

Optimizing quantum algorithms on bipotent architectures

YanJun Ji^{1,*}, Kathrin F. Koenig^{2,3,†} and Ilia Polian^{1,‡}

¹*Institute of Computer Architecture and Computer Engineering, University of Stuttgart, Pfaffenwaldring 47, 70569 Stuttgart, Germany*

²*Fraunhofer Institute for Applied Solid State Physics IAF, Tullastrasse 72, 79108 Freiburg, Germany*

³*Faculty of Engineering, University of Freiburg, Georges-Köhler-Allee 101, 79110 Freiburg, Germany*



(Received 23 March 2023; accepted 25 July 2023; published 17 August 2023)

Vigorous optimization of quantum gates has led to bipotent quantum architectures, where the optimized gates are available for some qubits but not for others. However, such gate-level improvements limit the application of user-side pulse-level optimizations, which have proven effective for quantum circuits with a high level of regularity, such as the ansatz circuit of the quantum approximate optimization algorithm (QAOA). In this paper, we investigate the trade-off between hardware-level and algorithm-level improvements on bipotent quantum architectures. Our results for various QAOA instances on two quantum computers offered by IBM indicate that the benefits of pulse-level optimizations currently outweigh the improvements due to vigorously optimized monolithic gates. Furthermore, our data indicate that the fidelity of circuit primitives is not always the best indicator for the overall algorithm performance; their gate type and schedule duration should also be taken into account. This effect is particularly pronounced for QAOA on dense portfolio optimization problems, since their transpilation requires many SWAP gates, for which efficient pulse-level optimization exists. Our findings provide practical guidance on optimal qubit selection on bipotent quantum architectures and suggest the need for improvements of those architectures, ultimately making pulse-level optimization available for all gate types.

DOI: [10.1103/PhysRevA.108.022610](https://doi.org/10.1103/PhysRevA.108.022610)

I. INTRODUCTION

Computational capabilities of today's quantum architectures are severely limited by comparatively high error rates of their hardware components. This shortcoming is commonly addressed on two levels: by providing improved implementations of qubits and quantum gates, and by adapting a given quantum circuit to error mechanisms and error rates of a specific quantum computer. A recent example from the first category is IBM's direct-CX gate that exhibits a shorter duration and a better fidelity than the previous entangling gate design based on the echo cross-resonance (ECR) principle [1]. Approaches from the second category include noise-aware transpilation methods [2] and pulse-level optimization of algorithm primitives [3,4].

Vigorous optimizations of quantum gates can lead to architectures where the optimized gate is not available for all existing qubits or qubit connections. For example, Fig. 1 shows the topology of the quantum processing unit (QPU) `ibmq_ehningen`. In this diagram, qubit pairs on which a direct-CX gate implementation is available share a blue (dark gray) edge, whereas qubit pairs on which only the ECR-CX gate implementation is available share an orange (medium gray) edge. While direct-CX gates tend to have better raw error rates, ECR-CX gates offer a significant advantage: they support bespoke pulse-level optimization, which is currently not

available for direct-CX gates unless additional calibration is performed. Pulse-level optimizations have proven to be highly effective for regular quantum circuits, such as the quantum approximate optimization algorithm (QAOA) [5], where frequently occurring subcircuits are mapped to efficient, bespoke sequences of pulses [3,6].

In this paper, we investigate the qubit mapping of QAOA onto quantum computers that provide two different implementations of the same gate on different pairs of qubits. We call such quantum architectures *bipotent*. Figure 1 is an example of a bipotent architecture, where two different CX gate types are available on different qubit pairs. Similarly, if there are more than two implementations of the same gate, we call such architectures *multipotent*. The reason for adopting a bipotent architecture may be the need to balance the quality of all qubits in engineering or manufacturing and the difficulty of having equally high-quality qubits due to various noises.

The qubit mapping must balance between hardware-level and algorithm-level improvements, that is, trade the lower error rate of direct-CX gates for pulse-level optimization supported by ECR-CX gates. We focus on a specific quantum algorithm: QAOA for portfolio optimization (PortOpt) and maximum-cut (MaxCut) instances with denser connectivity.

We start by algorithm-agnostically benchmarking direct-CX versus ECR-CX gates, generating data specific to the platforms being used. We then introduce pulse-level optimizations for QAOA circuits applied to PortOpt. In particular, mapping a dense PortOpt problem to a rather sparse topology map such as in Fig. 1 will require many SWAP gates, leading to a substantial contribution to the entire circuit's error

*yanjun.ji@informatik.uni-stuttgart.de

†kathrin.koenig@iaf.fraunhofer.de

‡ilia.polian@informatik.uni-stuttgart.de




FIG. 1. Bipotent architecture of `ibmq_ehningen`. It contains two types of CX gates: direct-CX, represented by blue or dark gray edges, and ECR-CX, represented by orange or medium gray edges. Qubits are depicted as circles, and the lines between the qubits are connections where a CX gate can be applied. When only one type of CX gate can be applied, the qubit has the same color as the CX gate type, otherwise it is green (light gray).

rate. Good pulse-level optimizations are known for such SWAP gates. Note that finding optimal pulse sequences is computationally expensive; for example, the compilation time of *gradient ascent pulse engineering* (GRAPE) scales exponentially in the size of the quantum algorithm [7]. For this reason, we only optimize the pulses of frequently occurring primitives of QAOA. The size of the primitives does not scale with the size of the algorithm. Moreover, we focus on improvements that rely on available calibration data and use pulse-level optimizations where possible, but do not consider solutions that require additional calibrations.

We report a comprehensive set of demonstrations on two of IBM’s bipotent architectures. We investigate not only the fidelities of QAOA circuits, but also their gate types and actual schedule durations on a bipotent architecture. Our results indicate that, for today’s error rates, pulse-level optimization leads to a stronger error-rate decrease than using improved but monolithic native gates in most cases. This comes with a few unexpected observations. For example, the relative impact of pulse-level optimization and monolithic gates varies for circuits of different sizes; pulse-level optimizations tend to outperform direct-CX gates for medium-scale circuits. Moreover, the control-target polarity of CX gates, which is usually ignored by implementing algorithms, has a significant influence on the quality of the results.

The key contributions of this paper are as follows.

(1) The features of a bipotent architecture are studied and used to optimize the implementation of algorithms on the QPUs.

(2) We improve the performance of QAOA by selectively optimizing partial gates implemented on some specific instead of all qubit pairs. This optimization technique can also be applied to other quantum algorithms.

(3) We describe how vigorously optimized quantum gates lead to bipotent quantum architectures and identify the conflict between using highly calibrated monolithic gates on the one hand and pulse-level optimization on the other.

(4) We extensively study the performance of QAOA on noisy quantum computers for different problem sizes and depths using different strategies for selecting direct-CX or/and ECR-CX gates.

(5) We provide practical advice on how to map QAOA instances onto a given bipotent quantum architecture, thus improving the performance of QAOA on said architectures.

(6) All methods presented in this paper do not require any additional calibration.

The remainder of this article is organized as follows. Section II reviews the implementation of QAOA on noisy intermediate-scale quantum (NISQ) devices and the challenges of noise. Section III provides the features of a bipotent architecture. Section IV discusses pulse-level optimizations in QAOA. Section V C demonstrates the performance of QAOA with different optimizations on two QPUs. In Sec. VI, we discuss the case beyond QAOA and error mitigation techniques. Finally, Sec. VII concludes this paper. Additional details about the QPUs used in this study are presented in the Appendix.

II. BACKGROUND

A. NISQ computing

Quantum computers form a new paradigm of computing that is directly based on the laws of quantum mechanics [8]. The basic unit of a quantum computer, the qubit, differs from the classical bit because it can be in a quantum *superposition* of the states 0 and 1, and can furthermore be *entangled* with other qubits. These quantum mechanical effects make quantum computers more powerful than classical computers in theory. There are various physical realizations of quantum computers already existing today, such as those based on superconducting circuits [9], photonics [10], quantum dots [11], trapped ions [12], and neutral atoms [13].

However, it proves to be experimentally extremely challenging to shield qubits from noise. Furthermore, technology has not advanced to the point that the errors induced by this noise can be fully detected and corrected during a quantum computation. Additionally, the currently available quantum computers have a limited number of qubits, and many platforms, including superconducting transmons considered in this work, offer restricted connectivity, i.e., for every qubit, two-qubit gates are only possible with a limited number (typically two to four) of other qubits. These limited devices are commonly referred to as NISQ devices [14].

Despite their limitations, NISQ devices may already solve problems that are intractable for any classical computers, and hence provide a so-called *quantum advantage*. Because of the noise, it is paramount to keep the quantum circuit depth of any NISQ algorithm to an absolute minimum, which becomes challenging because of the limited connectivity. Furthermore, the average number of errors incurred by the noisy two-qubit gates differs starkly from qubit pair to qubit pair. Hence, in the compilation of quantum algorithms, there should be a preference for executing gates on the pairs that induce few errors.

B. Noise-aware transpilation

In general, quantum circuits might include gates chosen by the designer and involving arbitrary qubits, whereas actual QPUs offer only a restricted set of quantum gates, called *basis gates*, and limited qubit connectivity. Therefore, a circuit must be *transpiled* to a given QPU: its operation must be mapped to

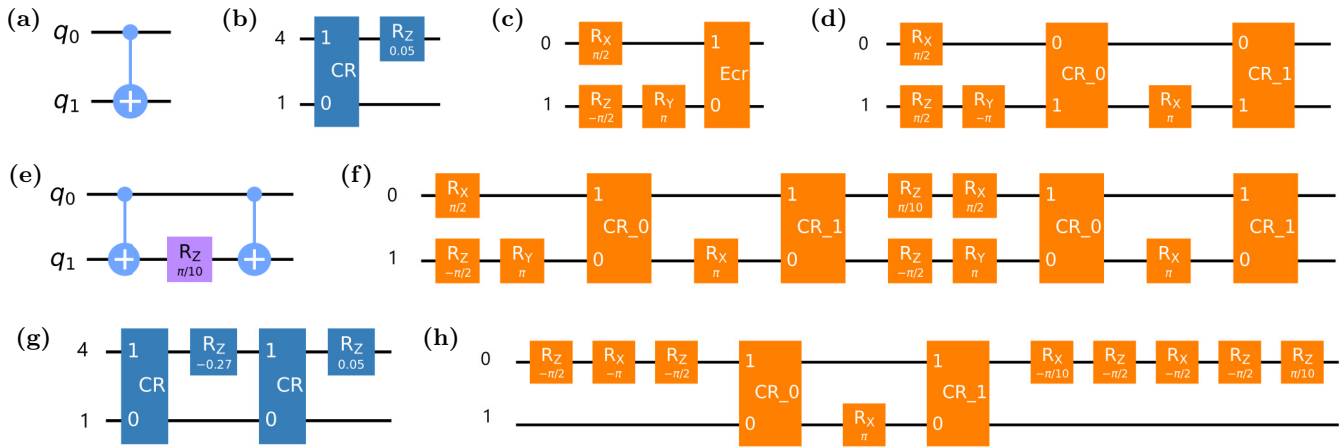


FIG. 2. Different implementations of CX and ZZ gates on IBM's QPUs. (a) Symbol of CX gate. (b) direct-CX gate on qubits [1, 4]. (c) ECR-CX gate on qubits [1, 0]. (d) Hardware implementation of the ECR-CX gate. (e) Decomposition of ZZ gate. ZZ gate implemented by (g) direct-CX gates on qubits [1, 4] and (f) ECR-CX gates on qubits [1, 0]. (h) ZZ_{OPT} gate, the pulse-level optimization of (f).

basis gates and connectivity violations must be addressed by, e.g., adding SWAP gates. Different qubits of a QPU are affected by noise to a different extent, and transpilation should be *noise aware*, i.e., evaluate current calibration data and select the qubits with the highest quality. It is worth noting that, for the considered QPUs, the added SWAP gates contribute to noise significantly, because each such gate includes three CX gates, which are a dominant error source. Therefore, transpilation must balance between selecting better qubits and requiring fewer additional SWAP gates.

Mapping the algorithm's qubits to quantum computer's qubits (the qubit mapping problem) has been shown to be NP-hard. Various transpilation approaches [15–22] were designed. Furthermore, many heuristic [2,23–28] and exact [2,23,29–33] methods were developed to solve the qubit mapping problem. While heuristics compile the circuit faster, the exact methods produce solutions that are optimal with respect to a user-specified objective (e.g., circuit depth or fidelity).

In this paper, we use an exact and scalable approach [33] on a linear topology aiming to minimize the circuit depth to pretranspile the circuit to satisfy the connectivity constraints. Then, we map the pretranspiled circuit to different qubits according to the circuit fidelity, gate type, and schedule duration.

C. Quantum gate design and pulse-level optimization

Quantum gates are a crucial component of quantum algorithms, and their quality is paramount to algorithm performance. Researchers have explored various methods for enhancing gate quality, including investigating different measurement durations to improve readout fidelity [34] and implementing pulse-level optimizations to enhance gate fidelity. For example, GRAPE can be used to translate quantum algorithms directly into optimal hardware control pulses with shorter pulse lengths than a gate-based compilation, thereby improving fidelity, but with exponentially increasing compilation times. A partial compilation of variational algorithms [35] was developed to achieve the pulse speedups of GRAPE by precomputing optimal pulses for parametrization-independent

blocks of gates. In addition, aggregating small gates into larger operations can reduce the compilation latency [36].

Compared to directly optimizing algorithms that require extensive computation, optimizing specific gates can improve the performance of the algorithm more effectively. In Ref. [3], the authors show a pulse-efficient transpilation approach that requires no additional calibration efforts by scaling cross-resonance (CR) pulses and removing redundant single-qubit operations to reduce the overall schedule duration of the gate and thus improves its fidelities. Faster and more reliable SWAP gates were developed at the pulse level with CR native gates on IBM's QPUs [37]. Moreover, using hardware primitive gates can reduce errors and run times, resulting in improved performance [4]. In addition, a hybrid gate-pulse model [38] was proposed to improve the performance of QAOA on IBM's QPUs. Recently, the authors in Ref. [39] demonstrated an improvement in performance by achieving shorter pulse schedules through the optimization of pulse amplitude and duration.

We now introduce the pulse-level structure of CX gates. Figure 2(a) is a general symbol of the CX gate on IBM's QPUs. The direct-CX and ECR-CX gates are indistinguishable at the gate level. However, there are differences in the pulse level. Figures 2(b) and 2(c) show the implementations (up to a global phase) of a direct-CX gate on qubits [1, 4] and an ECR-CX gate on qubits [1, 0], respectively. While the direct-CX consists of one CR gate and one RZ gate, the ECR-CX consists of one ECR gate and three single-qubit gates. A closer look into the hardware implementation shows that the ECR gate is implemented by two CR gates and one RX gate on the control qubit, as shown in Fig. 2(d). Their corresponding schedules consisting of several microwave pulses are shown in Fig. 3. The direct-CX with a schedule duration of 245.3 ns is implemented by one Gaussian square pulse on the connecting channel U3, a parallel echoed Gaussian square pulse on the drive channel D4 of target qubit 4, and a virtual (digital) RZ rotation (denoted by a \odot). In comparison, the ECR-CX gate with a duration of 320 ns consists of two CR pulses and four single-qubit gates. The schedule duration of the virtual RZ gate is 0 ns, while that of RY and RX gates is 32 ns.

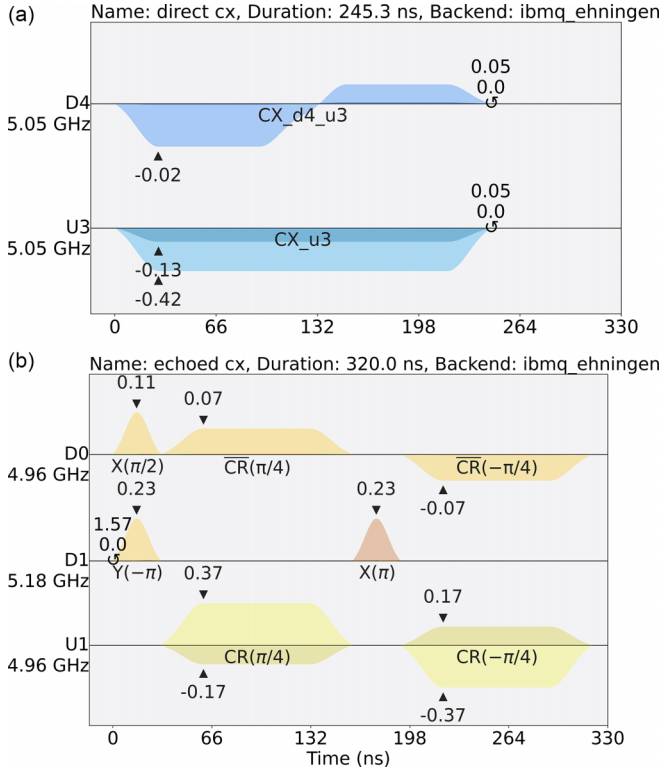


FIG. 3. Schedules of (a) direct-CX on qubits [1, 4] and (b) ECR-CX on qubits [1, 0].

Although direct-CX has a shorter schedule duration than ECR-CX, gates implemented by ECR-CX gates may perform better than those implemented by direct-CX, since they can be further optimized directly at the pulse level by employing existing techniques in QISKIT [40]. Specifically, the ZZ gate implemented by ECR-CX gates allows pulse-level optimization without any additional calibration [3]. Compared to this, optimizing gates implemented by direct-CX gates currently requires additional calibration (e.g., Ref. [41]). Figures 2(e)–2(h) show the different implementations of ZZ gate on IBM’s QPUs. It achieves a schedule duration of 490 ns on qubits [1, 4] using direct-CX gates, compared to a value of 640 ns on qubits [1, 0] using ECR-CX gates. However, the ZZ gate implemented by ECR-CX gates can be optimized at the pulse level, resulting in a reduced duration of 241.8 ns. We refer to the ZZ gate optimized at the pulse level as ZZ_{OPT}. In addition, single-pulse-gates RX and RY in ECR-CX enable gate cancellations, while a virtual RZ gate with 0-ns schedule duration in direct-CX cannot be further optimized.

As shown in Ref. [1], the direct-CX gate seems to have better quality. At the same time, the ECR-CX gate gives the user flexibility for pulse-level optimization [3], which will be discussed in Sec. III.

D. QAOA

Variational quantum algorithms (VQAs) hold promise for quantum advantage and have been developed in many applications [42]. The QAOA is among the most important VQAs and is promising for solving combinatorial optimization

problems on NISQ devices [5]. In this section, we describe the implementation of QAOA for PortOpt and MaxCut problems.

1. QAOA for PortOpt

The PortOpt aims to select the best portfolio from all portfolios in order to maximize expected returns and minimize financial risk. The QAOA holds promise for solving this problem [43–46].

The PortOpt problem needs to be first transformed into a *quadratic unconstrained binary optimization* (QUBO) [47], an NP-hard problem. The PortOpt with the QUBO representation is then converted to quantum operators so that it can be executed by NISQ computers.

Given n available assets and B assets to be selected, the cost function C can be represented in terms of quantum operators [44]:

$$\hat{C} = \sum_{i=1}^{n-1} \sum_{j=i+1}^n \frac{\lambda}{2} (q\sigma_{ij} + A)\hat{Z}_i\hat{Z}_j - \sum_{i=1}^n k_i\hat{Z}_i, \quad (1)$$

with

$$k_i = \frac{\lambda}{2} \left[A(2B - n) + (1 - q)\mu_i - q \sum_{j=1}^n \sigma_{ij} \right],$$

where λ , q , σ_{ij} , A , and μ_i are the global scaling factor, risk preference, covariance matrix, penalty factor, and expected return, respectively. $\hat{Z}_i\hat{Z}_j$ and \hat{Z}_i represent the ZZ interaction on qubits $[i, j]$ and Pauli Z operator acting on qubit i , respectively.

The implementation of QAOA starts from an initial state generated by Hadamard gates. After that, ZZ gates are performed with the definition

$$ZZ(\gamma) = e^{-i\frac{\gamma}{2}ZZ}, \quad (2)$$

where γ is the rotation angle, followed by single-qubit rotation gates RZ and RX. To optimize the parameters in QAOA, we utilize the constrained optimization by linear approximation (COBYLA) optimizer and QISKIT’s QASM simulator.

To evaluate the performance of QAOA for PortOpt, we define the approximation ratio (AR) and success probability (SP). A higher value of AR or SP implies better performance. The postselected results that satisfy the budget constraint (all feasible solutions) are used to calculate the AR of n assets $\{z_1, \dots, z_n\}$, defined as follows:

$$r(z_1, \dots, z_n) = C(z_1, \dots, z_n)/C_{\text{opt}}, \quad (3)$$

where $C(z_1, \dots, z_n)$ and C_{opt} are the mean value found by QAOA and optimal value, respectively. The SP is defined as the probability of optimal solution.

2. QAOA for MaxCut

Another application of QAOA is to find the maximum cut of a graph, i.e., to partition the nodes of the graph into two sets such that the number of edges between the sets is maximized, which is known as the MaxCut problem [5,48,49].




FIG. 4. Properties of ECR-CX (orange, medium gray) and direct-CX (blue, dark gray) gates on `ibmq_ehningen`. (a) Gate error rate. (b) Execution time in nanoseconds.

The objective function C can be reformulated as quantum operators:

$$\hat{C} = \frac{1}{2} \sum_{i,j} (1 - \hat{Z}_i \hat{Z}_j), \quad (4)$$

where vertices i, j share an edge. For a given graph with n edges $\{z_1, \dots, z_n\}$, the AR of QAOA for MaxCut is defined by

$$r(z_1, \dots, z_n) = C(z_1, \dots, z_n) / C_{\max}, \quad (5)$$

with $C(z_1, \dots, z_n)$ the mean value found by QAOA and $C_{\max} = \max C(z_1, \dots, z_n)$ the optimal value.

III. DIRECT-CX VS ECR-CX GATES

We now investigate the features of a bipotent architecture, namely, `ibmq_ehningen`, as shown in Fig. 1, on which two types of CX gates are present: direct-CX and ECR-CX. Additionally, we denote the qubits as Q-ECR (orange, medium gray), Q-direct (blue, dark gray), and Q-bipotent (green, light gray), depending on whether the qubit is connected to only ECR-CX, only direct-CX, or both.

We collected hourly calibration data of `ibmq_ehningen` throughout September 2022. Figure 4 shows the CX gate error rate and gate execution time of ECR-CX and direct-CX. While the error rate varies over time, the gate time is constant. Their average values are summarized in Table I. Compared to ECR-CX, the average error rate of direct-CX is slightly reduced by 4.82%, while the average gate time is significantly reduced by 32.79%. The data show that direct-CX is generally of better quality than ECR-CX.

TABLE I. Properties of ECR-CX and direct-CX gates.

Parameter	ECR-CX	direct-CX	Reduction (%)
Gate error (%)	0.83	0.79	4.82
Gate time (ns)	382.22	256.89	32.79

Figures 5(a) and 5(b) show the decoherence times T_1 and T_2 , respectively, of the 27 qubits on `ibmq_ehningen`. Qubits connected to two types of CX gates (Q-bipotent) have a higher average value of T_1 and T_2 compared to qubits connected only to ECR-CX (Q-ECR) or only to direct-CX gates (Q-direct). While the precise reason for this trend is unknown, it appears logical that more robust qubits are considered suitable for implementing both types of gates and are designated Q-bipotent. This implies the advantage of using a combination of both types of CX gates.

Single-qubit gate error and readout error are shown in Figs. 5(c) and 5(d), respectively. The mean gate errors are comparable for three qubit types. However, the Q-ECR has a higher readout error overall than the Q-direct and the Q-bipotent. Throughout this paper, we use the default readout error mitigation provided by QISKIT [40] to reduce the impact of measurements on results. The averages of decoherence times and error rates of qubits are summarized in Table II.

The study on bipotent architecture shows that direct-CX gates have significantly reduced schedule duration. In addition, Q-bipotent has the highest decoherence times T_1 and T_2 overall, while on average Q-direct has a better T_1 than Q-ECR that has a better T_2 .

IV. PULSE-LEVEL OPTIMIZATIONS

A. Algorithm-native gate set

Studying the properties of an algorithm is crucial to improve the performance of its implementation. One of the most

TABLE II. Properties of qubits.

Parameter	Q-ECR	Q-direct	Q-bipotent
T_1 (μ s)	122.59	154.29	161.16
T_2 (μ s)	132.94	118.02	175.16
Gate error (%)	0.021	0.023	0.026
Readout error (%)	1.386	1.009	1.077

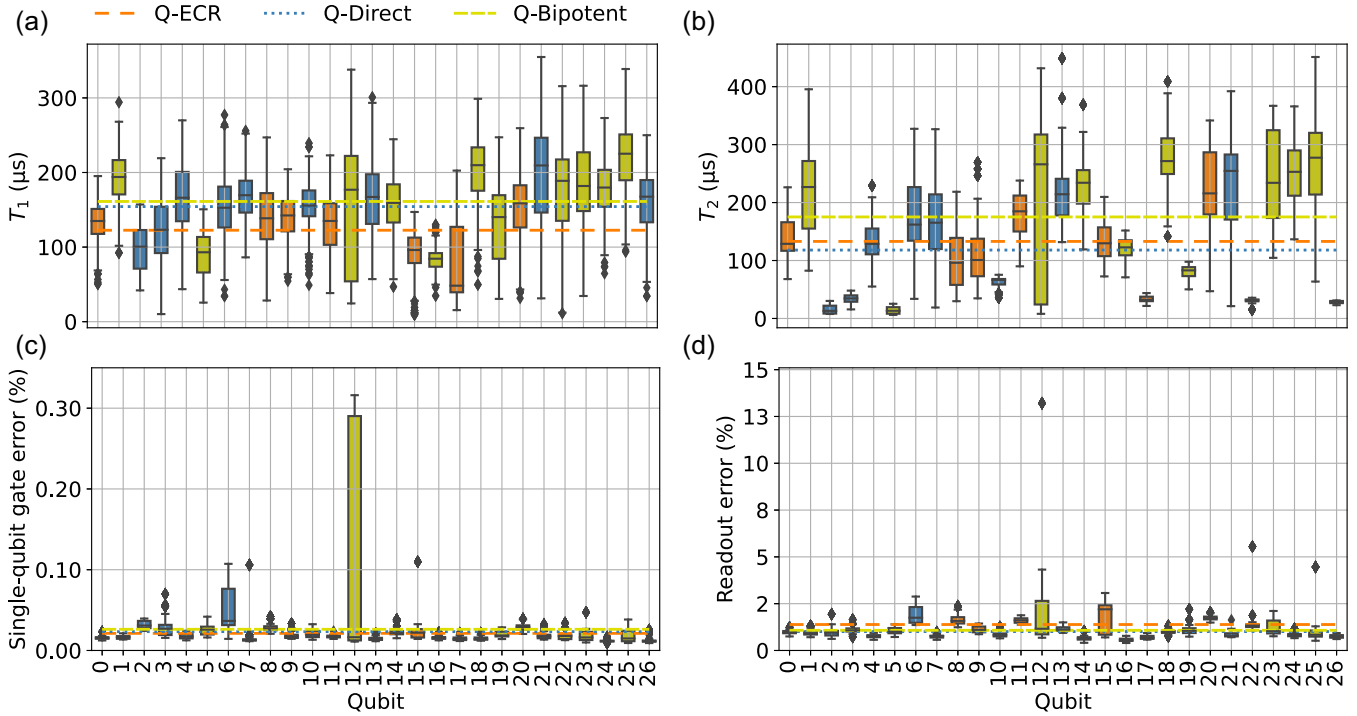


FIG. 5. Properties of 27 qubits on `ibmq_ehningen` used for experiments in this article. (a) T_1 and (b) T_2 decoherence times. A higher value is better. (c) Single-qubit gate error and (d) readout error. A lower value is better. We label qubits connected to two types of CX gates as Q-bipotent (green, light gray) and qubits connected only to ECR-CX or direct-CX gates as Q-ECR (orange, medium gray) or Q-direct (blue, dark gray), respectively.

important properties is the algorithm-native gate set (ANGS), which is the set of gates used to construct the algorithm. Efficient implementations of these gates will facilitate transpilation, enable exploitation of pulse-level optimization, and ultimately result in better performance.

Figure 6 shows the pretranspiled circuit of QAOA for PortOpt with five qubits and $p = 1$ on a linear topology using the strategy in Ref. [33] that provides optimal and scalable solutions. The resulting circuit has a structure similar to a swap network (see e.g., Ref. [18]), but two layers are skipped per p , namely, SWAP gates after the first and last layers of ZZ gates in every p . Thus, in the gate set $\{H, RX, RZ, ZZ, ZZ\text{-SWAP}\}$, the circuit with n qubits has a depth of $2 + (n + 2)p$ including initial state and measurement operators.

With the overall structure of QAOA circuit in Fig. 6, the ANGS of QAOA for PortOpt is $\{H, ZZ, ZZ\text{-SWAP}, RZ, RX\}$,

while that of QAOA for MaxCut is $\{H, ZZ, ZZ\text{-SWAP}, RX\}$, as no RZ is required for Eq. (4). Therefore, we focus on improvements of these gate types based on different CX primitives.

B. Analyzing two-qubit gates

We now analyze two-qubit gates. The CZ, ZZ, and ZZ-SWAP gates are undirected, which means that a more efficient implementation is possible with the native CX gate on the QPU. Additionally, gate cancellation at the pulse level provides opportunities for optimizations of gates implemented by ECR-CX gates.

Figure 7(b) shows the default implementation of a CZ gate using direct-CX on qubits [1, 4] with a schedule duration of 309.3 ns. Although the implementation using ECR-CX on qubits [1, 0] shown in Fig. 7(c) has a longer duration of 384 ns, it can be optimized at the pulse level (CZ_{OPT}) by reducing three




FIG. 6. Pretranspiled circuit of QAOA for PortOpt with five qubits and $p = 1$ on a linear topology [33]. For higher p , the circuit between the initial state and measurement is repeated p times, each time with new parameters.

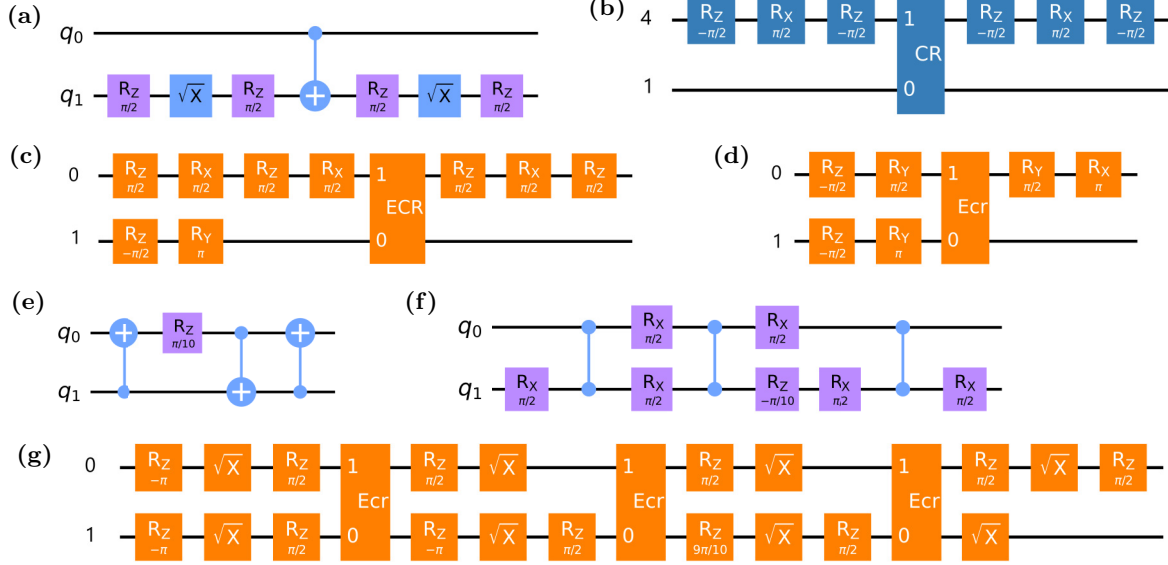


FIG. 7. Different implementations (up to a global phase) of CZ and ZZ-SWAP gates on IBM's QPUs. (a) Decomposition of CZ gate. CZ gate implemented by (b) direct-CX gate on qubits [1, 4] and (c) ECR-CX gate on qubits [1, 0]. (d) CZ_{OPT} gate, the pulse-level optimization of (c). (e) Decomposition of ZZ-SWAP gate. (f) Alternative decomposition of the ZZ-SWAP gate into CZ gates. (g) $ZZ-SWAP_{OPT}$ gate, the implementation of (f) on qubits [1, 0] using ECR-CX gates with pulse-level optimization

single-qubit gates, resulting in a duration of 352 ns, as shown in Fig. 7(d).

The decompositions of the ZZ-SWAP gate into CX gates and CZ gates [50] are shown in Figs. 7(e) and 7(f), respectively. In contrast to the CZ gate, the CX gate is directed. The ZZ-SWAP gate has a schedule duration of 992 ns on qubits [1, 0] and 800 ns on qubits [1, 4]. To implement the ZZ-SWAP gate with CZ gates, we use CZ_{OPT} implemented by ECR-CX since the CZ gate is not a basis gate of IBM's QPUs, and then take advantage of single-qubit cancellation at the pulse level, resulting in the same duration as the standard ZZ-SWAP gate. The optimized ZZ-SWAP gate ($ZZ-SWAP_{OPT}$) is shown in Fig. 7(g).

The data show that implementing gates using vigorously optimized monolithic gates results in a shorter duration, while using ECR-CX gates allows for user-friendly pulse-level optimization without requiring additional calibration.

C. Experimental results

We have discussed the pulse-level optimization of gates implemented by ECR-CX gates. We now use the *process fidelity* to benchmark the performance of default and optimized gates on the qubits that support ERC-CX gates. The *process fidelity* can be determined by quantum process tomography (QPT) using QISKIT's standard implementation. QPT is a method to characterize the actual behavior of quantum gates on QPUs. For QPT experiments, several initial states need to be prepared and then evolved, followed by numerous measurements at different measurement bases. The quantum channel is then reconstructed from the measurement data. The process fidelity between two quantum channels η and ξ is given by

$$F(\eta, \xi) = \text{Tr}[\sqrt{\sqrt{\rho_\eta} \rho_\xi \sqrt{\rho_\eta}}]^2, \quad (6)$$

where ρ_η and ρ_ξ are the normalized *Choi* matrices for the channel η and ξ , respectively. In our case, ρ_η is the *Choi* matrix of the ideal operator, while ρ_ξ is the *Choi* matrix determined by the QPT process on the QPU.

The ZZ and ZZ-SWAP denote gates implemented by the default basis gates of IBM's QPUs, while the ZZ_{OPT} and $ZZ-SWAP_{OPT}$ refer to gates with pulse-level optimizations. In addition, we employ the notations CT and TC to distinguish between the gate implemented by CX gates with control target (CT, in hardware native CX direction) and target control (TC, opposite to hardware native CX direction).

Figure 8(a) shows the infidelity ($1 - F$) of the ZZ gate on qubits [5, 8] (left) with a gate repetition of 1 and on qubits [1, 0] (right) with a gate repetition of 10, respectively, as a function of angle. The schedule duration of ZZ_{OPT} -CT and ZZ_{OPT} -TC increases with angle, while that of ZZ-CT and ZZ-TC is constant since the CX gate has a fixed duration and the RZ gate is virtual. ZZ_{OPT} -CT has the best overall performance as its duration is the shortest.

Figure 8(b) shows the infidelity of the ZZ-SWAP gate on qubits [0, 1] with gate repetitions of 10 (left) and 20 (right), respectively. We benchmark the performance of the ZZ-SWAP gate implemented by default CX gates considering CT and TC and by CZ_{OPT} gates considering CT. The schedule duration of ZZ-SWAP-ST and ZZ-SWAP-TC are the same, while that of ZZ-SWAP-CT is longer. The $ZZ-SWAP_{OPT}$ gate yields the best performance by decomposing the ZZ-SWAP into CZ gates, performing pulse-level optimization, and taking into account the control-target polarity of CX gates.

The results confirm that using hardware native polarity control-target to implement undirected gates, such as ZZ and ZZ-SWAP, is more efficient than using target-control polarity, which results in a longer duration. In addition, gates optimized at the pulse level perform better than their default

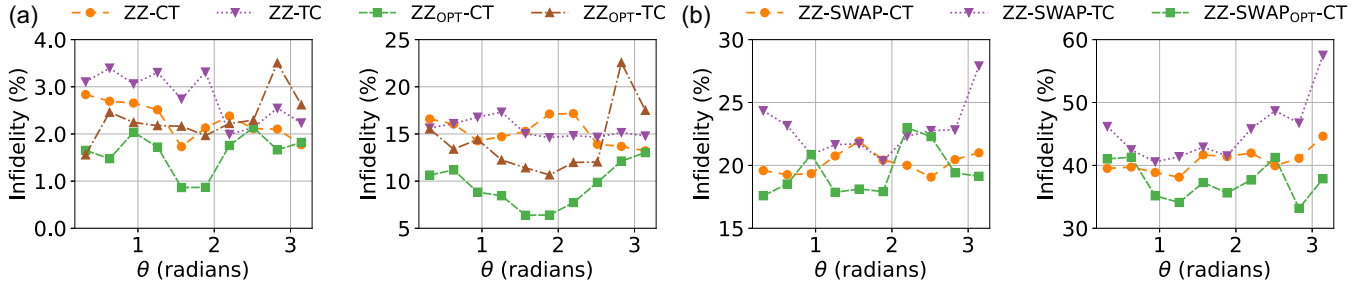


FIG. 8. Infidelity determined by QPT on the `ibmq_ehningen`. For (a), the ZZ gate acts on qubits [5, 8] with a gate repetition of 1 (left) and [1, 0] with a gate repetition of 10 (right). For (b), the ZZ-SWAP gate acts on qubits [1, 0] with a repeated number of gates of 10 (left) and 20 (right).

implementation. In the benchmarking below, we implement the undirected gates with control target.

V. BENCHMARKING WITH QAOA

A. Benchmarks and performance metrics

We consider as benchmarks the circuits with different numbers of qubits and depths p . The maximum number of qubits that can be connected linearly using only ECR-CX gates is five, e.g., [9, 8, 11, 14, 16], as illustrated in Fig. 1. Therefore, we limit the number of qubits to five to allow a full comparison of qubits with different CX gate types.

The five-asset PortOpt instance that runs on five qubits (5Q) is taken from Ref. [44], and its first three and four assets are used for the three- (3Q) and four-qubit (4Q) instances, respectively. The values of the variables q , B , A , and λ vary across the three cases considered. For 3Q, they are 0.33, 2, 0, and 20.97, respectively. For 4Q, they are 0.33, 2, 0.13, and 17.99, respectively. Finally, for 5Q, they are 0.33, 3, 0.07, and 17.51, respectively. For MaxCut we consider the *complete graph*, which results in the same connectivity requirements as the dense PortOpt problem instances.

The AR and SP defined in Sec. IID are used to evaluate the performance. We simulate the circuits in the noiseless case with the QASM simulator to get the values of AR and SP of QAOA for PortOpt and MaxCut that can be used as a baseline. For the demonstrations on QPUs, we set the number of shots to 50 000 for each circuit. Table III presents the simulated values of the AR and SP of QAOA for PortOpt with $p \in \{1, 2, \dots, 5\}$, while Table IV displays the results of QAOA for MaxCut. The data show that QAOA with five qubits for PortOpt exhibits smaller AR and SP values compared to MaxCut for the specific instances considered.

TABLE III. AR and SP of QAOA for PortOpt without noise.

p	1	2	3	4	5
3Q-AR	0.993	0.996	0.998	1.000	1.000
3Q-SP (%)	99.34	99.57	99.83	99.99	99.98
4Q-AR	0.525	0.674	0.892	0.930	0.935
4Q-SP (%)	21.18	29.04	44.30	45.61	46.97
5Q-AR	0.513	0.819	0.828	0.898	0.909
5Q-SP (%)	17.70	48.61	45.28	55.33	58.35

In addition to AR and SP, we use the schedule duration and the number of CX gates to study the properties of the circuit. A lower schedule duration means that the circuit executes faster, while a reduced number of CX gates implies pulse-level optimizations.

B. Methodology

To implement QAOA, we first prepare pretranspiled circuits as described in Sec. IVA. This process ensures that the circuits satisfy the connectivity constraints on a linear topology. The next crucial step is to identify high-quality qubits to execute them. To evaluate the qubit quality, we first consider only the circuit fidelity [15]. As illustrated in Fig. 9, we consider four families of QAOA circuits for evaluation. ECR, direct, and global circuits connected with violet (dark gray) lines use the qubits selected by MAPOMATIC [51] aiming to maximize the fidelity of the circuit. In comparison, the bipotent circuits connected with green (light gray) lines take into account the schedule duration, fidelity, and gate type. Moreover, ECR circuits utilize only the ECR-CX gates (with or without pulse-level optimizations). Direct circuits are composed of only direct-CX gates, while global and bipotent circuits contain mixtures of both gate types.

To construct an ECR circuits with k qubits, MAPOMATIC [51] selects a maximum-fidelity linear arrangement of k qubits such that ECR-CX gates are available on neighboring qubits. We use three versions of such circuits: ECR-default that employs default implementations of ZZ and ZZ-SWAP gates; ECR-ZZ_{OPT} that utilizes pulse-level optimized ZZ gate performed with the reduced CR gate; and ECR-ZZ-SWAP_{OPT} that leverages pulse-level optimizations for ZZ gates and reconstructs ZZ-SWAP gates with CZ_{OPT} gates. We maintain the same set of qubits for each data point if optimizations are implemented.

Direct circuits exclusively utilize direct-CX gates, and best-fidelity linear arrangements of qubits that support such gates are selected, again using MAPOMATIC. It is worth noting

TABLE IV. AR and SP of QAOA for MaxCut without noise.

p	1	2	3	4	5
5Q-AR	0.914	0.972	0.996	0.997	0.997
5Q-SP (%)	87.79	96.99	99.45	99.62	99.72

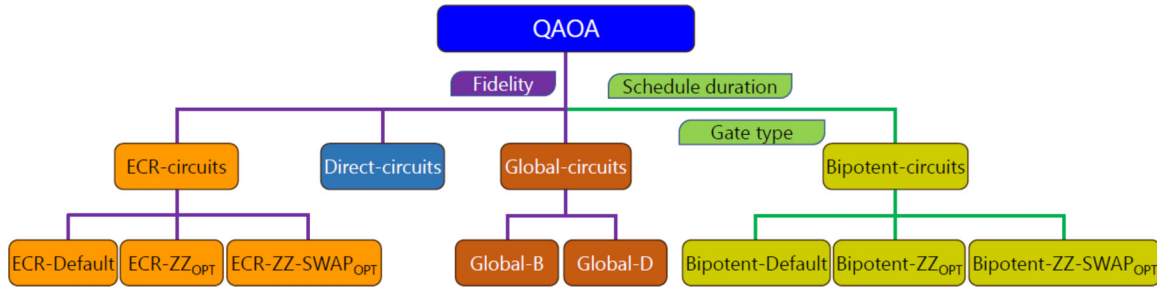


FIG. 9. QAOA circuits. ECR, direct, and global circuits consider fidelity and restrict the CX gate type in the circuit to only ECR-CX, only direct-CX, and regardless of its type, respectively, whereas Bipotent-circuits have a trade-off between fidelity, schedule duration, and gate type.

that there are no pulse-level optimizations in direct circuits. Global circuits are best-fidelity linear arrangements of arbitrary qubits, regardless of whether their connections support ECR-CX or direct-CX. It is possible, but not necessary, that all selected qubits of a global circuit have only direct-CX links. We call such configurations Global-D (for “direct”), and configurations with mixtures of direct-CX and ECR-CX Global-B (for “bipotent”). We observe that due to the higher error rate of ECR-CX gates, as shown in Fig. 4, only ECR-CX-linked qubits are not selected in global circuits.

In addition to those circuits where qubits are selected based on the circuit fidelity, we consider a further bipotent arrangement, called bipotent circuits. As the schedule duration of a circuit on the selected qubits has a fixed value and can be determined directly by QISKIT, bipotent circuit optimizes duration while controlling its fidelity. We focus only on qubits with below-average single-qubit and two-qubit error rates. Among these qubits, we select a linear arrangement on which the pulse-level optimized circuit has the shortest schedule duration. Moreover, we enforce bipotency by requiring at least one ECR-CX and at least one direct-CX gate. Pulse-level optimizations can be applied to ECR-CX gates within such circuits; the resulting circuits are labeled bipotent default, bipotent ZZ_{OPT}, and bipotent ZZ-SWAP_{OPT}, analogously to ECR circuits above.

C. Benchmarking results

1. QAOA for PortOpt

In this section, we first benchmark the performance of QAOA for PortOpt with ECR, global, and direct circuits. Then, we demonstrate the performance of QAOA with bipotent circuits and explore the pulse-level optimizations.

The results of QAOA-PortOpt with 3Q, 4Q, and 5Q are shown in Figs. 10(a), 10(b), and 10(c), respectively. The ECR-Default has the largest schedule duration. In the ECR-ZZ_{OPT}, the number of CX gates is reduced as the ZZ gates are optimized. Furthermore, the ZZ-SWAP_{OPT} gate is implemented by CZ_{OPT} gates, therefore there are no CX gates in the ECR-ZZ-SWAP_{OPT}.

In Fig. 10(a), the qubits used in the ECR circuits are [16, 14, 11], while in Global-D they are [18, 21, 23]. The ECR-ZZ-SWAP_{OPT} has the best performance, although the average CX gate error rate in ECR circuits is 1.6 times higher than that in the Global-D. Optimizing both ZZ and ZZ-SWAP

gates yields better and more stable AR and SP than optimizing only ZZ gates.

Figure 10(b) shows the results of QAOA using ECR circuits with qubits [16, 14, 11, 8], Global-B with qubits [18, 21, 23, 24], and direct circuit with qubits [1, 4, 7, 6]. Despite having the highest fidelity, Global-B does not perform the best. In comparison, ECR-ZZ-SWAP_{OPT} achieves the highest AR values, although its average CX gate error rate is 1.8 times higher and the duration is much longer. We believe this advantage comes from the combination of ZZ_{OPT} and ZZ-SWAP_{OPT} gates. The Direct-circuit has a shorter schedule duration than Global-B, resulting in comparable AR values and better SP values for higher p .

Figure 10(c) shows a performance comparison of 5Q-QAOA for PortOpt using ECR circuits with qubits [16, 14, 11, 8, 9], Global-B with qubits [17, 18, 21, 23, 24] including two ECR-CX gates, and direct circuit with qubits [6, 7, 4, 1, 2]. While the Global-B and direct circuit exhibit better fidelity and shorter schedule durations, ECR-ZZ-SWAP_{OPT} outperforms them with the highest AR values. Therefore, despite its longer duration, ECR-ZZ-SWAP_{OPT} can be considered as a promising alternative for achieving higher accuracy in quantum computing. A possible explanation is that combining ZZ_{OPT} and ZZ-SWAP_{OPT} produces a pulse that is more resilient to noise than the others.

We now benchmark the performance of QAOA with bipotent circuits (optimized for schedule duration rather than for fidelity). Figure 11(a) shows the results of 3Q-QAOA on ibmq_ehningen. The qubits used in bipotent circuits are [22, 25, 26] including one ECR-CX on qubits [22, 25] and one direct-CX on qubits [25, 26]. Compared to Bipotent-Default and Global-D, Bipotent-ZZ_{OPT} demonstrates better performance, as indicated by its higher AR and SP values. The decrease in the number of CX gates achieved through the use of optimized ZZ gates suggests that partial ZZ gates originally implemented with default CX gates have been optimized at the pulse level. This optimization leads to a reduction in schedule duration.

The bipotent circuits in Fig. 11(b) employ a set of five qubits, namely, [11, 14, 13, 12, 10], among which only a single pair of qubits is capable of performing an ECR-CX gate. The optimization of ZZ gates implemented using ECR-CX gates in Bipotent-ZZ_{OPT} leads to a reduction in the number of CX gates and a shorter schedule duration. Bipotent-ZZ-SWAP_{OPT} further optimizes ZZ-SWAP gates by replacing certain ECR-CX-based ZZ-SWAP gates with CZ_{OPT}-based ZZ-SWAP_{OPT} gates,

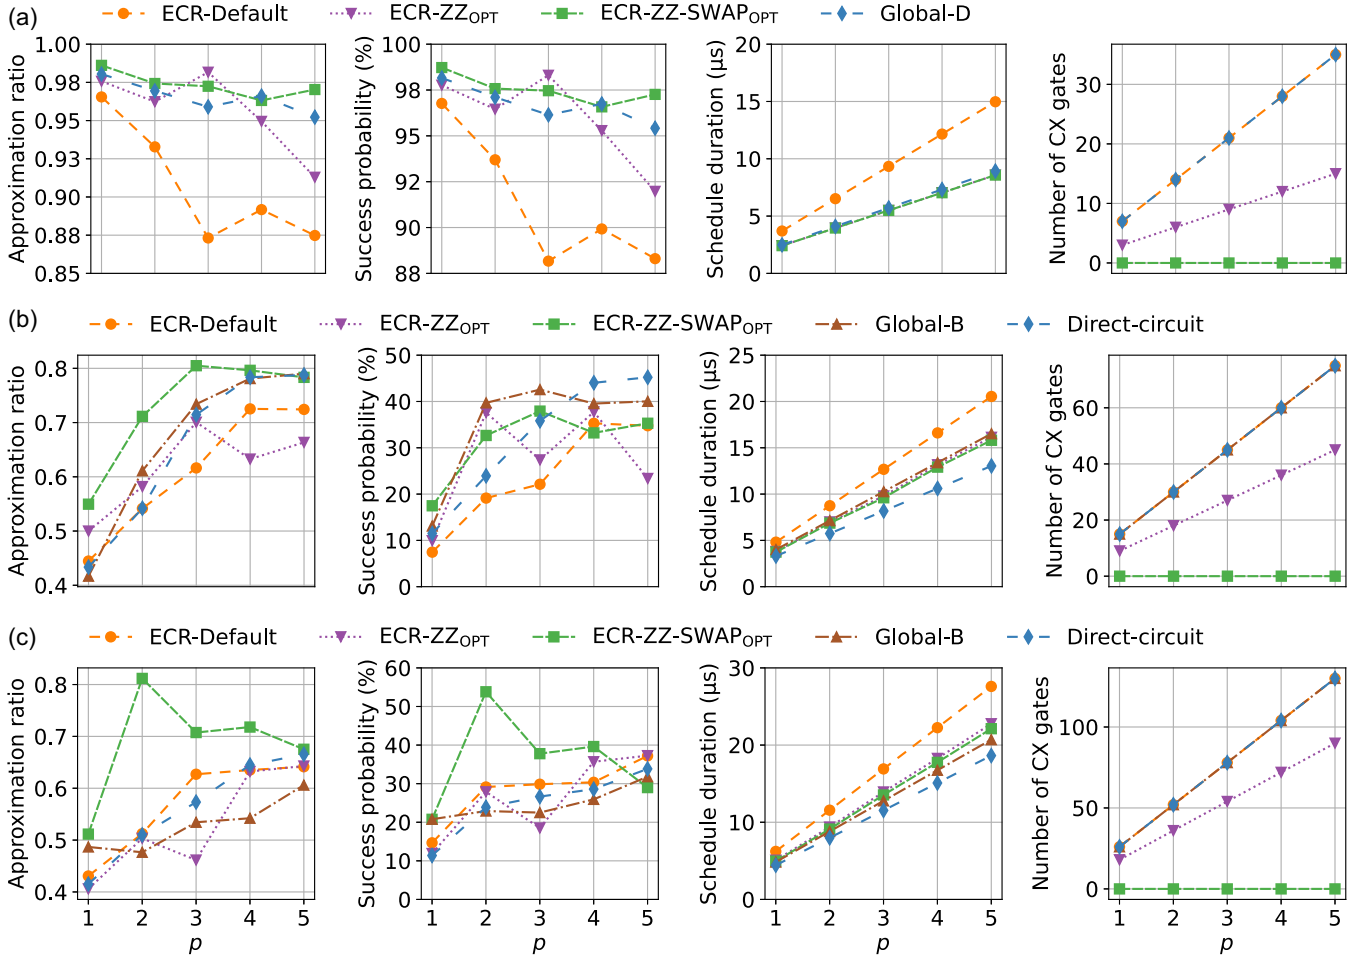


FIG. 10. Approximation ratio, success probability, schedule duration, and CX gate count of QAOA for PortOpt with p from 1 to 5 using ECR, global, and direct circuits on `ibmq_ehningen` for (a) 3Q, (b) 4Q, and (c) 5Q. The average CX gate error rates are 0.48% in all global circuits, they are 0.78%, 0.87%, and 0.89% in ECR circuits for 3Q, 4Q, and 5Q, respectively, while 0.48% and 0.60% in direct circuits for 4Q and 5Q, respectively.

thus further decreasing the number of CX gates. Although Global-B has better fidelity, the demonstrations show that bipotent circuits outperform it for $p \geq 2$. Overall, among all the methods tested, Bipotent-ZZ-SWAP_{OPT} exhibits the best performance in terms of the improvement in both the AR and SP metrics, making it a promising approach for quantum circuit optimization.

Figure 11(c) presents the benchmark results of 5Q-QAOA using qubits [11, 8, 5, 3, 2] on `ibmq_auckland`. It is worth noting that `ibmq_auckland` has a different bipotent architecture compared to `ibmq_ehningen`, as illustrated in Fig. 12. The results on `ibmq_auckland` are consistent with the trend of 5Q-QAOA for PortOpt observed on `ibmq_ehningen`. Specifically, we observe that QAOA with Bipotent-ZZ-SWAP_{OPT} outperforms the others. These results highlight the potential of the proposed optimization approach in improving the performance of quantum algorithms on other QPUs.

The data show that leveraging a combination of ZZ_{OPT} and ZZ-SWAP_{OPT} gates can significantly enhance both AR and SP, surpassing the performance achieved by solely optimizing ZZ gates. This approach also outperforms global circuits with the highest fidelity. Specifically, when compared to global circuits, we observe an improvement in AR

of up to 70%, on average 29%, for ECR-ZZ-SWAP_{OPT} at $p = 2$. Similarly, on `ibmq_auckland`, Bipotent-ZZ-SWAP_{OPT} shows an improvement of up to 55%, on average 27%, compared to Bipotent-Default. On `ibmq_ehningen`, Bipotent-ZZ-SWAP_{OPT} achieves an improvement of up to 34%, on average 14%, compared to Global-B.

2. QAOA for MaxCut

We now investigate the performance of QAOA for MaxCut with five qubits. The qubits used in ECR circuits, Global-B, and direct circuit are [16, 14, 11, 8, 9], [6, 7, 4, 1, 0], and [6, 7, 4, 1, 2], respectively. Figure 13(a) shows that Global-B, which has the highest fidelity, performs better for smaller p , whereas direct circuit, which has the lowest schedule duration, performs better for larger p . These results suggest that, when using default IBM gates, fidelity is more critical for medium-scale circuits, while a shorter schedule duration is more important for larger circuits. One possible explanation is that, as the size of the circuit increases, the effects of decoherence become dominant and prevail over the variability in qubit error rates. ECR circuits with higher error rates and longer schedule durations lead to lower AR compared to other

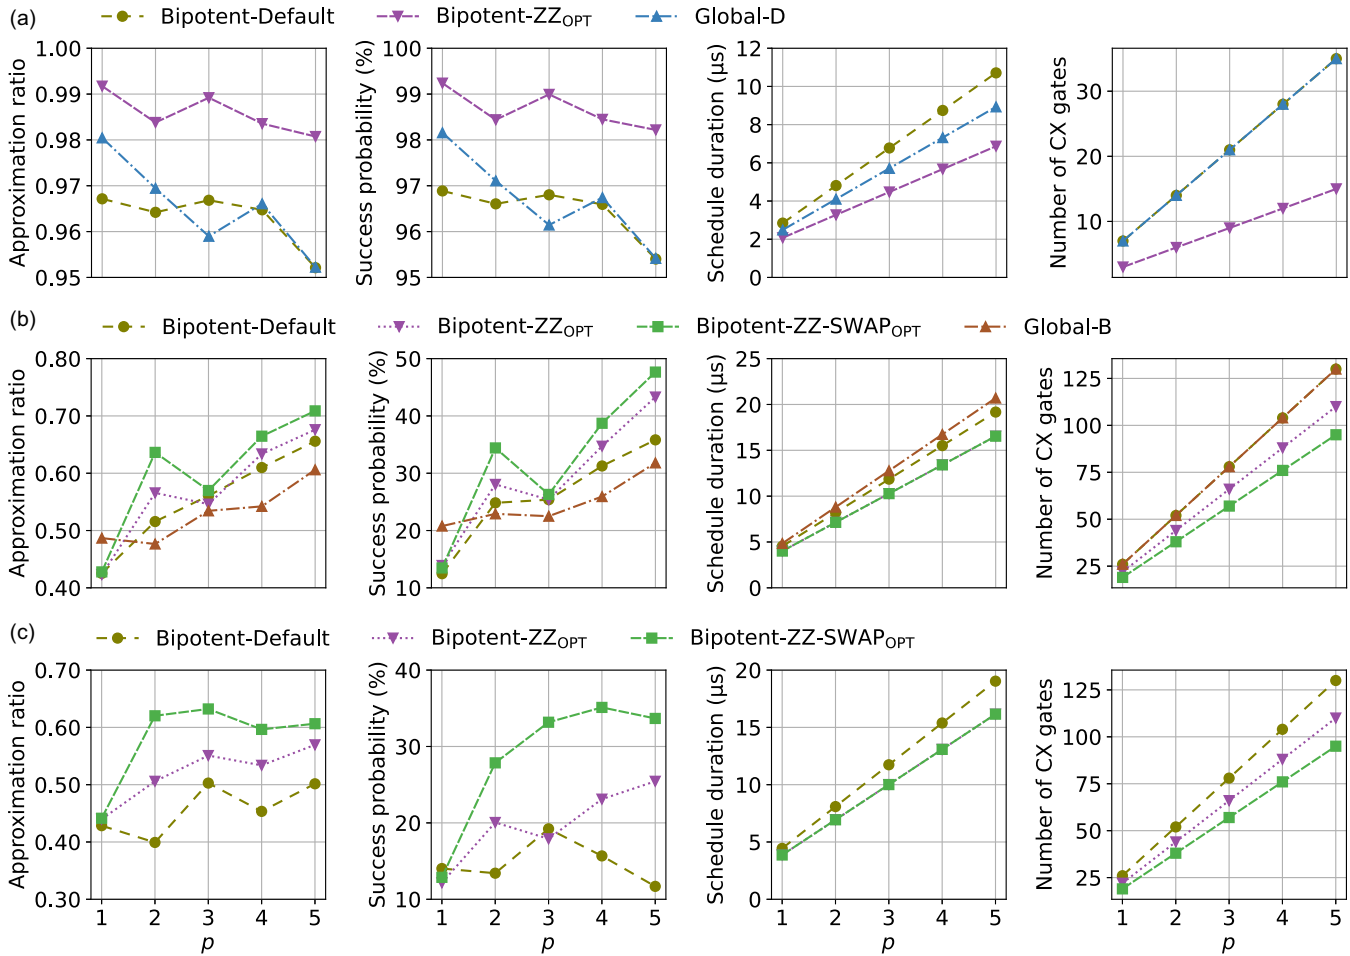


FIG. 11. Benchmarking results of QAOA for PortOpt using bipotent and global circuits on `ibmq_ehningen` for (a) 3Q and (b) 5Q, and (c) on `ibmq_auckland` for 5Q. The average CX gate error rates for each data point in (a), (b), and (c) are 0.60%, 0.65%, and 0.67%, respectively.

methods. However, by combining ZZ_{OPT} and $ZZ-SWAP_{OPT}$, $ECR-ZZ-SWAP_{OPT}$ achieves stable performance in both AR and SP.

Figure 13(b) shows the results of QAOA using bipotent circuits with qubits [17, 18, 21, 23, 24] including two $ECR-CX$ and two $direct-CX$ gates. While all three types of bipotent circuits have the same schedule duration, $Bipotent-ZZ-SWAP_{OPT}$ exhibits a clear advantage, as demonstrated by significant improvements in both AR and SP. However, $Global-B$ exhibits better performance for larger values of p due to its shorter duration.




FIG. 12. Bipotent architecture of `ibmq_auckland`. It has the same label as `ibmq_ehningen`.

Our study highlights the importance of efficient gate implementation, achieved through increasing fidelity and reducing schedule duration, in enhancing the algorithm’s performance. However, we found that the use of different gate types in a circuit can result in significantly divergent performances, even when the duration is identical. Therefore, careful consideration must be given to the selection of gate types to optimize the performance.

VI. EXTENSIONS

A. Beyond QAOA

While the primary focus of this paper is on optimizing QAOA on bipotent architectures without requiring additional calibration, the insights presented here have broader implications for other quantum algorithms and architectures. Specifically, some two-qubit gates in an algorithm may only achieve high fidelity after a time-consuming gate calibration, resulting in a notable increase in overhead costs [52]. However, our study suggests a *selective calibration* approach that can reduce the potential overhead by calibrating only the two-qubit gates implemented on specific qubit pairs with a higher CX gate error rate, instead of calibrating all two-qubit gates. This approach enhances algorithm performance by enabling more efficient resource utilization, leading to faster

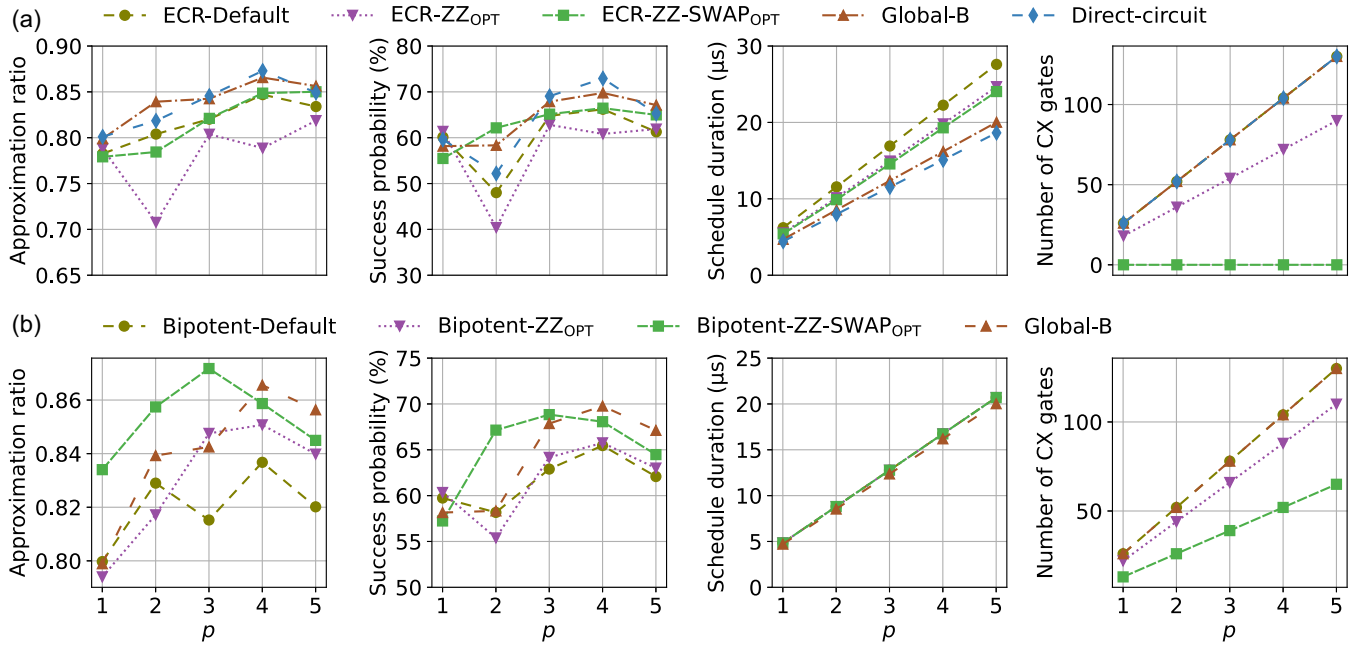


FIG. 13. Benchmarking results of 5Q-QAOA for MaxCut on `ibmq_ehningen` using two sets of circuits: (a) ECR, global, and direct circuits, and (b) bipotent and global circuits. The average CX gate error rates in the ECR circuits, Global-B, and direct circuit are 0.89%, 0.48%, and 0.60%, respectively, while that of each data point in (b) is 0.48%.

and more accurate results. Moreover, pulse-level optimization and various gate decomposition strategies can be combined to significantly improve algorithm performance. By investigating the resilience of different pulse shapes to noise, we can identify more robust pulse shapes that are less affected by noise and errors in the computation. Consequently, by leveraging these resilient pulse shapes, we can achieve better performance of quantum algorithms.

We optimize the algorithm on cross-resonance based hardware. The optimization is also applicable to other architectures, such as tunable couplers [53] that support interactions like ISWAP and CZ gates [54,55]. The implementation of ZZ or ZZ-SWAP gates can be realized using three sets of gates: CZ gates, ISWAP gates, and a combination of CZ and ISWAP gates. By tailoring distinct gate decompositions to the properties of the two-qubit gates on the target hardware, such as fidelity and schedule duration, we can achieve improved performance. Furthermore, integrating selective calibration at the pulse level makes this approach particularly promising for near-term quantum computing. Additionally, our study shows that decomposing ZZ-SWAP gates into CZ gates helps achieve improved results. Therefore, we expect that architectures with a native CZ gate may benefit even more from this optimization. Importantly, these techniques are not restricted to a particular algorithm and can be employed in various other quantum algorithms.

B. Error mitigation

In this paper, we applied readout error mitigation in QISKIT to mitigate measurement errors. However, there are additional techniques that can be utilized to suppress other types of errors in bipotent architectures. Here, we discuss two such

techniques: dynamical decoupling [56] and zero-noise extrapolation (ZNE) [57].

Dynamical decoupling is a widely used strategy to suppress decoherence errors [58] by applying pulse sequences to idle qubits. In bipotent architectures, the presence of two distinct CX gate types with unique pulse schedules can lead to significant variations in pulse lengths across different qubit pairs within the algorithm. Applying dynamical decoupling pulses to the idle regions of shorter pulses can be particularly advantageous for such architectures. Moreover, diverse gate decompositions on different qubit pairs and selective calibration can also result in noticeable differences in pulse lengths and larger idle regions. This underscores the importance of incorporating dynamical decoupling into this optimization strategy.

The ZNE technique estimates the noiseless result by utilizing expectation values measured at different noise levels. In bipotent architectures, two types of CX gates introduce different errors, which can be suppressed using ZNE. Additionally, ZNE can help identify the optimal gate decomposition for each qubit pair among various decomposition strategies used on different qubit pairs.

Thorough analysis, experimentation, and adaptation of these techniques to the unique characteristics of bipotent architectures are crucial for achieving improved performance on such architectures.

VII. CONCLUSIONS AND FUTURE WORK

Bipotent QPU architectures, which offer improved functions on some but not all of their qubits, provide new opportunities for quantum computations, but they also introduce new complexities. In this paper, we have investigated

TABLE V. Calibration data at the time of the demonstration presented in Sec. V C 1.

3Q-QAOA-PortOpt	ECR circuit	Global-D	Bipotent circuit
Qubits used	[16, 14, 11]	[18, 21, 23]	[22, 25, 26]
T_1 (μ s)	[194.12, 163.4, 150.43]	[214.69, 300.02, 222.14]	[195.66, 257.93, 195.59]
T_2 (μ s)	[240.94, 215.18, 174.7]	[333.28, 155.0, 234.21]	[31.47, 354.96, 27.85]
Frequency (GHz)	[5.022, 5.177, 5.119]	[4.996, 4.94, 4.805]	[4.725, 4.95, 5.151]
Anharmonicity (GHz)	[−0.3435, −0.3408, −0.3405]	[−0.343, −0.3456, −0.3471]	[−0.3464, −0.3457, −0.3391]
Prob meas0 prep1 (%)	[0.64, 0.92, 1.66]	[1.14, 1.0, 0.84]	[1.4, 1.16, 0.8]
Prob meas1 prep0 (%)	[0.56, 0.68, 1.4]	[0.86, 0.78, 0.76]	[1.12, 0.66, 0.66]
Readout length (ns)	[846.22, 846.22, 846.22]	[846.22, 846.22, 846.22]	[846.22, 846.22, 846.22]
Readout error (%)	[0.6, 0.8, 1.53]	[1.0, 0.89, 0.8]	[1.26, 0.91, 0.73]
Single-qubit gate error (%)	[0.016, 0.033, 0.022]	[0.026, 0.021, 0.013]	[0.016, 0.024, 0.014]
CX gate error (%)	[0.66, 0.89]	[0.44, 0.52]	[0.62, 0.57]
4Q-QAOA-PortOpt	ECR circuit	Global-B	Direct circuit
Qubits used	[16, 14, 11, 8]	[18, 21, 23, 24]	[1, 4, 7, 6]
T_1 (μ s)	[194.12, 163.4, 150.43, 145.9]	[214.69, 300.02, 222.14, 244.38]	[243.57, 134.77, 205.84, 105.37]
T_2 (μ s)	[240.94, 215.18, 174.7, 157.35]	[333.28, 155.0, 234.21, 225.4]	[172.38, 111.81, 266.71, 183.16]
Frequency (GHz)	[5.022, 5.177, 5.119, 5.174]	[4.996, 4.94, 4.805, 5.074]	[5.182, 5.054, 4.978, 4.89]
Anharmonicity (GHz)	[−0.3435, −0.3408, −0.3405, −0.3399]	[−0.343, −0.3456, −0.3471, −0.3416]	[−0.34, −0.3426, −0.344, −0.3448]
Prob meas0 prep1 (%)	[0.64, 0.92, 1.66, 1.18]	[1.14, 1.0, 0.84, 0.88]	[1.04, 1.02, 0.74, 1.48]
Prob meas1 prep0 (%)	[0.56, 0.68, 1.4, 1.46]	[0.86, 0.78, 0.76, 0.66]	[0.64, 0.58, 0.82, 2.02]
Readout length (ns)	[846.22, 846.22, 846.22, 846.22]	[846.22, 846.22, 846.22, 846.22]	[846.22, 846.22, 846.22, 846.22]
Readout error (%)	[0.6, 0.8, 1.53, 1.32]	[1.0, 0.89, 0.8, 0.77]	[0.84, 0.8, 0.78, 1.75]
Single-qubit gate error (%)	[0.016, 0.033, 0.022, 0.029]	[0.026, 0.021, 0.013, 0.017]	[0.026, 0.025, 0.016, 0.02]
CX gate error (%)	[0.66, 0.89, 1.06]	[0.44, 0.52, 0.49]	[0.68, 0.4, 0.35]
5Q-QAOA-PortOpt	ECR circuit	Global-B	
Qubits used	[16, 14, 11, 8, 9]	[17, 18, 21, 23, 24]	
T_1 (μ s)	[194.12, 163.4, 150.43, 145.9, 178.39]	[143.0, 214.69, 300.02, 222.14, 244.38]	
T_2 (μ s)	[240.94, 215.18, 174.7, 157.35, 195.4]	[39.76, 333.28, 155.0, 234.21, 225.4]	
Frequency (GHz)	[5.022, 5.177, 5.119, 5.174, 4.993]	[5.136, 4.996, 4.94, 4.805, 5.074]	
Anharmonicity (GHz)	[−0.3435, −0.3408, −0.3405, −0.3399, −0.3441]	[−0.341, −0.343, −0.3456, −0.3471, −0.3416]	
Prob meas0 prep1 (%)	[0.64, 0.92, 1.66, 1.18, 1.0]	[0.88, 1.14, 1.0, 0.84, 0.88]	
Prob meas1 prep0 (%)	[0.56, 0.68, 1.4, 1.46, 0.74]	[0.56, 0.86, 0.78, 0.76, 0.66]	
Readout length (ns)	[846.22, 846.22, 846.22, 846.22, 846.22]	[846.22, 846.22, 846.22, 846.22, 846.22]	
Readout error (%)	[0.6, 0.8, 1.53, 1.32, 0.87]	[0.72, 1.0, 0.89, 0.8, 0.77]	
Single-qubit gate error (%)	[0.016, 0.033, 0.022, 0.029, 0.02]	[0.018, 0.026, 0.021, 0.013, 0.017]	
CX gate error (%)	[0.66, 0.89, 1.06, 0.94]	[0.47, 0.44, 0.52, 0.49]	
5Q-QAOA-PortOpt	Direct circuit	Bipotent circuit	
Qubits used	[6, 7, 4, 1, 2]	[11, 14, 13, 12, 10]	
T_1 (μ s)	[42.89, 148.29, 108.74, 154.56, 75.75]	[127.56, 107.65, 147.9, 183.56, 115.74]	
T_2 (μ s)	[183.16, 266.71, 111.81, 172.38, 19.75]	[74.74, 196.1, 271.77, 442.33, 57.65]	
Frequency (GHz)	[4.89, 4.978, 5.054, 5.182, 5.127]	[5.119, 5.177, 4.926, 4.725, 4.835]	
Anharmonicity (GHz)	[−0.3448, −0.344, −0.3426, −0.34, −0.3403]	[−0.3405, −0.3408, −0.344, −0.3484, −0.3471]	
Prob meas0 prep1 (%)	[1.48, 0.74, 1.02, 1.04, 1.04]	[1.54, 1.08, 1.74, 1.0, 0.74]	
Prob meas1 prep0 (%)	[2.02, 0.82, 0.58, 0.64, 0.66]	[1.34, 0.44, 0.96, 0.9, 0.46]	
Readout length (ns)	[846.22, 846.22, 846.22, 846.22, 846.22]	[846.22, 846.22, 846.22, 846.22, 846.22]	
Readout error (%)	[1.75, 0.78, 0.8, 0.84, 0.85]	[1.44, 0.76, 1.35, 0.95, 0.6]	

TABLE V. (Continued.)

5Q-QAOA-PortOpt	Direct circuit	Bipotent circuit
Single-qubit gate error (%)	[0.02, 0.016, 0.025, 0.026, 0.05]	[0.025, 0.031, 0.018, 0.022, 0.017]
CX gate error (%)	[0.35, 0.4, 0.68, 0.96]	[0.84, 0.59, 0.66, 0.5]
5Q-QAOA-PortOpt	Bipotent circuit (ibm_auckland)	
Qubits used	[11, 8, 5, 3, 2]	
T_1 (μ s)	[168.1, 177.64, 280.77, 157.44, 126.14]	
T_2 (μ s)	[138.88, 97.6, 96.99, 219.48, 196.01]	
Frequency (GHz)	[5.055, 5.204, 4.993, 4.897, 5.006]	
Anharmonicity (GHz)	[-0.3422, -0.3407, -0.3445, -0.3455, -0.3434]	
Prob meas0 prep1 (%)	[0.84, 1.08, 1.14, 1.48, 1.28]	
Prob meas1 prep0 (%)	[0.5, 0.7, 0.74, 0.84, 0.98]	
Readout length (ns)	[757.33, 757.33, 757.33, 757.33, 757.33]	
Readout error (%)	[0.67, 0.89, 0.94, 1.16, 1.13]	
Single-qubit gate error (%)	[0.017, 0.019, 0.049, 0.016, 0.018]	
CX gate error (%)	[0.63, 0.98, 0.59, 0.49]	

the choice between improved direct-CX gates with better fidelity and duration and conventional ECR-CX gates that support pulse-level optimizations without additional calibration. Based on a careful validation of different circuits on two IBM's QPUs, we found that ECR-CX with pulse-level optimizations used to a maximum extent outperforms even seemingly globally optimal circuits with the highest fidelity. When comparing two approaches to constructing bipotent circuits, circuits that consider both fidelity and duration dominate circuits that focus only on fidelity.

We believe that the hardware-software codesign is an essential issue in emerging bipotent architectures. For instance,

the QAOA circuits investigated in this paper have a regular structure with relatively few quantum gate types and well-understood pulse-level optimizations. For other types of circuits with less pronounced effects of pulse-level optimizations, the improvements due to better basis gates may be more advantageous. An interesting avenue for future work is the development of universal and fully automatic methods to map a given quantum circuit to a bipotent QPU, thereby maximizing its probability of success and minimizing the duration of the resulting schedule. We believe that the existence of such methods will be a prerequisite for the practical success of bipotent quantum architectures. It would also be interesting

TABLE VI. Calibration data at the time of the demonstration presented in Sec. V C 2.

5Q-QAOA-MaxCut	ECR circuit	Global-B
Qubits used	[16, 14, 11, 8, 9]	[6, 7, 4, 1, 0]
T_1 (μ s)	[197.21, 207.15, 104.89, 122.12, 135.3]	[186.94, 217.29, 129.35, 143.59, 197.13]
T_2 (μ s)	[240.94, 215.18, 174.7, 157.35, 195.4]	[183.16, 266.71, 111.81, 172.38, 201.31]
Frequency (GHz)	[5.022, 5.177, 5.119, 5.174, 4.993]	[4.89, 4.978, 5.054, 5.182, 4.961]
Anharmonicity (GHz)	[-0.3435, -0.3408, -0.3405, -0.3399, -0.3441]	[-0.3448, -0.344, -0.3426, -0.34, -0.3445]
Prob meas0 prep1 (%)	[0.64, 0.92, 1.66, 1.18, 1.0]	[1.48, 0.74, 1.02, 1.04, 1.04]
Prob meas1 prep0 (%)	[0.56, 0.68, 1.4, 1.46, 0.74]	[2.02, 0.82, 0.58, 0.64, 0.88]
Readout length (ns)	[846.22, 846.22, 846.22, 846.22, 846.22]	[846.22, 846.22, 846.22, 846.22, 846.22]
Readout error (%)	[0.6, 0.8, 1.53, 1.32, 0.87]	[1.75, 0.78, 0.8, 0.84, 0.96]
Single-qubit gate error (%)	[0.016, 0.033, 0.022, 0.029, 0.02]	[0.02, 0.016, 0.025, 0.026, 0.018]
CX gate error (%)	[0.66, 0.89, 1.06, 0.94]	[0.35, 0.4, 0.68, 0.51]
5Q-QAOA-MaxCut	Direct circuit	Bipotent circuit
Qubits used	[6, 7, 4, 1, 2]	[17, 18, 21, 23, 24]
T_1 (μ s)	[42.89, 148.29, 108.74, 154.56, 75.75]	[143.0, 214.69, 300.02, 222.14, 244.38]
T_2 (μ s)	[183.16, 266.71, 111.81, 172.38, 19.75]	[39.76, 333.28, 155.0, 234.21, 225.4]
Frequency (GHz)	[4.89, 4.978, 5.054, 5.182, 5.127]	[5.136, 4.996, 4.94, 4.805, 5.074]
Anharmonicity (GHz)	[-0.3448, -0.344, -0.3426, -0.34, -0.3403]	[-0.341, -0.343, -0.3456, -0.3471, -0.3416]
Prob meas0 prep1 (%)	[1.48, 0.74, 1.02, 1.04, 1.04]	[0.88, 1.14, 1.0, 0.84, 0.88]
Prob meas1 prep0 (%)	[2.02, 0.82, 0.58, 0.64, 0.66]	[0.56, 0.86, 0.78, 0.76, 0.66]
Readout length (ns)	[846.22, 846.22, 846.22, 846.22, 846.22]	[846.22, 846.22, 846.22, 846.22, 846.22]
Readout error (%)	[1.75, 0.78, 0.8, 0.84, 0.85]	[0.72, 1.0, 0.89, 0.8, 0.77]
Single-qubit gate error (%)	[0.02, 0.016, 0.025, 0.026, 0.05]	[0.018, 0.026, 0.021, 0.013, 0.017]
CX gate error (%)	[0.35, 0.4, 0.68, 0.96]	[0.47, 0.44, 0.52, 0.49]

to explore the potential applicability of the proposed method on other QPUs such as Google's Sycamore processor [59] with a two-qubit hardware-native gate set {CZ, \sqrt{i} SWAP, Sycamore}.

Combining the proposed optimization strategies with error characterization, selective calibration, error correction, and various error mitigation techniques would be valuable in further enhancing the performance of algorithms on bipotent quantum architectures. Error characterization allows for a comprehensive understanding of the specific noise sources and patterns within the system. Selective calibration fine-tunes individual qubit pairs based on their unique characteristics, optimizing their performance. Error correction protects quantum information from errors caused by interactions with the environment [60]. Additionally, error mitigation suppresses errors and improves performance. We have discussed the potential benefits of utilizing dynamical decoupling and ZNE based on the properties of the two CX gate types specifically for bipotent architectures. Integrating these techniques improves the performance and reliability of algorithms on bipotent quantum architectures, enabling more accurate and efficient quantum computations.

ACKNOWLEDGMENTS

The authors would like to thank Joris Kattemölle and Xi Chen for their constructive criticism of the manuscript. The authors would also like to thank Thomas Wellens, John P. Hayes, Sebastian Brandhofer, and Andreas Ketterer for their useful discussions. This work was supported in part by the project QORA within the Competence Center Quantum Computing Baden-Württemberg (funded by the Ministerium für Wirtschaft, Arbeit und Tourismus Baden-Württemberg).

APPENDIX: CLOUD PLATFORM DETAILS

Here we provide more details about the two cloud-accessible QPUs `ibmq_ehningen` and `ibmq_auckland` used in our research. Both QPUs have identical layouts, as shown in Figs. 1 and 12. The basis gates of both QPUs include ID, RZ, SX, X, and CX, where ID is the identity gate, X is the Pauli X gate, and SX is the square root of the X gate. Calibration data at the time of the demonstration of QAOA for PortOpt and MaxCut are summarized in Tables V and VI, respectively.

-
- [1] P. Jurcevic, A. Javadi-Abhari, L. S. Bishop, I. Lauer, D. F. Bogorin, M. Brink, L. Capelluto, O. Günlük, T. Itoko, N. Kanazawa *et al.*, Demonstration of quantum volume 64 on a superconducting quantum computing system, *Quantum Sci. Technol.* **6**, 025020 (2021).
 - [2] 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 *Proceedings of the 24th International Conference on Architectural Support for Programming Languages and Operating Systems* (Association for Computing Machinery, Providence, RI, USA, 2019), pp. 1015–1029.
 - [3] N. Earnest, C. Tornow, and D. J. Egger, Pulse-efficient circuit transpilation for quantum applications on cross-resonance-based hardware, *Phys. Rev. Res.* **3**, 043088 (2021).
 - [4] P. Gokhale, A. Javadi-Abhari, N. Earnest, Y. Shi, and F. T. Chong, Optimized quantum compilation for near-term algorithms with OpenPulse, in *53rd Annual IEEE/ACM International Symposium on Microarchitecture, MICRO 2020, Athens, Greece, October 17–21, 2020* (IEEE, 2020), pp. 186–200.
 - [5] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, [arXiv:1411.4028](https://arxiv.org/abs/1411.4028) (2014).
 - [6] K. N. Smith, G. S. Ravi, T. Alexander, N. T. Bronn, A. Carvalho, A. Cervera-Lierta, F. T. Chong, J. M. Chow, M. Cubeddu, A. Hashim, L. Jiang, O. Lanes, M. J. Otten, D. I. Schuster, P. Gokhale, N. Earnest, and A. Galda, Summary: Chicago Quantum Exchange (CQE) pulse-level quantum control workshop, [arXiv:2202.13600](https://arxiv.org/abs/2202.13600) (2022).
 - [7] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, and S. J. Glaser, Optimal control of coupled spin dynamics: Design of NMR pulse sequences by gradient ascent algorithms, *J. Magn. Reson.* **172**, 296 (2005).
 - [8] F. Schwabl, *Quantum Mechanics* (Springer Science & Business Media, Berlin, 2007).
 - [9] J. Clarke and F. K. Wilhelm, Superconducting quantum bits, *Nature (London)* **453**, 1031 (2008).
 - [10] J. L. O'Brien, A. Furusawa, and J. Vučković, Photonic quantum technologies, *Nat. Photonics* **3**, 687 (2009).
 - [11] L. Jacak, P. Hawrylak, and A. Wojs, *Quantum Dots* (Springer Science & Business Media, Berlin, 2013).
 - [12] C. D. Bruzewicz, J. Chiaverini, R. McConnell, and J. M. Sage, Trapped-ion quantum computing: Progress and challenges, *Appl. Phys. Rev.* **6**, 021314 (2019).
 - [13] L. Henriët, L. Beguin, A. Signoles, T. Lahaye, A. Browaeys, G.-O. Reymond, and C. Jurczak, Quantum computing with neutral atoms, *Quantum* **4**, 327 (2020).
 - [14] J. Preskill, Quantum computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).
 - [15] Y. Ji, S. Brandhofer, and I. Polian, Calibration-aware transpilation for variational quantum optimization, in *2022 IEEE International Conference on Quantum Computing and Engineering (QCE)* (IEEE, Broomfield, Colorado, USA, 2022), pp. 204–214.
 - [16] S. Sivarajah, S. Dilkes, A. Cowtan, W. Simmons, A. Edgington, and R. Duncan, $t|ket\rangle$: A retargetable compiler for NISQ devices, *Quantum Sci. Technol.* **6**, 014003 (2021).
 - [17] M. Amy and V. Gheorghiu, staq-A full-stack quantum processing toolkit, *Quantum Sci. Technol.* **5**, 034016 (2020).
 - [18] M. P. Harrigan, K. J. Sung, M. Neeley, K. J. Satzinger, F. Arute, K. Arya, J. Atalaya, J. C. Bardin, R. Barends, S. Boixo *et al.*, Quantum approximate optimization of non-planar graph problems on a planar superconducting processor, *Nat. Phys.* **17**, 332 (2021).
 - [19] D. S. Steiger, T. Häner, and M. Troyer, ProjectQ: An open source software framework for quantum computing, *Quantum* **2**, 49 (2018).

- [20] 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 *Proceedings of the 46th International Symposium on Computer Architecture* (Association for Computing Machinery, Phoenix, Arizona, 2019), pp. 527–540.
- [21] N. Khammassi, I. Ashraf, J. Someren, R. Nane, A. Krol, M. A. Rol, L. Lao, K. Bertels, and C. G. Almudever, OpenQL: A portable quantum programming framework for quantum accelerators, *ACM J. Emerging Technol. Comput. Syst.* **18**, 1 (2021).
- [22] V. Bergholm, J. Izaac, M. Schuld, C. Gogolin, S. Ahmed, V. Ajith, M. S. Alam, G. Alonso-Linaje, B. AkashNarayanan, A. Asadi *et al.*, PennyLane: Automatic differentiation of hybrid quantum-classical computations, [arXiv:1811.04968](https://arxiv.org/abs/1811.04968) (2022).
- [23] M. Y. Siraichi, V. F. d. Santos, C. Collange, and F. M. Q. Pereira, Qubit allocation as a combination of subgraph isomorphism and token swapping, *Proc. ACM Program. Lang.* **3**, 1 (2019).
- [24] 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.* **38**, 1226 (2018).
- [25] 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 *Proceedings of the 24th International Conference on Architectural Support for Programming Languages and Operating Systems* (Association for Computing Machinery, Providence, RI, USA, 2019), pp. 987–999.
- [26] A. M. Childs, E. Schoute, and C. M. Unsal, Circuit Transformations for Quantum Architectures, in *14th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2019)* (Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2019), Vol. 135, pp. 3:1–3:24.
- [27] Z.-T. Li, F.-X. Meng, Z.-C. Zhang, and X.-T. Yu, Qubits’ mapping and routing for NISQ on variability of quantum gates, *Quant. Info. Proc.* **19**, 378 (2020).
- [28] S. Niu, A. Suau, G. Staffelbach, and A. Todri-Sanial, A hardware-aware heuristic for the qubit mapping problem in the NISQ era, *IEEE Trans. Quantum Eng.* **1**, 3101614 (2020).
- [29] D. Bhattacharjee and A. Chattopadhyay, Depth-optimal quantum circuit placement for arbitrary topologies, [arXiv:1703.08540](https://arxiv.org/abs/1703.08540) (2017).
- [30] A. Shafaei, M. Saeedi, and M. Pedram, Qubit placement to minimize communication overhead in 2D quantum architectures, in *2014 19th Asia and South Pacific Design Automation Conference (ASP-DAC)* (IEEE, SunTec, Singapore, 2014), pp. 495–500.
- [31] M. Pedram and A. Shafaei, Layout optimization for quantum circuits with linear nearest neighbor architectures, *IEEE Circuits Syst. Mag.* **16**, 62 (2016).
- [32] B. Tan and J. Cong, Optimal layout synthesis for quantum computing, in *Proceedings of the 39th International Conference on Computer-Aided Design* (Association for Computing Machinery, 2020), pp. 1–9.
- [33] Y. Ji and I. Polian, Algorithm-aware qubit mapping, [TechRxiv 10.36227/techrxiv.21786146.v1](https://arxiv.org/abs/10.36227/techrxiv.21786146.v1) (2022).
- [34] Y. Ji, Modellierung der messung von flussqubits mittels quanten-fluss-parametron, [arXiv:2301.11870](https://arxiv.org/abs/2301.11870) (2022).
- [35] P. Gokhale, Y. Ding, T. Propson, C. Winkler, N. Leung, Y. Shi, D. I. Schuster, H. Hoffmann, and F. T. Chong, Partial compilation of variational algorithms for noisy intermediate-scale quantum machines, in *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture* (Association for Computing Machinery, Columbus, OH, USA, 2019), pp. 266–278.
- [36] Y. Shi, N. Leung, P. Gokhale, Z. Rossi, D. I. Schuster, H. Hoffmann, and F. T. Chong, Optimized compilation of aggregated instructions for realistic quantum computers, in *Proceedings of the 24th International Conference on Architectural Support for Programming Languages and Operating Systems* (Association for Computing Machinery, Providence, RI, USA, 2019), pp. 1031–1044.
- [37] P. Gokhale, T. Tomesh, M. Suchara, and F. T. Chong, Faster and more reliable quantum swaps via native gates, [arXiv:2109.13199](https://arxiv.org/abs/2109.13199) (2021).
- [38] Z. Liang, Z. Song, J. Cheng, Z. He, J. Liu, H. Wang, R. Qin, Y. Wang, S. Han, X. Qian, and Y. Shi, Hybrid gate-pulse model for variational quantum algorithms, [arXiv:2212.00661](https://arxiv.org/abs/2212.00661) (2022).
- [39] D. J. Egger, C. Capecchi, B. Pokharel, P. K. Barkoutsos, L. E. Fischer, L. Guidoni, and I. Tavernelli, A study of the pulse-based variational quantum eigensolver on cross-resonance based hardware, [arXiv:2303.02410](https://arxiv.org/abs/2303.02410) (2023).
- [40] Qiskit: An Open-Source Framework for Quantum Computing, <https://qiskit.org>, Accessed 2022-10-31.
- [41] E. C. Peterson, L. S. Bishop, and A. Javadi-Abhari, Optimal synthesis into fixed XX interactions, *Quantum* **6**, 696 (2022).
- [42] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio *et al.*, Variational quantum algorithms, *Nat. Rev. Phys.* **3**, 625 (2021).
- [43] J. S. Baker and S. K. Radha, Wasserstein solution quality and the quantum approximate optimization algorithm: A portfolio optimization case study, [arXiv:2202.06782](https://arxiv.org/abs/2202.06782) (2022).
- [44] S. Brandhofer, D. Braun, V. Dehn, G. Hellstern, M. Hüls, Y. Ji, I. Polian, A. S. Bhatia, and T. Wellens, Benchmarking the performance of portfolio optimization with QAOA, *Quant. Info. Proc.* **22**, 25 (2023).
- [45] D. J. Egger, C. Gambella, J. Marecek, S. McFaddin, M. Mevissen, R. Raymond, A. Simonetto, S. Woerner, and E. Yndurain, Quantum computing for finance: State-of-the-art and future prospects, *IEEE Trans. Quantum Eng.* **1**, 3101724 (2020).
- [46] N. N. Hegade, P. Chandarana, K. Paul, X. Chen, F. Albarrán-Arriagada, and E. Solano, Portfolio optimization with digitized counterdiabatic quantum algorithms, *Phys. Rev. Res.* **4**, 043204 (2022).
- [47] F. Glover, G. Kochenberger, R. Hennig, and Y. Du, Quantum bridge analytics I: A tutorial on formulating and using QOBO models, *Ann. Oper. Res.* **314**, 141 (2022).
- [48] J. Wurtz and P. Love, Maxcut quantum approximate optimization algorithm performance guarantees for $p > 1$, *Phys. Rev. A* **103**, 042612 (2021).
- [49] P. Chandarana, N. N. Hegade, K. Paul, F. Albarrán-Arriagada, E. Solano, A. del Campo, and X. Chen, Digitized-counterdiabatic quantum approximate optimization algorithm, *Phys. Rev. Res.* **4**, 013141 (2022).
- [50] A. Hashim, R. Rines, V. Omole, R. K. Naik, J. M. Kreikebaum, D. I. Santiago, F. T. Chong, I. Siddiqi, and P. Gokhale, Optimized SWAP networks with equivalent circuit averaging for QAOA, *Phys. Rev. Res.* **4**, 033028 (2022).

- [51] Mapomatic: Automatic mapping of compiled circuits to low-noise sub-graphs, <https://github.com/Qiskit-Partners/mapomatic>, Accessed 2022-10-31.
- [52] L. Lao, P. Murali, M. Martonosi, and D. Browne, Designing calibration and expressivity-efficient instruction sets for quantum computing, in *2021 ACM/IEEE 48th Annual International Symposium on Computer Architecture (ISCA)* (IEEE, 2021), pp. 846–859.
- [53] D. C. McKay, S. Filipp, A. Mezzacapo, E. Magesan, J. M. Chow, and J. M. Gambetta, Universal Gate for Fixed-Frequency Qubits via a Tunable Bus, *Phys. Rev. Appl.* **6**, 064007 (2016).
- [54] M. Ganzhorn, G. Salis, D. J. Egger, A. Fuhrer, M. Mergenthaler, C. Müller, P. Müller, S. Paredes, M. Pechal, M. Werninghaus, and S. Filipp, Benchmarking the noise sensitivity of different parametric two-qubit gates in a single superconducting quantum computing platform, *Phys. Rev. Res.* **2**, 033447 (2020).
- [55] Y. Sung, L. Ding, J. Braumüller, A. Vepsäläinen, B. Kannan, M. Kjaergaard, A. Greene, G. O. Samach, C. McNally, D. Kim, A. Melville, B. M. Niedzielski, M. E. Schwartz, J. L. Yoder, T. P. Orlando, S. Gustavsson, and W. D. Oliver, Realization of High-Fidelity CZ and ZZ-free iSWAP Gates with a Tunable Coupler, *Phys. Rev. X* **11**, 021058 (2021).
- [56] B. Pokharel, N. Anand, B. Fortman, and D. A. Lidar, Demonstration of Fidelity Improvement Using Dynamical Decoupling with Superconducting Qubits, *Phys. Rev. Lett.* **121**, 220502 (2018).
- [57] K. Temme, S. Bravyi, and J. M. Gambetta, Error Mitigation for Short-Depth Quantum Circuits, *Phys. Rev. Lett.* **119**, 180509 (2017).
- [58] L. Viola and S. Lloyd, Dynamical suppression of decoherence in two-state quantum systems, *Phys. Rev. A* **58**, 2733 (1998).
- [59] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. Brandao, D. A. Buell *et al.*, Quantum supremacy using a programmable superconducting processor, *Nature (London)* **574**, 505 (2019).
- [60] C. H. Bennett, D. P. DiVincenzo, J. A. Smolin, and W. K. Wootters, Mixed-state entanglement and quantum error correction, *Phys. Rev. A* **54**, 3824 (1996).