The 30th International Conference on Principles and Practice of Constraint Programming

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

Jihye Jung jihye.jung@gatech.edu
Kevin Dalmeijer dalmeijer@gatech.edu
Pascal Van Hentenryck pvh@gatech.edu
H. Milton Stewart School of Ind. and Syst. Engineering, Georgia Institute of Technology

1. Introduction (1/3)

Motivation

- Efficient quantum circuit design has become an important area of quantum computing to mitigate current hardware errors.
- *Primary challenge*: To implement a target function using gates from a preset gate library to minimize the circuit costs according to a given metric.
- Our problem scope:
  - **A. Target function — Reversible Boolean function**
    : A key component that embeds the input data in most quantum algorithms.
  - **B. Preset gate library — Multiple-Control Toffoli (MCT) gate**
    : A typical high-level gate commonly used to represent reversible Boolean functions.
  - **C. Circuit cost metric — Quantum cost**
    : The number of low-level quantum gates required to realize the high-level gates in the circuit.

1. Introduction (2/3)

Previous Studies

- Algorithms utilizing preconfigured circuit templates and post-synthesis
- Algorithms leveraging on representations of reversible Boolean functions
  : e.g., Cycle representation, Reed-Muller expansion
- Various heuristic algorithms
  - Quantum multiple-valued decision diagram, A* algorithm, Isomorphic subgraph matching
  - Genetic algorithm, Genetic programming, Tabu search, Particle swarm optimization
- **Exact algorithms that guarantee optimality for given evaluating metrics**
  - Iterative satisfiability problems, Quantified Boolean formula satisfiability
    ⇒ To minimize the number of high-level gates
  - Mixed-integer programming
    ⇒ To minimize the total costs of high-level logical gates (= the number of low-level gates)

1. Introduction (3/3)

Contributions

- A new optimization model and new symmetry-breaking constraints.
  : Significantly expedites the solving with both CP and MIP solvers 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.
- Several new best-known circuits for well-known benchmarks.
- Extensive comparison with other synthesis approaches.
  : Shows that optimization approaches may require more time but can provide superior circuits with guaranteed optimality.

2. Terminologies (1/4)

A. Qubits

- Analogous to classical bits in classical computers.
- Classical bits assume values of 0 or 1 to define a single basis state (i.e., a binary vector).
- Qubits store superposed states (i.e., a complex vector) formed as a convex combination of the basis states.

B. Quantum Gates

- 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.
- Multiple quantum gates may be combined into a quantum circuit to represent more complicated functions.

2. Terminologies (2/4)

C. Reversible Boolean Function

- A bijective function where inputs and outputs are provided as binary strings of fixed length (i.e., typically, the number of qubits in the system).
- Considered fundamental operators in quantum computing.
- Corresponds to a unique permutation.
- Some instances are incompletely specified with *don't care* qubits ('-').

[Completely Specified]

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

[Incompletely Specified]

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

2. Terminologies (3/4)

D. Multiple Control Toffoli (MCT) Circuit

- MCT circuits consist of a sequence of MCT gates.
- One target qubit (⊕ symbol) + zero or more control qubits (• symbol).

  If all the control qubits are in state 1, then the target qubit flips the input state.
- Control qubits do not have to be adjacent.
- A vertical line connect the control qubits to the target qubit.
- Example implementation

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

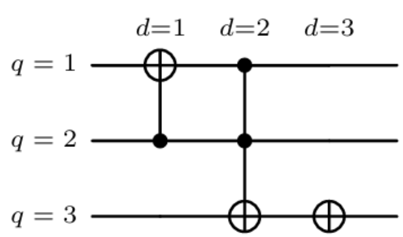

2. Terminologies (4/4)

E. Quantum Costs

- 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.
- Quantum cost $f(c)$ for an MCT gate that uses a total of $c \geq 0$ control qubits.

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	·

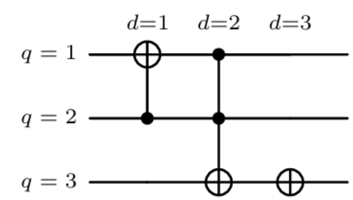

Total 1 + 5 + 1 = **7**

3. Problem Description (1/2)

A. Circuit Design

- Set of qubits $Q = \{1, \dots, n\}$ / Set of gates $D = \{1, \dots, m\}$
- t_q^d and w_q^d: Binary variables that indicate whether (q, d) contains a target or control qubit

B. Quantum Costs

- An MCT gate with $c \geq 0$ control qubits incurs a quantum cost of $f(c)$
- y_j^d: A binary indicator that takes value one if gate $d \in D$ contains exactly $j \in Q$ target and control qubits, or zero otherwise.

C. Flow Networks

- Indicates which state transitions are available depending on the design of the circuit

Case 1: Gate d flips some qubit $\bar{q} \in Q$

[$\bar{q}$ is the target qubit **] AND [** None of the controls are on qubits with value 0 in state σ **]**

Case 2: Gate d keeps state σ the same.

[No target qubit **] OR [** A target qubit $\bar{q}$, but at least one of the controls is on a zero state **]**

3. Problem Description (2/2)

C. Flow Networks (Cont'd)

- **Example:** Input state **010** → Output state **11–**
  - Gate $d = 1$ carries out the transition $010 \rightarrow 110$, i.e., vertex $(010, 1)$ to $(110, 2)$
  - Gate $d = 2$ carries out the transition $110 \rightarrow 111$, i.e., vertex $(110, 2)$ to $(111, 3)$
- A total transition of **010** → **111**: Aligning with output specification **11-** (110 or 111)

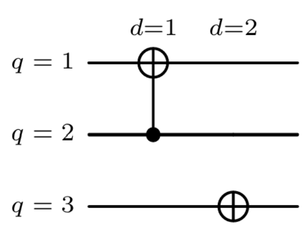

Input	Output	Impl. 2b	Input	Output	Impl. 2b
000	00-	001	100	101	101
001	00-	000	101	100	100
010	11-	111	110	011	011
011	---	110	111	010	010

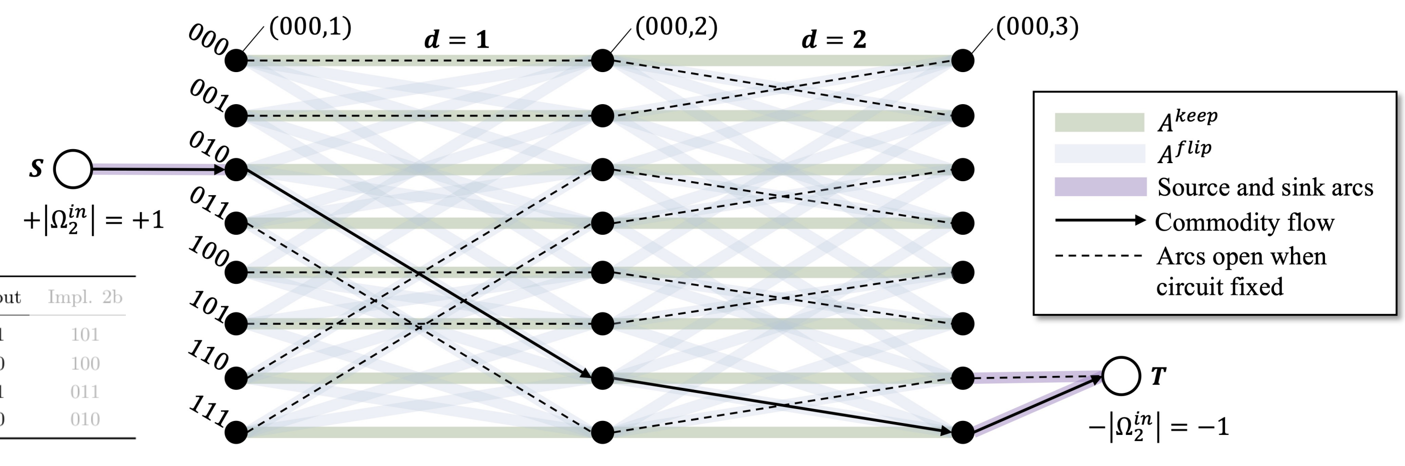

4. Optimization Model

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

Objective (Minimize Quantum Cost)

$$\text{s.t.} \quad t_q^d + w_q^d \le 1 \qquad \forall q\in Q, d\in D, \tag{1b}$$

Circuit Design

$$\sum_{q\in Q} t_q^d \le 1 \qquad \forall d\in D, \tag{1c}$$

$$w_q^d \le \sum_{r\in Q} t_r^d \qquad \forall q\in Q, d\in D, \tag{1d}$$

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

Quantum Cost

$$\sum_{j\in Q} y_j^d \le 1 \qquad \forall d\in D, \tag{1f}$$

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

Flow Networks

$$\bigvee_{q^0\in Q^0_{\sigma(a)}} \left(w_{q^0}^{d(a)} = 1\right) \bigvee \left(t_{q(a)}^{d(a)} = 0\right) \implies x_a^k = 0 \qquad \forall k\in K, a\in A_k^{flip}, \tag{1h}$$

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

Linear Reformulation

$$x_a^k \le t_{q(a)}^{d(a)} \qquad \forall k\in K, a\in A_k^{flip},$$

$$x_a^k \le 1 - w_{q^0}^{d(a)} \qquad \forall k\in K, a\in A_k^{flip}, q^0\in Q^0_{\sigma(a)},$$

$$x_a^k \le 1 - \sum_{q\in Q} t_q^{d(a)} + \sum_{q^0\in Q^0_{\sigma(a)}} w_{q^0}^{d(a)} \qquad \forall k\in K, a\in A_k^{keep}.$$

$$t_q^d, w_q^d, y_j^d \in \{0,1\} \qquad \forall q,j\in Q, d\in D, \tag{1j}$$

Binary Variables

$$x_a^k \in \{0,1\} \qquad \forall k\in K, a\in A_k. \tag{1k}$$

5. Symmetry-Breaking Constraints

Gate Swaps for Symmetry-Breaking

- Swap 1 (*Empty gate.*) If gate d is empty and $d+1$ is full, then swap two gates.
- Swap 2 (*Different target.*) If the target qubit q of gate d is at a higher line than the target qubit r of gate $d+1$ ($q > r$), and the target qubits do not neighbor a control qubit in each other gates, then swap two gates.
- Swap 3 (*Same Target.*) If gate d and gate $d+1$ have the same target qubit and gate d has fewer control bits, then swap the two gates.

$$\sum_{q\in Q} t_q^d \geq \sum_{q\in Q} t_q^{d+1} \qquad \forall d \in D, \tag{3a}$$

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

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

6. Computational Experiments (1/7)

Experiment Settings

- **Language:** Python 3.11
- **OS/Machine:** Linux / Dual Intel Xeon Gold 6226 CPUs (24 cores in total) / PACE Phoenix cluster
- **CP Solver:** CP-SAT 9.8.3296 with 24 workers (threads)
- **MIP Solver:** Gurobi 11.0.0
- **Instances:** RevLib (Wille et al., 2008)
  ** 49 functions with up to seven qubits that have known circuit implementations in fewer than 100 gates*
- **Time limit:** 3600 seconds per instance

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. URL: http://www.revlib.org, doi:10.1109/ismvl.2008.43.

6. Computational Experiments (2/7)

Performance New Optimization Model (vs. MIP)

- 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 the MIP solver.

	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

[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.

6. Computational Experiments (3/7)

Performance New Model with CP on Large Instances

- All instances with up to $m = 7$ gates can now be solved in a matter of minutes on average.
- Average runtime rises sharply at $m = 8$.
- More work remains to be done to solve the largest 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

6. Computational Experiments (4/7)

Performance New Optimization Model with CP on Large Instances

- The final solution status (*infeasible*, *optimal*, *suboptimal*, or *timeout*) for each of the instances.
- If no circuit is found, or if optimality cannot be proven within the time limit, then the time limit is reported as the runtime.

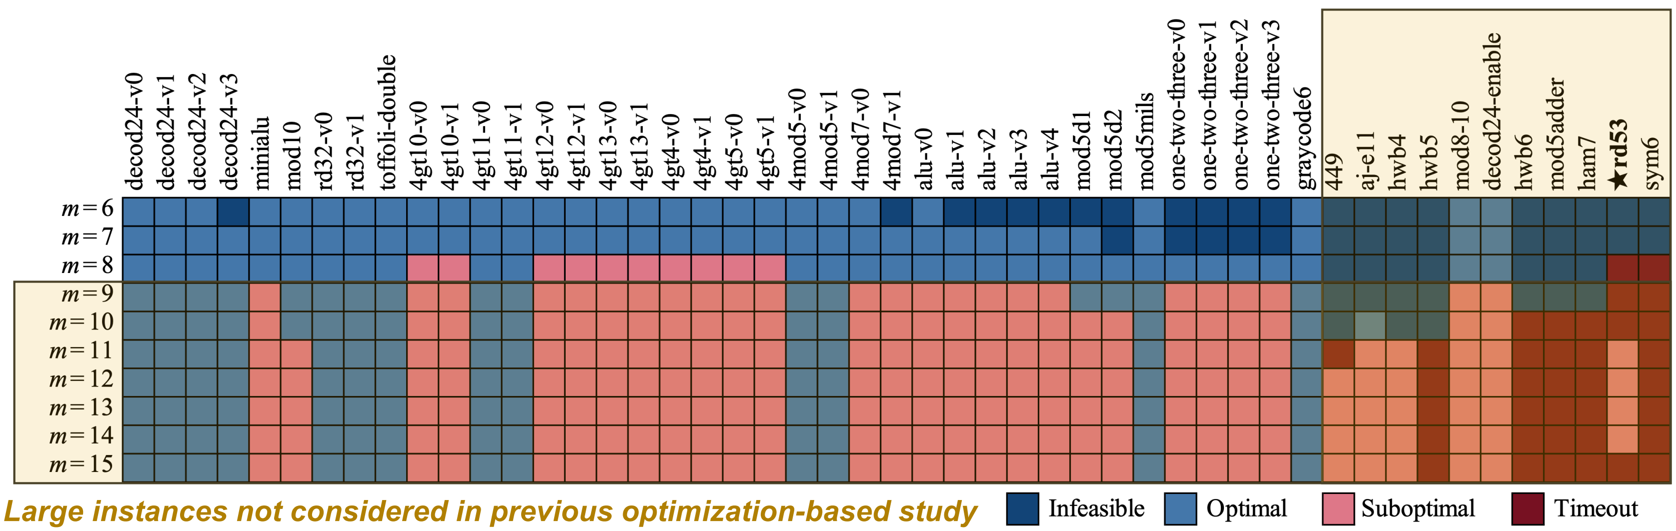

6. Computational Experiments (5/7)

Effect of Symmetry-Breaking Constraints

- For $m \geq 8$ gates, the difference in solvability becomes apparent.
- Out of the largest instances with $m = 15$ gates, only one instance can be solved without breaking symmetries, while 13 instances can be solved when the constraints are included.
- Our symmetry-breaking constraints outperforms the built-in symmetry detection in CP-SAT.

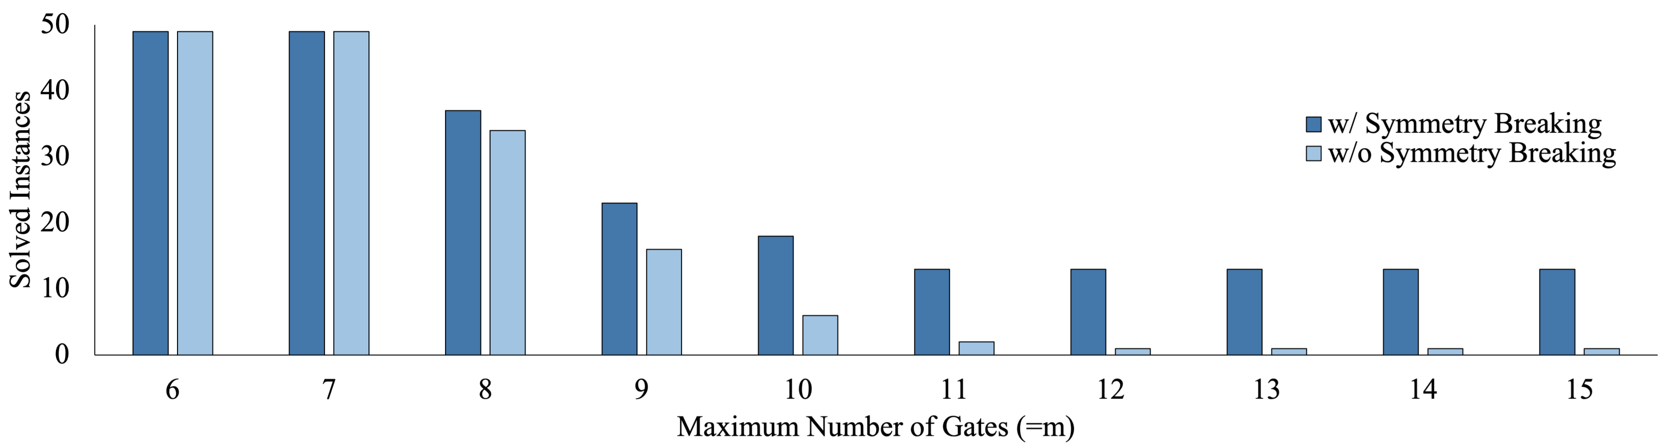

6. Computational Experiments (6/7)

Comparative Analysis: Five benchmark studies selected

- Studies that propose synthesis for the entire circuit from scratch
- Studies that report quantum cost and computation time for every experiment
- Studies where the benchmark suite overlaps significantly with our work.

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

6. Computational Experiments (7/7)

Comparative Analysis

- *Time-Quantum cost* plane
  : i.e., lower-left bubbles implies best performance
- Size of each bubble
  : The number of qubits used
- Upper-right blue box
  : CP performs the best
- Blue-faced circles
  : Optimality proven

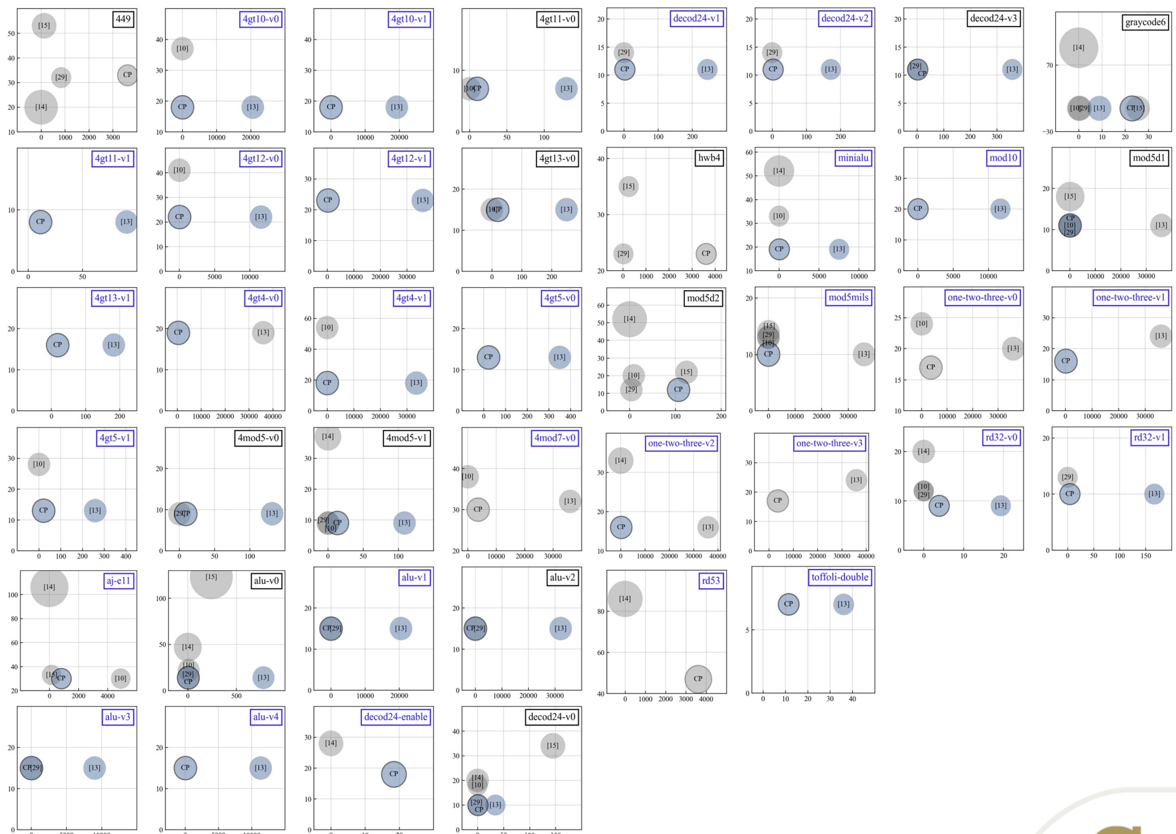

7. Ongoing Future Works

- To apply the decomposition method to utilize the decomposable structure of the new model.
- To extend the optimization model to different high-level gate libraries.
- To directly optimize over elementary quantum gates instead of high-level gates.

Thank you for listening

Appendix: Nomenclature

Symbol	Definition
Circuit Design: (1b)-(1d), (1j)	
Q	$= \{1, \ldots, n\}$ set of qubits.
D	$= \{1, \ldots, m\}$ 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	$= \{0_{(2)}, \ldots, (2^n - 1)_{(2)}\}$ 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 Ω_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}	$\subset A_k$ set of arcs for commodity $k \in K$ that represent a transition that flips a qubit.
A_k^{keep}	$\subset A_k$ 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.

Georgia Tech

Appendix: Optimization Model

Comparison to Previous Optimization Study (Jung and Choi, 2022)

A. Number of cases specified for state transition
   - Previous Study: Four cases depending on the number of target qubits/control qubits.
   - New model: Captures all cases withstate transition.

B. How the circuit is connected to opening and closing the flow arcs
   - Previous Study: $O(2nm)$ binary variables identifying whether the gate modifies state σ gate for each state $\sigma \in \Omega$ and gate $d \in D$.
   - New model: A much more direct way to close arcs through improved constraints, resulting in fewer variables and decomposable constraint structure.
     *(*A block-angular structure may be exploited by decomposition methods in future work*)

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.

Appendix: Symmetry-Breaking

Proposition on Symmetry-Breaking

Any swappable circuit can be turned into *an unswappable circuit* by repeatedly applying the defined Swap 1, 2, and 3.

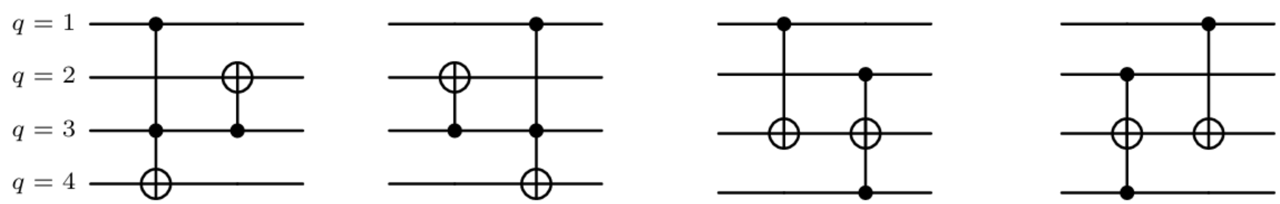

(a) Before Swap 2. **(b)** After Swap 2. **(c)** Before Swap 3. **(d)** After Swap 3.

Appendix: New Best-Known Circuits

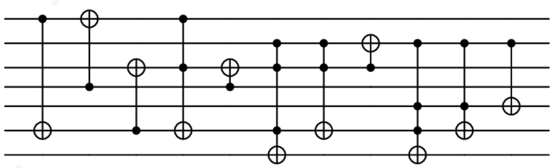

rd53
$m = 11$
Quantum cost: 47
Optimality proven: X

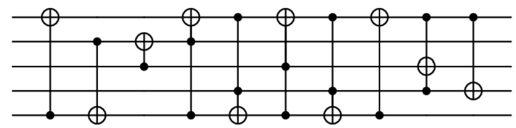

4mod7-v0
$m = 10$
Quantum cost: 30
Optimality proven: X

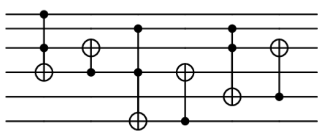

decod24-enable
$m = 6$
Quantum cost: 18
Optimality proven: O

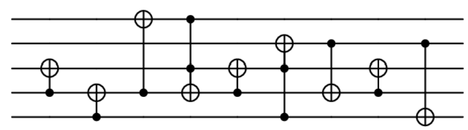

one-two-three-v0
$m = 9$
Quantum cost: 17
Optimality proven: X

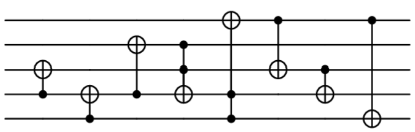

one-two-three-v1
$m = 8$
Quantum cost: 16
Optimality proven: O

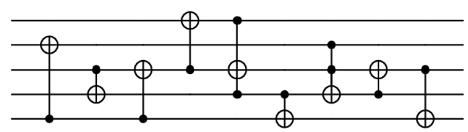

one-two-three-v3
$m = 9$
Quantum cost: 17
Optimality proven: X