

AUTOQ 2.0: From Verification of Quantum Circuits to Verification of Quantum Programs

Yu-Fang Chen¹ , Kai-Min Chung¹ , Min-Hsiu Hsieh² , Wei-Jia Huang² ,
Ondřej Lengál³ , Jyun-Ao Lin⁴ , and Wei-Lun Tsai¹ 

¹ Institute of Information Science, Academia Sinica, Taipei, Taiwan

² Hon Hai Quantum Computing Research Center, Taipei, Taiwan

³ Faculty of Information Technology, Brno University of Technology, Brno, Czech Republic

⁴ Innovation Frontier Institute of Research for Science and Technology & Department of
Computer Science and Information Engineering,
National Taipei University of Technology, Taipei, Taiwan

Abstract. We present a verifier of quantum programs called AUTOQ 2.0. Quantum programs extend quantum circuits (the domain of AUTOQ 1.0) by classical control flow constructs, which enable users to describe advanced quantum algorithms in a formal and precise manner. The extension is highly non-trivial, as we needed to tackle both theoretical challenges (such as the treatment of measurement, the normalization problem, and lifting techniques for verification of classical programs with loops to the quantum world), and engineering issues (such as extending the input format with a support for specifying loop invariants). We have successfully used AUTOQ 2.0 to verify two types of advanced quantum programs that cannot be expressed using only quantum circuits: the *repeat-until-success* (RUS) algorithm and the weak-measurement-based version of Grover’s search algorithm. AUTOQ 2.0 can efficiently verify all our benchmarks: all RUS algorithms were verified instantly and, for the weak-measurement-based version of Grover’s search, we were able to handle the case of 100 qubits in ~20 minutes.

1 Introduction

Quantum *programs* are an extension of quantum *circuits* that provide users with greater control over quantum computing by allowing them to use more complex programming constructs like branches and loops. Some of the most advanced quantum algorithms cannot be defined by quantum circuits alone. For example, certain class of programs, such as the *repeat-until-success* (RUS) algorithms [41] (which are commonly used in generating special quantum gates) and the weak-measurement-based version [6] of Grover’s search algorithm [31], use a loop with the condition being a classical value (0 or 1) obtained by measuring a particular qubit. This added expressivity presents new challenges, particularly in terms of verification. The additional complexity comes from the *measurement* operation, where a particular qubit is measured to obtain a classical value (and the quantum state is partially collapsed, which might require *normalization*), and reasoning about control flow induced by *branches* and *loops*.

In classical program verification, a prominent role is played by *deductive verification* [30,34,32], represented, e.g., by the tools Dafny [37], KeY [4], Frama-C [8],

VeriFast [36], VCC [23], and many more. These tools only require the users to provide specifications in the form of pre- and post-conditions, along with appropriate loop invariants. The rest of the proving process is entirely (in the ideal case) automated. Unfortunately, in the realm of quantum computing, similar fully automated deductive verification tools are, to the best of our knowledge, missing. Advanced tools for analysis and verification of quantum programs—based on, e.g., *quantum Hoare logic* and the tool CoqQ [50] or the *path-sum* formalism [5] and the tool QBRICKS [14]—are quite powerful but require a significant amount of human effort.

To bridge this gap, we present AutoQ 2.0, a major update over AutoQ 1.0 [19] with an added support for *quantum programs* (AutoQ 1.0 only supported quantum *circuits*). In AutoQ 1.0, given a triple $\{P\} C \{Q\}$, where P and Q are the pre- and post-conditions recognizing sets of (pure) quantum states (represented by *tree automata*) and C is a quantum circuit, we can verify if all quantum states in P reach some state in Q after executing C . In AutoQ 2.0, we addressed several key challenges to make the support of quantum programs possible. First, we need to handle *branch* statements. The key issue here is to handle *measurement* of quantum states whose value is used in a branch condition. For this we developed automata-based algorithms to compute the quantum states after the measurement (Section 5). The second challenge is the handling of *loop* statements. Similarly to deductive verification of classical programs, we require the users to provide *an invariant* for each loop. With the loop invariant provided, we developed a framework handling the rest of the verification process fully automatically. Moreover, we show that a naive implementation of the measurement operation will encounter the *probability amplitude normalization* problem. This is handled by designing a new algorithm for *entailment testing* (Section 6).

Under this framework, the preconditions, postconditions, and invariants are all described with a new automata model called *level-synchronized tree automata (LSTAs)* [1]. LSTAs are specifically designed to efficiently encode quantum states and gate operations. As the core data structure of the tool, we provide a formal definition of LSTAs in Section 2.2 to facilitate the presentation of our new entailment testing approach.

We used AutoQ 2.0 to verify various quantum programs using the *repeat-until-success* (RUS) paradigm [41], as well as the weak-measurement-based version [6] of Grover’s search [31] (Section 7). AutoQ 2.0 can efficiently verify all our benchmarks. The verification process for all RUS algorithms was instantaneous and for the weakly measured versions of Grover, we were able to handle the case of 100 qubits in ~ 20 min. To the best of our knowledge, AutoQ 2.0 is currently the only tool for verification of quantum programs with such a degree of automation.

Related work. Our work aligns with Hoare-style verification of quantum programs, a topic extensively explored in prior studies [51,43,48,27,39]. This approach, inspired by D’Hondt and Panangaden, utilizes diverse Hermitian operators as quantum predicates, resulting in a robust and comprehensive proof system [25]. However, specifying properties with Hermitian operators is often non-intuitive and difficult for automation due to their vast matrix sizes. Consequently, these methods are typically implemented using proof assistants like Coq [9], ISABELLE [45], or standalone tools built on top of Coq,

like CoqQ [50]. These tools require substantial manual effort in the proof search. The QBRICKS approach [15] addresses the challenge of proof search by combining cutting-edge theorem provers with decision procedures, leveraging the Why3 platform [29]. Nevertheless, this approach still demands considerable human intervention.

In the realm of automatic quantum software analysis tools, *circuit equivalence checkers* [5,22,33,47,23] prove to be efficient but less flexible in specifying desired properties, primarily focusing on equivalence. These tools are valuable in compiler validation, with notable examples being QCEC [13], FEYNMAN [5], and SLIQEC [17,44]. *Quantum model checking*, supporting a rich specification language (various temporal logics [28,40,46]), is, due to its limited scalability, more suited for verifying high-level protocols [7]. QPMC [28] stands out as a notable tool in this category. Quantum abstract interpretation [49,42] over-approximates the reachable state space to achieve better scalability, but so far handles only circuits. The work in [52,26] aims at the verification of *parameterized quantum programs* like *variational quantum eigensolver (VQE)* or *quantum approximate optimization algorithm (QAOA)*. However, the correctness properties they focused are very different from what AutoQ 2.0 can handle. While the mentioned tools are fully automated, they serve different purposes or address different phases of the development cycle compared to AutoQ 2.0.

2 Background

Before we start, we first provide a minimal background needed for this work.

2.1 The Tree-View of Quantum States

In a traditional computer system with n bits, a state is represented by n Boolean values 0 or 1. An n -qubit *quantum state* can be viewed as a “probabilistic distribution” over n -bit classical states. Here we often refer to each classical state as a *computational basis state* or *basis state* for short. Hence a quantum state can be represented by a *binary tree* whose branches correspond to the computational basis states and leaves correspond to *complex probability amplitudes*⁵.

In Fig. 1(a), we can see an example of a 2-qubit state q that maps basis states $|00\rangle$, $|01\rangle$, $|10\rangle$, $|11\rangle$ to probability amplitudes a_1 , a_2 , a_3 , and a_4 , respectively. The left-going dashed line denotes the 0-branch, and the right-going solid line denotes the 1-branch. A quantum state can also be represented as a formal sum of basis states multiplied by their amplitudes, e.g., we can represent the state q as $a_1 |00\rangle + a_2 |01\rangle + a_3 |10\rangle + a_4 |11\rangle$.

Quantum gates are fundamental operations performed on quantum states. Basic quantum gates and their effects on the state q from Fig. 1(a) are shown in Figs. 1(b) to 1(d). To specify the target qubit to which a single qubit gate U is applied, a subscript number i is used. For example, X_i denotes the application of the X gate, which is also known as the quantum “negation” gate, to the i -th qubit. The effect of this gate is to swap the 0- and 1-subtrees under all x_i nodes (cf. Fig. 1(b)). On the other hand, for a controlled gate, a superscript number i is used to indicate the control qubit, while a subscript number j is used for the target qubit. The most notable example is the controlled- X gate CX_j^i , which

⁵ A state with complex amplitude $a + bi$ has the probability $|a + bi|^2 = a^2 + b^2$ of being observed.




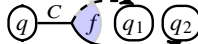
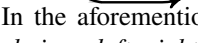
Fig. 1. The effect of applying selected quantum gates to state q

applies the X_j gate to 1-subtrees under (assuming $i < j$) all x_i nodes (cf. Fig. 1(c)). Observe that after applying an H gate (the *Hadamard* gate, which creates a quantum superposition; cf. Fig. 1(d)), there is a *normalization factor* $\frac{1}{\sqrt{2}}$ on all leaves to keep the sum of probabilities one. This normalization factor can be derived from the tree leaf values $(a_1 + a_3)$, $(a_2 + a_4)$, $(a_1 - a_3)$, and $(a_2 - a_4)$ and the fact that $\sum_{i=1}^4 |a_i|^2 = 1$.

2.2 Level-Synchronized Tree Automata

As we mentioned in Section 1, the new algorithms introduced in this work are built on top of LSTAs, making it essential to provide a formal definition. Readers may choose to skim this section initially and refer back to it for details as needed later.

Trees. Our framework is based on the concept of perfect binary trees. A *perfect binary tree* T is a map from *tree nodes* $\bigcup_{0 \leq i \leq n} \{0, 1\}^i$, for some $n \in \mathbb{N}_0 := \mathbb{N} \cup \{0\}$, to a nonempty set of symbols Σ , i.e., $T: \bigcup_{0 \leq i \leq n} \{0, 1\}^i \rightarrow \Sigma$. All nodes $v \in \bigcup_{0 \leq i < n} \{0, 1\}^i$ are *internal* and have children nodes $v.0$ (left) and $v.1$ (right) where ‘.’ denotes concatenation (we denote the empty string by ϵ). All nodes $v \in \{0, 1\}^n$ are *leaves* and have no children. A node v ’s *height* is its word length, denoted $|v|$. A node v is at tree level i when $|v| = i$. We denote T ’s height by n . Perfect binary trees can be used to represent quantum states or vectors of the size 2^n . For instance, the quantum state of Fig. 1(a) corresponds to a perfect binary tree $T = \{\epsilon \mapsto x_1, 0 \mapsto x_2, 1 \mapsto x_2, 00 \mapsto a_1, 01 \mapsto a_2, 10 \mapsto a_3, 11 \mapsto a_4\}$ of height 2. Children of the node 1 are 10 and 11, and the leaf node 10 has no children.

LSTAs. A (*symbolic*) *level-synchronized tree automaton (LSTA)* [1] is a tuple $\mathcal{A} = \langle Q, \mathbb{N} \cup \text{term}(\mathbb{C}, \mathbb{X}), \Delta, \mathcal{R}, \varphi \rangle$ where Q is a finite set of *states*, $\mathcal{R} \subseteq Q$ is a set of *root states*, $\text{term}(\mathbb{C}, \mathbb{X})$ is a set of terms obtained from complex numbers \mathbb{C} and a set of *complex variables* \mathbb{X} using function symbols from some fixed theory (in this paper, we will use \mathbb{N} for internal node symbols and $\text{term}(\mathbb{C}, \mathbb{X})$ for leaf symbols). Δ is a set of *transitions* of the form $\delta_i = q \xrightarrow{f|C} (q_1, q_2)$ (*internal transitions*) and $\delta_\ell = q \xrightarrow{f|C} ()$ (*leaf transitions*), where $q, q_1, q_2 \in Q$, $f \in \text{term}(\mathbb{C}, \mathbb{X})$, and $C \subseteq \mathbb{N}$ is a finite set of *choices*. In the rest of the paper, we also draw the internal transition δ_i and leaf transition δ_ℓ as  and , respectively, to provide a more intuitive presentation. In the aforementioned form, we call q , f , C , q_1 , q_2 , and $\{q_1, q_2\}$ the *top*, *symbol*, *choices*, *left*, *right*, and *bottom*, respectively, of the transition δ_i , and denote them by $\text{top}(\delta_i)$, $\text{sym}(\delta_i)$, $\text{ch}(\delta_i)$, $\text{left}(\delta_i)$, $\text{right}(\delta_i)$, and $\text{bot}(\delta_i)$, respectively. Needless to




Fig. 2. Intermediate states during the execution of Algorithm 1

say, $\text{bot}(\delta_\ell) = \emptyset$. We further require the choices of different transitions with the same top state to be disjoint, i.e., $\forall \delta_1 \neq \delta_2 \in \Delta: \text{top}(\delta_1) = \text{top}(\delta_2) \implies \text{ch}(\delta_1) \cap \text{ch}(\delta_2) = \emptyset$. Finally, the *global constraint* φ is a formula used to restrict the values of \mathbb{X} to those that satisfy φ (if not stated, it is assumed to be “true”). For instance, when $\mathbb{X} = \{a, b\}$, we can set $\varphi = |a| > |b|$ to restrict the allowed valuation of a and b .

Semantics of LSTAs. A *run* of an LSTA \mathcal{A} on a tree T is a map $\rho: \text{dom}(T) \rightarrow \Delta$ from tree nodes to transitions of \mathcal{A} such that for each node $v \in \text{dom}(T)$, when v is an internal node, $\rho(v)$ is of the form $q \xrightarrow{T(v)|C} (\text{top}(\rho(v.0)), \text{top}(\rho(v.1)))$. When v is a leaf node, $\rho(v)$ is of the form $q \xrightarrow{T(v)|C} ()$.

We give a *run* ρ of the LSTA \mathcal{A} in Fig. 3 on the tree T in Fig. 2(d) (used later in Algorithm 1) as follows:

$$\begin{aligned}
 \rho(\epsilon) &= (p \xrightarrow{\{1\}} x_1 \xrightarrow{\{1\}} q_0 \xrightarrow{\{1\}} q_-), & \rho(0) &= (q_0 \xrightarrow{\{1\}} x_2 \xrightarrow{\{1\}} r_0 \xrightarrow{\{1\}} r_0), \\
 \rho(1) &= (q_- \xrightarrow{\{1\}} x_2 \xrightarrow{\{1\}} r_3 \xrightarrow{\{1\}} r_4), & \rho(00) = \rho(01) &= (r_0 \xrightarrow{\{1\}} 0), \\
 \rho(10) &= (r_3 \xrightarrow{\{1\}} -\frac{a_1}{\sqrt{2}}), & \rho(11) &= (r_4 \xrightarrow{\{1\}} -\frac{a_0}{\sqrt{2}}).
 \end{aligned}$$

A run ρ is *accepting* if $\text{top}(\rho(\epsilon)) \in \mathcal{R}$ and all transitions used at the same level share some common choice, i.e., $\forall 0 \leq i \leq n: \bigcap \{\text{ch}(\delta) \mid \delta \in \{\rho(v) \mid v \in \{0, 1\}^i\}\} \neq \emptyset$ (this is the essential difference from standard tree automata that gives LSTAs the power to compactly represent some classes of quantum states). The *language* of \mathcal{A} , denoted as $\mathcal{L}(\mathcal{A})$ is a set of trees T over $\mathbb{N} \cup \text{term}(\mathbb{C}, \mathbb{X})$ such that there exists an accepting run of \mathcal{A} over T . Given a tree T over $\mathbb{N} \cup \text{term}(\mathbb{C}, \mathbb{X})$ and an assignment $\sigma: \mathbb{X} \rightarrow \mathbb{C}$, we use $T[\sigma]$ to denote the tree obtained from T by (i) substituting all occurrence of variables $x \in \mathbb{X}$ in T by $\sigma(x)$ and (ii) evaluating all terms in the resulting tree into values $c \in \mathbb{C}$.

3 Overview

In this section, we provide an overview of automata-based quantum program verification with a running example (chosen for its simplicity). In the example, the quantum program creates the effect of a non-standard quantum gate “ $-X$ ” (applying the X gate and negating all amplitude values) using the standard gates X , H , and CX (Algorithm 1). The program operates on a 2-qubit system and performs the “ $-X$ ” gate on the second qubit when the

first qubit is measured to be 1; otherwise, (the first qubit is 0 after measurement) the state stays unchanged.

For all $a_0, a_1 \in \mathbb{C}$, we verify Algorithm 1 against the precondition $\{a_0 |10\rangle + a_1 |11\rangle\}$, which allows the state with the first qubit $|1\rangle$ and the second qubit $a_0 |0\rangle + a_1 |1\rangle$, and the postcondition $\{a_0 |10\rangle + a_1 |11\rangle, -a_1 |10\rangle - a_0 |11\rangle\}$, which includes the original state and the state after the “ $-X_2$ ” gate.

Algorithm 1: “ $-X_2$ ” if $M_1 = 1$

```

1 Pre:  $\{a_0 |10\rangle + a_1 |11\rangle\}$ ;
2  $H_1; CX_2^1$ ;
3 if  $M_1 = 0$  then  $\{X_1\}$ ;
4 Post:  $\{a_0 |10\rangle + a_1 |11\rangle,$ 
5    $-a_1 |10\rangle - a_0 |11\rangle\}$ ;

```

Our approach first constructs two LSTAs P and Q that can recognize the states (binary trees) of the pre- and post-conditions, respectively, and then computes another LSTA by executing the gates $H_1; CX_2^1$ from P (see Fig. 2(a) for the only quantum state accepted by P) with the gate algorithm introduced in [1]. This results in an LSTA P_1 that recognizes the state shown in Fig. 2(b). After applying the operator M_1 to measure x_1 (Line 3), P_1 splits into two copies. One copy, P_2 , accepts the only quantum state shown in Fig. 2(c), where the first qubit is measured to be 0. The other copy, P_3 , accepts the only quantum state shown in Fig. 2(d), where the first qubit is measured to be 1.

It is important to note that the probability amplitudes of the quantum states from Figs. 2(c) and 2(d) have not been normalized yet. To do that, we need to multiply all leaves with the normalization factor $\sqrt{2}$. This will ensure that the square sum of their absolute amplitude values becomes 1. Although the quantum states are not yet normalized, we, however, still have sufficient information to obtain the corresponding normalized states. In AUTOQ 2.0, we choose to ignore all normalization factors and design a new entailment testing algorithm (Section 6) that can detect the equivalence of two non-normalized states. After both the true and false paths of the **if** statement in the example are processed, we obtain two LSTAs P_2 and P_3 capturing all reachable states. We then construct their union and test if all states in the union are included in the post-condition (recognized by Q) by testing entailment.

A drawback of Algorithm 1 is that the desired effect “ $-X$ ” manifests only if $M_1 = 1$. In the case $M_1 = 0$, we need to run the same algorithm again until we get a measurement of 1. To achieve this, we can use a **while** loop statement, as shown in Algorithm 2. The loop allows us to repeatedly execute the same branch statement until the desired outcome is achieved.

Algorithm 2: “ $-X_2$ ”

```

1 Pre:  $\{a_0 |10\rangle + a_1 |11\rangle\}$ ;
2  $H_1; CX_2^1$ ;
3 Inv:
    $\{\frac{a_0}{\sqrt{2}} |00\rangle + \frac{a_1}{\sqrt{2}} |01\rangle - \frac{a_1}{\sqrt{2}} |10\rangle - \frac{a_0}{\sqrt{2}} |11\rangle\}$ ;
4 while  $M_1 = 0$  do  $\{X_1; H_1; CX_2^1\}$ ;
5 Post:  $\{-a_1 |10\rangle - a_0 |11\rangle\}$ ;

```

To verify that the loop works correctly, we require the user to provide a loop invariant in the form of an LSTA. The invariant here is $\{\frac{a_0}{\sqrt{2}} |00\rangle + \frac{a_1}{\sqrt{2}} |01\rangle - \frac{a_1}{\sqrt{2}} |10\rangle - \frac{a_0}{\sqrt{2}} |11\rangle\}$ (cf. Fig. 2(b)). The verification process then involves checking if the invariant is *inductive*, covers all reachable states before entering the loop, and does not contain any state that would violate the post-condition. More details on the verification process will be given in Section 5.3. With the loop invariant provided, we can ensure that the algorithm ends up with a state where the “ $-X$ ” gate is applied to the second qubit when it terminates.




Fig. 4. The LSTA for Bell states and its textual description

In AUTOQ 2.0, preconditions, postconditions, and invariants are represented as sets of quantum states, encoded using the LSTA model. Therefore, it is important for users to understand how to encode a set of quantum states with an LSTA. Below, we provide two examples to give a basic understanding of the process. In the first example, we show how to encode the postcondition of Algorithm 1, $\{a_0 |10\rangle + a_1 |11\rangle, -a_1 |10\rangle - a_0 |11\rangle\}$.

The corresponding LSTA is shown in Fig. 3. The LSTA constructs trees that depict quantum states beginning from the initial state p at the root. It continues to build the tree by choosing transitions to explore new child states at each step, and this process continues until it reaches the leaves. For instance, the tree in Fig. 2(c) can be generated by first picking the transition $p \xrightarrow{\{1\}} x_1 \rightarrow q_+ \rightarrow q_0$, then the two transitions $q_+ \xrightarrow{\{1\}} x_2 \rightarrow r_1$ and $q_0 \xrightarrow{\{1\}} x_2 \rightarrow r_0$, and ending with the three leaf




Fig. 3. The LSTA recognizing the postcondition of Algorithm 1

transitions $r_1 \xrightarrow{\{1\}} \frac{a_0}{\sqrt{2}}, r_2 \xrightarrow{\{1\}} \frac{a_1}{\sqrt{2}},$ and $r_0 \xrightarrow{\{1\}} 0$. Similar to the conventional tree automata (TAs) model, LSTAs utilize disjunctive branches to represent various states that share a common structure. In Fig. 3, the state p has two possible outgoing transitions. If one picks the other transition $p \xrightarrow{\{2\}} x_1 \rightarrow q_-$ at the beginning, we can generate the tree shown in Fig. 2(d).

The previous example does not fully demonstrate why incorporating a set of choices (the numbers in the curly brackets) into the design of LSTAs is beneficial. Let us consider another well-known example: the set of Bell states $\{|00\rangle \pm |11\rangle, |01\rangle \pm |10\rangle\}$, generated by the LSTA in Fig. 4(a). Without the restriction that all transitions at the same level

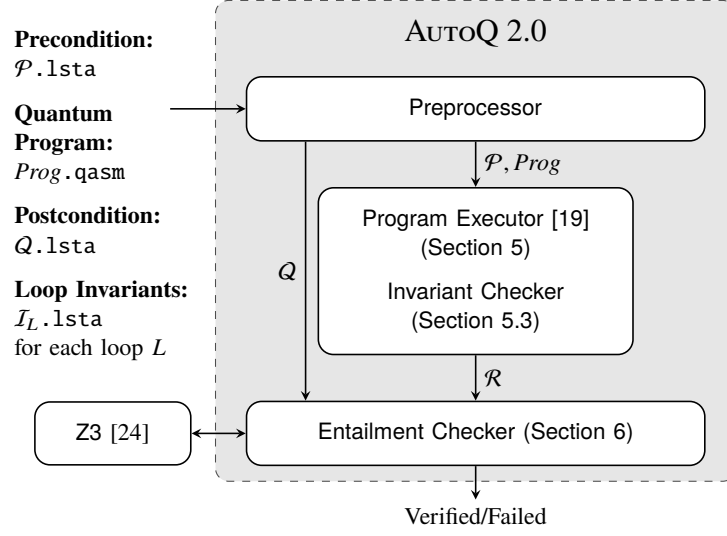


Fig. 5. The architecture of AutoQ 2.0

must share a common choice, this LSTA would generate eight different trees (since q_+ , q_- , and r_\pm each have two outgoing transitions), which correspond to the quantum states $\{|00\rangle \pm |11\rangle, |01\rangle \pm |10\rangle, |00\rangle \pm |10\rangle, |01\rangle \pm |11\rangle\}$. However, only four of these trees correspond to the Bell states, meaning the others are unintended. The LSTA uses the labeled choices to rule out the unintended trees. More specifically, at level 2, the two transitions labeled $\{1\}$ can be used simultaneously, as they share the common choice 1. Similarly, the two transitions labeled $\{2\}$ can be used together due to their common choice 2. In contrast, a transition labeled $\{1\}$ cannot be used alongside one labeled 2, as their choice sets are disjoint. At level 3, the transitions from r_\pm can be used freely with those from r_+ and r_0 , since $\{1\}$ (and $\{2\}) \cap \{1, 2\} \neq \emptyset$. There are two valid combinations of transitions at levels 2 and 3, and this LSTA generates exactly the four Bell states using the nine transitions shown in the figure. The corresponding `.lsta` file, which illustrates the input format for AutoQ 2.0, is shown in Fig. 4(b). In `.lsta` files, transitions are labeled with pairs $[a, b]$, where a indicates the symbol x_a and b is the set of choices.

4 System Architecture

We present the architecture of AutoQ 2.0 in Fig. 5. The tool is written in C++ and uses the SMT solver Z3 [24] for satisfiability and entailment checking of constraints. We allow the use of any theory that is supported by Z3. In our experiments, we used NIRA (non-linear integer and real arithmetic). While this logic is generally undecidable, Z3 always quickly solved the formulae we presented to it in our experiments.

Similar to verifiers for classical programs, in order to use AutoQ 2.0, the user needs to provide the following: (i) a *quantum program* in the OPENQASM 3.0 format, (ii) *pre- and post-conditions* in the `.lsta` format along with SMT formulae in the `.smt` format

Algorithm 3: Algorithm for measurement

Input: LSTA $\mathcal{A} = \langle Q, \Sigma, \Delta, \mathcal{R}, \varphi \rangle$, target qubit x_i , measurement outcome b

Output: The LSTA $M_i^{=b}(\mathcal{A})$

- 1 $Q' := \{q' \mid q \in Q\};$
 - 2 $\Delta' := \{q' \xrightarrow{f|C} (q'_1, q'_2) \mid q \xrightarrow{f|C} (q_1, q_2) \in \Delta\} \cup \{q' \xrightarrow{0|C} () \mid q \xrightarrow{f|C} () \in \Delta\};$
 - 3 $\Delta^{\neq x_i} := \{q \xrightarrow{f|C} (q_0, q_1) \mid q \xrightarrow{f|C} (q_0, q_1) \in \Delta \wedge f \neq x_i\} \cup \{q \xrightarrow{f|C} () \mid q \xrightarrow{f|C} () \in \Delta\};$
 - 4 **if** $b = 0$ **then** $\Delta^{=x_i} := \{q \xrightarrow{x_i|C} (q_0, q'_1) \mid q \xrightarrow{x_i|C} (q_0, q_1) \in \Delta\};$
 - 5 **else** $\Delta^{=x_i} := \{q \xrightarrow{x_i|C} (q'_0, q_1) \mid q \xrightarrow{x_i|C} (q_0, q_1) \in \Delta\};$
 - 6 **return** $M_i^{=b}(\mathcal{A}) = \langle Q \cup Q', \Sigma, \Delta^{=x_i} \cup \Delta^{\neq x_i} \cup \Delta', \mathcal{R}, \varphi \rangle;$
-

to constrain the terms, and (iii) *invariant* for each loop in the .lsta format together with an SMT formula. The specification (pre- and post-conditions) and the invariant (for each loop) can be written as an LSTA (an .lsta file). Once these files are provided, AutoQ 2.0 will process them and report either “Verified” or “Failed”.

Compared to AutoQ 1.0, there are several major changes in AutoQ 2.0. Firstly, it features a new input interface to facilitate the use of quantum programs (instead of only circuits) and uses LSTA as the back-end model (instead of standard tree automata). Additionally, Program Executor now supports measurement and branch statements. Another significant addition is the new Invariant Checker component, which handles loop invariants. The Invariant Checker also uses the Entailment Checker to verify the *inductiveness* of the invariant, which we do not explicitly show in the figure.

5 Handling Branch, Measurement, and Loop

We will start by presenting the syntax of quantum programs that AutoQ 2.0 can handle and then informally describe their semantics. We use a flavor of quantum programs that is similar to the one in [48], which is captured by the following grammar:

$$P ::= U \mid P; P \mid \textbf{while } (M_i = b) \textbf{ do } \{P\} \mid \textbf{if } (M_i = b) \textbf{ then } \{P\} \textbf{ else } \{P\}$$

where P is a quantum program, U is a quantum gate annotated with its control and target qubits (e.g., CX_1^2), $b \in \{0, 1\}$, and M_i is the measured value of the i -th qubit. AutoQ 2.0 supports standard non-parameterized quantum gates that allow (approximate) universal computation [12,3], including Clifford gates (H , S , and CX), T , Z , $SWAP$, Toffoli, etc.

The execution of quantum gate U updates a quantum state (tree) in the standard way [20]. The language allows sequential composition ($P; P$) of gate operations, branches (**if** . . . **else** . . .), and loops (**while** . . .). When using **if** and **while** statements, the condition $M_i = b$ (denoting that the value obtained from measuring x_i was b) is used to determine in which path to continue.

5.1 Handling Measurement

The key part of handling branch statements in AutoQ 2.0 is how measurement changes the quantum states and how we should update the LSTA capturing the set of reachable

states. As mentioned in Section 3, if the measured value of x_i is 1, then we should update the probability of the 0-subtrees below all x_i nodes to 0. Examples can be found in Fig. 2(b) (before measuring x_1) and Figs. 2(c) and 2(d) (after measuring x_1 as 0 and 1 respectively). An important design choice was that we do not normalize the probability amplitudes and simply just make all leaves of one of the subtrees zero in this step, leaving the task of *matching non-normalized states* to the entailment test (cf. Section 6).

In some cases, the measurement can generate an LSTA whose language contains a tree where all leaves are 0. This can happen, e.g., when we compute the tree representing the quantum state obtained from the state in Fig. 2(a) by measuring $x_1 = 0$. We do not consider such a tree to represent a quantum state. To handle such cases, our entailment test $\mathcal{R} \models^{uts} Q$ (formally defined in Section 6) adds a 0-labeled tree to the language of Q before the test. We use $M_i^{=b}(\mathcal{A})$ to represent the LSTA obtained from \mathcal{A} after measuring x_i for the outcome b . The procedure for computing $M_i^{=b}(\mathcal{A})$ is given in Algorithm 3.

The goal of the algorithm is to update all leaf values of \bar{b} -subtrees under x_i to 0, where $\bar{b} = 1 - b$. Lines 1 and 2 of Algorithm 3 create a primed copy of the input LSTA and update all leaf values to 0 (Line 2). Lines 3 to 5 construct the new transition system: only transitions labeled with x_i are modified (Lines 4 and 5), while others remain unchanged (Line 3). The key steps are in Lines 4 and 5, which control the subtrees of the measured qubit. When $b = 0$ (Line 4), all leaves of the 1-subtree are modified to 0, and thus, we update q_1 in the original transition to q'_1 (symmetrically for $b = 1$ in Line 5).

5.2 Handling Branch Statements

Given an LSTA \mathcal{A} that recognizes a set of quantum states, we can precisely compute the set of states that result from executing a branch statement **if** ($M_i = b$) **then** $\{P_1\}$ **else** $\{P_0\}$ as follows (assuming that P_0 and P_1 do not involve loops): (i) Create two LSTAs $M_i^{=1}(\mathcal{A})$ and $M_i^{=0}(\mathcal{A})$. (ii) Compute the result after executing P_b from $M_i^{=b}(\mathcal{A})$ for $b \in \{0, 1\}$, following the gates' semantics and recursively trigger the procedure for branches. We use \mathcal{A}^0 and \mathcal{A}^1 to denote the LSTAs after executing P_0 and P_1 , respectively. (iii) Construct an LSTA recognizing the union of \mathcal{A}^0 and \mathcal{A}^1 and return it as the final result of this procedure. In principle, our approach can handle nested control flow. We are, however, not aware of any real-world quantum program that uses a nested control structure, and, therefore, for simplicity, AutoQ 2.0 now only supports programs with non-nested control flow.

5.3 Handling Loop Statements

If we come across a loop statement **while** ($M_i = b$) **do** $\{B\}$ with B being the loop body, we require the user to provide a loop invariant in the form of an LSTA. We refer to the invariant as \mathcal{I} ; it needs to satisfy the following properties: (i) It contains all reachable states, captured by an LSTA \mathcal{R} , before entering the loop. That is, $\mathcal{R} \models^{uts} \mathcal{I}$. (ii) It is *inductive*, i.e., $B(M_i^{=b}(\mathcal{I})) \models^{uts} \mathcal{I}'$, where $B(\mathcal{A})$ denotes an LSTA recognizing the set of quantum states after executing B from the quantum states in \mathcal{A} and \mathcal{I}' is an LSTA obtained from \mathcal{I} whose variables and constraints are updated to a primed version. The inductiveness guarantees that if we take any state accepted by \mathcal{I} and perform B on the

state, the result will also be accepted by \mathcal{I} . Together with the condition that \mathcal{I} covers all reachable states before entering the loop, \mathcal{I} *over-approximates* all reachable states at the loop entrance. We can, therefore, use $M_i^{\#b}(\mathcal{I})$ to over-approximate all reachable states at the loop exit.

6 Testing Entailment up to Scaling

In our approach, we use a special entailment test at some points, which we call *entailment up to scaling*, denoted as $\mathcal{A} \models^{uts} \mathcal{B}$. The reason for a special entailment relation is that—as mentioned before—at measurements, we do not perform normalization. Intuitively, given two LSTAs $\mathcal{A} = \langle Q_{\mathcal{A}}, \Sigma, \Delta_{\mathcal{A}}, \mathcal{R}_{\mathcal{A}}, \varphi_{\mathcal{A}} \rangle$ and $\mathcal{B} = \langle Q_{\mathcal{B}}, \Sigma, \Delta_{\mathcal{B}}, \mathcal{R}_{\mathcal{B}}, \varphi_{\mathcal{B}} \rangle$, the relation $\mathcal{A} \models^{uts} \mathcal{B}$ holds if and only if for every tree $T_{\mathcal{A}}$ in the language of \mathcal{A} and assignment to the variables occurring in $T_{\mathcal{A}}$, we can find a linearly scaled copy of $T_{\mathcal{A}}$ in the semantics of \mathcal{B} (such that the values of variables occurring in both $T_{\mathcal{A}}$ and $T_{\mathcal{B}}$ match). Formally,

$$\begin{aligned} \mathcal{A} \models^{uts} \mathcal{B} \iff & (\forall T_{\mathcal{A}} \in \mathcal{L}(\mathcal{A})) (\forall \sigma_{\mathcal{A}} : \text{vars}(T_{\mathcal{A}}) \rightarrow \mathbb{C}) : \\ & (\exists T_{\mathcal{B}} \in \mathcal{L}(\mathcal{B})) (\exists \sigma_{\mathcal{B}} : (\text{vars}(T_{\mathcal{B}}) \setminus \text{vars}(T_{\mathcal{A}})) \rightarrow \mathbb{C}) : \\ & (\exists r \in \mathbb{R} \setminus \{0\}) : T_{\mathcal{A}}[\sigma_{\mathcal{A}}] = r \cdot T_{\mathcal{B}}[\sigma_{\mathcal{A}} \cup \sigma_{\mathcal{B}}], \end{aligned}$$

where $r \cdot T$ denotes the tree with the same structure as T with all numbers in leaves multiplied by r and $\text{vars}(\gamma)$ for any mathematical object γ (a term, a tree with terms in leaves, a set of terms, etc.) denotes the set of free variables occurring in the object. The \models^{uts} relation is central to our approach.

Enumerating all trees of \mathcal{A} and looking for their scaled copies in \mathcal{B} would be too inefficient and even impossible in the case of LSTAs recognizing infinitely many finite trees (such as those modelling invariants of parameterized quantum programs [1]). Therefore, we modified the algorithm for LSTA language inclusion test presented in [1]. We note that language inclusion testing for LSTAs is more involved than for standard TAs (cf. [11,2,35,38]). In the modification, we allow to relate the leaf values with a linear factor for scaling (in contrast to only by identity as done in the original inclusion testing algorithm), so that it tests the entailment $\mathcal{A} \models^{uts} \mathcal{B}$.

The algorithm makes use of the following essential property of trees generated by an LSTA \mathcal{A} : if two nodes at the same level of a tree T are labelled by the same state in an accepting run of \mathcal{A} on T , then the subtrees rooted in these nodes are identical (this follows from the semantics of LSTAs and the restriction on transitions, cf. Section 2.2).

Intuitively, the algorithm works as follows. It starts in the root states of \mathcal{A} and \mathcal{B} and performs a downward traversal through the LSTAs, level by level, remembering, at each level, how states from \mathcal{A} can map to the states in \mathcal{B} . Moreover, the algorithm also remembers how the terms in the leaves of the tree from \mathcal{A} map to the terms in the leaves of the tree from \mathcal{B} . The downward successors of each level are computed from transitions leaving states at the level that need to be synchronized on their choice. The algorithm explores the space of all of the reachable mappings until it reaches a point such that the tree from \mathcal{A} has all branches terminated. At this moment, we check whether the terms from the leaf transitions of \mathcal{A} can be mapped (up to scaling) to the corresponding terms from the leaf transitions in \mathcal{B} . If not, conclude that the entailment does not hold.

Formally, the algorithm performs a search in the directed graph (V, E) (constructed on the fly), where each vertex $v \in V$ is of the form $v = (D, \{(f_1, g_1), \dots, (f_m, g_m)\})$ for some m where $D \subseteq Q_{\mathcal{A}}$ is called the *domain*, each $f_i: D \rightarrow 2^{Q_{\mathcal{B}}}$ is a map that assigns every state of \mathcal{A} from the domain D a set of states of \mathcal{B} , and each $g_i: \text{term}(\mathbb{C}, \mathbb{X}) \rightarrow 2^{\text{term}(\mathbb{C}, \mathbb{X})}$ is a (partial) mapping from terms to sets

of terms. Intuitively, D represents the set of states of \mathcal{A} at one level of \mathcal{A} 's run ρ , and every f_i represents the same level of some possible run ρ_i of \mathcal{B} on the same tree and the way it can match the run ρ of \mathcal{A} . For instance, in Section 6, the state q of \mathcal{A} corresponds to the states r, s of \mathcal{B} because they are used in the same tree level and the same tree nodes, so we have $f_i(q) = \{r, s\}$. Due to the property that all occurrences of a state at the same level in a run generate the same subtree mentioned above, we only need to maintain encountered states and their alignment with each another. The term mappings g_i are used to remember how the terms from the leaves of \mathcal{A} are mapped to terms in the leaves of \mathcal{B} . For instance, if some term t from a leaf of \mathcal{A} is mapped by two different terms t_1 and t_2 of \mathcal{B} , we will later need to check whether there is a scaling factor r such that $t = r \cdot t_1$ and, at the same time, $t = r \cdot t_2$ (and the global constraints of \mathcal{A} and \mathcal{B} are satisfied). We give a general algorithm here; for LSTAs that accept only perfect trees (as is the case for the ones encoding quantum states), all branches of the accepted trees terminate at the same time so there is no need to remember the term mappings g across different levels and the scaling compatibility could be checked only locally.

The graph search begins from the *source* vertices, one for each state $q \in \mathcal{R}_{\mathcal{A}}$, of the form $(\{q\}, \{(\{q \mapsto \{r_1\}\}, \emptyset), \dots, (\{q \mapsto \{r_k\}\}, \emptyset)\})$, where $\{r_1, \dots, r_k\} = \mathcal{R}_{\mathcal{B}}$, corresponding to the root states of both \mathcal{A} and \mathcal{B} (the \emptyset 's denote empty term mappings). If the search finds a *terminal* vertex (\emptyset, F) , where $(\emptyset, g) \notin F$, meaning that an accepting run of \mathcal{A} has been found, but there is no corresponding matching run of \mathcal{B} (for any g), we can conclude that the entailment test failed (it represents the case when \mathcal{A} finished reading all branches of the tree but \mathcal{B} did not). On the other hand, if there is $(\emptyset, g) \in F$, we still need to check that the set of terms g is compatible. The graph's edges represent generating the next level of runs for both \mathcal{A} and \mathcal{B} and how the respective states align with each other. The specific construction of the edges from a vertex $v = (D, \{(f_1, g_1), \dots, (f_m, g_m)\})$, where $D \neq \emptyset$, to each lower-level vertex $v' = (D', \{(f'_1, g'_1), \dots, (f'_n, g'_n)\})$ follows.

First, we compute possible successors of the D -component of v . To do this, we need to explore all *feasible* sets $\Gamma_{\mathcal{A}}$ of transitions from D in \mathcal{A} . More concretely, in each set $\Gamma_{\mathcal{A}}$, we select exactly one downward transition $\delta_{q_{\mathcal{A}}}$ originating from each $q_{\mathcal{A}} \in D$, such that all transitions in $\Gamma_{\mathcal{A}}$ share a common choice (as required by the definition of an accepting run (cf. Section 2.2)). Formally, given $D = \{q_1, \dots, q_k\}$, we consider all sets of transitions $\Gamma_{\mathcal{A}} = \{\delta_1, \dots, \delta_k\}$ such that the following formula holds:

$$(\forall 1 \leq i \leq k: \delta_i \in \Delta \wedge \text{top}(\delta_i) = q_i) \quad \wedge \quad \bigcap \{\text{ch}(\delta_i) \mid 1 \leq i \leq k\} \neq \emptyset. \quad (1)$$




Fig. 6. Computing f'_j from f_i

Algorithm 4: FindAllMappings($\Gamma_{\mathcal{A}}, f$)

Input: Set of transitions $\Gamma_{\mathcal{A}} \subseteq \Delta_{\mathcal{A}}$, map $f: \{\text{top}(\delta) \mid \delta \in \Gamma_{\mathcal{A}}\} \rightarrow 2^{\mathcal{Q}_{\mathcal{B}}}$

Output: The set of pairs (f', g') of mappings compatible with $\Gamma_{\mathcal{A}}$ obtained from f

```
1  $Q \leftarrow \bigcup \text{img}(f); F' \leftarrow \emptyset;$ 
2 foreach  $\Gamma_{\mathcal{B}} \in \text{FTrans}(\mathcal{B}, Q)$  do
3    $f' \leftarrow \emptyset; g' \leftarrow \emptyset; \text{failed} \leftarrow \text{false};$ 
4   foreach  $\delta_{\mathcal{A}} \in \Gamma_{\mathcal{A}}$  and  $q \in f(\text{top}(\delta_{\mathcal{A}}))$  do
5      $\{\delta_{\mathcal{B}}\} \leftarrow \{\delta_{\mathcal{B}} \in \Gamma_{\mathcal{B}} \mid q = \text{top}(\delta_{\mathcal{B}})\};$ 
6     if  $a = \text{sym}(\delta_{\mathcal{A}})$  and  $b = \text{sym}(\delta_{\mathcal{B}})$  are both leaf symbols then  $g'(a).\text{insert}(b)$  ;
7     else if  $\text{sym}(\delta_{\mathcal{A}}) = \text{sym}(\delta_{\mathcal{B}})$  is an internal symbol then
8        $f'(\text{left}(\delta_{\mathcal{A}})).\text{insert}(\text{left}(\delta_{\mathcal{B}})); f'(\text{right}(\delta_{\mathcal{A}})).\text{insert}(\text{right}(\delta_{\mathcal{B}}));$ 
9     else  $\text{failed} \leftarrow \text{true}; \text{break};$ 
10  if  $\neg \text{failed}$  then  $F'.\text{insert}((f', g'))$  ;
11 return  $F'$ ;
```

We will denote the set of all such $\Gamma_{\mathcal{A}}$'s from a set of states D as $\text{FTrans}(\mathcal{A}, D)$.

Next, we will show how to construct the set of all feasible pairs of mappings $\{(f'_1, g'_1), \dots, (f'_n, g'_n)\}$ for some $\Gamma_{\mathcal{A}}$ and a pair of upper level mappings (f, g) . Each pair of mappings (f'_j, g'_j) at the lower level is derived from some pair of mappings (f_i, g_i) at the upper level and a set of downward transitions $\Gamma_{\mathcal{B}}$ from \mathcal{B} . The construction process is described in Algorithm 4, where we use (f, g) to denote an upper level state and term mapping respectively and (f', g') to denote their lower level counterparts.

The basic idea of the algorithm is simple: (i) we use the upper level mapping f and transitions $\Gamma_{\mathcal{A}}$ to compute the set of top states Q (Line 1), (ii) then, we find all feasible transitions $\Gamma_{\mathcal{B}}$ from Q (Line 2), and, finally, (iii) for each pair $(\Gamma_{\mathcal{A}}, \Gamma_{\mathcal{B}})$, we construct one pair of lower level state and term mappings (f', g') (Lines 4–9). Specifically, in step (iii), it must hold that for each transition $\delta_{\mathcal{A}} \in \Gamma_{\mathcal{A}}$, every state $q \in f(\text{top}(\delta_{\mathcal{A}}))$ needs to be able to match $\delta_{\mathcal{A}}$ by a transition from $\Gamma_{\mathcal{B}}$. To check this, for every such q we select from $\Gamma_{\mathcal{B}}$ the transition $\delta_{\mathcal{B}}$, which is the transition of $\Gamma_{\mathcal{B}}$ with q as its top (it follows from FTrans that there is exactly one). There are three possible cases:

- *Both $\delta_{\mathcal{A}}$ and $\delta_{\mathcal{B}}$ are leaf transitions (Line 6):* we remember in g' that the symbol of $\delta_{\mathcal{B}}$ need to be able to match the symbol of $\delta_{\mathcal{A}}$, which will be checked later for all such matchings together.
- *The symbols of both $\delta_{\mathcal{A}}$ and $\delta_{\mathcal{B}}$ are internal (Line 7):* in this case, we add new entries to the lower level mapping f' .
- *One transition is internal while the other is leaf (Line 9):* the pair $(\Gamma_{\mathcal{A}}, \Gamma_{\mathcal{B}})$ cannot form a feasible lower level mapping.

Finally, the main entailment testing routine is summarized in Algorithm 5. Line 2 creates the set of source vertices of the explored graph. Lines 3–4 pick a vertex (D, F) that has not been processed yet. Lines 6–10 check whether (D, F) is a terminal vertex and conclude that the entailment test fails when such a vertex is reached. This check consists

Algorithm 5: Checking if $\mathcal{A} \models^{uts} \mathcal{B}$

Input: LSTAs $\mathcal{A} = \langle Q_{\mathcal{A}}, \mathbb{N} \cup \text{term}(\mathbb{C}, \mathbb{X}), \Delta_{\mathcal{A}}, \mathcal{R}_{\mathcal{A}}, \varphi_{\mathcal{A}} \rangle$,

$\mathcal{B} = \langle Q_{\mathcal{B}}, \mathbb{N} \cup \text{term}(\mathbb{C}, \mathbb{X}), \Delta_{\mathcal{B}}, \mathcal{R}_{\mathcal{B}} = \{r_1, \dots, r_k\}, \varphi_{\mathcal{B}} \rangle$.

Output: **true** if $\mathcal{A} \models^{uts} \mathcal{B}$, **false** otherwise

```

1 processed  $\leftarrow \emptyset$ ;
2 workset  $\leftarrow \{(\{q\}, \{(\{q \mapsto \{r_1\}\}, \emptyset), \dots, (\{q \mapsto \{r_k\}\}, \emptyset)\}) \mid q \in \mathcal{R}_{\mathcal{A}}\}$ ;
3 while  $\exists (D, F) \in \text{workset}$  do
4   workset.remove( $(D, F)$ );
5   processed.insert( $(D, F)$ );
6   if  $D = \emptyset$  then
7      $\mathbb{Y} \leftarrow \{\text{vars}(u_1) \mid (u_1 \mapsto U_2) \in g \wedge (\emptyset, g) \in F\} \cup \text{vars}(\varphi_{\mathcal{A}})$ ;
8      $\mathbb{Z} \leftarrow \{\text{vars}(U_2) \mid (u_1 \mapsto U_2) \in g \wedge (\emptyset, g) \in F\} \cup \text{vars}(\varphi_{\mathcal{B}}) \setminus \mathbb{Y}$ ;
9     if  $\neg \forall \mathbb{Y}. \varphi_{\mathcal{A}} \implies \exists \mathbb{Z}. \varphi_{\mathcal{B}} \wedge \bigvee_{(\emptyset, g) \in F} \exists r \in \mathbb{R} \setminus \{0\}. \bigwedge_{\substack{(u_1 \mapsto U_2) \in g \\ u_2 \in U_2}} u_1 = r \cdot u_2$  then
10      return false;
11   foreach  $\Gamma_{\mathcal{A}} \in \text{FTrans}(\mathcal{A}, D)$  do
12      $D' \leftarrow \{q \in \text{bot}(\delta) \mid \delta \in \Gamma_{\mathcal{A}}\}$ ;  $F' \leftarrow \emptyset$ ;
13     foreach  $(f, g) \in F$  do
14        $U \leftarrow \text{FindAllMappings}(\Gamma_{\mathcal{A}}, f)$ ;
15        $F' \leftarrow F' \cup \{(f', g \uplus g') \mid (f', g') \in U\}$ ;
16   if  $(D', F') \notin \text{processed} \cup \text{workset}$  then workset.insert( $(D', F')$ );
17 return true;
```

of looking at all pairs (\emptyset, g) in F and checking whether there is a way how the \mathcal{B} -terms from the g 's can together cover (modulo a scaling factor) the behaviour of the \mathcal{A} -terms. Lines 11–16 are the edge construction procedure. Lines 11–12 enumerate all feasible $\Gamma_{\mathcal{A}}$ and use them to create the next vertex (D', F') . Specifically, D' are the bottom states from $\Gamma_{\mathcal{A}}$ (Line 12), and F' are the union of all feasible mappings of $\Gamma_{\mathcal{A}}$ (Lines 13–16). The successor term mapping is computed from the upper-level one and from the one returned by `FindAllMappings` by, for each term of \mathcal{A} , merging corresponding sets of terms of mappings: $g \uplus g' = \{x \mapsto Y \mid x \in \text{dom}(g \cup g'), Y = g(x) \cup g'(x)\}$ (Line 15).

A crucial part of the algorithm is the term mapping g check between the terms from \mathcal{A} and the terms from \mathcal{B} (Lines 7–10). Here, we want to check that for every possible value that we can obtain from a term t on the left-hand side of the entailment (satisfying \mathcal{A} 's global constraint $\varphi_{\mathcal{A}}$) and every term t_i from $g(t) = \{t_1, \dots, t_k\}$, there is a way to obtain a value (satisfying \mathcal{B} 's global constraint $\varphi_{\mathcal{B}}$) such that for all these values, there is a common scaling factor r to make them equal to the value from t . We emphasize the way how we need to deal with the quantified variables. Variables \mathbb{Y} occurring on the left-hand side (and possibly also on the right-hand side, since we may need to synchronize the values) are quantified universally, while variable \mathbb{Z} occurring only on the right-hand side of the entailment are quantified existentially. This scaling check requires invoking an SMT solver with a formula in the NIRA (non-linear integer and real arithmetic) logic. In the special case where all leaf symbols are constants, r is the only variable, and global constraints are **true**, the problem reduces to a simple QF_LRA (quantifier-free linear real arithmetic) formula.

Theorem 1 (Soundness). *When Algorithm 5 terminates, it returns true iff $\mathcal{A} \models^{us} \mathcal{B}$.*

Theorem 2 (Termination). *When the terms in leaf symbols and the global constraints of \mathcal{A} and \mathcal{B} use a decidable theory, the algorithm always terminates.*

Proof. Since the number of states and terms occurring in \mathcal{A} and \mathcal{B} is finite, the constructed graph is also finite. Further, since the underlying theory for the terms and global constraints is assumed to be decidable, the check at Line 9 always terminates. \square

7 Experimental Results

We demonstrate the use of AutoQ 2.0 [21] on two real-world use cases consisting of quantum programs with loops that were proposed in [41,6]. We ran all experiments on a server running Ubuntu 22.04.3 LTS with an AMD EPYC 7742 64-core processor (1.5 GHz), 1,152 GiB of RAM, and a 1 TB SSD.

7.1 The Weakly Measured Version of Grover’s Algorithm

Grover’s algorithm [31], introduced in 1996, is a quantum algorithm that performs an unstructured search. Given an *oracle function* (which can say whether a particular binary assignment is a solution), Grover’s algorithm can efficiently find a solution (with high probability). The algorithm requires approximately $O(\sqrt{N/k})$ evaluations of the oracle function, where N is the size of the function’s domain

(usually 2^n for n qubits), and k is the number of solutions. The number of solutions is, however, not always known, making it difficult to determine the algorithm’s parameters (the algorithm is sensitive to the number of evaluations; in particular, doing more evaluations may make the probability of finding the solution smaller). To address this issue, a variation of Grover’s search, called the *weakly measured* version (cf. Algorithm 6), was recently proposed [6]. The weakly measured version eliminates the need for knowing the number of solutions, making the algorithm more applicable.

To explain the algorithm, we first introduce some of its key components. The algorithm works over qubits x_1, \dots, x_{n+2} . Line 2 first applies multiple *Hadamard* gates in parallel to obtain the superposition on all qubits other than x_1 and x_2 (which are two ancillas), obtaining the state in Fig. 7(a). The *oracle* circuit, denoted as $O_{2,\dots,(n+2)}$, works from qubits x_2 to x_{n+2} , where x_2 is the ancilla qubit and x_3 to x_{n+2} are the working qubits. As shown from Figs. 7(a) and 7(b) (and also from Figs. 7(c) and 7(d)), the effect of the oracle circuit is to flip the ancilla qubit of the computational bases corresponding to the solutions. That is, it swaps the amplitude values of $|0s\rangle$ and $|1s\rangle$, for all solutions s . The oracle circuit can be constructed using gates supported in AutoQ 2.0.

Algorithm 6: A Weakly Measured Version of Grover’s algorithm (solution $s = 0^n$)

```

1 Pre:  $\{1|0^{n+2}\rangle + 0|*\rangle\}$ ;
2  $H_3; H_4; \dots; H_{n+2}$ ;
3  $O_{2,\dots,(n+2)}; CK_1^2; O_{2,\dots,(n+2)}$ ;
4 Inv:  $\{v_{sol1}|000^n\rangle + v_k|000^{n-1}1\rangle + \dots +$ 
5  $v_k|001^n\rangle + v_{sol2}|100^n\rangle + 0|*\rangle\}$ ;
6 while  $M_1 = 0$  do
7    $\{G_{2,\dots,(n+2)}; O_{2,\dots,(n+2)}; CK_1^2; O_{2,\dots,(n+2)}\}$ ;
8 Post:  $\{1|10s\rangle + 0|*\rangle\}$ ;

```

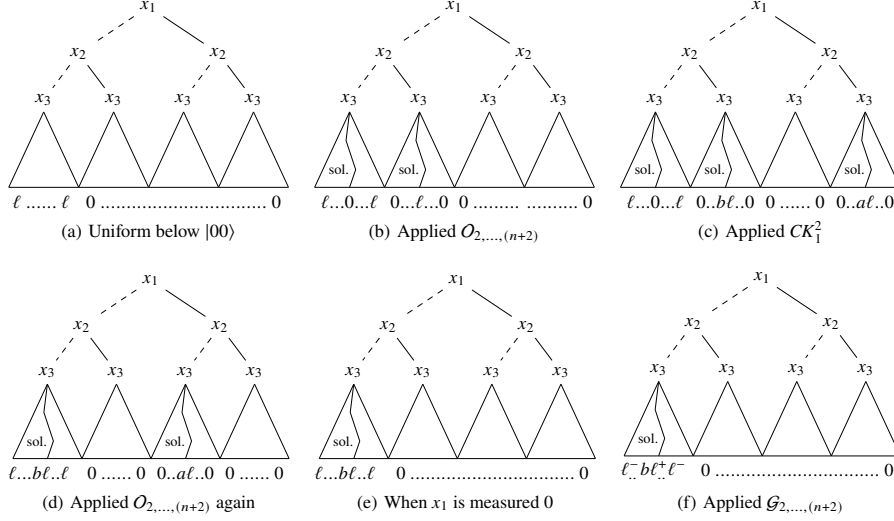


Fig. 7. Intermediate states of Algorithm 6

The controlled rotation gate CK_j^i is a special gate supported in **AUTOQ 2.0**. In this algorithm, the gate always applies to a target qubit whose value is $|0\rangle$ when the controlled qubit is $|1\rangle$, and it updates the target qubit to $a|1\rangle + b|0\rangle$ with $a^2 + b^2 = 1$, for some very small b . In **AUTOQ 2.0**, we use $a = \frac{21}{221}$ and $b = \frac{220}{221}$. We demonstrate the behavior of CK_1^2 from Figs. 7(b) and 7(c), where a small portion ($\frac{21^2}{221^2} \approx 1\%$) of the probability under the branch $|01\rangle$ is moved to the branch $|11\rangle$, as shown in Fig. 7(c). After applying $O_{2,\dots,(n+2)}$ again, we obtain the state in Fig. 7(d) (this state is captured by the loop invariant). Here, we can already measure the qubit x_1 and if the result is 1, this collapses the probability of the left sub-tree of x_1 in Fig. 7(d) to 0, so the only non-zero probability basis is the solution $|10s\rangle$.

Otherwise (the result of measuring x_1 was 0), we enter the loop, which contains the *Grover iteration* circuit, denoted as $\mathcal{G}_{2,\dots,(n+2)}$, which also uses $O_{2,\dots,(n+2)}$ as a component. The effect of $\mathcal{G}_{2,\dots,(n+2)}$ is to increase the probabilities of basis states for the solutions and decrease others, as shown in Figs. 7(e) and 7(f). After $\mathcal{G}_{2,\dots,(n+2)}$, we execute the same sequence *oracle-rotation-oracle* as above to obtain a state resembling Fig. 7(d). We keep repeating the above procedure until we measure $x_1 = 1$, in which case we terminate with a solution.

The results of the verification of weakly measured Grover’s search are in the left-hand side of Table 1: **AUTOQ 2.0** was able to verify the program w.r.t. the specification even for larger numbers of qubits in reasonable time.

7.2 Unitaries as Repeat-Until-Success Circuits

Repeat-until-success programs are a general framework that was introduced to simplify quantum circuit decomposition (we introduced an example of generating the “ $-X$ ” gate via the RUS framework in Section 3). RUS programs have been shown to be

Table 1. Results of verifying some real-world examples with AutoQ 2.0. The number x in WMGrover(x) indicates that the number of items to be searched is 2^x .

<i>Weakly Measured Grover's Search [6]</i>						<i>Repeat-Until-Success [41]</i>					
program	qubits	gates	result	time	memory	program	qubits	gates	result	time	memory
WMGrover (03)	7	50	OK	0.0s	42MB	$(2X + \sqrt{2}Y + Z)/\sqrt{7}$	2	29	OK	0.0s	7MB
WMGrover (10)	21	169	OK	0.2s	42MB	$(I + i\sqrt{2}X)/\sqrt{3}$	2	17	OK	0.0s	7MB
WMGrover (20)	41	339	OK	0.8s	42MB	$(I + 2iZ)/\sqrt{5}$	2	27	OK	0.0s	6MB
WMGrover (30)	61	509	OK	2.3s	43MB	$(3I + 2iZ)/\sqrt{13}$	2	43	OK	0.0s	7MB
WMGrover (40)	81	679	OK	5.4s	43MB	$(4I + iZ)/\sqrt{17}$	2	77	OK	0.0s	6MB
WMGrover (50)	101	849	OK	11s	44MB	$(5I + 2iZ)/\sqrt{29}$	2	69	OK	0.0s	7MB

more efficient (in terms of circuit depth) than ancilla-free techniques when it comes to synthesizing single-qubit gates (cf. [41,10]). We present the results of verification of RUS programs for generating various non-standard gates in the right-hand side of Table 1. Note that AutoQ 2.0 can verify these programs instantaneously.

8 Conclusion and Future Work

We presented a major extension of AutoQ 1.0 [19] with an added support for control flow constructs and evaluated its feasibility on a family of programs for the weak-measurement-based version of Grover's algorithm and on implementations of a number of non-standard quantum gates using repeat-until-success circuits. In the future, we wish to extend the framework with automating invariant generation (e.g., using a modification of the symbolic-execution-based technique from [16]) and add support for dealing with more complex loops that give rise to mixed states.

Acknowledgements.

We thank the reviewers for their useful remarks that helped us improve the quality of the paper. This work was supported by the Czech Ministry of Education, Youth and Sports project LL1908 of the ERC.CZ programme, the Czech Science Foundation project 25-18318S, the FIT BUT internal project FIT-S-23-8151, the National Science and Technology Council, Taiwan, projects under Grant no. NSTC 113-2222-E-027-007- and 112-2222-E-001-002-MY3, Foxconn research project 05T-1120327-1C-, Academia Sinica Grand Challenge Seeding Project, and Academia Sinica Principal Investigator Project no. AS-IV-114-M07-ASSA.

Data Availability Statement.

An environment with the tools and data used for the experimental evaluation in the current study is available at [18].

References

1. Abdulla, P.A., Chen, Y.G., Chen, Y.F., Holík, L., Lengál, O., Lin, J.A., Lo, F.Y., Tsai, W.L.: Verifying quantum circuits with level-synchronized tree automata. *Proceedings of the ACM on Programming Languages* **9**(POPL), 923–953 (2025)

2. Abdulla, P.A., Chen, Y., Holík, L., Mayr, R., Vojnar, T.: When simulation meets antichains. In: Esparza, J., Majumdar, R. (eds.) Tools and Algorithms for the Construction and Analysis of Systems, 16th International Conference, TACAS 2010, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2010, Paphos, Cyprus, March 20-28, 2010. Proceedings. Lecture Notes in Computer Science, vol. 6015, pp. 158–174. Springer (2010). https://doi.org/10.1007/978-3-642-12002-2_14, https://doi.org/10.1007/978-3-642-12002-2_14
3. Aharonov, D.: A simple proof that Toffoli and Hadamard are quantum universal (2003). <https://doi.org/10.48550/arxiv.quant-ph/0301040>, <https://arxiv.org/abs/quant-ph/0301040>
4. Ahrendt, W., Beckert, B., Bubel, R., Hähnle, R., Schmitt, P.H., Ulbrich, M. (eds.): Deductive Software Verification - The KeY Book - From Theory to Practice, Lecture Notes in Computer Science, vol. 10001. Springer (2016). <https://doi.org/10.1007/978-3-319-49812-6>, <https://doi.org/10.1007/978-3-319-49812-6>
5. Amy, M.: Towards large-scale functional verification of universal quantum circuits. In: Quantum Physics and Logic (2018)
6. Andrés-Martínez, P., Heunen, C.: Weakly measured while loops: peeking at quantum states. Quantum Science and Technology 7(2), 025007 (feb 2022). <https://doi.org/10.1088/2058-9565/ac47f1>, <https://dx.doi.org/10.1088/2058-9565/ac47f1>
7. Anticoli, L., Piazza, C., Taglialegne, L., Zuliani, P.: Towards quantum programs verification: From Quipper circuits to QPMC. In: Devitt, S.J., Lanese, I. (eds.) Reversible Computation - 8th International Conference, RC 2016, Bologna, Italy, July 7-8, 2016, Proceedings. Lecture Notes in Computer Science, vol. 9720, pp. 213–219. Springer (2016). https://doi.org/10.1007/978-3-319-40578-0_16, https://doi.org/10.1007/978-3-319-40578-0_16
8. Baudin, P., Bobot, F., Bühler, D., Correnson, L., Kirchner, F., Kosmatov, N., Maroneze, A., Perrelle, V., Prevosto, V., Signoles, J., Williams, N.: The dogged pursuit of bug-free C programs: the Frama-C software analysis platform. Commun. ACM 64(8), 56–68 (2021). <https://doi.org/10.1145/3470569>
9. Bertot, Y., Castéran, P.: Interactive theorem proving and program development: Coq’Art: the calculus of inductive constructions. Springer Science & Business Media (2013)
10. Bocharov, A., Roetteler, M., Svore, K.M.: Efficient synthesis of universal repeat-until-success quantum circuits. Phys. Rev. Lett. 114, 080502 (Feb 2015). <https://doi.org/10.1103/PhysRevLett.114.080502>, <https://link.aps.org/doi/10.1103/PhysRevLett.114.080502>
11. Bouajjani, A., Habermehl, P., Holík, L., Touili, T., Vojnar, T.: Antichain-based universality and inclusion testing over nondeterministic finite tree automata. In: Ibarra, O.H., Ravikumar, B. (eds.) Implementation and Applications of Automata, 13th International Conference, CIAA 2008, San Francisco, California, USA, July 21-24, 2008. Proceedings. Lecture Notes in Computer Science, vol. 5148, pp. 57–67. Springer (2008). https://doi.org/10.1007/978-3-540-70844-5_7, https://doi.org/10.1007/978-3-540-70844-5_7
12. Boykin, P.O., Mor, T., Pulver, M., Roychowdhury, V.P., Vatan, F.: A new universal and fault-tolerant quantum basis. Inf. Process. Lett. 75(3), 101–107 (2000). [https://doi.org/10.1016/S0020-0190\(00\)00084-3](https://doi.org/10.1016/S0020-0190(00)00084-3), [https://doi.org/10.1016/S0020-0190\(00\)00084-3](https://doi.org/10.1016/S0020-0190(00)00084-3)
13. Burgholzer, L., Wille, R.: Advanced equivalence checking for quantum circuits. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems 40(9), 1810–1824 (2020)

14. 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: 30th European Symposium on Programming, ESOP 2021, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2021, Luxembourg City, Luxembourg, March 27–April 1, 2021, Proceedings 30. pp. 148–177. Springer International Publishing (2021)
15. Chareton, C., Bardin, S., Bobot, F., Perrelle, V., Valiron, B.: An automated deductive verification framework for circuit-building quantum programs. In: Yoshida, N. (ed.) ESOP. Lecture Notes in Computer Science, vol. 12648, pp. 148–177. Springer International Publishing, Cham (March 2021)
16. Chen, T., Chen, Y., Jiang, J.R., Lengál, O., Jobranová, S.: Accelerating quantum circuit simulation with symbolic execution and loop summarization. In: Proc. of ICCAD’24. ACM (2024)
17. Chen, T., Jiang, J.R., Hsieh, M.: Partial equivalence checking of quantum circuits. In: 2022 IEEE International Conference on Quantum Computing and Engineering (QCE). pp. 594–604 (2022). <https://doi.org/10.1109/QCE53715.2022.00082>
18. Chen, Y.F., Chung, K.M., Hsieh, M.H., Huang, W.J., Lengal, O., Lin, J.A., Tsai, W.L.: AutoQ 2.0: An automata-based quantum program verifier (TACAS artifact evaluation record). <https://zenodo.org/records/14114791> (2024)
19. Chen, Y.F., Chung, K.M., Lengál, O., Lin, J.A., Tsai, W.L.: AutoQ: An automata-based quantum circuit verifier. In: International Conference on Computer Aided Verification. pp. 139–153. Springer (2023)
20. Chen, Y., Chung, K., Lengál, O., Lin, J., Tsai, W., Yen, D.: An automata-based framework for verification and bug hunting in quantum circuits. In: 44th ACM SIGPLAN Conference on Programming Language Design and Implementation—PLDI’23. ACM (2023)
21. Chen, Y.F., Lengál, O., Tsai, W.L.: AutoQ 2.0: An automata-based C++ tool for quantum program verification (Nov 2024), <https://github.com/alan23273850/AutoQ/tree/tacas25>
22. Coecke, B., Duncan, R.: Interacting quantum observables: categorical algebra and diagrammatics. *New Journal of Physics* **13**(4), 043016 (apr 2011). <https://doi.org/10.1088/1367-2630/13/4/043016>
23. Cohen, E., Dahlweid, M., Hillebrand, M.A., Leinenbach, D., Moskal, M., Santen, T., Schulte, W., Tobies, S.: VCC: A practical system for verifying concurrent C. In: Berghofer, S., Nipkow, T., Urban, C., Wenzel, M. (eds.) Theorem Proving in Higher Order Logics, 22nd International Conference, TPHOLs 2009, Munich, Germany, August 17–20, 2009. Proceedings. Lecture Notes in Computer Science, vol. 5674, pp. 23–42. Springer (2009). https://doi.org/10.1007/978-3-642-03359-9_2, https://doi.org/10.1007/978-3-642-03359-9_2
24. De Moura, L., Bjørner, N.: Z3: An efficient SMT solver. In: Tools and Algorithms for the Construction and Analysis of Systems: 14th International Conference, TACAS 2008, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2008, Budapest, Hungary, March 29–April 6, 2008. Proceedings 14. pp. 337–340. Springer (2008)
25. D’Hondt, E., Panangaden, P.: Quantum weakest preconditions. *Mathematical Structures in Computer Science* **16**(3), 429–451 (2006)
26. Fang, W., Ying, M., Wu, X.: Differentiable quantum programming with unbounded loops. *ACM Trans. Softw. Eng. Methodol.* **33**(1) (nov 2023). <https://doi.org/10.1145/3617178>, <https://doi.org/10.1145/3617178>
27. Feng, Y., Ying, M.: Quantum Hoare logic with classical variables. *ACM Transactions on Quantum Computing* **2**(4), 1–43 (2021)

28. Feng, Y., Yu, N., Ying, M.: Model checking quantum Markov chains. *J. Comput. Syst. Sci.* **79**(7), 1181–1198 (2013). <https://doi.org/10.1016/j.jcss.2013.04.002>, <https://doi.org/10.1016/j.jcss.2013.04.002>
29. Filliâtre, J.C., Paskevich, A.: Why3—where programs meet provers. In: *Programming Languages and Systems: 22nd European Symposium on Programming, ESOP 2013, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2013, Rome, Italy, March 16-24, 2013. Proceedings* 22. pp. 125–128. Springer (2013)
30. Floyd, R.W.: Assigning meanings to programs. In: *Mathematical Aspects of Computer Science*. pp. 19–32. *Proceedings of Symposia in Applied Mathematics*, AMS (1967), <https://mathscinet.ams.org/mathscinet/article?mr=235771>
31. Grover, L.K.: A fast quantum mechanical algorithm for database search. In: Miller, G.L. (ed.) *Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing*, Philadelphia, Pennsylvania, USA, May 22-24, 1996. pp. 212–219. ACM (1996). <https://doi.org/10.1145/237814.237866>, <https://doi.org/10.1145/237814.237866>
32. Hähnle, R., Huisman, M.: Deductive software verification: From pen-and-paper proofs to industrial tools. In: Steffen, B., Woeginger, G.J. (eds.) *Computing and Software Science - State of the Art and Perspectives*, *Lecture Notes in Computer Science*, vol. 10000, pp. 345–373. Springer (2019). https://doi.org/10.1007/978-3-319-91908-9_18, https://doi.org/10.1007/978-3-319-91908-9_18
33. Hietala, K., Rand, R., Hung, S.H., Wu, X., Hicks, M.: Verified optimization in a quantum intermediate representation. *arXiv preprint arXiv:1904.06319* (2019)
34. Hoare, C.A.R.: An axiomatic basis for computer programming. *Commun. ACM* **12**(10), 576–580 (1969). <https://doi.org/10.1145/363235.363259>, <https://doi.org/10.1145/363235.363259>
35. Holík, L., Lengál, O., Šimáček, J., Vojnar, T.: Efficient inclusion checking on explicit and semi-symbolic tree automata. In: Bultan, T., Hsiung, P. (eds.) *Automated Technology for Verification and Analysis*, 9th International Symposium, ATVA 2011, Taipei, Taiwan, October 11-14, 2011. *Proceedings. Lecture Notes in Computer Science*, vol. 6996, pp. 243–258. Springer (2011). https://doi.org/10.1007/978-3-642-24372-1_18, https://doi.org/10.1007/978-3-642-24372-1_18
36. Jacobs, B., Smans, J., Piessens, F.: A quick tour of the VeriFast program verifier. In: Ueda, K. (ed.) *Programming Languages and Systems - 8th Asian Symposium, APLAS 2010, Shanghai, China, November 28 - December 1, 2010. Proceedings. Lecture Notes in Computer Science*, vol. 6461, pp. 304–311. Springer (2010). https://doi.org/10.1007/978-3-642-17164-2_21, https://doi.org/10.1007/978-3-642-17164-2_21
37. Leino, K.R.M.: Dafny: An automatic program verifier for functional correctness. In: Clarke, E.M., Voronkov, A. (eds.) *Logic for Programming, Artificial Intelligence, and Reasoning - 16th International Conference, LPAR-16, Dakar, Senegal, April 25-May 1, 2010, Revised Selected Papers. Lecture Notes in Computer Science*, vol. 6355, pp. 348–370. Springer (2010). https://doi.org/10.1007/978-3-642-17511-4_20, https://doi.org/10.1007/978-3-642-17511-4_20
38. Lengál, O., Šimáček, J., Vojnar, T.: VATA: A library for efficient manipulation of non-deterministic tree automata. In: *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*. pp. 79–94. Springer (2012)
39. Liu, J., Zhan, B., Wang, S., Ying, S., Liu, T., Li, Y., Ying, M., Zhan, N.: Formal verification of quantum algorithms using quantum Hoare logic. In: *International conference on computer aided verification*. pp. 187–207. Springer (2019)

40. Mateus, P., Ramos, J., Sernadas, A., Sernadas, C.: Temporal Logics for Reasoning about Quantum Systems, p. 389–413. Cambridge University Press (2009).
<https://doi.org/10.1017/CBO9781139193313.011>
41. Paetznick, A., Svore, K.M.: Repeat-until-success: Non-deterministic decomposition of single-qubit unitaries. *Quantum Info. Comput.* **14**(15–16), 1277–1301 (nov 2014)
42. Perdrix, S.: Quantum entanglement analysis based on abstract interpretation. In: International Static Analysis Symposium. pp. 270–282. Springer (2008)
43. Unruh, D.: Quantum Hoare logic with ghost variables. In: 2019 34th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS). pp. 1–13. IEEE (2019)
44. Wei, C.Y., Tsai, Y.H., Jhang, C.S., Jiang, J.H.R.: Accurate bdd-based unitary operator manipulation for scalable and robust quantum circuit verification. In: Proceedings of the 59th ACM/IEEE Design Automation Conference. p. 523–528. DAC '22, Association for Computing Machinery, New York, NY, USA (2022).
<https://doi.org/10.1145/3489517.3530481>,
<https://doi.org/10.1145/3489517.3530481>
45. Wenzel, M., Paulson, L.C., Nipkow, T.: The Isabelle framework. In: Theorem Proving in Higher Order Logics: 21st International Conference, TPHOLs 2008, Montreal, Canada, August 18–21, 2008. Proceedings 21. pp. 33–38. Springer (2008)
46. Xu, M., Fu, J., Mei, J., Deng, Y.: Model checking QCTL plus on quantum Markov chains. *Theor. Comput. Sci.* **913**, 43–72 (2022).
<https://doi.org/10.1016/j.tcs.2022.01.044>,
<https://doi.org/10.1016/j.tcs.2022.01.044>
47. Xu, M., Li, Z., Padon, O., Lin, S., Pointing, J., Hirth, A., Ma, H., Palsberg, J., Aiken, A., Acar, U.A., et al.: Quartz: superoptimization of quantum circuits. In: Proceedings of the 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation. pp. 625–640 (2022)
48. Ying, M.: Floyd-Hoare logic for quantum programs. *ACM Transactions on Programming Languages and Systems (TOPLAS)* **33**(6), 1–49 (2012)
49. Yu, N., Palsberg, J.: Quantum abstract interpretation. In: Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation. pp. 542–558 (2021)
50. Zhou, L., Barthe, G., Strub, P.Y., Liu, J., Ying, M.: Coqq: Foundational verification of quantum programs. *Proceedings of the ACM on Programming Languages* **7**(POPL), 833–865 (2023)
51. Zhou, L., Yu, N., Ying, M.: An applied quantum Hoare logic. In: Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation. pp. 1149–1162 (2019)
52. Zhu, S., Hung, S.H., Chakrabarti, S., Wu, X.: On the principles of differentiable quantum programming languages. In: Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation. p. 272–285. PLDI 2020, Association for Computing Machinery, New York, NY, USA (2020).
<https://doi.org/10.1145/3385412.3386011>,
<https://doi.org/10.1145/3385412.3386011>