








Received 17 August 2023; revised 24 November 2023; accepted 11 December 2023; date of publication 15 December 2023;
date of current version 12 January 2024.

Digital Object Identifier 10.1109/TQE.2023.3343625

Quantum Vulnerability Analysis to Guide Robust Quantum Computing System Design

FANG QI¹ , KAITLIN N. SMITH² , TRAVIS LECOMPTE³ ,
NIAN-FENG TZENG⁴  (Life Fellow, IEEE),
XU YUAN⁵  (Senior Member, IEEE), FREDERIC T. CHONG⁶  (Fellow, IEEE),
AND LU PENG¹  (Senior Member, IEEE)

¹Tulane University, New Orleans, LA 70118 USA

²Northwestern University, Evanston, IL 60208 USA

³Louisiana State University, Baton Rouge, LA 70803 USA

⁴University of Louisiana at Lafayette, Lafayette, LA 70503 USA

⁵University of Delaware, Newark, DE 19716 USA

⁶University of Chicago, Chicago, IL 60637 USA

Corresponding author: Lu Peng (e-mail: lpeng3@tulane.edu).

The work of Travis LeCompte was supported by a Louisiana Board of Regents Graduate Fellowship. This work was supported in part by the National Science Foundation (NSF) under Grant OIA-2019511, in part by Enabling Practical-scale Quantum Computing (EPiQC), an NSF Expedition in Computing, under Award CCF-1730449, in part by Software-Tailored Architecture for Quantum (STAQ) under Award NSF Phy-1818914, in part by the US Department of Energy (DOE) Office of Advanced Scientific Computing Research, Accelerated Research for Quantum Computing Program, in part by the NSF Quantum Leap Challenge Institute for Hybrid Quantum Architectures and Networks under NSF Award 2016136, in part by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, in part by the Army Research Office under Grant W911NF-23-1-0077, and in part by the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725.

ABSTRACT While quantum computers provide exciting opportunities for information processing, they currently suffer from noise during computation that is not fully understood. Incomplete noise models have led to discrepancies between quantum program success rate (SR) estimates and actual machine outcomes. For example, the estimated probability of success (ESP) is the state-of-the-art metric used to gauge quantum program performance. The ESP suffers poor prediction since it fails to account for the unique combination of circuit structure, quantum state, and quantum computer properties specific to each program execution. Thus, an urgent need exists for a systematic approach that can elucidate various noise impacts and accurately and robustly predict quantum computer success rates, emphasizing application and device scaling. In this article, we propose quantum vulnerability analysis (QVA) to systematically quantify the error impact on quantum applications and address the gap between current success rate (SR) estimators and real quantum computer results. The QVA determines the cumulative quantum vulnerability (CQV) of the target quantum computation, which quantifies the quantum error impact based on the entire algorithm applied to the target quantum machine. By evaluating the CQV with well-known benchmarks on three 27-qubit quantum computers, the CQV success estimation outperforms the estimated probability of success state-of-the-art prediction technique by achieving on average six times less relative prediction error, with best cases at 30 times, for benchmarks with a real SR rate above 0.1%. Direct application of QVA has been provided that helps researchers choose a promising compiling strategy at compile time.

INDEX TERMS Quantum computing, resilience, success rate (SR), vulnerability analysis.

I. INTRODUCTION

Excitement surrounds quantum computation due to the great theoretical potential both fault-tolerant [16], [44] and near-term [30] quantum computers have to solve high-impact

problems. By carefully leveraging quantum superposition, interference, and entanglement, quantum computers are projected to be applied to computational tasks that are currently intractable on today's most powerful supercomputers.




FIGURE 1. Quantum experiment workflow with QVA. One complete cycle ranges from idea conception to real machine run.

Recent progress in quantum hardware has allowed many prototype quantum computers to emerge, and superconducting circuits [5], [6], [45] are gaining popularity as one of the forefront quantum computing technologies. Compared to other quantum hardware, superconducting quantum computers have advantages in scalability, microwave control, and nanosecond-scale gate operation [7], [9], [10].

While promising, superconducting quantum architectures are currently too error-prone to support programs targeted for large-scale applications. Near-term superconducting quantum computers suffer from various noise channels that degrade both quantum information and computation. This noise is difficult to fully characterize and causes retention and operational errors that vary both across-chip and between quantum computers [46] significantly. Quantum error correction was developed to accommodate occasional errors during quantum computation, but current noisy intermediate-scale quantum (NISQ) era machines do not have the operator precision or device scale to implement this routine [40]. Therefore, NISQ quantum machines perform noisy operations as errors can happen on any physical qubit at any time during program execution according to error rates characterized by randomized benchmarking. Fig. 1 shows the full quantum computing flow that transforms a research idea into an experiment on a real quantum computer.

Improving circuit success rate (SR) is a popular research topic, but few studies examine the quantum computer and circuit-dependent errors that result in specific SR. A robust noise model is essential for accurate SR prediction and developing SR boost technologies such as error reduction, bypass, and compression. Unfortunately, existing methods like statistical fault injection [15], [37], [41] and estimated probability of success (ESP) [47] have accuracy and scaling challenges. These models neglect the impact of circuit composition on reported error rates from randomized benchmarking. Therefore, this article proposes a systematic approach to explain different errors and provide accurate SR prediction.

While studying ESP, we found that there are gates whose impact on the output is much more significant than their error rates. We propose the quantum vulnerability analysis (QVA) to address these gate error inconsistent behaviors. QVA is a systematic methodology that performs error modeling based on randomized benchmarking to determine the success of a compiled circuit on a targeted quantum computer: the vulnerability metric cumulative quantum vulnerability (CQV)

quantifies the compiled circuit performance, and 1-CQV provides an SR estimation that closely models actual quantum computer performance. By performing the QVA during compilation, researchers will have a clear and direct view of how gate error behaviors couple to the circuit structure to influence runtime performance. The blue box in Fig. 1 reveals the brief structure of the QVA and how CQV would be used in an experimental flow to estimate the quantum computer SR before real machine evaluation.

We extensively validated the accuracy and stability of QVA in predicting the SR of any given compiled circuits. This assertion is backed by over 160 K experiments conducted on three state-of-the-art 27-qubit IBM quantum computers [1], spanning six distinguished algorithms with varying qubit sizes. The compiled circuits were produced using a variety of compilation strategies and incorporated multiple error-mitigation techniques to address the diverse noise profiles encountered over months of experimentation. This comprehensive approach bolsters the extensibility of our analysis to a broader range of circuits.

All results show that QVA maintains a stable SR estimation via 1-CQV, providing on average 6x improvements (30x in the best case) in relative prediction error over the widely implemented ESP estimator. Additionally, we conduct a case study to provide the quantum community with a direct application of the presented method by choosing promising compiling strategies at compile time. Further, our QVA module provides instructions for reconstructing the model based on a particular quantum device's topology and error behavior. It ensures its compatibility with various superconducting quantum computers, irrespective of their vendor or technology. Below are the contributions of our article.

- 1) We design and build a lightweight error modeling scheme based on QVA and are the first that quantify error rate impact propagating across the CNOT gates with the error rate reported by randomized benchmarking.
- 2) We implement and evaluate a framework that calculates the unique CQV for an algorithm/machine pairing. We show that the proposed SR estimator, 1-CQV, outperforms the current state-of-the-art SR estimation, ESP, by an average 6x less relative prediction error.
- 3) We highlight the scaling potential of CQV: as an algorithm reaches and surpasses the quantum volume of a device, CQV-based methods experience more than 10x improvement in relative prediction error rate compared with the state of the art.

II. BACKGROUND AND RELATED WORK

A. NISQ ERA QUANTUM BASICS AND ERROR CHARACTERIZATION

The flow for generating a quantum executable from a given algorithm and evaluating it on a quantum chip is illustrated in Fig. 1. The quantum compiler will be given information about the target quantum chip and compiler strategy, such as optimization levels, initial layout method, mapping method,

etc. Based on that input, the compiler will follow all intermediate compiling steps to generate a compiled circuit for execution on a quantum computer.

When operating a superconducting quantum computer, operations may fail due to poor environmental conditions, inaccurate control, state decoherence, and more. The error rate associated with an operation, operational error rates, is closely estimated by randomized benchmarking, described in detail below, which approximates the extent of failures without revealing the exact source. The lifetime of a qubit, or its ability to retain a quantum state, is determined by its relaxation time (T1) and decoherence time (T2), so-called retention errors. The decoherence and relaxation time represent the qubit's average time to retain its energized and superimposed states, respectively.

Randomized Benchmarking: Operator performance must be accurately characterized to use a quantum computer effectively. Unfortunately, quantum computer noise models are complex, and it is unscalable to completely characterize system noise via process tomography [39]. In addition to scaling considerations, characterization procedures must separate noise associated with quantum gates from errors stemming from state preparation and measurement to ensure that computation quality can be adequately estimated. Randomized benchmarking [26], [27] is a method of assessing quantum computer hardware that achieves an average error rate for operations through a process known as twirling. At the high level, twirling implements long sequences of random gate operations and fits the resulting data to a curve to determine the average error. Because randomized benchmarking considers only the exponential decay of sequences of random gates, sensitivity to measurement noise in the resulting average error is minimized. Meanwhile, the T1 and T2 errors on gate operation are included as part of the gate's error rate reported by the randomized benchmarking [36]. Randomized benchmarking, while applicable to systems of any dimension, is predominantly employed for single-qubit or two-qubit gates [11]. This preference arises from the exponential growth in the number of required gates as the system's dimension increases, making the method less practical for larger systems. Despite this, the technique can be adapted to pinpoint errors due to unintended crosstalk [12]. Notably, IBM utilizes randomized benchmarking in each calibration cycle to ascertain error rates for their quantum computer's single and two-qubit gates, as reflected in the system's properties

$$\text{Success Rate}(SR) = \frac{\text{Trial counts of correct output}}{\text{Total trial counts.}} \quad (1)$$

Success Rate: The SR is used to gauge quantum program performance on a quantum computer. We compute SR by dividing the number of correct outputs by total executions, as shown in (1). For more details on quantum computing, we refer to [35].

B. RELATED WORK ON QUANTUM SR ESTIMATION

Early quantum computing research was focused on designing quantum hardware [28], instruction set architecture [8], and quantum computer microarchitecture [9], [10], [32], [33]. Afterward, the temporal and spatial noise variation challenges of SC quantum computers were studied to discover mapping and allocation-enhanced compilations to make algorithm execution more robust to diverse errors [20], [21], [22], [31], [46]. Additionally, works such as [18] contribute to the understanding of how noise, fidelity, and computational cost interplay in quantum processing, enriching the broader discourse on quantum system performance. Currently, the focus of quantum computing is on optimizing the SR by applying different compiler strategies, such as mitigating the effect of errors by enhancing the quantum instructions [14], [43], decreasing measurement errors [13], [48], mitigating crosstalk errors [7], [34], combining preexecution and postexecution software approaches to improve performance [19], [47], and compiling with specific constraints [7], [23].

While many studies have improved quantum program performance, two areas have been underexplored: 1) accurate SR prediction for specific compiled circuits and 2) better modeling of error/algorithm relationships in current quantum systems. Regarding SR prediction methods, popular alternatives to noisy quantum computer simulations include using machine learning for SR prediction, which treats the entire computation as a black box [25], developing detailed noise models, and methods like statistical fault injection [17], [37], [41] and estimated success probability (ESP) [36], [46], [47]

$$g_i^s = (1 - g_i^e) \quad m_i^s = (1 - m_i^e) \quad (2)$$

$$\text{ESP} = \prod_{i=1}^{N_{\text{gates}}} g_i^s * \prod_{i=1}^{N_{\text{Measurement}}} m_i^s. \quad (3)$$

III. MOTIVATION

A. LIMITATION OF CURRENT SR ESTIMATOR

1) MACHINE LEARNING BASED

The machine learning-based success rate (SR) prediction method, referenced in [25], simplifies quantum computations into a black box model. This approach can blur distinctions between circuits with varied gate parameters and requires retraining when adapting to different quantum machine sizes. Data collection for larger machines is resource-intensive, especially given the method's requirement to gather data within a single calibration period. Furthermore, its validation, limited to specific compiled circuits, may not account for the complexities of larger machines or diverse error-mitigation strategies.

2) FAULT INJECTION-BASED

The statistical fault injection method employs classical computers to simulate full-state quantum computations. During this process, errors are systematically injected into each basis

gate based on specific triggering probabilities [15]. A very recent work [37] employs fault injection methods to evaluate quantum vulnerabilities. This study proposes the use of quantum fault injection to scrutinize circuit vulnerabilities, with a specific emphasis on radiation-induced errors. We note, however, that the fault injection approach may face challenges when applied to large machines. This complexity is accentuated when sampling a broad spectrum of circuits that exhibit variations in qubit size, circuit depth, and concurrent fault counts, especially in the context of larger machines and intricate algorithms. Moreover, the quantum vulnerability factor (QVF) metric proposed in [37], can face challenges when assessing circuits reaching the quantum supremacy, typically seen in machines with 50+ qubits [4]. The QVF calculation hinges on the contrast function, which necessitates prior knowledge of $P(A)$, the expected correct state. In the absence of this knowledge, the correct state must be determined through a noise-free simulation. Relying on classical computing for full-state quantum circuit simulation poses a significant challenge, as such methods approach the limits of classical computational capabilities. Alternatively, our research introduces a different approach, formulating a noise model addressing 1- and 2-qubit gate errors, measured errors, and crosstalk errors. Our methodology is crafted with a focus on scalability.

3) ESTIMATED SUCCESS PROBABILITY-BASED

The estimated success probability (ESP), shown in (3), predicts the correct output trial probability by multiplying the SR, or fidelity, of each gate (g_i^e) and measurement (m_i^e) operations, generated by one minus the gate (g_i^e) and measurement (m_i^e) error rate in (2). While ESP considers all circuit operations, the product treats all gate errors that contribute to the final SR estimation equally when some gate errors influence the final circuit outcome more or less than others. As a basic demonstration of the inaccuracy of ESP modeling, the gate success products in (3) commute, whereas most operations in quantum circuits are fixed in ordering [29]. Based on such position differences of the gates on the compiled circuits, gate errors will contribute differently based on their propagation path to the measurement, which influence the estimated SR differently than their original gate error rates. The detailed analysis is presented in Section III-B. On the other hand, the simplicity of the ESP metric has made it frequently applied in quantum compiler design and circuit optimization efforts as a method to predict quantum program success on quantum computers [2], [3], [24], [32], [36], [38], [47]. However, if a better SR estimator was available, the effectiveness of the aforementioned quantum computer optimizations could potentially experience significant improvements.

We used statistical fault injection, ESP_CP [46], and ESP, as illustrated in Fig. 2, to estimate the success rate (SR) of the Bernstein-Vazirani (BV), and quantum Fourier transform (QFT) algorithms across varying scales on three distinct quantum computers. The BV algorithm was selected due to




FIGURE 2. Current success rate estimators performance.




FIGURE 3. ESP for different circuits.

its relatively shallow depth, allowing us to demonstrate the effects of significant algorithm size increments. Conversely, the QFT, characterized by deeper circuits, exhibits only modest size increases. For a comprehensive evaluation, both algorithms were tested with three different input sizes, yielding actual SRs ranging from 70% to 5%. ESP_CP, a variant of ESP, focuses solely on multiplying the SRs of gates situated on the critical path. Unfortunately, the predicted outcomes from both methods show a significant deviation from the actual SR, with discrepancies ranging from 25% to 60% and relative error rates ranging between 70% and 470%. Further compounding the issue, both methods produce SR estimations that diverge sharply from the real machine results as the circuits increase in size, indicating that their tendency toward scaling is not well performed.

B. SOME ERRORS MATTER, WHILE OTHERS DO NOT

The ESP model, upon closer examination, presents potential sources of inaccuracies in SR prediction. Illustrated in Fig. 3, the ESP model accurately predicts the SR only for the middle circuit. When considering a compiled circuit, its final SR results from the product of individual qubit SRs, which are premeasured. Therefore, only errors impacting a measurement gate influence the final output.

For scenarios akin to the first circuit, the ESP model tends to overestimate error rates. Here, the error from the red-boxed Z gate does not influence the measurement gate, meaning the Z gate’s error only impacts the SR of Q2. Yet, this is not captured by the subsequent Q1 measurement gate.

Contrastingly, for situations resembling the third circuit, the ESP model tends to underestimate error rates. Errors originating from the two red-boxed H gates not only affect the measurements of their associated qubits but also influence other qubits via the CNOT gate. Despite this, the gate error to SR transformation, as outlined in (2), only accounts for this error once. Any subsequent error impacts on other qubit measurements arising from error propagation are disregarded in the ESP model.

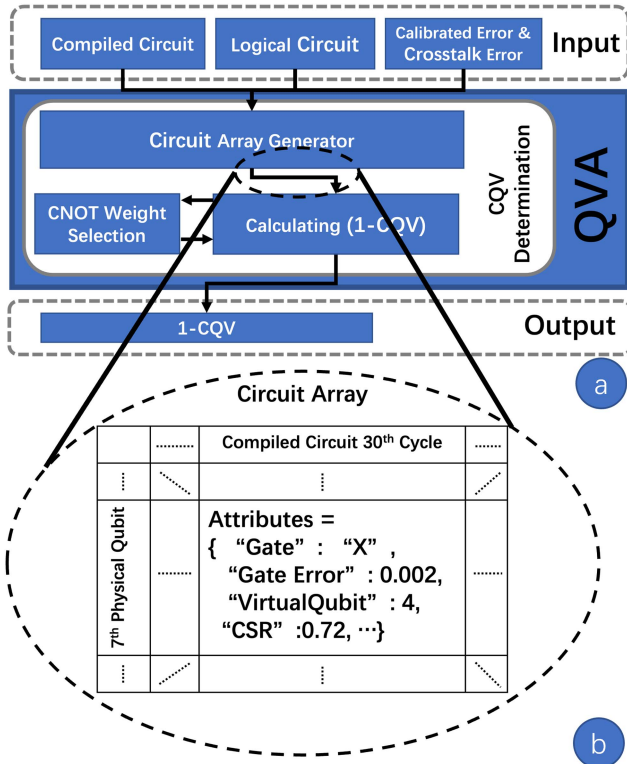


FIGURE 4. QVA workflow and circuit array example.

Such observations underscore that certain errors exert more influence on the final output than others, highlighting the need for a refined approach. This analysis emphasizes the importance of taking into account circuit structure and architectural vulnerabilities when estimating SR on actual quantum computers.

IV. QUANTUM VULNERABILITY ANALYSIS

A. QVA OVERVIEW

In Section III-B, we discover that the SR of quantum computation is the outcome of the SR of each qubit being measured. Each measured qubit's correctness is influenced by gate errors propagated to it. Our proposed QVA is a systematic methodology that follows error propagation. The QVA will estimate the vulnerability of the compiled circuit by performing error modeling based on the error rates from randomized benchmarking calibration.

The QVA generates the cumulative quantum vulnerability (CQV) metric. Definition of CQV: *CQV presents the final circuit's vulnerability by predicting the failure rate (FR) for the compiled circuit.* We emphasize that the CQV will not predict the correct result but the possibility of an incorrect result during runtime. The calculation of $1 - CQV$ represents the estimated SR of a compiled circuit on the target machine calculated with the CQV. In Fig. 4(a), we present the complete workflow of QVA.

B. CQV DETERMINATION

1) CIRCUIT ARRAY GENERATOR

To understand the error propagation path within a quantum circuit during runtime, a connection between when an error occurs and how much it affects the compiled circuit must be established. For more granularity, we quantify the compiled circuit to a finer degree by representing the algorithm at the cycle level. Cycle-level representation for a quantum circuit is analogous to the classical electrical circuit diagram to replace the previous analysis at the level of the complete compiled circuit. The circuit array generator block transfers the compiled circuit further to a 2-D array where each element represents the attributes of the physical qubit at that cycle, as shown in Fig. 4(b). The circuit array records each physical qubit's attribution in every cycle, including the gate type, gate error, associated virtual qubit, its cumulative SR, etc. Based on the cycle level compiled circuit, a snapshot of the operating quantum chip at a given cycle can be linked with the corresponding cycle in the compiled circuit.

2) CALCULATING (1 - CQV)

The CQV methodology aims to predict the failure rate for a given compiled circuit on a designated quantum chip by effectively modeling the errors based on its propagation. The Calculating $1 - CQV$ block of Fig. 4(a) will first receive a circuit array with attributes filled. Next, we perform a crosstalk error calibration based on [34] and update it to the *gate error* by multiplying their SR based on (2). To determine the SR estimation, which is $1 - CQV$, we introduce an algorithm (referenced as Algorithm 1). This algorithm progressively updates the cumulative success rate (CSR) of each physical qubit based on the circuit array, considering the propagation of errors.

The algorithm initiates by setting the CSR for all entries in the circuit array to a perfect score, i.e., $CSR = 1$ (100% success), as depicted in line 3. Subsequent steps, from lines 3 to 16, loop through all the gates, updating CSR values. For single-qubit gates, lines 6 and 7 modify the CSR for that gate by multiplying its total SR with the preceding CSR value of the same qubit from the last cycle.

Complex operations, like the CNOT gate, necessitate a deeper understanding. Here, errors from one qubit can cascade to itself and affect the paired qubit. In lines 8–11, for every occurrence of such two-qubit interactions in a given compiled circuit, we introduce a weight w . This weight, which lies between 0 and 1, signifies the fraction of cumulative error originating from the paired qubit that might propagate via the CNOT gate.

This error propagation model stems from the constraints imposed by the error rates disclosed through randomized benchmarking. The intricacy lies in the fact that these reported error rates cannot be disaggregated into individual types, such as phase, bit, or decoherence errors. Each type behaves differently when channeled through the CNOT gate.

Algorithm 1: Calculate $(1 - CQV)$.

Input: Physical qubits QP ; Compiled Circuit Cycles C ; Weight w ; Circuit Array $CA = [QP][C + 1][Attr.]$;
Output: $(1 - CQV)$

- 1: $(1 - CQV) = 1$
- 2: let $Attr_{qp,c} = CA[qp][c][Attr.]$ for all Attributes
- 3: Initialize $CSR_{qp,0} = 1.00$ for every qubit at first cycle.
- 4: **for** each cycle c from 1 to $C + 1$ **do**
- 5: **for** each physical qubit qp in QP **do**
- 6: **if** the $gate_{qp,c}$ is 1-qubit gate **then**
- 7: $CSR_{qp,c} = g_{qp,c}^s * CSR_{qp,c-1}$
- 8: **else if** the $gate_{qp,c}$ is CNOT gate **then**
- 9: $crosserror = (1 - CSR_{qp',c}) * w$
- 10: $CSR_{qp,c} = g_{qp,c}^s * CSR_{qp,c-1} * (1 - crosserror)$
- 11: **end if**
- 12: **end for**
- 13: **if** any qp in the final swap cycle **then**
- 14: swap the $CSR_{qp,c}$ and $CSR_{qp',c}$ for all swap pairs
- 15: **end if**
- 16: **end for**
- 17: $1 - CQV = \prod_{gate_{qp,c}=Measure} CSR_{qp,c}$
- 18: **return** $1 - CQV$

For a target qubit involved in a CNOT gate, its CSR is computed as a product of its own SR (g^s), its preceding CSR, and the SR inherited from its paired qubit. This inherited SR factors in only the weighted portion of the cumulative error. It is crucial to remember that the SR (or the associated error rate) for any quantum element (qubit, gate, or circuit) can be deduced using (2), which is based on the complementary relationship between success and error rates.

For idle cycles in the quantum circuit, we attribute an error value of zero. In the case of repeated gates, either compiler-introduced or manually inserted by the programmer, we appoint the error value based on the known error rates of such gates in the target machine.

In conclusion, the $1 - CQV$ value, which represents the overall SR prediction, is derived by multiplying the CSRs of all measurement gates. Upon examining Algorithm 1, it becomes evident that its complexity is $O(G)$, where G represents total gate counts for all types. This complexity arises because the algorithm iteratively checks each physical qubit's gate in every cycle to update the corresponding cumulative SR. Consequently, the algorithm scales linearly with the gate count, encompassing all gate types, including ideal gates. This linear scalability of our noise model offers a significant advancement, facilitating both current NISQ-era and future quantum computing endeavors without being constrained by the limitations of classical computational power.

3) CNOT WEIGHT SELECTION

In Section V, we have provided a study to observe the weight selection impact for various compiled circuits with




FIGURE 5. Average success rate prediction comparison between weight as 0 and 1 for benchmarks on IBMQ_Montreal.

different error profiles. Then we used a machine learning-based method to learn the weight value and used it to infer the proper weight in the CQV calculation. For more details. Please refer to Section V.

V. DETERMINING CNOT WEIGHT

To accurately predict the real SR, QVA requires a proper weight value, between zero and one, at compile time to assist with the CQV calculation. Before identifying the value of the weight, we first demonstrate the prediction performance when the weight is set equal to zero or one representing no error or full error crossover the CNOT gate, respectively. The compiled circuits of the experiments are generated from different combinations of compiler settings for four benchmarks at two different algorithm sizes. The CQV calculation is performed using the calibration error and execution results from IBMQ_Montreal on April 1, 2022. As shown in Fig. 5, though the CQV results with weight set to zero, $1 - CQV_0$, are closer to the real SR than ESP, there is still a nontrivial gap between $1 - CQV_0$ and the real SR meaning that some errors are not well represented. Meanwhile, $1 - CQV_1$ sets the weight to one, which makes the predictions close to zero all the time and lose track of the real SR, meaning the errors are being overestimated. The experimental results show that, for those benchmarks, using zero or one as the weight will lead to inaccurate predictions. When brute-force performing the CQV prediction for all the weights with 1% granularity, we found the correlation between the weight value and its corresponding 1-CQV prediction is approximately -1 . In other words, among all the weight values, there will always be one and only one weight value that returns a SR prediction closest to the real SR, which will be labeled as the best weight.

Based on such observation, we calculated the best weight for all compiled circuits generated from combining all the different algorithms, target machines, and compiled strategies. As shown in Fig. 6, we have plotted the best weight against the CNOT count of the compiled circuit. The result is consistent with our expectation—the best weight will be very arbitrary when the CNOT count is low, but as the CNOT count of the compiled circuit increases, the best weight begins to approach zero and shows an overall decreasing trend.




FIGURE 6. Comparing the best weight among all experiments.




FIGURE 7. Graphic neural network-based model layout.




FIGURE 8. GNN training performance for IBMQ_Mumbai.

After analysis, we found that many factors, such as machine error properties, compiled circuit properties, etc., influence the best weight value. To fulfill the need of taking the graph-like machine information and circuit features into consideration, we choose a graph neural network (GNN) [42] and combine it with feed-forward networks to perform the best weight prediction shown in Fig. 7. The node matrix represents the information for every physical qubit, including single-qubit operation error rates. The edge matrix in the figure presents the CNOT error rates for each physical qubit pair. The circuit matrix contains information for each qubit’s operation count, measurement info, and two-qubit operations count. As shown in Fig. 8, by leveraging the GNN model, we can provide the weight value with an average 2% difference from the best weight, which strongly supports an accurate SR prediction in later execution. The trained model will be used to infer the best weight for CQV calculation.

TABLE 1. Benchmarks and Quantum Computer Description

Item	Description
BV	Bernstein-Vazirani
DJ	Deutsch-Jozsa
HS	Hidden Shift
GHZ	Greenberger-Horne-Zeilinger
QFT	Quantum Fourier Transform
QPE	Quantum Phase Estimation
ibmq_montreal	27-qubits with Hexagon
ibmq_toronto	27-qubits with Hexagon
ibmq_mumbai	27-qubits with Hexagon

VI. IMPLEMENTATION

We employ Qiskit [15], a renowned open-source framework for quantum computing, as the foundation for implementing and assessing our QVA methodology. Our work extends Qiskit version 0.34.2, enabling it to execute QVA and compute the (1-CQV) for any specified compiled circuit. Our QVA approach is meticulously crafted to provide an accurate and efficient estimation of the SR for compiled circuits on specific quantum machines. To ensure a diverse set of compiled circuits, we utilize various combinations of compiling policies for each benchmark-machine pairing. This approach integrates a spectrum of error-mitigation techniques at every stage, ensuring comprehensive and robust evaluations. Table 1 describes the backends and benchmarks used in our evaluation. We repeatedly perform all the algorithms that range in scale from 4 input qubits to 15 input qubits over months, which generated 160 K distinctive compiled circuits captured with diverse noise calibration profiles. To focus on meaningful results, we ignore experiments that return a real SR below 0.1%. We chose the state-of-the-art SR estimator ESP for the baseline and ignored the ESP_CP since it underestimates error.

In the experimental flow, we first run the given compiled circuits on their target quantum computers for the default 8192 trials and log the outputs. Then, we use the Qiskit simulator to capture the correct result and generate the SR of the experiment based on the logged output on the classical computer. For scenarios approaching quantum supremacy, where classical computers struggle with full-state quantum circuit simulations, we introduce an alternative method to ascertain correct results, as detailed in Section VII-B. Now the experiment’s real performance is known and named real SR. For each machine, we used the first ten days of data to perform the GNN training on weight selection and to infer the weight to assist CQV prediction for the rest of the experiment data. This offline training is a one-time procedure, and in our experiments, it took approximately an hour on an NVIDIA 3060 GPU. Then we perform ESP and 1-CQV, which estimate the real SR based on the calibrated error of the experiment on the classical computer. The estimated SR generated from ESP and 1-CQV will be compared with the real SR. In addition to quantifying the estimation accuracy directly by performing the absolute difference between the real and estimated SR, we also use the relative prediction




FIGURE 9. Predicted and real success rates of QPE with IBMQ_Montreal Quantum machine. For each compile configuration listed on the x-axis, the variables separated by an underscore are ordered as algorithm input qubits, optimization level, allocation method, and routing method.




FIGURE 10. Average relative predict error for all the benchmarks on single Quantum machine IBMQ_Montreal.

error metric, which uses the absolute difference divided by the real SR, to present the accuracy trend of the method while compensating for increases in the algorithm size resulting in a decrease in real SR.

VII. RESULTS

A. CQV ACCURACY

1) ALGORITHM SIZE WITHIN QUANTUM VOLUME

Here, we present the CQV prediction performance for all six benchmarks on the Quantum machines IBMQ_Montreal, IBMQ_Toronto, and IBMQ_Mumbai. As shown in Fig. 9, we presented the 1-CQV prediction accuracy compared with ESP with varying compiled circuits for a five-qubit QPE algorithm on IBMQ_Montreal on April 5, 2022. The different configurations are guidelines for the compiler to generate the final compiled circuit based on the given logical circuit and target device.

We observe that $1 - CQV$ is much closer to the real SR than the ESP. After being shown to the right of Fig. 9, the average absolute error rate for ESP over all the configurations compared to ground truth SR is 29.8%, while $1 - CQV$ achieves an average error of 4.8%. Such an error rate difference means that $1 - CQV$ achieves an 84% error reduction compared to the ESP, which also means the CQV calculation is adaptable to the variation of errors across different calibration periods and provides excellent predictions.

As shown in Fig. 10, we present the 1-CQV prediction for all the benchmarks on the single machine IBMQ_Montreal. We include experiments with an SR higher than 0.1%. The relative prediction error is produced by dividing the absolute




FIGURE 11. Average relative predict error across all benchmarks and backends.

prediction error by the SR. The trend emerges that 1-CQV prediction outperforms the ESP by achieving an average of six times less relative prediction error rate and the best improvement of 30 times. Meanwhile, the relative prediction error jumps clearly when the algorithm size increases, which follows the nature of the benchmark by employing significantly more two-qubit gates. When increasing algorithm size, not only does the number of two-qubit gates increase, but the number of swapping operations also increases. As shown in Fig. 11, 1-CQV presents a stable and accurate average relative prediction error across all machines and benchmarks. Based on the results, we conclude that the CQV achieved the goal of designing a more precise SR estimator consistently across different dates, machines, and algorithms than the state-of-the-art SR estimator.

2) ALGORITHM SIZE BEYOND QUANTUM VOLUME

From the results shown in Figs. 9–11, $1 - CQV$ proves to predict SR with a closer distance to the real SR, even when it falls in the 10%–0.1% range. The primary reason for the low SR is that the size of the compiled circuits equals or exceeds the desired quantum volume of the target quantum machine. Quantum volume can be defined as the product of the number of virtual qubits and maximum circuit depth supported by the machine.

After making a full-spectrum comparison among all the backends and benchmarks, from the observations of the results, we can say that the CQV has better error modeling than the ESP noise model. Fig. 10 demonstrates that the CQV prediction is stable across different algorithms with a much

TABLE 2. Execution Time Comparison (Seconds)

Algorithm	CX count	ESP	Qiskit	CQV
QFT_5	59	0.00299	1.31142	0.01830
QFT_10	408	0.02004	1.86264	0.05310
QFT_20	2657	0.09802	> 10 mins	0.25227
QFT_50	26408	0.71115	N/P	6.01701
QFT_100	157428	1.83842	N/P	7.82150
QFT_120	208260	3.10630	N/P	20.4157

lower relative error rate. Additionally, the results show that CQV performs much better when the algorithm reaches or exceeds the quantum volume, which is a valuable property when the limited quantum volume is the bottleneck of the current NISQ era. Therefore, we conclude that the $1 - CQV$ prediction is accurate across the full spectrum of algorithm sizes and SRs.

B. SCALABILITY ANALYSIS OF CQV

In the rapidly evolving landscape of quantum computing, the intricacy of quantum circuits is escalating, bringing us closer to addressing real-world challenges. As we navigate this frontier, the need for prediction models that are both accurate and scalable becomes paramount. Our study, therefore, focuses on the scalability of the QVA, examining its efficiency across a spectrum of quantum circuit sizes.

The QFT was chosen as the benchmark for our evaluation. With input sizes spanning from 5 to 120 qubits, it is noteworthy that machines operating with around 50 qubits are on the threshold of quantum supremacy [4]. Our scalability experiments for QVA were conducted based on the topology and error profile of the IBMQ_Washington machine, a state-of-the-art 127-qubit quantum computer.

To ensure our execution time assessment was both transparent and unbiased, we considered only the inference time of the GNN, which provides the CNOT weight value, and the execution time of the CQV noise model. This choice was made because the GNN training is an offline process, executed just once. As shown in Table 2, our CQV approach demonstrated linear scalability, with execution time increasing proportionally with the CNOT gate counts. It is worth mentioning that the execution of the CQV takes place on the CPU, including the pretrained GNN to infer the weight value w . This linear trajectory underscores the potential of our model to efficiently manage complex quantum circuits in the future.

However, the full-state quantum circuit simulator-based approaches face challenges in scalability, particularly with circuits that exceed 32 qubits. This limitation highlights the simulator’s constraints when tasked with emulating large-scale, real-world quantum systems. In contrast, our approach requires only a fraction of the computational power for estimating SRs. To both validate our predictions and refine our noise model, particularly for circuits approaching quantum supremacy, we have devised a novel method: by merging the original circuit with its inverse and using the input states as a benchmark, we can effectively measure the circuit’s




FIGURE 12. Choose compiler policies for HS on IBMQ_Toronto.

performance and fine-tune our noise model. It is pertinent to note potential overfitting for specific circuits due to the reliance on reverse circuits. Nevertheless, we believe our approach adds a valuable technique to the toolbox for exploring circuit vulnerabilities.

Furthermore, while our preliminary results on QVA’s scalability are promising, they represent just the tip of the iceberg. Our model’s design is inherently versatile, unencumbered by specific error rates, gate types, or circuit structures. This adaptability not only allows for the integration of partial error correction techniques in larger devices but also hints at the vast potential for future optimization strategies, further refining the SR estimation process.

VIII. CASE STUDY: CHOOSING COMPILING STRATEGY

The current access modes for quantum computers are either limited free access to small machines or expensive hourly institutional subscriptions to large devices. Naturally, users will want the highest SR with as few executions as possible to save time, money, and access. However, finding the best combination among all the available machines and compiler configurations to achieve the best performance is challenging. The search space will grow enormously when also considering compilation optimizations. Without performing a brute-force execution of all the combinations, identifying the best strategy is challenging. Since CQV is more accurate than ESP, we would naturally ask, could CQV be used to suggest the compiling configuration with the optimal performance? To answer that, we performed both ESP and CQV for all the combinations of compiler configuration at compile time for the HS benchmark and picked the two configurations with the highest estimated SR for both ESP and CQV. Based on the result, the CQV outperforms the ESP across all algorithm sizes. One example of HS at level 2 optimization and six input states is shown in Fig. 12. It is clear that the two choices with the highest 1-CQV estimation not only have less prediction error but also result in the highest real SR. In contrast, the top two high-ranking ESP configurations result in the 4th and 5th best in real SR out of nine combinations. Moreover, it will be the 7th and 8th choice of ESP to pick up the compiler configurations for the highest real SR. For the computation overhead, executing CQV prediction is acceptable since the calculation is done on a classical computer. Furthermore,

since the execution of all the SR predictions is independent of each other, it is possible to perform parallel computing for different compiler strategies, and the individual prediction overhead is discussed in Section VII-B. In this case, the CQV can guide a user to choose a more effective and reliable compiler strategy with a higher SR than ESP. No prediction can be perfect (that would be computationally intractable), but CQV improves prediction enough to be usable for compiler decisions.

IX. CONCLUSION

In the rapidly evolving field of quantum computing, predicting the SR of a quantum circuit remains a challenging task. Existing methodologies often fall short, either by oversimplifying the error model or by not adequately accounting for the intricacies of error propagation within complex quantum circuits. Recognizing this gap, we present the QVA in this article, a robust systematic approach to determining a given computation's CQV. The QVA offers a nuanced, detailed method to estimate the failure rate of a given compiled circuit, considering the effects of individual gate errors, their cumulative influence, and the unique properties of quantum gates such as CNOT. To establish the efficacy of QVA, we subjected it to rigorous validation on cutting-edge quantum machines using well-known benchmarks. The results demonstrated that QVA consistently outperformed the prevalent SR estimator, ESP, and showcased linear scalability. On average, our model exhibited a sixfold reduction in the relative prediction error rate compared to ESP. Such accuracy not only bolsters confidence in our method but also has far-reaching implications at both the hardware and software levels.

ACKNOWLEDGMENT

The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

REFERENCES

- [1] "IBM quantum systems," Accessed: Jul. 28, 2022. [Online]. Available: <https://quantum-computing.ibm.com/services?systems=all>
- [2] N. Acharya and S. M. Saeed, "A lightweight approach to detect malicious/unexpected changes in the error rates of NISQ computers," in *Proc. IEEE ACM Int. Conf. Comput. Aided Des.*, 2020, pp. 1–9, doi: [10.1145/3400302.3415684](https://doi.org/10.1145/3400302.3415684).
- [3] N. Acharya and S.-M. Saeed, "Automated flag qubit insertion for reliable quantum circuit output," in *Proc. IEEE Comput. Soc. Annu. Symp.*, 2021, pp. 431–436, doi: [10.1109/ISVLSI51109.2021.00085](https://doi.org/10.1109/ISVLSI51109.2021.00085).
- [4] F. Arute et al., "Quantum supremacy using a programmable superconducting processor," *Nature*, vol. 574, no. 7779, pp. 505–510, 2019, doi: [10.1038/s41586-019-1666-5](https://doi.org/10.1038/s41586-019-1666-5).
- [5] J. Clarke and F. K. Wilhelm, "Superconducting quantum bits," *Nature*, vol. 453, no. 7198, pp. 1031–1042, 2008, doi: [10.1038/nature07128](https://doi.org/10.1038/nature07128).
- [6] L. DiCarlo et al., "Demonstration of two-qubit algorithms with a superconducting quantum processor," *Nature*, vol. 460, no. 7252, pp. 240–244, 2009, doi: [10.1038/nature08121](https://doi.org/10.1038/nature08121).
- [7] Y. Ding, P. Gokhale, S. F. Lin, R. Rines, T. Propson, and F. T. Chong, "Systematic crosstalk mitigation for superconducting qubits via frequency-aware compilation," in *Proc. 53rd Annu. IEEE/ACM Int. Symp. Microarchitecture*, 2020, pp. 201–214, doi: [10.1109/MICRO50266.2020.00028](https://doi.org/10.1109/MICRO50266.2020.00028).
- [8] X. Fu et al., "eQASM: An executable quantum instruction set architecture," in *Proc. IEEE Int. Symp. High Perform. Comput. Archit.*, 2019, pp. 224–237, doi: [10.1109/HPCA.2019.00040](https://doi.org/10.1109/HPCA.2019.00040).
- [9] X. Fu et al., "A microarchitecture for a superconducting quantum processor," *IEEE Micro*, vol. 38, no. 3, pp. 40–47, May/Jun. 2018, doi: [10.1109/MM.2018.032271060](https://doi.org/10.1109/MM.2018.032271060).
- [10] X. Fu et al., "An experimental microarchitecture for a superconducting quantum processor," in *Proc. 50th Annu. IEEE/ACM Int. Symp. Microarchitecture*, 2017, pp. 813–825, doi: [10.1145/3123939.3123952](https://doi.org/10.1145/3123939.3123952).
- [11] J. P. Gaebler et al., "Randomized benchmarking of multiqubit gates," *Phys. Rev. Lett.*, vol. 108, no. 26, 2012, Art. no. 260503, doi: [10.1103/PhysRevLett.108.260503](https://doi.org/10.1103/PhysRevLett.108.260503).
- [12] J. M. Gambetta et al., "Characterization of addressability by simultaneous randomized benchmarking," *Phys. Rev. Lett.*, vol. 109, no. 24, 2012, Art. no. 240504, doi: [10.1103/PhysRevLett.109.240504](https://doi.org/10.1103/PhysRevLett.109.240504).
- [13] P. Gokhale et al., "Optimization of simultaneous measurement for variational quantum eigensolver applications," in *Proc. IEEE Int. Conf. Quantum Comput. Eng.*, 2020, pp. 379–390, doi: [10.1109/QCE49297.2020.00054](https://doi.org/10.1109/QCE49297.2020.00054).
- [14] P. Gokhale, A. Javadi-Abhari, N. Earnest, Y. Shi, and F. T. Chong, "Optimized quantum compilation for near-term algorithms with openpulse," in *Proc. 53rd Annu. IEEE/ACM Int. Symp. Microarchitecture*, 2020, pp. 186–200, doi: [10.1109/MICRO50266.2020.00027](https://doi.org/10.1109/MICRO50266.2020.00027).
- [15] I. Q. Group, "Open-source quantum development," Accessed: Apr. 16, 2021. [Online]. Available: <https://qiskit.org/>
- [16] L. K. Grover, "A fast quantum mechanical algorithm for database search," in *Proc. 28th Annu. ACM Symp. Theory Comput.*, 1996, pp. 212–219, doi: [10.1145/237814.237866](https://doi.org/10.1145/237814.237866).
- [17] IBM, "IBM quantum," Accessed: Apr. 16, 2021. [Online]. Available: <https://quantum-computing.ibm.com/>
- [18] K. Kechedzhi et al., "Effective quantum volume, fidelity and computational cost of noisy quantum processing experiments," *Future Gen. Comput. Syst.*, vol. 153, pp. 431–441, Apr. 2023, doi: [10.1016/j.future.2023.12.002](https://doi.org/10.1016/j.future.2023.12.002).
- [19] R. LaRose et al., "Mitiq: A software package for error mitigation on noisy quantum computers," *Quantum*, vol. 6, 2022, Art. no. 774, doi: [10.22331/q-2022-08-11-774](https://doi.org/10.22331/q-2022-08-11-774).
- [20] T. LeCompte, F. Qi, and L. Peng, "Robust cache-aware quantum processor layout," in *Proc. Int. Symp. Reliable Distrib. Syst.*, 2020, pp. 276–287, doi: [10.1109/SRDS51746.2020.00035](https://doi.org/10.1109/SRDS51746.2020.00035).
- [21] T. LeCompte, F. Qi, X. Yuan, N.-F. Tzeng, M. H. Najaf, and L. Peng, "Graph neural network assisted quantum compilation for qubit allocation," in *Proc. ACM Great Lakes Symp. VLSI*, 2023, pp. 415–419, doi: [10.1145/3583781.3590300](https://doi.org/10.1145/3583781.3590300).
- [22] T. LeCompte, F. Qi, X. Yuan, N.-F. Tzeng, M. H. Najafi, and L. Peng, "Machine learning-based qubit allocation for error reduction in quantum circuits," *IEEE Trans. Quantum Eng.*, vol. 4, 2023, Art. no. 3101414, doi: [10.1109/TQE.2023.3301899](https://doi.org/10.1109/TQE.2023.3301899).
- [23] G. Li, Y. Ding, and Y. Xie, "Tackling the qubit mapping problem for NISQ-era quantum devices," in *Proc. 24th Int. Conf. Architectural Support Program. Lang. Operating Syst.*, 2019, pp. 1001–1014, doi: [10.1145/3297858.3304023](https://doi.org/10.1145/3297858.3304023).
- [24] G. Li, A. Wu, Y. Shi, A. Javadi-Abhari, Y. Ding, and Y. Xie, "Paulihedral: A generalized block-wise compiler optimization framework for quantum simulation kernels," in *Proc. 27th ACM Int. Conf. Architectural Support Program. Lang. Operating Syst.*, 2022, pp. 554–569, doi: [10.1145/3503222.3507715](https://doi.org/10.1145/3503222.3507715).
- [25] J. Liu and H. Zhou, "Reliability modeling of NISQ-era quantum computers," in *Proc. IEEE Int. Symp. Workload Characterization*, 2020, pp. 94–105, doi: [10.1109/IISWC50251.2020.00018](https://doi.org/10.1109/IISWC50251.2020.00018).
- [26] E. Magesan, J. Gambetta, and J. Emerson, "Robust randomized benchmarking of quantum processes," 2010, *arXiv:1009.3639*, doi: [10.1103/PhysRevLett.106.180504](https://doi.org/10.1103/PhysRevLett.106.180504).
- [27] E. Magesan, J. M. Gambetta, and J. Emerson, "Characterizing quantum gates via randomized benchmarking," *Phys. Rev. A*, vol. 85, no. 4, 2012, Art. no. 042311, doi: [10.1103/PhysRevA.85.042311](https://doi.org/10.1103/PhysRevA.85.042311).
- [28] J. Majer et al., "Coupling superconducting qubits via a cavity bus," *Nature*, vol. 449, no. 7161, pp. 443–447, 2007, doi: [10.1038/nature06184](https://doi.org/10.1038/nature06184).

- [29] A. Matsuo, W. Hattori, and S. Yamashita, "Reducing the overhead of mapping quantum circuits to IBM Q system," in *Proc. IEEE Int. Symp. Circuits Syst.*, 2019, pp. 1–5, doi: [10.1109/ISCAS.2019.8702439](https://doi.org/10.1109/ISCAS.2019.8702439).
- [30] N. Moll et al., "Quantum optimization using variational algorithms on near-term quantum devices," *Quantum Sci. Technol.*, vol. 3, no. 3, 2018, Art. no. 030503, doi: [10.1088/2058-9565/aab822](https://doi.org/10.1088/2058-9565/aab822).
- [31] P. Murali, J. M. Baker, A. Javadi-Abhari, F. T. Chong, and M. Martonosi, "Noise-adaptive compiler mappings for noisy intermediate-scale quantum computers," in *Proc. 24th Int. Conf. Architectural Support Program. Lang. Operating Syst.*, 2019, pp. 1015–1029, doi: [10.1145/3297858.3304075](https://doi.org/10.1145/3297858.3304075).
- [32] P. Murali, N. M. Linke, M. Martonosi, A. J. Abhari, N. H. Nguyen, and C. H. Alderete, "Architecting noisy intermediate-scale quantum computers: A real-system study," *IEEE Micro*, vol. 40, no. 3, pp. 73–80, May/June 2020, doi: [10.1109/MM.2020.2985683](https://doi.org/10.1109/MM.2020.2985683).
- [33] P. Murali, N. M. Linke, M. Martonosi, A.-J. Abhari, N. H. Nguyen, and C. H. Alderete, "Full-stack, real-system quantum computer studies: Architectural comparisons and design insights," in *Proc. ACM/IEEE 46th Annu. Int. Symp. Comput. Archit.*, 2019, pp. 527–540, doi: [10.1145/3307650.3322273](https://doi.org/10.1145/3307650.3322273).
- [34] P. Murali, D. C. McKay, M. Martonosi, and A. Javadi-Abhari, "Software mitigation of crosstalk on noisy intermediate-scale quantum computers," in *Proc. 25th Int. Conf. Architectural Support Program. Lang. Operating Syst.*, 2020, pp. 1001–1016, doi: [10.1145/3373376.3378477](https://doi.org/10.1145/3373376.3378477).
- [35] M. A. Nielsen and I. Chuang, *Quantum Computation and Quantum Information*. MA, USA: Cambridge, Univ. Press, 2002, doi: [10.1017/CBO9780511976667](https://doi.org/10.1017/CBO9780511976667).
- [36] S. Nishio, Y. Pan, T. Satoh, H. Amano, and R. Van Meter, "Extracting success from IBM's 20-qubit machines using error-aware compilation," *ACM J. Emerg. Technol. Comput. Syst.*, vol. 16, no. 3, pp. 1–25, 2020, doi: [10.1145/3386162](https://doi.org/10.1145/3386162).
- [37] D. Oliveira, E. Giusto, B. Baheri, Q. Guan, B. Montrucchio, and P. Rech, "A systematic methodology to compute the quantum vulnerability factors for quantum circuits," *IEEE Trans. Dependable Secure Comput.*, early access, Sep. 12, 2023, doi: [10.1109/TDSC.2023.3313934](https://doi.org/10.1109/TDSC.2023.3313934).
- [38] T. Patel, B. Li, R. B. Roy, and D. Tiwari, "UREQA: Leveraging operation-aware error rates for effective quantum circuit mapping on NISQ-era quantum computers," in *Proc. USENIX Annu. Tech. Conf.*, 2020, pp. 705–711. [Online]. Available: <https://www.usenix.org/system/files/atc20-patel.pdf>
- [39] J. Poyatos, J. I. Cirac, and P. Zoller, "Complete characterization of a quantum process: The two-bit quantum gate," *Phys. Rev. Lett.*, vol. 78, no. 2, 1997, Art. no. 390, doi: [10.1103/PhysRevLett.78.390](https://doi.org/10.1103/PhysRevLett.78.390).
- [40] J. Preskill, "Quantum computing in the NISQ era and beyond," *Quantum*, vol. 2, 2018, Art. no. 79, doi: [10.22331/q-2018-08-06-79](https://doi.org/10.22331/q-2018-08-06-79).
- [41] S. Resch, S. Tannu, U. R. Karpuzcu, and M. Qureshi, "A day in the life of a quantum error," *IEEE Comput. Archit. Lett.*, vol. 20, no. 1, pp. 13–16, Jan.–Jun. 2020, doi: [10.1109/LCA.2020.3045628](https://doi.org/10.1109/LCA.2020.3045628).
- [42] F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini, "The graph neural network model," *IEEE Trans. Neural Netw.*, vol. 20, no. 1, pp. 61–80, Jan. 2008, doi: [10.1109/TNN.2008.2005605](https://doi.org/10.1109/TNN.2008.2005605).
- [43] Y. Shi et al., "Optimized compilation of aggregated instructions for realistic quantum computers," in *Proc. 24th Int. Conf. Architectural Support Program. Lang. Operating Syst.*, 2019, pp. 1031–1044, doi: [10.1145/3297858.3304018](https://doi.org/10.1145/3297858.3304018).
- [44] P. W. Shor, "Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer," *SIAM Rev.*, vol. 41, no. 2, pp. 303–332, 1999, doi: [10.1137/S0036144598347011](https://doi.org/10.1137/S0036144598347011).
- [45] R. Stassi, M. Cirio, and F. Nori, "Scalable quantum computer with superconducting circuits in the ultrastrong coupling regime," *npj Quantum Inf.*, vol. 6, no. 1, pp. 1–6, 2020, doi: [10.1038/s41534-020-00294-x](https://doi.org/10.1038/s41534-020-00294-x).
- [46] S. Tannu and M. Qureshi, "Not all qubits are created equal: A case for variability-aware policies for NISQ-era quantum computers," in *Proc. 24th Int. Conf. Architectural Support Program. Lang. Operating Syst.*, 2019, pp. 987–999, doi: [10.1145/3297858.3304007](https://doi.org/10.1145/3297858.3304007).
- [47] S. S. Tannu and M. Qureshi, "Ensemble of diverse mappings: Improving reliability of quantum computers by orchestrating dissimilar mistakes," in *Proc. 52nd Annu. IEEE/ACM Int. Symp. Microarchitecture*, 2019, pp. 253–265, doi: [10.1145/3352460.3358257](https://doi.org/10.1145/3352460.3358257).
- [48] S. S. Tannu and M. K. Qureshi, "Mitigating measurement errors in quantum computers by exploiting state-dependent bias," in *Proc. 52nd Annu. IEEE/ACM Int. Symp. Microarchitecture*, 2019, pp. 279–290, doi: [10.1145/3352460.3358265](https://doi.org/10.1145/3352460.3358265).