

# Lawrence Berkeley National Laboratory

## LBL Publications

### Title

Test Points for Online Monitoring of Quantum Circuits

### Permalink

<https://escholarship.org/uc/item/3qh3v75z>

### Journal

ACM Journal on Emerging Technologies in Computing Systems, 18(1)

### ISSN

1550-4832

### Authors

Acharya, Nikita  
Urbanek, Miroslav  
De Jong, Wibe A  
[et al.](#)

### Publication Date

2022-01-31

### DOI

10.1145/3477928

Peer reviewed

# Test Points for Online Monitoring of Quantum Circuits

NIKITA ACHARYA, City College of New York, City University of New York, USA

MIROSLAV URBANEK, Computational Research Division, Lawrence Berkeley National Laboratory, USA

WIBE A. DE JONG, Computational Research Division, Lawrence Berkeley National Laboratory, USA

SAMAH MOHAMED SAEED, City College of New York, City University of New York, USA

Noisy Intermediate-Scale Quantum (NISQ) computers consisting of tens of inherently noisy quantum bits (qubits) suffer from reliability problems. Qubits and their gates are susceptible to various types of errors. Due to limited numbers of qubits and high error rates, quantum error correction cannot be applied. Physical constraints of quantum hardware including the error rates are used to guide the design and the layout of quantum circuits. The error rates determine the selection of qubits and their operations. The resulting circuit is executed on the quantum computer.

This study explores the risk of unexpected changes in the error rates of NISQ computers post-calibration. We show that unexpected changes in error rates can alter the output state of a quantum circuit. To detect these changes, we propose the insertion of test points into the quantum circuit to enable an online monitoring of the physical qubit behavior. We utilize classical, superposition, and uncompute test points. Furthermore, we use a gate error coverage metric to assess the quality of the tests. We verify the effectiveness of the proposed scheme on different IBM quantum computers (IBM Q), in addition to a noisy simulation that shows the scalability of the proposed approach.

CCS Concepts: • **Computer systems organization** → **Quantum computing**;

Additional Key Words and Phrases: Quantum circuit, quantum circuit test point, Noisy Intermediate-Scale Quantum (NISQ) computer, side-channel information, mapping, reliability, security

## ACM Reference Format:

Nikita Acharya, Miroslav Urbanek, Wibe A. de Jong, and Samah Mohamed Saeed. 2021. Test Points for Online Monitoring of Quantum Circuits. *ACM J. Emerg. Technol. Comput. Syst.* 1, 1, Article 1 (January 2021), 18 pages. <https://doi.org/10.1145/3477928>

## 1 INTRODUCTION

Quantum computing is among the most significant advancements of the 21st century, with the potential to solve many important problems ranging from science to finance. Near-term quantum computers, referred to as Noisy Intermediate-Scale Quantum (NISQ) computers, are expected to work as accelerators for solving various problems including data analytics, optimization, and material and drug discovery. While NISQ computers are very promising, they are fragile. Quantum systems consist of error-prone quantum bits (qubits). Their error rates vary from one qubit to another and change over time. Obtaining accurate results from a NISQ device is challenging in the absence of quantum error correction due to the limited number of qubits and their high error rates.

---

Authors' addresses: Nikita Acharya, City College of New York, City University of New York, New York, NY, USA, [nachary000@citymail.cuny.edu](mailto:nachary000@citymail.cuny.edu); Miroslav Urbanek, Computational Research Division, Lawrence Berkeley National Laboratory, Berkeley, CA, USA, [urbanek@lbl.gov](mailto:urbanek@lbl.gov); Wibe A. de Jong, Computational Research Division, Lawrence Berkeley National Laboratory, Berkeley, CA, USA, [wadejong@lbl.gov](mailto:wadejong@lbl.gov); Samah Mohamed Saeed, City College of New York, City University of New York, New York, NY, USA, [ssaeed@ccny.cuny.edu](mailto:ssaeed@ccny.cuny.edu).

---

ACM acknowledges that this contribution was authored or co-authored by an employee, contractor, or affiliate of the United States government. As such, the United States government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for government purposes only.

© 2021 Association for Computing Machinery.

1550-4832/2021/1-ART1 \$15.00

<https://doi.org/10.1145/3477928>

To reduce the impact of noise on quantum circuits executed on a NISQ computer, a noise model of the quantum hardware should be incorporated into the process of designing a quantum circuit. The noise model can be characterized using tomography [40] or randomized benchmarking [26] during the hardware calibration process. The measured error rates along with other physical characteristics of the quantum device are provided as inputs to a quantum compiler to minimize the circuit error rates. The quantum compilation process includes gate decomposition, optimization, and mapping of the logical circuit to a physical circuit, which consists of elementary quantum gates of the target hardware. On one hand, the noise characterization during the calibration process may fail to capture errors in real applications. On the other hand, any changes to the expected qubit behavior (error rates) can alter the physical circuit structure, and thus, significantly affect the reliability of the quantum circuit output.

In this work<sup>1</sup>, we show the implications of the variation of the qubit behavior post-calibration on the output of a quantum circuit. We consider both natural error drifts in the quantum hardware as well as malicious enforcement of artificial error rates. We propose an approach to detect changes in the qubit behavior using test points injected into a quantum circuit. We also propose a gate coverage metric to show the quality of the generated quantum test circuits. Experimental evaluation confirms the effectiveness of the proposed approach for a variety of IBM NISQ computers. The contributions of this paper are as follows:

- (1) We study the impact of qubit allocation on the fidelity of the quantum circuit output in the presence of variable error rates.
- (2) We experimentally demonstrate the unexpected changes in the qubit behavior post-calibration.
- (3) We propose an approach to detect dynamic changes in the qubit behavior using a limited number of quantum circuits based on test point insertion and develop a gate coverage metric to assess the quality of the tests.
- (4) We validate the effectiveness and the scalability of the proposed approach using various quantum circuits executed on several superconducting quantum devices provided by IBM, in addition to a noisy simulator used for larger quantum circuits.

The remainder of this paper is organized as follows: Section 2 provides a background on quantum circuit design and compilation. Section 3 describes related works on applying dynamic information for quantum circuit debugging and compilation. Section 4 motivates the proposed work. In Section 5, we present our proposed scheme. In section 6, we conduct a variety of experiments that show the effectiveness of the proposed method. In section 7, we discuss different applications of our proposed scheme. We conclude the paper in Section 8.

## 2 BACKGROUND

### 2.1 Quantum circuits

A quantum circuit is a collection of quantum gates that operate on qubits. We use bra–ket notation [31] to denote the state of qubits. A state of a single qubit can be described by its state vector  $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ , where  $|0\rangle$  and  $|1\rangle$  are the basis states, and  $\alpha$  and  $\beta$  are complex coefficients such that  $|\alpha|^2 + |\beta|^2 = 1$ . Similarly, an  $n$ -qubit state can be in a superposition state of  $2^n$  basis states and described by a state vector  $|\psi\rangle = \sum_{x \in \{0,1\}^n} \alpha_x |x\rangle$ , where  $\alpha_x$  are complex numbers such that  $\sum_{x \in \{0,1\}^n} |\alpha_x|^2 = 1$ . Measurement of the system produces only a single outcome. The probability of measuring outcome  $|x\rangle$  is given by amplitude  $|\alpha_x|^2$ .

The action of a quantum gate on qubits can be represented as a linear transformation of the state vector. A quantum gate is described by a unitary matrix  $U$ , i.e., a complex matrix whose conjugate

<sup>1</sup>A preliminary version of this paper has appeared at IEEE/ACM International Conference On Computer Aided Design [5].




Fig. 1. CU gate decomposition using the KAK decomposition [12] in which A, B, and C represent different single-qubit gates.

transpose  $U^\dagger$  is also its inverse ( $UU^\dagger = U^\dagger U = I$ , where  $I$  is an identity matrix). For  $n$  qubits,  $U$  is a unitary  $2^n \times 2^n$  matrix. Universal quantum computers support single-qubit and multi-qubits gates. A single-qubit gate can be generalized as a unitary operation

$$U3(\theta, \phi, \lambda) = \begin{pmatrix} \cos \frac{\theta}{2} & -e^{i\lambda} \sin \frac{\theta}{2} \\ e^{i\phi} \sin \frac{\theta}{2} & e^{i(\phi+\lambda)} \cos \frac{\theta}{2} \end{pmatrix},$$

where  $\theta$ ,  $\phi$ , and  $\lambda$  are control parameters (rotation angles). Examples of single-qubit gates include the bit-flip gate  $X = U3(\pi, 0, \pi)$ , the Hadamard gate  $H = U3(\pi/2, 0, \pi)$  that can generate a superposition of basis states, and the phase-shift-by- $\pi/4$  gate  $T = U3(0, 0, \pi/4)$ . A quantum gate that operates on two qubits can create entanglement between them. An example of a two-qubit gate is the controlled NOT (CNOT) gate, which operates on a target and a control qubit. The CNOT gate acts as the XOR gate, which inverts the target qubit if the control qubit is set to one. Otherwise, the target qubit remains unchanged. The matrix representation of the CNOT gate is given by

$$\text{CNOT} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}.$$

Quantum circuits can contain complex quantum operations including SWAP, Toffoli, and arbitrary gates. SWAP is a two-qubit operation that switches states of two qubits. Toffoli gate is an example of an  $n$ -qubit gate, which consists of  $n - 1$  control qubits and a single target qubit. Arbitrary gates include multi-qubit gates such as the controlled U gate (CU), which applies the U gate to the target qubit only if the control qubit is set to one.

The measurement operation, which collapses a superposition state, determines the output of a quantum circuit. On NISQ computers, a quantum circuit is executed in a large number of trials. The resulting output distribution is used to determine the output of the algorithm. The Probability of Successful Trial (PST) has been proposed as a metric to measure the reliability of single-output quantum circuits [39]. It is defined as  $\frac{\text{Number Of Successful Trials}}{\text{Total Number Of Trials}}$ . To identify the correct output while using PST, the probability of each incorrect output should be less than PST. In other words, the correct output should be the most frequent output pattern. For example, if PST is equal to 40%, while the probability of each of the remaining output combinations is less than 40%, we can extract the correct output of the quantum circuit.

## 2.2 Quantum circuit compilation

A quantum compiler converts an abstract quantum circuit to a circuit that consists of operations compatible with the quantum architecture such as single-qubit and CNOT gates. It supports gate decomposition, optimization, and mapping to satisfy the quantum hardware constraints. Complex quantum operations represented by  $2^n \times 2^n$  unitary matrices, such as CU gates, can be decomposed into matrix and tensor products of elementary gates, represented by smaller unitary matrices [3, 6, 8, 10, 21, 31, 35] as illustrated in Fig. 1. A quantum compiler applies a series of




Fig. 2. The coupling graph of IBM Q 16 Melbourne and its gate and measurement error rates. All error rates are given in units of  $10^{-3}$  and were measured on Oct. 19, 2020.

optimization techniques to reduce the gate count by using rule-based and propagation-based optimization for gate cancellation [4, 19, 27, 30]. Next, physical qubits are allocated through the quantum circuit mapping, which transforms the optimized logical circuit to a physical circuit ready to be executed on the quantum computer [17, 18, 23, 24, 33, 36, 41, 45]. Superconducting qubits have been adopted by different quantum computing companies (e.g., IBM [2], Google [22], and Rigetti [1]). They have a limited qubit-to-qubit connectivity, which is restricted to neighboring qubits. The qubit connectivity is described using a Coupling Graph (CG), where each node represents a qubit and each edge corresponds to an entanglement capability between two qubits. To enable multi-qubit gates on nonadjacent physical qubits, the compiler can apply SWAP operations, each of which can be decomposed into three CNOT gates.

### 2.3 Qubit reliability

Qubits are susceptible to gate, coherence, crosstalk, and measurement errors. Qubits suffer from decoherence caused by natural relaxation time ( $T_1$ ) and dephasing time ( $T_2$ ) due to environmental effects.  $T_1$  coherence time, referred to as amplitude damping, causes a qubit at state  $|1\rangle$  to collapse to the ground state (state  $|0\rangle$ ). On the other hand,  $T_2$  coherence time, referred to as phase damping, moves the qubit to a different state. These errors impose limitations on the depth of a quantum circuit. Errors in qubit operations can spawn from imperfect calibrations or from a leakage of qubit states to higher excitations during a computation. The gate error is modeled using Pauli error channel, which consists of Pauli operators ( $X, Y, Z, I$ ) that operate on qubits. The error rates vary from one qubit to another. Measurement operations are another source of errors. Crosstalk can create correlations between independent qubits of the quantum circuit [34]. The likelihood of crosstalk errors depends on the operations applied to adjacent physical qubits. NISQ quantum computers are calibrated very frequently. For example, different IBM Q computers are calibrated once, twice, or even more per day. In the calibration process, the properties of the quantum hardware such as the qubit frequency are computed. Furthermore, the qubit error rates are characterized. Randomized benchmarking can characterize the behavior of qubits and their gates [26]. A randomized benchmark consists of a sequence of quantum gates, which form the identity operation. Quantum state tomography [40] can be used also to reconstruct the quantum state of the circuits by performing sequences of measurements in different bases.

**EXAMPLE 1.** Fig. 2 shows the CG of IBM Q 16 Melbourne quantum architecture. The edge label indicates the CNOT gate error rate for two connected qubits, while the upper and the lower node labels indicate the single-qubit gate and measurement error, respectively.

The variation in the error rates has a significant impact on the reliability of circuit output [15, 29]. Thus, the physical constraints of the quantum architecture including the coupling constraints, gate times, gate error rates, and qubit coherence times are provided as compile-time information

which determines the structure of the physical quantum circuit. Quantum mapping approaches based on the physical constraints of the quantum hardware can enhance the success probability of the correct output by allocating physical qubits with the lowest CNOT error rates [39] and limiting the quantum circuit depth according to the qubits coherence time [7, 28]. The mapping process can rely on the Estimated Success Probability (ESP) of the quantum circuit defined as  $ESP = \prod_{i=0}^G (1 - e_{g_i}) \times \prod_{i=0}^Q (1 - e_{m_i})$ , where  $G$  and  $Q$  are the number of gates and measured qubits, while  $e_{g_i}$  and  $e_{m_i}$  are gate and measurement error rates, respectively [32].

### 3 RELATED WORK

Due to variable error rates of qubits and their gate operations, quantum circuit mapping approaches that rely on dynamic information to improve the reliability of quantum circuit outputs have been proposed. To enhance the success probability of outputs, various physical qubit allocations are explored and used in different trials/runs of the physical quantum circuit to reduce the impact of correlated errors. Their output distributions are then merged to reduce the probability of incorrect outcomes [37]. Another approach is based on inverting quantum states prior to a measurement operation to reduce measurement errors since measuring an all-zero state has a higher fidelity compared to other states [38]. The quantum circuit is executed several times to predict the suitable qubit inversions. Multiprogramming quantum algorithms can benefit from efficient utilization of quantum hardware resources to increase system throughput by running them concurrently on a NISQ computer [14]. Another approach suggests a very frequent characterization of the error rates of the quantum hardware prior to the execution of the quantum circuit using randomized benchmarking to capture the variation in the error rates [43].

Quantum circuit assertions have been proposed to detect bugs in the circuit, which can affect the qubit state such as classical, superposition, and entanglement assertions. While some of the assertion based techniques apply statistical test [20] or projection-based techniques [25], others rely on executing the quantum circuits including the assertions to not only detect bugs in the circuit but also improve the reliability of the circuit output [44]. In the later set of assertions, ancillary qubits are used for debugging purposes to test for classical, superposition, and entangled qubit conditions, which can be inserted at different circuit locations. In the absence of the error, the measurement operation applied to an ancillary qubit does not affect the functionality (or the entanglement) of the circuit, but instead it is used to enhance the success probability of the output. For example, to check an arbitrary superposition state ( $U3(\theta, \phi, \lambda)$ ) of qubit  $q_0$ , we use ancillary qubit  $q_1$  and apply the following gates supported by IBM Q computers:

$$H q_1; \quad CU3(2 \cdot \theta, \phi, \pi - \phi) q_1 q_0; \quad H q_1.$$

followed by the measurement operation on  $q_1$ . A uniform superposition assertion can be simplified to:

$$H q_1; \quad CNOT q_1 q_0; \quad H q_1.$$

In this paper, we identify changes in circuit outputs associated with unexpected variations in qubit error rates and propose a dynamic detection mechanism based on test point insertion, which can be integrated with quantum mapping approaches.

### 4 IMPACT OF THE VARIATION IN QUBIT ERROR RATES

Due to the noisy nature of NISQ computers, the construction of physical quantum circuits is primarily dependent on the error rates of the qubits and their operations. Without accurate mapping policies that minimize the error probability of a quantum circuit, output state fidelities can be significantly degraded.

The impact of different qubit allocations that yield the same circuit structure and gate count depends on the error probabilities of the quantum hardware. We use the standard deviation of the quantum hardware error rates to study the impact of the variation in the error rates on the quantum circuit output. A qubit is susceptible to many different error types. We focus on two-qubit gate error rates to compute the error rate standard deviation since SWAP operations are heavily used in quantum circuits to cope with the coupling constraints of a quantum architecture. Given the mean of the two-qubit gate errors  $m_e$ , we compute the standard deviation of the two-qubit gate error as  $\sqrt{\sum_g (e_g - m_e)^2 / K}$ , where  $e_g$  is a two-qubit gate error and  $K$  is the number of two-qubit gates. We can also compute standard deviations of measurement errors and coherence times. A high standard deviation indicates a diverse error rate distribution, which increases the impact of qubit allocation on the output state fidelity.

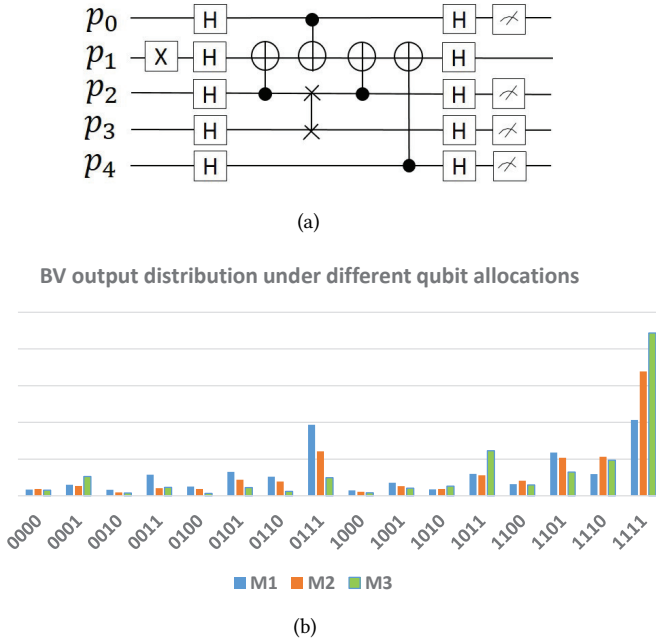


Fig. 3. (a) Bernstein-Vazirani quantum circuit that satisfies the constraints of the IBM Q 16 architecture, and (b) its output distributions for three different physical qubit allocations M1, M2, and M3.

**EXAMPLE 2.** To illustrate the impact of the variation in the qubit error rates, we consider the Bernstein-Vazirani (BV) quantum algorithm [9]. In this example, BV consists of 5 qubits and outputs the 1111 string. We map the BV quantum circuit to satisfy IBM Q 16 Melbourne coupling constraints. Three different physical qubit allocations, denoted as M1, M2, and M3 are selected from the CG of IBM Q 16 Melbourne architecture in Fig. 2 to generate physical quantum circuits. The resulting BV circuit post-mapping using any of the three qubit allocations is shown in Fig. 3(a). For  $p_0, p_1, p_2, p_3$ , and  $p_4$  logical qubits,  $M1 = \{q_6 \leftarrow p_0; q_5 \leftarrow p_1; q_4 \leftarrow p_2; q_3 \leftarrow p_3; q_9 \leftarrow p_4\}$ ,  $M2 = \{q_8 \leftarrow p_0; q_9 \leftarrow p_1; q_{10} \leftarrow p_2; q_{11} \leftarrow p_3; q_5 \leftarrow p_4\}$ , and  $M3 = \{q_3 \leftarrow p_0; q_4 \leftarrow p_1; q_5 \leftarrow p_2; q_6 \leftarrow p_3; q_{10} \leftarrow p_4\}$ . The output distributions of M1, M2, and M3 quantum circuits when executed with 8192 shots (measurements) are provided in Fig. 3(b). M2 and M3 mappings provide the correct output (1111) as the most frequent output, while M1 mapping results in two output states (0111 and 1111) with almost equal frequency (19.3% and 20.6%). Similarly, the standard deviation of CNOT error rates used in M2 and M3 qubit allocations (0.012) is lower than the corresponding standard deviation of CNOT error rates used in the M1 and M2 qubit allocations (0.014).

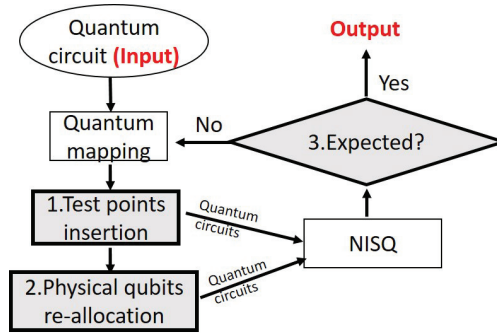


Fig. 4. Overview of the proposed steps highlighted in the gray color. The input is a quantum circuit and the output is the expected (correct) circuit output.

*This indicates that the more diverse the error rates are, the higher the impact of using different qubit allocations on the circuit output.*

The variation in the qubit error rates impacts the reliability of the quantum circuit output. The lack of an accurate failure model to capture changes in the qubit behavior reduces the output reliability. While the error rates change naturally most of the time, they can also be modified intentionally to open a backdoor into quantum systems. For example, a malicious compiler can alter error rates of the quantum architecture used as compile-time information to allocate very noisy qubits. The end result is unreliable and incorrect output extraction.

## 5 PROPOSED SCHEME

We propose to monitor quantum circuit error rates at runtime using test points. The goal is to capture dynamic changes in the error rates. Our test points are inserted into physical quantum circuits post-mapping and prior to an execution on the quantum computer. We first identify viable test points of the quantum circuit. Second, we allocate another set of physical qubits that generates a quantum circuit with a low ESP according to the physical constraints of the quantum architecture given at the compile time to provide diverse error rates. Test points are inserted into both the initial and the newly generated physical circuit. The resulting circuits are executed on the quantum hardware. Finally, the output distributions of the test points for both qubit allocations are used to detect any unexpected changes in the qubit error rates. Our proposed scheme can also be used to select between different qubit allocations that result in circuits with distinct output distributions when executed on the quantum computer. In the presence of unlimited access to the quantum computer, our proposed approach can trigger the need for recalibration using randomized benchmarking. The overview of our approach is shown in Fig. 4.

### 5.1 Test point insertion

Test point insertion involves observing (measuring) the state of the qubits at different points in a circuit. To identify the circuit error rate at each test point, the expected noise-free qubit state should be known, which may not be feasible for large circuits using classical computers. Thus, we rely on test circuits that produce a known output. Our test types include classical, superposition, and uncompute tests.

Classical tests require measuring circuit qubits at their classical states including helper qubits used to implement complex quantum operations (oracles) after they are uncomputed to their initial state. Similar to the superposition assertion in [44], the superposition test entangles an



ancillary qubit with a superposition qubit through a CU gate based on the qubit superposition control parameters. The ancillary qubit is measured to test for the qubit superposition state. The location and the number of the classical and superposition test points depend on the structure of the quantum circuit. Then, we apply uncompute test for a given quantum subcircuit to quantify the subcircuit gate error rates. Uncompute test involves applying the conjugate transpose  $U^\dagger$  gates of a given quantum subcircuit in a reverse order such that  $UU^\dagger = U^\dagger U = I$ . Then, all the physical qubits, which should ideally be in their initial state, are measured.

Our test point insertion algorithm traces the quantum circuits to identify all possible test point locations suitable for each test type, while satisfying the physical constraints of the quantum architecture used at the compile time. Superposition and classical tests are injected after the circuit SWAP operations, which are applied to qubits post initialization, to capture the two-qubit gate errors. In order to insert a test point  $x_{ij}$  at cycle  $t$  for testing a superposition state of qubit  $q_j$ , a free ancillary qubit  $q_i$  adjacent to  $q_j$  is required, so that a CNOT gate can be directly applied to  $q_i$  and  $q_j$ . While a superposition test depends on the coupling constraints of the quantum architecture, an uncompute test should instead satisfy the qubit coherence time. The uncompute insertion algorithm is shown in Algorithm 1. Overall, the number of test points depends on the structure of the physical quantum circuit post-mapping.

**Input:**  $C$  = The physical quantum circuit,  $T_1$  and  $T_2$  of all circuit physical qubits,  $N$  = User defined number of uncompute tests

**Output:**  $L$  = List of all test circuits,  $Cov$  = The gate coverage

- 1  $T_{min}$  = The minimum coherence time of all the circuit qubits,  $M$  = is the maximum number of the physical quantum circuit gate layers in any uncompute test,  $T_{C_M}$  = The subcircuit  $C_M$  time
- 2 **for each physical qubit  $q_i$  in the circuit do**
- 3     **if**  $T_{min} > \text{Min}(T_{1i}, T_{2i})$  **then**
- 4          $T_{min} = \text{Min}(T_{1i}, T_{2i});$
- 5     **end**
- 6 **end**
- 7 Find the maximum number of gate layers  $M$  such that  $T_{min} \geq 2 \times T_{C_M}$ ;
- 8 Generate  $N$  different uncompute tests where the depths of the test circuits are multiples of  $M/N$ ;
- 9 Compute  $Cov_i$  for each test  $i$  and the overall  $Cov$  for all tests;
- 10 Return  $L$ ;

**Algorithm 1:** Uncompute test insertion.

To evaluate the efficiency of the uncompute tests, we propose a gate coverage metric defined as the percentage of two-qubit gates operated up to that point in the generated circuit. It is computed as  $\frac{\#E_{Cov}}{K} \times 100$ , where  $\#E_{Cov}$  is the number of tested unique two-qubit gates and  $K$  is the total number of unique two-qubit gates in the circuit. We compute the coverage metric for each test as well as the overall coverage metric of all the tests. Using our proposed coverage metric, we can 1) eliminate redundant test points, and 2) assess the quality of the test outcome, especially in the presence of the coherence requirement which limits the test circuit depth.

**EXAMPLE 3.** Fig. 5 shows the classical, superposition, and uncompute test points that can be applied to the BV quantum circuit generated in Fig. 3(a) using the M3 qubit allocation. The classical test is applied by measuring all the quantum circuit qubits ( $q_3, q_4, q_5, q_6,$  and  $q_{10}$ ). We can test 4 superposition states at point 2, 3, 4, and 5 using  $q_9, q_2, q_9,$  and  $q_9$  ancillary qubits, respectively, which can directly interact with the tested qubits in the superposition state using a single CNOT gate. The only exception is the  $q_4$  qubit which doesn't have a free adjacent ancillary qubit to conduct the superposition test. Given the qubit coherence time, we can uncompute not only part of the circuit (point 6) but also the

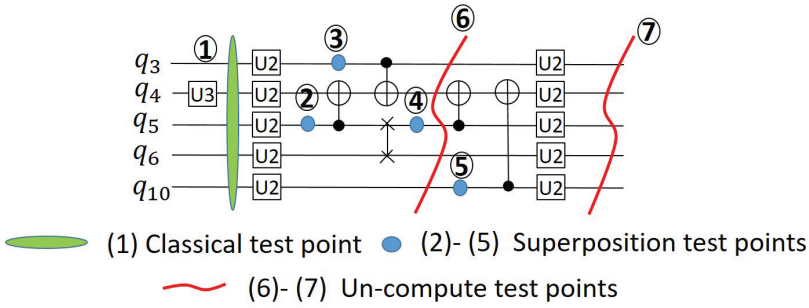


Fig. 5. The location of classical, superposition, and uncompute test points of the BV circuit generated in Fig. 3(a) using the M3 qubit allocation.

entire circuit at point 7. The two-qubit gate coverages of the uncompute test at point 6 and 7 are 75% and 100%, respectively.

Classical and superposition tests check few qubit states and require additional ancillary qubits, which limit their applicability to different quantum algorithms. Thus, we also focus on the uncompute test circuit type as a standalone testing mechanism to monitor the circuit error rates. Particularly, we show in the experimental section (1) the effectiveness of applying only the uncompute test circuits compared to all the test circuit types, and (2) the small number of the required uncompute test circuits with respect to different quantum algorithms executed on NISQ computers.

## 5.2 Physical qubit allocation selection

Due to variable error rates of quantum hardware, predicting the success rate of each test point given its physical constraints is very challenging. Thus, we compare test points under different qubit allocations which are expected to have different error rates. We ensure the generation of the same test circuit structure under different qubit allocations. We allocate another set of qubits such that the resulting quantum circuit shares the same structure as the initial physical circuit generated by the mapping procedure, but has different error rates. This problem is equivalent to finding a subgraph of CG, which is isomorphic to the initial qubit allocation subgraph, but has a lower ESP. We use VF2 algorithm [11] to find isomorphic subgraphs. We compute the ESP for the resulting quantum circuit of each isomorphic subgraph and select the one with the lowest ESP value. The time complexity of the algorithm is  $O(V^2)$  in the best case and  $O(V!V)$  in the worst case, where  $V$  is the number of nodes in the CG. The complexity can be reduced by imposing constraints on the graph topology according to the coupling constraints of the quantum architecture. For example, many of the current NISQ devices support grid-like or nearest-neighbor architectures.

## 5.3 Offline analysis

We use the ESP to estimate the fidelity of the quantum test circuits including the additional gates and qubits used for testing. We compare success rates of the test circuits after executing them on the quantum computer by comparing the test outcome using ESP of each pair of test circuits under the two different qubit allocations with the corresponding test outcome using PST. We count the number of Effective Tests (ET) that comply with the ESP of the test circuits. For quantum circuits  $C$  and  $C'$ , generated using the physical qubit allocations with high and low ESP, a test  $t_i$  is considered effective when the following conditions are satisfied:

$$\begin{aligned}
 \text{ESP}(t_{i_C}) > \text{ESP}(t_{i_{C'}}) &\Rightarrow \text{PST}(t_{i_C}) > \text{PST}(t_{i_{C'}}), \text{ and} \\
 \text{ESP}(t_{i_C}) < \text{ESP}(t_{i_{C'}}) &\Rightarrow \text{PST}(t_{i_C}) < \text{PST}(t_{i_{C'}}).
 \end{aligned}$$

Table 1. Properties of quantum circuit benchmarks.

Benchmarks	# Qubits	#			Depth	Expected output
		U	CX	M		
BV_5	6	12	4	5	7	11110
BV_6	7	14	5	6	8	111110
BV_7	8	16	4	7	7	0011011
BV_8	9	18	5	8	8	01110011
Adder	3	18	17	2	27	0100
Decoder (decod24-v0_38)	4	24	21	4	29	11
Graycode (graycode6_47)	6	0	5	6	5	000000
Grover	3	13	7	2	15	10
-----	-----	-----	-----	-----	-----	-----
Tof_3	5	24	18	5	31	00000
Tof_4	7	40	30	7	50	0000000
Tof_5	9	56	42	9	69	00000000
QFT	10	99	90	15	62	00000000000000

Thus, the Percentage of Effective Tests (PET) is computed as  $PET = \frac{ET}{N} \times 100$ , where  $N$  is the total number of tests.

## 6 EXPERIMENTAL EVALUATION

We present the experimental results in this section. Three sets of experiments have been performed. In the first set, we study the variation in qubit errors and its impact on the output state fidelity. In the second set, we demonstrate the impact of two-qubit gate errors on qubit allocation. In the third set, we show the effectiveness of test points in detecting major changes in qubit behavior.

### 6.1 Experimental setup

We use Qiskit SDK [4] for quantum circuit decomposition, optimization, and mapping. The inputs to Qiskit are various quantum circuits described using the quantum assembly language (QASM) [12]. We consider the following benchmarks:

- **Grover Search.** A circuit that finds the value of  $x_1$  and  $x_2$  which set the  $f(x_1, x_2) = x_1x_2$  oracle to 1 [16].
- **Bernstein–Vazirani (BV).** Multiple BV circuits with varying numbers of qubits that generate different strings [9].
- **Quantum Fourier Transform (QFT).** A QFT circuit with 10 qubits.
- **Other reversible circuits.** Other reversible circuits (1-bit adder, graycode, and decoder) obtained from [13] and RevLib [42], in addition to Toffoli gates with varying number of control qubits.

Table 1 provides detailed information about the quantum circuit benchmarks prior to the mapping process including the number of qubits, single-qubit (U) gates, CNOT (CX) gates, and measurement operations (M), the depth of each quantum circuit, and its expected output. While many of these benchmarks are executed on actual quantum computers, larger benchmarks have been simulated on a noisy simulator. We list all the simulated benchmarks below the dashed horizontal line in Table 1. The output of Qiskit is a physical quantum circuit that satisfies the architectural constraints of the underlying hardware. The applied heuristic mapping approach provided by Qiskit minimizes the gate and readout errors. We refer to the resulting quantum circuit as *good mapping* circuit. In each of the three experiments, PST, and thus the percentage of the expected output state ( $PST \times 100$ ), can infer the correct output if it is the most frequent output pattern.

Table 2. Properties of different quantum computers

Company	Device	Number of Qubits
IBM	IBM Q 16 Melbourne	15
IBM	IBM Q Cambridge	28
IBM	IBM Q Paris	27

In the first set of experiments, we study the possibility of unpredictable changes in the qubit behavior. We execute a subset of the benchmarks provided in Table 1 (BV\_6, BV\_7, BV\_8, Decoder, Graycode) on the IBM Q 16 Melbourne quantum chip. We generate the good mapping circuit for each benchmark according to the physical constraints of the quantum architecture. Each good mapping circuit is executed twice a day (morning and evening) to monitor any changes in the qubit behavior prior to the next calibration, where a single execution consists of 1024 shots. We repeated the experiment over a period of 6 months. The quantum computer was calibrated once a day during the 6 months period of the experiment, which occurs very early in the morning.

In the second and third set of experiments, we use our qubit reallocation script to generate another physical qubit allocation that has the lowest ESP and results in an isomorphic subgraph of the good mapping circuit. We refer to this qubit allocation as *bad mapping*. In the second set of experiments, we run both the good and bad mapping circuits on the IBM Q 16 Melbourne quantum chip. Each run consists of 8192 shots.

In the third set of experiments, we developed the test point insertion algorithm described in Section 5.1, which scans the circuit, identifies all possible superposition, classical, and uncompute test points locations, injects the test points into the good and bad mapping circuits, and computes the coverage of these test points. We conduct our analysis on various IBM quantum computers, where each run consists of 8192 shots. Table 2 shows the quantum computers that we use in our third set of experiments. We provide two sets of data. The first one is generated by executing quantum circuits on several quantum computers, while the second one is generated on a noisy simulator. We utilize the Qiskit noise model based on the properties of a particular device. The noise model is based on the backend calibration data which are passed as inputs to the quantum compilation and simulation tools. In both scenarios, we show the effectiveness of the test points in identifying unexpected changes in qubit behavior.

In all our experiments, we download the calibration data when the quantum circuit is executed on a NISQ computer to ensure that 1) the good and bad mapping circuits and their test point circuits are executed on the same day, and 2) ESP of the good and bad mapping circuits and their test point circuits are computed based on the most up-to-date calibration data. Thus, we can still apply our approach to quantum circuits that are waiting in a queue for more than a day as long as the good and bad mapping circuits and their corresponding test circuits are executed on the same day.

## 6.2 Variation in qubit error rates

This class of experiments shows the impact of the variation in the qubit error rates post-calibration on the output state fidelity of the quantum circuit output. We select a subset of the good mapping quantum circuits, which provide the expected (correct) output as the most frequent output when executed on IBM Q 16 Melbourne in either one of the two runs (morning or evening) per day. We compare the percentage of the expected output state of the morning and evening runs of each good mapping quantum circuit executed every day.

Fig. 6 (a) provides the properties of the physical quantum circuits post-mapping including the average number of single-qubit (U) and two-qubit (CX) gates, the average depth, and the average ESP. Fig. 6 (b) shows the percentage of good mapping quantum circuits that do not meet the requirement

for correct output extraction defined in Section 2.1 either in the morning or the evening. Out of all five circuits, Decoder has the largest number of executed circuits with different morning and evening extracted outputs. We observe up to 47% difference in the morning and evening percentage of the expected output. Our results show that qubit behavior can change throughout the day and prior to the next calibration. We also observe that the more the number of quantum gates and the larger the circuit depth, the higher the impact of the error drifts on the quantum circuit output. However, the error drifts can either improve or reduce the fidelity of the quantum circuit output. We found that only 49% of our collected data through the 6 months period show that morning executions yield a higher success rate than the corresponding evening executions. Thus, there is no precise error behavior that explains the impact of error drifts on the output state fidelity.

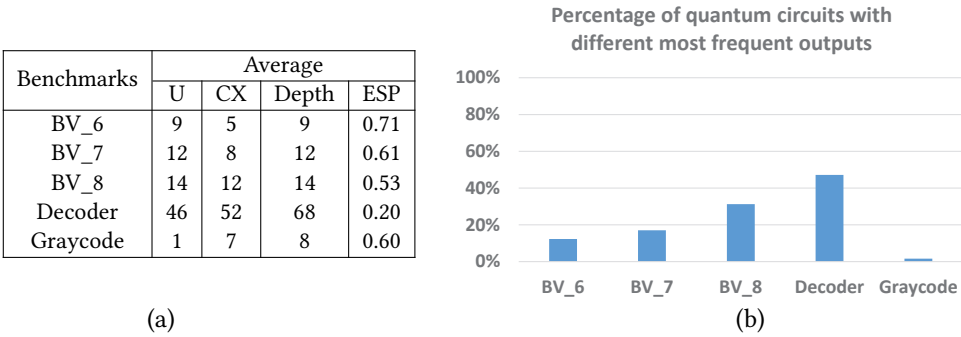


Fig. 6. (a) The properties of quantum circuits post-mapping, and (b) the relative number of days where morning and evening runs yield different (most frequent) outputs.

### 6.3 Impact of two-qubit gate errors on qubit allocation

This class of experiments shows the importance of considering different qubit allocations in the presence of highly variable error rates of NISQ systems. As our test point insertion targets covering as many two-qubit gate errors as possible, we study the correlation between the variation in the CNOT error rates and the impact of the qubit allocation on the quantum circuit output. We compute the standard deviation of the two-qubit gate errors of the IBM Q 16 Melbourne quantum computer when executing quantum circuits generated using the good and bad mappings. As our proposed approach restricts qubit allocations to isomorphic subgraphs, we compute the standard deviation of all two-qubit gate errors used in the good and the bad qubit mappings. To show the impact of the variation of the two-qubit gate errors on the circuit output, we select physical quantum circuits, which are generated on different days but share the same number of CNOT gates post-mapping.

Fig. 7 illustrates the correlation between the percentage of the expected output of the good and bad mapping Grover Search physical quantum circuits and the standard deviation of the error rates of CNOT gates used in the mappings. For each standard deviation value, we show the percentage of the expected output of the good and bad mapping circuit. Each of these two percentages can provide the expected output (Good mapping\_Exp in blue color, Bad mapping\_Exp in green color) or not (Good mapping\_Unexp in yellow color, Bad mapping\_Unexp in orange color). In Fig. 7, the good mapping quantum circuits always provide the expected output.

Fig. 7 shows that the more diverse the CNOT error rates are, the higher the impact of the qubit allocation on the quantum circuit output, which motivates the use of qubit reallocation to detect unexpected changes in the qubit behavior. Specifically, as the standard deviation of the CNOT error rates increases, the percentage of the expected output of bad mapping qubit allocation is

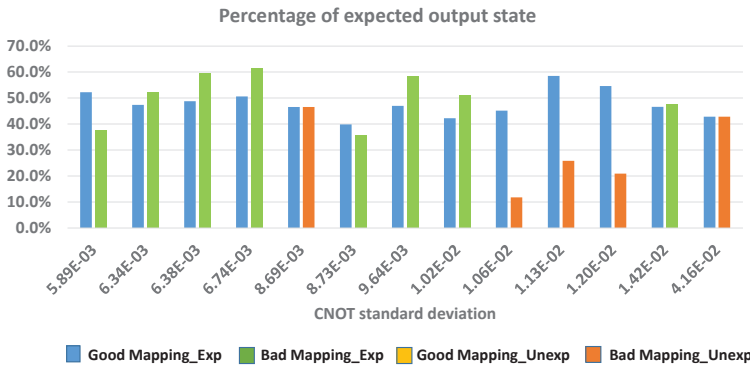


Fig. 7. The percentage of the expected output of the Grover Search quantum circuits generated using good and bad qubit allocations in the presence of different standard deviation of the CNOT error rates.

reduced up to the point where the bad mapping qubit allocation fails to provide the expected output. An exception is the result for the standard deviation of 1.42E-02. This can happen due to evenly distributed CNOT gates with diverse error rates over the good and bad mapping quantum circuits.

### 6.4 Online monitoring of quantum circuits using test points

This class of experiments shows the effectiveness of our proposed online monitoring approach in selecting qubit allocation, which enhances the circuit output state fidelity. For each benchmark, we run good and bad mapping quantum circuits and their test point circuits on different quantum computers several times. The total number of executions of good and bad mapping circuit pairs on all the selected quantum computers is 104 executions. We omit 21% of these executions in which both good and bad mapping circuits fail to provide the expected output due to high error rates. We select a subset of the remaining results to illustrate the effectiveness of our approach.

Table 3 shows the effectiveness of the test points in selecting between good and bad mapping qubit allocations. Column 1 indicates the name of the quantum computer. Column 2 provides the name of the quantum circuit. Columns 3-5 show the number of classical, superposition, and uncompute tests, respectively. Columns 6-7 indicate whether the outputs of the quantum circuit generated using the good mapping (G\_Map) and the bad mapping (B\_Map) are expected (correct). Columns 8-9 provide PET (Percentage of Effective Tests) of the good mapping circuits compared to the bad mapping circuits considering all the tests and the uncompute tests only, respectively. Column 10 shows the percentage of the two-qubit gate (CNOT) coverage of the test circuits. Column 11 indicates the number of CNOT gates post-mapping. We highlighted the cases in which the good and bad mapping circuits provide conflicting output states.

The number of superposition, classical, and uncompute tests vary from one physical quantum circuit to another depending on the circuit structure and the coherence time of the physical qubits. The small the number of superposition tests is due to the need for a free ancillary qubit adjacent to the target qubit in the superposition state in both good and bad mapping circuits. While each circuit has its own free adjacent physical qubits, the relative location of these qubits with respect to the circuit may vary from good to bad mapping circuits. Therefore, the location of the superposition test points in good and bad mappings may vary, which necessitates discarding these test points. Thus, if there is a free qubit adjacent to the  $p_0$  logical qubit in the good mapping circuit, while there is no free qubit adjacent to  $p_0$  in the bad mapping circuit, we ignore  $p_0$  superposition test.

Table 3 shows that when good and bad mapping circuits provide different output states, PET can be used to select the output of the quantum circuit as long as the coverage metric is very high.

Table 3. Test points comparison of good and bad mapping quantum circuits executed on NISQ computers. C, S, and UN denote the classical, superposition, and uncompute tests. G\_Map and B\_Map are the good and bad mapping circuits.

Quantum Computer	Benchmarks	# Tests			Most Frequent Output is Correct?		PET		Gate Coverage	CNOT Count
		C	S	UN	G_Map	B_Map	All_tests	UN		
IBM Q Cambridge	BV_5	2	0	8	Yes	Yes	80%	70%	100%	23
	BV_6	1	0	8	Yes	Yes	20%	0%	100%	25
	BV_7	1	4	9	No	No	21%	12%	84%	27
	<b>BV_8</b>	<b>1</b>	<b>2</b>	<b>9</b>	<b>Yes</b>	<b>No</b>	<b>67%</b>	<b>62%</b>	<b>100%</b>	<b>27</b>
	<b>Decoder</b>	<b>2</b>	<b>1</b>	<b>14</b>	<b>Yes</b>	<b>No</b>	<b>94%</b>	<b>100%</b>	<b>83%</b>	<b>54</b>
	Grover	2	2	11	Yes	Yes	87%	100%	100%	8
	Adder	3	0	13	Yes	Yes	100%	100%	100%	29
IBM Q Paris	<b>BV_5</b>	<b>2</b>	<b>0</b>	<b>14</b>	<b>Yes</b>	<b>No</b>	<b>100%</b>	<b>100%</b>	<b>100%</b>	<b>36</b>
	BV_6	3	0	12	No	No	100%	100%	100%	55
	BV_7	2	1	12	Yes	Yes	94%	100%	100%	21
	BV_8	2	1	12	Yes	Yes	94%	100%	100%	35
	Decoder	1	1	12	No	No	8%	0%	100%	54
	Grover	3	1	11	Yes	Yes	60%	100%	100%	11
	<b>Adder</b>	<b>3</b>	<b>2</b>	<b>12</b>	<b>Yes</b>	<b>No</b>	<b>89%</b>	<b>100%</b>	<b>100%</b>	<b>34</b>
IBM Q 16 Melbourne	BV_5	2	1	6	Yes	Yes	33%	0%	100%	14
	BV_6	3	1	6	No	No	30%	20%	50%	30
	<b>BV_7</b>	<b>1</b>	<b>1</b>	<b>6</b>	<b>No</b>	<b>Yes</b>	<b>0%</b>	<b>0%</b>	<b>70%</b>	<b>20</b>
	BV_8	1	0	7	No	No	75%	83%	50%	28
	Decoder	2	1	6	Yes	Yes	11%	0%	48%	56
	Grover	1	0	6	Yes	Yes	57%	40%	100%	7
	Adder	3	2	13	Yes	Yes	94%	100%	100%	26

Specifically, PET computed based on all the tests and the corresponding PET computed based on the uncompute tests only agree on the selection of the qubit allocation that yields a better output. Furthermore, a low PET value indicates changes in the qubit error rates.

Fig. 8 provides the percentage of the expected output for good and bad mapping circuits as well as PET based on all the tests, also listed in Table 3. The figure shows that when the percentage of the expected output of the good mapping circuit outperforms the bad mapping circuit, PET value will be higher than 50% most of the time. The same observation holds if PET is computed based on the uncompute tests only. There are few exceptions (such as BV\_5 circuits executed on the Melbourne quantum computer) in which good mapping circuit provides a higher percentage of the expected output but the value of PET is lower than 50%. However, in these cases both good and bad mapping circuits agree on the most frequent output pattern, which is the expected output provided in Table 3.

Our results show a correlation between the percentage of the expected output, or PST, and PET across different qubit allocations. Thus, our approach can be used to identify incorrect outputs and suggest better quantum circuit mapping policies.

To show the scalability of the proposed approach, we apply it to larger quantum circuits and evaluate the effectiveness of test points using a noisy simulator. We use the physical properties and the qubit error rates of IBM Q 16 Melbourne to imitate its noisy behavior using a noisy simulator. PET is computed based on all the test types. Table 4 and Fig. 9 show that if the percentage of the expected output of one qubit allocation is higher than the other one, the majority of the test points of the former qubit allocation will also be higher, and vice versa. This confirms our results for

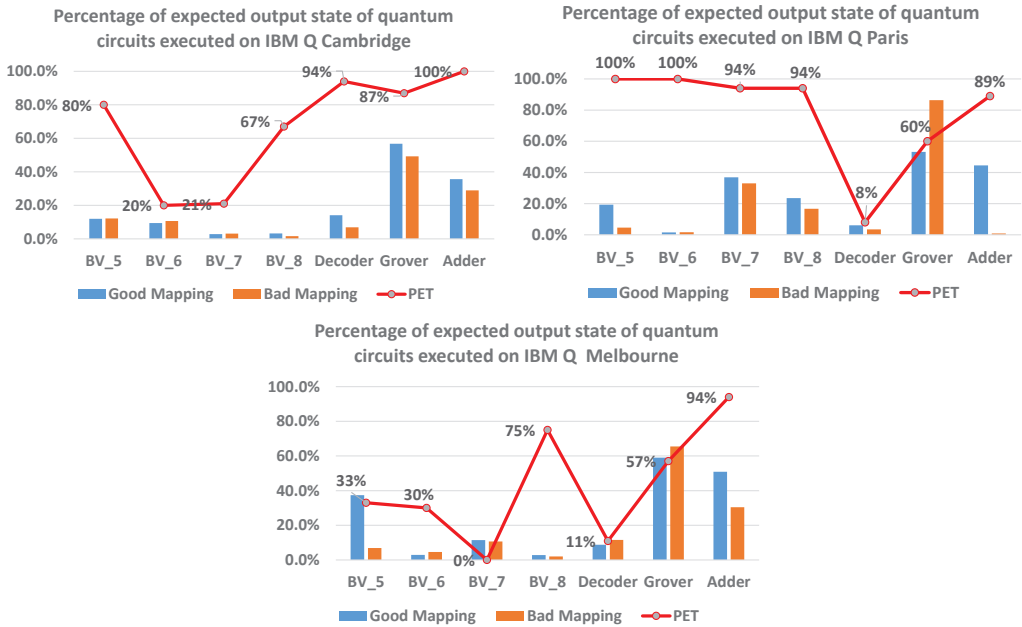


Fig. 8. Expected output percentage under good and bad mappings and their corresponding PET.

Table 4. Test points comparison of good and bad mapping quantum circuits simulated on a noisy simulator. C, S, and UN denote the classical, the superposition, and the uncompute tests. G\_Map and B\_Map are the good and the bad mapping circuits.

Noisy Simulator	Benchmarks	# Tests			Most Frequent Output is Corr. Out?		PET	Gate Coverage	CNOT Count
		C	S	UN	G_Map	B_Map			
IBM Q 16 Melbourne	Tof_3_a	4	1	13	Yes	Yes	100%	78%	49
	Tof_3_b	4	1	15	Yes	Yes	95%	100%	43
	Tof_4_a	5	1	16	Yes	Yes	82%	56%	73
	Tof_4_b	6	0	15	Yes	Yes	81%	56%	72
	Tof_5_a	6	0	16	Yes	Yes	86%	45%	126
	Tof_5_b	6	1	17	Yes	Yes	96%	32%	129
	QFT	3	0	16	Yes	Yes	58%	24%	140

circuit executions on the quantum hardware. For example, when the percentage of the expected output of the good mapping circuit of a Toffoli gate with 4 control lines (Tof\_4\_b) is higher than the percentage of the expected output of the bad mapping circuit, the corresponding PET is 81%.

Finally, to show the limited overhead of the proposed online monitoring scheme of quantum circuits in terms of the number of additional executions in the quantum computer, we show the relationship between the number of uncompute test circuits and the PET based on uncompute tests only of good mapping circuit with respect to the bad mapping circuit in Figure 10. For two different physical quantum circuits (Adder and BV\_5), which vary in the circuit depth (30 to 70) and executed on both IBM Q Melbourne and Paris, the PET value is stable regardless of the number of uncompute tests. Thus, a limited number of uncompute test circuits is sufficient to detect any changes in the circuit error rates.






Fig. 9. Expected output percentage and the corresponding PET obtained on a noisy simulator.




Fig. 10. The impact of the number of uncompute tests on the PET accuracy.

## 7 APPLICATIONS OF THE PROPOSED SCHEME

Our proposed online monitoring scheme can be leveraged for several applications. It can be used to build accurate data driven reliability models of quantum circuits using machine learning algorithms, which take into account not only the error rates obtained during the calibration process but also the error drifts using the proposed test circuit outcomes. These models can improve the accuracy of the estimated success probability of the correct outcome and predict whether the correct outcome can be extracted. Thus, our proposed test circuit design can select between different mapping policies of quantum circuits. It can also be integrated with other quantum compilation approaches such as the one in [37] to find diverse and ensemble set of qubit allocations that will conceal the probability of incorrect outcomes.

## 8 CONCLUSION

In this paper, we addressed the problem of unstable qubits that significantly degrade the output state fidelity of quantum circuits due to malicious activities or drift in hardware error rates. We applied flexible combinations of assertion-based and uncompute testing approaches to monitor quantum circuit errors. The case studies were executed on several IBM quantum computers. Our experimental results show not only the variation in qubit error rates over time but also the effectiveness of our proposed scheme in detecting these changes and aiding the selection of qubit allocations that enhance the success probability of the circuit output. The proposed scheme can also be used to detect incorrect outputs.

## ACKNOWLEDGMENT

MU and WAdJ were supported, and NA was partially supported, through an LBNL Sustainable Horizons internship, by the U.S. Department of Energy (DOE) under Contract No. DE-AC02-05CH11231, through the Office of Advanced Scientific Computing Research Accelerated Research for Quantum Computing and Quantum Algorithms Team Programs. This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract No. DE-AC05-00OR22725.

## REFERENCES

- [1] Rigetti Computing 2018. *Rigetti 16Q Aspen*. Rigetti Computing. <https://www.rigetti.com/gpu/>
- [2] IBM 2019. *IBM Q Systems*. IBM. <https://www.research.ibm.com/ibm-q/technology/devices/>
- [3] N. Abdessaied, M. Amy, M. Soeken, and R. Drechsler. 2016. Technology Mapping of Reversible Circuits to Clifford+T Quantum Circuits. In *2016 IEEE 46th International Symposium on Multiple-Valued Logic (ISMVL)*. 150–155.
- [4] Héctor Abraham and et al. 2019. Qiskit: An Open-source Framework for Quantum Computing. <https://doi.org/10.5281/zenodo.2562110>
- [5] Nikita Acharya and Samah Mohamed Saeed. 2020. A Lightweight Approach to Detect Malicious/Unexpected Changes in the Error Rates of NISQ Computers. In *Proceedings of the IEEE/ACM International Conference on Computer-Aided Design (ICCAD '20)*. Association for Computing Machinery, Article 153, 9 pages.
- [6] M. Amy, D. Maslov, and M. Mosca. 2014. Polynomial-Time T-Depth Optimization of Clifford+T Circuits Via Matroid Partitioning. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 33, 10 (2014), 1476–1489.
- [7] Abdullah Ash-Saki, Mahabubul Alam, and Swaroop Ghosh. 2019. QURE: Qubit Re-allocation in Noisy Intermediate-Scale Quantum Computers. In *Proceedings of ACM/IEEE Design Automation Conference*. New York, NY, USA, 141:1–141:6.
- [8] Adriano Barenco, Charles H. Bennett, Richard Cleve, David P. DiVincenzo, Norman Margolus, Peter Shor, Tycho Sleator, John A. Smolin, and Harald Weinfurter. 1995. Elementary gates for quantum computation. *Phys. Rev. A* 52 (1995), 3457–3467. Issue 5.
- [9] Ethan Bernstein and Umesh Vazirani. 1997. Quantum Complexity Theory. *SIAM J. Comput.* 26, 5 (Oct. 1997), 1411–1473. <https://doi.org/10.1137/S0097539796300921>
- [10] Stephen S. Bullock and Igor L. Markov. 2003. Arbitrary two-qubit computation in 23 elementary gates. *Phys. Rev. A* 68 (Jul 2003), 012318. Issue 1. <https://doi.org/10.1103/PhysRevA.68.012318>
- [11] L. P. Cordella, P. Foggia, C. Sansone, and M. Vento. 2004. A (sub)graph isomorphism algorithm for matching large graphs. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 26, 10 (2004), 1367–1372.
- [12] A. W. Cross, L. S. Bishop, J. A. Smolin, and J. M. Gambetta. 2017. Open quantum assembly language. *CoRR* arXiv:1707.03429v2 (2017).
- [13] Steven A. Cuccaro, Thomas G. Draper, Samuel A. Kutin, and David Petrie Moulton. 2004. A new quantum ripple-carry addition circuit. *arXiv e-prints*, Article quant-ph/0410184 (Oct. 2004), quant-ph/0410184 pages. arXiv:quant-ph/quant-ph/0410184
- [14] Poulami Das, Swamit S. Tannu, Prashant J. Nair, and Moimuddin Qureshi. 2019. A Case for Multi-Programming Quantum Computers. In *Proceedings of the 52Nd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO '52)*. ACM, New York, NY, USA, 291–303. <https://doi.org/10.1145/3352460.3358287>
- [15] Will Finigan, Michael Cubeddu, Thomas Lively, Johannes Flick, and Prineha Narang. 2018. Qubit Allocation for Noisy Intermediate-Scale Quantum Computers. arXiv:quant-ph/1810.08291
- [16] Lov K. Grover. 1996. A fast quantum mechanical algorithm for database search. In *Theory of computing*. 212–219.
- [17] Gian Giacomo Guerreschi and Jongsoo Park. 2018. Two-step approach to scheduling quantum circuits. *Quantum Science and Technology* 3, 4 (2018), 045003.
- [18] Wakaki Hattori and Shigeru Yamashita. 2018. Quantum Circuit Optimization by Changing the Gate Order for 2D Nearest Neighbor Architectures. In *Reversible Computation*, Jarkko Kari and Irek Ulidowski (Eds.). Springer International Publishing, Cham, 228–243.
- [19] Keshu Hietala, Robert Rand, Shih-Han Hung, Xiaodi Wu, and Michael Hicks. 2019. A Verified Optimizer for Quantum Circuits. arXiv:cs.PL/1912.02250
- [20] Yipeng Huang and Margaret Martonosi. 2019. Statistical Assertions for Validating Patterns and Finding Bugs in Quantum Programs. In *Proceedings of the 46th International Symposium on Computer Architecture (ISCA '19)*. Association for Computing Machinery, New York, NY, USA, 541–553. <https://doi.org/10.1145/3307650.3322213>
- [21] Raban Iten, Roger Colbeck, Ivan Kukuljan, Jonathan Home, and Matthias Christandl. 2016. Quantum circuits for isometries. *Phys. Rev. A* 93 (Mar 2016), 032318. Issue 3. <https://doi.org/10.1103/PhysRevA.93.032318>

- [22] J. Kelly. 2018. *A Preview of Bristlecone, Google's New Quantum Processor*. Google Quantum AI lab. <https://ai.googleblog.com/2018/03/a-preview-of-bristlecone-googles-new.html/>
- [23] Janusz Kusyk, Samah M. Saeed, and Muharrem Umüt Uyar. 2021. Survey on Quantum Circuit Compilation for Noisy Intermediate-Scale Quantum Computers: Artificial Intelligence to Heuristics. *IEEE Transactions on Quantum Engineering* 2 (2021), 1–16.
- [24] Gushu Li, Yufei Ding, and Yuan Xie. 2019. Tackling the qubit mapping problem for nisq-era quantum devices. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*. ACM, 1001–1014.
- [25] Gushu Li, Li Zhou, Nengkun Yu, Yufei Ding, Mingsheng Ying, and Yuan Xie. 2020. Proq: Projection-based Runtime Assertions for Debugging on a Quantum Computer. [arXiv:cs.PL/1911.12855](https://arxiv.org/abs/1911.12855)
- [26] Easwar Magesan, Jay M. Gambetta, and Joseph Emerson. 2012. Characterizing quantum gates via randomized benchmarking. *Phys. Rev. A* 85 (Apr 2012), 042311. Issue 4. <https://doi.org/10.1103/PhysRevA.85.042311>
- [27] D. M. Miller, D. Maslov, and G. W. Dueck. 2003. A transformation based algorithm for reversible logic synthesis. In *Design Automation Conference*. 318–323.
- [28] Prakash Murali, Jonathan M Baker, Ali Javadi-Abhari, Frederic T Chong, and Margaret Martonosi. 2019. Noise-adaptive compiler mappings for noisy intermediate-scale quantum computers. In *Proceedings of ASPLOS*. ACM, 1015–1029.
- [29] Prakash Murali, Norbert Matthias Linke, Margaret Martonosi, Ali Javadi Abhari, Nhung Hong Nguyen, and Cinthia Huerta Alderete. 2019. Full-Stack, Real-System Quantum Computer Studies: Architectural Comparisons and Design Insights. In *Proceedings of the 46th International Symposium on Computer Architecture (ISCA '19)*. Association for Computing Machinery, New York, NY, USA, 527–540. <https://doi.org/10.1145/3307650.3322273>
- [30] Yunseong Nam, Neil J. Ross, Yuan Su, Andrew M. Childs, and Dmitri Maslov. 2018. Automated optimization of large quantum circuits with continuous parameters. *npj Quantum Inf.* 4, 1 (10 May 2018), 23.
- [31] M. Nielsen and I. Chuang. 2000. *Quantum Computation and Quantum Information*. Cambridge Univ. Press.
- [32] Shin Nishio, Yulu Pan, Takahiko Satoh, Hideharu Amano, and Rodney Van Meter. 2019. Extracting Success from IBM's 20-Qubit Machines Using Error-Aware Compilation. *CoRR* abs/1903.10963 (2019). [arXiv:1903.10963](https://arxiv.org/abs/1903.10963) <http://arxiv.org/abs/1903.10963>
- [33] Riccardo Rasconi and Angelo Oddi. 2019. An Innovative Genetic Algorithm for the Quantum Circuit Compilation Problem. *Proceedings of the AAAI Conference on Artificial Intelligence* 33 (07 2019), 7707–7714.
- [34] Mohan Sarovar, Timothy Proctor, Kenneth Rudinger, Kevin Young, Erik Nielsen, and Robin Blume-Kohout. 2019. Detecting crosstalk errors in quantum information processors. [arXiv:quant-ph/1908.09855](https://arxiv.org/abs/1908.09855)
- [35] V. V. Shende, S. S. Bullock, and I. L. Markov. 2006. Synthesis of quantum-logic circuits. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 25, 6 (2006), 1000–1010.
- [36] Marcos Yukio Sraichi, Vinicius Fernandes dos Santos, Sylvain Collange, and Fernando Magno Quintão Pereira. 2018. Qubit allocation. In *Proceedings of CGO*. ACM, 113–125.
- [37] Swamit S. Tannu and Moinuddin Qureshi. 2019. Ensemble of Diverse Mappings: Improving Reliability of Quantum Computers by Orchestrating Dissimilar Mistakes. In *Proceedings of MICRO*. 253–265.
- [38] Swamit S. Tannu and Moinuddin K. Qureshi. 2019. Mitigating Measurement Errors in Quantum Computers by Exploiting State-Dependent Bias. In *Proceedings of the 52Nd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO '52)*. ACM, New York, NY, USA, 279–290. <https://doi.org/10.1145/3352460.3358265>
- [39] Swamit S. Tannu and Moinuddin K. Qureshi. 2019. Not All Qubits Are Created Equal: A Case for Variability-Aware Policies for NISQ-Era Quantum Computers. In *Proceedings of ASPLOS*. ACM, 987–999.
- [40] R. T. Thew, K. Nemoto, A. G. White, and W. J. Munro. 2002. Qudit quantum-state tomography. *Phys. Rev. A* 66 (Jul 2002), 012303. Issue 1. <https://doi.org/10.1103/PhysRevA.66.012303>
- [41] Robert Wille, Lukas Burgholzer, and Alwin Zulehner. 2019. Mapping Quantum Circuits to IBM QX Architectures Using the Minimal Number of SWAP and H Operations. In *Proceedings of DAC*. 142–1.
- [42] R. Wille, D. Grosse, L. Teuber, G. W. Dueck, and R. Drechsler. 2008. RevLib: An Online Resource for Reversible Functions and Reversible Circuits. In *Proceedings of International Symposium on Multiple Valued Logic*. 220–225.
- [43] E. Wilson, S. Singh, and F. Mueller. 2020. Just-in-time Quantum Circuit Transpilation Reduces Noise. In *2020 IEEE International Conference on Quantum Computing and Engineering (QCE)*. 345–355. <https://doi.org/10.1109/QCE49297.2020.00050>
- [44] Huiyang Zhou and Gregory T. Byrd. 2019. Quantum Circuits for Dynamic Runtime Assertions in Quantum Computation. *IEEE Computer Architecture Letters* 18, 2 (Jul 2019), 111–114. <https://doi.org/10.1109/lca.2019.2935049>
- [45] Alwin Zulehner and Robert Wille. 2017. Make It Reversible: Efficient Embedding of Non-reversible Functions. In *Design Automation and Test in Europe*. 458–463.