

How Many Quantum Circuit Identities Are Needed to Generate All Others?

Anonymous Authors

Anonymous Institution

Abstract. Quantum compiler optimization relies on rewriting rules derived from equivalences between quantum circuits, yet prior work has identified thousands of such identities, creating substantial challenges for their storage, management, and effective application. For many widely used unitary gate sets, including Clifford+T, this apparent complexity is largely redundant, raising a fundamental question: *How many quantum circuit identities are actually needed to generate all others?* In this work, we provide strong evidence that a small set of independent identities suffices to generate all circuit equivalences of bounded depth. Surprisingly, for circuits on up to nine qubits in which each side of an equality has depth at most ten, fewer than twenty independent identities are sufficient to derive all others, and for circuits on up to five qubits with depth at most ten, only 17 rules—each involving at most three qubits—are enough. These results enable significantly more compact and efficient rewriting systems for quantum compiler optimization and reveal underlying algebraic structure in common gate sets, showing that the vast majority of known circuit identities are consequences of a small foundational basis.

Keywords: Quantum Computing · Rewriting · Equational Theory.

1 Introduction

Quantum circuit optimization is essential for improving the performance and reliability of near-term devices, as reducing circuit size directly increases execution fidelity. Traditional compilers rely on manually designed rewrite rules and fixed optimization schedules, whereas recent automated approaches enumerate candidate circuits and verify equivalences to synthesize rewrite rules. Superoptimizers such as QUARTZ generate transformations via exhaustive enumeration and search-based rewriting [23], while other methods synthesize large collections of symbolic rules and apply probabilistic verification to produce thousands of rules—or even complete optimizers—for specific architectures in short time [22]. These advances demonstrate the power of automation, but also expose fundamental limitations in scalability and interpretability.

Despite this progress, two critical gaps remain. First, synthesis-based methods produce large and highly redundant rule sets whose application requires costly pattern matching over circuit graphs and scales poorly in practice; moreover, the practical value of most discovered rules—their frequency of use, generality across benchmarks, and mutual interactions—remains largely unexplored. Second, existing strategies treat rule discovery primarily as a combinatorial search problem, driven by enumeration or

randomized sampling, without exploiting the algebraic structure underlying quantum circuit identities. As a result, exploration is inefficient and provides little insight into why so many rules arise or how they relate to one another.

Classical Boolean circuits offer a striking contrast. Circuits over AND, OR, and NOT admit a compact and transparent equational theory in which all rewrites follow from a small set of fundamental laws such as commutativity, associativity, distributivity, and De Morgan’s rules. Quantum circuits, by contrast, appear far more complex due to superposition, entanglement, and phase interactions. Indeed, prior work reports over 3,000 rewrite rules even for seven-qubit circuits, suggesting that quantum circuit identities form a fragmented and opaque collection rather than a coherent algebraic system. Moreover, composing even two simple equalities can generate seemingly new identities, revealing a substantial degree of redundancy and an apparent combinatorial explosion in the space of circuit equivalences.

Together, these observations raise a fundamental question.

How many quantum circuit identities are actually needed to generate all others?

Progress has been made in equational theories for quantum circuits [7,6], establishing minimal and complete axiom systems and showing that completeness in unrestricted settings typically requires infinitely many rules. Although mathematically elegant, these infinitely many rules offer limited insight for realistic bounded regimes. It remains unclear which equalities are necessary and sufficient and whether a small, manageable rule set can capture most circuit behavior in practice.

We adopt a complementary empirical perspective, asking whether a small set of natural identities suffices to capture the circuit equalities that matter in practice. To this end, we (i) analyze outputs of modern synthesis tools to expose systematic redundancy among discovered rules, and (ii) identify a small set of natural algebraic identities that form the basis of compact equational theories. In summary, this paper makes the following contributions:

- We uncover small, finite, and near-complete equational theories for several common quantum gate sets, bridging the gap between infinite but impractical axioms and large, redundant rule sets from synthesis tools.
- We demonstrate that these theories are relatively complete, capturing all circuit equalities up to the relevant size bound and reproducing all rewrite rules generated by existing synthesizers.
- We show that the resulting theories are concise, structured, and human-interpretable, providing both theoretical and practical insights into the minimal algebraic structure of quantum circuit optimization.

Together, these results suggest that equality-theory-driven synthesis and optimization provides a parsimonious and interpretable alternative to blind enumeration, while building a production-quality optimizer on top of such finite theories remains future work. Additionally, our generated rules provide principled guidance and empirical guarantees for designing effective rewriting strategies, capturing the most frequently applicable transformations in quantum circuit optimizations.

As a concrete illustration, Figure 1 presents a generated equational theory for Clifford+T circuits consisting of only 17 rules that derive all five-qubit identities up to




Fig. 1: An equational theory for Clifford+T circuits consisting of 17 rules that can derive all 5-qubit identities up to 16 gates. We simply use a quantum circuit C to represent an identity $C = I$ in the equational theory.

16 gates, with the largest rule containing 11 gates. Remarkably, each rule involves at most three qubits, yet together they suffice to generate all five-qubit circuit equalities. In other words, the seemingly complex space of five-qubit identities used for circuit optimization can be fully captured by only three-qubit equalities, revealing an unexpectedly compact and structured underlying system. In contrast to prior approaches that either yield infinite but impractical axiom systems [7] or thousands of highly redundant rules [22], our method discovers small, finite theories that are nearly complete and practically useful. Moreover, the resulting rules are concise, structured, and human-interpretable, revealing unexpected simplicity in common quantum gate sets. For example, two rules in Figure 1 can be written as

$$CZ \cdot CX = (X \otimes I) \cdot CX \cdot CZ, \quad S^\dagger \otimes I = CX \cdot CY \cdot CZ.$$

These transparent identities not only facilitate efficient reasoning about circuit equivalence but also illuminate the underlying algebraic structure of quantum circuits.

The remainder of the paper is organized as follows. Section 2 introduces the necessary preliminaries. Section 3 presents our synthesis algorithm. Section 4 describes our method for pruning and minimizing equational theories. Section 5 evaluates our approach on multiple gate sets. Section 6 reviews related work, and Section 7 concludes. We defer most proofs to the appendix.

2 Preliminaries

2.1 Quantum States and Operations

A *quantum state* describes the physical state of a quantum system. A single-qubit pure state $|\psi\rangle$ is a unit vector in a two-dimensional Hilbert space and may be written as a superposition of $|0\rangle$ and $|1\rangle$. An n -qubit system is represented in a 2^n -dimensional Hilbert space, enabling superposition and entanglement.

Quantum operations are modeled as unitary transformations on this Hilbert space. A unitary matrix U satisfies $U^\dagger U = I$ and preserves state norms. It acts on pure and mixed states as $|\psi\rangle \mapsto U|\psi\rangle$ and $\rho \mapsto U\rho U^\dagger$, respectively. Common single-qubit gates include H , T , and the Pauli gates $\{I, X, Y, Z\}$, while standard two-qubit gates include SWAP and CNOT.

$$I = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix},$$

$$T = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{pmatrix}, \quad H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}, \quad \text{SWAP} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad \text{CNOT} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}.$$

2.2 Quantum Circuits

Quantum circuits implement quantum computations by composing quantum gates. A circuit is a sequence of gates acting on qubits, typically depicted as wires connected by gate symbols, and its semantics is given by the composition of the corresponding unitary operations. In this work, we adopt the port graph formalism to obtain a minimal and canonical representation of quantum circuits. Under this formalism, a circuit is modeled as a directed graph whose vertices represent quantum gates and whose edges represent qubit wires. Each vertex is equipped with a set of *ports* that serve as connection points for edges and correspond to the qubits on which the gate operates. To capture the reversibility of quantum operations, each port functions simultaneously as an input and an output. We formalize this representation as follows.

Definition 1 (Quantum Circuit). A quantum circuit over a quantum gate set $\mathcal{G}$ is a tuple $C = (V, g, m)$. Specifically,

- V is a finite set of vertices (nodes), each representing a quantum gate.
- $g : V \rightarrow \mathcal{G}$ labels each vertex with a gate from $\mathcal{G}$.
- $m : \text{Ports}(C) \leftrightarrow \text{Ports}(C)$ is a bijection mapping (perfect matching) that connect each port to its associated port.
- The set of all ports $\text{Ports}(C)$ is defined as $\{(v, k) \mid v \in V, 1 \leq k \leq g(v).\text{arity}\} \cup \{(io, k) \mid 1 \leq k \leq n\}$ where the first part represents all the ports associated with the gates and in the second part, we use special label io to denote the input/output qubit ports of the circuit. Each port is identified by a pair consisting of a vertex (or io) and a port index.

- m connects ports in forward direction, and the inverse mapping m^{-1} , connect ports in backward direction. For example, m will map (i_0, k) ports to the initial gate applied on qubit k , and m^{-1} will map (i_0, k) to the final gate applied on k .
- For the well-formedness of quantum circuits, we require that the corresponding directed graph (V, E) is acyclic, where $E = \{(v_1, v_2) \mid (v_1, k_1) \in P, m(v_1, k_1) = (v_2, k_2)\}$ is the set of directed edges generated by ports. This guarantees there exists a topological order of executing gates in V , so that no gate will be applied before its input qubits are ready.

For convenience, we call $|C| = |V|$ the size of the quantum circuit. Similarly, we denote the set of all n -qubit quantum circuits over gate set $\mathcal{G}$ as $\text{Circ}_n(\mathcal{G})$. We specify the size using superscripts: $\text{Circ}_n^k(\mathcal{G})$ for exactly k gates, and $\text{Circ}_n^{\leq k}(\mathcal{G})$ for at most k gates. We also use diagrammatic representation for quantum circuits, where we represent the circuit C using a conventional quantum circuit diagram with n wires, as shown below, where each gate represents a vertex, each line represents a connection between ports, and the left/right endpoints represents input/output ports:

$$\begin{array}{c} i_{01} \text{ --- } \\ \vdots \\ i_{0n} \text{ --- } \end{array} \left[\begin{array}{c} \text{---} \\ \vdots \\ \text{---} \end{array} \right] C \begin{array}{c} \text{---} \\ \vdots \\ \text{---} \end{array} \begin{array}{c} i_{01} \\ \vdots \\ i_{0n} \end{array} = \text{---}^n \boxed{C} \text{---}$$

The reason of using this representation of circuits is its topological treatment of permutation gates: rather than appearing as distinct nodes, these operations are implicitly encoded through the permutation of edges. Consequently, this yields a canonical form where the underlying graph structure remains invariant regardless of the explicit placement or arrangement of permutation gates. To show this, we define the following operations for quantum circuits.

- Composition $\circ : \text{Circ}_n(\mathcal{G}) \times \text{Circ}_n(\mathcal{G}) \rightarrow \text{Circ}_n(\mathcal{G})$ is graphically depicted as

$$\text{---}^n \boxed{C_2} \text{---} \circ \text{---}^n \boxed{C_1} \text{---} = \text{---}^n \boxed{C_1} \text{---} \boxed{C_2} \text{---}$$

- Tensor product $\otimes : \text{Circ}_n(\mathcal{G}) \times \text{Circ}_m(\mathcal{G}) \rightarrow \text{Circ}_{n+m}(\mathcal{G})$ is depicted as

$$\text{---}^n \boxed{C_1} \text{---} \otimes \text{---}^m \boxed{C_2} \text{---} = \begin{array}{c} \text{---}^n \boxed{C_1} \text{---} \\ \text{---}^m \boxed{C_2} \text{---} \end{array}$$

- Identity circuits $id_n : \text{Circ}_n(\mathcal{G})$ which is graphically identical to ---^n .
- Permutation circuits $\sigma^c : \text{Circ}_n(\mathcal{G})$ which represent the permutation of qubits according to the symmetry $\sigma \in \text{Sym}(n)$, formally defined as $\sigma^c = (\emptyset, \emptyset, m)$ where $m(i_0, k) = (i_0, \sigma(k))$ for $1 \leq k \leq n$.
- Adjoint circuits $(\cdot)^\dagger : \text{Circ}_n(\mathcal{G}) \rightarrow \text{Circ}_n(\mathcal{G})$ which is defined as $(V, g, m)^\dagger = (V, g^\dagger, m^{-1})$ where $g^\dagger(v) = g(v)^\dagger$ for $v \in V$.

- Generator circuits $G(k_1, \dots, k_m) : \text{Circ}_n^1(\mathcal{G})$ which represents the application of gate $G \in \mathcal{G}$ on qubits $k_1, \dots, k_m$ (where $m = G.\text{arity}$), formally defined as $G(k_1, \dots, k_m) = (\{v\}, \lambda v.G, m)$ where

$$\begin{aligned} m(\text{io}, k_i) &= (v, i) & m(v, i) &= (\text{io}, k_i) & \text{for } 1 \leq i \leq m, \\ m(\text{io}, j) &= (\text{io}, j) & & & \text{for } j \notin \{k_1, \dots, k_m\}. \end{aligned} \quad (2.1)$$

We denote the set of all generator circuits as $\text{Gen}_n(\mathcal{G})$, note that $\text{Gen}_n(\mathcal{G}) \subsetneq \text{Circ}_n^1(\mathcal{G})$ for $n > 1$ because additional permutation can be embedded in $\text{Circ}_n^1(\mathcal{G})$.

These operations induce several basic properties of our circuit representation.

1. Our representation of quantum circuits remain the same regardless of different composition of permutations gates, which can be easily observed from

$$(\sigma_2 \circ \sigma_1)^c = \sigma_2^c \circ \sigma_1^c \text{ and } (\sigma^{-1})^c = (\sigma^c)^\dagger. \quad (2.2)$$

2. Generator circuits can pass through σ^c by permuting the qubits it acts on, i.e.,

$$\sigma^c \circ G(k_1, \dots, k_m) = G(\sigma(k_1), \dots, \sigma(k_m)) \circ \sigma^c. \quad (2.3)$$

3. Finally, $\sigma^c \circ C \circ (\sigma^{-1})^c$ forms a circuit of C with only the qubit wires permuted in the circuit diagram, without really adding any additional permutation gates in the circuit diagram.

These properties imply that our representation of quantum circuits effectively captures the essential structure of quantum operations while abstracting away the superficial arrangement of permutation gates. This canonical form simplifies the analysis and manipulation of quantum circuits, particularly in the context of equational reasoning. Now let us define the equational logic for quantum circuits.

2.3 Quantum Circuit Equational Logic

In this section, we establish the foundations of equational logic and equational theory as applied to quantum circuits. We begin by defining the syntax for quantum circuit equations, which serves as the basis for our equational proof system.

Definition 2 (QCEL-Equation). A quantum circuit equation over a quantum gate set $\mathcal{G}$ is an n -qubit quantum circuit $C \in \text{Circ}_n(\mathcal{G})$ in the canonical form $C \approx id_n$. For convenience, we also write $C_1 \approx C_2$ as a shorthand for $C_1 \circ C_2^\dagger \approx id_n$ where $C_1, C_2 \in \text{Circ}_n(\mathcal{G})$.

With this notation, we then define the equational proof system for quantum circuits via the inference rules presented in Figure 2.

Definition 3 (QCEL-Derivability). Let Γ be a set of quantum circuit equations over a quantum gate set $\mathcal{G}$, and let C be an n -qubit quantum circuit. The judgment

$$\Gamma \vdash C \approx id_n$$

signifies that $C \approx id_n$ is derivable from Γ using the inference rules in Figure 2.

EQ	ASSUME	ADJOINT	EXTEND
$\frac{}{\Gamma \vdash id_n \approx id_n}$	$\frac{C \approx id_n \in \Gamma}{\Gamma \vdash C \approx id_n}$	$\frac{\Gamma \vdash C \approx id_n}{\Gamma \vdash C^\dagger \approx id_n}$	$\frac{\Gamma \vdash C \otimes id_m \approx id_{n+m}}{\Gamma \vdash C \otimes id_k \approx id_{n+k}}$
ROTATE	PERMUTATION	TRANSITIVITY	
$\frac{\Gamma \vdash C_1 \circ C_2 \approx id_n}{\Gamma \vdash C_2 \circ C_1 \approx id_n}$	$\frac{\Gamma \vdash C \approx id_n \quad \sigma \in \text{Sym}(n)}{\Gamma \vdash \sigma^c \circ C \circ (\sigma^c)^\dagger \approx id_n}$	$\frac{\Gamma \vdash C_1 \approx C_2 \quad \Gamma \vdash C_2 \approx C_3}{\Gamma \vdash C_1 \approx C_3}$	

Fig. 2: Equational Proof Rules for Quantum Circuits

To provide intuition for the inference rules in Figure 2:

- The ROTATE rule permits the cyclic rotation of circuits within an equation. Physically, this holds because $\llbracket C_1 \rrbracket \llbracket C_2 \rrbracket = I$ implies $\llbracket C_2 \rrbracket = \llbracket C_1 \rrbracket^\dagger$, which consequently implies $\llbracket C_2 \rrbracket \llbracket C_1 \rrbracket = I$.
- The PERMUTATION rule exploits the symmetry of qubits. If a circuit C reduces to the equation, any permutation of the wires in C (conjugated by that permutation) remains an equation operation.
- The TRANSITIVITY rule enables the chaining of equations and circuit substitution. For instance, $C_1 \circ C_2 \circ C_3 \approx id_n$ is equivalent to $C_2 \approx C_1^\dagger \circ C_3^\dagger$. Thus, if $C_2 \approx C_4$, we may deduce $C_4 \approx C_1^\dagger \circ C_3^\dagger$, or equivalently $C_1 \circ C_4 \circ C_3 \approx id_n$.

With this in mind, we can now define the semantics for our notation of quantum circuit equations. Since our paper doesn't involve symbolic circuits, we directly define the semantics using the quantum gates' unitary matrix interpretation.

Definition 4 (QCEL-Identity). Let $C \approx id_n$ be an n -qubit quantum circuit equation over a quantum gate set $\mathcal{G}$, we call $C \approx id_n$ an *identity* if it holds that:

$$\models C \approx id_n \iff \llbracket C \rrbracket = e^{i\alpha} I^{\otimes n}$$

where $\llbracket C \rrbracket$ is the unitary matrix on $U(2^n)$ derived using the quantum gate set $\mathcal{G}$, and $\alpha \in \mathbb{R}$ represents a global phase that may arise during the computation.

Similar to circuits, we denote the set of n -qubit quantum circuit identities over $\mathcal{G}$ as $\text{Id}_n(\mathcal{G})$, and use $\text{Id}_n^k(\mathcal{G})$ and $\text{Id}_n^{\leq k}(\mathcal{G})$ to specify equations with exactly k gates and at most k gates, respectively. We also call an equation set Γ as an *equational theory* if all equations in Γ are identities, denoted as $\models \Gamma$.

In this paper, we adopt a simple definition of the soundness and completeness of the equational logic, defined as follows.

Theorem 1 (QCEL Soundness). For any equational theory Γ over a quantum gate set $\mathcal{G}$ and any n -qubit quantum circuit C , if $\Gamma \vdash C \approx id_n$ and $\models \Gamma$, then $\models C \approx id_n$.

Definition 5 (QCEL Relative Completeness). An equational theory Γ over gate set $\mathcal{G}$ is complete with respect to an equation set $\mathcal{I}$ if, for any equation $C \approx id_n \in \mathcal{I}$, $\models C \approx id_n \iff \Gamma \vdash C \approx id_n$.

3 Identity Synthesis Algorithm

We present an algorithm to synthesize all quantum circuit identities for a given gate set and size bound. To evaluate empirical completeness, we systematically enumerate circuits and test equivalence using canonical representatives, following [3]. In contrast to prior work limited to Clifford circuits, our approach extends to universal gate sets by relying only on equivalence classes generated by permutation gates.

3.1 Overview

First we define our representative reduce function $\text{REDUCE} : \text{U}(2^n) \rightarrow \text{U}(2^n)$ that maps a quantum operator U to its canonical representative in the equivalence class $\{\llbracket \sigma_2^c \rrbracket U \llbracket \sigma_1^c \rrbracket \mid \sigma_1, \sigma_2 \in \text{Sym}(n)\}$. In other words, REDUCE satisfies the following:

$$\text{REDUCE}(U_1) = \text{REDUCE}(U_2) \iff \exists \sigma_1, \sigma_2 \in \text{Sym}(n) \text{ s.t. } U_1 = \llbracket \sigma_2^c \rrbracket U_2 \llbracket \sigma_1^c \rrbracket \quad (3.1)$$

A specific implementation of the function of REDUCE is described in Section 3.2. Now let us define the result of the algorithm. Given a quantum gate set $\mathcal{G}$, a circuit size k , and number of qubits n , our synthesis algorithm aims to compute $\mathcal{R}_n^k[\mathcal{G}] : \text{U}(2^n) \rightarrow \mathcal{P}(\text{Circ}_n^k(\mathcal{G}))$ defined as follows:

$$\mathcal{R}_n^k[\mathcal{G}](U) = \left\{ C \mid \llbracket C \rrbracket = \text{REDUCE}(\llbracket C \rrbracket) = U, C \in \text{Circ}_n^k(\mathcal{G}) \right\} \quad (3.2)$$

We also write $\mathcal{R}_n^{\leq k}[\mathcal{G}](U) = \bigcup_{1 \leq i \leq k} \mathcal{R}_n^i[\mathcal{G}](U)$ similar to circuits and identities.

Intuitively, $\mathcal{R}_n^k[\mathcal{G}]$ contains all circuits of size k over n qubits evaluate to the canonical representative U . Note that $\mathcal{R}_n^k[\mathcal{G}]$ provides a compact representation of all pair of equivalent circuits in $\text{Circ}_n^k(\mathcal{G})$ and their unitaries, according to the following lemma:

Lemma 1. $\text{Circ}_n^k(\mathcal{G}) = \left\{ \sigma_2^c \circ C \circ \sigma_1^c \mid C \in \mathcal{R}_n^k[\mathcal{G}](U), U \in \text{U}(2^n), \sigma_1, \sigma_2 \in \text{Sym}(n) \right\}.$

Proof. ($\supseteq$) Trivial from the definition of $\text{Circ}_n^k(\mathcal{G})$. ($\subseteq$) For any circuit $C' \in \text{Circ}_n^k(\mathcal{G})$, let $U = \text{REDUCE}(\llbracket C' \rrbracket)$. By the definition of REDUCE , there exist $\sigma_1, \sigma_2 \in \text{Sym}(n)$ such that $U = \llbracket \sigma_2^c \rrbracket \llbracket C' \rrbracket \llbracket \sigma_1^c \rrbracket$. Let $C = \sigma_2^c \circ C' \circ \sigma_1^c$, then $\llbracket C \rrbracket = U$ and $C \in \mathcal{R}_n^k[\mathcal{G}](U)$. Therefore, $C' \in \mathcal{R}_n^k[\mathcal{G}]$. $\square$

Then we also need to prove that $\mathcal{R}_n^k[\mathcal{G}]$ can be used to naturally converted to a set of quantum circuit identities by combining every two circuits that evaluate to the same unitary. Specifically, for each unitary U such that $\mathcal{R}_n^k[\mathcal{G}](U) \neq \emptyset$, we can generate the identities $C_1 \approx C_2$ for all distinct pairs $C_1, C_2 \in \mathcal{R}_n^{\leq k}[\mathcal{G}](U)$. The following lemma states that this conversion yields all circuit identities of size up to $2k$.

Lemma 2. $\text{Id}_n^{\leq 2k}(\mathcal{G}) = \left\{ \sigma^c \circ C_1 \circ C_2^\dagger \circ (\sigma^{-1})^c \approx id_n \mid C_1, C_2 \in \mathcal{R}_n^{\leq k}[\mathcal{G}](U), \sigma \in \text{Sym}(n) \right\}.$

Algorithm 1: Quantum Circuit Synthesis Algorithm with REDUCE

Input : Quantum gate set $\mathcal{G}$, Maximum circuit size k , and number of qubits n .
Output : A mapping $\mathcal{R}_n^{\leq k}[\mathcal{G}] : \mathcal{U}(2^n) \rightarrow \mathcal{P}(\text{Circ}_n^{\leq k}(\mathcal{G}))$.

```

1 fn SYNTH( $\mathcal{G}, k, n$ ):
2    $R^1 = \mathcal{R}_n^1[\mathcal{G}]$ 
3   for  $i \leftarrow 2$  to  $k$ :
4     invariant  $R^{i-1} = \mathcal{R}_n^{i-1}[\mathcal{G}]$  //  $\mathcal{R}_n^k[\mathcal{G}]$  (3.2), C.f. Theorem 2
5     Initialize  $R^i = \lambda U. \emptyset$ 
6     for  $U \in \mathcal{U}(2^n)$  that  $R^i(U) \neq \emptyset$ :
7       for  $G \in \text{Gen}_n(\mathcal{G})$ : //  $\text{Gen}_n(\mathcal{G})$  (2.1)
8          $U' = \text{REDUCE}(\llbracket G \rrbracket U)$  // See Section 3.2
9         for  $\sigma_1, \sigma_2 \in \text{Sym}(n)$  that  $\exists \alpha. U' = e^{i\alpha} \llbracket \sigma_2^c \rrbracket \llbracket G \rrbracket U \llbracket \sigma_1^c \rrbracket$ :
10           $R^i(U') = R^i(U') \cup \{\sigma_2^c \circ G \circ C \circ \sigma_1^c \mid C \in R^{i-1}(U)\}$ 
11   return  $R^1 \cup R^2 \cup \dots \cup R^k$ 

```

Using these notation, we can now describe our synthesis algorithm as shown in Algorithm 1. The algorithm iteratively constructs the mappings $\mathcal{R}_n^1[\mathcal{G}], \mathcal{R}_n^2[\mathcal{G}], \dots, \mathcal{R}_n^k[\mathcal{G}]$. In each iteration i , it initializes an empty mapping R^i and populates it by extending circuits from the previous iteration R^{i-1} with one additional gate from the gate set $\mathcal{G}$. For each unitary U in the previous mapping R^{i-1} , it considers all possible single-gate extensions G and computes the new unitary $U' = \text{REDUCE}(\llbracket G \rrbracket U)$. It then finds all permutations σ_1, σ_2 such that U' can be expressed as $\llbracket \sigma_2^c \rrbracket \llbracket G \rrbracket U \llbracket \sigma_1^c \rrbracket$. Finally, it updates the mapping R^i by adding the new circuits formed by $\sigma_2^c \circ G \circ C \circ \sigma_1^c$ for each circuit C in $R^{i-1}(U)$. The following theorem states the correctness of the algorithm.

Theorem 2. Algorithm 1 correctly computes the mapping $\mathcal{R}_n^{\leq k}[\mathcal{G}]$ for the given quantum gate set $\mathcal{G}$, circuit size k , and number of qubits n .

3.2 Implementation of REDUCE

Now we describe our implementation of the REDUCE function used in Algorithm 1. The goal of REDUCE is to find the canonical representative of a unitary U under the equivalence class generated by two different permutation gates on both sides: $\{\llbracket \sigma_2^c \rrbracket U \llbracket \sigma_1^c \rrbracket \mid \sigma_1, \sigma_2 \in \text{Sym}(n)\}$. To achieve this, we first introduce the notion of *seemingly ordered quantum state* that helps us to reduce the search space of possible symmetries.

Definition 6 (Seemingly Ordered Quantum State). A n -qubit quantum state $|\psi\rangle$ is called to be *seemingly ordered* if for any two qubit $1 \leq i < j \leq n$, we have:

$$(\langle \psi | 2^i \rangle, \langle \psi | 2^n - 1 - 2^i \rangle) \sqsubseteq (\langle \psi | 2^j \rangle, \langle \psi | 2^n - 1 - 2^j \rangle)$$

where $\sqsubseteq$ is the lexicographical order on complex numbers defined as: for any two complex numbers $a = a_r + ia_i$ and $b = b_r + ib_i$, $a \sqsubseteq b$ if and only if either $a_r < b_r$, or $a_r = b_r$ and $a_i \leq b_i$. We denote the set of all symmetries $\sigma \in \text{Sym}(n)$ that converts the quantum state $|\psi\rangle$ into a seemingly ordered state $\llbracket \sigma^c \rrbracket |\psi\rangle$ as $\text{Sym}^{\approx}(\psi)$.

Algorithm 2: Implementation of the REDUCE function

Parameters: ψ_l a randomized quantum state on symmetric qubits. ψ_r a fully randomized quantum state.

Input : A unitary U on n qubits.

Output : $U_{\min} = \text{REDUCE}(U)$ and $\Sigma = \{(\sigma_l, \sigma_r) \mid \exists \alpha. \sigma_l U \sigma_r = e^{i\alpha} U_{\min}\}$.

```

1 fn REDUCEIMPL( $U$ ) :
2   Initialize  $\psi_{\min} = \top$ ,  $U_{\min} = U$ ,  $\Sigma = \emptyset$ 
3    $\psi_0 = \text{dephase } U^\dagger |\psi_l\rangle$  // dephase (3.3), C.f. Sec 3.3 (5)
4   foreach  $\sigma_r \in \text{Sym}^\approx(\psi_0)$  : //  $\text{Sym}^\approx$  Definition 6
5      $|\psi\rangle = \text{dephase } U \llbracket \sigma_r^c \rrbracket^\dagger |\psi_r\rangle$ 
6     assert  $|\text{Sym}^\approx(\psi)| = 1$  // See Theorem 3
7      $\{\sigma_l\} = \text{Sym}^\approx(\psi)$ 
8     if  $\llbracket \sigma_l^c \rrbracket \psi \sqsubset \psi_{\min}$  :
9        $\Sigma = \emptyset$ ;  $\psi_{\min} = \llbracket \sigma_l^c \rrbracket \psi$ ;  $U_{\min} = \llbracket \sigma_l^c \rrbracket U \llbracket \sigma_r^c \rrbracket^\dagger$ 
10    if  $\llbracket \sigma_l^c \rrbracket \psi = \psi_{\min}$  :
11      assume  $U_{\min} = e^{i\alpha} \llbracket \sigma_l^c \rrbracket U \llbracket \sigma_r^c \rrbracket^\dagger$  for some  $\alpha \in \mathbb{R}$  // See Theorem 4
12       $\Sigma = \Sigma \cup \{(\sigma_l, \sigma_r^\dagger)\}$ 
13  return  $U_{\min}, \Sigma$ 

```

Algorithm 2 shows our implementation of the REDUCE function. The key idea is to apply U twice, once backward, to find possible right-side symmetries, and once forward to find the corresponding left-side symmetries. To do this, we need two initial quantum states: the *left-initial* state $|\psi_l\rangle$ is a randomized quantum state on symmetric qubits, i.e., $\llbracket \sigma^c \rrbracket |\psi_l\rangle = |\psi_l\rangle$ for any $\sigma \in \text{Sym}(n)$; and the *right-initial* state $|\psi_r\rangle$ is a fully randomized quantum state, i.e., $|\text{Sym}^\approx(\psi_r)| = 1$. We require the left-initial state to be symmetric so that after applying U from the backward, the resulting state $|\psi_0\rangle$ can help us to reduce the search space of possible right-side symmetries. The right-initial state is fully randomized to ensure that $|\psi\rangle$ has a unique left-side symmetry mapping it to a canonical form after applying U and a right-side symmetry.

In Algorithm 2, we first apply U from the backward. We compute $U^\dagger |\psi_l\rangle$ and then normalize the state using the following phase removing function:

$$\text{dephase } |\psi\rangle = \frac{|\langle \psi | 0 \rangle|}{\langle \psi | 0 \rangle} |\psi\rangle \quad (3.3)$$

This normalization ensures that the amplitude of the $|0\rangle$ basis state is always a non-negative real number, which helps to eliminate any redundant phase factor arisen from quantum circuits. Then, we obtain state $|\psi_0\rangle$, to be used to reduce search space of right-side symmetries. Next, we iterate through all symmetries $\sigma_r \in \text{Sym}^\approx(\psi_0)$. Each of such a symmetry σ_r represents a potential permutation on the right side of U .

For each such symmetry, we compute forward by applying $U \llbracket \sigma_r^c \rrbracket^\dagger$ to the right-initial state $|\psi_r\rangle$. This results in a new quantum state $|\psi\rangle$. Since $|\psi_l\rangle$ is fully-randomized and independent to $U \llbracket \sigma_r^c \rrbracket^\dagger$, very likely, $|\psi\rangle$ has a unique left-side symmetry that sorts it into a seemingly ordered state. We assume this because: 1) This assertion holds up to floating-point numerical precision, as shown in Theorem 3; 2) in our experiments, we

have never encountered a case where $|\text{Sym}^\approx(\psi)| > 1$. Therefore, we directly extract the unique left-side symmetry σ_l from $\text{Sym}^\approx(\psi)$.

We then compare the resulting state $[\![\sigma_l^c]\!]|\psi\rangle$ with the current minimum state $\psi_{\min}$ using the lexicographical order $\sqsubset$. If it is smaller, we update $\psi_{\min}$ and $U_{\min}$. If it is smaller or equal, record the corresponding symmetries (σ_l, σ_r) . After enumerating all σ_r , we return the minimal unitary $U_{\min}$ and its symmetry set Σ . Σ will be used at line 9 in Algorithm 1 to enumerate all equivalent circuits of U .

Note that we assume at line 11 that if two states are equal, then their corresponding unitaries are also equal. In Theorem 4, we prove this observation only fails due to the inaccuracies of floating-point arithmetic. To verify the assumption at line 11 in Algorithm 2 practically holds, we observed this fact that if this assumption is incorrect, the incorrect Σ will cause line 9 in Algorithm 1 to introduce incorrect equivalences in $\mathcal{R}_n^k[\mathcal{G}]$. However, in our experiments reported in Section 5, especially the correctness validation in Section 5.2, we didn't find any incorrect identities synthesized by our algorithm. Thus, we are confident about the practical correctness of this assumption.

We then discuss the correctness of Algorithm 2. First, we justify the validity of the assertion at line 6 and the assumption at line 11. The following theorems hold almost surely under sufficient numerical precision:

Theorem 3. Given a unitary U on n qubits, and a fully randomized quantum state $|\psi_r\rangle$ independent to U , $\Pr(|\text{Sym}^\approx(\text{dephase } U|\psi_r\rangle)| > 1) = 0$.

Intuitively, the condition of $\text{Sym}^\approx$ is pretty hard for random $|\psi_r\rangle$ to satisfy.

Theorem 4. Given two unitaries U_1 and U_2 , and a fully randomized quantum state $|\psi_r\rangle$ independent to both U_1 and U_2 , if $\text{dephase } U_1|\psi_r\rangle = \text{dephase } U_2|\psi_r\rangle$ holds, then the posterior probability $\Pr(\exists \alpha. U_1 = e^{i\alpha} U_2) = 1$.

It is very hard for $|\psi_r\rangle$ to be an eigenvector of $U_2^\dagger U_1$.

Next, we need to prove that Algorithm 2 practically returns a canonical representative of U under the equivalence class generated by double permutation gates. This can be shown by the following theorem, which proves the correctness of Algorithm 2.

Theorem 5. Algorithm 2 returns $U_{\min}$ satisfying (3.1) under the assertion and assumption.

Theorem 6. If the assertion and assumption hold, Algorithm 2 returns

$$\Sigma = \{(\sigma_l, \sigma_r) \mid \exists \alpha, \sigma_l U \sigma_r = e^{i\alpha} U_{\min}\}.$$

3.3 Implementation Details

We provide additional implementation notes for our synthesis algorithm here:

1. We implement our algorithm using floating-point arithmetic, which can introduce numerical errors. However, as demonstrated in Section 5.2, such numerical errors can be avoided by choosing sufficiently precise floating-point representations.

2. To minimize the execution time within REDUCE, we don't construct the full unitary matrix U of the input circuit. Instead, we only simulate the action of U using a corresponding circuit C by applying its gates sequentially on the input state.
3. To save the space to store all U in $\mathcal{R}_n^k[\mathcal{G}]$, we only store the hash values of U using a high-quality hash function as keys in the mapping. In Section 5.2, we didn't find any hash collision that causes the equivalence classes of $\mathcal{R}_n^k[\mathcal{G}](U)$ not provable using our equational theory minimization procedure.
4. In Algorithm 1, maintaining a set for all circuits at line 10 is too costly. Instead, for each $\mathcal{R}_n^k[\mathcal{G}]$, we only keep track of each quadruple $(\sigma_1, G, U, \sigma_2)$. Similarly, we also filtered out all duplicate circuits in the equivalence class of each $\mathcal{R}_n^k[\mathcal{G}](U)$, that simply adds additional gates to existing equations, which simplifies our equational theory minimization process.
5. In the dephase function, the amplitude $\langle\psi|0\rangle$ may be zero during backward evaluation. When this occurs, we instead remove the phase of the sum of amplitudes $\sum_{i=0}^{2^n-1} \langle\psi|i\rangle$ from the state $|\psi\rangle$. This guarantees the normalization step is always symmetric to different order of qubits, i.e. $\text{dephase}[\llbracket\sigma^c\rrbracket]|\psi\rangle = \llbracket\sigma^c\rrbracket \cdot \text{dephase}|\psi\rangle$. Therefore, the correctness of REDUCE is preserved.

4 Equational Theory Pruning

In this section, we describe our algorithm for pruning a given equational theory of quantum circuits. The goal of the theory pruning is to find a small set of circuit identities that can derive all circuit identities in $\text{Id}_n^{\leq k}(\mathcal{G})$. Similar to prior work on circuit optimization [23], we utilized an algorithm similar to TASO [16], which iteratively applies rewrite rules to test if an identity can be derived from other identities.

Some rules like PERMUTATION and ROTATE in Definition 3 are easy to deduce using the following canonical representative: CANONC and CANONI. Intuitively, CANONC maps a circuit to its canonical representative in the equivalence class defined by the PERMUTATION rules, while CANONI maps an identity to its canonical representative in the equivalence class defined by the PERMUTATION and ROTATE rules, which satisfy the following properties:

$$\text{CANONC}(C_1) = \text{CANONC}(C_2) \iff C_1 \approx \text{id}_n \xleftrightarrow{\text{PERMUTATION}} C_2 \approx \text{id}_n \quad (4.1)$$

$$\text{CANONI}(I_1) = \text{CANONI}(I_2) \iff I_1 \xleftrightarrow[\text{ROTATE}]{\text{PERMUTATION}} I_2 \quad (4.2)$$

We implement these two canonical representative by finding the lexicographically smallest circuit and identity in their respective equivalence classes. This can be efficiently done by enumerating all permutations of qubits and rotations of circuits, which is feasible since the number of qubits n is small in our setting.

With these two canonical representative, we can now present our equational theory pruning algorithm in Algorithm 3. The algorithm takes as input the mapping $\mathcal{R}_n^{\leq k}[\mathcal{G}]$ generated by our synthesis algorithm in Section 3 and outputs a small set of circuit identities that can derive all identities in $\text{Id}_n^{\leq 2k}(\mathcal{G})$. The main idea of the algorithm is to first convert $\mathcal{R}_n^{\leq k}[\mathcal{G}]$ into a set of circuit identities $\mathcal{I}$, then iteratively try

to prove each identity in $\mathcal{I}$ using previously proved identities and a set of rewrite rules derived from $\mathcal{R}_n^{\leq k}[\mathcal{G}]$. If an identity cannot be proved within a certain search limit, it is added to the set of assumed identities, which will be included in the final output.

Algorithm 3: Equational Theory Pruning Algorithm

Input : $R = \mathcal{R}_n^{\leq k}[\mathcal{G}]$ generated by Algorithm 1.
Output : A small set of circuit identities that can derive all identities in $\text{Id}_n^{\leq 2k}(\mathcal{G})$.

```

1 fn PRUNE( $R$ ) :
2   Initialize  $rules = \lambda C. \emptyset$ ,  $proved = \emptyset$ ,  $assumed = \emptyset$ ,  $\mathcal{I} = \emptyset$ 
3   for  $U$  if  $R(U) \neq \emptyset$  :
4     for  $C_1, C_2 \in R(U)$ ,  $C_1 \neq C_2$  :
5       Let  $\llbracket \sigma^c \rrbracket C_1 \llbracket \sigma^c \rrbracket^\dagger = \text{CANONC}(C_1)$  // CanonC (4.1)
6        $rules(\text{CANONC}(C_1)) \leftarrow rules(\text{CANONC}(C_1)) \cup \{\llbracket \sigma^c \rrbracket C_2 \llbracket \sigma^c \rrbracket^\dagger\}$ 
7        $\mathcal{I} = \mathcal{I} \cup \{\text{CANONI}(C_1 \approx C_2)\}$  // CanonI (4.2)
8   for  $I \in \mathcal{I}$  increasing by  $|I|$  :
9     invariant  $assumed \vdash proved$  // C.f. Theorem 7
10     $proved = proved \cup \text{TASOSEARCH}(I)$ 
11  fn TASOSEARCH( $I$ ) :
12     $Q = \text{PriorityQueue}(I)$ 
13     $visited = \emptyset$ 
14    while  $|Q| > 0 \wedge |visited| \geq \text{searchlimit}$  :
15       $I_0 = Q.\text{pop}()$ 
16      foreach  $\text{CANONI}(C_1 \approx C_2) = I_0$  :
17        for  $C_3 \in rules(\text{CANONC}(C_1))$  and  $\text{CANONI}(C_3 \approx C_1) \in proved$  :
18           $I_1 = \text{CANONI}(C_3 \approx C_2)$ 
19          if  $\neg visited.\text{insert}(I_1)$  : continue
20          else if  $I_1 \in proved$  : return  $visited$ 
21          else if  $|I_1| < |I|$  : return  $visited \cup \text{TASOSEARCH}(I_1)$ 
22          else :  $Q.\text{push}(I_1)$ 
23     $assumed \leftarrow assumed \cup \{I\}$ 
24    return  $visited$ 
25  return  $assumed$ 

```

In Algorithm 3, the outer function PRUNE first constructs the set of circuit identities $\mathcal{I}$ and the mapping $rules$ from $\mathcal{R}_n^{\leq k}[\mathcal{G}]$. The mapping $rules$ maps each canonical circuit $\text{CANONC}(C)$ to a set of circuits that are equivalent to C given by $\mathcal{R}_n^{\leq k}[\mathcal{G}]$. Then, our algorithm iterates every identity in set $\mathcal{I}$ in increasing order and tries to prove it using the inner function TASOSEARCH. The function TASOSEARCH performs a best-first search on the space of circuit identities, starting from the input identity I . In each iteration, it pops a minimum identity I_0 from the priority queue Q and tries to find a proved rewrite rule from $rules$ that can be applied to I_0 to generate a new identity I_1 . If I_1 has already been proved, then the function returns successfully. If I_1 is shorter than I , then

Table 1: Number of Rules

k	2	3	4	5	6	7	8	9
L	6	8	8	8	8	8	8	8
C	11	16	17	17	17	17	17	
$C+T$	10	13	14	16	17	17	17	
C^*+T	30	28	28	28				
$C+\sqrt{T}$	15	21	22	22	22	22		
$C+R_Z(\pi/3)$	15	21	22	22	22	22		

the function recursively calls itself on I_1 to skip all identities larger than I . Otherwise, it pushes I_1 back into the priority queue for further exploration. If no proof is found within the search limit, the identity I is included in *assumed*.

TASOSEARCH returns the set of all rules explored during the search, which is then added to the set of proved identities *proved* in the outer function PRUNE. Finally, after all identities in $\mathcal{I}$ have been processed, the algorithm returns the set of assumed identities *assumed* as the pruned equational theory.

Theorem 7. Algorithm 3 correctly computes an equational theory that can derive all identities in $\text{Id}_n^{\leq 2k}(\mathcal{G})$.

5 Experimental Result

We implemented our synthesis and minimization algorithms in Rust¹ and evaluated them on a machine with 20-core Intel(R) Xeon(R) E5 CPU and 128 GB RAM. We set the floating-point precision to 64 bits and the search limit in Algorithm 3 to 50000. We evaluated our algorithms on six different gate sets, defined as follows:

$$\begin{aligned}
L &= \{X, CX\} & C^*+T &= \{X, Y, Z, CX, CY, CZ, H, T, T^\dagger, S, S^\dagger\} \\
C &= \{X, CX, H, S, S^\dagger\} & C+\sqrt{T} &= \{X, CX, H, S, S^\dagger, \sqrt{T}, \sqrt{T}^\dagger\} \\
C+T &= \{X, CX, H, T, T^\dagger\} & C+R_Z(\pi/3) &= \left\{X, CX, H, S, S^\dagger, R_Z\left(\frac{\pi}{3}\right), R_Z\left(-\frac{\pi}{3}\right)\right\}
\end{aligned}$$

5.1 Synthesis and Pruning Results

In this section, we present our primary experimental results. We ran our synthesis and pruning algorithms on 5 qubits for all six gate sets. Due to computational limitations, the maximum circuit size k explored varies for each gate set, ranging from $k = 5$ for C^*+T to $k = 9$ for L. For each gate set and circuit size, we recorded the number of rewrite rules, synthesized ECCs, and canonical identities. The results appear in the tables below.

Table 1 presents the core result of our work: the size of the pruned equational theory for each gate set. A key observation is the rapid convergence to a small, constant

¹ Will release the code upon publication

Table 3: Synthesized Canonical Identities

k	2	3	4	5	6	7	8	9
L	6	19	140	705	3595	16055	46654	66904
C	12	42	306	2221	18259	168611	1652070	
$C+T$	11	36	243	1775	14349	129274	1225169	
C^*+T	64	645	9583	172251				
$C+\sqrt{T}$	18	72	496	3989	36330	366327		
$C+R_Z(\pi/3)$	18	77	499	4085	37040	372194		

number of rules. For instance, the L gate set stabilizes at 8 rules for $k \geq 3$, and the $C+T$ gate set stabilizes at 17 rules for $k \geq 6$; these 17 rules are shown in Figure 1. This convergence across different gate sets strongly supports our central hypothesis that a small, finite equational theory is sufficient for describing the vast majority of circuit identities up to a practical size. The larger number of rules for gate sets like C^*+T (30 rules at $k = 5$) is expected due to the richer set of available gates.

Table 3 lists the total number of canonical circuit identities before pruning (in $\mathcal{I}$ of Algorithm 3). When compared with Table 1, the data reveals the power of our pruning algorithm. For example, at $k = 8$ for the $C+T$ gate set, our method discovered 848,880 canonical identities, which were then proven to be derivable from a core set of just 17 rules. This massive reduction in size demonstrates that the vast majority of identities found through brute-force enumeration are redundant. Our approach effectively distills this redundancy into a compact and powerful equational theory.

5.2 Correctness

Ensuring the correctness of the synthesis and pruning process is paramount, especially given its reliance on floating-point arithmetic to represent unitary matrices. Therefore, we conducted several validation experiments to verify the accuracy and reliability of our implementation. These experiments include a comparison with a naive implementation, proving existing quantum optimizers, assessing the impact of floating-point precision, and generating exportable rules and proofs.

Comparison with a Naive Implementation We first compared our optimized implementation with a naive baseline. The baseline uses standard 64-bit floating-point numbers (f64) and a fixed-epsilon equality test for unitaries, and it does not use the equivalence classes from Section 3. Tables 7, and 5 report the results. Because the naive method is slow and memory-intensive, we could only run it on smaller circuit sizes; entries that are skipped due to our resource limits are marked with “-”.

For each synthesized ECCs of the naive method, we checked if it exists in the results of our optimized implementation. We found that all unitaries and circuits synthesized by the naive method has a corresponding entry in our optimized results for all gate sets and circuit sizes where the naive method completed. This cross-validation confirms our optimized synthesis algorithm matches the naive approach.

Table 5 shows the number of synthesized canonical identities from the naive method. We observe it is a lot larger than those from our optimized implementation (Table 1).

Table 5: Number of Naive-Synthesized Canonical Identities

k	2	3	4	5	6	7
L	5	16	202	2958	37563	438482
C	11	36	372	5449	79392	-
$C+T$	10	30	311	4758	71001	-
C^*+T	73	821	21493	-		
$C+\sqrt{T}$	17	63	558	8000	-	-
$C+R_Z(\pi/3)$	17	68	563	8088	-	-

Table 7: Number of Irreducible Rules after Adding Naive Rules

k	2	3	4	5	6	7
L	6	8	8	8	8	8
C	12	16	17	17	17	-
$C+T$	10	13	14	16	17	-
C^*+T	30	28	28	-		
$C+\sqrt{T}$	16	21	22	22	-	-
$C+R_Z(\pi/3)$	16	21	22	22	-	-

Table 9: Proving Rules from Quartz [23]

Quartz($q=3$)	L	C	$C+T$	C^*+T	$C+\sqrt{T}$	$C+R_Z(\pi/3)^2$
k	7	7	7	6	6	6
ECCs	340	103135	65106	610684	27437	580196
Identities	525	116040	70876	785177	30059	712823
Synth Time (s)	0.34	298.81	367.84	1101.27s	162.91	9974.40
Prove Time (s)	17.91	3466.2	1950.6	30917	519.67	11638.2

However, the rules synthesized by the naive method are all derivable from those synthesized by our optimized implementation for all gate sets and circuit sizes where the naive method completed. This validates the correctness of our pruning algorithm.

After we obtain our synthesized and pruned result, we continue to add naive-synthesized identities to see if any of them can not be pruned. Table 7 shows we found that for $k \geq 3$, no additional irreducible rules from the naive method can be added into our optimized pruned set. The additional irreducible rules for $k = 2$ are due to the limited pruning power caused by the rewrite rule set, *rules*, being too small in Algorithm 3. Specifically, if rewrite rules from 3 gates to 1 gates are not added into *rules*, many optimizations cannot be applied in our pruning algorithm.

Proving Existing Quantum Optimizations Second, we validated our synthesized equational theories by using them to prove the rewrite rules employed in an existing quantum circuit optimizer, Quartz [23]. By demonstrating that our synthesized rules can derive those in Quartz, we provide strong evidence for the practical utility of our approach. We extracted the rewrite rules used in Quartz for various gate sets and attempted to prove each rule using our synthesized equational theories.

Table 9 summarizes the results of this validation. We ran Quartz on 3 qubits and at most 7 gates for most gate sets, as it is the setting used in the original Quartz paper; however, the last 3 gate sets give out-of-memory for this setting, so we $k = 6$ for them. We then attempted to prove each of the rewrite rules used in Quartz using our synthesized equational theories. We proved all the rules used in Quartz for all gate sets, which showcases the level of completeness of our synthesized equational theories. The synthesis and proving times are also reported in Table 9, showing that our methods are efficient enough to handle practical sets of rewrite rules.

² Due to bugs in the GitHub version of Quartz for this gate set, we use the artifact version without pruning, which produces many more ECCs compared to $C+\sqrt{T}$.

Table 11: Number of ECCs for $C+T(k=7)$ under f64/f128 and different comparison precisions (p).

p	8	16	24	32	40	48
f64	5658695	5658695	5658695	5658695	5660176	6431388
f128	5658695	5658695	5658695	5658695	5658695	5658695

Impact of Floating-Point Precision Third, we investigate the impact of floating-point precision on the synthesis results. In our main experiments, we used 64-bit floating-point numbers to represent unitary matrices, which may raise concerns about numerical stability and correctness. To address this, we ran experiments on the $C+T$ gate set with $k=7$ using both standard 64-bit (f64) and 128-bit (f128) floating-point numbers. In these experiments, we varied the comparison precision threshold, p , which dictates how many of the most significant bits of the floating-point representation of the state vectors must match to be considered equal in Algorithm 2. This parameter is crucial for balancing comparison precision against numerical noise sensitivity.

The results are shown in Table 11. We observe that for comparison precisions p up to 32 bits, both f64 and f128 yield the same number of synthesized ECCs (5658695). This consistency indicates that our synthesis results are robust to floating-point precision within this range. However, as we increase p beyond 32 bits, discrepancies begin to emerge. At $p=40$, f64 produces a slightly different count (5660176) compared to f128 (5658695), and at $p=48$, the difference becomes more pronounced (6,431,388 for f64 vs. 5,658,695 for f128). These are caused by numerical inaccuracies in f64 arithmetic that lead to incorrect equivalence classifications.

Exportable Rules and Proofs A key benefit of our approach is that the final, pruned set of rules is small enough for manual inspection. For example, the 17 rules generated for the $C+T$ gate set (shown in Figure 1) are simple and can be individually verified. Since our framework proves that all other synthesized identities can be derived from this core set, verifying these base rules provides strong evidence for the correctness of the thousands of identities they represent.

To provide stronger confidence, and to inspect the correctness of our pruning algorithm (Algorithm 3) itself, our implementation can generate a formal, human-readable proof for any pruned identity. This feature makes the entire process transparent and allows for independent verification, ensuring that our results are both trustworthy and understandable. An example of such a proof is provided in Appendix B.

5.3 Synthesis and Pruning Time

In addition to the synthesis and pruning results presented earlier, we also measured the time taken for both processes across different gate sets and circuit sizes, focusing on the $C+T$ gate set as a representative example. Table 13 summarizes the time taken (in seconds) for the pruning and synthesis stages of our algorithm, as well as the corresponding times for a naive implementation that does not utilize equivalence classes or optimized floating-point comparisons.

From the data, we observe that our optimized implementation significantly outperforms the naive approach, especially as the circuit size k increases. For instance,

Table 13: Time for Each Stages of our Algorithm for gate set $C+T$. (s)

k	2	3	4	5	6	7	8
Synth	0.01	0.03	0.17	1.17	10.76	109.13	1119.40
Prune	0.03	0.83	1.47	12.59	258.15	7274.39	149602.22
Naive Synth	0.00	0.06	1.32	23.21	391.60	-	-
Check Equiv	0.18	2.35	42.12	772.77	1145.84	-	-
Naive Prune	0.01	0.11	1.42	42.05	2483.19	-	-
Quartz Synth	0.08	0.84	15.68	327.43	-	-	-

at $k = 6$, our pruning process takes approximately 258 seconds, while the naive method takes over 2375 seconds. The synthesis times also show a similar trend, with our method being substantially faster. Our method is also more memory efficient, allowing us to complete synthesis for $k = 7$ and $k = 8$, where the naive approach runs out of memory. We also report the synthesis time of Quartz [23] in the same settings (5 qubits). Quartz runs out of memory for $k \geq 6$ and is much slower than the naive baseline, likely due to general matrix multiplication rather than our SV-Sim-style gate-by-gate approach [17].

6 Related Work

Recent years have seen numerous quantum circuit optimizers, including heuristic methods [20,15,5] and formal methods [23,11,13,12,4]. Our work is most closely related to formal methods that leverage rewrite rules, such as Quartz [23] and QUESO [22]. These approaches typically target circuit-level optimization by synthesizing rewrite rules or equivalent circuits, without understanding and reasoning about the underlying equational theory. In contrast, we focus on synthesizing a minimal set of identities that can derive all equalities up to a given size, reducing the complexity of the underlying equational theory.

Theoretical work has also explored complete equational theories for quantum circuits [2,10,18,19,1,14,9,8,6,7], establishing minimal and complete generators for various gate sets. However, algebraic completeness does not guarantee practical utility for optimization or synthesis. Our work bridges this gap by identifying which rules are essential and sufficient in realistic settings, showing that a small, tractable set of identities suffices to capture the circuit behavior relevant for practical applications.

7 Conclusion

We presented algorithms for synthesizing and pruning equational theories of quantum circuits. Our experiments show that, for several finite gate sets, a small set of identities suffices to derive all equalities up to practically relevant sizes, revealing a surprisingly compact structure underlying circuit identities. These findings not only shed light on practical quantum compiling but also open the door to a deeper exploration of the algebraic structure of quantum circuit equalities. Future work includes (i) applying our methods to practical quantum compilers to evaluate their optimization impact, and (ii) investigating the deeper structure of equalities, including questions of finiteness and redundancy.

References

1. Amy, M., Chen, J., Ross, N.J.: A finite presentation of CNOT-dihedral operators. In: Coecke, B., Kissinger, A. (eds.) *Proceedings of the 14th International Conference on Quantum Physics and Logic, QPL 2017, Nijmegen, The Netherlands, 3–7 July 2017. Electronic Proceedings in Theoretical Computer Science*, vol. 266, pp. 84–97 (2018). <https://doi.org/10.4204/EPTCS.266.5>, arXiv:1701.00140
2. Bian, X., Selinger, P.: Generators and relations for 2-qubit Clifford+T operators. In: *Proceedings of the 18th International Conference on Quantum Physics and Logic, QPL 2022. Electronic Proceedings in Theoretical Computer Science*, vol. 394, pp. 13–28 (2023). <https://doi.org/10.4204/EPTCS.394.2>, arXiv:2204.02217
3. Bravyi, S., Latone, J.A., Maslov, D.: 6-qubit optimal clifford circuits. *npj Quantum Information* **8**(1), 79 (2022). <https://doi.org/10.1038/s41534-022-00583-7>, <https://doi.org/10.1038/s41534-022-00583-7>
4. Chareton, C., Bardin, S., Bobot, F., Perrelle, V., Valiron, B.: An automated deductive verification framework for circuit-building quantum programs. In: *Programming Languages and Systems (ESOP 2021). Lecture Notes in Computer Science*, vol. 12648, pp. 148–177. Springer (2021). https://doi.org/10.1007/978-3-030-72019-3_6, arXiv:2003.05841
5. Cirq Developers: Cirq (2024). <https://doi.org/10.5281/zenodo.4062499>, <https://github.com/quantumlib/Cirq>, see <https://quantumai.google/cirq>
6. Clément, A., Delorme, N., Perdrix, S.: Minimal equational theories for quantum circuits. In: *Proceedings of the 39th Annual ACM/IEEE Symposium on Logic in Computer Science. LICS '24, Association for Computing Machinery, New York, NY, USA (2024)*. <https://doi.org/10.1145/3661814.3662088>, <https://doi.org/10.1145/3661814.3662088>
7. Clément, A., Heurtel, N., Mansfield, S., Perdrix, S., Valiron, B.: A complete equational theory for quantum circuits. In: *2023 38th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pp. 1–13 (2023). <https://doi.org/10.1109/LICS56636.2023.10175801>
8. Cockett, J.R.B., Comfort, C.: The category TOF. In: Selinger, P., Chiribella, G. (eds.) *Proceedings of the 15th International Conference on Quantum Physics and Logic, QPL 2018, Halifax, Canada, 3–7 June 2018. Electronic Proceedings in Theoretical Computer Science*, vol. 287, pp. 67–84 (2019). <https://doi.org/10.4204/EPTCS.287.4>, arXiv:1804.10360
9. Cockett, J.R.B., Comfort, C., Srinivasan, P.V.: The category CNOT. In: Coecke, B., Kissinger, A. (eds.) *Proceedings of the 14th International Conference on Quantum Physics and Logic, QPL 2017, Nijmegen, The Netherlands, 3–7 July 2017. Electronic Proceedings in Theoretical Computer Science*, vol. 266, pp. 258–293 (2018). <https://doi.org/10.4204/EPTCS.266.18>, arXiv:1707.02348
10. Coecke, B., Wang, Q.: ZX-rules for 2-qubit Clifford+T quantum circuits. In: *Reversible Computation – 10th International Conference, RC 2018, Leicester, UK, September 12–14, 2018, Proceedings. Lecture Notes in Computer Science*, vol. 11106, pp. 144–161. Springer (2018). https://doi.org/10.1007/978-3-319-99498-7_10, arXiv:1804.05356
11. Cowtan, A., Dilkes, S., Duncan, R., Krajenbrink, A., Simmons, W., Sivarajah, S.: On the qubit routing problem. In: *14th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2019). Leibniz International Proceedings in Informatics (LIPIcs)*, vol. 135, pp. 5:1–5:32 (2019). <https://doi.org/10.4230/LIPIcs.TQC.2019.5>, arXiv:1902.08091
12. Davis, M.G., Smith, E., Tudor, A., Sen, K., Siddiqi, I., Iancu, C.: Towards optimal topology aware quantum circuit synthesis. In: *2020 IEEE International Conference on Quantum Computing and Engineering (QCE)*, pp. 223–234. IEEE (2020). <https://doi.org/10.1109/QCE49297.2020.00036>, best Paper Award at IEEE Quantum Week 2020

13. Hietala, K., Rand, R., Hung, S.H., Wu, X., Hicks, M.: A verified optimizer for quantum circuits. *Proceedings of the ACM on Programming Languages* **5**(POPL), 1–29 (2021). <https://doi.org/10.1145/3434318>, arXiv:1912.02250, Distinguished Paper at POPL 2021
14. Iwama, K., Kambayashi, Y., Yamashita, S.: Transformation rules for designing CNOT-based quantum circuits. In: *Proceedings of the 39th Annual Design Automation Conference, DAC 2002*, New Orleans, LA, USA, June 10–14, 2002. pp. 419–424. ACM (2002). <https://doi.org/10.1145/513918.514026>
15. Javadi-Abhari, A., Treinish, M., Krsulich, K., Wood, C.J., Lishman, J., Gacon, J., Martiel, S., Nation, P.D., Bishop, L.S., Cross, A.W., Johnson, B.R., Gambetta, J.M.: Quantum computing with Qiskit. arXiv preprint arXiv:2405.08810 (2024), arXiv:2405.08810
16. Jia, Z., Padon, O., Thomas, J., Warszawski, T., Zaharia, M., Aiken, A.: Taso: Optimizing deep learning computation with automatic generation of graph substitutions. In: *Proceedings of the 27th ACM Symposium on Operating Systems Principles*. pp. 47–62. SOSP '19, Association for Computing Machinery, New York, NY, USA (2019). <https://doi.org/10.1145/3341301.3359630>, <https://doi-org.ezproxy.lib.purdue.edu/10.1145/3341301.3359630>
17. Li, A., Krishnamoorthy, S.: Sv-sim: Scalable pgas-based state vector simulation of quantum circuits. In: *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis* (2021)
18. Makary, J., Ross, N.J., Selinger, P.: Generators and relations for real stabilizer operators. In: Heunen, C., Backens, M. (eds.) *Proceedings of the 18th International Conference on Quantum Physics and Logic, QPL 2021*, Gdansk, Poland, and online, 7–11 June 2021. *Electronic Proceedings in Theoretical Computer Science*, vol. 343, pp. 14–36 (2021). <https://doi.org/10.4204/EPTCS.343.2>, arXiv:2109.05655
19. Ranchin, A., Coecke, B.: Complete set of circuit equations for stabilizer quantum mechanics. *Physical Review A* **90**(1), 012109 (2014). <https://doi.org/10.1103/PhysRevA.90.012109>, arXiv:1310.7932
20. Sivarajah, S., Dilkes, S., Cowtan, A., Simmons, W., Edgington, A., Duncan, R.: t|ket): A re-targetable compiler for NISQ devices. *Quantum Science and Technology* **5**(1), 014003 (2020). <https://doi.org/10.1088/2058-9565/ab8e92>, arXiv:2003.10611
21. Sivarajah, S., Dilkes, S., Cowtan, A., Simmons, W., Edgington, A., Duncan, R.: t|ket): A re-targetable compiler for NISQ devices. *Quantum Science and Technology* **5**(1), 014003 (2020). <https://doi.org/10.1088/2058-9565/ab8e92>, arXiv:2003.10611
22. Xu, A., Molavi, A., Pick, L., Tannu, S., Albarghouthi, A.: Synthesizing quantum-circuit optimizers. vol. 7. Association for Computing Machinery, New York, NY, USA (Jun 2023). <https://doi.org/10.1145/3591254>, <https://doi.org/10.1145/3591254>
23. Xu, M., Li, Z., Padon, O., Lin, S., Pointing, J., Hirth, A., Ma, H., Palsberg, J., Aiken, A., Acar, U.A., Jia, Z.: Quartz: Superoptimization of quantum circuits. In: *Proceedings of the 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation*. p. 625–640. PLDI 2022, Association for Computing Machinery, New York, NY, USA (2022). <https://doi.org/10.1145/3519939.3523433>, <https://doi.org/10.1145/3519939.3523433>

A Missing Proofs

Theorem 1 (QCEL Soundness). For any equational theory Γ over a quantum gate set $\mathcal{G}$ and any n -qubit quantum circuit C , if $\Gamma \vdash C \approx id_n$ and $\models \Gamma$, then $\models C \approx id_n$.

Proof (Proof of Theorem 1). The proof proceeds by induction on the derivation of $\Gamma \vdash C \approx id_n$. We show that each inference rule in Figure 2 preserves validity.

- EQ: Trivial because $\llbracket id_n \rrbracket = I^{\otimes n}$.
- ASSUME: If $C \approx id_n \in \Gamma$, since $\models \Gamma$, we have $\models C \approx id_n$ by definition.
- ADJOINT: Assume $\models C \approx id_n$, i.e., $\llbracket C \rrbracket = e^{i\alpha}I$. Then $\llbracket C^\dagger \rrbracket = \llbracket C \rrbracket^\dagger = (e^{i\alpha}I)^\dagger = e^{-i\alpha}I$, so $\models C^\dagger \approx id_n$.
- EXTEND: Assume $\models C \otimes id_m \approx id_{n+m}$. Then $\llbracket C \rrbracket \otimes I^{\otimes m} = e^{i\alpha}I^{\otimes(n+m)} = (e^{i\alpha}I^{\otimes n}) \otimes I^{\otimes m}$. This implies $\llbracket C \rrbracket = e^{i\alpha}I^{\otimes n}$. Consequently, $\llbracket C \otimes id_k \rrbracket = \llbracket C \rrbracket \otimes I^{\otimes k} = e^{i\alpha}I^{\otimes(n+k)}$, so the conclusion holds.
- ROTATE: Assume $\models C_1 \circ C_2 \approx id_n$, so $\llbracket C_1 \rrbracket \llbracket C_2 \rrbracket = e^{i\alpha}I$. Multiplying by $\llbracket C_1 \rrbracket^\dagger$ on the left gives $\llbracket C_2 \rrbracket = e^{i\alpha} \llbracket C_1 \rrbracket^\dagger$. Then $\llbracket C_2 \rrbracket \llbracket C_1 \rrbracket = e^{i\alpha} \llbracket C_1 \rrbracket^\dagger \llbracket C_1 \rrbracket = e^{i\alpha}I$, so $\models C_2 \circ C_1 \approx id_n$.
- PERMUTATION: Assume $\models C \approx id_n$. Then $\llbracket \sigma^c \circ C \circ (\sigma^c)^\dagger \rrbracket = \llbracket \sigma^c \rrbracket \llbracket C \rrbracket \llbracket \sigma^c \rrbracket^\dagger = \llbracket \sigma^c \rrbracket (e^{i\alpha}I) \llbracket \sigma^c \rrbracket^\dagger = e^{i\alpha} \llbracket \sigma^c \rrbracket \llbracket \sigma^c \rrbracket^\dagger = e^{i\alpha}I$.
- TRANSITIVITY: Assume $\models C_1 \circ C_2^\dagger \approx id_n$ and $\models C_2 \circ C_3^\dagger \approx id_n$. Then $\llbracket C_1 \rrbracket \llbracket C_2 \rrbracket^\dagger = e^{i\alpha}I$ and $\llbracket C_2 \rrbracket \llbracket C_3 \rrbracket^\dagger = e^{i\beta}I$. Thus $\llbracket C_1 \rrbracket = e^{i\alpha} \llbracket C_2 \rrbracket$ and $\llbracket C_2 \rrbracket = e^{i\beta} \llbracket C_3 \rrbracket$. Substituting gives $\llbracket C_1 \rrbracket = e^{i(\alpha+\beta)} \llbracket C_3 \rrbracket$, or $\llbracket C_1 \rrbracket \llbracket C_3 \rrbracket^\dagger = e^{i(\alpha+\beta)}I$, which means $\models C_1 \approx C_3$. $\square$

Proof (Proof of Lemma 2). ($\supseteq$) Trivial based on definition. ($\subseteq$) Consider any identity $C' \approx id_n$ in $\text{Id}_n^{2k}(\mathcal{G})$. By definition, $\llbracket C' \rrbracket = id_n$ and the size of C' is at most $2k$. We can decompose C' into two circuits C'_1 and C'_2 such that $C' = C'_1 \circ C'_2$ and both C'_1 and C'_2 have size at most k . Let $U = \text{REDUCE}(\llbracket C'_1 \rrbracket) = \text{REDUCE}(\llbracket C'_2 \rrbracket)$. By Lemma 1, there exist circuits $C_1 \in \mathcal{R}_n^{\leq k}[\mathcal{G}](U)$ and symmetries $\sigma_1 \sigma_2 \in \text{Sym}(n)$ such that $C'_1 = \sigma_1^c \circ C_1 \circ \sigma_2^c$. One can also check that C_2 defined by $C'_2 = \sigma_1^c \circ C_2 \circ \sigma_2^c$, also satisfying $C_2 \in \mathcal{R}_n^{\leq k}[\mathcal{G}](U)$ since $\llbracket C'_1 \rrbracket = \llbracket C'_2 \rrbracket$. Therefore, using C_1 and C_2 , we have $C' = \sigma_1^c \circ C_1 \circ \sigma_2^c \circ (\sigma_2^c)^\dagger \circ C_2^\dagger \circ (\sigma_1^c)^\dagger = \sigma_1^c \circ C_1 \circ C_2^\dagger \circ (\sigma_1^{-1})^c$ using (2.2). Thus, $C' \approx id_n$ can be expressed in the desired form. $\square$

Theorem 2. Algorithm 1 correctly computes the mapping $\mathcal{R}_n^{\leq k}[\mathcal{G}]$ for the given quantum gate set $\mathcal{G}$, circuit size k , and number of qubits n .

Proof (Proof of Theorem 2). We prove the correctness by induction on the circuit size i . For the base case $i = 1$, the algorithm initializes R^1 to contain all single-gate circuits from $\mathcal{G}$, which is correct by definition. For the inductive step, assume that R^{i-1} correctly computes $\mathcal{R}_n^{i-1}[\mathcal{G}]$. We need to show that R^i correctly computes $\mathcal{R}_n^i[\mathcal{G}]$.

($\subseteq$) Consider any unitary $U' \in \text{U}(2^n)$ such that $R^i(U') \neq \emptyset$. By the construction of R^i , there exists a gate $G \in \text{Circ}_n^1(\mathcal{G})$ and a unitary $U \in \text{U}(2^n)$ such that $U' = \text{REDUCE}(\llbracket G \rrbracket U)$ and $R^{i-1}(U) \neq \emptyset$. By the inductive hypothesis, $R^{i-1}(U)$ contains all circuits of size $i-1$ that evaluate to the canonical representative of U . Therefore, for

each circuit $C \in R^{i-1}(U)$, the circuit $\sigma_2^c \circ G \circ C \circ \sigma_1^c$ evaluates to U' and has size i . Thus, all circuits in $R^i(U')$ are of size i and evaluate to the canonical representative of U' .

($\supseteq$) Consider any circuit $C' \in \text{Circ}_n^i(\mathcal{G})$ such that $\llbracket C' \rrbracket = \text{REDUCE}(\llbracket C' \rrbracket) = U'$. By the definition of circuit construction, there exists a gate $G' \in \text{Gen}_n(\mathcal{G})$ and a circuit C'' such that $C' = G' \circ C''$. Since $C'' \in \text{Circ}_n^{i-1}(\mathcal{G})$, by Lemma 1, there exists a circuit C and symmetries σ_1, σ_2 such that $C'' = \sigma_2^c \circ C \circ \sigma_1^c$ and $C \in \mathcal{R}_n^{i-1}[\mathcal{G}](U)$ where $U = \text{REDUCE}(C'')$. Let $G = (\sigma_2^{-1})^c \circ G' \circ \sigma_2^c$, by (2.3), $G \in \text{Gen}_n(\mathcal{G})$. Then we have $C' = \sigma_2^c \circ G \circ C \circ \sigma_1^c$. By the inductive hypothesis, $C \in R^{i-1}(U)$. Therefore, during the construction of R^i , the circuit C' will be added to $R^i(U')$ where $U' = \text{REDUCE}(\llbracket G \rrbracket U)$. Thus, all circuits of size i that evaluate to the canonical representative of U' are included in $R^i(U')$.

By induction, the algorithm correctly computes $\mathcal{R}_n^{\leq k}[\mathcal{G}]$. $\square$

Theorem 3. Given a unitary U on n qubits, and a fully randomized quantum state $|\psi_r\rangle$ independent to U , $\Pr(|\text{Sym}^\approx(\text{dephase } U|\psi_r)\rangle| > 1) = 0$.

Proof (Proof of Theorem 3). Given a Haar-random quantum state $|\psi_r\rangle \sim \text{Haar}(\mathbb{C}_2^{\otimes n})$. Let $|\psi\rangle = U|\psi_r\rangle$, obviously, $|\psi\rangle$ is also Haar-random. Since $|\text{Sym}^\approx(\text{dephase } U|\psi_r)\rangle| > 1$ implies that there exist two different qubits $1 \leq i < j \leq n$ such that one of $\langle \psi | 2^i \rangle = \langle \psi | 2^j \rangle$, and $\langle \psi | 2^n - 1 - 2^i \rangle = \langle \psi | 2^n - 1 - 2^j \rangle$ holds. Note that the dephase operation does not affect these equalities. Since $|\psi\rangle$ is Haar-random, the probability of this event happening is zero. Therefore, we have $\Pr(|\text{Sym}^\approx(\text{dephase } U|\psi_r)\rangle| > 1) = 0$. $\square$

Theorem 4. Given two unitaries U_1 and U_2 , and a fully randomized quantum state $|\psi_r\rangle$ independent to both U_1 and U_2 , if $\text{dephase } U_1|\psi_r\rangle = \text{dephase } U_2|\psi_r\rangle$ holds, then the posterior probability $\Pr(\exists \alpha. U_1 = e^{i\alpha} U_2) = 1$.

Proof (Proof of Theorem 4). Let events $A = \{\text{dephase } U_1|\psi_r\rangle = \text{dephase } U_2|\psi_r\rangle\}$, $B = \{\exists \alpha. U_1 = e^{i\alpha} U_2\}$, and $V_1, \dots, V_m$ are eigenspaces of $U_2^\dagger U_1$. We have,

$$\begin{aligned} \Pr(A|\neg B) &= \Pr(\exists \alpha. U_1|\psi_r\rangle = e^{i\alpha} U_2|\psi_r\rangle | \neg B) = \Pr(\exists \alpha. U_2^\dagger U_1|\psi_r\rangle = e^{i\alpha} |\psi_r\rangle | \neg B) \\ &= \sum_{k=1}^m \Pr(|\psi_r\rangle \in V_k | \neg B) = 0 \end{aligned}$$

The last equality holds because if $\neg B$ holds, then $U_2^\dagger U_1 - e^{i\alpha}$ is not a zero operator for any α , thus each eigenspace V_k has a dimension less than 2^n . Thus, by Bayes' theorem, we have $\Pr(B|A) = 1$. $\square$

Theorem 5. Algorithm 2 returns $U_{\min}$ satisfying (3.1) under the assertion and assumption.

Proof (Proof of Theorem 5). ($\Rightarrow$) Trivial due to line 9. ($\Leftarrow$) For any $\sigma_1, \sigma_2 \in \text{Sym}(n)$, let $U' = \llbracket \sigma_2^c \rrbracket U \llbracket \sigma_1^c \rrbracket$, we compare each line of Algorithm 2 to show that $\text{REDUCE}(U') = \text{REDUCE}(U)$. Variables in the algorithm of $\text{REDUCE}(U')$ are denoted with $\cdot'$, to distinguish from those in $\text{REDUCE}(U)$.

- At line 3, we can see $|\psi'_0\rangle = \text{dephase } \llbracket \sigma_1^c \rrbracket^\dagger U^\dagger \llbracket \sigma_2^c \rrbracket^\dagger |\psi_l\rangle = \llbracket \sigma_1^c \rrbracket^\dagger \text{dephase } U^\dagger |\psi_l\rangle = \llbracket \sigma_1^c \rrbracket^\dagger |\psi_0\rangle$, based on the definition of dephase and ψ_l .

- At line 4, from the definition of $\text{Sym}^\sim$, we can see that $\text{Sym}^\sim(\llbracket \sigma_1^c \rrbracket^\dagger \psi_0) = \text{Sym}^\sim(\psi_0) \cdot \sigma_1^c$. Therefore, for each $\sigma'_r \in \text{Sym}^\sim(\psi'_0)$, there exists a unique $\sigma_r \in \text{Sym}^\sim(\psi_0)$ such that $\sigma'_r = \sigma_r \cdot \sigma_1$, then we compare every two corresponding iterations based on σ_r and σ'_r .
- At line 5, we have $|\psi'\rangle = \text{dephase} \llbracket \sigma_2^c \rrbracket U \llbracket \sigma_1^c \rrbracket \llbracket \sigma_1^c \rrbracket^\dagger \llbracket \sigma_r^c \rrbracket^\dagger |\psi_r\rangle = \llbracket \sigma_2^c \rrbracket |\psi\rangle$.
- At line 7, using the definition of $\text{Sym}^\sim$ again, we have $\sigma'_l = \sigma_l \sigma_2^{-1}$.
- Finally, the algorithm yields the same set of $\llbracket \sigma_l^c \rrbracket \psi$ and $\llbracket \sigma_l^c \rrbracket U \llbracket \sigma_r^c \rrbracket$. Using the assumption at line 11, the same minimum unitary $U'_{\min} = U_{\min}$ is returned. $\square$

Theorem 6. If the assertion and assumption hold, Algorithm 2 returns

$$\Sigma = \{(\sigma_l, \sigma_r) \mid \exists \alpha, \sigma_l U \sigma_r = e^{i\alpha} U_{\min}\}.$$

Proof (Proof of Theorem 6). ($\subseteq$) Trivial due to assumption at line 11. ($\supseteq$) For those symmetries $\sigma_r \notin \text{Sym}^\sim(\psi_0)$ filtered out at line 4, since $\text{dephase } U_{\min}^\dagger |\psi_l\rangle$ is seemingly ordered while $\text{dephase} \llbracket \sigma_r^c \rrbracket U^\dagger \llbracket \sigma_l^c \rrbracket^\dagger |\psi_l\rangle$ is not, the corresponding unitary $\llbracket \sigma_l^c \rrbracket U \llbracket \sigma_r^c \rrbracket^\dagger$ are guaranteed to be different from $U_{\min}$ up to a certain phase factor. For any symmetries $\sigma_l \in \text{Sym}(n)$, $\sigma_r \in \text{Sym}^\sim(\psi)$ that doesn't satisfy the condition at line 10, because $\text{dephase } U_{\min} |\psi_r\rangle \neq \text{dephase} \llbracket \sigma_l^c \rrbracket U \llbracket \sigma_r^c \rrbracket^\dagger |\psi_r\rangle$, $U_{\min}$ and $\llbracket \sigma_l^c \rrbracket U \llbracket \sigma_r^c \rrbracket^\dagger$ are also different up to a certain phase factor. $\square$

Theorem 7. Algorithm 3 correctly computes an equational theory that can derive all identities in $\text{Id}_n^{\leq 2k}(\mathcal{G})$.

Proof (Proof of Theorem 7). First we prove the invariant at line 9, i.e., all identities in *proved* can be derived from the identities in *assumed*. This is true at the beginning since both sets are empty. Now consider any identity I added to *proved* during the iteration. In TASOSEARCH, every I_1 added into *visited* set can be derived from and to I using the proved identities in *proved*. Therefore, 1) if one of the I_1 is already in *proved*, the entire *visited* can also be derived from the identities in *assumed*. 2) If the search exceeds the limit, adding I to *assumed* also leads to *assumed* $\vdash$ *visited*. 3) If I_1 is smaller than I , the *visited* set inside $\text{TASOSEARCH}(I_1)$ procedure are also equivalent to proving I . Therefore, *assumed* $\vdash$ *visited* and the invariant holds.


Since $\mathcal{I} \subseteq \text{proved}$, and all identities in $\text{Id}_n^{\leq 2k}(\mathcal{G})$ can be derived using identities in $\mathcal{I}$ using PERMUTATION and ROTATE rules due to Lemma 2, all identities in $\text{Id}_n^{\leq 2k}(\mathcal{G})$ can also be derived from the identities in *assumed*. $\square$

B Example of an Exported QCEL Proof


For simplicity, we only show TRANSITIVITY rule application in this example. We leave the application of other rules implicit to the readers.

Rule	Canonical Identity
$0 \leftarrow 1, 26$	
$1 \leftarrow 2, 23$	
$2 \leftarrow 3, 25$	
$3 \leftarrow 4, 25$	
$4 \leftarrow 5, 19$	


$5 \leftarrow 6, 11$




$6 \leftarrow 7, 18$




$7 \leftarrow 8, 15$




$8 \leftarrow 9, 14$




$9 \leftarrow 10, 13$




10 $\leftarrow$ 11, 12




11 $\leftarrow$ ASSUME




12 $\leftarrow$ ASSUME




13 $\leftarrow$ ASSUME









14 $\leftarrow$ 13, 12




15 $\leftarrow$ 16, 11




16 $\leftarrow$ 10, 1717 $\leftarrow$ ASSUME18 $\leftarrow$ ASSUME19 $\leftarrow$ 20, 1220 $\leftarrow$ 21, 1221 $\leftarrow$ 22, 12


$22 \leftarrow 8, 23$




$23 \leftarrow 24, 12$




$24 \leftarrow 11, 12$



$25 \leftarrow 16, 12$



$26 \leftarrow 27, 12$



$27 \leftarrow 24, 17$

