

A Framework for Debugging Quantum Programs

Damian Rovara*

Lukas Burgholzer*[†]

Robert Wille*^{‡†}

*Chair for Design Automation, Technical University of Munich, Germany

[†]Munich Quantum Software Company GmbH, Garching near Munich, Germany

[‡]Software Competence Center Hagenberg GmbH (SCCH), Austria

damian.rovara@tum.de

lukas.burgholzer@tum.de

robert.wille@tum.de

<https://www.cda.cit.tum.de/research/quantum>

Abstract—Recent advancements in quantum computing software are gradually increasing the scope and size of quantum programs being developed. At the same time, however, these larger programs provide more possibilities for functional errors that are harder to detect and resolve. Meanwhile, debugging tools that could aid developers in resolving these errors are still barely existent and far from what we take for granted in classical design automation and software engineering. As a result, even if one manages to identify the incorrect behavior of a developed quantum program, detecting and resolving the underlying errors in the program remains a time-consuming and tedious task. Moreover, the exponential growth of the state space in quantum programs makes the efficient manual investigation of errors radically difficult even for respectively simple algorithms, and almost impossible as the number of qubits increases. To address this problem, this work proposes a debugging framework, available as an open-source implementation at <https://github.com/cda-tum/mqt-debugger>. It assists developers in debugging errors in quantum programs, allowing them to efficiently identify the existence of errors and diagnose their causes. Users are given the ability to place assertions in the code that test for the correctness of a given algorithm and are evaluated using classical simulations of the underlying quantum program. Once an assertion fails, the proposed framework employs different diagnostic methods to point towards possible error causes. This way, the debugging workload for quantum programs is drastically reduced.

I. INTRODUCTION

Due to the complexity of quantum computing, fully understanding larger quantum algorithms and the impact of all individual operations within them can be a challenging task, even for domain experts. Because of this, identifying errors and finding their underlying root causes often requires a large effort. Operations rarely commute, the exponential state space requires developers to keep track of a large number of amplitudes, and consequences of incorrect operations may sometimes only materialize at a much later step of execution.

Because of this, strategies to assist developers in checking and debugging quantum programs are becoming a rapidly growing area of research [1]–[7]. Here, proven methods from classical debugging, such as interactive debuggers [8], [9] and assertions [10]–[19], are employed to find and resolve errors in quantum programs efficiently. However, they are still limited in their scope and usability compared to their classical counterparts. More precisely, while the state of the art in quantum program debugging offers some promising methods to check quantum algorithms for their correctness, they still leave the task of identifying the underlying error to the developers—requiring further intensive work and extensive domain knowledge to resolve the problem.

To address this issue, this work proposes a comprehensive framework for debugging quantum programs, providing the following features:

- The framework employs *assertions* (as introduced in [10]–[19]) to check for the correctness of a given quantum program.

- The framework employs *classical simulation* to execute the quantum program and check the assertions. This way, the framework supports detailed insight into the system’s quantum state during execution, circumventing limitations encountered on physical quantum devices, such as the *no-cloning theorem* [20], that would otherwise greatly limit the possibilities of interaction with the current state.
- Finally, the framework employs automated *diagnosis techniques* to point towards possible error causes. Through continuous access to the full quantum state, the framework is able to analyze the program and search for anomalies that may be related to failing assertions. At the same time, the framework inspects dependencies of individual instructions and interactions of specific qubits, allowing it to reason about causes for misconfiguration by analyzing the corresponding dependency graph.

Through its modular nature, the proposed framework can support a variety of different simulation backends, allowing it to take advantage of various optimized simulators to debug larger quantum circuits. While there is an inevitable limit to the sizes of circuits that can be simulated, this approach can be employed to test circuits on a smaller scale and extrapolate their correctness for larger instances. The framework is fully open-source and available at <https://github.com/cda-tum/mqt-debugger> as part of the Munich Quantum Toolkit (MQT) [21].

The remainder of this work is structured as follows: Section II reviews core concepts and related work in the fields of quantum computing and debugging. Section III then illustrates the considered problem and devises a general idea to mitigate it. Following this, Section IV proposes a framework for efficient debugging of quantum programs based on this general idea, highlighting its main features, including the different types of supported assertions and its employed diagnosis methods. Section V then evaluates the applicability and limitations of this framework for larger problem instances. Finally, Section VI suggests possible future improvements and concludes this work.

II. BACKGROUND

This section describes existing work in the area of classical debugging and related techniques for quantum programs.

A. Classical Debugging

Debugging is a crucial part of the classical development workflow. Interactive debugging tools, such as GDB [22], are commonly employed by developers to efficiently find and resolve errors in programs. To make such debugging tools more accessible, the *Debug Adapter Protocol* [23] specifies a standard for interactive debugging software.

Since these software tools still require manual work from developers, automated debugging and analysis approaches are

another popular area of research that attempt to infer potential problems in software without active user interaction. In many cases, they are based on *static analysis* [24], examining programs without actively running the code and instead inspecting the individual instructions and their relations on a general level. In contrast, dynamic analysis [25] focuses on analyzing the properties of programs while executing them.

Automated debugging methods typically attempt to reduce the number of lines that need to be inspected to allow efficient error searches even in large programs. One way to achieve this is through *program slicing* [26], a method used to reduce a program by inspecting its control and data flow. In 1999, Zeller [27] introduced *delta debugging*, an efficient method of finding error causes by inspecting changes between a correct original implementation and an update that introduces an error.

The concept of debugging can also be applied to classical circuits. Automated circuit verification is a commonly employed method to test circuits for their correctness [28]–[30] and several methods have been proposed to automatically find error causes in different types of circuits [31]–[37] once verification has discovered the presence of errors.

B. Assertions and Debugging of Quantum Programs

Assertions are a powerful tool in the process of debugging and testing software [38]. They are represented by expressions that can be embedded in software code. During runtime, these expressions are evaluated and, depending on the evaluation results, the execution is aborted, and errors are reported.

In the context of quantum computing, several approaches to define quantum-state assertions for verification and runtime testing on real devices have been proposed [10]–[19]. When evaluating assertions on quantum devices, however, additional care must be taken to prevent them from impacting the system’s state. Measurements required to extract the state of qubits may cause side effects, collapsing parts of the system.

To circumvent this problem, these approaches propose methods that leave the state of the system unchanged at the cost of increasing the complexity of the circuit—by introducing ancilla qubits [17], [18] or performing projective measurements [16]. These trade-offs increase the depth and width of the corresponding circuits, diminishing the algorithms’ performance, especially on *Noisy Intermediate-Scale Quantum* (NISQ) devices. Huang et al. [15] proposed a framework for statistical assertions that measures the probability distribution of the current state and, then, terminates the execution, comparing the measured results with the expected values.

Different types of assertions have been proposed that allow developers to test a program for several different properties. Examples include precise and approximate assertions on the system’s full state or its individual amplitudes [17], [18], assertions testing for entanglement or superposition properties of the system [15], [19], as well as assertions that require the current state to be contained in the span of a set of state vectors [12].

Despite the advantages of these quantum assertions, their deployment still requires manual work from experts in the field. Developers must place assertions manually, which is a time-consuming task that requires a deep understanding of the full algorithm. Moreover, even if assertions detect an issue, identifying the precise location of the underlying cause is not trivial and often requires further manual work from experts to find it. Witharana et al. [19] propose *quAssert*, a framework that employs static analysis to extract locations for

efficient assertions from a quantum circuit definition. However, this approach requires the existence of a “ground truth” circuit functioning without errors to be employed. Similarly, Li et al. [2] propose a method to compute fitting locations for projective measurements to monitor the correctness of a given operator. Both of these approaches, however, only assist in the placement of assertions, not in the task of error location.

At the same time, alternative debugging approaches that do not directly employ assertions have been gaining more traction recently. Taxonomies of bugs in quantum programs [39], [40] are being collected to establish a common understanding for frequently encountered issues. Verification methods [10], [41] allow developers to compare updated programs, such as results of optimizations, with existing ground truth programs. Recently, Ramalho et al. [42] have proposed a roadmap for possible advances in the testing and debugging of quantum programs until 2030, suggesting interesting areas of research, proposing solutions for existing issues and reviewing previous work in the field.

III. MOTIVATION AND GENERAL IDEA

This section further highlights shortcomings and challenges in developing and debugging quantum algorithms. Based on this, it then devises a general idea to reduce the required workload and assist developers in the debugging process.

A. Considered Problem

Even with the help of existing debugging tools for quantum programs, locating error causes remains a challenging task. This process is further complicated by the increased availability of quantum software tools for the automated design of quantum computing applications (such as [43]–[50]), as well as the ongoing development of quantum programming languages (such as [51]–[55]). These developments allow users to create increasingly complex quantum programs and, hence, lead to a rapidly increasing potential for errors. While existing evaluations show that many near-term algorithms only require quantum programs of small sizes [56], adequately supporting the development of quantum programs also requires assistance in the research of new large-scale approaches. Due to the more complex nature of quantum algorithms compared to their classical counterparts, finding errors in such situations can be deceptively difficult, even in comparatively simple circuits.

Example 1. Consider the instance of Grover’s algorithm [57] with the target state $|111\rangle$ as defined in Listing 1, which will act as a running example throughout this paper. After two Grover iterations, one would expect to measure $|111\rangle$ with a high probability. However, this implementation is incorrect. Moreover, even for this considerably small quantum algorithm composed of only 23 lines of code, it is not easy to find the errors at first glance.

The first error is located in the `ccz` instruction in Line 2. Clearly, the oracle function, responsible for marking the target state $|111\rangle$ should take all three qubits of the data register q into account, not just $q1$ and $q2$. To resolve this issue, Line 2 should be replaced by `cccz q0, q1, q2, flag`.

The second error is hidden on a deeper level, as it is not caused by an incorrect configuration but by the absence of a gate. For Grover’s algorithm to work correctly, the quantum register q has to be put into an equal superposition using Hadamard gates (H q_i) before Line 17. Only after these gates are added will the algorithm yield the expected result.

```

1 gate oracle q0, q1, q2, flag {
2   ccz q1, q2, flag;
3 }
4
5 gate diffusion q0, q1, q2 {
6   h q0; h q1; h q2;
7   x q0; x q1; x q2;
8   ccz q0, q1, q2;
9   x q2; x q1; x q0;
10  h q2; h q1; h q0;
11 }
12
13 qreg q[3]; qreg flag[1]; creg c[3];
14
15 x flag;
16 // missing: h q; h flag;
17 oracle q[0], q[1], q[2], flag;
18 diffusion q[0], q[1], q[2];
19
20 oracle q[0], q[1], q[2], flag;
21 diffusion q[0], q[1], q[2];
22
23 measure q -> c;

```

Listing 1. An incorrect implementation of Grover’s algorithm using 4 qubits. This instance of the algorithm is supposed to search for the state $|111\rangle$.

Both of the issues in the above example are frequently occurring problems in the implementation of quantum algorithms. The misconfiguration of controlled quantum gates, e.g., by switching or missing individual controls, can often go unnoticed, especially in larger code bases containing multiple such gates. Missing individual gates, on the other hand, may also lead to larger issues. This is even more so the case, when the consequences are only encountered in different parts of the code, such as in custom gate definitions. These custom gates often require certain pre-conditions to function—effectively forcing developers to search the entire code base for possible sources of errors.

As a result, manually inspecting quantum algorithms for possible errors is a difficult and time-consuming task that requires an exact mathematical understanding of the entire system at each point in time. Furthermore, as the state space grows exponentially with respect to the number of qubits, the limitations of manual processing are easily reached, even for small examples. While it may still be possible to check the above example by hand, adding just four more qubits increases the number of amplitudes that have to be considered to 256. Therefore, automated tools that support developers in the search for error causes are crucial for real-world applications.

B. General Idea

To assist developers in this process, this work introduces a framework to aid in the identification and resolution of errors in quantum programs. Its methods are based on the *classical simulation* of quantum circuits. Compared to execution on physical quantum devices, this allows it to avoid physical limitations of quantum computing, such as the *no-cloning theorem* [20] and provide continuous access to the full quantum state of the system during each step of the execution.

In order to test the current state of the system during execution, the proposed framework employs *assertions*. These assertions are then examined during the simulation process, and their results are used to perform automated diagnosis on the program.

While the placement of assertions still needs to be conducted manually, the proposed framework supports developers by providing diagnosis methods once an assertion has failed via automatic analysis passes. By investigating the previous

interactions of qubits involved in a failed assertion and analyzing the state of the system during the simulation, the proposed solution *automatically* infers possible causes for the error. While finding the exact location of errors may not always be feasible or even possible, the framework takes advantage of the reported results to minimize the number of instructions that need to be manually inspected by the developer. To demonstrate the benefits of this approach, the ideas are illustrated using the running example again: In some cases, this effort can be further reduced by incorporating methods for the automatic generation of assertions [19], [58]. While these methods may not be applicable to all cases, they can still lead to substantial reductions in the manual debugging effort and the required domain knowledge to verify quantum programs.

```

1 gate oracle q0, q1, q2, flag {
2   ⚡ assert-sup q0, q1, q2;
3   ccz q1, q2, flag;
4   ⚡ assert-ent q0, q1, q2;
5 }
6
7 gate diffusion q0, q1, q2 {
8   h q0; h q1; h q2;
9   x q0; x q1; x q2;
10  ccz q0, q1, q2;
11  x q2; x q1; x q0;
12  h q2; h q1; h q0;
13 }
14
15 qreg q[3]; qreg flag[1]; creg c[3];
16
17 x flag;
18
19 oracle q[0], q[1], q[2], flag;
20 diffusion q[0], q[1], q[2];
21 ⚡ assert-eq 0.8, q { 0, 0, 0, 0, 0, 0, 0, 1 }
22
23 oracle q[0], q[1], q[2], flag;
24 diffusion q[0], q[1], q[2];
25 ⚡ assert-eq 0.9, q { 0, 0, 0, 0, 0, 0, 0, 1 }
26
27 measure q -> c;

```

Listing 2. The incorrect implementation of Grover’s algorithm from Listing 1, including new quantum-state *assertions* and suggested problem causes. Here, red wavy underlines indicate possible error locations, and yellow lightning symbols highlight failing assertions.

Example 2. Returning to the instance of Grover’s algorithm considered in Listing 1, we place assertions on Lines 2, 4, 21, and 25, as illustrated in Listing 2. The assertion on Line 2 is used as a pre-condition to enforce that the given qubits are in a superposition state. This ensures that the oracle is not used on invalid input states. The assertion on Line 4, on the other hand, checks whether the oracle function itself acts as expected. Finally, the assertions on Lines 21 and 25 test the functionality of the overall circuit.

Simulating the algorithm while evaluating these assertions and employing the corresponding possible diagnosis methods (detailed more thoroughly later in Section IV) reveals a potential error in Line 3. In particular, static analysis reveals the absence of interactions between qubit q_0 and the remainder of the data register q , which prevents entanglement. Furthermore, the failure of the pre-condition assertion on Line 2 and subsequent analysis at runtime indicate the possibility of a missing gate before Line 19. This way, instead of 23 lines, the developer only has to inspect 4 lines to resolve all errors (all of them explicitly highlighted by the proposed framework).

In the following, the proposed approach is described in more detail.

IV. PROPOSED FRAMEWORK

In this section, we describe the framework introduced above. This includes a more detailed definition of the required inputs, as well as all the analysis methods employed for diagnosis, including a *cone of influence analysis*, an *interaction analysis*, and a *control value analysis*. For additional information, we refer to its open-source implementation, available at <https://github.com/cda-tum/mqt-debugger>.

A. Input Specification

The proposed framework assumes the specification of a quantum program (currently based on the OpenQASM [59] standard), which includes manually placed *assertions* to test for its correctness. The assertions are assumed to be expressed as pre-defined OpenQASM-style instructions. At the moment, the proposed framework supports three types of assertions on quantum states:

- *entanglement-assertions* (`assert-ent`) test, whether a given set of qubits is fully entangled,
- *superposition-assertions* (`assert-sup`) test, whether the given set of qubits is in a superposition state, and
- *equality-assertions* (`assert-eq`) test, whether the quantum state of a system is equal to the provided reference state (up to some similarity margin).

These types of assertions can only be evaluated on simulated quantum devices, as they require perfect insight into the underlying quantum state. On real quantum computers, access to the state is only limited, rendering the required checks infeasible without sacrificing accuracy.

Example 3. *The quantum program in Listing 2 includes four assertions on Lines 2, 4, 21, and 25. The superposition assertion on Line 2 enforces that the three qubits of the data register are in some superposition state. At this point, we expect them to be in the state $\frac{1}{2\sqrt{2}} \sum_{x=0}^7 |x\rangle$, which clearly constitutes a superposition. However, due to the incorrect initialization, their state will instead be $|000\rangle$ at the first execution, causing the assertion to fail.*

Line 4, on the other hand, requires that the data register is fully entangled, i.e., non-separable. However, as the data register is still in the state $|000\rangle$ at its first execution, the assertion will fail since it can trivially be separated into $|0\rangle \otimes |0\rangle \otimes |0\rangle$.

Finally, the equality assertions on Lines 21 and 25 both test whether the data register q equals the target state $|111\rangle$. The assertion on Line 21 requires a similarity of 0.8, while the assertion on Line 25 requires a similarity of 0.9. With more iterations of the Grover circuit, we expect q to converge toward the target state $|111\rangle$. However, due to the errors in the program, Grover's algorithm is not performed correctly and both assertions fail.

The failures of any of these assertions provide valuable information on the underlying errors in the code. To leverage this information, the proposed framework provides a set of diagnosis methods, that can be used to locate the corresponding error causes. They are described in the following.

B. Cone of Influence Analysis

By statically analyzing the data dependencies of a failed assertion, the program's instructions can be partitioned into two sets, based on whether an instruction influences any of the qubits contained in the failed assertion (also called the *cone of influence*) or not. These results can then be employed in several

ways: By highlighting relevant instructions in an integrated development environment, users can manually skim through the smaller subset of instructions, searching for incorrect instructions inside the cone of influence, or for instructions that are missing from it. Furthermore, this cone of influence analysis can also be taken as the basis for further automated diagnosis methods.

Example 4. *Returning to Listing 2, each of the failing assertions acts solely on the data register q . The cone of influence analysis will highlight all lines above the corresponding assertion that directly or indirectly affect q . For the assertion on Line 2, this includes Lines 1, 15, and 19. This allows the developer, as well as the remaining analysis steps, to investigate just these four lines for errors, or, alternatively, conclude that a crucial gate may be missing.*

As the execution progresses, the number of gates inside the cone of influence tends to increase. The assertion on Line 4 includes Lines 3 and 17 in its cone of influence, while the cone of influence for the assertion on Line 21 includes every instruction from Line 1 to Line 20. The proposed framework highlights the lines included in the cone of influence automatically, aiding the developer in inspecting the code.

The applicability of this method strongly depends on the nature of a given quantum program. Algorithms that require a large number of interactions between qubits can easily cause a majority of the instructions to be included in a given cone of influence, especially at later steps of execution. However, as discussed later in Section V, certain types of algorithms can profit from this approach even at larger sizes, and any improvement of performance greatly increases debugging capabilities.

C. Interaction Analysis

For two qubits to be entangled with each other, they must interact directly or indirectly in some previous instructions. Interaction analysis searches for qubits used in an entanglement assertion that do not interact with each other in any previous instruction, reporting them to the user. While this method can only be employed specifically for entanglement assertions, it can be a helpful tool specifically for the error type of misconfigured controlled gates, easily detecting gates that are missing individual controls.

For this purpose, the interaction of two qubits x and y is defined recursively, such that x interacts with y if and only if there exists a gate that acts on both x and y , or if there exists a gate acting on x and some qubit z , where z interacts with y . To evaluate this, we perform a breadth-first search on the instructions included in the failed assertion's cone of influence, constructing an interaction graph for the different qubits, as illustrated in Fig. 1. After the first controlled gate, the interaction between q_1 and q_2 is added to the interaction graph. Then, the controlled gate between q_0 and q_1 adds an additional connection to the graph. As there now exists a path from q_0 to q_2 , the qubits are proven to interact with each other.

Example 5. *After encountering the failing entanglement assertion on Line 4 in Listing 2, the proposed framework performs an interaction analysis between the three data qubits. Iterating over their cone of influence, it finds an interaction between q_1 and q_2 in Line 3, but no interactions between q_0 and other data qubits. The full interaction graph, therefore, consists of only one edge between q_1 and q_2 , similarly to*




Fig. 1. The process of constructing interaction graphs by collecting gates acting on multiple qubits. Two qubits are said to interact if there exists a path between them on the interaction graph.

the state after the first gate has been considered in Fig. 1. As no path exists between q_0 and the other qubits, this approach concludes that a controlled gate interacting with q_0 might be missing, or that Line 3 may be missing a control qubit.

As deep circuits commonly re-use qubits and this definition of interaction does not allow for interactions to be canceled without resetting the corresponding qubits, this diagnosis method is once again only partially applicable to larger circuits. However, especially during early steps of the execution, this method still provides valuable assistance to the developer. Especially in programs that employ several custom gate definitions, it shields developers from being required to follow the execution through all the required jumps and points them easily toward possible errors.

D. Control Value Analysis

During execution, the proposed framework keeps track of all controlled gates that can only perform the identity operation. This happens if any of the control qubits used in a controlled gate is always zero. As the proposed framework has access to the state vector during the entire execution process, this dynamic analysis task can be applied easily. When encountering a controlled gate for which all control qubits are zero, the framework marks them as potential problems. Should a later execution of the same instruction no longer have this problem, then the mark will be removed. If a failing assertion is encountered, it will investigate whether any of the instructions within its cone of influence have been marked this way, and report those instructions. Once again, this method can easily help developers in finding misconfigured controlled gates. In particular, cases in which control and target have been swapped, as well as missing instruction during state preparation can often be found by this analysis method.

Example 6. When the `ccz` gate in Line 3 is first encountered, it will note that the control qubits are in state $|00\rangle$ at the time and mark this gate as a potential error candidate. Once any assertion that includes Line 3 in its cone of influence is encountered, the automatically performed control value analysis will highlight that the controls in Line 3 are always zero, and infer that a gate may be missing before Line 3 in the cone of influence. When the `oracle` gate is called a second time, the data register will have been modified by previous instructions. Because of this, the second execution of Line 3 will remove the mark from the instruction, and, consequently, the control value analysis employed after the failing assertion on Line 25 will not highlight this issue again.

This analysis method can also be employed efficiently on deep circuits. As controlled gates typically only act on a small number of qubits, the number of variables to be investigated for each gate remains limited. Furthermore, due to the crucial role controlled gates play in quantum circuits, this analysis

TABLE I
CHARACTERISTICS OF THE PROGRAMS USED FOR THIS EVALUATION.

program	dj-130	random-13
qubits	130	13
instructions	391	432

method may be employed to find error causes for all types of assertions and can find applications in all types of quantum programs.

V. APPLICATION AND EVALUATION

The methods proposed above have been implemented and made available as an open-source framework at <https://github.com/cda-tum/mqt-debugger>. The running example and discussions above already demonstrated its applicability and usefulness. However, to confirm the corresponding benefits also on larger instances, further tests and evaluations have been conducted. The corresponding results are summarized in the following.

Due to page limitation, we hereby focus on two representative instances (the availability as an open-source framework easily allows to try and test further instances), namely the programs `dj-130`, implementing the Deutsch-Jozsa algorithm [60] as well as `random-13`, representing a random circuit. The characteristics of the employed programs are further described in Table I. Both of these circuits were generated using the benchmark library MQT Bench [50] and were selected to properly illustrate the benefits (but also limitations) of the proposed methods. For all considered evaluations, the analysis methods were performed in real-time. Even as the program size increased, the algorithms efficiently performed their analysis tasks, with the primary performance bottleneck remaining the classical quantum circuit simulation. To this end, we have employed a dedicated simulator backend using decision diagrams as a data structure [61], [62].

In the following, we discuss the performance of the framework and, in particular, each analysis method individually, highlighting their advantages as well as possible reasons for their limitations in more detail.

1) *Cone of Influence Analysis:* For both programs, the proposed framework performed a cone of influence analysis over an equality assertion involving three qubits at the end of execution. Fig. 2 illustrates the partitions of the input program returned by this analysis. Each of the boxes represents the code of the corresponding program. Black sections represent lines of code that are outside of the cone of influence, whereas orange sections represent lines of code inside the cone of influence.

The results confirm the expectation that dedicated quantum algorithms exhibit significant local structure, which is advantageous for this type of analysis. Out of the 391 lines of the `dj-130` program, only 11 are included in the cone of influence, reducing the number of lines to be considered for debugging by more than 97%.

However, the situation is different for the `random-13` program. Given that deep random circuits typically involve pairwise interactions between all qubits, the cone of influence grows significantly. As a result, out of the 432 lines of `random-13`, 416 are included in the cone of influence, merely reducing the number of lines to be considered by 4%.

Clearly, the cone of influence analysis performs differently depending on the characteristics of the program. Particularly




Fig. 2. The partitions inside (orange) and outside (black) the cone of influence for an assertion over three qubits at the end of the evaluation programs.




Fig. 3. The interaction graphs for both evaluation programs, as computed by the proposed framework after 50% of their operations have been applied. For dj-130, qubits 32 to 128 have not yet interacted with any other qubits.

in deep circuits, the reduction of lines to be considered for locating errors is only minimal. However, in many near-term algorithms, circuits are desired to be shallow, providing promising applications for this method to help find errors.

2) *Interaction Analysis*: Interaction analysis was performed on both programs to understand the connectivity and interaction patterns among qubits. Performing this analysis at the end of the programs yields a strongly connected interaction graph in both cases, as all qubits either directly or indirectly interact with other qubits at some point. However, typically, this analysis is conducted on failing entanglement assertions placed at intermediate positions in the code to identify whether qubits that should be entangled are indeed interacting. Hence, Fig. 3 shows the interaction graphs for both programs computed after half of their instructions have been processed.

For the dj-130 program, the interaction graph exhibited a simple star topology, connecting only 33 of the 130 qubits. This simplified structure makes it easier to identify expected interactions and potential issues. For example, the failure of an assertion requiring qubits 128 and 129 to be entangled can be attributed to the missing interaction of the two qubits in the interaction graph. In contrast, the random-13 program also showed a strongly connected interaction graph halfway through its execution, reflecting the complexity and randomness of the circuit. In these cases, the analysis remains inconclusive for identifying possible causes of failing assertions.

Overall, the interaction analysis supports developers by allowing them to identify potentially missing interactions between qubits, such as missing controls or entirely missing gates.

3) *Control Value Analysis*: Control value analysis was conducted to evaluate the framework’s performance on large inputs. Given the number of qubits in the dj-130 program, accessing the state vector regularly for this analysis is impractical in this case. Therefore, we focused on the random-13 program, modifying a single, randomly chosen instruction 100 times and running the analysis to determine how often it provides a suggestion that allows the developer to identify the error. If the randomly chosen instruction was a controlled gate, we flipped the control and target qubits. If it was another type of instruction, it was removed.

The control value analysis managed to correctly suggest error locations in 22 out of 100 test cases. In particular, it triggered in 31% of the test instances when an instruction was removed, while it only triggered in 10% of the instances when a controlled gate was flipped. This discrepancy likely arises because removing gates, especially early in the program, can disrupt important state preparation steps, leaving qubits in the $|0\rangle$ state longer and allowing the analysis to trigger more frequently. Additionally, the analysis showed a 75.9% success rate in identifying errors introduced before Line 100. Errors introduced after Line 100 were not correctly identified in any test case, as the system’s state became too disordered due to the random nature of the circuit. At the same time, this approach also resulted in a large number of false positives. In total, running random-13 resulted in 21 suggested locations of potential errors that were not related to the introduced error. However, this is an expected issue of random circuits. We do not expect the same properties for programs that were not generated randomly, as developers have no reason to add controlled gates that do not affect the underlying state.

Overall, the control value analysis is especially helpful for identifying errors early in the program, pointing developers toward their possible causes.

VI. CONCLUSION

In this work, we proposed a framework that allows developers to efficiently debug quantum programs. Given a set of assertions, it simulates provided programs and checks the assertions for correctness. Upon encountering an unexpected state, it uses several automated analysis methods to diagnose the problem and find the underlying causes. Evaluations showed that, in many cases, these methods can be employed to efficiently diagnose problems in programs consisting of several hundred instructions. Future work involves devising further analysis methods and extending the workflow by using analysis results to suggest better locations for assertions when diagnosis fails. The proposed debugging framework is available as an open-source implementation at <https://github.com/cdatum/mqt-debugger>.

ACKNOWLEDGMENTS

This work received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement No. 101001318), was part of the Munich Quantum Valley, which is supported by the Bavarian state government with funds from the Hightech Agenda Bayern Plus, and has been supported by the BMWK on the basis of a decision by the German Bundestag through project QuaST, as well as by the BMK, BMDW, and the State of Upper Austria in the frame of the COMET program (managed by the FFG).

REFERENCES

- [1] O. Di Matteo, “On the need for effective tools for debugging quantum programs,” in *International Workshop on Quantum Software Engineering*, 2024.
- [2] Y. Li and M. Ying, “Debugging quantum processes using monitoring measurements,” *Physical Review A*, 2014.
- [3] Y. Huang and M. Martonosi, “Statistical assertions for validating patterns and finding bugs in quantum programs,” in *International Symposium on Computer Architecture*, 2019.
- [4] S. A. Metwalli and R. Van Meter, “A tool for debugging quantum circuits,” in *Int’l Conf. on Quantum Computing and Engineering*, 2022.
- [5] A. Miransky, L. Zhang, and J. Doliskani, *On testing and debugging quantum software*, 2021. arXiv: 2103.09172[quant-ph].
- [6] A. García de la Barrera, I. García-Rodríguez de Guzmán, M. Polo, and M. Piattini, “Quantum software testing: State of the art,” *Journal of Software: Evolution and Process*, 2023.
- [7] S. A. Metwalli and R. Van Meter, “Testing and debugging quantum circuits,” *IEEE Transactions on Quantum Engineering*, 2024.
- [8] QunTag, *Openqasm debugger extension*, <https://github.com/quantag/qasm-adapter-vscode>, 2024.
- [9] N. Jiang, Z. Wang, and J. Wang, *Debugging quantum programs using probabilistic quantum cloning*, 2023.
- [10] J. Liu, B. Zhan, S. Wang, S. Ying, T. Liu, Y. Li, M. Ying, and N. Zhan, “Formal verification of quantum algorithms using quantum Hoare logic,” in *Computer Aided Verification*, 2019.
- [11] T. Häner, T. Hoefler, and M. Troyer, “Assertion-based optimization of quantum programs,” *ACM on Programming Languages*, 2020.
- [12] N. Yu and J. Palsberg, “Quantum abstract interpretation,” in *International Conference on Programming Language Design and Implementation*, 2021.
- [13] M. Ying, *Birkhoff-von Neumann quantum logic as an assertion language for quantum programs*, 2022. arXiv: 2205.01959[quant-ph].
- [14] F. Bauer-Marquart, S. Leue, and C. Schilling, “symQV: Automated symbolic verification of quantum programs,” in 2023. arXiv: 2212.02267[quant-ph].
- [15] Y. Huang and M. Martonosi, “Statistical assertions for validating patterns and finding bugs in quantum programs,” in *International Symposium on Computer Architecture*, 2019.
- [16] G. Li, L. Zhou, N. Yu, Y. Ding, M. Ying, and Y. Xie, *Proj: Projection-based runtime assertions for debugging on a quantum computer*, 2020. arXiv: 1911.12855[quant-ph].
- [17] J. Liu, G. T. Byrd, and H. Zhou, “Quantum circuits for dynamic runtime assertions in quantum computation,” in *International Conference on Architectural Support for Programming Languages and Operating Systems*, 2020.
- [18] J. Liu and H. Zhou, “Systematic approaches for precise and approximate quantum state runtime assertion,” in *International Symposium on High-Performance Computer Architecture (HPCA)*, 2021.
- [19] H. Witharana, D. Volya, and P. Mishra, *quAssert: Automatic generation of quantum assertions*, 2023. arXiv: 2303.01487[quant-ph].
- [20] W. K. Wootters and W. H. Zurek, “A single quantum cannot be cloned,” *Nature*, 1982.
- [21] R. Wille, L. Berent, T. Forster, J. Kunasikaran, K. Mato, T. Peham, et al., “The MQT Handbook: A Summary of Design Automation Tools and Software for Quantum Computing,” in *Int’l Conf. on Quantum Software*, 2024. arXiv: 2405.17543[quant-ph].
- [22] R. Stallman, R. Pesch, S. Shebs, et al., “Debugging with gdb,” *Free Software Foundation*, 1988.
- [23] Microsoft, “Debug adapter protocol.” (2024).
- [24] N. Ayewah, D. Hovemeyer, J. D. Morgenthaler, and J. Penix, “Using static analysis to find bugs,” *IEEE Software*, 2008.
- [25] T. Ball, “The concept of dynamic analysis,” *ACM SIGSOFT Software Engineering Notes*, 1999.
- [26] M. Weiser, “Program slicing,” *IEEE Transactions on Software Engineering*, 1984.
- [27] A. Zeller, “Yesterday, my program worked. today, it does not. why?” *ACM SIGSOFT Software Engineering Notes*, 1999.
- [28] Bryant, “Graph-based algorithms for boolean function manipulation,” *IEEE Trans. Comput.*, 1986.
- [29] A. Kuehlmann, V. Paruthi, F. Krohm, and M. Ganai, “Robust boolean reasoning for equivalence checking and functional property verification,” *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, 2003.
- [30] P. Molitor and J. Mohnke, *Equivalence checking of digital circuits: Fundamentals, principles, methods*. Springer, 2010.
- [31] A. Veneris and I. N. Hajj, “Design error diagnosis and correction via test vector simulation,” *IEEE Trans. on CAD of Integrated Circuits and Systems*, 1999.
- [32] A. Smith, A. Veneris, M. F. Ali, and A. Viglas, “Fault diagnosis and logic debugging using Boolean satisfiability,” *IEEE Trans. on CAD of Integrated Circuits and Systems*, 2005.
- [33] M. F. Ali, S. Safarpour, A. Veneris, M. S. Abadir, and R. Drechsler, “Post-verification debugging of hierarchical designs,” in *Int’l Conf. on CAD*, 2005.
- [34] R. Wille, D. Große, S. Frehse, G. W. Dueck, and R. Drechsler, “Debugging reversible circuits,” *Integration, the VLSI Journal*, 2011.
- [35] E. Strassnick, M. Agrawala, and S. Follmer, “Coupling simulation and hardware for interactive circuit debugging,” in *Conference on Human Factors in Computing Systems*, 2021.
- [36] R. Wille, M. Soeken, and R. Drechsler, “Debugging of inconsistent UML/OCL models,” in *Design, Automation and Test in Europe*, 2012.
- [37] D. Grosse, R. Wille, R. Siegmund, and R. Drechsler, “Contradiction analysis for constraint-based random simulation,” in *Forum on Specification and Design Languages*, 2008.
- [38] D. Rosenblum, “A practical approach to programming with assertions,” *IEEE Transactions on Software Engineering*, 1995.
- [39] P. Zhao, J. Zhao, Z. Miao, and S. Lan, “Bugs4Q: A Benchmark of Real Bugs for Quantum Programs,” in *IEEE/ACM International Conference on Automated Software Engineering (ASE)*, 2021.
- [40] Y. Huang and M. Martonosi, “QDB: From Quantum Algorithms Towards Correct Quantum Programs,” *OASlcs, Volume 67, PLATEAU 2018*, 2019. arXiv: 1811.05447[quant-ph].
- [41] L. Burgholzer and R. Wille, “Advanced equivalence checking for quantum circuits,” *IEEE Trans. on CAD of Integrated Circuits and Systems*, 2021.
- [42] N. C. Leite Ramalho, H. Amario de Souza, and M. Lordello Chaim, “Testing and Debugging Quantum Programs: The Road to 2030,” *ACM Trans. Softw. Eng. Methodol.*, 2025.
- [43] Qiskit contributors, *Qiskit: An open-source framework for quantum computing*, 2023.
- [44] V. Bergholm, J. Izaac, M. Schuld, C. Gogolin, S. Ahmed, V. Ajith, et al., “PennyLane: Automatic differentiation of hybrid quantum-classical computations,” 2018. arXiv: 1811.04968[quant-ph].
- [45] N. Quetschlich, L. Burgholzer, and R. Wille, “Towards an automated framework for realizing quantum computing solutions,” in *Int’l Symp. on Multi-Valued Logic*, 2023. arXiv: 2210.14928[quant-ph].
- [46] O. D. Matteo, J. Izaac, T. R. Bromley, A. Hayes, C. Lee, M. Schuld, A. Száva, C. Roberts, and N. Killoran, “Quantum computing with differentiable quantum transforms,” *ACM Transactions on Quantum Computing*, 2023.
- [47] D. Rovara, N. Quetschlich, and R. Wille, *A framework to formulate pathfinding problems for quantum computing*, 2024. arXiv: 2404.10820[quant-ph].
- [48] D. Volpe, N. Quetschlich, M. Graziano, G. Turvani, and R. Wille, “Towards an Automatic Framework for Solving Optimization Problems with Quantum Computers,” in *IEEE International Conference on Quantum Software (QSW)*, 2024. eprint: 2406.12840[quant-ph].
- [49] B. Apak, M. Bandic, A. Sarkar, and S. Feld, “Ketgpt-dataset augmentation of quantum circuits using transformers,” in *International Conference on Computational Science*, 2024.
- [50] N. Quetschlich, L. Burgholzer, and R. Wille, “MQT Bench: Benchmarking Software and Design Automation Tools for Quantum Computing,” *Quantum*, 2023. arXiv: 2204.13719[quant-ph].
- [51] R. Seidel, S. Bock, R. Zander, M. Petrič, N. Steinmann, N. Tcholtchev, and M. Hauswirth, *Qrisp: A framework for compilable high-level programming of gate-based quantum computers*, 2024. arXiv: 2406.14792[quant-ph].
- [52] A. Cross, A. Javadi-Abhari, T. Alexander, N. De Beaudrap, L. S. Bishop, S. Heide, et al., “Openqasm 3: A broader and deeper quantum assembly language,” *ACM Transactions on Quantum Computing*, 2022.
- [53] B. Bichsel, M. Baader, T. Gehr, and M. Vechev, “Silq: A high-level quantum language with safe uncomputation and intuitive semantics,” in *Conference on Programming Language Design and Implementation*, 2020.
- [54] R. S. Smith, M. J. Curtis, and W. J. Zeng, *A practical quantum instruction set architecture*, 2017. arXiv: 1608.03355[quant-ph].
- [55] Microsoft, *Azure quantum development kit*, <https://github.com/microsoft/qsharp>, 2024.
- [56] M. Schönberger, M. Franz, S. Scherzinger, and W. Mauerer, “Peel | Pile? Cross-Framework Portability of Quantum Software,” in *19th International Conference on Software Architecture Companion (ICSA-C)*, 2022.
- [57] L. K. Grover, “A fast quantum mechanical algorithm for database search,” in *ACM symposium on Theory of computing*, 1996.
- [58] D. Rovara, L. Burgholzer, and R. Wille, *Automatically refining assertions for efficient debugging of quantum programs*, 2024. arXiv: 2412.14252[quant-ph].
- [59] A. W. Cross, L. S. Bishop, J. A. Smolin, and J. M. Gambetta, *Open quantum assembly language*, 2017. arXiv: 1707.03429[quant-ph].
- [60] D. Deutsch and R. Jozsa, “Rapid solution of problems by quantum computation,” *Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences*, 1992.
- [61] A. Zulehner and R. Wille, “Advanced simulation of quantum computations,” *IEEE Trans. on CAD of Integrated Circuits and Systems*, 2019.
- [62] R. Wille, S. Hillmich, and B. Lukas, “Tools for quantum computing based on decision diagrams,” *ACM Transactions on Quantum Computing*, 2022.