Adaptive Clifford+T Decomposition of Large Toffoli Gates with One Clean Ancilla

Abhoy Kole
Cyber-Physical Systems
DFKI GmbH
Bremen, Germany
abhoy.kole@dfki.de

Majd Assaad
Cyber-Physical Systems
DFKI GmbH
Bremen, Germany
majd.assaad@dfki.de

Till Schnittka
Cyber-Physical Systems
DFKI GmbH
Bremen, Germany
till.schnittka@dfki.de

Rolf Drechsler
Institute of Computer Science
University of Bremen / DFKI GmbH
Bremen, Germany
drechsler@uni-bremen.de

***Abstract*—Multi-controlled Toffoli gates are fundamental building blocks in quantum computation, with applications in quantum arithmetic, simulation, and search algorithms. In fault-tolerant architectures, their realization is constrained by the high cost of non-Clifford resources, particularly in terms of T-count and T-depth. Recent advances have demonstrated that the use of ancillary qubits, relative-phase Toffoli gates, and dynamic circuit techniques can substantially reduce this overhead.**

In this work, we investigate the decomposition of large Toffoli gates using 3- and 4-input relative-phase Toffoli gates in the presence of a single clean ancilla and conditionally clean ancillas. We derive explicit resource bounds for Clifford+T implementations incorporating dynamic-circuit-based uncomputation and measurement-conditioned corrections. Our analysis emphasizes T-depth reduction under fixed CX and T-count overhead, ensuring relevance for near-term devices. We show that introducing 4-input relative-phase Toffoli gates enables significant T-depth reductions through enhanced parallelism while maintaining favorable ancilla requirements. We further validate our theoretical results through experimental evaluation and comparative analysis with existing approaches.

***Index Terms*—Toffoli gate, decomposition, fault-tolerant quantum computing, dynamic circuits**

I. Introduction

The Toffoli gate and its multi-controlled extensions are fundamental primitives in quantum computation, with applications in quantum arithmetic [1], quantum simulation [2], and Grover-oracle implementations [3]. However, efficient synthesis of large Toffoli gates remains challenging, as standard Clifford+T decompositions require high T-count, large circuit depth, and extensive CX-based entangling structures [4]. These overheads are particularly restrictive for Noisy Intermediate-Scale Quantum (NISQ) devices due to noise, decoherence, and limited qubit resources. In fault-tolerant architectures, Toffoli gates are typically realized via magic-state-based constructions, making both T-count and T-depth critical cost metrics [5].

Designing resource-efficient circuits for large Toffoli gates with minimal ancilla overhead therefore remains an active area of research. In many constructions, reductions in T-gate resources can be achieved by introducing additional ancilla qubits [6], [7]. In particular, the use of conditionally clean ancillas has been shown to be especially beneficial for realizing large Toffoli gates using fewer non-Clifford operations [8], [9]. Such techniques lead to improved resource bounds when one or two clean ancillas are available.

Recent advances in dynamic quantum circuits, which permit intermediate measurements, qubit resets, and classically controlled operations, have further expanded the design space for resource-efficient quantum computation [10], [11]. Unlike static circuits, dynamic circuits enable adaptive behavior in which measurement outcomes determine subsequent gate executions. These measurement-driven feed-forward mechanisms can replace costly multi-qubit unitary subcircuits [12]. Moreover, the availability of clean ancillas enables economical uncomputation through intermediate measurements, thereby reducing overhead [13], [14].

In decomposing large Toffoli gates, relative-phase Toffoli gates have been shown to reduce both T-count and ancilla requirements [4]. Motivated by these results, we study the use of 3- and 4-input relative-phase Toffoli gates with a single clean ancilla, combined with conditionally clean ancillas and dynamic uncomputation, to improve T-resource bounds. Although more T-depth-optimized constructions exist [15], [16], they typically incur higher CX overhead. We therefore adopt simpler realizations [4] with constant CX cost to enable fair comparison and practical relevance for NISQ devices. The main contributions of this work are summarized as follows:

1) We derive resource bounds for large Toffoli networks using 3- and 4-input relative-phase Toffoli gates with one clean and conditionally clean ancillas.
2) We show that 4-input relative-phase Toffoli gates reduce T-Depth by up to $2\lfloor (n-3)/6 \rfloor$ gate layers.
3) We incorporate dynamic-circuit-based uncomputation into the resource analysis.
4) We experimentally evaluate the proposed constructions.

The remainder of this paper is organized as follows. Section II reviews background concepts and related decomposition techniques. Section III presents the resource analysis for large Toffoli decompositions using one clean ancilla and 3-

This work was partly funded by the Federal Ministry of Research, Technology and Space (BMFTR) through the EASEPROFIT project (grant no. 16KIS2127) and the EQeS project (grant no. 13N17378), and by the German Research Foundation (DFG) through the CONAD-QC project (grant no. 559888852). This research was conducted within the scope of the DFG Priority Programme 2514: Quantum Software, Algorithms and Systems.

and 4-input relative-phase Toffoli gates. Section IV reports experimental results, and Section V concludes the paper.

II. Background

A. Quantum Computation and Clifford+T Gate Set

Quantum computation is based on the manipulation of quantum bits (qubits), which are represented by unit vectors in a two-dimensional complex Hilbert space $\mathcal{H}$. The computational basis states are

$$|0\rangle = \begin{pmatrix}1\\0\end{pmatrix}, \quad |1\rangle = \begin{pmatrix}0\\1\end{pmatrix}, \tag{1}$$

and an arbitrary single-qubit state is given by $|\psi\rangle = \alpha_0 |0\rangle + \alpha_1 |1\rangle$, where $\alpha_0, \alpha_1 \in \mathbb{C}$ and $|\alpha_0|^2 + |\alpha_1|^2 = 1$. An n-qubit system is described by a superposition of computational basis states and evolves under unitary operations (quantum gates) arranged in a quantum circuit.

A widely used universal gate set is the Clifford+T basis. The Clifford group is generated by the Hadamard (H), phase (S), and controlled-NOT (CX) gates, while universality is achieved by supplementing these with the non-Clifford T gate. In fault-tolerant architectures, Clifford gates are relatively inexpensive, whereas T gates require costly magic-state distillation and injection. Consequently, non-Clifford operations dominate the cost of fault-tolerant circuits. This cost is commonly quantified using the *T-count*, which measures the total number of T gates, and the *T-depth*, which denotes the number of sequential layers of T gates.

B. Dynamic Quantum Circuits

Dynamic quantum circuits allows intermediate measurements, qubit resets, and classically controlled operations during circuit execution. Unlike static circuits, in which all gates are fixed in advance, dynamic circuits enable adaptive behavior where subsequent operations depend on measurement outcomes. This capability is particularly valuable in fault-tolerant and resource-constrained settings, where measurement-based feedback can be used to reduce circuit depth and ancilla overhead.

In many quantum algorithms, ancillary qubits are introduced to facilitate intermediate computations and must be uncomputed before being reused. Traditional uncomputation relies on reversing previously applied unitary operations, which often incurs substantial additional cost. In contrast, measurement-based uncomputation [10], [11] exploits intermediate measurements on ancilla qubits, followed by classically conditioned corrections, to erase unwanted entanglement more efficiently. By replacing coherent uncomputation circuits with measurement-conditioned operations, this approach can significantly reduce the number of required non-Clifford gates.

C. Multi-Controlled Toffoli Gates

The Toffoli gate (CCX) is a three-qubit reversible gate that flips its target qubit t only when both control qubits: c_1 and c_2 are in the state $|1\rangle$. More generally, its n-control extension, denoted by $\mathrm{C}^n\mathrm{X}$, is defined as

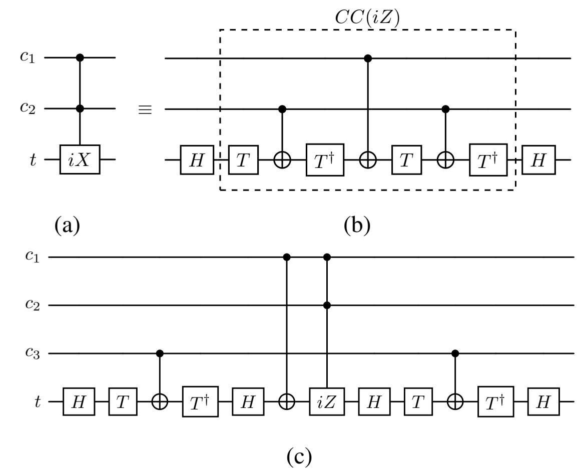

Fig. 1: Relative-phase Toffoli gate implementations: (a) the CC(iX) gate, (b) its Clifford+T decomposition, and (c) the Clifford+T realization of the C^3(iX) gate.

$$\mathrm{C}^n\mathrm{X}\,|c_1 \cdots c_n\rangle\,|t\rangle = |c_1 \cdots c_n\rangle\,|t \oplus (c_1 \wedge \cdots \wedge c_n)\rangle\,. \tag{2}$$

Such gates enable the implementation of complex Boolean functions and are fundamental components in many quantum algorithms. In the Clifford+T framework, Toffoli gates are non-Clifford operations and require decomposition into Clifford and T gates. Standard constructions often lead to high T-count and T-depth, making large Toffoli gates expensive in fault-tolerant settings.

As reported in [17], the Clifford+T realization of the three-qubit Toffoli (CCX) gate requires 7 CX gates and 7 T gates with a T-depth of 3. A T-depth of 1 can be achieved using the construction of [15], at the cost of four ancilla qubits. Various approaches have been proposed to optimize the realization of $\mathrm{C}^n\mathrm{X}$ gates, including ancilla-assisted constructions [6]–[9], relative-phase implementations [4], and measurement-based techniques [10], [11]. In this work, we build upon these ideas to reduce the resource cost of large Toffoli gates using relative-phase gates and dynamic circuit techniques, while adopting simpler realizations with constant CX overhead to ensure fair comparison and practical relevance for NISQ devices.

III. One-Clean Ancilla Decomposition

A. Clifford+T Optimization with CCiX and C3iX

Relative-phase Toffoli gates have proven advantageous for realizing higher-order Toffoli gates $\mathrm{C}^n\mathrm{X}$ ($n > 2$), as demonstrated in [4]. Figure 1 illustrates the Clifford+T implementations of the 3-input CC(iX) and 4-input C^3(iX) relative-phase Toffoli gates. The C^3(iX) construction exactly doubles the Clifford+T resources required for CC(iX), that is,

$$\mathcal{C}_{\text{ost}}\big(\mathrm{C}^3(\mathrm{iX})\big) = 2\ \mathcal{C}_{\text{ost}}\big(\mathrm{CC(iX)}\big). \tag{3}$$

Nevertheless, the use of C^3(iX) gates is beneficial for minimizing the number of ancilla qubits.

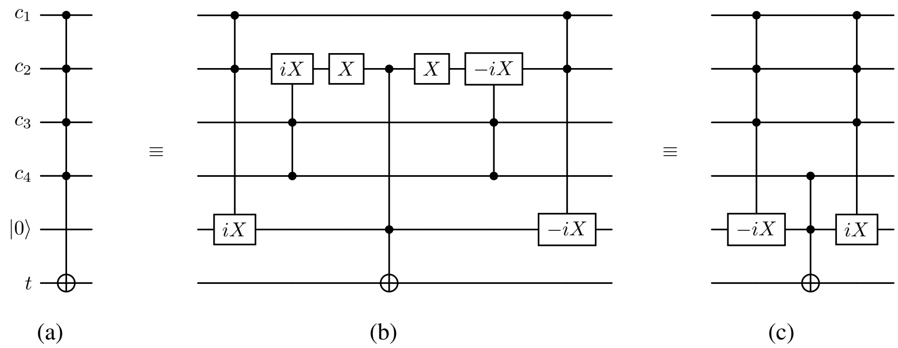

Fig. 2: (a) The realization of a 5-input Toffoli (C^4X) gate with one clean ancilla, implemented through two equivalent decompositions: (b) with four CC(iX) gates, and (c) with two C^3(iX) gates.

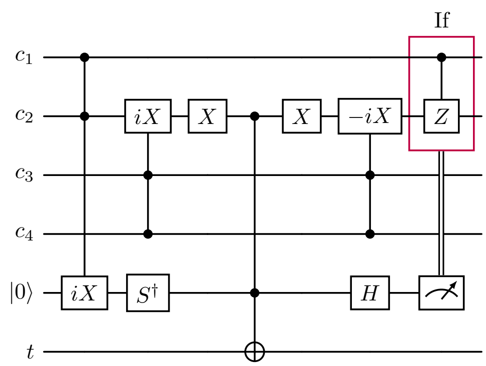

Fig. 3: A dynamic realization of the C^4X gate employing three CC(iX) gates and a conditional CZ operation triggered by the ancilla measurement.

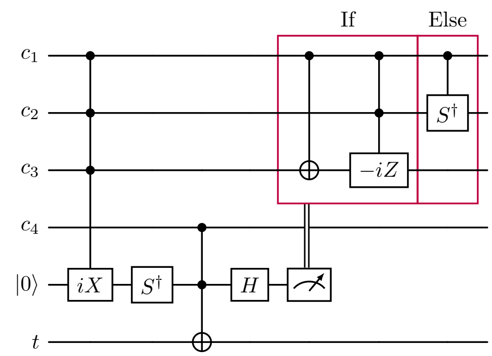

Fig. 4: A dynamic implementation of the C^4X gate using one C^3(iX) gate together with measurement-conditioned operations: CX·CC(iZ) when the ancilla outcome is 1, and $CS^\dagger$ when it is 0.

Using conditionally clean ancilla qubits [8], [9], a C^4X gate can be implemented as a network of CC(iX) gates with a single clean ancilla, as shown in Fig. 2b. Since CC(iX) requires 3 CXs, 4 T gates, and has T-depth 4 (see Fig. 1b), the resulting construction using four CC(iX) gates and one CCX gate has the total cost

$$\begin{aligned} &4 \times (3\text{ CX},\ 4\text{ T-count},\ 4\text{ T-depth}) \\ &\quad + (7\text{ CX},\ 7\text{ T-count},\ 3\text{ T-depth}). \end{aligned} \tag{4}$$

The two additional X gates enable the control qubit c_2 to act as a conditionally clean ancilla.

An alternative realization based on a single clean ancilla was proposed in [4], using C^3(iX) gates. As illustrated in Fig. 2c, this construction employs two C^3(iX) gates and yields the cost

$$\begin{aligned} &2 \times (6\text{ CX},\ 8\text{ T-count},\ 8\text{ T-depth}) \\ &\quad + (7\text{ CX},\ 7\text{ T-count},\ 3\text{ T-depth}). \end{aligned} \tag{5}$$

Both approaches in Eqs. (4) and (5) therefore incur identical overall costs. However, the C^3(iX)-based construction avoids the additional X gates, as shown in Fig. 2.

Further reductions can be achieved using dynamic circuits. According to [13], a fault-tolerant CCX gate can be implemented with a clean ancilla using only four T gates. Applying this approach to the C^4X decomposition in Fig. 3, and replacing the final CC(iX) gate by an ancilla phase correction ($S^\dagger$), a Hadamard-basis measurement, and measurement-conditioned phase restoration, yields

$$\begin{aligned} &3 \times (3\text{ CX},\ 4\text{ T-count},\ 4\text{ T-depth}) \\ &\quad + (7\text{ CX},\ 7\text{ T-count},\ 3\text{ T-depth}) + [1\text{ CX}]. \end{aligned} \tag{6}$$

Here, [1 CX] denotes an optional CX applied only when the Hadamard-basis measurement outcome of the clean ancilla is one, i.e. $M_H = 1$. Compared with Eqs. (4) and (5), this construction reduces the T-count and T-depth by four and the CX count by two.

A similar dynamic approach can be applied to the C^3(iX)-based network in Fig. 4, leading to

$$\begin{aligned} &(6\text{ CX},\ 8\text{ T-count},\ 8\text{ T-depth}) \\ &\quad + (7\text{ CX},\ 7\text{ T-count},\ 3\text{ T-depth}) \\ &\quad + \begin{cases} (4\text{ CX},\ 4\text{ T-count},\ 4\text{ T-depth}), & M_H = 1 \\ (2\text{ CX},\ 3\text{ T-count},\ 2\text{ T-depth}), & M_H = 0. \end{cases} \end{aligned} \tag{7}$$

It replaces the final C^3(iX) gate by an ancilla phase correction ($S^\dagger$), a Hadamard-basis measurement, and measurement-

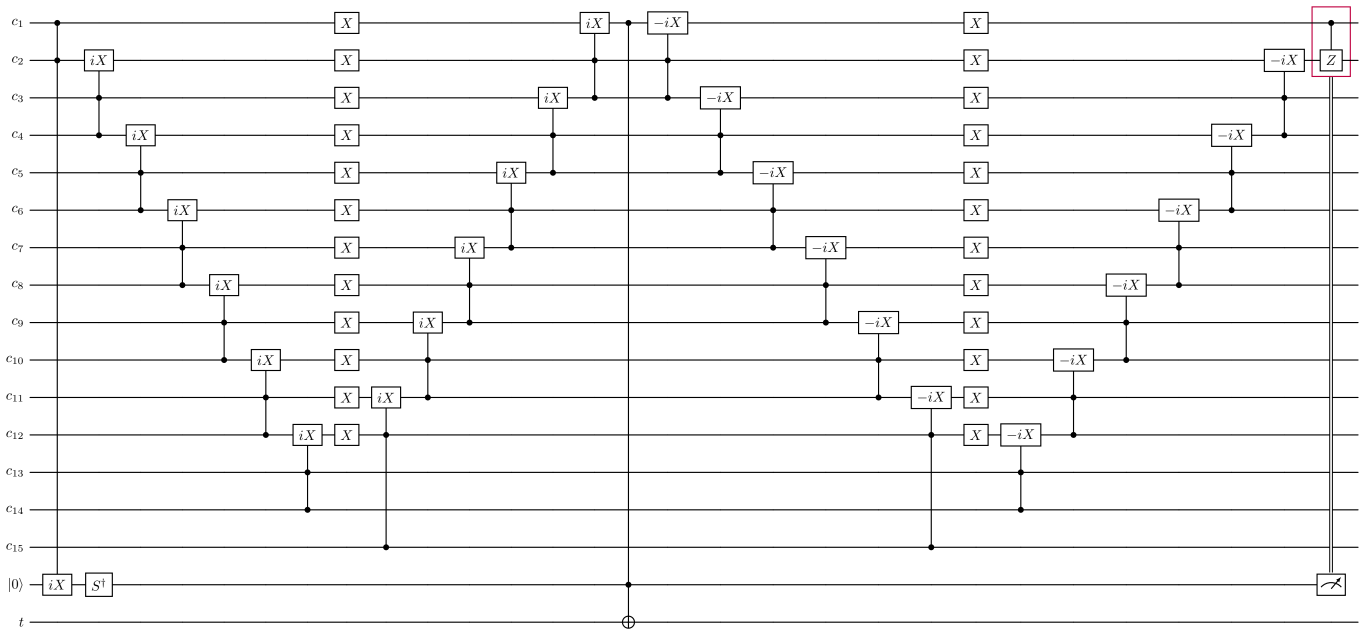

Fig. 5: A dynamic realization of the $C^{15}X$ gate with one clean ancilla, using a network of 25 CC(iX) gates, one CCX gate, and a conditional CZ operation triggered by ancilla measurement, derived from [8].

conditioned phase restoration of the form CX · CC(iZ) for $M_H = 1$ and $CS^\dagger$ for $M_H = 0$.

When $M_H = 1$, both dynamic constructions in Eqs. (6) and (7) have comparable costs (ignoring X gates). When $M_H = 0$, however, the construction in Fig. 4 is more efficient, requiring one fewer CX and T gate and reducing the T-depth by two.

B. Depth Minimization with C3iX over CCiX

According to [8], the use of conditionally clean ancilla enables the realization of a C^nX gate using $2n-3$ CCX gates when one clean ancilla is used. The use of CC(iX) gates in the network enables the realization to be obtained using $2n-4$ CC(iX) gates and one CCX gate. Replacing the final CC(iX) gate by an ancilla phase correction ($S^\dagger$), a Hadamard-basis measurement, and measurement-conditioned phase restoration further improves the resource bounds of the one-clean-ancilla-based realization:

$$\begin{aligned}&(2n-5) \times (3\text{ CX},\ 4\text{ T-count},\ 4\text{ T-depth})\\ &\quad + (7\text{ CX},\ 7\text{ T-count},\ 3\text{ T-depth})\ +\ [1\text{ CX}]. \qquad (8)\end{aligned}$$

As an illustration, Fig. 5 shows the realization of a $C^{15}X$ gate as a network of 25 CC(iX) gates.

The decomposition using the C^3(iX) gate does not reduce the overall gate count, as it requires twice the resources of the CC(iX) gate (see Eq. (3)). However, its use increases the scope for parallel gate scheduling and thereby reduces both the T-depth and the overall circuit depth in realizing larger Toffoli operations. This improvement arises from the fact that C^3(iX) provides three conditionally clean ancilla, compared to two in the case of CC(iX). The benefits of such parallelism

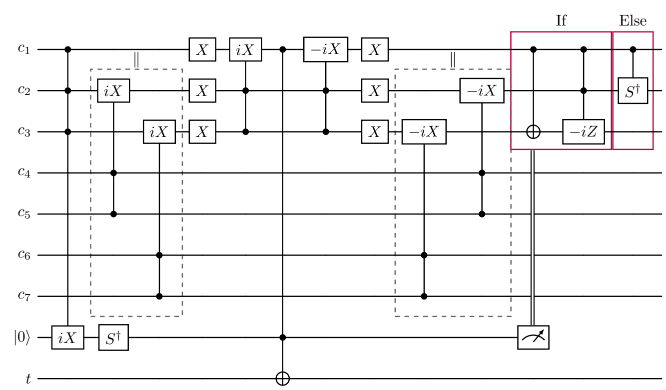

Fig. 6: A dynamic realization of the C^7X gate with reduced T-depth and one clean ancilla. The use of the C^3(iX) gate enables enables two successive parallel pairs of CC(iX) gates acting on the same set of qubits. The final C^3(iX) gate is replaced by measurement-conditioned operations: CX·CC(iZ) when the ancilla outcome is 1, and $CS^\dagger$ when it is 0.

become evident when decomposing a C^nX gate for $n \geqslant 7$, as shown in Fig. 6.

By shifting a CC(iX) gate to serve as a measurement-conditioned phase correction CC(iZ), the static[1] cost of the C^3(iX)-based dynamic decomposition is reduced by one CC(iX) gate compared to the CC(iX)-based dynamic decomposition. According to Eq. (8), the realization of a C^7X gate

[1]Here, *static* refers to circuit components that are not conditioned on classical measurement outcomes.

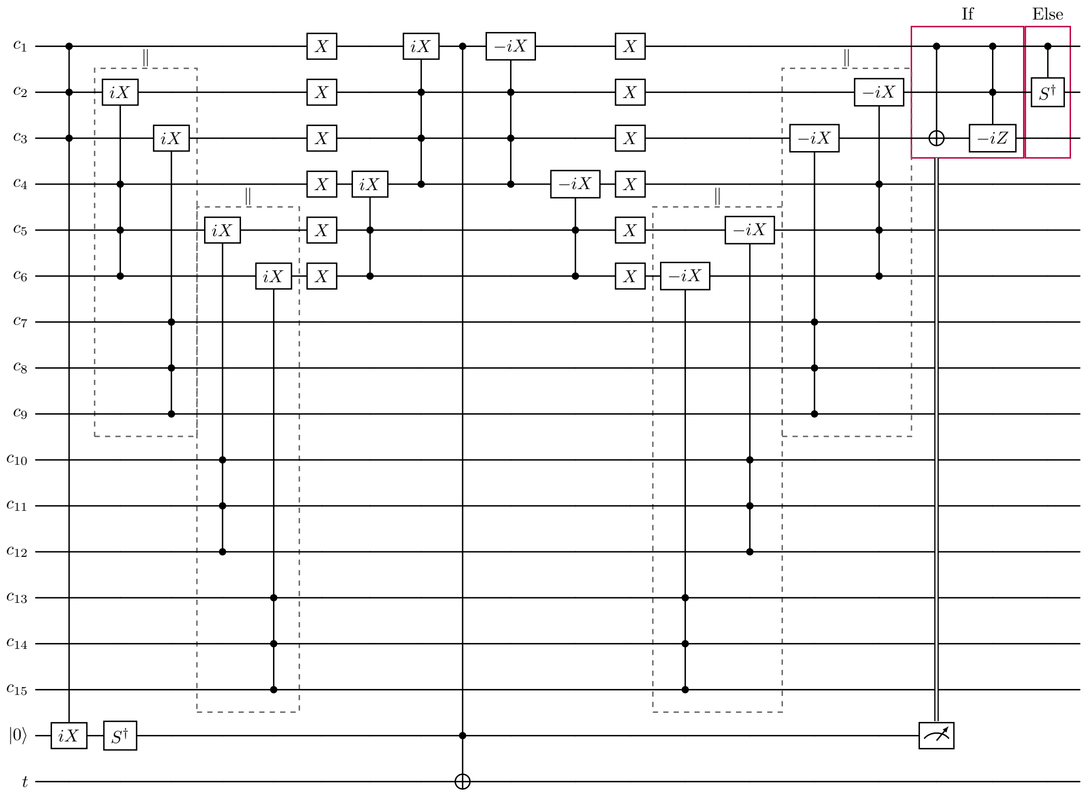

Fig. 7: A dynamic realization of the $C^{15}X$ gate with one clean ancilla. The use of the C^3(iX) gate enables four parallel pairs of C^3(iX) gates. The final C^3(iX) gate is replaced by measurement-conditioned operations: CX · CC(iZ) when the ancilla outcome is 1, and $CS^\dagger$ when it is 0.

requires 9 CC(iX) gates, whereas the use of C^3(iX) gates reduces this requirement to 6 CC(iX) gates (see Fig. 6). However, since the cost of a C^3(iX) gate is twice that of a CC(iX) gate (see Eq. (3)), the overall static cost of the C3(iX)-based dynamic decomposition corresponds to 8 CC(iX) gates. Nevertheless, due to the parallel execution of two successive pairs of CC(iX) gates (see Fig. 6), the T-depth is reduced by $2 \times$ T-depth(CC(iX)). For $n \geqslant 7$, the realization admits an additional T-depth reduction of

$$2\,\mathbf{1}_{\{\, n \bmod 6 \in \{1,2\}\,\}} \times \text{T-depth(CC(iX))}, \tag{9}$$

i.e., two successive pairs of CC(iX) gates can be scheduled in parallel when $n = 6k+1$, $k \geqslant 1$.

For $n \geqslant 9$, the T-depth reduction can be further increased, since the realization of a C^nX gate enables the parallel scheduling of at least two successive pairs of C^3(iX) gates. As C^3(iX) is twice as costly as CC(iX), parallel execution of these pairs yields an additional factor of two improvement in T-depth. Since each pair of C^3(iX) gates operates on six distinct qubits, the maximum T-depth reduction for $n \geqslant 9$ is given by

$$2\left\lfloor \frac{n-3}{6} \right\rfloor \times \text{T-depth}(\text{C}^3\text{(iX)}). \tag{10}$$

Consequently, the C^3(iX)-based dynamic decomposition improves the resource bounds of the C^nX realization as

$$\begin{aligned}
&(2n-6) \times (3\text{ CX},\ 4\text{ T-count},\ 4\text{ T-depth}) \\
&\quad + (7\text{ CX},\ 7\text{ T-count},\ 3\text{ T-depth}) \\
&\quad + \begin{cases} (4\text{ CX},\ 4\text{ T-count},\ 4\text{ T-depth}), & M_H = 1, \\ (2\text{ CX},\ 3\text{ T-count},\ 2\text{ T-depth}), & M_H = 0, \end{cases} \\
&\quad - 2\left\lfloor \frac{n-3}{6} \right\rfloor \times \text{T-depth}(\text{C}^3\text{(iX)}) \\
&\quad - 2\,\mathbf{1}_{\{\, n \bmod 6 \in \{1,2\}\,\}} \times \text{T-depth(CC(iX))}.
\end{aligned} \tag{11}$$

Thus, compared to Eq. (8), the introduction of C^3(iX) yields a more T-depth-efficient realization. As an illustration, Fig. 7 shows the realization of a $C^{15}X$ gate as a static network of 11 C^3(iX) gates and 2 CC(iX) gates, enabling four pairs of parallel C^3(iX) gate executions.

IV. Experimental Results

For experimental evaluation, we consider C^nX gates with $4 \leqslant n \leqslant 16$. As a baseline, we adopt the one-clean-ancilla realization proposed in [8], whose resource cost is given by

$$\begin{aligned}
&(2n-4) \times (3\text{ CX},\ 4\text{ T-count},\ 4\text{ T-depth}) \\
&\qquad + (7\text{ CX},\ 7\text{ T-count},\ 3\text{ T-depth}).
\end{aligned} \tag{12}$$

TABLE I: Results for the dynamic realization of the Toffoli operation using one clean ancilla, comparing the method of [8] with the proposed optimization.

	Static [8]			CC(iX)			Impr.(%)			CC(iX) + C^3(iX)			Impr.(%)		
n	CX	T_c	T_d	CX	T_c	T_d	CX	T_c	T_d	CX	T_c	T_d	CX	T_c	T_d
4	19	23	19	17	19	15	10.53	17.39	21.05	17	19	15	10.53	17.39	21.05
5	25	31	27	23	27	19	8.00	12.90	29.63	23	27	20	8.00	12.90	25.93
6	31	39	35	29	35	27	6.45	10.26	22.86	29	35	28	6.45	10.26	20.00
7	37	47	43	35	43	31	5.41	8.51	27.91	35	43	29	5.41	8.51	32.56
8	43	55	51	41	51	39	4.65	7.27	23.53	41	51	33	4.65	7.27	35.29
9	49	63	59	47	59	43	4.08	6.35	27.12	47	59	37	4.08	6.35	37.29
10	55	71	67	53	67	51	3.64	5.63	23.88	53	67	43	3.64	5.63	35.82
11	61	79	75	59	75	55	3.28	5.06	26.67	59	75	51	3.28	5.06	32.00
12	67	87	83	65	83	63	2.99	4.60	24.10	65	83	53	2.99	4.60	36.14
13	73	95	91	71	91	67	2.74	4.21	26.37	71	91	55	2.74	4.21	39.56
14	79	103	99	77	99	75	2.53	3.88	24.24	77	99	59	2.53	3.88	40.40
15	85	111	107	83	107	79	2.35	3.60	26.17	83	107	63	2.35	3.60	41.12
16	91	119	115	89	115	87	2.20	3.36	24.35	89	115	67	2.20	3.36	41.74

For fair comparison, we implement both the CC(iX)-based and the combined CC(iX) + C^3(iX)-based constructions and evaluate their resource costs in terms of CX-count (CX), T-count (T_c), and T-depth (T_d). Moreover, we consider the worst-case measurement-driven uncomputation cost, corresponding to the case $M_H = 1$, in all evaluations.

The resulting performance metrics are summarized in Table I. As observed from the table, while the CX-count and T-count remain identical, the combined CC(iX) + C^3(iX) construction achieves greater T-depth reductions than the CC(iX)-based approach for C^nX gates with $n \geqslant 7$, in accordance with the theoretical bounds derived in Eqs. (8) and (11).

V. Conclusion

In this work, we studied the efficient realization of large Toffoli gates using relative-phase constructions, conditionally clean ancillas, and dynamic circuit techniques. Using 3- and 4-input relative-phase Toffoli gates with a single clean ancilla, we derived resource bounds and analyzed their impact on T-count and T-depth. Our results show that higher-input relative-phase Toffoli gates enable significant T-depth reductions through enhanced parallelism, while dynamic uncomputation and measurement-conditioned corrections further reduce overhead. Experimental results confirm consistent improvements over existing methods, highlighting the effectiveness of combining relative-phase and dynamic-circuit techniques for scalable quantum computation.

Future work will investigate constructions that further optimize T-depth by allowing controlled increases in CX overhead, thereby exploring the tradeoff between two-qubit gate cost and circuit latency.

References

[1] K. Oonishi, T. Tanaka, S. Uno, T. Satoh, R. Van Meter, and N. Kunihiro, "Efficient construction of a control modular adder on a carry-lookahead adder using relative-phase toffoli gates," *IEEE Transactions on Quantum Engineering*, vol. 3, pp. 1–18, 2022.

[2] D. Rocca, C. L. Cortes, J. F. Gonthier, P. J. Ollitrault, R. M. Parrish, G.-L. Anselmetti, M. Degroote, N. Moll, R. Santagati, and M. Streif, "Reducing the runtime of fault-tolerant quantum simulations in chemistry through symmetry-compressed double factorization," *Journal of Chemical Theory and Computation*, vol. 20, no. 11, pp. 4639–4653, Jun 2024.

[3] A. Kole, M. E. Djeridane, L. Weingarten, K. Datta, and R. Drechsler, "qSAT: Design of an efficient quantum satisfiability solver for hardware equivalence checking," *J. Emerg. Technol. Comput. Syst.*, Apr. 2025.

[4] D. Maslov, "Advantages of using relative-phase Toffoli gates with an application to multiple control Toffoli optimization," *Phys. Rev. A*, vol. 93, p. 022311, Feb 2016.

[5] D. Litinski, "A Game of Surface Codes: Large-Scale Quantum Computing with Lattice Surgery," *Quantum*, vol. 3, p. 128, Mar. 2019.

[6] S. Dutta, S. Wang, A. Baksi, A. Chattopadhyay, and S. Maitra, "Exact space-depth trade-offs in multicontrolled toffoli decomposition," *Phys. Rev. A*, vol. 111, p. 052611, May 2025.

[7] S. Balauca and A. Arusoaie, "Efficient constructions for simulating multi controlled quantum gates," in *Computational Science – ICCS 2022*, D. Groen, C. de Mulatier, M. Paszynski, V. V. Krzhizhanovskaya, J. J. Dongarra, and P. M. A. Sloot, Eds. Cham: Springer International Publishing, 2022, pp. 179–194.

[8] T. Khattar and C. Gidney, "Rise of conditionally clean ancillae for efficient quantum circuit constructions," *Quantum*, vol. 9, p. 1752, May 2025.

[9] J. Nie, W. Zi, and X. Sun, "Quantum circuit for multi-qubit toffoli gate with optimal resource," *arXiv preprint arXiv:2402.05053 [quant-ph]*, 2024.

[10] L. Shirizly, L. C. G. Govia, and D. C. McKay, "Randomized benchmarking protocol for dynamic circuits," *Phys. Rev. A*, vol. 111, p. 012611, Jan 2025.

[11] A. Kole, K. Datta, and R. Drechsler, "Design automation challenges and benefits of dynamic quantum circuit in present nisq era and beyond: (invited paper)," in *2024 IEEE Computer Society Annual Symposium on VLSI (ISVLSI)*, 2024, pp. 601–606.

[12] A. Kole, A. Deb, K. Datta, and R. Drechsler, "Dynamic realization of multiple control Toffoli gate," in *Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2024, pp. 1–6.

[13] C. Jones, "Low-overhead constructions for the fault-tolerant Toffoli gate," *Phys. Rev. A*, vol. 87, p. 022328, Feb 2013.

[14] C. Gidney, "Halving the cost of quantum addition," *Quantum*, vol. 2, p. 74, Jun. 2018.

[15] P. Selinger, "Quantum circuits of T-depth one," *Phys. Rev. A*, vol. 87, p. 042302, Apr 2013.

[16] A. Paradis, J. Dekoninck, B. Bichsel, and M. Vechev, "Synthetiq: Fast and versatile quantum circuit synthesis," *Proc. ACM Program. Lang.*, vol. 8, no. OOPSLA1, Apr. 2024.

[17] M. Amy, D. Maslov, M. Mosca, and M. Roetteler, "A meet-in-the-middle algorithm for fast synthesis of depth-optimal quantum circuits," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 32, no. 6, pp. 818–830, 2013.