

# Improved Synthesis of Clifford+T Quantum Functionality

Philipp Niemann\*  
\*Cyber-Physical Systems  
DFKI GmbH  
D-28359 Bremen, Germany  
Email: philipp.niemann@dfki.de

Robert Wille\*†  
†Institute for Integrated Circuits  
Johannes Kepler University Linz  
A-4400 Linz, Austria  
Email: robert.wille@jku.at

Rolf Drechsler\*‡  
‡Group for Computer Architecture  
University of Bremen  
D-28359 Bremen, Germany  
Email: drechsler@uni-bremen.de

**Abstract**—The Clifford+T library provides robust and fault-tolerant realizations for quantum computations. Consequently, (logic) synthesis of Clifford+T quantum circuits became an important research problem. However, previously proposed solutions are either only applicable to very small quantum systems or lead to circuits that are far from being optimal—mainly caused by a local, i.e. column-wise, consideration of the underlying transformation matrix to be synthesized. In this paper, we suggest an improved approach that considers the matrix globally and, by this, overcomes many of these drawbacks. Preliminary evaluations show the promises of this direction.

## I. INTRODUCTION

Due to its potential for significant speed-ups, quantum computation [1] has attracted significant research attention in the recent years. Here, operations are conducted by means of quantum bits (*qubits*) rather than conventional bits, which can not only assume Boolean basis states, but also superpositions of them. This enables significant speed-ups for several interesting and practically relevant problems such as factorization, database search, or the simulation of chemical dynamics [1]. To this end, complex quantum operations are usually described in terms of a cascade of *quantum gates* which eventually form a *quantum circuit*.

In this regard, the *Clifford+T gate library* [2] provides a convenient set of quantum gates because it is universal (i.e. any quantum functionality can be realized with it up to an arbitrary small error  $\epsilon$ ), fault-tolerant (i.e. robust to many faults which easily can occur in quantum computations), and physical implementations are expected to be available for the most promising and large-scale quantum computation technologies. Consequently, how to synthesize the desired quantum functionality in terms of elementary quantum gates from this library became an important research problem.

First approaches for synthesis followed a two-stage scheme in which the desired quantum functionality is first synthesized into a quantum circuit composed of arbitrary one-qubit gates and so-called controlled NOT gates (using solutions as e.g. proposed in [3], [4]) and, afterwards, the resulting gate structure is individually *compiled* to Clifford+T gates (using solutions as e.g. proposed in [5]). But this does not only require a tremendous number of gates, but also yields just an approximation of the desired quantum functionality. In order to overcome this drawback, researchers considered *exact* synthesis of Clifford+T quantum circuits in [6]–[8], i.e. the

desired quantum functionality is realized without any rounding errors rather than being approximated with respect to an  $\epsilon$ .

However, while guaranteeing exactness, the respective schemes are either applicable to very small quantum systems only (e.g. [6], [7]) or synthesize the given transformation matrix (representing the desired quantum functionality) in a local fashion, i.e. column by column (e.g. [8]). The latter is disadvantageous since the manipulation of a single column frequently makes the remaining columns harder to synthesize and often several columns could be considered much more efficiently in one step rather than in several individual steps. Overall, this frequently leads to quantum circuits which are “far from optimal”—as confirmed by the authors of [8].

In this work, we investigate these drawbacks and suggest an improved approach that considers the matrix globally rather than locally, i.e. a solution is suggested which always keeps track of the entire matrix. This eventually yields an alternative and improved synthesis algorithm which has strong potential to realize the desired quantum functionality significantly cheaper than before. A preliminary implementation of the proposed idea shows the promises of this direction and motivates a more detailed investigation of a global matrix consideration during synthesis.

The remainder of this work is structured as follows: The next section reviews the background on quantum circuits and the considered Clifford-T library. Then, Section III discusses the drawbacks of previously proposed synthesis approaches for Clifford-T quantum circuits followed by suggestions for improving this state-of-the-art. Section IV summarizes preliminarily obtained results showing the promises of the proposed method. Finally, the paper is concluded in Section V.

## II. BACKGROUND

Quantum systems are composed of *qubits*, which can be in one of the *basis states*  $|0\rangle$  and  $|1\rangle$  or within a *superposition*  $\alpha|0\rangle + \beta|1\rangle$  for complex-valued  $\alpha, \beta$  with  $|\alpha|^2 + |\beta|^2 = 1$ . Accordingly, an  $n$ -qubit quantum system can be in one of the  $2^n$  basis states ( $|0\dots 00\rangle, |0\dots 01\rangle, \dots, |1\dots 11\rangle$ ) or a superposition of these states (represented by a *state vector* of dimension  $2^n$ ). The evolution of a quantum state due to a quantum operation can be described by a *unitary transformation matrix*.

Realizations of (complex) quantum operations are represented by *quantum gates*  $g_i$  which eventually form a *quantum circuit*  $G = g_1 \dots g_d$  with  $1 \leq i \leq d$ . The unitary matrix of the entire circuit is computed as the matrix product of the individual gate matrices (in reversed order). The *Clifford+T gate library* [2] represents a set of quantum gates which is universal (any quantum operation, i.e. unitary transformation matrix, can be realized up to an arbitrary precision) as well as *fault-tolerant* (i.e. robust implementations of these gates are known for most technologies that are considered promising for large-scale quantum computers).

The most elementary gates in the Clifford+T library comprise the *Hadamard* operation  $H$  (setting a qubit into a superposition), the *NOT* operation  $X$  (flipping the basis states  $|0\rangle$  and  $|1\rangle$ ), as well as the *phase shift* operations  $T$  ( $\pi/4$  gate),  $S = T^2$  (Phase gate), and  $Z = S^2$ . The corresponding unitary matrices are defined as

$$H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}, \quad X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad T = \begin{pmatrix} 1 & 0 \\ 0 & \omega \end{pmatrix}$$

where  $\omega = \frac{1+i}{\sqrt{2}} = e^{i\pi/4}$ .

Besides that, also controlled operations on multiple qubits are required to achieve a universal gate library. Here, the state of the additional *control qubits* determines which operation is performed on the target qubit. The most elementary representative for this is the *controlled NOT* (CNOT) operation on two qubits whose transformation matrix is defined by

$$\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}.$$

This operation performs a NOT operation on the target qubit if, and only if, the control is in the  $|1\rangle$ -state.

By employing additional controls, the basic Clifford+T operations can be applied on a dedicated subset of basis states. Transferring this to the concatenation of quantum gates (matrix multiplication) corresponds to altering only a certain subset of columns in the original matrix. However, these constructions have a high realization cost in terms of primitive Clifford+T gates. In fact, they require many  $T$  gates which are significantly more complex to be realized than Clifford group gates ( $H$ ,  $S$ , CNOT). A detailed overview on the costs of multiple-controlled Clifford+T gates is given in Table I (in terms of  $T$ -depth, assuming the availability of one additional *ancilla* qubit, based on [6], [8]).

### III. SYNTHESIS OF CLIFFORD+T CIRCUITS

In this work, we consider the synthesis of quantum functionality to elementary circuit descriptions based on the Clifford+T library. More precisely, the task is considered in which a quantum functionality represented in terms of a transformation matrix  $F$  is decomposed into a sequence  $G = g_0 \dots g_d$  of elementary quantum operations (i.e. quantum gates  $g_i$  with  $0 \leq i < d$ ). The resulting sequence eventually forms a quantum circuit and is supposed to be composed of Clifford+T

TABLE I  
COST METRIC FOR MULTIPLE-CONTROLLED CLIFFORD+T GATES

#CONTROLS	CNOT	H	T
0	0	0	1
1	0	1	5
2	2	7	21
3	12	21	33
4	32	33	69
5	68	69	105
6	104	105	129
$\vdots$	+24 per add. control		

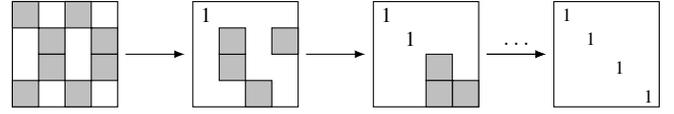


Fig. 1. Local, i.e. column-wise, synthesis scheme [8].

gates as reviewed above. In the following, we re-visit the current state-of-the-art approach which has been proposed for this purpose and discuss its main drawbacks. Based on this, we sketch the general idea of an improved quantum circuit synthesis proposed in this work.

#### A. State-of-the-Art and Motivation

Fig. 1 sketches the current state-of-the-art in the synthesis of Clifford+T quantum circuits as proposed by Giles and Selinger [8]. The main idea is to apply quantum gates so that, eventually, the given matrix to be synthesized (sketched on the left-hand side of Fig. 1) is transformed to the identity matrix (sketched on the right-hand side of Fig. 1). To this end, the matrix is transformed column by column (as sketched in the middle of Fig. 1). For each column, three steps are applied:

- Eliminate superposition*, i.e. apply quantum gates so that all multiple non-zero matrix entries in the column are combined to a single non-zero entry.
- Move to diagonal*, i.e. apply quantum gates which move the remaining non-zero entry to the matrix' diagonal.
- Remove phase shifts*, i.e. apply quantum gates which transform the diagonal entry to 1—eventually yielding a column of the identity matrix.

Each of these steps is achieved by using so-called *two-level operations* that modify pairs of entries in the current column. More precisely, the following operations are utilized:

- *Combine entries*:  $(a, b) \Rightarrow \frac{1}{\sqrt{2}}(a + b, a - b)$
- *Exchange entries*:  $(a, b) \Rightarrow (b, a)$
- *Modify phase shift*:  $(a, b) \Rightarrow (a, b \cdot \omega)$  where  $\omega = e^{i\pi/4}$ .

**Example 1.** Consider the matrix in Fig. 2a which represents a quantum operation on three qubits  $x_0, x_1, x_2$ . The four non-zero entries in the first column can be combined pairwise from  $(\frac{-1}{2}, \frac{1}{2})$  to  $(0, \frac{-1}{\sqrt{2}})$ . The resulting pair  $(\frac{-1}{\sqrt{2}}, \frac{-1}{\sqrt{2}})$  is combined to  $(-1, 0)$ . Finally, the  $-1$  is exchanged with the 0 entry in the first row and a phase shift by  $-1$  is applied. This leads to the matrix shown in Fig. 2b where the first column is of the desired form (note the extracted scalar factor of  $\frac{1}{2}$ ).

$$\begin{array}{c}
x_2 \\
x_1 \\
x_0
\end{array}
\begin{array}{c}
000 \\
001 \\
010 \\
011 \\
100 \\
101 \\
110 \\
111
\end{array}
\begin{array}{c}
\begin{array}{cccccccc}
1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 1
\end{array} \\
\begin{array}{cccccccc}
2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 1 & 0 & 0 & 1 & 0 & \omega^3 & -1 & 0 \\
0 & \frac{1}{\sqrt{2}} & \frac{-i}{\sqrt{2}} & \frac{-1}{\sqrt{2}} & \frac{1}{\sqrt{2}} & 0 & 0 & 0 & \frac{1}{\sqrt{2}} \\
0 & 1 & 0 & 0 & 1 & 0 & \omega^7 & 1 & 0 \\
0 & \frac{1}{\sqrt{2}} & \frac{i}{\sqrt{2}} & \frac{-1}{\sqrt{2}} & \frac{1}{\sqrt{2}} & 0 & 0 & 0 & \frac{1}{\sqrt{2}} \\
0 & 0 & 0 & 0 & 0 & -1 & \omega^7 & -1 & 1 \\
0 & -1 & 0 & 1 & 1 & 1 & 0 & 0 & 1 \\
0 & 0 & 0 & 0 & i & \omega & -i & -i
\end{array}
\end{array}$$

(a) Initial Matrix                      (b) After synthesizing first column

Fig. 2. Application of the local, i.e. column-wise, synthesis scheme.

However, this approach has two major drawbacks:

- 1) Two-level operations do not solely have an effect on that particular column, but on all columns of the matrix. Consequently, locally synthesizing one column without taking the global view, i.e. the remaining columns, into consideration may significantly worsen the degree of superposition in these columns (as already became evident in the previous example where the number of non-zero entries in columns 001 to 111 increased from Fig. 2a to Fig. 2b).
- 2) Two-level operations rely on multiple-controlled Clifford+T gates: Hadamard for combining, CNOT for exchanging, and  $T$  for modifying the phase shift between entries. These gates have many controls (in fact,  $n - 1$  controls where  $n$  is the number of qubits). Since the number of control lines significantly increases the costs of those gates (cf. Table I), this leads to high costs of the resulting circuits.

### B. General Idea of Proposed Approach

The key idea for improving logic synthesis of Clifford+T circuits is to particularly address these two drawbacks. Regarding the first drawback, it seems to be promising to globally consider the whole matrix at once and to only apply operations that do not lead to a worsening in any other column. Overall, this may lead to fewer steps to be conducted in order to derive the desired identity matrix.

Regarding the second drawback, it can be observed that Clifford+T matrices often exhibit similar structures occurring globally throughout the matrix. In many cases, the corresponding two-level operations applied to these structures are similar and can be combined to a joint, again global, operation that can be realized with lower costs. That is, controls (which enforce a local change in the matrix only) often can be saved as the correspondingly more global change can be conducted at once.

**Example 2.** The first two combine operations from the previous example correspond to multiple-controlled Hadamard gates with a target on qubit  $x_2$  and controls on  $x_0, x_1$ . The only difference between the gates is that the control connection on  $x_0$  is positive in one case and negative in the other case. Thus, this control can be dropped and both gates can be compressed to a single-controlled Hadamard gate. Hence, the

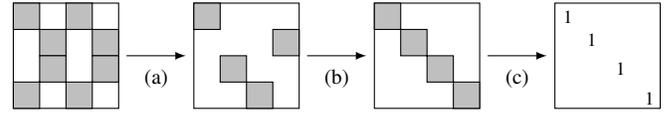


Fig. 3. Proposed global synthesis scheme.

two steps cannot only be conducted with one gate rather than two; but the required gate is even significantly cheaper. Overall, this reduces the cost from  $2 \cdot 7 = 14$  to 1 (cf. Table I).

Inspired by these observations, the general scheme of the proposed approach can be summarized as follows: Globally consider multiple columns simultaneously and establish recurrent structures that allow to combine as many two-level operations as possible. More precisely, the three steps (a)-(c) that have so far been conducted locally for the individual columns are now globally performed on the entire matrix as sketched in Fig. 3:

- (a) Regarding the elimination of superposition, we aim to gradually reduce superposition in the whole matrix, i.e. establish recurrent structures that allow to reduce the superposition in all columns at once. To this end, all entries of the matrix have to be rearranged and (possibly) phase shifted in such a way that all entries have a suitable partner to be combined with.

**Example 3.** Consider again the matrix from Fig. 2a. After exchanging rows 010 and 110 and applying a phase shift by  $i$  to row 111, all entries have a suitable partner to be combined with. Thus, an uncontrolled Hadamard gate can be applied on qubit  $x_1$ . As a consequence, the number of non-zero entries in the matrix can be reduced by a factor of 2 in one step. The resulting matrix is shown in Fig. 4a. Here, only rows 100 and 101 need to be swapped before another Hadamard gate on qubit  $x_0$  eliminates the remaining superposition and leads to the matrix shown in Fig. 4b.

- (b) Regarding the movement of entries to the diagonal, there is a large body of research on how to achieve this for permutation matrices (i.e. transformation matrices with Boolean entries only). In fact, this problem is known as *reversible circuit synthesis* and most of the approaches employed for this purpose (see e.g. [9]–[11]) are based on multiple-controlled Toffoli gates (which are exactly realizable in the Clifford+T library; see e.g. [8, Sec. 5]). Taking into account that potential phase shifts within in the matrix do not affect the applicability of these (highly optimized) approaches, any of these can be utilized here.

**Example 4.** Consider the resulting matrix from the previous step (shown in Fig. 4b). Applying a CNOT gate with control qubit  $x_1$  and target qubit  $x_0$  exchanges columns 010 and 011 with 110 and 111, respectively. This establishes zero sub-matrices in the upper-left and lower-right quadrant of the matrix and moves the  $-1$  in row 110 to the diagonal. Afterwards, a multiple-controlled NOT on  $x_2$  (with a negative control on  $x_0$  and positive control on  $x_1$ ) swaps columns 010

