

A New Optimization Model for Multiple-Control Toffoli Quantum Circuit Design

Jihye Jung¹  

H. Milton Stewart School of Ind. and Syst. Engineering, Georgia Institute of Technology, Atlanta, GA, USA

Kevin Dalmeijer  

H. Milton Stewart School of Ind. and Syst. Engineering, Georgia Institute of Technology, Atlanta, GA, USA

Pascal Van Hentenryck  

H. Milton Stewart School of Ind. and Syst. Engineering, Georgia Institute of Technology, Atlanta, GA, USA

Abstract

As quantum technology is advancing, the efficient design of quantum circuits has become an important area of research. This paper provides an introduction to the MCT quantum circuit design problem for reversible Boolean functions without assuming a prior background in quantum computing. While this is a well-studied problem, optimization models that minimize the true objective have only been explored recently. This paper introduces a new optimization model and symmetry-breaking constraints that improve solving time by up to two orders of magnitude compared to earlier work when a Constraint Programming solver is used. Experiments with up to seven qubits and using up to 15 quantum gates result in several new best-known circuits, obtained by any method, for well-known benchmarks. Finally, an extensive comparison with other approaches shows that optimization models may require more time but can provide superior circuits with optimality guarantees.

2012 ACM Subject Classification Theory of computation → Constraint and logic programming

Keywords and phrases Constraint Programming, Quantum Circuit Design, Reversible Circuits

Digital Object Identifier 10.4230/LIPIcs.CP.2024.16

Related Version *Full Version*: <https://doi.org/10.48550/arXiv.2404.14384>

Funding This research was partly funded by the NSF AI Institute for Advances in Optimization (Award 2112533).

1 Introduction

As quantum technology is advancing, the efficient design of quantum circuits has become an important area of research. The primary challenge in quantum circuit design is to implement a *target function* using gates from a *preset gate library* to minimize the *circuit costs* according to some metric. This paper focuses on three choices that are often considered in the literature:

1. Target function: *Reversible Boolean function*, a key component that embeds the input data in most quantum algorithms.
2. Preset gate library: *Multiple-Control Toffoli (MCT) gate*, a typical high-level gate commonly used to represent reversible Boolean functions.
3. Circuit cost: *Quantum cost*, the number of low-level quantum gates required to realize the high-level gates in the circuit.

¹ Corresponding author



© Jihye Jung, Kevin Dalmeijer, and Pascal Van Hentenryck;
licensed under Creative Commons License CC-BY 4.0

30th International Conference on Principles and Practice of Constraint Programming (CP 2024).

Editor: Paul Shaw; Article No. 16; pp. 16:1–16:20



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

These concepts will be introduced in detail in Section 2, without assuming a prior background in quantum computing. The goal of this paper is to introduce a new optimization model to design quantum circuits within the setting defined above. For brevity, this problem will be referred to as the *MCT quantum circuit design problem*.

Literature Review. In similar settings to the one described above, early methods for quantum circuit design were developed based on intuitive observations and preconfigured circuit templates. Studies in this stage constructed a base library of small-scaled circuits to heuristically synthesize larger circuits for reversible Boolean functions [17, 9]. Post-synthesis algorithms, such as relocation algorithms [1, 22, 18, 19], were introduced to further improve these circuits, although the improved results cannot guarantee optimality either.

Several papers have used different representations of reversible Boolean functions to develop efficient synthesis algorithms. Cycle representation was used to devise several decomposition-based approaches [24, 31]. In particular, reference [24] reports an average cost improvement of 20% for benchmark functions with up to 20 qubits. Other approaches have leveraged the Reed-Muller expansion to decompose reversible boolean functions into exclusive-OR terms of Boolean products. The Reed-Muller expansion and the corresponding decision diagrams have appeared in reference [11] and [15] to address functions with up to 30 and 15 qubits, respectively. Reference [15] demonstrates a cost improvement of approximately 35% compared to previous studies, using a time limit of 600 seconds. A comparative analysis of decision diagram approaches for Reed-Muller expansion is also proposed [27].

Another heuristic approach was proposed via a quantum multiple-valued decision diagram, an efficient representation for matrices, to handle both reversible and irreversible functions [32]. The authors demonstrate the high scalability of the algorithm, handling functions with states up to length 156. A* algorithm was applied to the problem with an approximate heuristic function deduced from observations on state transitions [6], while others used heuristics based on isomorphic subgraph matching [14] and window optimization [26]. Evolutionary algorithms such as genetic algorithms [4], adaptive genetic algorithms [25], genetic programming [2], tabu search [8], and particle swarm optimization [7] were also suggested to obtain near-optimal solutions. While these methods offer different trade-offs between computation time and circuit quality, they are all heuristic and do not provide optimality guarantees.

Meanwhile, exact synthesis methodologies have been proposed to obtain optimal quantum circuits. Reference [10] iteratively solves satisfiability problems to obtain a quantum circuit with the minimum number of gates. This exact approach handles benchmark functions with three up to six qubits within a maximum of 5,000 seconds of computing time. The method based on quantified Boolean formula satisfiability, a generalized version of Boolean satisfiability, is also proposed to handle the same problem [29]. The authors report results for functions with four up to six qubits using a 2,000-second time limit. The exact methods notably handle relatively small functions but find better solutions within this space. Both references [10] and [29] report results in terms of quantum cost, which is the total number of low-level quantum gates required to implement a sequence of high-level logical gates. However, they do not directly minimize this objective; instead, they minimize the number of high-level gates as a proxy.

An optimization model was introduced to directly minimize the quantum cost in reference [13]. The authors use a multi-commodity flow-based model that is solved with a Mixed Integer Programming (MIP) solver. The method is applied to functions with three to six qubits, and significant improvements in quantum cost between 18.8% and 68.6% are observed. This provides a strong motivation to optimize quantum costs directly. While good results are obtained for small functions, the method does not scale well beyond seven gates due to the exponential number of binary variables in the model.

Contributions. This paper introduces a new optimization model to minimize quantum cost directly. Compared to [13], the new model is easier to implement, requires exponentially fewer binary variables, and has a beneficial block-angular structure. Furthermore, the paper demonstrates the advantage of Constraint Programming (CP) in solving this model. The key contributions can be summarized as follows:

- The paper introduces a new optimization model and new symmetry-breaking constraints for MCT quantum circuit design.
- The new model allows both CP and MIP solvers to significantly improve solving time, with up to *two orders of magnitude* speedup when the CP solver is used.
- Experiments with up to seven qubits and using up to 15 quantum gates result in several new best-known circuits for well-known benchmarks.
- An extensive comparison with other approaches shows that optimization models may require more time but can provide superior circuits with guaranteed optimality.

The remainder of the paper is organized as follows. Section 2 presents the necessary terminology and provides the problem description. Section 3 introduces the new optimization model, while Section 4 introduces new symmetry-breaking constraints. The computational results are presented by Section 5, and Section 6 concludes the paper.

2 Terminology

As discussed in the introduction, this paper designs quantum circuits for *reversible Boolean functions* using *MCT gates* to minimize the *quantum costs* of the resulting circuit. The relevant definitions are introduced here without assuming a prior knowledge of quantum computing. Example 1 presents a running example that is used throughout this section.

Basics of Quantum Computing. A state of a quantum system is represented by *qubits*, analogous to classical bits in classical computers. While bits assume values of 0 or 1 to define a single *basis state* (i.e., a binary vector), qubits may represent a *superposed state* (i.e., a complex vector) formed as a convex combination of the basis states. A *quantum gate* operates on qubits to transition the system to a new state based on the specification. Not every state transition can be realized by a single elementary gate, and multiple quantum gates may be combined into a *quantum circuit* to represent more complicated functions.


Reversible Boolean Function. A reversible Boolean function is a bijective function where inputs and outputs are provided as binary strings of fixed length. This function corresponds to a unique permutation and is often presented in the form of a truth table. It is noteworthy that reversible Boolean functions have been recognized as fundamental operators in quantum computing, thus explored extensively in prior research on quantum circuit synthesis [23].

Example 1a provides an example of a three-qubit reversible Boolean function. The specification defines a one-to-one mapping for each of the $2^3 = 8$ basis states. For instance, the input state (qubit 1, qubit 2, qubit 3) = (1, 1, 0) is mapped to the output state (qubit 1, qubit 2, qubit 3) = (0, 1, 1), or $110 \rightarrow 011$ for short. It is sufficient to only specify the function for the basis states: when superposed states are involved, they can simply be decomposed into a convex combination of basis states, after which the function can be applied to each basis state according to the specification.

16:4 A New Optimization Model for MCT Quantum Circuit Design

Input	Output	Input	Output
000	001	100	101
001	000	101	100
010	110	110	011
011	111	111	010

(a) Truth Table (completely specified).



(b) Implementing Circuit.

■ **Example 1** Specification and Implementing Circuit (interactive: algassert.com/quirk).

Multiple Control Toffoli (MCT) Circuit. *MCT circuits* consist of a sequence of *MCT gates*. Example 1b provides a circuit that meets the specification of Example 1a. It has three horizontal lines (one for each qubit q) and three MCT gates (one per column d). An MCT gate consists of one *target qubit* with the \oplus symbol and zero or more *control qubits* with the \bullet symbol. Control qubits do not have to be adjacent, and vertical lines connect the control qubits to the target qubit. For a given input, the circuit is read from left to right, and the MCT gates are applied one at a time. Transitions follow the following rule: *if all the control qubits are in state 1, then the target qubit is flipped*.

For example, consider the input 110 and start from the very left of the given circuit Example 1b. The top line has value 1, the middle line has value 1 and the bottom line has value 0. The first gate has one control qubit on line two. It follows that all control qubits are in state 1. As a result, the target qubit (qubit 1) is flipped, changing the state to 010 after the first gate. The second gate has two control qubits, but they are not all in state 1 (qubit 1 is in state 0) so nothing happens. The third gate does not have any control bits, so the target qubit is flipped. This results in the output 011. That is, the total transition is $110 \rightarrow 010 \rightarrow 010 \rightarrow 011$, meeting the specification. It can be checked that the circuit meets the specifications for the other input states as well.

An important property of MCT circuits is that they are *reversible*, i.e., they perform the inverse operation when read from right to left [24]. Therefore, MCT circuits are a natural candidate to represent reversible Boolean functions. In fact, it is well-known that every reversible Boolean function can be represented in this way.

Quantum Costs. To implement an MCT circuit in practice, each MCT gate is decomposed into elementary quantum gates. The number of elementary quantum gates is a well-established proxy for the cost of the MCT circuit, known as the *quantum cost*. Table 1 summarizes the best-known quantum cost $f(c)$ for an MCT gate that uses a total of $c \geq 0$ control qubits [5, 16, 10]. Note that the costs change based on the number of *slack qubits* that are available and that are not used in the MCT gate otherwise. The cost of the circuit in Example 1 is $f(1) + f(2) + f(0) = 7$. Note that the costs in the table may go down in the future as better decompositions are found. It can also be seen that the quantum cost of an MCT gate tends to increase rapidly as more control qubits are added.


Remarks on Incomplete Specification. In Example 1, the truth table was completely specified, but this does not always have to be the case: depending on the application, there may be specific qubits that are used in the computation but for which the output is uninteresting (*don't care* qubits). However, every circuit *implementation* still represents a bijective function that assigns specific states to the don't cares, due to the reversible nature of quantum operators. Note that don't cares apply only to the outputs, whereas the inputs are completely specified in practice.

■ **Table 1** Quantum Costs for MCT Gates (dots indicate the same cost as above; slack qubits are qubits that are available but not used in the MCT gate).

Slack qubits	Control qubits p							
	0	1	2	3	4	5	6	≤ 7
0	1	1	5	13	29	62	125	$2^{p+1} - 3$
1	52	80	.
2	26	.	.	.
3	38	.	.
≤ 4	50	.

Input	Output	Impl. 2b	Input	Output	Impl. 2b
000	00-	001	100	101	101
001	00-	000	101	100	100
010	11-	111	110	011	011
011	---	110	111	010	010

(a) Incompletely Specified Truth Table.



(b) Implementing Circuit.

■ **Example 2** Specification and Implementing Circuit (interactive: algassert.com/quirk).

Example 2 turns the complete specification of Example 1 into an incomplete specification by replacing some of the output qubits by “-” (don’t care) instead of 0 or 1. Note that the circuit in Example 1b is still valid, but with the additional freedom it might be possible to find a circuit with a cost lower than 7. Example 2b shows that a better circuit can indeed be found. The outputs of this implementation are added to Example 2a under “Impl. 2b”. These outputs are different from the circuit in Example 1, but they both meet the incomplete specification. At a cost of only $f(1) + f(0) = 2$, the new circuit is a better implementation.

3 Optimization Model

This section presents a new optimization model to design minimum-cost MCT quantum circuits that meet a given specification. The new model is provided as Model (1) in Figure 3, and the different components will be introduced over the next several paragraphs. Constraints (1h)-(1i) are introduced as logical constraints first to clearly state their intent, after which a linear implementation is provided as just one possible implementation. An overview of the nomenclature is provided in Appendix A for convenience. While [13] and this paper both use network flows to model quantum state transitions, there are significant differences between the models that will be discussed at the end of this section.

Circuit Design. The design of a quantum circuit is modeled through variables that describe a quantum circuit diagram such as Example 1b or Example 2b. The set of qubits $Q = \{1, \dots, n\}$ describes the rows, while the set of gates $D = \{1, \dots, m\}$ describes the columns. Note that the maximum number of gates m is an input to the problem. Variables t_q^d and w_q^d are binary variables in Model (1) that indicate whether (q, d) , contains a target or control qubit, respectively, for any $q \in Q$, $d \in D$. Constraints (1b) state that each spot contains a target qubit, a control qubit, or neither; but not both. Constraints (1c) enforce at most one target qubit per gate. Gates without a target qubit are forced to be empty through Constraints (1d). The variables are defined in Equation (1j).

$$\min \sum_{d \in D} \sum_{j \in Q} f(j-1)y_j^d, \quad (1a)$$

$$\text{s.t.} \quad t_q^d + w_q^d \leq 1 \quad \forall q \in Q, d \in D, \quad (1b)$$

$$\sum_{q \in Q} t_q^d \leq 1 \quad \forall d \in D, \quad (1c)$$

$$w_q^d \leq \sum_{r \in Q} t_r^d \quad \forall q \in Q, d \in D, \quad (1d)$$

$$\sum_{j \in Q} j y_j^d = \sum_{q \in Q} t_q^d + \sum_{q \in Q} w_q^d \quad \forall d \in D, \quad (1e)$$

$$\sum_{j \in Q} y_j^d \leq 1 \quad \forall d \in D, \quad (1f)$$

$$\sum_{a \in \delta_k^+(v)} x_a^k - \sum_{a \in \delta_k^-(v)} x_a^k = \begin{cases} |\Omega_k^{in}| & \text{if } v = S \\ -|\Omega_k^{in}| & \text{if } v = T \\ 0 & \text{else} \end{cases} \quad \forall k \in K, v \in V, \quad (1g)$$

$$\bigvee_{q^0 \in Q_{\sigma(a)}^0} (w_{q^0}^{d(a)} = 1) \bigvee (t_{q(a)}^{d(a)} = 0) \implies x_a^k = 0 \quad \forall k \in K, a \in A_k^{flip}, \quad (1h)$$

$$\bigwedge_{q^0 \in Q_{\sigma(a)}^0} (w_{q^0}^{d(a)} = 0) \bigwedge \left(\sum_{q \in Q} t_q^{d(a)} = 1 \right) \implies x_a^k = 0 \quad \forall k \in K, a \in A_k^{keep}, \quad (1i)$$

$$t_q^d, w_q^d, y_j^d \in \{0, 1\} \quad \forall q, j \in Q, d \in D, \quad (1j)$$


$$x_a^k \in \{0, 1\} \quad \forall k \in K, a \in A_k. \quad (1k)$$

■ **Figure 3** Optimization Model.

Quantum Cost. An MCT gate with $c \geq 0$ control qubits incurs a quantum cost of $f(c)$, as defined in Table 1. Let y_j^d be a binary indicator that takes value one if gate $d \in D$ contains exactly $j \in Q$ target and control qubits, or zero otherwise. Objective (1a) then calculates the total quantum cost over all gates. Note that the input $q-1$ subtracts the single target qubit, as $f(\cdot)$ is defined in terms of control qubits only. When gate $d \in D$ is empty, all indicators y_j^d are zero (note $j \geq 1$) and there is no contribution to the objective. The y -variables are forced to take the correct values by Constraints (1e) and (1f). The former ensures that the indicators together represent the total number of target and control qubits, and the latter ensures that at most one indicator is active.

Intuition Flow Networks. The new model uses network flows to model the state transitions that are caused by the circuit. Before formally introducing this part of the model, this paragraph aims to give some intuition through Figure 4. The graph in this figure has vertices (σ, d) that indicate being in state $\sigma \in \Omega$ before gate $d \in D$ is applied. The arrows show the transitions when state 010 is provided as an input to Circuit 2b. That is, gate $d = 1$ carries out the transition $010 \rightarrow 110$ (vertex $(010, 1)$ to $(110, 2)$), and gate $d = 2$ carries out the transition $110 \rightarrow 111$ (vertex $(110, 2)$ to $(111, 3)$), for a total transition of $010 \rightarrow 111$.

Which state transitions are available depends on the design of the circuit. The dotted lines in Figure 4 show the state transitions that are allowed when the circuit design variables represent Circuit 2b. Because of the reversible nature of MCT gates, the dotted lines create



■ **Figure 4** Flow Network Example for $k = 2$ corresponding to Example 2.

bijections between the states. As such, flow entering vertex $(\sigma, 1)$ is pushed to a unique vertex $(\sigma', m + 1)$, which represents the transition $\sigma \rightarrow \sigma'$ that is caused by the circuit. To make sure that the state transitions match the specification, a source S and a sink T are added. The source pushes the flow to a particular input state, while the sink only accepts flow from the correct output states. In the example, input state 010 has output specification 11– (Table 2a). As such, the source pushes flow to 010, and the sink only accepts flow from 110 and 111. If no flow from S to T is possible, then the circuit fails to meet the specification, and the current assignment of the circuit design variables is infeasible.

A separate flow network is introduced for each input state to ensure that the circuit meets all the specifications. The model is further improved by grouping input states with the same output specification into *commodities*. For a completely specified function, each commodity is associated with a single input state. In contrast, for an incomplete function, multiple input states can be grouped together based on the output specification. The source pushes a unit flow to each of the input states in the group. The flows remain separated due to the bijections, and are eventually collected by the sink. Note that this grouping is only possible when input states have the same output specification, and therefore the arcs connecting to the sink are the same.

Quantum States and Flow Commodities. Given the intuition, this paragraph formalizes the quantum states and the flow commodities. The set Ω of basis states is given by the 2^n binary vectors of length n . For example, $\Omega = \{000, 001, 010, 011, 100, 101, 110, 111\}$ in Example 2. States with the same (possibly incomplete) output specification will be modeled together and grouped into commodities $k \in K$. Each commodity is represented by the set $\Omega_k^{in} \subseteq \Omega$ of corresponding input states. Again, using Example 2: $\Omega_1^{in} = \{000, 001\}$ (output specification 00–), $\Omega_2^{in} = \{010\}$ (output specification 11–), etc. In addition, the set Ω_k^{out} defines the permitted output states for each commodity $k \in K$. In the example this yields $\Omega_1^{out} = \{000, 001\}$ (both match 00–), $\Omega_2^{out} = \{110, 111\}$ (both match 11–), etc. Note that the sets Ω_k^{in} partition Ω by definition, while the sets Ω_k^{out} may overlap and together cover Ω .

Flow Networks. Next, flow networks are defined that model the state transitions throughout the circuit. Figure 4 continues to provide a running example for this paragraph. A graph $G_k = (V, A_k)$ is defined for each commodity $k \in K$. The vertex set V consists of a source S , a sink T , and the vertices (σ, d) for $\sigma \in \Omega$, $d \in D \cup \{m + 1\}$. Each vertex (σ, d) represents being in state σ right after gate $d - 1$ or right before gate d , i.e., gate d is to be applied next. The arc set A_k of commodity $k \in K$ consists of four components:

- *Source arcs* that connect the source to the commodity input states: $(S, (\sigma, 1)) \forall \sigma \in \Omega_k^{in}$.
- *Sink arcs* that connect permitted output states to the sink: $((\sigma, m+1), T), \forall \sigma \in \Omega_k^{out}$.
- *Flip arcs* A_k^{flip} that represent the cases when gate $d \in D$ acts on state $\sigma \in \Omega$ by flipping bit $q \in Q$. Let $\sigma \oplus q$ denote state σ with bit q flipped. Then A_k^{flip} consists of arcs $((\sigma, d), (\sigma \oplus q, d+1)), \forall d \in D, \sigma \in \Omega, q \in Q$.
- *Keep arcs* A_k^{keep} that represent the cases when gate $d \in D$ keeps state $\sigma \in \Omega$ the same. That is, the keep arcs are given by $((\sigma, d), (\sigma, d+1)), \forall d \in D, \sigma \in \Omega$.

The state transitions for commodity $k \in K$ are represented by a network flow in G_k . Equation (1k) defines flow variables x_a^k that take on value one if arc $a \in A_k$ is used, and zero otherwise. The required flow is a flow from the source to the sink of size $|\Omega_k^{in}|$. This flow is distributed to all the input states Ω_k^{in} by the source arcs, after which the flip and keep arcs model the state transitions. The only way to reach the sink is through the sink arcs that start from one of the permitted output states Ω_k^{out} . The flow balance constraints (1g) enforce that the x_a^k variables represent such a flow. Here $\delta_k^+(v)$ and $\delta_k^-(v)$ denote the out-arcs and in-arcs of vertex $v \in V$, respectively.

It remains to connect the circuit design decisions to the network flows, i.e., ensure that arcs can only be used if they match the circuit specification. For convenience, the following shorthands are used for properties of flip and keep arcs $a \in A_k^{flip} \cup A_k^{keep}$: $d(a) \in D$ is the gate associated with arc a , $q(a) \in Q$ is the qubit that is flipped by arc a (flip arcs only), and $\sigma(a) \in \Omega$ is the state that the arc a transitions from. Furthermore, for a given state $\sigma \in \Omega$ it will be convenient to define Q_σ^0 as the set of qubits that are in state 0, e.g., $Q_{010}^0 = \{1, 3\}$. Constraints (1h) and (1i) eliminate the flow from all flip and keep arcs that do not match the circuit specification. This can be seen by considering the outgoing arcs of an arbitrary vertex (σ, d) , $\sigma \in \Omega$, $d \in D$. That is, the set of one keep arc and n flip arcs that model the state transition due to gate d .

- *Case 1: Gate d flips some qubit $\bar{q} \in Q$.* Based on the transition rule (Section 2) this means that \bar{q} is the target qubit ($t_{\bar{q}}^d = 1$) and all controls are on qubits with value 1 in state σ . Or alternatively, none of the controls are on qubits with value 0 in state σ ($w_{q^0}^d = 0 \forall q^0 \in Q_\sigma^0$). It follows that the antecedent of (1i) holds and that the keep arc is excluded as expected. Flip arcs a are eliminated by (1h) as soon as they flip the wrong qubit, i.e., $q(a) \neq \bar{q}$, which implies $t_{q(a)}^d = 0$. The arc that flips \bar{q} is the only flip arc that is *not* excluded by (1h), as the target is in the right place ($t_{q(a)}^d = 1$) and all controls are on the zero states ($w_{q^0}^d = 0$).
- *Case 2: Gate d keeps state σ the same.* Either there is no target qubit, in which case all flip arcs have $t_{q(a)}^d = 0$ and get eliminated by (1h) while the keep arc is unaffected by (1i). Or there is a target qubit \bar{q} , but at least one of the controls is on a zero state, i.e., $w_{q^0}^d = 1$ for some $q^0 \in Q_\sigma^0$. It follows again that all flip arcs are eliminated while the keep arc is unaffected.

It is concluded that Constraints (1h)-(1i) close the correct arcs to match the design of the circuit, which completes the model. It should also be noted that the flow variables may be relaxed to the continuous domain $x_a^k \in [0, 1] \forall k \in K, a \in A_k$. This follows from the fact that for fixed values of the t , w , and y -variables the remaining problem decomposes into $|K|$ independent minimum-cost flow problems, which are known to have the integrality property [3]. This means that the new model requires only $\mathcal{O}(nm)$ binaries.

Implementation of (1h)-(1i). There are multiple ways to implement Constraints (1h)-(1i), depending on the solver. This paper reformulates the implications as linear constraints, which are widely supported. The conversion is straightforward (e.g., see [30]), and results in the following inequalities that replace (1h)-(1i):

$$x_a^k \leq t_{q(a)}^{d(a)} \quad \forall k \in K, a \in A_k^{flip}, \quad (2a)$$

$$x_a^k \leq 1 - w_{q^0}^{d(a)} \quad \forall k \in K, a \in A_k^{flip}, q^0 \in Q_{\sigma(a)}^0, \quad (2b)$$


$$x_a^k \leq 1 - \sum_{q \in Q} t_q^{d(a)} + \sum_{q^0 \in Q_{\sigma(a)}^0} w_{q^0}^{d(a)} \quad \forall k \in K, a \in A_k^{keep}. \quad (2c)$$

Comparison to [13]. The new optimization model differs from [13] in a number of significant ways. First, [13] separately treat four cases for each state transition: empty gate (no flip), zero control qubits (flip), one or more control qubits with flip, one or more control qubits without flip. The new model captures all these cases in the same framework, significantly simplifying the formulation. Another important difference is in how the circuit is connected to opening and closing the flow arcs. [13] define a binary variable for each state $\sigma \in \Omega$ and gate $d \in D$ that identifies whether this state is modified by the gate, which introduces $\mathcal{O}(2^n m)$ binary variables. The new model provides a much more direct way to close arcs through Constraints (1h) and (1i). Compared to [13], this multiplies the number of constraints by a factor of $\mathcal{O}(n^2)$, *but no additional binary variables are necessary*. As a result, the new model requires only $\mathcal{O}(nm)$ binary variables. The number of flow variables remains at $\mathcal{O}(2^n nm|K|)$. However, as mentioned before, the new model decomposes into smaller independent minimum-cost flow problems for a fixed design. This implies that the formulation has a block-angular structure that may be exploited by decomposition methods in future work.

4 Symmetry-Breaking Constraints

It has been observed previously that the optimal quantum circuit design is not necessarily unique. The paper by [13] observes that empty gates do not affect the overall circuit, and constraints are added to force empty gates to appear at the end. The work by [12] presents multiple transformations that lead to new circuits with equivalent outputs (but not necessarily the same quantum cost). Inspired by these observations, this section defines three different swap operations – *Swap 1*, *Swap 2*, and *Swap 3* – that result in a different but equivalent circuit with the same cost. It will be proven that repeatedly applying these operations eventually results in a circuit that is *unswappable*, i.e., no further swaps of these types can be applied. This makes it possible to introduce symmetry-breaking constraints that limit the search to unswappable circuits, without loss of optimality. After all, every swappable quantum circuit is associated with an unswappable quantum circuit that has the same cost. The swaps used in this paper are introduced below, and future works may expand the list to eliminate additional symmetries.

- *Swap 1: Empty Gate.* If gate $d \in D$ is empty ($\sum_{q \in Q} t_q^d = 0$) and gate $d + 1$ is not empty ($\sum_{q \in Q} t_q^{d+1} = 1$), then swap the two gates.
- *Swap 2: Different Target.* If the target qubit $q \in Q$ of gate $d \in D$ is at a higher line than the target qubit $r \in Q$ of gate $d + 1$ ($q > r$, lower in the diagram), and the target qubits do not neighbor a control qubit ($w_q^{d+1} = 0$ and $w_r^d = 0$), then swap the two gates. It was observed by [12] that the target qubits do not affect the control qubits in either direction, and hence the gates can safely be swapped. Example 5a-5b provides a visualization.



■ **Example 5** Swap Operations.

■ *Swap 3: Same Target.* If gate $d \in D$ and gate $d + 1$ have the same target qubit $q \in Q$ ($t_q^d = 1$ and $t_q^{d+1} = 1$), and gate d has fewer control bits ($\sum_{r \in Q} w_r^d < \sum_{r \in Q} w_r^{d+1}$), then swap the two gates. Again, the target qubits do not affect the neighboring control qubits, which justifies the swap. A visualization is provided by Example 5c-5d.

► **Proposition 1.** *Any swappable circuit can be turned into an unswappable circuit by repeatedly applying Swap 1-3.*

Proof. See Appendix B. ◀

Constraints. Proposition 1 justifies symmetry-breaking constraints that enforce that the circuit is unswappable. The three swaps are translated into the following three classes of inequalities:

$$\sum_{q \in Q} t_q^d \geq \sum_{q \in Q} t_q^{d+1} \quad \forall d \in D, \quad (3a)$$

$$t_q^d + t_r^{d+1} \leq 1 + w_q^{d+1} + w_r^d \quad \forall d \in D, q, r \in Q, q > r, \quad (3b)$$

$$\sum_{r \in Q} w_r^d - \sum_{r \in Q} w_r^{d+1} \geq (n - 1)(t_q^d + t_q^{d+1} - 2) \quad \forall d \in D, q \in Q. \quad (3c)$$

Constraints (3a) prevent Swap 1 by forcing gate $d \in D$ to be non-empty when gate $d + 1$ is non-empty. A similar constraint is used in [13]. Constraints (3b) model that if the target bits are set up correctly for Swap 2 (left-hand side of the equation equals two), then $w_q^{d+1} = 1$ or $w_r^d = 1$. This is necessary because $w_q^{d+1} = w_r^d = 0$ would allow Swap 2 to be applied. The constraint is automatically satisfied when the left-hand side is less than two. Finally, consider Constraints (3c). If the gates have targets on the same qubit ($t_q^d = t_q^{d+1} = 1$) then the constraint reduces to $\sum_{r \in Q} w_r^d \geq \sum_{r \in Q} w_r^{d+1}$ to prevent Swap 3. The term $n - 1$ (the maximum number of control bits per gate) is sufficiently large to make the constraint inactive when the targets are not on the same qubit.

5 Computational Experiments

Computational experiments on well-known benchmarks are presented to demonstrate the performance of the new optimization model. The new model and the symmetry-breaking constraints are implemented in Python 3.11 and run on a Linux machine with dual Intel Xeon Gold 6226 CPUs (24 cores in total) on the PACE Phoenix cluster [20]. CP-SAT 9.8.3296 [21] is used as the CP solver with 24 workers (threads), and Gurobi 11.0.0 is used for the MIP approach. The instances are sourced from RevLib [28], a common benchmark for reversible and quantum circuit design. This paper selects boolean functions with up to seven qubits that have known circuit implementations in fewer than 100 gates. After removing easy three-qubit functions, a benchmark suite of 49 functions remains. A time limit of 3600

■ **Table 2** Performance New Optimization Model compared to [13].

	Average Runtime (s)				Solved Instances		
	$m = 6$	$m = 7$	$m = 8$	Limit	$m = 6$	$m = 7$	$m = 8$
[13] (MIP)	6,614	21,126	29,895	36,000	36/38	20/38	7/38
New Model (MIP)	160	1252	2541	3,600	38/38	38/38	15/38
New Model (CP)	12	115	1193	3,600	38/38	38/38	28/38

■ **Table 3** Performance New Optimization Model with CP on Large Instances.


	$m = 6$	$m = 7$	$m = 8$	$m = 9$	$m = 10$
Average Runtime (s)	14	111	1,101	2,140	2,502
Solved Instances	49/49	49/49	37/49	23/49	18/49
	$m = 11$	$m = 12$	$m = 13$	$m = 14$	$m = 15$
Average Runtime (s)	2,754	2,757	2,753	2,761	2,758
Solved Instances	13/49	13/49	13/49	13/49	13/49

seconds per instance is imposed in all experiments. The circuit design problem is considered to be *solved* if an optimal circuit implementation is found and proven to be optimal, or if it is proven that the problem is infeasible. If no circuit is found, or if optimality cannot be proven within the time limit, then the time limit is reported as the runtime.

Performance New Optimization Model. The performance of the new optimization model is first compared to the optimization model by [13]. To the best of our knowledge, [13] is the first MIP approach to the MCT quantum circuit design problem and therefore closests to the current work. Results are compared for the 38 reversible functions that overlap with the benchmark suite in this paper and experiments are conducted for $m \in \{6, 7, 8\}$ number of gates.

Table 2 demonstrates that the new optimization model completely outperforms previous work. Even accounting for the difference in hardware (6 cores vs. 24 cores), the new model is an order of magnitude faster when solved with Gurobi. When solved with CP-SAT, the new runtimes even improve by another magnitude. For $m = 6$ for example, CP is over 500x times faster than the runtime reported in [13]. It can also be seen that the new model solves significantly more instances. All benchmarks with $m = 7$ gates can now be solved, where only 20 out of 38 were solved previously. Further inspection of the results reveals that every instance that is solved by [13] is solved by the new model as well, regardless of whether MIP or CP was used. When CP fails to solve the problem (which only happens for $m = 8$), the best feasible solution is never worse than the best feasible solution obtained by [13]. The best performance is clearly obtained by the new model with CP, and this is the setting that is used for the remainder of the experiments.

Larger Instances. The improved performance of the new optimization model motivates experiments on the full benchmark suite for $m = 6$ up to $m = 15$ gates. This includes reversible functions such as `rd53` and `decod24-enable` that were not previously considered, and a number of gates that far exceeds the $m = 8$ gates in the experiments by [13].



■ **Figure 6** Solution Status of New Model with CP at 1-Hour Time Limit.

Table 3 shows that a big step was made in handling larger instances. All instances with up to $m = 7$ gates can now be solved in a matter of minutes on average, including the more complicated circuits. While most instances with $m = 8$ gates can still be solved within one hour, the average runtime starts to rise sharply at this point. It is clear that more work remains to be done to solve the largest instances, although 13 out of the 49 instances with $m = 15$ gates can already be solved.

Figure 6 details the final solution status (infeasible, optimal, suboptimal, or timeout) for each of the large instances. The newly added reversible functions are 449 up to `sym6` on the right side of the figure. It is interesting to observe that these functions require a relatively large number of gates. For example, `rd53` (marked with a \star) is proven infeasible for $m = 6$ and $m = 7$, and the solver was unable to find a feasible solution for $m \in \{8, 9, 10\}$, which are presumably infeasible. At $m = 11$, however, the solver fails to prove optimality but does find a feasible circuit. This circuit, shown in Appendix C, improves the state of the art with a quantum cost of 47, using only seven qubits. Compare this to [14], for example, who obtain a circuit with quantum cost 86 that requires 12 qubits. This demonstrates the benefit of using a model that can handle a larger number of gates, as good results may be obtained even when optimality cannot be proven.

Effect of Symmetry-Breaking Constraints. Figure 7 demonstrates the benefit of using symmetry-breaking constraints (3a)-(3c). Without symmetry-breaking constraints, all instances with $m = 6$ and $m = 7$ gates can still be solved, but the average solution times are 29% and 139% longer, respectively. For $m \geq 8$ gates, the difference in solvability becomes apparent. Out of the largest instances with $m = 15$ gates, only `graycode6` can be solved without breaking symmetries, while 13 instances can be solved when the constraints are included. It is not surprising that the symmetry-breaking constraints are more effective for longer circuits (i.e., large m values), as they are expected to have more symmetric solutions. More interestingly, adding the constraints outperforms the built-in symmetry detection in CP-SAT, which suggests that the symmetries observed in this paper are not obvious to detect automatically.

Comparative Analysis. A comparative analysis is provided to show how optimization-based methods fit in with other methods considered in the literature. Papers are selected that synthesize the entire circuit from scratch (as opposed to post-processing), that report quantum cost and computation time for every experiment, and for which the benchmark suite overlaps significantly with the current paper. This results in five studies that are summarized by




Figure 7 Comparison Symmetry-Breaking Constraints.

Table 4 Summary of Papers for Comparative Analysis.

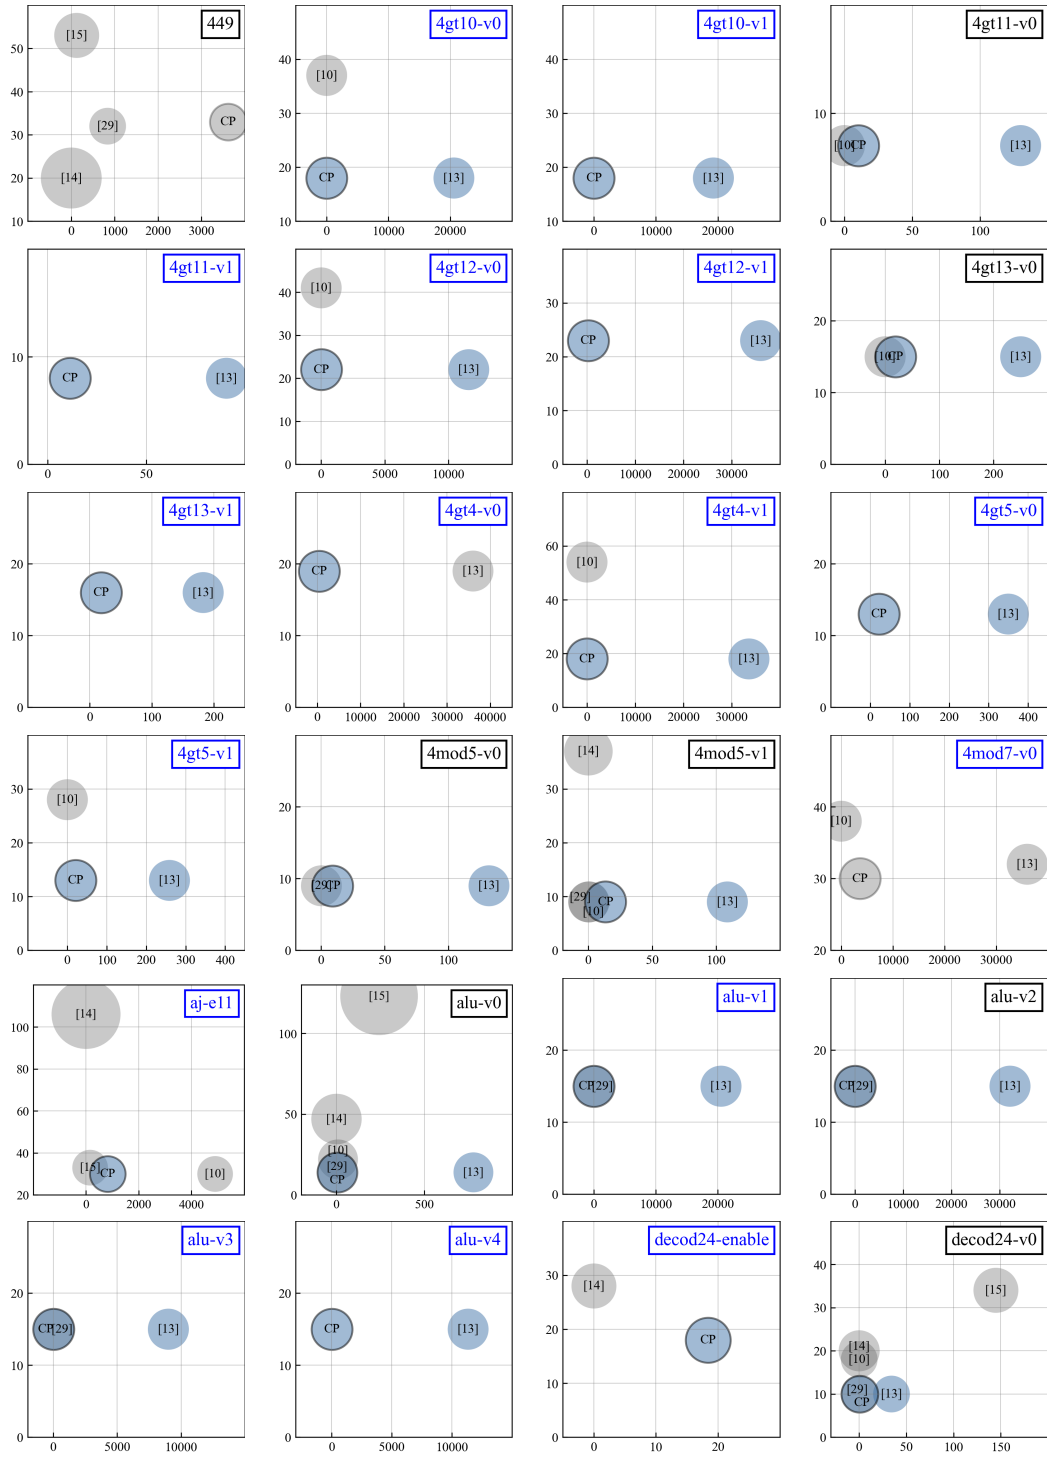
Paper	Method	Objective	Type	Gate Lib.	Max Time
[15]	Reed-Muller + decision diagram	Gate count	Heuristic	MCT	600s
[14]	Subgraph matching + decision diagram	Qubit count	Heuristic	MCT up to two controls	<1s
[10]	Satisfiability problem	Gate count	Exact	MCT	5,000s
[29]	Quantified Boolean satisfiability problem	Gate count	Exact	MCT	2,000s
[13]	Optimization model + MIP solver	Quantum cost	Exact	MCT	36,000s
Current	Optimization model + CP solver	Quantum cost	Exact	MCT	3,600s

Table 4. The papers provide a mix of exact and heuristic methods that provide different trade-offs in terms of solution time and solution quality. Also note that while all papers report quantum cost, the methods themselves often use a different objective function as a proxy, such as minimizing the number of gates or the number of qubits.

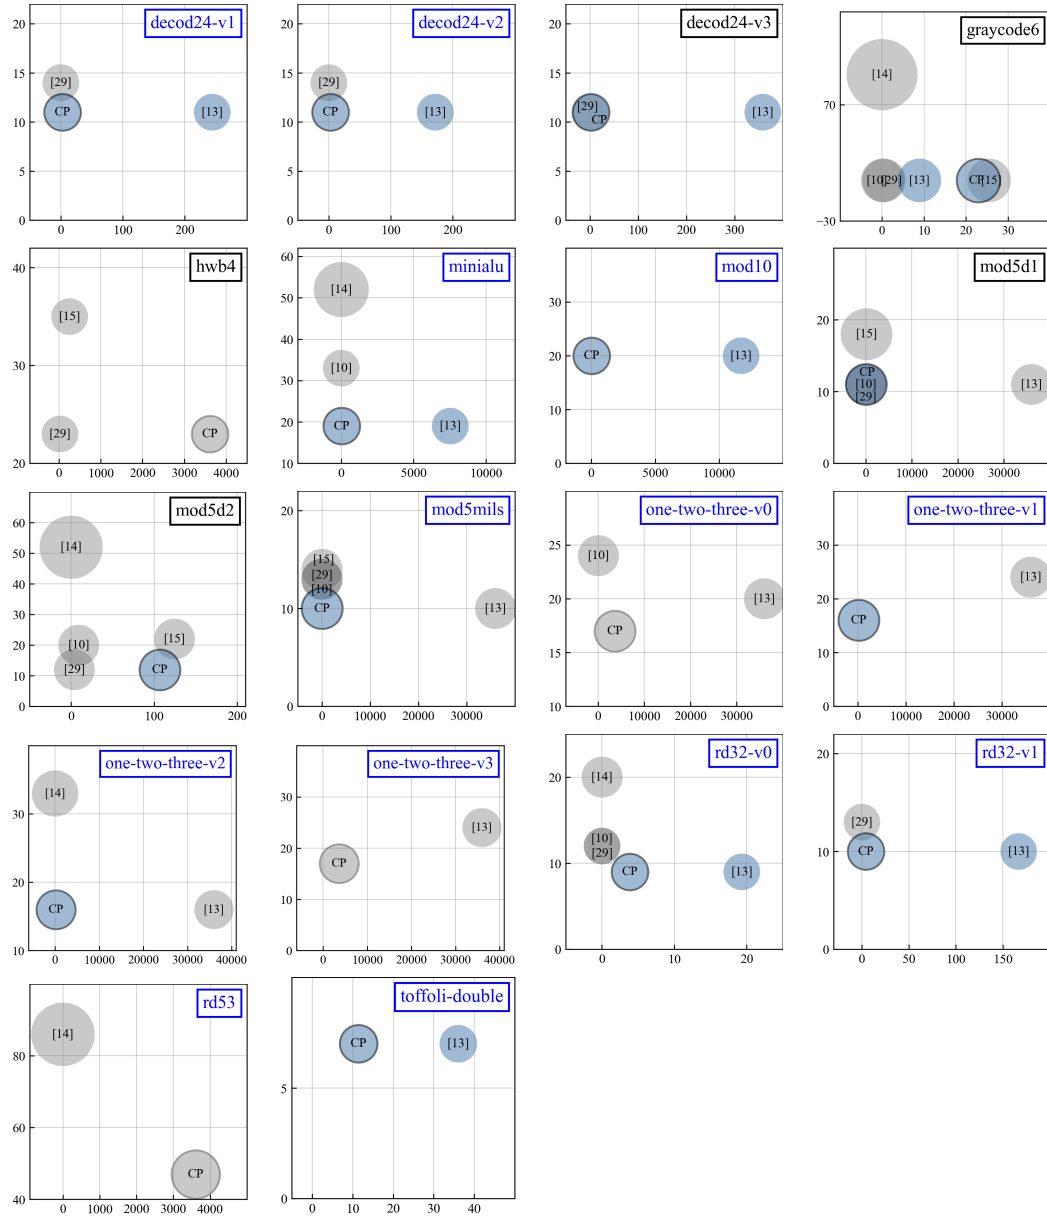
Figure 8 includes 42 plots that show the performance of each method with a circle on the two-dimensional *time-quantum cost* plane. As for the results of this study, we select a solution with the smallest m value that brings the lowest quantum costs regardless of the proven optimality. Thus, circles with black borders indicate the current paper, and blue circles indicate that the solution was proven to be optimal in the selected m in terms of quantum cost, which is only optimized directly by [13] and this paper. While the current paper only considers circuits for a given number of qubits, [14] and [15] introduce additional qubits to obtain a feasible design. The number of qubits used is indicated by the relative size of the circle in Figure 8. A circuit is considered better if it has a lower quantum cost, and methods are preferred when they have a shorter solution time and introduce fewer ancilla qubits. That is, small circles in the lower-left are preferred.

Figure 8 shows that the current method outperforms the other methods in quantum cost (or in solution time if quantum cost is tied) in 29 out of 42 cases. These function names are marked by a blue box in the top right corner. Even when the new model does not reduce quantum cost or solution time, it can still provide a benefit of guaranteed optimality (e.g., 4gt11-v0, 4gt11-v0, 4mod5-v0, 4mod5-v1, alu-v2, decod24-v0, decod24-v3, graycode6, mod5d1, mod5d2) or fewer ancilla qubits (e.g., 449). It should be noted, however, that the improved performance comes at the cost of a longer solution time: if only limited time is available, some of the other methods are better suited to provide good solutions quickly.

16:14 A New Optimization Model for MCT Quantum Circuit Design



■ **Figure 8** Comparison of Best Results with Previous Studies.



■ **Figure 8** Comparison of Best Results with Previous Studies. (*cont.*).

In particular, [29] provides high-quality solutions in a short amount of time, but improvements are still possible for some instances. For `decod24-v1` for example, [29] presents a circuit with six gates and a quantum cost of 14. This solution is found by first minimizing the number of gates, and then minimizing the quantum cost when the number of gates is fixed to six. The new optimization model, however, can directly solve the case with seven gates to find a circuit with quantum cost 11. This indicates that even for a small circuit of only four qubits, significant savings may still be obtained: a 21% cost reduction in this case. Overall, six new circuits have been found with new best-known quantum costs. These circuits are presented in Appendix C.

6 Conclusion

This paper introduced a new optimization model and symmetry-breaking constraints for the MCT quantum circuit design problem. The new model simplifies earlier work and drastically reduces the number of binary variables in the formulation. Computational experiments have shown that the new model allows both CP and MIP solvers to significantly improve solving time, with up to two orders of magnitude speedup when the CP solver is used. Experiments with larger instances of up to seven qubits and 15 gates have resulted in six new circuits with a lower cost than the previously best known. It was also shown that the symmetry-breaking constraints are very effective, especially in larger instances. Finally, a detailed comparison demonstrated that optimization models may require more time, but can provide superior circuits with guaranteed optimality.

There are several directions that may be explored in future work. While the new optimization model is effective, technical work remains to be done to scale to instances with more gates and more qubits. One opportunity would be to apply decomposition methods, as the structure of the proposed optimization model is such that the problem decomposes into independent minimum-cost flow problems when the binary variables are fixed. Another potential direction would be to extend the optimization model to different gate libraries or to directly optimize over elementary quantum gates instead of MCT gates.

References

- 1 Nabila Abdessaied, Mathias Soeken, Robert Wille, and Rolf Drechsler. Exact Template Matching Using Boolean Satisfiability. In *International Symposium on Multiple-Valued Logic*, pages 328–333, 2013. doi:10.1109/ismvl.2013.26.
- 2 Mustapha Y. Abubakar, Low Tang Jung, Nordin Zakaria, Ahmed Younes, and Abdel-Haleem Abdel-Aty. Reversible circuit synthesis by genetic programming using dynamic gate libraries. *Quantum Information Processing*, 16(6):1–24, 2017. doi:10.1007/s11128-017-1609-8.
- 3 Ravindra K. Ahuja, Thomas L. Magnanti, and James B. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice-Hall, 1993.
- 4 Mohammad AlFailakawi, Imtiaz Ahmad, Laila AlTerkawi, and Suha Hamdan. Depth optimization for topological quantum circuits. *Quantum Information Processing*, 14(2):447–463, 2015. doi:10.1007/s11128-014-0867-y.
- 5 Adriano Barenco, Charles H. Bennett, Richard Cleve, David P. DiVincenzo, Norman Margolus, Peter Shor, Tycho Sleator, John A. Smolin, and Harald Weinfurter. Elementary gates for quantum computation. *Physical Review A*, 52(5):3457–3467, 1995. doi:10.1103/physreva.52.3457.
- 6 Kamalika Datta, Gaurav Rathi, Indranil Sengupta, and Hafizur Rahaman. Synthesis of Reversible Circuits Using Heuristic Search Method. In *International Conference on VLSI Design*, pages 328–333, 2012. doi:10.1109/vlsid.2012.92.

- 7 Kamalika Datta, Indranil Sengupta, and Hafizur Rahaman. Particle Swarm Optimization Based Circuit Synthesis of Reversible Logic. In *International Symposium on Electronic System Design*, pages 226–230, 2012. doi:10.1109/ised.2012.33.
- 8 Alexandre A. A. de Almeida, Gerhard W. Dueck, and Alexandre C. R. da Silva. Reversible Circuit Optimization Based on Tabu Search. In *International Symposium on Multiple-Valued Logic*, pages 103–108, 2018. doi:10.1109/ismvl.2018.00026.
- 9 Oleg Golubitsky and Dmitri Maslov. A Study of Optimal 4-Bit Reversible Toffoli Circuits and Their Synthesis. *IEEE Transactions on Computers*, 61(9):1341–1353, 2011. doi:10.1109/tc.2011.144.
- 10 Daniel Große, Robert Wille, Gerhard W. Dueck, and Rolf Drechsler. Exact Multiple-Control Toffoli Network Synthesis With SAT Techniques. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 28(5):703–715, 2009. doi:10.1109/tcad.2009.2017215.
- 11 Pallav Gupta, Abhinav Agrawal, and Niraj K. Jha. An Algorithm for Synthesis of Reversible Logic Circuits. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 25(11):2317–2330, 2006. doi:10.1109/tcad.2006.871622.
- 12 Kazuo Iwama, Yahiko Kambayashi, and Shigeru Yamashita. Transformation rules for designing CNOT-based quantum circuits. In *Design Automation Conference*. ACM Press, 2002. doi:10.1145/513918.514026.
- 13 Jihye Jung and In-Chan Choi. A multi-commodity network model for optimal quantum reversible circuit synthesis. *PLOS ONE*, 16(6):e0253140, 2021. doi:10.1371/journal.pone.0253140.
- 14 Mridul Krishna and Anupam Chattopadhyay. Efficient Reversible Logic Synthesis via Isomorphic Subgraph Matching. In *International Symposium on Multiple-Valued Logic*, pages 103–108, 2014. doi:10.1109/ismvl.2014.26.
- 15 Chia-Chun Lin and Niraj K. Jha. RMDDS: Reed-Muller decision diagram synthesis of reversible logic circuits. *ACM Journal on Emerging Technologies in Computing Systems*, 10(2):1–25, 2014. doi:10.1145/2564923.
- 16 Dmitri Maslov and Gerhard W. Dueck. Improved quantum cost for n-bit Toffoli gates. *Electronics Letters*, 39(25):1790, 2003. doi:10.1049/el:20031202.
- 17 Dmitri Maslov, Gerhard W. Dueck, and D. Michael Miller. Toffoli network synthesis with templates. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 24(6):807–817, 2005. doi:10.1109/tcad.2005.847911.
- 18 Dmitri Maslov, Gerhard W. Dueck, and D. Michael Miller. Techniques for the synthesis of reversible Toffoli networks. *ACM Transactions on Design Automation of Electronic Systems*, 12(4):42:1–42:28, 2007. doi:10.1145/1278349.1278355.
- 19 Dmitri Maslov, Gerhard W. Dueck, D. Michael Miller, and Camille Negrevergne. Quantum Circuit Simplification and Level Compaction. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 27(3):436–444, 2008. doi:10.1109/tcad.2007.911334.
- 20 PACE. *Partnership for an Advanced Computing Environment (PACE)*, 2017. URL: <http://www.pace.gatech.edu>.
- 21 Laurent Perron and Vincent Furnon. Google OR-Tools v9.8. <https://developers.google.com/optimization/>, 2023.
- 22 Aditya K. Prasad, Vivek V. Shende, Igor L. Markov, John P. Hayes, and Ketan N. Patel. Data structures and algorithms for simplifying reversible circuits. *ACM Journal on Emerging Technologies in Computing Systems*, 2(4):277–293, 2006. doi:10.1145/1216396.1216399.
- 23 Mehdi Saeedi and Igor L. Markov. Synthesis and optimization of reversible circuits—a survey. *ACM Computing Surveys*, 45(2):1–34, 2013. doi:10.1145/2431211.2431220.
- 24 Mehdi Saeedi, Morteza Saheb Zamani, Mehdi Sedighi, and Zahra Sasanian. Reversible circuit synthesis using a cycle-based approach. *ACM Journal on Emerging Technologies in Computing Systems*, 6(4):1–26, 2010. doi:10.1145/1877745.1877747.

- 25 Trailokya Nath Sasamal, Ashutosh Kumar Singh, and Anand Mohan. Reversible Logic Circuit Synthesis and Optimization Using Adaptive Genetic Algorithm. *Procedia Computer Science*, 70:407–413, 2015. doi:10.1016/j.procs.2015.10.054.
- 26 Mathias Soeken, Robert Wille, Gerhard W. Dueck, and Rolf Drechsler. Window optimization of reversible and quantum circuits. In *IEEE Symposium on Design and Diagnostics of Electronic Circuits and Systems*, pages 341–345, 2010. doi:10.1109/ddecs.2010.5491754.
- 27 Robert Wille and Rolf Drechsler. Effect of BDD Optimization on Synthesis of Reversible and Quantum Logic. *Electronic Notes in Theoretical Computer Science*, 253(6):57–70, 2010. doi:10.1016/j.entcs.2010.02.006.
- 28 Robert Wille, Daniel Große, Lisa Teuber, Gerhard W. Dueck, and Rolf Drechsler. RevLib: An Online Resource for Reversible Functions and Reversible Circuits. In *International Symposium on Multiple-Valued Logic*, pages 220–225, 2008. doi:10.1109/ismvl.2008.43.
- 29 Robert Wille, Hoang M. Le, Gerhard W. Dueck, and Daniel Große. Quantified synthesis of reversible logic. In *Conference on Design, automation and test in Europe*, pages 1015–1020, 2008. doi:10.1145/1403375.1403620.
- 30 H. Paul Williams. *Model Building in Mathematical Programming*. John Wiley & Sons, 2013.
- 31 Wei Zhu, Zhiqiang Li, Gaoman Zhang, Suhan Pan, and Wei Zhang. A Reversible Logical Circuit Synthesis Algorithm Based on Decomposition of Cycle Representations of Permutations. *International Journal of Theoretical Physics*, 57(8):2466–2474, 2018. doi:10.1007/s10773-018-3768-5.
- 32 Alwin Zulehner and Robert Wille. Skipping Embedding in the Design of Reversible Circuits. In *International Symposium on Multiple-Valued Logic*, pages 173–178, 2017. doi:10.1109/ismvl.2017.19.

A

 Nomenclature

Symbol	Definition
Circuit Design: (1b)-(1d), (1j)	
Q	$= \{1, \dots, n\}$ set of qubits.
D	$= \{1, \dots, m\}$ set of gates.
t_q^d	variable with value 1 if qubit $q \in Q$ is the target qubit of gate $d \in D$, and 0 otherwise.
w_q^d	variable with value 1 if qubit $q \in Q$ is a control qubit of gate $d \in D$, and 0 otherwise.
Quantum Cost: (1a), (1e)-(1f), (1j)	
$f(c)$	quantum cost of a single MCT gate with $c \geq 0$ control qubits.
y_j^d	variable with value 1 if gate $d \in D$ consists of a total of $j \in Q$ target and control qubits, zero otherwise.
Quantum States and Flow Commodities: (1g)-(1i), (1k)	
Ω	$= \{0_{(2)}, \dots, (2^n - 1)_{(2)}\}$ set of pure quantum states.
Q_σ^0	$= \{q \in Q : \sigma_q = 0\}$ set of qubits that are zero in state $\sigma \in \Omega$.
K	set of indices of the flow commodities; each commodity represents a set of input quantum states that have the same (possibly incomplete) output specification.
Ω_k^{in}	$\subseteq \Omega$ set of input quantum states that represent commodity $k \in K$; together the sets $\Omega_k^{in} \forall k \in K$ provide a partition of Ω .
Ω_k^{out}	$\subseteq \Omega$ set of quantum states that meet the (possibly incomplete) output specification associated with commodity $k \in K$; the sets Ω_k^{out} may overlap, and together cover Ω .
Flow Networks: (1g)-(1i), (1k)	
V	set of vertices in each flow network; consists of source S , sink T , and nodes (σ, d) $\forall \sigma \in \Omega, d \in D \cup \{m+1\}$.
A_k	set of arcs in the flow network of commodity $k \in K$.
A_k^{flip}	$\subseteq A_k$ set of arcs for commodity $k \in K$ that represent a transition that flips a qubit.
A_k^{keep}	$\subseteq A_k$ set of arcs for $k \in K$ that represent a transition that keeps the state the same.
x_a^k	variable with value 1 if commodity $k \in K$ uses arc $a \in A_k$, and 0 otherwise.
$\delta_k^+(v)$	$\subseteq A_k$ set of arcs for $k \in K$ coming out of vertex $v \in V$.
$\delta_k^-(v)$	$\subseteq A_k$ set of arcs for $k \in K$ coming into vertex $v \in V$.
$d(a)$	$\in D$ shorthand for the gate associated with arc $a \in A_k^{flip} \cup A_k^{keep}$.
$q(a)$	$\in Q$ shorthand for the qubit that is flipped by arc $a \in A_k^{flip}$.
$\sigma(a)$	$\in \Omega$ shorthand for the state that arc $a \in A_k^{flip} \cup A_k^{keep}$ transitions from.


B

 Proof Proposition 1


► **Proposition 1.** *Any swappable circuit can be turned into an unswappable circuit by repeatedly applying Swap 1-3.*

Proof. The number of permutations of the gates is finite, so it is sufficient to prove that the swaps can be applied in a way that avoids cycling. First repeatedly apply Swap 1, which results in all empty gates moving to the end of the circuit. These empty gates are not affected by Swap 2 and 3, and hence will stay in place and Swap 1 will not be applicable again. Let τ be a vector of length m that contains the target qubit of each gate, or zero otherwise. Every Swap 2 strictly decreases τ in lexicographical order, which avoids cycling. For example, the swap in Example 5a-5b changes τ from $(4, 2)$ to $(2, 4)$. Whenever no Swaps 2 can be made, repeatedly applying Swap 3 has the effect of sorting groups with the same target qubit by the number of control qubits, and does not introduce cycles. Swap 3 does not affect the vector τ , so no cycles are introduced when returning to Swap 2 after Swap 3 is exhausted. Eventually, an unswappable circuit is obtained after finitely many steps. ◀


C New Best-Known Circuits




■ **Figure 9** rd53 for $m = 11$ (quantum cost: 47, optimality proven: no).




■ **Figure 10** 4mod7-v0 for $m = 10$ (quantum cost: 30, optimality proven: no).




■ **Figure 11** decod24-enable for $m = 6$ (quantum cost: 18, optimality proven: yes).



■ **Figure 12** one-two-three-v0 for $m = 9$ (quantum cost: 17, optimality proven: no).



■ **Figure 13** one-two-three-v1 for $m = 8$ (quantum cost: 16, optimality proven: yes).



■ **Figure 14** one-two-three-v3 for $m = 9$ (quantum cost: 17, optimality proven: no).