T-depth Optimization for Fault-Tolerant Quantum Circuits

Philipp Niemann*, Anshu Gupta†, and Rolf Drechsler*‡
*Cyber-Physical Systems, DFKI GmbH, Bremen, Germany
†Indian Institute of Engineering Science and Technology, Shibpur, India
‡Department of Computer Science, University of Bremen, Bremen, Germany
philipp.niemann@dfki.de, drechsler@uni-bremen.de

***Abstract*—The Clifford+T gate library consisting of Hadamard, T, and $CNOT$ gates has attracted much interest in quantum circuit synthesis, particularly due to its applicability to fault tolerant realizations. Since fault tolerant implementations of the T gate have very high latency, recent work in this area is aiming at minimizing the number of T stages, referred to as the T-*depth*. In this paper, we present an approach to exploit additional ancilla qubits in the mapping of reversible circuits consisting of multiple controlled Toffoli gates (MCT gates) into Clifford+T quantum circuits, with the primary optimization objective to minimize the T-depth. Our proposed approach takes advantage of and generalizes earlier work on corresponding mapping algorithms. An experimental evaluation shows that our approach leads to a significant T-depth reduction compared to earlier approaches.**

I. Introduction

Quantum computing [1] is a flourishing and very attractive research area. Inheriting properties from quantum mechanics, quantum computing allows for solving certain problems exponentially faster than any known classical algorithm. In contrast to Boolean logic, quantum bits (qubits) can not only represent the classical 0 and 1 states, but also any linear combination of both (superposition). Hence, quantum circuits have an advantage of being able to perform massively parallel computations in one time step leading to a significant speed-up for several interesting and practically relevant problems such as integer factorization [2], database search [3], etc.

Due to this fact, the synthesis of quantum circuits, i.e. the (automatic) generation of quantum circuits realizing a given quantum functionality, has become an active research area and many theoretical implementations for this kind of circuits have been presented. As quantum logic synthesis is a very complex and challenging problem, Boolean functions—which constitute a major component in many quantum algorithms—are usually treated separately using a two-step approach: the desired Boolean function is first realized in terms of a reversible circuit, i.e., by means of classical reversible logic gates, after which the resulting circuit is transformed to a functionally equivalent quantum circuit by mapping each reversible gate to a corresponding cascade of quantum gates.

For the synthesis of, or mapping to, quantum circuits, several gate libraries were introduced in the past like the 1-qubit and CNOT library, the NCV library, and the Clifford+T library. Recently, there has been particular interest in quantum circuits using the Clifford+T gate library [4] because it is *universal* (i.e., it can not only realize Boolean functions, but also general quantum functionality with arbitrary precision) as well as *robust* (i.e. fault-tolerant implementations of the gates are known for most technologies that are considered promising for large-scale quantum computing).

In the optimization of quantum circuits based on Clifford+T gate library, major objectives are to minimize the T-count (number of T gates, [5]) and particularly the T-depth of the circuit (the number of T-stages where each stage consists of one or more T or $T^\dagger$ gates that can operate simultaneously on separate qubits) [6]–[8]. This is due to the high cost of fault tolerant implementations of the T gate, exceeding the cost of Clifford group gates ($CNOT, H, S$ gates) by as much as a factor of a hundred or more [8].

Recent work has shown the prospects of dedicated mapping approaches for reversible circuits [9]. However, while all such approaches require additional helper qubits (so-called *ancilla* qubits) to perform the mapping, only a few special configurations of those have been considered thus far and in all of these either *clean* ancillae (i.e. helper qubits initialized to a certain pre-defined state) or *dirty* ancillae (i.e. helper qubits with an arbitrary quantum state) are used.

In this paper, we adapt previous mapping algorithms and generalize their ideas with the motive to further reduce the T-depth of the circuits. We present:

1) An improved algorithm to map multiple-controlled Toffoli (MCT) gates into Clifford+T circuits using *any number of ancilla*.
2) Moreover, the algorithm—for the first time—-supports an arbitrary *combination of clean and dirty ancillae* which allows us to "borrow" temporarily unused qubits in order to further reduce the cost the mapping.

This algorithm is then used to map reversible circuits consisting of MCT gates into Clifford+T quantum circuits, with the help of the specified number of ancilla available at each reversible gate. Our approach for the improved reversible circuit mapping allows for significant T-depth reduction compared to state-of-the-art technology mapping methods proposed earlier for Clifford+T circuits. As confirmed by an experimental evaluation, improvements of the T-depth of up to 60% can be observed. This clearly demonstrates the efficiency of our approach.

The remainder of this paper is structured as follows. The next section introduces notations and preliminaries. Section III discusses related work. Based on this, Section IV presents our decomposition method which generates Clifford+T quantum circuits for MCT gates. Experimental results are presented in Section V. Finally, the paper is concluded in Section VI.

II. Background and Preliminaries

To keep the paper self-contained, this section briefly introduces the basics of reversible and quantum circuits.

A. Reversible Functions and Circuits

A multi-output Boolean function $f\colon \{0,1\}^n \to \{0,1\}^n$ is called *reversible* if f is *bijective*, i.e., if each input pattern is mapped to a unique output pattern and vice versa.

Reversible functions on n bits are realized by reversible circuits consisting of at least n lines (carrying binary values). These reversible circuits are cascades of reversible gates belonging to a particular gate library, with no fan-out or feedback. The most popular gate library is given by multiple-controlled Toffoli (MCT) gates which are defined as follows:

Definition 1 (Multiple-Controlled Toffoli gate)**.** *Given a set of circuit lines $X = \{x_1, x_2, \ldots, x_n\}$, an m-controlled Toffoli gate $T(C;t)$ is given by a set of control lines $C = \{x_{c_1}, \ldots, x_{c_m}\} \subset X$ (where $|C| = m$), and a target line $t \in X \setminus C$. On the target line, the gate performs the mapping $t \mapsto t \oplus (x_{c_1} \wedge \ldots \wedge x_{c_m})$, i.e. the target line is inverted if, and only if, all control lines are in the 1-state. All other lines pass through unaltered.*

An MCT gate with no control line always inverts the target line and is thus the well-known *NOT gate*. An MCT gate with a single control line is called a *controlled-NOT (CNOT) gate* (also called a *Feynman gate*). The case of two control lines is the original *Toffoli* gate. Here in this paper, we follow the normal convention of using $\bigoplus$ to indicate target lines and $\bullet$ to indicate control connections.[1]

B. Quantum Circuits

A quantum circuit is a model of quantum computation representing a sequence of quantum operations. Each operation is represented by a quantum gate and the circuit is a cascade of quantum gates where the circuit lines represent the *qubits (quantum bits)* of a quantum system.

In contrast to classical bits which can only assume two discrete states, *qubits* can represent any combination of the classical Boolean values. More precisely, the state space of a qubit is a 2-dimensional Hilbert space such that all possible states can be written as $|\psi\rangle = a|0\rangle + b|1\rangle = \binom{a}{b}$ where $|0\rangle, |1\rangle$ denote the computational basis states (associated with the classical Boolean values) and $a, b \in \mathbb{C}$ such that $|a|^2 + |b|^2 = 1$. Analogously, the state space of an n-qubit quantum system has 2^n basis states ($|0\ldots00\rangle, |0\ldots01\rangle, \ldots, |1\ldots11\rangle$) such that the state vector has 2^n dimensions.

[1] In this paper, we only consider "positive" controls (sensitive to the 1-state). The extension to "negative" controls (sensitive to the 0-state) is straightforward and has no impact on the T-gate cost of the mappings.

Symbol	*Name*	*Unitary Matrix*
H	Hadamard gate	$\frac{1}{\sqrt{2}}\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$
X	NOT gate	$\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$
	$CNOT$ gate	$\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}$
T	T gate	$\begin{pmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{pmatrix}$
S	S gate	$\begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix}$
αZ	αZ gate ($\alpha \in \{\pm i, \pm 1\}$)	$\begin{pmatrix} \alpha & 0 \\ 0 & -\alpha \end{pmatrix}$

Fig. 1. Quantum Operations

Definition 2 (Quantum Gate)**.** *In general, a quantum gate acting on n qubits/lines is represented by a $2^n \times 2^n$ unitary matrix. As in reversible circuits, circuit lines may be designated as control lines, which have the effect of applying the gate to the target qubit whenever (possibly in superposition) the control qubit is in the $|1\rangle$-state.*

In this work, we consider the Clifford+T gate library which consists of $NOT, CNOT, H, S, S^\dagger, T, T^\dagger$ gates. It is to be noted that since the matrices of interest are unitary, the adjoint of a matrix (denoted by $\dagger$) is its inverse. Commonly used quantum operations in the Clifford+T gate library and their corresponding unitary matrices are shown in Fig. 1. Quantum gates can be applied in parallel (in one stage) if they act on disjoint sets of qubits. The resulting matrix of this parallelized execution is computed using the Kronecker product of the individual gate matrices.

Two important metrics for the cost of a Clifford+T circuit are given by the T-count, i.e. the total number of T and $T^\dagger$ gates in the circuit, and the *T-depth*, i.e. the number of stages of the circuit that contain one or more T or $T^\dagger$ gates performed concurrently on separate qubits.

Example 1. *Figure 2a shows the realization of the Toffoli gate, i.e. a two-controlled MCT gate, in Clifford+T [8, Fig. 7(a)]. The gates are to be applied from left to right. The circuit has a T-count of seven and a T-depth of three.*

Additional helper qubits, so-called *ancilla* qubits, are frequently employed for storing intermediate results during quantum computation. These can either be initialized to a predefined state (usually $|0\rangle$, denoted as *clean* ancillae) or have an arbitrary quantum state (denoted as *dirty* ancillae). In both cases, ancillae should be returned back to their original state after the computation in order to re-use them for future computations.

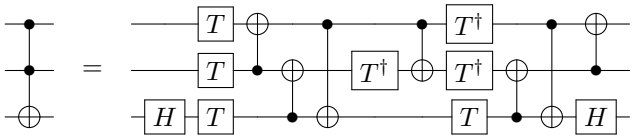

(a) Mapping without ancilla qubits

(b) Mapping with one clean ancilla

Fig. 2. Mapping of Toffoli gate into Clifford+T circuits [8]

Example 2. *Consider the quantum circuit in Fig. 2b (taken from [8, Fig. 15]). It also realizes the Toffoli gate, though with the use of one additional clean ancilla. This ancilla allows us to reduce the T-depth to two, while the T-count remains seven.*

III. Related Work

The mapping of reversible MCT circuits to quantum circuits is usually conducted in two steps. First, the MCT gates with $c \geq 3$ controls are decomposed into MCT gates with less than three controls, i.e. NOT, CNOT, and Toffoli (NCT) gates. Afterwards, each NCT gate is mapped individually to an equivalent cascade of quantum gates.

The complexity and cost of this decomposition as well as the mapping of NCT gates to Clifford+T is linear in the number of controls. More precisely, using the well-known decompositions of MCT into NCT gates proposed by Barenco et al. [10] as well as the realization of the Toffoli gate in Clifford+T shown in Fig. 2a, we obtain the following

Lemma 1 (Lemma 7.2 from [10])**.** *An MCT gate with $c \geq 3$ controls can be mapped directly to a circuit that consists of $4(c-2)$ Toffoli gates using $(c-2)$ dirty ancilla (see Fig. 4a). The resulting circuit has a T-depth of $12(c-2)$.*

Lemma 2 (Lemma 7.3 from [10])**.** *A MCT gate with $c \geq 3$ controls can be mapped directly to a circuit consisting of 4 MCT gates with fewer controls with the help of only one dirty ancilla (see Fig. 4b). These 4 MCT gates have enough dirty ancilla to be further mapped to Toffoli gates using Lemma 1.*

As noted by Nielsen and Chuang [1], the availability of clean ancillae instead of dirty ancillae can significantly reduce the cost of the above decompositions. More precisely, the presence of one clean ancilla allows for omitting the last gate in the construction from Lemma 2, while the construction from Lemma 1 simplifies to $2(c-2)+1$ Toffoli gates, i.e. a T-depth of $6(c-2)+3$, for $(c-2)$ clean ancillae. E.g., Fig. 4c shows the simplified version of the decomposition of the 4-controlled MCT gate shown in Fig. 4a.

Further improvements of the mappings with better T-depth have been proposed in [9], based on structural investigations at the quantum gate level:

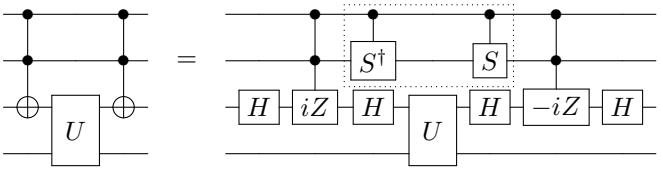

Fig. 3. Reducing T-depth of Toffoli gate pairs using controlled $\pm iZ$-gates

Lemma 3 (Lemma 1 in [9])**.** *An MCT gate with $c \geq 3$ controls can be realized with a T-depth of $4(c-1)$ using $(c-2)$ dirty ancillae.*

Note that the original formulation of Lemma 3 does not cover the case $c = 3$, but it is also valid for $c = 3$ according to the results from [11].

Lemma 4 (Lemma 2 in [9])**.** *An MCT gate with $c > 3$ controls can be realized with a T-depth of:*

- $8(c-2)-4$, *using one dirty ancilla.*
- $6(c-2)+2$, *using one clean ancilla (if c is odd).*
- $6(c-2)$, *using one clean ancilla (if c is even).*

Note that the original formulation of Lemma 4 only covers the case $c \geq 5$. However, it can be extended to the case $c = 4$ using the "Miller mapping" from [9, Fig. 3]. For $c = 4$, the four MCT gates are in fact two pairs of Toffoli gates with the same set of controls. As nothing happens to these control qubits between the application of the corresponding Toffoli gates, the circuit re-writing from Fig. 3 can be performed to replace the Toffoli gates by controlled $\pm iZ$-gates together with controlled $S/S^\dagger$ gates (adapted from [9, Proof of Lemma 1]).

The controlled-$S^\dagger/S$ gates cancel, while a two-controlled $\pm iZ$-gate can be realized in T-depth 2 [9]. Thus, the T-depth for realizing one pair of Toffoli gates is reduced from 6 to 4. Overall, the T-depth of the complete 4-controlled MCT gate becomes $12 = 6(4-2) = 8(4-2)-4$ as stated in Lemma 4. Finally, a 3-controlled MCT gate can be realized with T-depth 6 using one clean ancilla according to [11].

IV. Proposed Decomposition and Mapping Scheme

In this section we propose an approach for the decomposition of MCT gates into smaller MCT gates with fewer controls using a definite number of ancillae. The MCT gate is decomposed in such a way that after mapping it into Clifford+T circuits, it leads to further reduction of T-depth which is our main objective.

In contrast to earlier mappings (c.f. Lemma 1–4), where only the cases of exactly one or $(c-2)$ ancillae have been considered for the decomposition, our approach generalizes to any number of ancilla ranging from one to ∞, taking advantage of the inverse proportionality between number of ancilla and T-depth.

The general idea of our approach is to split the decomposition into several stages and re-use the mapping from Lemma 3 in each stage. In the following, let k denote the number of clean ancillae and a_i denote the i-th clean ancilla qubit ($i = 1, \ldots, k$), while d denotes the number of additionally available dirty ancilla qubits.

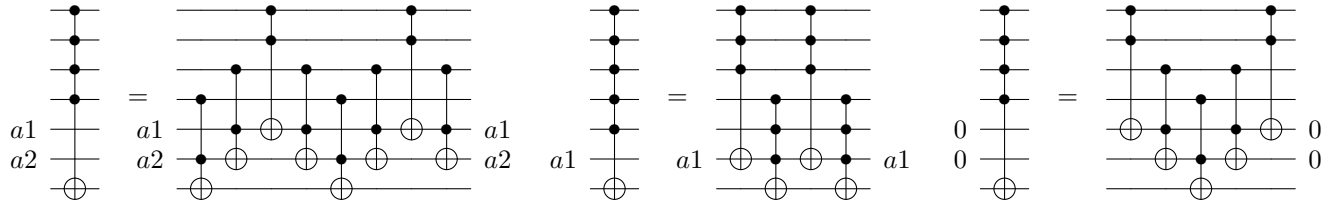

(a) Decomposition of 4-controlled MCT gate using 2 dirty ancilla as [10, Lemma 7.2]

(b) Decomposition of 5-controlled MCT gate using 1 dirty ancilla [10, Lemma 7.3].

(c) Decomposition of 4-controlled MCT gate using 2 clean ancilla [1]

Fig. 4. Decompositions of MCT to NCT gates from Barenco et al. [10] and Nielsen & Chuang [1]

A. Using $1 \leq k \leq \frac{c}{2}$ clean ancillae

The decomposition is divided into three stages.

1) In the first stage, the control lines of a c-controlled MCT gate $T(C, t)$ are divided into $k+1$ groups $C_1, \ldots, C_{k+1}$ of (adjacent) control lines. We apply k MCT gates $T(C_i, a_i)$ $(i = 1, \ldots, k)$ with targets on the clean ancillae. These gates can be applied in parallel as they do not share any control or target.
2) In the second stage, the remaining control lines from C_{k+1} are taken into account using an MCT gate with $(k+|C_{k+1}|)$ controls, namely $T(C_{k+1} \cup \{a_1, \ldots, a_k\}, t)$.
3) The third stage is a copy of the first stage in order to restore the values on the ancilla qubits.

Example 3. *Figure 5a shows the decomposition of a 9-controlled MCT gate using $k = 2$ clean ancillae with the three stages discussed above where the MCT gates in the first/third stage have 3 controls and the gate MCT gate in the center has 5 controls. The overall T-depth of this decomposition if $8 + 16 + 8 = 32$.*

In order to use the mapping from Lemma 3 in the second stage, we require $(k + |C_{k+1}|) - 2$ dirty ancillas. As all clean ancillae have been occupied in the first stage, we may only use the control qubits from the first stage as well as the d additional dirty ancillae for this purpose. Hence,

$$d + |C_1| + \ldots + |C_k| \geq (k + |C_{k+1}|) - 2.$$

Analogously, only the qubits from C_{k+1} and the target qubit t may be used as (dirty) ancilla qubits for the first/third stage in addition to the d extra ancillae. Hence, to employ the mapping from Lemma 3 also in these stages, we need that

$$\begin{aligned} d + |C_{k+1}| + 1 &\geq (|C_1| - 2) + \ldots + (|C_k| - 2) \\ &= |C_1| + \ldots + |C_k| - 2k. \end{aligned}$$

Using that $|C_1| + \ldots + |C_k| + |C_{k+1}| = c$, we obtain

$$\frac{c + 2 - k + d}{2} \geq |C_{k+1}| \geq \frac{c - 2k - d - 1}{2} \tag{1}$$

Example 4. *Consider the case $c = 9$, $k = 2$, $d = 0$. According to Eqn. (1), we have $4.5 \geq |C_{k+1}| \geq 2$. The case $|C_{k+1}| = 3$ is shown in Fig. 5a (T-depth 32 as calculated in Example 3), while the cases $|C_{k+1}| = 2$ and $|C_{k+1}| = 4$ yield T-depths of $12 + 12 + 12 = 36$ and $8 + 20 + 8 = 36$. If a dirty ancilla is available ($d \geq 1$), the case $|C_{k+1}| = 5$ (shown in Fig. 5b) becomes possible which reduces the T-depth to $2 \cdot 2 + 24 = 28$.*

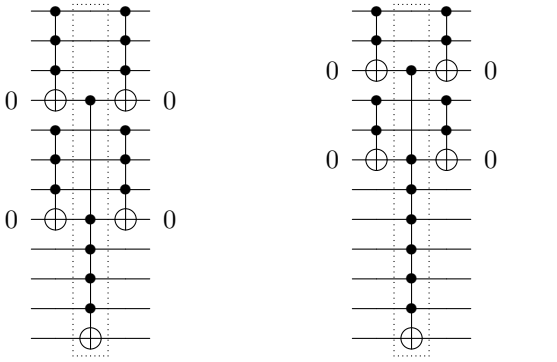

(a) without dirty ancillae

(b) with $d \geq 1$ dirty ancillae

Fig. 5. Decomposition of 9-controlled MCT gate using $k = 2$ clean ancillae

As the T-depth of the second stage is directly correlated with $|C_{k+1}|$, it is beneficial to make C_{k+1} as small as possible and "consume" more controls already in the first/third stage. On the other hand, since all gates in these stages are executed in parallel, the T-depth of the first/third stage is correlated with the maximum number of qubits in $|C_1|, \ldots, |C_k|$ such that this maximum should only be increased if this would allow for transferring a sufficient number of controls from the central stage.

Accordingly, we propose the following algorithm for determining a beneficial partition of the controls into $C_1, \ldots, C_{k+1}$:

1) Start with the largest possible $|C_{k+1}|$ (according to Eq. 1) and distribute the remaining $c - |C_{k+1}|$ controls equally among $C_1, \ldots, C_k$ such that they differ by at most one control.
2) Successively move controls from C_{k+1} to $C_1, \ldots, C_k$ as long as this does not increase the maximum number of controls in $C_1, \ldots, C_k$ and Eq. (1) remains valid.
3) If $k \geq 2$ and at least three controls can be removed from C_{k+1} without violating Eq. (1), move controls to C_1, C_2 (one each) and repeat Step 2. Else, terminate.

Remark 1. *Note that each control that can be moved in Step 2 reduces the T-depth of the second stage by 4 without increasing the T-depth of the other stages. The third step reduces the T-depth of the second stage by 8, while the T-depth of the first as well as the third stage is increased by 4 each (if the previous maximum of the $|C_i|$ was at least 3 controls). In the case that the previous maximum was 2 (i.e. only Toffoli gates in the first stage), the T-depth of the first*

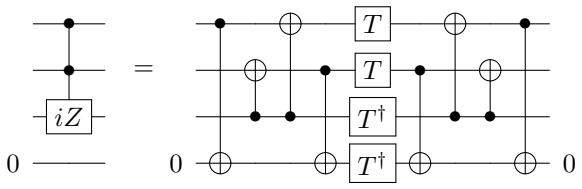

Fig. 6. Mapping of two-controlled iZ-gate into Clifford+T circuit using one clean ancilla (giving a T-depth of 1)

and third stage can be reduced to 2 each (using the same replacement with $\pm iZ$-gates as in Fig. 3), such that the T-depth increases from 2 to 8 per stage. This is why we require at least three movable controls (which in total can reduce the cost of the second stage by 12) in the first place in order to ensure that Step 3 does not worsen the overall cost.

Overall, since $k \leq \frac{c}{2}$, we are always able to find a partition such that all gates in the first/third stage have at least two controls. For the corner case $k = \frac{c}{2}$, we obtain a decomposition with a T-depth of $2 + 4(k-2) + 2 = 4(k-1) = 2(c-2)$.

B. Using $\frac{c}{2} < k \leq (c-2)$ clean ancillae

If $k > \frac{c}{2}$, we are not able to find such a partition and would end up with some C_i only containing one control. In this case, it is more promising to use the corresponding clean ancilla lines in a different way. More precisely, we propose to use $\lfloor \frac{c}{2} \rfloor$ of the clean ancillae as targets for Toffoli gates with controls from C in the first/third stage. Again, these gates can be executed in parallel, have T-depth 2, and consume $2\lfloor \frac{c}{2} \rfloor$ of the original controls.

For the second stage, we recursively apply the strategy from either this or the previous subsection for an MCT gate with $c' = c - \lfloor \frac{c}{2} \rfloor$ controls (out of which $c - 2\lfloor \frac{c}{2} \rfloor$ are the remaining original controls and $\lfloor \frac{c}{2} \rfloor$ are new controls on the already used clean ancillae), $k' = k - \lfloor \frac{c}{2} \rfloor$ clean ancillae and $d' = d + 2\lfloor \frac{c}{2} \rfloor$ dirty ancillae (using the already consumed controls as such).

For the case $k = (c-2)$, this leads to $2\lceil \log_2(c) \rceil - 1$ stages consisting only of Toffoli gates which form a V-shape/binary tree structure, such that the $\pm iZ$-gate replacement can be applied for all but the central stage, hence yielding a T-depth of $4\lceil \log_2(c) \rceil - 1$ in contrast to $6(c-2)+3$ using the naive Nielsen&Chuang mapping discussed above.

C. Using more than $(c-2)$ clean ancillae

If there are still clean ancillae left after applying the decomposition into Toffoli gates only as described above for the case $k = (c-2)$, we can employ these ancillae to further reduce the cost of the Toffoli/$\pm iZ$-gates using the mappings shown in Figs. 2b and 6, respectively. In fact, using one additional ancilla, the T-depth of a Toffoli and a $\pm iZ$-gate can be reduced from 3 to 2 and from 2 to 1, respectively.

Consequently, the cost of a stage can be reduced by 1 if, and only if, the number of unused clean ancillae is greater than or equal to the number of Toffoli/$\pm iZ$-gates in that stage. E.g., for the first stage this is the case when $k - \lfloor \frac{c}{2} \rfloor \geq \lfloor \frac{c}{2} \rfloor$. The maximum reduction that can be achieved in this way is a T-depth of $2\lceil \log_2(c) \rceil$.

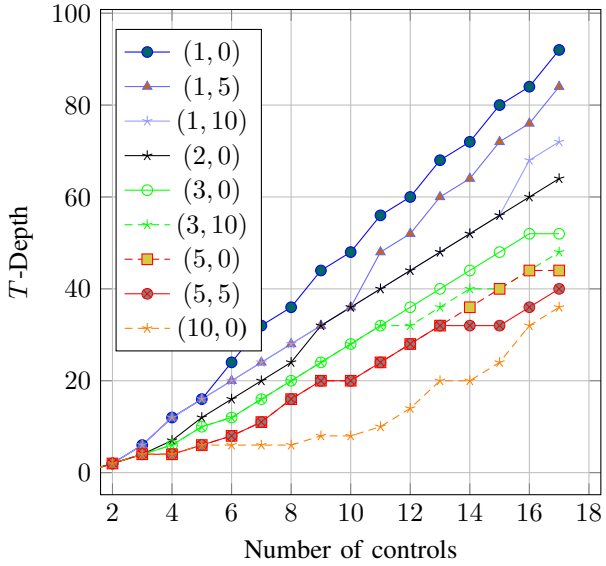

Fig. 7. T-depth of MCT gate decompositions for various combinations of clean/dirty ancillae (k, d)

V. Experimental Evaluation

In this section, we evaluate the proposed decomposition scheme in comparison with previous work. To this end, the proposed scheme has been implemented in C and has been tested on a suite of benchmarks taken from RevLib [12].

A. Impact of clean and dirty ancillae

Figure 7 shows the T-depth of the resulting decompositions for MCT gates with up to $c = 17$ controls using different combinations of clean and dirty ancillae (denoted as (k, d)). Here, the dataset $(1, 0)$ represents the state-of-the-art for $k = 1$ clean ancilla and no dirty ancillae as given by Lemma 4. It can be seen that additional dirty ancillae make it possible to gradually reduce the cost by around 10–25% (c.f. datasets $(1, 5)$ and $(1, 10)$). The minimum T-depth of $4(c-1)$ is achieved for $d = c - 5$ dirty ancillae which can be explained by the fact that in this setting a single Toffoli gate in the outer stages is sufficient and the central MCT gate has $(c - 2 + 1) = c - 1$ which requires $c - 3 = (c - 5 + 2)$ dirty ancillae in order to apply the Lemma 3 mapping.

However, as one can also see from Fig. 7, as second clean ancilla is sufficient to achieve the same reduction—without the need for any additional dirty ancilla (c.f. dataset $(2, 0)$).

Employing a higher number of $k > 2$ clean ancillae allows us to easily outperform the Lemma 3 mapping, while additional dirty ancillae only show an effect for $c > 10$ controls. In fact, for $k = 10$ clean ancillae, dirty ancillae do not have any effect for the considered number of controls. Because of this, we only show the dataset $(10, 0)$.

Moreover, we observed that for $c \leq 30$ the minimal cost was already achieved for $k = c$ clean ancillae, i.e. two additional clean ancilla reduce the cost by 50% compared to the $k = c-2$ mapping (from $4\lceil \log_2(c) \rceil - 1$ to $2\lceil \log_2(c) \rceil$).

TABLE I
EXPERIMENTAL RESULTS

Benchmark				T-depth (and improvement in %)											
ID	L	G	$\max c$	Conv.	(1,0)	Δ_{conv}	(1,5)	$\Delta_{(1,0)}$	(2,0)	$\Delta_{(1,0)}$	(3,0)	$\Delta_{(1,0)}$	(5,0)	$\Delta_{(1,0)}$
9symml_195	10	129	9	2144	1896	-11.6	1640	-13.5	1600	-15.6	1128	-40.5	846	-55.4
max46_240	10	107	8	1640	1384	-15.6	1272	-8.1	1220	-11.8	850	-38.6	612	-55.8
urf3_279	10	14075	8	13355	10666	-20.1	10506	-1.5	8700	-18.4	8082	-24.2	7086	-33.6
sqn_258	10	76	6	751	674	-10.3	674	-0.0	501	-25.7	402	-40.4	268	-60.2
sym9_148	10	210	4	2184	2016	-7.7	2016	-0.0	1218	-39.6	1092	-45.8	840	-58.3
sym10_262	11	194	10	3756	3516	-6.4	2860	-18.7	2828	-19.6	2038	-42.0	1568	-55.4
cm152a_212	12	16	4	120	114	-5.0	114	-0.0	68	-40.4	60	-47.4	44	-61.4
sao2_257	14	88	10	2199	1718	-21.9	1662	-3.3	1662	-3.3	1230	-28.4	982	-42.8
pm1_249	14	35	4	182	166	-8.8	166	-0.0	102	-38.6	92	-44.6	72	-56.6
co14_215	15	30	13	952	952	-0.0	784	-17.6	672	-29.4	560	-41.2	448	-52.9
ham15_109	15	109	4	62	48	-22.6	48	-0.0	35	-27.1	34	-29.2	32	-33.3
inc_237	16	93	7	914	898	-1.8	898	-0.0	654	-27.2	536	-40.3	351	-60.9
cnt3-5_180	16	20	3	70	50	-28.6	50	-0.0	40	-20.0	40	-20.0	40	-20.0
t481_263	17	21	4	108	104	-3.7	104	-0.0	64	-38.5	56	-46.2	40	-61.5
cmb_214	20	18	12	240	176	-26.7	176	-0.0	176	-0.0	128	-27.3	112	-36.4
alu1_198	20	32	3	132	98	-25.8	98	-0.0	68	-30.6	68	-30.6	68	-30.6
mux_246	22	35	6	432	432	-0.0	432	-0.0	340	-21.3	262	-39.4	170	-60.6
frg1_234	31	212	23	5179	4976	-3.9	4920	-1.1	4863	-2.3	3636	-26.9	2925	-41.2
simple_321	235	4227	6	64491	63628	-1.3	63628	-0.0	48335	-24.0	39276	-38.3	25126	-60.5
alu_319	235	15764	36	698563	697322	-0.2	697322	-0.0	675805	-3.1	494900	-29.0	375658	-46.1
Average Improvement (in %)						**-11.1**		**-3.2**		**-21.8**		**-36.0**		**-49.2**

B. Mapping of Reversible Circuits

We evaluated the proposed approach on a suite of benchmarks taken from [12]. Due to page limitations, we only report the results for a restricted set of benchmarks in Table I.

Here, the first four columns describe the benchmark in terms of its name (ID), number of qubits (L), number of gates (G) as well as the maximum number of controls of all gates ($\max c$). The next column (Conv.) denotes the cost obtained using the state-of-the-art mapping approach from [9]. More precisely, we assume the availability of one clean ancilla and map a c-controlled MCT gate using Lemma 3, if there are at least $c-2$ spare qubits, and by applying Lemma 4 (for one clean ancilla) otherwise. In the remaining columns we list the cost of the decompositions obtained by the proposed approach for different combinations of clean and dirty ancillae (denoted as (k,d)) as well as the improvement w.r.t. the state-of-the-art (Δ_{conv}) or w.r.t. the proposed $(1,0)$ mapping ($\Delta_{(1,0)}$).

The proposed $(1,0)$ mapping in many cases shows an improvement of 10–20% over the state-of-the-art mapping with one clean ancilla (11,1% on average). In most cases, however, adding additional dirty ancillae does not have a significant impact (3,2% improvement on average for $(1,5)$ over $(1,0)$). This can be explained by the fact that in most circuits there are already enough unused qubits that can be employed as dirty ancillae anyway. Finally, one can observe that the effect of extra clean ancillae is rather small if all MCT gates in a circuit have only few controls ($\max c = 3,4$), but grows to an average improvement of almost 50% for $k=5$ clean ancillae if $\max c$ is greater.

VI. CONCLUSIONS

In this work, we have enhanced the existing decomposition and mapping approaches for MCT gates with the goal of reducing the T-depth which is one main cost parameter in fault-tolerant quantum computation. Our mapping approach for the first time allows us to take into account an arbitrary number and combination of clean and dirty ancillae. An experimental evaluation confirms the benefits of this approach especially for a large number of controls.

REFERENCES

[1] M. Nielsen and I. Chuang, *Quantum Computation and Quantum Information*. Cambridge Univ. Press, 2000.

[2] P. W. Shor, "Algorithms for quantum computation: discrete logarithms and factoring," *Foundations of Computer Science*, pp. 124–134, 1994.

[3] L. K. Grover, "A fast quantum mechanical algorithm for database search," in *Theory of computing*, 1996, pp. 212–219.

[4] P. O. Boykin, T. Mor, M. Pulver, V. Roychowdhury, and F. Vatan, "A new universal and fault-tolerant quantum basis," *Information Processing Letters*, vol. 75, no. 3, pp. 101–107, 2000.

[5] D. Gosset, V. Kliuchnikov, M. Mosca, and V. Russo, "An algorithm for the T-count," *Quantum Information & Computation*, vol. 14, no. 15-16, pp. 1261–1276, 2014.

[6] N. Abdessaied, M. Soeken, and R. Drechsler, "Quantum circuit optimization by Hadamard gate reduction," in *Reversible Computation*, 2014, pp. 149–162.

[7] D. M. Miller, M. Soeken, and R. Drechsler, "Mapping NCV circuits to optimized Clifford+T circuits," in *Reversible Computation*, 2014, pp. 163–175.

[8] M. Amy, D. Maslov, M. Mosca, and M. Roetteler, "A meet-in-the-middle algorithm for fast synthesis of depth-optimal quantum circuits," *IEEE Trans. on CAD*, vol. 32, no. 6, pp. 818–830, 2013.

[9] N. Abdessaied, M. Amy, M. Soeken, and R. Drechsler, "Technology mapping of reversible circuits to Clifford+T quantum circuits," in *ISMVL*, 2016, pp. 150–155.

[10] A. Barenco, C. H. Bennett, R. Cleve, D. DiVinchenzo, N. Margolus, P. Shor, T. Sleator, J. Smolin, and H. Weinfurter, "Elementary gates for quantum computation," *The American Physical Society*, vol. 52, pp. 3457–3467, 1995.

[11] M. Amy, D. Maslov, and M. Mosca, "Polynomial-time T-depth optimization of Clifford+T circuits via matroid partitioning," *IEEE Trans. on CAD of Integrated Circuits and Systems*, vol. 33, no. 10, pp. 1476–1489, 2014.

[12] R. Wille, D. Große, L. Teuber, G. W. Dueck, and R. Drechsler, "RevLib: an online resource for reversible functions and reversible circuits," in *Int'l Symp. on Multiple-Valued Logic*, 2008, pp. 220–225, RevLib is available at http://www.revlib.org.