ allows the replacement of these multiple control Toffoli gates with their relative-phase versions implemented up to any desired unitary $V(X)$, such as illustrated next:

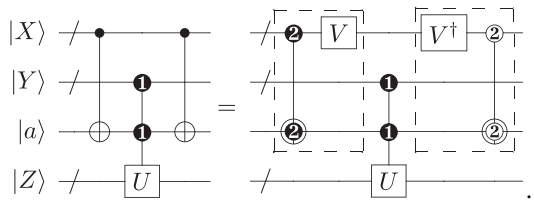

(1)

Proof. The proof is accomplished via the following set of circuit transformations:

$$\begin{aligned} &\mathrm{TOF}(X;a)R_1CU(Y,a;Z)\mathrm{TOF}(X;a)\\ &= \mathrm{TOF}(X;a)D_2(X;a)D_2^{-1}(X;a)\\ &\quad R_1CU(Y,a;Z)\mathrm{TOF}(X;a)\\ &= \mathrm{TOF}(X;a)D_2(X;a)R_1CU(Y,a;Z)\\ &\quad D_2^{-1}(X;a)\mathrm{TOF}(X;a)\\ &= \mathrm{R}_2\mathrm{TOF}(X;a)R_1CU(Y,a;Z)\mathrm{R}_2^{-1}\mathrm{TOF}(X;a)\\ &= \mathrm{R}_2\mathrm{TOF}(X;a)V(X)V^{-1}(X)R_1CU(Y,a;Z)\\ &\quad \mathrm{R}_2^{-1}\mathrm{TOF}(X;a)\\ &= [\mathrm{R}_2\mathrm{TOF}(X;a)V(X)]R_1CU(Y,a;Z)\\ &\quad \left[V^{-1}(X)R_2^{-1}\mathrm{TOF}(X;a)\right]. \end{aligned}$$

■

The result of Proposition 1 can be reduced to the following form once $\mathrm{RCU}(Y,a;Z)$ is set to implement the Toffoli-type gate, $\mathrm{TOF}(Y,a;z)$:

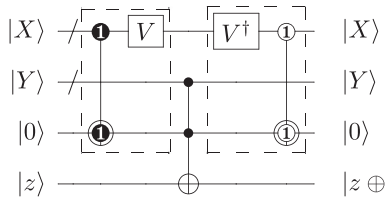

Indeed, the corresponding circuit on the left-hand side in (1) computes

$$\begin{aligned} &|X,Y,0,z\rangle \overset{\mathrm{TOF}(X;0)}{\mapsto} |X,Y,x,z\rangle \overset{\mathrm{TOF}(Y,x;z)}{\mapsto}\\ &|X,Y,x,z\oplus xy\rangle \overset{\mathrm{TOF}(X;x)}{\mapsto} |X,Y,0,z\oplus xy\rangle, \end{aligned}$$

which is indicated by the formulas on the output side. This, in turn, leads to the following corollary.

Corollary 1. An n-qubit Toffoli gate TOF^n can be implemented with the cost not exceeding the sum of twice the cost of an n-qubit relative-phase Toffoli gate RTOF^n and the cost of the CNOT gate, using one ancilla qubit set to and returned in the value $|0\rangle$. In other words, in the presence of such an ancilla,

$$\mathrm{Cost}(\mathrm{TOF}^n) \leqslant 2\times\mathrm{Cost}(\mathrm{RTOF}^n)+\mathrm{Cost}(\mathrm{CNOT}).$$

This corollary may be reformulated for a different choice of the middle gate, e.g., as follows: $\mathrm{Cost}(\mathrm{TOF}^n) \leqslant 2\times \mathrm{Cost}(\mathrm{RTOF}^{n-1})+\mathrm{Cost}(\mathrm{TOF}^3)$.

Other gate configurations are also supported by the relative-phase Toffoli gates. The following proposition complements the set of basic rules on which we base the proposed optimization approach.

Proposition 2. Consider the conjugation of a controlled-U gate $R_1CU(W;X,Y)$ implemented possibly up to some relative phase, by a pair of identical multiple control Toffoli gates, such as illustrated in (2) on the left-hand side. Then, the following circuit identity holds for any unitary transformation V over the qubit set $\{Z\cup a\}$ and any $\mathrm{S}^Y\mathrm{R}_2\mathrm{TOF}(Y,Z;a)$ (a type-Y special form relative-phase Toffoli gate):

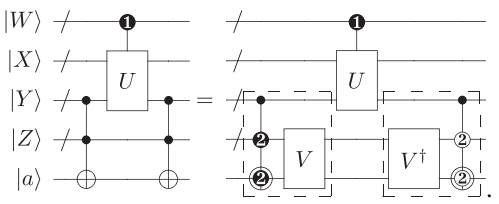

(2)

Proof. This proposition may be proved similarly to Proposition 1,

$$\begin{aligned} &\mathrm{TOF}(Y,Z;a)R_1CU(W;X,Y)\mathrm{TOF}(Y,Z;a)\\ &= \mathrm{TOF}(Y,Z;a)D_2(Z;a)D_2^{-1}(Z;a)\\ &\quad R_1CU(W;X,Y)\mathrm{TOF}(Y,Z;a) \end{aligned}$$

$$
\begin{aligned}
&= \mathrm{TOF}(Y,Z;a)D_2(Z;a)R_1CU(W;X,Y)\\
&\quad D_2^{-1}(Z;a)\mathrm{TOF}(Y,Z;a)\\
&= \mathrm{S}^y\mathrm{R}_2\mathrm{TOF}(Y,Z;a)R_1CU(W;X,Y)\\
&\quad \mathrm{S}^y\mathrm{R}_2^{-1}\mathrm{TOF}(Y,Z;a)\\
&= \mathrm{S}^y\mathrm{R}_2\mathrm{TOF}(Y,Z;a)V(Z;a)V^{-1}(Z;a)\\
&\quad R_1CU(W;X,Y)\mathrm{S}^y\mathrm{R}_2^{-1}\mathrm{TOF}(Y,Z;a)\\
&= [\mathrm{S}^y\mathrm{R}_2\mathrm{TOF}(Y,Z;a)V(Z;a)]R_1CU(W;X,Y)\\
&\quad [V^{-1}(Z;a)\mathrm{S}^y\mathrm{R}_2^{-1}\mathrm{TOF}(Y,Z;a)].
\end{aligned}
$$

An alternate proof may be constructed via restricting W,X,Y, and Z to contain at most a single qubit each, and multiplying the corresponding matrices [5]. The benefit of considering such a matrix multiplication is in the ability to show that the $\mathrm{S}^Y\mathrm{R}_2\mathrm{TOF}(Y,Z;a)$ turns out to be the relative-phase Toffoli gate that allows the most freedom in selecting relative phases for a general unitary U, allowing one to formulate this proposition as an "if-and-only-if" statement. Furthermore, looking at the matrices helps to expand the set of possible allowed relative-phase replacements once U is known. ■

The results of Propositions 1 and 2 may be generalized via introducing a control set P that controls all three gates on the left-hand side as well as all five gates on the right-hand side, and a control set Q that controls all gates except V.

Observe that between the two propositions they cover all situations when a relative-phase controlled U is conjugated by a pair of multiple control Toffoli gates such that the targets of those multiple control Toffoli gates do not intersect with the U, resulting in the ability to replace a pair of multiple control Toffoli gates with a pair of simpler gates. A similar circuit identity may be developed for the scenario when the target of the multiple control Toffoli gates intersects with the qubits used by the unitary U. This circuit identity relies on the special form relative-phase Toffoli gates. We have not yet found practical examples where such circuit identity would yield an advantage and the results of Propositions 1 and 2 do not apply, but formulate the statement of the respective proposition for completeness.

Proposition 3. The conjugation of the controlled unitary U implemented up to a relative phase, $R_1CU(X;Y,Z,a)$, by a pair of the multiple control Toffoli gates $\mathrm{TOF}(W,Z;a)$ allows the replacement of these multiple control Toffoli gates with the type-$\{Z\cup a\}$ special form relative-phase version [up to a multiplication by any desired unitary $V(W)$] and its inverse, as follows:

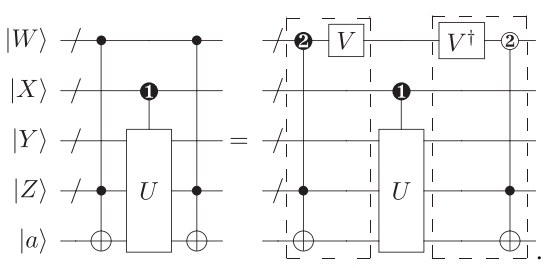

We do not include an explicit proof, but mention that it may be obtained similarly to that of Propositions 1 and 2. Furthermore, we note that the scenario where $R_1CU(X;Y,Z,a)$ is a diagonal gate, e.g., a controlled R_z implemented up to a possible relative phase, is better handled by applying Proposition 1 than Proposition 3 (e.g., see item III A, Sec. III A). Indeed, Proposition 1 uses the most generic unspecified type relative-phase Toffoli, and any controlled R_z may be thought of as a targetless gate ($|Z|=0$ in the statement of Proposition 1) or otherwise, one may introduce a new target qubit that applies a global phase (Fig. 4.5 in [4]).

A. Applications

The principal circuit equalities (1) and (2) suggest a circuit optimization procedure by which a suitable pair of the multiple control Toffoli gates can be replaced with their relative phase or special form relative-phase implementations up to the right-hand multiplication by any desired unitary over the proper qubit set. The rules may be used interchangeably and combined. In particular, we next illustrate how the above approach can be applied to optimize the most popular constructs used to implement/decompose the multiple control Toffoli gates into simpler gates. In the following discussions, we will omit unitaries V, with the understanding that if need be, they may be added back in.

Corollary 2. (Optimization of the construction reported in [3], Lemma 7.2.) A multiple control Toffoli gate TOF^n can be implemented by a circuit consisting of $4n-14$ relative-phase Toffoli gates RTOF^3 and a type-$\{y\}$ special form relative-phase Toffoli gate $\mathrm{S}^y\mathrm{RTOF}^3(x,y;z)$ and its inverse over a circuit with at least $2n-3$ qubits, such as illustrated in Fig. 2.

Proof. The numeric order of subscripts in the special form and relative-phase Toffoli gates indicates the order in which the circuit equalities (2) and (1) are applied to the original circuit reported in [3], Lemma 7.2, to obtain the desired simplified decomposition. Observe that when during this process a pair of Toffoli gates $\mathrm{TOF}^3(a,b;c)$ is replaced with a special form or a relative-phase implementation, the circuit in the middle may be equivalent to a combination of a suitable multiple control Toffoli gate—possibly up to a relative phase, and a transformation on the qubits outside the set $\{a,b,c\}$. This latter transformation may be factored out, thereby allowing all circuit alternations to retain the original functional correctness.

Finally, observe that the identities (1) and (2) may be used in a number of different ways, resulting in different constructions, and not just the particular one selected in the statement of the corollary. In Fig. 2(b) we used one of such constructions that minimizes the number of the special form relative-phase Toffoli gates to gain most freedom in substituting Toffoli gates with their relative-phase implementations. In Fig. 2(c) we furthermore restricted the number of potentially different RTOF gates via making the following assignments: $\mathrm{R}_4\mathrm{TOF}:=\mathrm{R}_3\mathrm{TOF}$, $\mathrm{R}_5\mathrm{TOF}:=\mathrm{R}_3^{-1}\mathrm{TOF}$, and $\mathrm{R}_6\mathrm{TOF}:=\mathrm{R}_3^{-1}\mathrm{TOF}$. This implementation will be used later in the paper. ■

Corollary 3. (Optimization of the construction reported in [3], Lemma 7.3.) A multiple control Toffoli gate TOF^n can be implemented by a circuit consisting of two relative-phase Toffoli gates RTOF^k and two special form relative-phase Toffoli gates SRTOF^{n-k+2} over a circuit with at least $n+1$

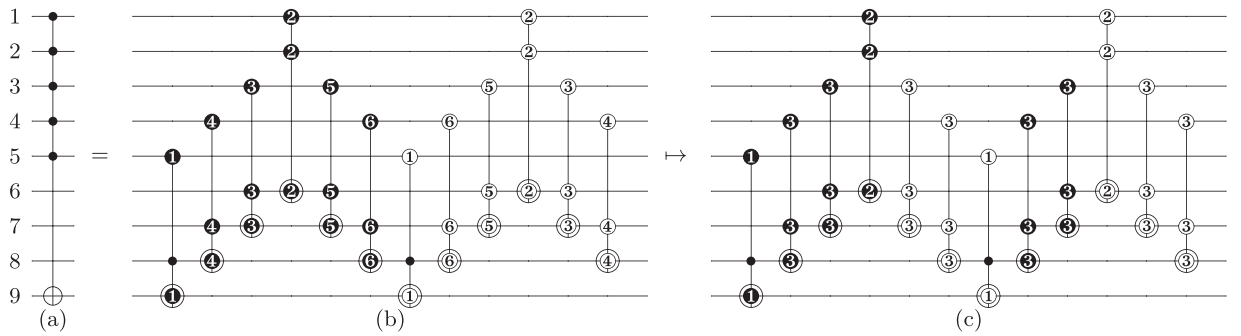

FIG. 2. (a) and (b) Implementation of TOF^6 on qubits 1–9 using $\mathrm{S}^{\{8\}}\mathrm{RTOF}^3$ and its inverse, and 10 RTOF^3 gates. (a)–(c) Implementation of TOF^6 using a narrow selection of the relative-phase and special form relative-phase Toffoli gates.

qubits, such as illustrated next:

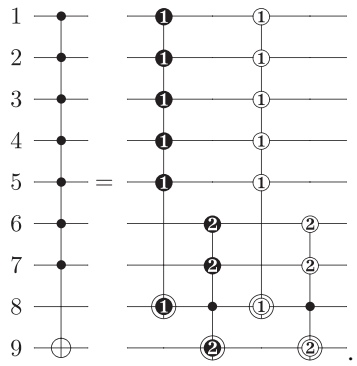

Proof. To obtain this construction, both circuit identities (1) and (2) need to be applied once, in any order. ■

Corollary 4. (Optimization of the construction in [4], page 184.) A multiple control gate C^nU can be implemented by a circuit consisting of $2n-2$ relative-phase Toffoli gates RTOF^3 and one CU gate over a circuit with at least $2n$ qubits of which some $n-1$ qubits are set to and returned in the value $|0\rangle$, such as illustrated next:

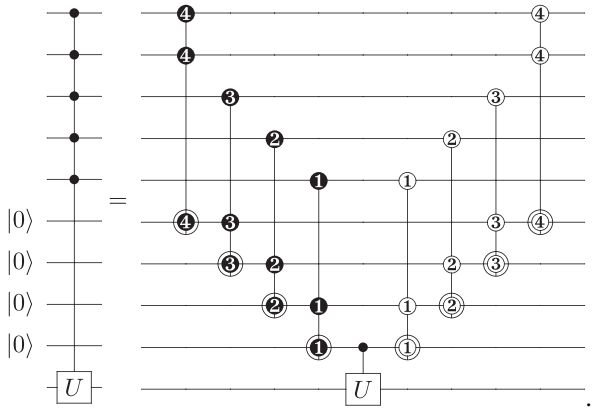

Proof. The circuit identity (1) is applied $n-1$ times. ■

The implementation in [7] optimizes the depth of the circuit ([4], page 184), but does not prevent our construction from being applied. We formalize this observation in the following corollary.

Corollary 5. (Optimization/generalization of the construction in Eq. (13) of [7].) A multiple control gate C^nU can be implemented by a circuit consisting of $2n-2$ relative-phase Toffoli gates RTOF^3 and one CU gate over a circuit with at least $2n$ qubits of which some $n-1$ qubits are set to and returned in the value $|0\rangle$, such as illustrated next:

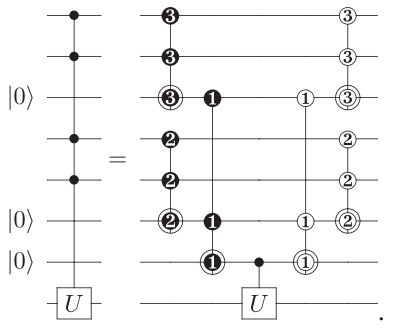

Some other optimizations include the following.

(1) Circuit in [3], Lemma 7.5, may rely on the simpler relative-phase multiple control Toffoli gate and its inverse, rather than two multiple control Toffoli gates (gates No. 2 and No. 4 on the right-hand side).

(2) Circuit in [3], Lemma 7.9, may rely on the simpler special form relative-phase multiple control Toffoli gate and its inverse, rather than two multiple control Toffoli gates (gates No. 2 and No. 4 on the right-hand side).

(3) Circuit in [3], Lemma 7.11, may rely on the simpler relative-phase multiple control Toffoli gate and its inverse, rather than two multiple control Toffoli gates (gates No. 1 and No. 3 on the right-hand side).

(4) Circuit ([6], Fig. 3) may rely on the simpler relative-phase Toffoli gates and their inverses, as is best seen via applying Proposition 1.

IV. OPTIMIZING IMPLEMENTATIONS OF THE MULTIPLE CONTROL TOFFOLI GATES USING THE EXISTING RELATIVE-PHASE TOFFOLI CIRCUITS

In this section we study in detail how to optimize the implementations of the multiple control Toffoli gates, show that all of the known optimized implementations can be explained by the means of the relative-phase Toffoli substitutions described in this work, and report some new optimized circuits.

A. Circuit cost

The question of the efficiency of implementing a certain transformation requires one to formally define a circuit cost.

Depending on the definition of cost, certain circuits will be preferred over other circuits.

There are a number of different definitions of the circuit cost used in the literature, each originating from considering certain specific requirements. At the highest abstraction level, first, one needs to determine if they are dealing with logical level or physical level circuits.

In the former case, one has to derive the protocols and compute the costs of the constructible fault-tolerant gates, given the selected approach to error correction. Within this framework Clifford+T circuits received a significant attention. This is because Clifford gates such as Pauli-X, Y, Z, Hadamard, Phase, and CNOT are believed to be relatively inexpensive to implement fault tolerantly on the logical level. The non-Clifford gate T, or any other constructible non-Clifford gate required for computational universality, is more difficult to generate. The known approaches employ state purification and gate teleportation as a means of generating the T gate, that can get quite costly in the realistic systems [8]. As a result, the cost of the implementation of a logical circuit can be very crudely approximated by the number of the T gates used.

In the case of physical level circuits, one is limited to the ability of the controlling apparatus to apply transformations to the physical quantum information processing system of choice. There is a great variety of the possibilities here. We consider a simple and popular weak interaction model, where the single-qubit gates can be implemented efficiently, and of the two-qubit gates, that take considerably more effort to implement, we have just the CNOT gate. The cost of the circuits can thus be evaluated via counting the number of the CNOT gates in the single-qubit and CNOT gate circuits. Despite apparent oversimplification, there is a specific promising quantum information processing approach, where exactly this formula describes the circuit cost at a high abstraction level. Indeed, trapped ions with the Molmer-Sorenson gate [9] operate in the weak coupling regime (two-qubit gates take roughly 10- to 20-fold effort to implement compared to arbitrary single-qubit gates), and the Molmer-Sorenson gate itself is equivalent to the CNOT up to a conjugation by a pair of $R_Z(a)$ and $R_Z(-a)$ gates on both qubits, for a proper choice of parameter a, and a few single-qubit Phase and Hadamard gates.

An advantage of measuring the cost of the circuit implementations by the T count and the CNOT count is due to the popularity of these circuit cost metrics in the literature, and the ability to compare relative-phase inspired implementations developed in this work to the known ones.

Disadvantages of using either one of these two circuit cost metrics are numerous. Neither circuit metric accounts for the following:

(1) the depth, that could be more important than the gate count, especially when one is, quite naturally, concerned with the speed of the computation given by a quantum circuit rather than just its size;

(2) the connectivity pattern of the qubits. Indeed, physical space spans only three dimensions, and every qubit cannot be connected to every other qubit in a scalable fashion within a finite-dimensional space; or

(3) the number of ancillary qubits used, that is particularly important on the physical level. The number of ancillary qubits used also influences the efficiency of connections between primary qubits. This is because both primary qubits and ancillary qubits share the same physical space and yet need to be as close to each other as possible for higher efficiency.

These are all very important practical considerations. However, our goal is to demonstrate the advantages of the framework introduced in this paper for designing efficient circuits, therefore we restrict the attention to the above two simplistic metrics. We furthermore encourage one to apply the techniques from this paper to designing efficient circuits in the scenario where the details of the circuit cost function are known.

B. Toffoli and Toffoli-4 gates up to a relative phase

First, recall a circuit implementing the Toffoli gate TOF(a,b;c) itself:

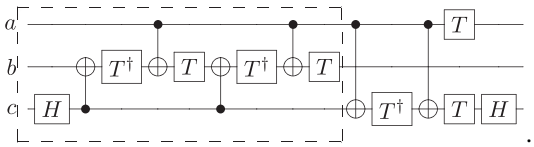

(3)

This circuit may be drawn in many different ways using no more than the minimal numbers of six CNOT gates and 7 $T/T^\dagger$ gates, however, we prefer this form since it has the largest number of gates operating on the qubits a and c after no more gates are being applied to the qubit b.

Literature encounters two apparently related implementations of the Toffoli gate up to a relative phase ([4], page 183, and [7]), that we summarize in one distilled picture; see Fig. 3. There are more symmetries and properties to this circuit than those that necessarily meet the eye on the first glance. In particular, we note the following.

(1) Gates 1–10 implement a type-{c} special form relative-phase Toffoli gate $\mathrm{S}^c\mathrm{RTOF}(a,b;c) = \mathrm{diag}\{1,1,1,1,1,1,\left(\begin{smallmatrix}0 & i\\ i & 0\end{smallmatrix}\right)\}$, whereas gates 2–10 implement a relative-phase Toffoli gate $\mathrm{RTOF}(a,b;c) = \mathrm{diag}\{1,1,1,1,1,-1,\left(\begin{smallmatrix}0 & -i\\ i & 0\end{smallmatrix}\right)\}$.

(2) The first gate, the controlled Z, can be moved to the end of the circuit, resulting in the construction of the type-{c} special form relative-phase Toffoli gate $\mathrm{S}^c\mathrm{RTOF}(a,b;c) = \mathrm{diag}\{1,1,1,1,1,1,\left(\begin{smallmatrix}0 & -i\\ -i & 0\end{smallmatrix}\right)\}$.

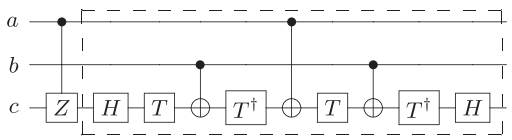

FIG. 3. Toffoli gate implemented up to a relative phase. Gates 1–10 implement a type-{c} special relative-phase Toffoli gate, known as the controlled-controlled-iX in [7], whereas circuit with gates 2–10 implement some generic relative-phase Toffoli gate. The controlled-Z gate $CZ(a;c)$ may commute through the Hadamard $H(c)$, at which point it will change into CNOT(a;c), and the circuit will show in an alternate form. It may be established, via applying the result of Corollary 1, that the CNOT count of the circuit with gates 2–10 is optimal.

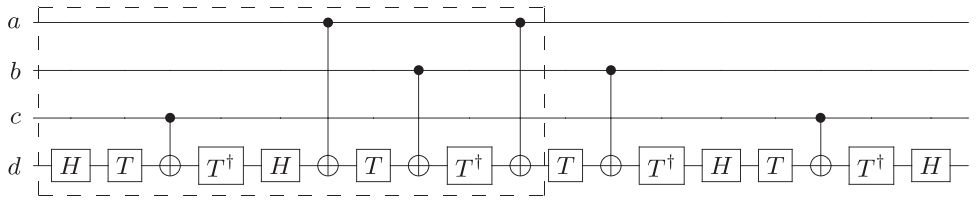

FIG. 4. Circuit implementing Toffoli-4 up to a relative phase, RTOF(a,b,c;d).

(3) Simultaneous substitution $T \mapsto T^\dagger$ and $T^\dagger \mapsto T$ allows constructing more circuits implementing a relative-phase Toffoli gate.

(4) The circuit given by gates 2–10 is self-inverse.

(5) Qubits a and b may be interchanged. Applying this operation gives modified relative-phase Toffoli circuits.

(6) Adding gates $T^m(a)$ and $T^n(b)$ (powers of the T gate), where $m,n \in \{0,1,...,7\}$, to both the beginning and the end of the circuit in Fig. 3 allows constructing more relative-phase Toffoli gates.

(7) Consider gates 3–9. Using the CNOT-T algebra terminology [7,10], the T gate is being applied to $\{c, -(b \oplus c), a \oplus b \oplus c, -(a \oplus c)\}$ (negative sign indicates the application of $T^\dagger$). Instead, we may apply the T gate to $\{c, b \oplus c, -(a \oplus b \oplus c), -(a \oplus c)\}$. Then, the circuit we obtain looks as follows:

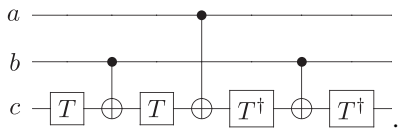

Observe how similar it is to [4], page 183—essentially, Y rotations are replaced by Z rotations. Optimality of the above circuit employing R_Y rotations in place of T (sometimes known as the Margolus gate) was shown in [11]. Conjugating this circuit by a pair of Hadamard gates on the qubit c allows one to obtain a relative-phase Toffoli RTOF(a,$\bar{b}$;c), where $\bar{b}$ denotes the negative control. Similarly, if the $T/T^\dagger$ gates of the circuit in Fig. 3, gates 3–9, were replaced with $R_Y(\pi/4)/R_Y(-\pi/4)$, as illustrated next,

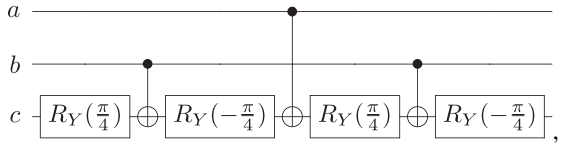

we would have obtained an RTOF(a,$\bar{b}$;c).

We found no relative-phase Toffoli-4 implementations in the literature, but realized that one may be constructed as follows. Consider circuit in Fig. 3, gates 2–10. Replace CNOT(a;c) with a type-{c} special form relative-phase Toffoli gate S^{c}RTOF(x,a;c); this operation introduces a new qubit, x. The result is an RTOF(x,a,b;c). Figure 4 illustrates the result of such a procedure for SRTOF selection per Fig. 3 (observe that the controlled Z was commuted through the Hadamard gate to obtain the CNOT). In the matrix form, the gate looks as follows, diag$\{1,1,1,1,1,1,1,1,1,1,1,1,1,i,-i,\left(\begin{smallmatrix}0 & 1\\ -1 & 0\end{smallmatrix}\right)\}$.

C. Results of the simplification

Since T-count optimal and CNOT-count optimal implementations of the three-qubit Toffoli gate are known, we will concentrate on the Toffoli-4 and larger gates. This section is not meant to report complete results of the optimization that is possible to obtain (indeed, there is no guarantee there are no better relative-phase Toffoli-4 gates to be used, and we did not look for the relative-phase Toffoli-5 and larger gates), rather show a clear advantage of using relative-phase and special form relative-phase Toffoli gates and motivate their further in-depth study.

Consider Toffoli-4 implementation via a circuit with Clifford + T gates. Using the matrix determinant argument, one may establish that the Toffoli-4 may not be implemented unless at least one ancilla qubit is available. This is because the determinant of the 16×16 matrix representing the Toffoli-4 evaluates to the number (-1), whereas the determinants of all Clifford + T library gates, when viewed as 16×16 matrices, are equal to 1. By composing the products of matrices with determinant 1 it is impossible to obtain a matrix with determinant (-1). As a result, at least one ancilla is required.

Once we have established that an ancilla qubit is required, there are two options for the kind of ancilla qubit it is. One, more restrictive, prescribes that the ancilla be available in the state $|0\rangle$; the other provides the ancilla in some unknown state $|x\rangle$. In both cases, when implementing Toffoli-4 with the help of an ancilla, special care needs to be taken to return the value of ancilla to its original state. We consider both cases next.

D. Optimization of Toffoli-4

Ancilla $|0\rangle$, minimizing T count. Literature encounters two results, [10] and [7], both based on the optimization of [4], page 184. In particular, [10] reports an optimized circuit with 15 T gates (down from unoptimized 21), and [7] observes that two Toffoli gates can be replaced with the relative-phase Toffoli called the controlled-controlled-iX (Fig. 3), which explains the optimization obtained in [10]. Our solution uses a somewhat simpler relative-phase Toffoli [see Fig. 3, dashed (gates 2–10)], to obtain TOF^4(a,b,c;d):

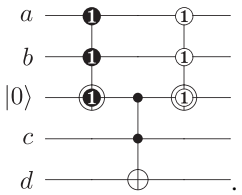

(4)

There is no advantage in the number of T gates. However, our solution explains both known circuits and features a smaller overall gate count.

Ancilla $|0\rangle$*, minimizing* CNOT *count.* Reference [7] uses controlled-controlled-iX to obtain an implementation with 14 CNOTs. To our knowledge, this was the best known result in the literature to date. Our construction, (4), requires only $12 = 3 + 6 + 3$ CNOT gates, since our relative-phase Toffoli (Fig. 3, dashed) requires one less CNOT gate. Observe, that per [13] the lower bound for the number of CNOT gates is eight. Therefore, our 12-CNOT construction may not be improved by more than four CNOT gates.

Arbitrary single-qubit ancilla, minimizing T count. The best known solution, [10], optimizes the 28-T-gate implementation from [3], Lemma 7.2. The result is a circuit with 16 T gates. Our solution matches this solution, and in fact explains how it works. Indeed, we obtain the desired $\mathrm{TOF}^4(a,b,c;d)$ as follows:

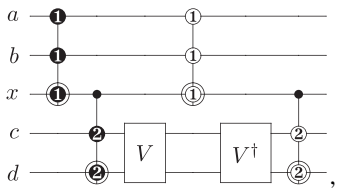

(5)

where x is the ancilla qubit in an unknown state, $\mathrm{R_1TOF}(a,b;x)$ is the relative-phase Toffoli per Fig. 3, dashed; and the $\mathrm{S^xR_2TOF}(x,c;d)V(c,d)$ pair is given by (3), dashed. Essentially, $V(c,d)$ is designed such as to undo all gates applied to the qubits c and d at the end of the implementation given by (3). We have not found a suitable special relative-phase Toffoli gate implementation that is different from the implementation of the Toffoli gate itself, per (3), and giving a better optimization once combined with proper $V(c,d)$. The resulting T count of our construction is thus $16 = 4 + (7-3) + 4 + (7-3)$. Apart from the matching number of T gates, our solution contains fewer Clifford gates (e.g., 14 CNOTs vs 54 CNOTs in [10]), and may also be rewritten as a T-depth four circuit (T depth 1 per each relative-phase Toffoli stage) at the cost of a higher number of ancillae and a higher number of CNOT gates.

Arbitrary single-qubit ancilla, minimizing CNOT *count.* Using CNOT-optimal implementation of the controlled-controlled-iX from [7] over [3], Lemma 7.2, would yield a circuit with 20 CNOT gates, as is done in [12]. The original circuit ([3], Lemma 7.2) uses 24 CNOT gates after each Toffoli is substituted with their CNOT-optimal implementation. Our construction, (5), contains $14 = 3 + 4 + 3 + 4$ CNOT gates.

Observe that the above implementations, if considered as circuits over the Clifford $+ T$ library, use the minimal number of ancillae, being one.

E. Optimization of Toffoli-5

One may once again apply the determinant argument to establish that the Toffoli-5 gate needs at least one ancilla to be available before it may be implemented as a Clifford $+ T$ circuit.

All ancillae in the state $|0\rangle$*, minimizing T count.* The best known solution is given by [10] via an optimization of the construction in [4], page 184, and explained by [7] to be a four controlled-controlled-iX and one Toffoli circuit. The T count is 23 and both known solutions use two ancillae. Our solution implementing $\mathrm{TOF}^5(a,b,c,d;e)$ is as follows:

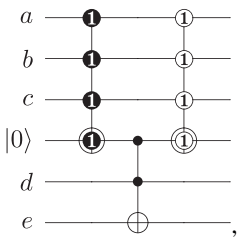

(6)

per $\mathrm{R_1TOF}^4$ implementation found in Fig. 4 and Toffoli implementation from (3). Our solution uses $23 = 8 + 7 + 8$ T gates, relies on only one ancilla, and has a smaller total number of gates compared to the previously known constructions.

All ancillae in the state $|0\rangle$*, minimizing* CNOT *count.* The construction from [7] gives the best known CNOT count of 22 over a circuit that uses two ancillae. Our circuit (6) contains 18 CNOTs and uses only one ancilla. Recall that the lower bound for the number of CNOT gates is 10 [13].

All ancillae in an unknown state, minimizing T count. The best known solution is given in [10] and features 28 T gates. Our solution implementing $\mathrm{TOF}^5(a,b,c,d;e)$ is as follows:

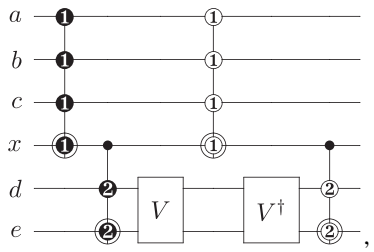

(7)

where x is the ancilla qubit in an unknown state, $\mathrm{R_1TOF}^4$ is the relative-phase Toffoli from Fig. 4, and the $\mathrm{S^xR_2TOF}(x,d;e)V(d,e)$ pair is given by (3), dashed. Observe, that the overall number of T gates is $24 = 8 + (7-3) + 8 + (7-3)$, we use one less ancilla compared to the best known construction, and a smaller overall number of the non-T gates.

We can furthermore explain how to obtain the solution with $12k - 20$ T gates to implement a k-controlled Toffoli gate using $k - 2$ unrestricted ancillae featured in [10] without resorting to a computer optimization. This is done via the use of the relative-phase Toffoli and Toffoli-V pair from Fig. 3 (dashed) and (3), dashed, over the construction reported in Corollary 2. We illustrate how this works using the circuit from Fig. 2(c) and observe that the arguments easily generalize to arbitrary k. Substituting relative-phase and special relative-phase-V pair implementations into the construction in Fig. 2(c) replaces each relative-phase Toffoli with a circuit containing 4 T gates. The total number of the T gates would thus be 48 (for arbitrary k, $16k - 32$), higher than 40 [10]. However, observe

that R_3TOF and R_3^{-1}TOF are inverses of each other. This means that the gates $T^\dagger$ and H on the target qubit that the R_3TOF ends with would cancel with H and T that the R_3^{-1}TOF begins with. This cancellation happens between all four such pairs $\{R_3TOF,R_3^{-1}TOF\}$ found in the circuit. The total reduction is thus by eight T gates (for arbitrary k, $4k-12$), leading to a circuit with 40 T gates (for arbitrary k, $16k-32-4k+12=12k-20$).

All ancillae in an unknown state, minimizing CNOT *count.* [12] includes an implementation where the controlled-controlled-iX is used within [3], Lemma 7.2, for all but two gates. This construction relies on 36 CNOT gates. For arbitrary n, the CNOT count is $16n-44$, which we further refer to as cc-iX implementation in Table I. Observe that [14] reports an implementation with 26 two-qubit gates using two ancillae. The optimization in [14] is motivated by a computational model where the two-qubit interaction given by $\mathrm{diag}\{I,X^{\pm t}\}$, where $t\in\mathbb{R}[0,1]$ and X is Pauli-X, is tunable and parametrized by time. Therefore, for example, a controlled-$\sqrt{\text{NOT}}$ would cost half as much as the CNOT, as it only needs to be evolved for half the time. In our calculations given here, we do not allow such things to happen, but observe that it would be interesting to apply the reported relative-phase Toffoli constructions within that framework. Controlled-$\sqrt{\text{NOT}}$ may be implemented as a 2-CNOT circuit (Fig. 4.6 in [4]). The 26 two-qubit gate circuit of [14] has 18 controlled-$\sqrt{\text{NOT}}$ gates and eight CNOT gates, therefore it would be transformed into one with 44 CNOT gates. Note, however, that it would make little sense from the point of view of the computational model considered in [14], as a length-0.5 interaction is being replaced with a length-2 interaction.

In comparison, our solution, given by (7), is a circuit with 20 $(=6+4+6+4)$ CNOT gates that uses only one ancilla—the latter being provably optimal within the framework of Clifford + T circuits.

We generalize the above examples of Toffoli-4 and Toffoli-5 optimization to any number of qubits in the following two propositions.

Proposition 4. A size $n\geqslant 4$ multiple control Toffoli gate TOF^n may be implemented using $\lceil\frac{n-3}{2}\rceil$ ancillary qubits, set to and returned in the value $|0\rangle$, by a circuit with:

$8n-17$ T gates,
$6n-12$ CNOT gates, and
$4n-10$ Hadamard gates.

Proof. The proof is by induction. The statement is clearly true for $n=4$ and $n=5$, as has been explicitly verified in the previous discussions. To prove the transition from an even $n=2k$ to the odd $n=2k+1$ observe that the middle gate TOF^3 can be replaced with the circuit (4). This introduces an RTOF^3 (Fig. 3, dashed) and its inverse. Note that a new ancillary qubit is being introduced on this step, and the gate counts increase by $8=4+4$ for T, by $6=3+3$ for CNOT, and by $4=2+2$ for Hadamard. The transition from an odd $n=2k+1$ to the even $n=2k+2$ is accomplished via replacing RTOF^3 with RTOF^4 (Fig. 4) and its inverse with the inverse of RTOF^4. Observe that the gate counts grow by 8/6/4 for T/CNOT/Hadamard, but no new ancilla is being introduced. ■

Note that [7] reports a circuit with $n-3$ $|0\rangle$ ancillae, $8n-17$ T gates, $8n-18$ CNOT gates, and $4n-10$ Hadamard gates.

Proposition 5. A size $n\geqslant 5$ multiple control Toffoli gate TOF^n may be implemented by a circuit using $\lceil\frac{n-3}{2}\rceil$ ancillary qubits residing in an arbitrary state and returned unchanged, by a circuit with

$8n-16$ T gates,
$8n-20$ CNOT gates, and
$4n-10$ Hadamard gates.

Proof. To assist with proving this proposition, define the following gates:

(1) $RTL(a,b,c)$ per Fig. 3, dashed. This is a relative-phase Toffoli gate. The implementation contains nine elementary gates: four T gates, three CNOTs, and two Hadamards.

(2) $RTS(a,b,c)$ per Fig. 3, gates 2–6. This is a relative-phase Toffoli followed by a $V(b,c)$ that removes the last four gates. The circuit contains five elementary gates: two T gates, two CNOTs, and one Hadamard.

(3) $SRTS(a,b,c)$ per circuit (3), dashed. This is a Toffoli gate (as such it is also a type-$\{b\}$ special form relative-phase Toffoli) followed by a $V(a,c)$ that removes the last six gates. $SRTS$ contains nine elementary gates: four T gates, four CNOTs, and one Hadamard.

(4) $RT4L(a,b,c,d)$ per Fig. 4. This is a four-qubit relative-phase Toffoli. It contains eight T gates, six CNOTs, and four Hadamards.

(5) $RT4S(a,b,c,d)$ per Fig. 4, dashed. This is a relative-phase Toffoli-4 $RT4L(a,b,c,d)$ followed by a $V(b,c,d)$ that removes the last eight gates. It is composed of the following elementary gates: four T gates, four CNOTs, and two Hadamards.

We first prove the proposition for the resource count of $n-3$ ancillae, $8n-16$ T gates, $8n-18$ CNOT gates, and $4n-10$ Hadamard gates, and then introduce the $RT4L/RT4S$ gates that further improve the ancilla and CNOT count. The proof relies on the construction found in Fig. 2(c). Assuming qubits are numbered 1 to $2n-3$ and we are attempting to implement $\text{TOF}^n(1,2,...,n-1;2n-3)$, select the gates in Fig. 2(c) as follows:

(1) First gate is $SRTS(n-1,2n-4,2n-3)$.

(2) Next $k=1..n-4$ gates are $RTS(2n-4-k,n-1-k,2n-3-k)$.

(3) Next gate is $RTL(1,2,n)$.

(4) Next $k=1..n-4$ gates are $RTS^{-1}(n-1+k,k+2,n+k)$ (inverses of the gates in item 2 read in reverse order).

(5) Next gate is $SRTS^{-1}(n-1,2n-4,2n-3)$ (this is the matching inverse pair for the gate in item 1).

(6) Next $k=1..n-4$ gates are $RTS(2n-4-k,n-1-k,2n-3-k)$ (same as item 2).

(7) Next gate is $RTL^{-1}(1,2,n)$ (this is the matching inverse for the gate in item 3).

(8) Last $k=1..n-4$ gates are $RTS^{-1}(n-1+k,k+2,n+k)$ (same as item 4).

Observe that the desired preliminary gate counts are satisfied. Next step is introducing $RT4L/RT4S$ gates to replace as many RTL and RTS as possible.

(1) First, replace the circuit $RTS(n,3,n+1)RTL(1,2,n)RTS^{-1}(n,3,n+1)$ (last gate in item 2, the gate in item 3, and first gate in item 4) with $RT4L(1,2,3,n+1)$ and $RTS(n,3,n+1)RTL^{-1}(1,2,n)\ RTS^{-1}(n,3,n+1)$ (last gate in item 6, the gate in item 7, and first gate in item 8) with $RT4L^{-1}(1,2,3,n+1)$. Note that this procedure may

TABLE I. Optimization of the multiple control Toffoli gates using RTOF[3] and RTOF[4] gates.

Gate	Source	Optimization goal	No. T	No. CNOT	No. H	No. P/Z	No. ancillae	Ancillae type
TOF[4]	[10]	T	15	35	6	3	1	$\|0\rangle$
	[7]	T	15	14	6	0	1	$\|0\rangle$
	Ours	T, CNOT	15	12	6	0	1	$\|0\rangle$
	[10]	T	16	54	6	6	1	$\|x\rangle$
	cc-iX [12]	CNOT	22	20	8	0	1	$\|x\rangle$
	Ours	T, CNOT	16	14	6	0	1	$\|x\rangle$
TOF[5]	[10]	T	23	63	10	6	2	$\|00\rangle$
	[7]	T	23	22	10	0	2	$\|00\rangle$
	Ours	T, CNOT	23	18	10	0	1	$\|0\rangle$
	[10]	T	28	90	10	13	2	$\|xx\rangle$
	cc-iX [12]	CNOT	38	36	16	0	2	$\|xx\rangle$
	Ours	T, CNOT	24	20	10	0	1	$\|x\rangle$
TOF[6]	[10]	T	31	94	14	9	3	$\|000\rangle$
	[7]	T	31	30	14	0	3	$\|000\rangle$
	Ours	T, CNOT	31	24	14	0	2	$\|00\rangle$
	[10]	T	40	132	14	20	3	$\|xxx\rangle$
	cc-iX [12]	CNOT	46	52	24	0	3	$\|xxx\rangle$
	Ours	T, CNOT	32	28	14	0	2	$\|xx\rangle$
TOF[11]	[10]	T	71	232	34	24	8	$\|00000000\rangle$
	[7]	T	71	70	34	0	8	$\|00000000\rangle$
	Ours	T, CNOT	71	54	34	0	4	$\|0000\rangle$
	[10]	T	100	328	34	55	8	$\|xxxxxxxx\rangle$
	cc-iX [12]	CNOT	134	132	64	0	8	$\|xxxxxxxx\rangle$
	Ours	T, CNOT	72	68	34	0	4	$\|xxxx\rangle$
TOF^n,$n \geqslant 5$	[10]	T	8n-17	N/A	N/A	N/A	n-3	$\|00...0\rangle$
	[7]	T	8n-17	8n-18	4n-10	0	n-3	$\|00...0\rangle$
	Ours	T, CNOT	8n-17	6n-12	4n-10	0	$\lceil\frac{n-3}{2}\rceil$	$\|00...0\rangle$
	[10]	T	12n-32	N/A	N/A	N/A	n-3	$\|xx...x\rangle$
	cc-iX [12]	CNOT	16n-42	16n-44	8n-24	0	n-3	$\|xx...x\rangle$
	Ours	T, CNOT	8n-16	8n-20	4n-10	0	$\lceil\frac{n-3}{2}\rceil$	$\|xx...x\rangle$

only apply for $n \geqslant 5$. It furthermore reduces the CNOT count from $7 = 2+3+2$ to 6 twice, for a total saving of two CNOTs. Finally, observe that the qubit n is no more used. Thus, we save one ancillary qubit worth of computational space.

(2) For $k = 1..\lceil\frac{n-6}{2}\rceil$ we introduce four $RT4S$ gates by replacing a pair of neighboring RTS on the left- and right-hand sides of the previous step. In particular, we replace $RTS(n+2k,\ 2k+3,\ n+2k+1)RTS(n-1+2k,\ 2k+2,\ n+2k)$ (item 2) with $RT4S(n-1+2k,\ 2k+2,\ 2k+3,\ n+2k+1)$ and $RTS^{-1}(n-1+2k, 2k+2, n+2k)RTS^{-1}(n+2k,\ 2k+3,\ n+2k+1)$ (item 4) with $RT4S^{-1}(n-1+2k,\ 2k+2,,\ 2k+3,\ n+2k+1)$, and similarly in the second half of the circuit (items 6 and 8). Observe that this operation does not change the gate counts, but frees up qubit $n+2k$ that is no more used, providing a reduction of one ancilla.

The total reductions from the above construction are a pair of CNOT gates, and $\lfloor\frac{n-3}{2}\rfloor$ qubits, leading to the resource counts as announced in the statement of the proposition.

Looking at the following circuit helps visualize all replacements and gate counts:

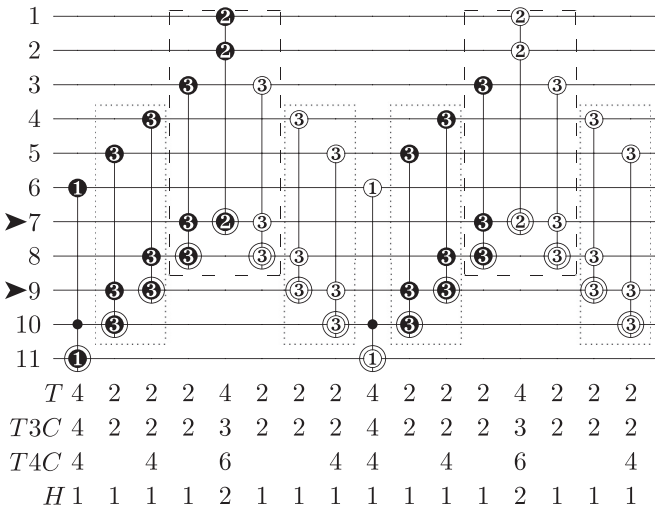

In the above, dashed gates are replaced with $RT4L(1,2,3,8)$ and its inverse, freeing qubit 7, and dotted gates are replaced with $RT4S(8,4,5,10)$ and its inverse, freeing qubit 9. The line starting with "T" reports the T count, the line starting with "$T3C$" reports the CNOT count when only RTOF[3] are being used, the line starting with "$T4C$" reports the CNOT count when

RTOF[4] are allowed, and the line starting with "H" reports the Hadamard gate count. ■

We summarize the results in Table I and compare them against best known. The names of the columns are self-explanatory. Observe that [14] features multiple control Toffoli implementations using $12n - 34$ two-qubit gates over a circuit with $n - 3$ ancillae. In comparison, our implementation uses $8n - 20$ CNOT gates over a circuit with only $\lceil \frac{n-3}{2} \rceil$ ancillae. It is furthermore interesting to highlight that in terms of implementing a multiple control Toffoli gate the cost of moving away from using unrestricted ancillae to ancillae residing in the state $|0\rangle$ is only one T gate, but in terms of the CNOTs, it is a noticeable term, $2n - 8$.

V. OPEN PROBLEMS

The problem of systematically synthesizing and analyzing multiple control relative-phase Toffoli implementations—both unrestricted as well as the special form, is important to address next. The results of such a search could be used directly to optimize implementations of the multiple control Toffoli gates, arithmetic parts of quantum algorithms, and reversible circuits.

How efficient may a relative-phase multiple control Toffoli gate implementation be? In the three-qubit case the answer is as follows: It requires at least three CNOTs as a circuit over CNOT and any single-qubit gates library, as otherwise, per Corollary 1, we would come to a contradiction with any lower CNOT gate count [13]. If it is established that the Toffoli gate requires 7 T gates in the presence of ancillae, a similar argument can be applied towards showing that any relative-phase Toffoli gate requires at least 4 T gates as a circuit over the Clifford + T library.

The reported constructions obtain best solutions simultaneously for two circuit cost metrics arising from different considerations, the CNOT count and the T count. It may be that this is not a coincidence. Is there a relation between these two resource counts?

VI. CONCLUSION

In this paper, we reported an approach for systematic optimization of quantum circuits via replacing suitable pairs of the multiple control Toffoli gates with their relative-phase implementations. This operation preserves the functional correctness. However, since the relative-phase Toffoli gates are easier to implement than their regular counterparts, the advantage can be witnessed through the optimized resource counts. We have furthermore illustrated the advantage via optimizing and, when applicable, explaining the nature of best known implementations of the multiple control Toffoli gates. Our demonstrated optimizations include a simultaneous optimization of the T count by a factor of $\frac{4}{3}$ in the leading constant, the CNOT count by a factor 2 in the leading constant, and the number of ancillary qubits by a factor of 2 in the leading constant. The above refers to the optimization of the circuit implementing the multiple control Toffoli gate using arbitrary ancillae, whose construction resulted from employing the relative-phase Toffoli gates.

ACKNOWLEDGMENTS

I wish to thank anonymous reviewers for their useful comments.

This material was based on work supported by the National Science Foundation, while working at the foundation. Any opinion, finding, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation.

[1] C. Monroe and J. Kim, Science **339**, 1164 (2013).
[2] M. H. Devoret and R. J. Schoelkopf, Science **339**, 1169 (2013).
[3] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, and H. Weinfurter, Phys. Rev. A **52**, 3457 (1995).
[4] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, New York, 2000).
[5] Helpful MATHEMATICA calculations are available online at http://www.umiacs.umd.edu/~dmaslov/papers/mathematicacomputations.txt.
[6] A. M. Childs, R. Cleve, E. Deotto, E. Farhi, S. Gutmann, and D. A. Spielman, *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing STOC '03* (ACM, New York, NY, 2003), pp. 59–68.
[7] P. Selinger, Phys. Rev. A **87**, 042302 (2013).
[8] S. Bravyi and A. Kitaev, Phys. Rev. A **71**, 022316 (2005).
[9] A. Sorensen and K. Molmer, Phys. Rev. Lett. **82**, 1971 (1999).
[10] M. Amy, D. Maslov, and M. Mosca, IEEE Trans. CAD **33**, 1476 (2014).
[11] G. Song and A. Klappenecker, Quantum Inf. Comput. **4**, 361 (2004).
[12] Computer code QUIPPER 0.5 [http://www.mathstat.dal.ca/~selinger/quipper/doc/frames.html].
[13] V. V. Shende and I. L. Markov, Quantum Inf. Comput. **9**, 461 (2009).
[14] D. Maslov, G. W. Dueck, D. M. Miller, and C. Negrevergne, IEEE Trans. CAD **27**, 436 (2008).