IEEE Transactions on Quantum Engineering

Received June 1, 2021; revised February 21, 2022; accepted March 2, 2022; date of publication March 23, 2022; date of current version May 10, 2022.

Digital Object Identifier 10.1109/TQE.2022.3161197

Pauli Error Propagation-Based Gate Rescheduling for Quantum Circuit Error Mitigation

VEDIKA SARAVANAN[1] (Student Member, IEEE), AND SAMAH M. SAEED[1] (Member, IEEE)

Department of Electrical Engineering, City College of New York, City University of New York, New York, NY 10031 USA

Corresponding author: Samah M. Saeed (e-mail:ssaeed@ccny.cuny.edu).

ABSTRACT Noisy intermediate-scale quantum algorithms, which run on noisy quantum computers, should be carefully designed to boost the output state fidelity. While several compilation approaches have been proposed to minimize circuit errors, they often omit the detailed circuit structure information that does not affect the circuit depth or the gate count. In the presence of spatial variation in the error rate of the quantum gates, adjusting the circuit structure can play a major role in mitigating errors. In this article, we exploit the freedom of gate reordering based on the commutation rules to show the impact of gate error propagation paths on the output state fidelity of the quantum circuit, propose advanced predictive techniques to project the success rate of the circuit, and develop a new compilation phase postquantum circuit mapping to improve its reliability. Our proposed approaches have been validated using a variety of quantum circuits with different success metrics, which are executed on IBM quantum computers. Our results show that rescheduling quantum gates based on their error propagation paths can significantly improve the fidelity of the quantum circuit in the presence of variable gate error rates.

INDEX TERMS Commutation rules, error propagation, gate rescheduling, noisy intermediate-scale quantum (NISQ) computer, Pauli errors, quantum circuit, quantum circuit mapping, reliability.

I. INTRODUCTION

Noisy intermediate-scale quantum (NISQ) computers, which have tens to hundreds of quantum bits (qubits), are anticipated to accelerate the computational power of classical computers for certain classes of problems. However, the noisy nature of their qubits and operations degrades the reliability of NISQ systems. Thus, errors should be mitigated to improve the fidelity of NISQ algorithms, which necessitates error characterization and circuit design approaches to minimize the impact of errors on the quantum circuit output. Randomized benchmarking is typically used to characterize different sources of noise in the quantum hardware, including gate, measurement, and decoherence errors [1]. The computed error rates are used by noise-aware quantum compilers to generate physical quantum circuits with high fidelity.

While quantum noise is very complex, simplified error models based on error rates collected during the calibration process are typically used for quantum circuit compilation [2]. Depending on the permissible access to the quantum computer, more advanced error characterization can be unavailable. A predictive analysis capable of delivering a projection with reasonable accuracy about the circuit error rates, despite the lack of a complete and accurate error model is also required. A predictive technique, which is computationally fast and applicable to many quantum computing technologies, and takes into account the circuit structure can enhance the compilation process [3].

In this article, we propose error propagation-based predictive techniques to project the quantum circuit success rate primarily based on gate errors. We take into account the propagation path of different gate errors in the circuit. While the existing estimated success probability (ESP) typically used by noise-aware quantum compilers to generate physical quantum circuits incorporates the gate error rates into its model, it ignores the order of the gates that do not affect the circuit depth or the gate count. Our proposed predictive techniques can differentiate between different physical implementations of the same quantum circuit, which use the same physical qubits, share the same set of quantum gates, and have the same depth, but vary in the order of quantum

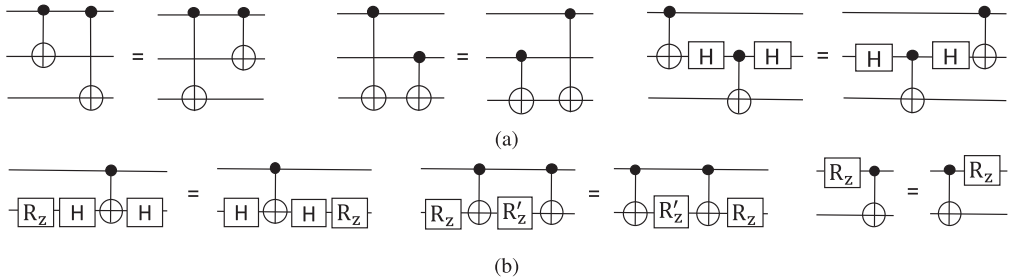

FIGURE 1. Commutation rules of (a) CNOT and (b) R_z gates [4].

gates in the circuit and their corresponding error rates. We experimentally show the advantage of considering the order of the physical quantum circuit gates in the presence of significant variation in the gate errors. Accordingly, we also propose gate rescheduling algorithms at different quantum circuit abstractions to maximize the fidelity of the quantum circuit. Our proposed approaches can be integrated with other gate rescheduling algorithms that reduce the gate count and the circuit depth. To the best of our knowledge, no such analysis has been proposed to incorporate the impact of the error propagation paths to the quantum gate scheduling procedure. The main contributions of this article are as follows.

1) We propose a weighted ESP (WESP) metric based on the quantum circuit structure.
2) We propose new quantum circuit gate rescheduling algorithms at different design levels, including complex and elementary gate-level, based on our proposed WESP metric to boost the fidelity of the quantum circuit output.
3) We experimentally validate the effectiveness of our proposed metrics and gate rescheduling algorithms using a variety of quantum circuits executed on different IBM quantum computers.

The rest of this article is organized as follows. Section II provides a background on the quantum circuit compilation, quantum hardware errors, an application of NISQ computers, and different success criteria for quantum circuit evaluation. Section III discusses related works on quantum circuit mapping and gate reordering approaches. Section IV shows the implication of the gate error propagation path on the output-state fidelity of the quantum circuit in the presence of different gate errors. Section V provides our proposed reliability metric and rescheduling algorithms. Section VI validates the effectiveness of our approaches through several experiments. Finally, Section VII concludes this article.

II. BACKGROUND

A. QUANTUM CIRCUIT COMPILATION

Quantum algorithms are described using quantum circuits, which are executed on the quantum computer. A quantum circuit comprises quantum gates, which change the state of the qubits. Single-qubit gates operate on a single qubit, such as the Hadamard gate, which creates a superposition state, and the R_z gate, which rotates the qubit around the z-axis. A multiqubit gate, such as controlled NOT (CNOT), entangles two qubits [5]. To enable a quantum circuit execution on the targeted NISQ computer, the circuit has to go through several compilation steps. Complex gates should be constructed using elementary gates supported by the NISQ architecture. To reduce the depth and the gate count of the circuit, several gate-level optimization techniques are applied, including template matching- [6] and gate reordering-based [7] techniques. In the former one, a cascade of quantum gates is substituted with a subcircuit with a lower gate count, while in the latter one, consecutive gates cancel each other based on the commutation rules of the quantum gates [4], which are defined as follows.

Definition 1: Let U_1 and U_2 be two unitary matrices. U_1 and U_2 are said to be commutative if $U_1U_2 = U_2U_1$ for any input state. Commutation rules of different quantum gates are shown in Fig. 1.

Next, the physical qubits of the quantum circuit are allocated, and their gates are scheduled to meet the constraints of the quantum architecture. This process is referred to as quantum circuit mapping. A main challenge in the mapping process is the restricted qubits connectivity in superconducting quantum architectures. Two qubit gates can be applied to certain pairs of physical qubits described using the coupling graph of the quantum architecture, in which nodes represent qubits and edges show the connectivity between different qubits. To enable arbitrary multiqubit gates in the presence of the coupling constraint, SWAP operations consisting of three CNOT gates are used for gate scheduling, resulting in an excessive gate count and a large circuit depth. Efficient mapping approaches are used to generate physical quantum circuits with as a minimum gate count as possible to be executed on the quantum computer. These approaches often represent the quantum circuit as a directed acyclic graph, referred to as a gate dependency graph, where each node represents a gate and each edge represents a direct dependency between two gates. The dependency graph is divided into levels, where each level consists of quantum gates applied simultaneously in the same circuit layer. All the nodes reachable from a given node i in the graph are gates, which depend on the g_i gate (g_i reachable gates).

Example 1: Fig. 2 shows an example of a quantum circuit and its corresponding gate dependency graph, in which the

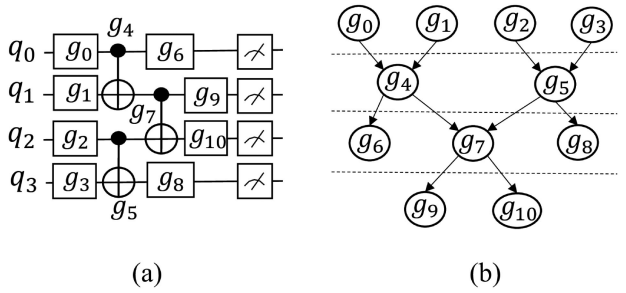

FIGURE 2. (a) Quantum circuit and (b) its corresponding gate dependency graph.

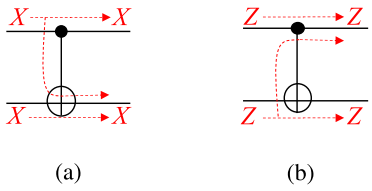

FIGURE 3. Propagation path of (a) *X* and (b) *Z* Pauli gate errors through a CNOT gate.

gates in each circuit layer are mapped to the graph nodes applied at the same level. For each quantum gate, the list of reachable gates is extracted from the gate dependency graph. For example, there are four reachable gates from g_4 gate, which are g_6, g_7, g_9, and g_{10} gates.

B. RELIABILITY OF QUANTUM DEVICES

Different sources of noise contribute to the noise complexity of NISQ systems. The elevated noise levels endanger the reliability of quantum circuits. Gate errors are a major source of quantum computing noise. They are modeled as Pauli gates that operate on a single qubit. They are represented as X (a π rotation around the x-axis), Y (a π rotation around the y-axis), Z (a π rotation around the z-axis), and I (an identity matrix) single-qubit gates [8]. The propagation of some of the Pauli errors through the CNOT gate is shown in Fig. 3. The measurement/read-out error can occur while measuring the output state of the quantum circuit. A qubit can lose its state after a period of time, resulting in decoherence errors. Quantum gates, which act on different qubits, simultaneously can cause crosstalk errors.

As NISQ computers do not support quantum error correction, they often rely on noise-aware quantum compilers to mitigate errors. The quantum compiler should maximize the success probability of the physical quantum circuit. An ESP can be used by a noise-aware quantum compiler to select the physical qubits of the quantum circuit and schedule the quantum gates. Given the gate (e_{g_i}) and measurement (e_{m_i}) errors for all the quantum circuit gates, and its qubits, ESP is computed as $\prod_{i=0}^{G-1}(1-e_{g_i}) \times \prod_{i=0}^{Q-1}(1-e_{m_i})$ [2]. Randomized benchmarking sequences are executed on the quantum hardware to compute single- and two-qubit gate errors based on the average sequence errors. The qubit measurement/read-out error is computed as the average measurement errors 0 and 1, in which the qubit is initially assigned to states 1 and 0, respectively.

C. QUANTUM APPROXIMATION OPTIMIZATION ALGORITHM

The quantum approximation optimization algorithm (QAOA) is a hybrid quantum-classical algorithm for solving optimization problems. It relies not only on the output of the quantum circuit but also on the classical optimizer that updates the circuit parameters to improve the solution of the optimization problem [9]. The objective of the algorithm is to minimize/maximize a cost function [C(X)] of a given problem described using the following Hamiltonian: $H = \sum_{x\in 0,1^n} C(x)|x\rangle\langle x|$. A QAOA quantum circuit contains two main components: 1) phase separation; and 2) mixing operations, which are applied p times repeatedly with (γ_1, β_1),..., (γ_p, β_p) parameters. The phase separation can be represented using ZZ gates while mixing operations are represented using single-qubit gate rotation around the x-axis (R_x). An example of a problem to be solved using the QAOA algorithm is the maximum cut (Max-Cut) problem. For a graph with n vertices, the solution of the Max-Cut problem divides the graph into two subsets such that the total weight of the edges between those subsets is as maximum as possible. The graph vertices are described as a string of n qubits, in which each qubit can be measured as 0 or 1. The QAOA circuit output provides a candidate solution to the Max-Cut problem. Using an iterative procedure, we can find the n-qubit string, which maximizes the total weight of the edges between the two subsets, referred to as the cost function.

D. QUANTUM CIRCUIT SUCCESS METRICS

Different success criteria/metrics have been proposed to show the impact of the quantum compilation approaches on the quantum circuit output. For a single-output quantum circuit, the probability of successful trails (PSTs) measures the probability of the correct output of the quantum circuit as $\frac{\text{Number of successful trials}}{\text{Total number of trials}}$ [10]. For hybrid quantum algorithms, application-specific success criteria are used. For example for QAOA, which maximizes or minimizes a cost function, the approximation ratio (AR) of the cost function quantifies the success of the quantum circuit [11], [12]. It is defined as $\frac{\text{Mean of cost function over all sampled output}}{\text{Maximum cost function value}}$. Given AR, we can compute the approximation ratio gap (ARG) to quantify how close is the output of the actual execution of QAOA circuit to the simulation result [11]. It is computed as $\frac{\text{AR of simulation - AR of execution}}{\text{AR of simulation}} \times 100$. The smaller the gap is, the more reliable the output of the QAOA circuit is.

III. RELATED WORK

A. QUANTUM CIRCUIT MAPPING

Since quantum circuit mapping is an NP-complete problem [13], traversing the search space is speculative for quantum architectures with a large number of qubits. Various

heuristics-based techniques have been proposed to address this problem [14]. Some of these techniques target minimizing the number of circuit SWAP operations [15]–[19], while others focus on minimizing the circuit error rates [10], [20]. For example, quantum circuit mapping can be achieved by using BRIDGE and SWAP operations as transformation rules inserted based on dynamic programming or heuristic-based look-ahead schemes [15]. The BRIDGE gates are composed of a sequence of four CNOT gates. SAT solver minimizes the number of SWAP operations in the physical quantum circuit at the cost of extensive computation overhead (e.g., [16] and [17]). The proposed technique in [16] divides the quantum circuit into subcircuits that require no SWAP operations within each subcircuit. Next, the SAT solver is used to construct the quantum circuit by inserting a minimum number of SWAP operations between the subcircuits that share qubits. In [17], a logical quantum circuit was mapped to a physical quantum circuit using the SAT solver too. Constraints are added to the SAT solver based on the coupling graph of the quantum architecture to reduce its complexity at the expense of the gate count. An A* algorithm searches for the shortest path in the coupling graph between qubits to implement a two-qubit gate, and thus, reduce the number of SWAP operations in the quantum circuit [18]. To address the difference in the error rates of various qubits, a mapping technique has been proposed based the on Dijkstra algorithm, which schedules quantum gates and allocates physical qubits with low error rates [10]. The success probability of the physical quantum circuit can be further improved by searching for isomorphic subgraphs in the quantum architecture's coupling constraint with the highest ESP [20]. Other mapping approaches address decoherence, correlated, and unexpected errors (e.g., [21]–[25]). To mitigate crosstalk errors, they should be characterized first, followed by an efficient mapping approach that selectively serializes circuit gates, while satisfying the required coherence time [26]. Most recently, quantum circuit mapping approaches that target distributed quantum architectures have also been proposed [27], [28].

B. GATE REORDERING OF QUANTUM CIRCUITS

A quantum circuit consists of layers of gates that can be applied in parallel. Consecutive commuting gates can be reordered while preserving the state of the quantum system [4], [29]. The commutation rules of the gates have been utilized to reduce the quantum circuit depth, gate count [4], [29]–[31], and optimize the control pulses [32]. A heuristics-based approach has been developed to perform a gate-level optimization by reordering quantum gates to reduce the circuit gate count [4]. A two-step approach has been proposed to reduce the depth of the quantum circuit by constructing a dependency graph based on the circuit gates and then swapping the gates that satisfy the commutation rules to reduce the circuit depth [29]. A compilation approach has been developed to enhance the performance of QAOA circuits by reordering the complex gates [11], [30]. Due to the commutation relation, most of the dependencies are ignored. Instead, a two-level search process is adopted to minimize the circuit depth using a breadth-first search algorithm, which finds a possible two-layer interchanges in each iteration until the desired depth is achieved. A layout synthesis algorithm has been proposed to minimize the circuit depth by reordering complex gates of QAOA circuits using a satisfiability modulo theories solver [31]. An optimized pulse-level compilation approach has been proposed to address the inefficiency of the standard gate compilation process that directly translates the logical instructions to control pulses [32]. In this approach, the gates are reordered based on commutation rules, and then a small set of gates are aggregated to a larger operation by finding the optimal control pulse based on the gradient descent method. Gate reordering has also been applied to minimize the impact of decoherence errors [33]. Faults in the form of additional gates due to decoherence errors are injected into the quantum circuit to analyze their impact on the circuit output. Accordingly, gates are reordered at the circuit locations, in which decoherence errors significantly affect the circuit output. In this article, we propose additional compilation layers applied to the physical quantum circuit postmapping. We restrict the search space of our gate rescheduling algorithms to maintain the gate count and the circuit depth of the optimized quantum circuit using the previously proposed compilation approaches. Thus, our work can be incorporated with the other compilation/optimization layers to further improve the output state fidelity of the quantum circuit.

IV. IMPLICATION OF THE GATE ERROR PROPAGATION PATH ON THE OUTPUT STATE FIDELITY

Stochastic Pauli noise is a primary factor to estimate the circuit error rates [2]. While tracking the precise impact of the error propagation on the circuit can be computationally expensive in the presence of universal gates [8], two-qubit gates enable error propagation across different qubits. Thus, the location of the quantum gates especially the two-qubit gates can affect the circuit error rates. For the sake of simplicity, we assume that the two-qubit gate (CNOT) error impacts the control and the target qubits, and thus, propagates to the gates applied next to these qubits.

Example 2: To illustrate the gate error propagation through the quantum circuit, let's consider the example provided in Fig. 4 . Two equivalent subcircuits using the commutation rules of quantum gates and their corresponding gate dependency graphs are provided. We use a dotted line to show the output qubits affected by the gates in the dependency graph. Every gate in the subcircuit can produce Pauli errors that propagate to the output qubits through the gate dependency graph [8]. In Fig. 4(a), q_0, q_1, and q_2 output qubits are affected by $\{g_0\}$, $\{g_0, g_1, g_2, g_3\}$, and $\{g_0, g_1, g_2\}$ gate errors, respectively. On the other hand, in Fig. 4(b), q_0, q_1, and q_2 output qubits are affected by $\{g_0, g_1, g_2, g_3\}$, $\{g_0, g_1, g_2, g_3\}$, and $\{g_1, g_2\}$ gate errors, respectively. Thus, while the two subcircuits in Fig. 4 are equivalent, their output

IEEE Transactions on Quantum Engineering

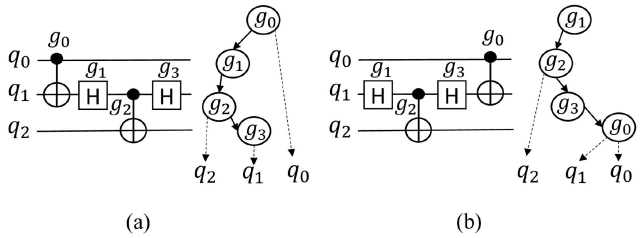

FIGURE 4. (a) Quantum subcircuit and its gate dependency graph and (b) its corresponding equivalent subcircuit and its gate dependency graph.

qubits q_0 and q_2 are susceptible to different sets of gate errors.

Due to the spatial variation in error rates of the quantum hardware, gates applied to different qubits result in different error rates. Although all gate errors affect the output fidelity of the quantum circuit, the two-qubit gate errors are higher than the single-qubit gate errors, and thus, have more impact on the output fidelity of the quantum circuit. Reordering quantum gates, while maintaining the circuit functionality, changes not only the error propagation paths but also the gate error impact on the output state of the circuit. In the previous example, if g_0 gate has the highest error rate in the circuit, which significantly deviates from other gate errors in the circuit, applying g_0 gate at the first layer of the subcircuit as shown in Fig. 4(a) will result in spreading g_0 error to all the output qubits, unlike the equivalent subcircuit in Fig. 4(b), in which g_0 error affects only q_0 and q_1 output qubits.

Example 3: Fig. 5 shows the impact of gate reordering on the quantum circuit output fidelity. A four-qubit BV circuit is mapped to satisfy the coupling constraint of the IBM Q Santiago architecture. Fig. 5(a) represents the coupling graph and the error rates of the IBM Q Santiago architecture, in which the values in the upper and lower part of each node represent the read-out and single-qubit errors, respectively, while the edge label is the two-qubit gate error applied to the corresponding pair of qubits. All error rates are multiplied by 10^{-3}. Two physical implementations of the BV quantum circuit: 1) M1 and 2) M2 are provided in Fig. 5(b) and (c), respectively. The difference between the two physical implementations of the circuit is the location of the first two CNOT gates of the quantum circuit, which can be reordered based on the commutation rules. The CNOT gate highlighted in red color has the highest two-qubit gate error rate. According to the order of the quantum gates, the error propagation path also changes as shown in Fig. 5(b) and (c). We demonstrate the impact of the gate reordering on the quantum circuit output by executing both M1 and M2 quantum circuits on the IBM Q Santiago architecture, and reporting their PST, as shown in Fig. 5(d). To eliminate the impact of gate reordering on decoherence errors, we insert barriers before and after applying all the two-qubit gates in the two circuit implementations in addition to a barrier after the second CNOT gate of M2 quantum circuit in Fig. 5(c), which ensure that the qubits in the two implementations are used for the same period of time. The PST values are 0.784 and 0.861 for the BV quantum circuit generated using the M1 and M2 mapping, respectively, and executed for 8192 shots/trials. Our observation indicates that reordering gates with higher error rates to earlier circuit layers will result in error propagation to a larger number of circuit gates, and thus, degrade the output state fidelity. While PST is different for both mapping, the ESP of both M1 and M2 mapping is 0.7945, which implies the need to consider the gate error propagation path for estimating the success rate of the circuit, and thus refining the quantum circuit mapping policies.

V. ERROR PROPAGATION-BASED GATE RESCHEDULING

A. WEIGHTED-ESTIMATED SUCCESS PROBABILITY

ESP enables quick and easy estimation of the circuit success probability. Yet, it overlooks the circuit structure since different orders of the quantum gates yield the same ESP. To account for the penalty of rescheduling the same set of physical quantum gates with variable error rates on the output state fidelity of the quantum circuit, we propose a new metric referred to as WESP, which can be used after generating physical quantum circuits using ESP. It is defined as

$$\text{WESP} = \prod_{i=0}^{G-1} (1 - (\lambda_i + e_{g_i})) \times \prod_{i=0}^{Q-1} (1 - e_{m_i}).$$

λ_i tunes the error rate of the corresponding g_i gate to account for the gate position as

$$\lambda_i = w_i(e_{g_i} - \min_{g_i \in G} e_{g_i})$$

where w_i and $\min_{g_i \in G} e_{g_i}$ are the weight of the ith gate and the minimum gate error rate in the quantum circuit, respectively. Instead of multiplying the weight of each gate by its error rate, which can significantly reduce the gate error of very noisy gates applied in later layers of the circuit, we multiply the weight with the difference between the current gate error and the minimum gate error in the circuit to account for the variation in the gate errors. The resulting number is added to the gate error. To compute w_i, the gate dependency graph of the circuit is constructed first. For each node i in the graph, we count all the nodes that can be visited from node i in the dependency graph. In other words, for each gate (g_i) in the circuit (graph), we identify the number of all subsequent dependent gates (reachable gates) that are applied to the same qubit or other qubits, which are directly or indirectly connected to the current gate through two-qubit gates. For a quantum circuit with G gates, the weight of the ith gate (w_i) is computed as $w_i = \frac{S_i}{G}$. We exclude the R_z gate since its error rate is zero. Thus, the term λ is used to approximate the impact of noisy quantum gates on the output state fidelity of the quantum circuit, given the number of subsequent dependent (reachable) gates. To simplify our analysis, we omit how Pauli errors propagate through different quantum gates in the circuit. Instead, we assume that for each gate, Pauli errors will propagate to its corresponding qubits. While the proposed WESP metric does not accurately characterize

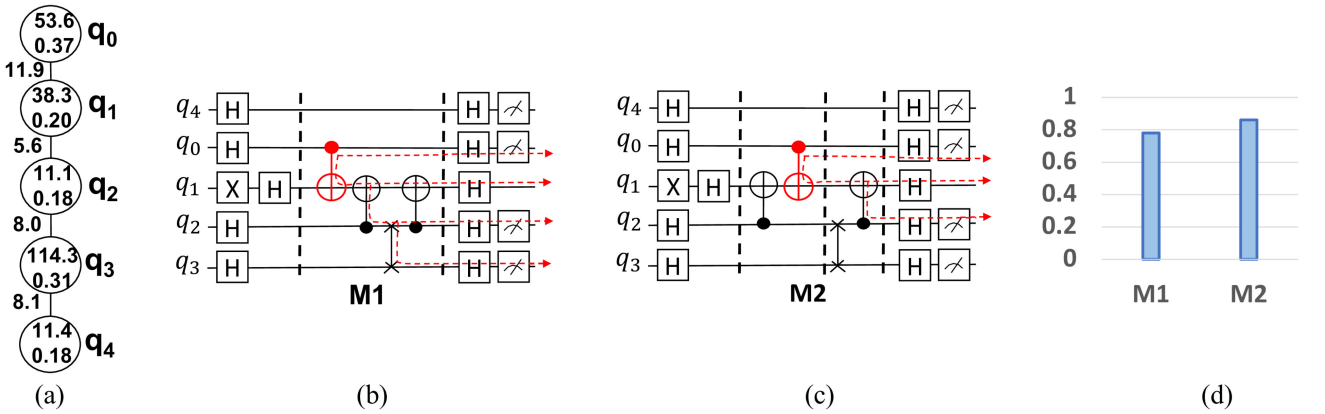

FIGURE 5. (a) IBM Q Santiago architecture, 4-qubit Bernstein–Vazirani (BV) implemented using (b) M1 and (c) M2 mappings and (d) their corresponding PST value.

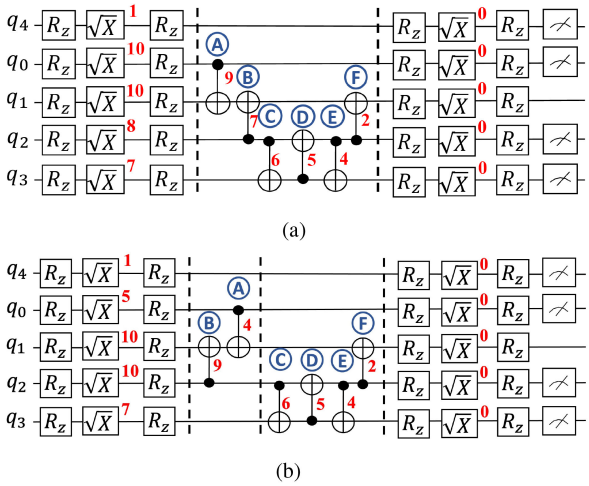

FIGURE 6. Number of reachable gates (S_i) for each quantum gate in (a) M1 and (b) M2 mappings.

the circuit error rate, it can be used to differentiate between different quantum circuit implementations, which share the same gate count and circuit depth, and maintain the same gate parallelism to avoid additional crosstalk or decoherence errors.

Example 4: Fig. 6(a) and (b) show the physical quantum circuits implemented using elementary gates supported by IBM quantum computers for M1 and M2 quantum circuits in Fig. 5, respectively. The S_i values for each single- and two-qubit gates are shown in red color. For example, in the physical circuit provided in Fig. 6(a), the $\sqrt{X}$ gate applied to qubit q_4 in layer 2 is connected to the next $\sqrt{X}$ gate applied to the same qubit at a later layer in the gate dependency graph. Hence, the value of S_i for the first $\sqrt{X}$ gate is 1. Similarly, the value of S_i for all the other gates except R_z gates are computed. The number of erroneous gates (G) is 16. Based on the device error rates, as shown in Fig. 5(a), the value of WESP for M1 mapping is 0.7811. On the other hand, WESP for M2 mapping is 0.7846, which indicates that M2 circuit provides a higher PST. This observation is aligned with the PST value of M1 and M2 circuits provided in Fig. 5(d).

B. ELEMENTARY GATE RESCHEDULING BASED ON WESP

We exploit the gate commutation rules to reduce the impact of Pauli errors while maintaining the depth of the circuit. We propose a gate rescheduling algorithm postquantum circuit mapping guided by the proposed WESP metric as a new optimization layer to reduce the circuit gate errors. Since the gates of the physical quantum circuit satisfy the coupling constraint of the quantum hardware, no additional SWAP operation is required. We develop a greedy approach for gate rescheduling based on the circuit dependency graph ($D(V, E)$) as described in Algorithm 1. The objective of our rescheduling algorithm is to maximize WESP while maintaining the circuit depth (Depth_C) to avoid any additional decoherence errors. An important condition is added to ensure that the updated gate scheduling does not change the depth of the circuit (No_Change_Depth). For g_i gate operating on q_x and q_y qubits at layer t and g_j gate operating on q_x and q_z qubits at layer $t + k$, g_i and g_j can be reordered if they are commuting, there is no non-commutable gate in between, and q_y and q_z are idle at layers $t + k$ and t, respectively. Our evaluation is repeated for every pair of immediate dependent gates in the gate dependency graph. For a quantum circuit with $D(V, E)$ gate dependency graph, in which V (nodes) represents the set of quantum circuit gates and E (edges) shows the dependency between different gates, the time complexity of calculating WESP is $O(|V| + |E|)$. Thus, the time complexity of our greedy heuristic, which is provided in Algorithm 1, is $O(|V|^2 \cdot (|V| + |E|))$. Our proposed gate rescheduling approach is computationally fast compared to the exhaustive approach that requires checking the entire search space, which can be very large for quantum circuits with a large number of commuting gates.

Example 5: To motivate the need and the effectiveness of our greedy approach, we exhaustively generate all possible ways of gate scheduling for our BV circuit provided in Figs. 5 and 6. We pick the circuit with the highest WESP

IEEE Transactions on Quantum Engineering

Algorithm 1: Gate Rescheduling.

Input: C = The physical quantum circuit
Output: $C^{'}$ = The updated physical quantum circuit
Construct the D graph of C;
Initialize $C^{'} = C$;
Initialize $L(i)$ = gates applied at layer $i \mid 0 < i < Depth_C$;
for *each layer* i **do**
  for *each gate* g_j *in* $L(i)$ **do**
    for *each immediate dependent gate* g_k **do**
      if *Are_commutative(D, g_j, g_k)* **then**
        if *No_Inc_Depth(D, g_j, g_k)* **then**
          Compute WESP of $C^{'}$ under rescheduled gates;
        end
      end
    end
    Select rescheduling of g_j with max. WESP;
    Update $C^{'}$;
  end
end
Return $C^{'}$;

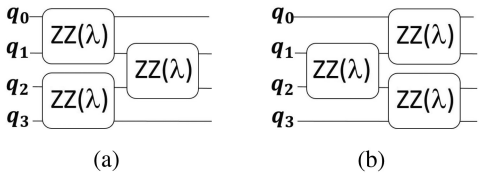

FIGURE 7. Subcircuits with different ZZ gate scheduling.

and compare it with the generated circuit using our greedy approach in terms of WESP and PST. Specifically, we label each CNOT gate in Fig. 6 to list all possible combinations of gate scheduling. The total number of all possible ways of gate scheduling, which maintain the depth of the circuit is 14. Among all the gate scheduling combinations, B–C–D–E–F–A has the highest WESP of 0.7768, while the resulting circuit based on our greedy algorithm has a CNOT combination of B–A–C–D–E–F with a WESP of 0.7764. The PST values of the quantum circuit based on the exhaustive and our greedy approach are 0.863 and 0.861, respectively. Accordingly, the gap between the exhaustive and the greedy solutions is very small.

C. COMPLEX GATE RESCHEDULING BASED ON WESP FOR QAOA

QAOA quantum circuits consist of single-qubit rotations and two-qubit phase gates implemented using ZZ complex gate. Since ZZ gates are commutative [34], [35], we exploit the freedom of the ZZ gate placement at higher quantum circuit design level to further reduce the quantum circuit errors. ZZ gates are reordered without violating the circuit depth and the gate count.

Example 6: An example of ZZ gate reordering is shown in Fig. 7, in which both subcircuits share the same gate count and circuit depth, but vary in the ZZ gate scheduling.

We propose a look-ahead approach that approximates the error rate of each complex gate in the circuit. The approximation entitles the generation of physical quantum circuit first, using an efficient quantum mapping approach that allocates physical qubits with minimum error rates and minimizes the number of SWAP operations. We scan the physical quantum circuit next to identify the corresponding two-qubit gates used for building each ZZ complex gate. Given the decomposed subcircuit of each complex gate, we compute the error rate of each complex gate.

We adjust the WESP metric to operate on the intermediate representation of the quantum circuit prior to complex gate decomposition. In the updated WESP, G is the number of ZZ gates, e_{g_i} is the ZZ gate error rate computed as the product of the error rates of its elementary gates, w_i is the ratio of dependent (reachable) ZZ gates, and $\min_{g_i \in G} e_{g_i}$ is the minimum error rate of the ZZ complex gate in the circuit. The ZZ gates are reordered based on the updated WESP using a similar rescheduling algorithm as the one provided in Algorithm 1, in which C is updated to be the intermediate representation of the quantum circuit, and every pair of quantum gates (g_i, g_j) corresponds to a pair of ZZ gates in the quantum circuit. The look-ahead approach requires scanning the entire physical circuit, including all its gates. The time complexity of the look-ahead approach is linear in the number of gates ($O(|V|)$). The time complexity of our ZZ gate rescheduling algorithm is $O(|V'|^2 \cdot (|V'| + |E'|))$ for the $D'(V', E')$ complex gate dependency graph, in which V' represents the set of ZZ complex gates and E' is the set of edges that show the dependency between different ZZ gates.

We emphasize that our rescheduling approaches, which are applied postmapping, can be integrated with other quantum compilation approaches that exploit the gate commutation rules to reduce the circuit depth and the gate count since our algorithms do not alter the circuit depth nor the gate count. Furthermore, our proposed approaches can be easily applied to quantum circuits executed on different quantum hardware in the presence of a significant variation in the hardware gate errors.

VI. EXPERIMENTAL EVALUATION

A. EXPERIMENTAL SETUP

We show the effectiveness of our proposed rescheduling algorithms based on WESP in improving the output state fidelity of different quantum circuits executed on NISQ computers. We conduct two experiments. In the first experiment, we use PST to evaluate our proposed elementary gate rescheduling approach applied to different quantum circuits, while in the second experiment, we show the impact of our proposed complex ZZ gate rescheduling on the AR of QAOA quantum circuits.

We use a total of 12 benchmark circuits with a different number of qubits and gate counts. The BV algorithm identifies a hidden string, which is encoded in the circuit [36].

TABLE I. Properties of Quantum Circuit Benchmarks

Benchmark	# Qubits	# U	# CNOT	# M	Depth	Expected Output
BV_3	4	9	2	3	6	110
Adder	3	30	17	2	28	0100
Grover	3	24	7	2	22	10
Toffoli	3	11	6	3	12	111
QFT_3	3	10	6	3	14	000
QFT_4	4	26	18	4	28	0000
Decoder	5	28	21	4	29	00100
QPE_3	3	19	7	2	17	01
QPE_4	4	24	14	3	31	001
QAOA_1	15	437	111	15	176	NA
QAOA_2	15	525	135	15	214	NA
QAOA_3	15	602	156	15	224	NA

TABLE II. Properties of Different IBM Quantum Computers

Quantum Computer	# Qubits	QV
IBM Q16 Melbourne	15	8
IBM Q Casablanca	7	32
IBM Q Santiago	5	32
IBM Q Rome	5	32

The discrete Fourier transform is represented using a quantum Fourier transform (QFT) circuit [5]. A shift operation is carried out for a given input Boolean function using the hidden shift (HS) algorithm [37]. The Grover search (Grover) algorithm provides a solution (x) of a function $f(x)$ equal to 1. Given a unitary operator, the eigenvalue of the eigenvector is estimated using the quantum phase estimation (QPE) algorithm [5]. QAOA solves optimization problems [9]. Other reversible quantum circuits, such as adder, Toffoli gate, and decoder, are also used. We obtained QFT and QPE circuits from Qiskit [38], HS and QAOA circuits from Cirq [39], adder from [40], decoder from RevLib [41], and BV, Toffoli, and Grover search by the manual construction. The properties of all the benchmark circuits are provided in Table 1, prior to the circuit mapping, which includes the number of qubits (# Qubits), single-qubit gates (U), CNOT gates (CNOT), measurement operations (M), the circuit depth (Depth), and the expected output. We use QAOA to solve the Max-Cut problem of three different graphs, which consist of 15 nodes and a different number of edges of random weight per node. The QAOA quantum circuits are QAOA_1, QAOA_2, and QAOA_3, which correspond to graphs with at most 5, 6, and 7 edges per node, respectively. To construct and test the QAOA circuits, we set p to 1 with default γ and β parameters provided by Cirq [39]. Unlike the other abovementioned quantum circuits, we consider the entire output distribution of QAOA circuits to evaluate their cost function.

We run our circuits on a variety of quantum architectures with different numbers of qubits and quantum volume (QV), as given in Table 2. The QV determines the largest size of random quantum circuits in terms of the number of qubits and the circuit depth, which can be successfully executed on the quantum computer with high fidelity [42]. Each execution of the quantum circuit consists of 8192 shots/trials.

All the benchmark circuits are mapped to the quantum architecture using the Qiskit software development kit based on the error rates of the quantum hardware [38]. We implement our rescheduling algorithms using Python programming language for easy integration with Qiskit. We run our algorithms on a 2.10 GHz Intel Xeon CPU E5-2620 processor with 62.8 GB memory.

We execute 150, 153, 121, and 117 pairs of base and reordered circuits on IBM Q Melbourne, IBM Q Santiago, IBM Q Rome, and IBM Q Casablanca, respectively, based on the availability of these devices. The results of all these circuits show that the reordered quantum circuits according to WESP provided in Section V yield a better solution than the base circuits generated by Qiskit. We select a subset of these circuits to show the effectiveness of our approach.

B. ELEMENTARY GATE RESCHEDULING RESULTS

In the first experiment, we study the impact of our proposed elementary gate rescheduling on the PST of different benchmark circuits. For each quantum circuit to be executed on a given quantum computer, we first generate the corresponding physical quantum circuit using the highest optimization level of Qiskit based on the calibration data of the quantum hardware. We refer to this approach as a *Base*. Next, we apply our proposed elementary gate rescheduling algorithm based on WESP to the physical quantum circuit to generate a new optimized physical quantum circuit. We refer to this process as *Prop. I*. We run the two physical quantum circuits on the quantum computer and evaluate their PST.

Fig. 8 provides a comparison of the PST of quantum circuits generated using Base and Prop. I and executed on different quantum computers. We also report the WESP of quantum circuits generated using Base and Prop. I, the number of gate reordering using Prop. I (R), and the depth of each physical quantum circuit (D) executed on different quantum computers in Table 3. We emphasize that the depth of the physical quantum circuit prior to and post applying Prop. I is always the same to avoid additional decoherence errors.

Fig. 8 shows that our proposed rescheduling algorithm improves the PST, and thus, the fidelity of the quantum circuit output. We observe significant improvement of PST for some quantum circuits more than others, such as adder, Toffoli, and Grover search quantum circuits for most of the quantum computers, which implies that the circuit structure can impact the effectiveness of our approach. We also observe that our proposed approach can also perform well for quantum computers with a larger QV, and thus, a better fidelity, such as IBM Q Santiago, IBM Q Casablanca, and IBM Q Rome. Furthermore, as our approach exploits the variation in gate error rates, which can result in different error rates propagating throughout different circuit paths, it is expected to perform better in the presence of higher gate error standard deviation of the quantum hardware. This is further illustrated in Fig. 8, for which the standard deviation of the two-qubit gate error of IBM Q16 Melbourne, IBM Q Casablanca, IBM Q Santiago, and IBM Q Rome on average are 1.31E-02, 1.86E-02, 1.33E-02, and 1.20E-02, respectively, which are considered relatively high.

IEEE Transactions on Quantum Engineering

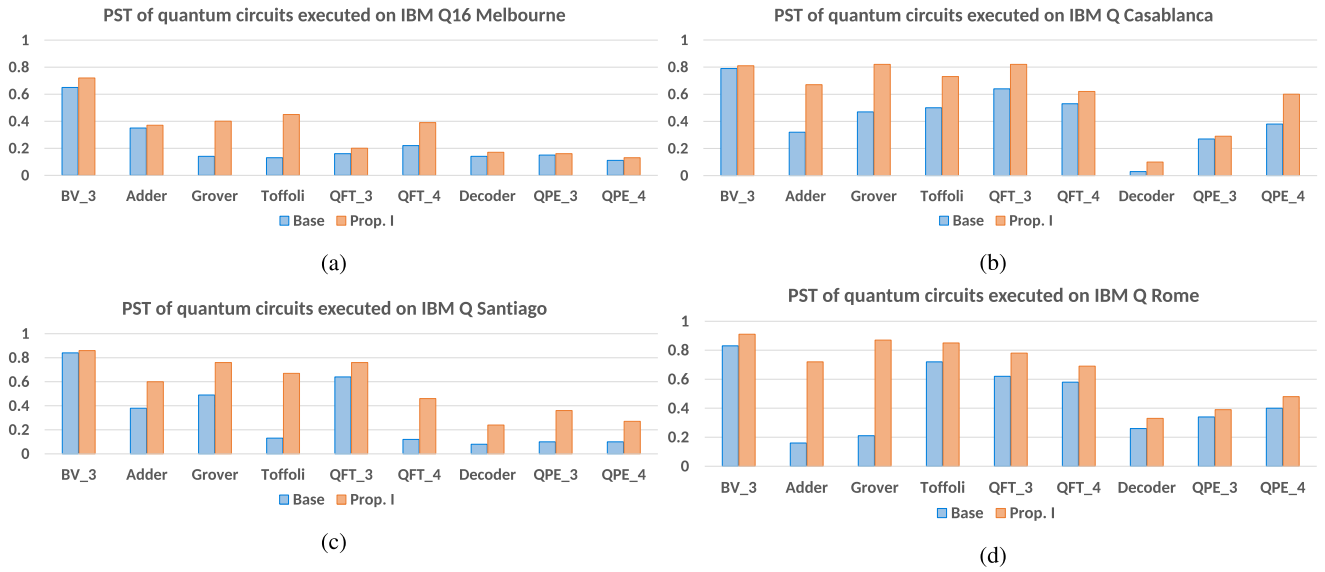

FIGURE 8. PSTs for different benchmarks executed on various IBMQ architectures. (a) PST of quantum circuits executed on IBM Q16 Melbourne. (b) PST of quantum circuits executed on IBM Q Casablanca. (c) PST of quantum circuits executed on IBM Q Santiago. (d) PST of quantum circuits executed on IBM Q Rome.

TABLE III. WESP of Quantum Circuits Generated Using Prop. I Gate Rescheduling With *R* Reordered Gates and *D* Circuit Depth

Benchmark	Melbourne				Casablanca				Santiago				Rome			
	WESP		R	D	WESP		R	D	WESP		R	D	WESP		R	D
	Base	Prop. I			Base	Prop. I			Base	Prop. I			Base	Prop. I		
BV_3	6.666E-1	6.672E-1	1	34	7.064E-1	7.071E-1	1	16	8.590E-1	8.591E-1	1	15	9.078E-1	9.079E-1	1	15
Adder	4.243E-1	4.945E-1	7	74	5.992E-1	6.442E-1	4	102	4.662E-1	5.339E-1	5	69	7.330E-1	7.743E-1	3	74
Grover	3.909E-1	3.940E-1	3	51	8.312E-1	8.432E-1	2	56	8.052E-1	8.218E-1	10	73	8.756E-1	8.831E-1	1	65
Toffoli	4.246E-1	4.292E-1	2	16	7.516E-1	7.870E-1	2	51	7.820E-1	7.998E-1	3	45	9.006E-1	9.107E-1	2	44
QFT_3	5.940E-1	6.278E-1	3	65	7.942E-1	8.127E-1	2	68	7.146E-1	7.284E-1	7	71	8.681E-1	8.833E-1	4	65
QFT_4	4.784E-1	4.789E-1	3	94	6.978E-1	7.235E-1	2	137	4.271E-1	4.572E-1	4	150	7.604E-1	7.800E-1	2	76
Decoder	2.886E-2	4.869E-2	5	141	2.648E-1	3.059E-1	3	200	2.117E-1	2.824E-1	6	215	5.079E-1	5.705E-1	6	229
QPE_3	3.430E-1	3.438E-1	1	77	7.664E-1	7.911E-1	2	71	7.422E-1	7.635E-1	9	92	8.374E-1	8.377E-1	1	92
QPE_4	1.528E-1	1.705E-1	3	127	6.691E-1	7.096E-1	3	134	3.467E-1	4.034E-1	5	117	6.842E-1	7.232E-1	8	102

As given in [33], the output state fidelity of the quantum circuit can be improved by rescheduling quantum gates as late as possible in the circuit to reduce the usage time of the qubit, and thus, reduce the qubit decoherence errors. To ensure that the improvement in PST shown in Fig. 8 is not due to a reduction in the qubit lifetime, and thus, decoherence errors, we identify the first layer in which each physical qubit is being used in each pair of quantum circuits generated using the Base and the Prop. I approach. For nonmeasured qubits, we also identify the last layer being used in both physical quantum circuits. For measured qubits, all measurement operations will be at the end of the circuit according to IBM quantum computers. We observe that none of the circuits used in Fig. 8 exhibits any reduction in their physical qubit lifetime after running the rescheduling algorithm, which confirms the effectiveness of our approach in reducing the impact of Pauli errors on the output state fidelity of the quantum circuit.

The run time of our proposed elementary gate rescheduling algorithm for each benchmark circuit to be executed on various quantum architectures is given in Table 4. Each quantum circuit requires only a small fraction of a second for gate rescheduling time.

TABLE IV. Run Time of the Prop. I Algorithm Applied to Different Quantum Circuits Mapped to Different IBM Q Computers

Benchmark	Prop. I Execution Time (msec)			
	Melbourne	Casablanca	Santiago	Rome
BV_3	0.54	0.17	0.17	0.23
Adder	1.25	2.17	1.37	1.52
Grover	1.14	1.27	1.52	1.34
Toffoli	0.20	0.53	0.85	0.90
QFT_3	1.12	1.38	1.49	1.26
QFT_4	1.47	2.59	2.92	1.57
Decoder	2.52	3.27	3.11	3.28
QPE_3	1.57	1.51	1.93	1.72
QPE_4	2.51	1.86	2.32	1.89

C. COMPLEX GATE RESCHEDULING RESULTS

In the second experiment, we show the effectiveness of our proposed ZZ complex gate rescheduling algorithm with respect to AR, and the corresponding ARG of different QAOA quantum circuits, namely QAOA_1, QAOA_2, and QAOA_3, executed on the IBM Q16 Melbourne quantum computer. We refer to the complex gate rescheduling algorithm as *Prop. II*. We also study the impact of applying Prop. II followed by the Prop. I approach to generate the physical QAOA quantum circuit, which implies the use of

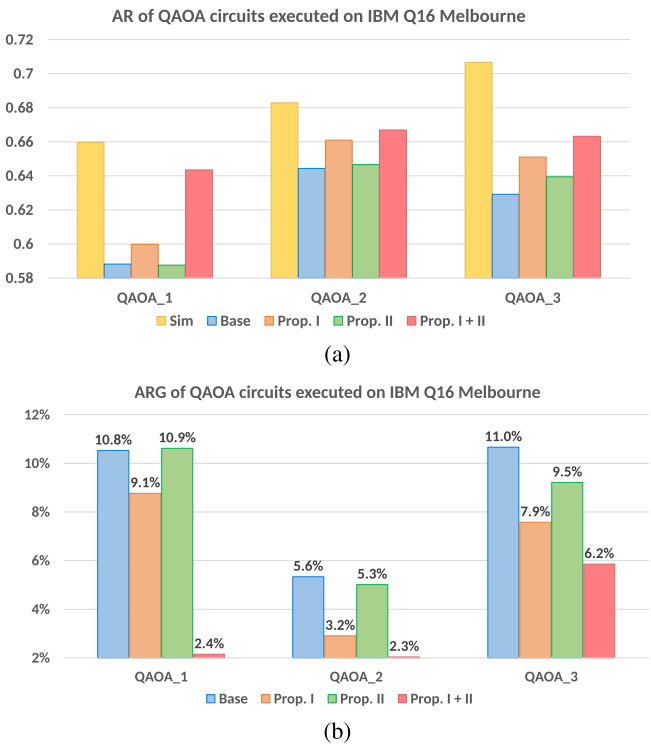

FIGURE 9. (a) AR of QAOA circuits. (b) Corresponding ARG of QAOA circuits.

TABLE V. WESP of QAOA Quantum Circuits Generated Using the Base, Prop.i, Prop. II, and Prop. I+II and Their Corresponding ESP Value

Benchmark	ESP	WESP (complex)		WESP (elementary)		
		Base	Prop. II	Base	Prop. I	Prop. I+II
QAOA_1	1.03E-04	5.01E-05	6.32E-05	1.69E-05	5.09E-05	7.19E-05
QAOA_2	6.81E-07	2.49E-07	4.58E-07	5.61E-08	2.31E-07	5.36E-07
QAOA_3	9.05E-10	3.78E-10	5.63E-10	1.66E-11	4.98E-10	7.92E-10

ZZ complex gate rescheduling algorithm first followed by the elementary gate rescheduling algorithm after complex gate decomposition based on the commutation rules in Fig. 1. This process is referred to as *Prop. I+II*.

Fig. 9(a) and (b) show a comparison of the AR and the corresponding ARG, respectively, of different QAOA circuits generated using Base, Prop. I, Prop. II, and Prop. I+II. The depth of physical QAOA_1, QAOA_2, and QAOA_3 quantum circuits are 315, 375, and 412, respectively, for all the proposed rescheduling approaches. The range of the number of reordering for QAOA_1, QAOA_2, and QAOA_3 quantum circuits are 3–35, 5–39, and 7–53, respectively.

The WESP computed at the complex and the elementary levels of the base and reordered QAOA quantum circuits as described in Section V and their corresponding ESP are provided in Table 5.

Our results show that while the ZZ gates rescheduling algorithm improves the AR of the QAOA circuits, and thus, reduces the ARG, applying elementary gate rescheduling provides a better output state fidelity. Our results also show that using both rescheduling algorithms at the complex and the elementary gate level delivers the best reduction in the ARG, and therefore, the best output state fidelity. Since both decoherence and gate errors contribute to the quantum circuit noise, by incorporating our gate rescheduling approaches that push very noisy gates to as later circuit layers as possible, we can improve the output state fidelity even for QAOA circuits with large depth. As we do not directly compare the output distribution of the simulated and the executed quantum circuits but their cost functions, we can achieve a good reduction in the ARG even with a small value of ESP and WESP.

TABLE VI. Run Time of Prop. I, Prop. II, and Prop. I+II Algorithms for QAOA Circuits Mapped to IBM Q16 Melbourne

Benchmark	Execution Time (msec)		
	Prop. I	Prop. II	Prop. I+II
QAOA_1	3.15	3.01	6.17
QAOA_2	3.22	3.11	6.35
QAOA_3	3.53	3.37	7.26

The run time of our proposed rescheduling approaches for QAOA_1, QAOA_2, and QAOA_3 quantum circuits to be executed on the IBM Q16 Melbourne quantum computer is given in Table 6. Our proposed approaches consume only a small fraction of a second despite the large circuit depth and the number of qubits.

VII. CONCLUSION

In this article, we proposed gate rescheduling algorithms based on the gate error propagation paths in the quantum circuits. Given the variation in the error rates of NISQ computers, we show that the location of the quantum gate can significantly affect the output state fidelity of the quantum circuit. Our proposed approaches can be easily integrated with other quantum compilation approaches, which ensure that the quantum circuit mapping process maximizes ESP. We also anticipated further improvements in the output state fidelity when applying our approaches with other compilation approaches that target decoherence and correlated errors. Our future work will further investigate the quantum circuit structures that benefit the most from our proposed rescheduling algorithms.

ACKNOWLEDGMENT

The authors would like to acknowledge the use of IBM Quantum services for this work. The views expressed are those of the authors and do not reflect the official policy or position of IBM or the IBM Quantum team.

REFERENCES

[1] E. Magesan, J. M. Gambetta, and J. Emerson, "Characterizing quantum gates via randomized benchmarking," *Phys. Rev. A*, vol. 85, 2012, Art. no. 042311, doi: 10.1103/PhysRevA.85.042311.

[2] 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.

[3] S. T. Flammia and J. J. Wallman, "Efficient estimation of Pauli channels," *ACM Trans. Quantum Comput.*, vol. 1, no. 1, Dec. 2020, Art. no. 3, doi: 10.1145/3408039.

[4] Y. Nam, N. J. Ross, Y. Su, A. M. Childs, and D. Maslov, "Automated optimization of large quantum circuits with continuous parameters," *npj Quantum Inf.*, vol. 4, no. 23, 2018, Art. no. 23, doi: 10.1038/s41534018-0072-4.

[5] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*. Cambridge, U.K.: Cambridge Univ. Press, 2019.

[6] D. Maslov, G. W. Dueck, and D. M. Miller, "Toffoli network synthesis with templates," *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 24, no. 6, pp. 807–817, Jun. 2005, doi: 10.1109/TCAD.2005.847911.

[7] K. Hietala, R. Rand, S.-H. Hung, X. Wu, and M. Hicks, "A verified optimizer for quantum circuits," *Proc. ACM Program. Lang.*, vol. 5, no. POPL, Jan. 2021, Art. no. 37, doi: 10.1145/3434318.

[8] S. Janardan, Y. Tomita, M. Gutierrez, and K. R. Brown, "Analytical error analysis of clifford gates by the fault-path tracer method," *Quantum Inf. Process.*, vol. 15, no. 8, pp. 3065–3079, 2016, doi: 10.1007/s11128016-1330-z.

[9] E. Farhi, J. Goldstone, and S. Gutmann, "A quantum approximate optimization algorithm," 2014, *arXiv:1411.4028*, doi: 10.48550/arXiv.1411.4028.

[10] S. S. Tannu and M. K. Qureshi, "Not all qubits are created equal: A case for variability-aware policies for NISQ-era quantum computers," in *Proc. 24th ACM Int. Conf. Architectural Support Program. Lang. Oper. Syst.*, 2019, pp. 987–999, doi: 10.1145/3297858.3304007.

[11] M. Alam, A. Ash-Saki, and S. Ghosh, "Circuit compilation methodologies for quantum approximate optimization algorithm," in *Proc. 53rd Annu. IEEE/ACM Int. Symp. Microarchit.*, 2020, pp. 215–228, doi: 10.1109/MICRO50266.2020.00029.

[12] L. Zhou, S.-T. Wang, S. Choi, H. Pichler, and M. D. Lukin, "Quantum approximate optimization algorithm: Performance, mechanism, and implementation on near-term devices," *Phys. Rev. X*, vol. 10, Jun. 2020, Art. no. 021067, doi: 10.1103/PhysRevX.10.021067.

[13] M. Y. Siraichi, V. F. d. Santos, S. Collange, and F. M. Q. Pereira, "Qubit allocation," in *Proc. 2018 Int. Symp. Code Gener. Optim.*, 2018, pp. 113–125, doi: 10.1145/3168822.

[14] J. Kusyk, S. M. Saeed, and M. U. Uyar, "Survey on quantum circuit compilation for noisy intermediate-scale quantum computers: Artificial intelligence to heuristics," *IEEE Trans. Quantum Eng.*, vol. 2, 2021, Art. no. 2501616, doi: 10.1109/TQE.2021.3068355.

[15] T. Itoko, R. Raymond, T. Imamichi, and A. Matsuo, "Optimization of quantum circuit mapping using gate transformation and commutation," *Integration*, vol. 70, pp. 43–50, Jan. 2020, doi: 10.1016/j.vlsi.2019.10.004.

[16] 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.

[17] R. Wille, L. Burgholzer, and A. Zulehner, "Mapping quantum circuits to IBM QX architectures using the minimal number of SWAP and H operations," in *56th ACM/EDAC/IEEE Des. Automat. Conf.*, 2019, Art. no. 142, doi: 10.1145/3316781.3317859.

[18] A. Zulehner, A. Paler, and R. Wille, "An efficient methodology for mapping quantum circuits to the IBM QX architectures," *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 38, no. 7, pp. 1226–1236, Jul. 2019, doi: 10.1109/TCAD.2018.2846658.

[19] 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. Oper. Syst.*, 2019, pp. 1001–1014, doi: 10.1145/3297858.3304023.

[20] A. Ash-Saki, M. Alam, and S. Ghosh, "QURE: Qubit re-allocation in noisy intermediate-scale quantum computers," in *Proc. ACM/IEEE Des. Automat. Conf.*, 2019, Art. no. 141, doi: 10.1145/3316781.3317888.

[21] M. Alam, A. Ash-Saki, and S. Ghosh, "Addressing temporal variations in qubit quality metrics for parameterized quantum circuits," in *Proc. IEEE/ACM Int. Symp. Low Power Electron. Des.*, 2019, pp. 1–6, doi: 10.1109/ISLPED.2019.8824907.

[22] 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. Microarchit.*, 2019, pp. 253–265, doi: 10.1145/3352460.3358257.

[23] N. Acharya and S. M. Saeed, "A lightweight approach to detect malicious/unexpected changes in the error rates of NISQ computers," in *Proc. 39th Int. Conf. Comput.-Aided Des.*, 2020, pp. 1–9, doi: 10.1145/3400302.3415684.

[24] S. M. Saeed, N. Mahendran, A. Zulehner, R. Wille, and R. Karri, "Identification of synthesis approaches for IP/IC piracy of reversible circuits," *ACM J. Emerg. Technol. Comput. Syst.*, vol. 15, no. 3, pp. 1–17, Apr. 2019, doi: 10.1145/3289392.

[25] N. Acharya and S. M. Saeed, "Automated flag qubit insertion for reliable quantum circuit output," in *Proc. IEEE Comput. Soc. Annu. Symp. VLSI*, 2021, pp. 431–436, doi: 10.1109/ISVLSI51109.2021.00085.

[26] P. Murali, D. C. Mckay, M. Martonosi, and A. Javadi-Abhari, "Software mitigation of crosstalk on noisy intermediate-scale quantum computers," in *Proc. Int. Conf. Architectural Support Program. Lang. Oper. Syst.*, 2020, pp. 1001–1016, doi: 10.1145/3373376.3378477.

[27] D. Cuomo *et al.*, "Optimized compiler for distributed quantum computing," 2021, *arXiv:2112.14139*, doi: 10.48550/arXiv.2112.14139.

[28] D. Ferrari, A. S. Cacciapuoti, M. Amoretti, and M. Caleffi, "Compiler design for distributed quantum computing," *IEEE Trans. Quantum Eng.*, vol. 2, 2021, Art. no. 4100720, doi: 10.1109/TQE.2021.3053921.

[29] G. G. Guerreschi and J. Park, "Two-step approach to scheduling quantum circuits," *Quantum Sci. Technol.*, vol. 3, no. 4, Jul. 2018, Art. no. 045003, doi: 10.1088/2058-9565/aacf0b.

[30] M. Alam, A. A. Saki, and S. Ghosh, "An efficient circuit compilation flow for quantum approximate optimization algorithm," in *Proc. 57th ACM/IEEE Des. Automat. Conf.*, 2020, pp. 1–6, doi: 10.1109/DAC18072.2020.9218558.

[31] B. Tan and J. Cong, "Optimal layout synthesis for quantum computing," in *Proc. 39th Int. Conf. Comput.-Aided Des.*, 2020, Art. no. 137, doi: 10.1145/3400302.3415620.

[32] Y. Shi *et al.*, "Optimized compilation of aggregated instructions for realistic quantum computers," in *Proc. 24th Int. Conf. Architectural Support Program. Lang. Oper. Syst.*, 2019, pp. 1031–1044, doi: 10.1145/3297858.3304018.

[33] 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. 2021, doi: 10.1109/LCA.2020.3045628.

[34] G. E. Crooks, "Performance of the quantum approximate optimization algorithm on the maximum cut problem," 2018, *arXiv:1811.08419*, doi: 10.48550/arXiv.1811.08419.

[35] D. Venturelli, M. Do, E. Rieffel, and J. Frank, "Compiling quantum circuits to realistic hardware architectures using temporal planners," *Quantum Sci. Technol.*, vol. 3, no. 2, Feb. 2018, Art. no. 025004, doi: 10.1088/2058-9565/aaa331.

[36] E. Bernstein and U. Vazirani, "Quantum complexity theory," *SIAM J. Comput.*, vol. 26, no. 5, pp. 1411–1473, Oct. 1997, doi: 10.1137/S0097539796300921.

[37] W. van Dam, S. Hallgren, and L. Ip, "Quantum algorithms for some hidden shift problems," in *Proc. 14th Annu. ACM-SIAM Symp. Discrete Algorithms*, 2003, pp. 489–498, doi: 10.1137/S009753970343141X.

[38] M. Treinish *et al.*, "Qiskit: An open-source framework for quantum computing," 2019, doi: 10.5281/zenodo.2573505.

[39] Cirq Developers, "quantumlib/Cirq: Cirq v0.9.1," Zenodo, Oct. 2020, doi: 10.5281/zenodo.4064322.

[40] S. A. Cuccaro, T. G. Draper, S. A. Kutin, and D. Petrie Moulton, "A new quantum ripple-carry addition circuit," Oct. 2004, *arXiv:quant-ph/0410184*, doi: 10.48550/arXiv.quant-ph/0410184.

[41] R. Wille, D. Grosse, L. Teuber, G. W. Dueck, and R. Drechsler, "Revlib: An online resource for reversible functions and reversible circuits," in *Proc. Int. Symp. Mult. Valued Log.*, 2008, pp. 220–225, doi: 10.1109/ISMVL.2008.43.

[42] A. W. Cross, L. S. Bishop, S. Sheldon, P. D. Nation, and J. M. Gambetta, "Validating quantum computers using randomized model circuits," *Phys. Rev. A*, vol. 100, no. 3, Sep. 2019, Art. no. 032328, doi: 10.1103/PhysRevA.100.032328.