

## Pulse-efficient circuit transpilation for quantum applications on cross-resonance-based hardware

Nathan Earnest,<sup>1</sup> Caroline Tornow,<sup>2,3</sup> and Daniel J. Egger<sup>3,\*</sup><sup>1</sup>IBM Quantum, IBM T. J. Watson Research Center, Yorktown Heights, New York 10598, USA<sup>2</sup>Institute for Theoretical Physics, ETH Zurich 8093, Switzerland<sup>3</sup>IBM Quantum, IBM Research, Zurich, Säumerstrasse 4, 8803 Rüschlikon, Switzerland

(Received 20 May 2021; revised 6 September 2021; accepted 10 September 2021; published 29 October 2021)

We show a pulse-efficient circuit transpilation framework for noisy quantum hardware. This is achieved by scaling cross-resonance pulses and exposing each pulse as a gate to remove redundant single-qubit operations with the transpiler. Crucially, no additional calibration is needed to yield better results than a CNOT-based transpilation. This pulse-efficient circuit transpilation therefore enables a better usage of the finite coherence time without requiring knowledge of pulse-level details from the user. As demonstration, we realize a continuous family of cross-resonance-based gates for SU(4) by leveraging Cartan's decomposition. We measure the benefits of a pulse-efficient circuit transpilation with process tomography and observe up to a 50% error reduction in the fidelity of  $R_{ZZ}(\theta)$  and arbitrary SU(4) gates on IBM Quantum devices. We apply this framework for quantum applications by running circuits of the quantum approximate optimization algorithm applied to MAXCUT. For an 11-qubit nonhardware native graph, our methodology reduces the overall schedule duration by up to 52% and errors by up to 38%.

DOI: [10.1103/PhysRevResearch.3.043088](https://doi.org/10.1103/PhysRevResearch.3.043088)

## I. INTRODUCTION

Quantum computers have the potential to impact a broad range of disciplines such as quantum chemistry [1], finance [2,3], optimization [4,5], and machine learning [6,7]. The performance of noisy quantum computers has been improving as measured by metrics such as the quantum volume [8,9] or the coherence of superconducting transmon-based devices [10–12], which has exceeded  $100 \mu\text{s}$  [13,14]. To overcome limitations set by the noise, several error mitigation techniques such as readout error mitigation [15,16] and Richardson extrapolation [17,18] have been developed. Gate families with continuous parameters further improve results [19–23] as they require less coherence time than circuits in which the CNOT is the only two-qubit gate. Aggregating instructions and optimizing the corresponding pulses, using, e.g., gradient ascent algorithms such as GRAPE [24], reduces the duration of the pulse schedules [25]. However, such pulses require calibration to overcome model errors [26,27], which typically needs closed-loop optimization [28,29] and sophisticated readout methods [30,31]. This may therefore be difficult to scale because calibration is time consuming and increasingly harder as the control pulses become more complex. Some of these limitations may be overcome with novel control methods [32] and error robust control [33,34].

Since calibrating a two-qubit gate is time-consuming, IBM Quantum [14] backends only expose a calibrated CNOT gate built from echoed cross-resonance pulses [35,36] with rotary tones [37]. Quantum circuit users must therefore transpile their circuits to CNOT gates, which often makes a poor usage of the limited coherence time. With the help of Qiskit pulse [38,39] users may extend the set of two-qubit gates [40–42]. Such gates can in turn generate other multiqubit gates more effectively than when the CNOT gate is the only two-qubit gate available [41]. However, creating these gates comes at the expense of additional calibration experiments, which is often impractical on a queue-based quantum computer. Furthermore, only a limited number of users can access these benefits due to the need for an intimate familiarity with quantum control. In Ref. [43] the authors show a pulse-scaling methodology to create the control pulses for the continuous gate set  $R_{ZX}(\theta)$ , which they leverage to create  $R_{YX}(\theta)$  gates and manually assemble into pulse schedules. Crucially, the scaled pulses improve gate fidelity without requiring any additional calibration experiments to determine the new values of the pulse parameters.

Here we extend the methodology of Ref. [43] to arbitrary SU(4) gates and show how to make pulse-efficient circuit transpilation available to general users without having to manipulate pulse schedules. In Sec. II we review the pulse-scaling methodology of Ref. [43] and carefully benchmark the performance of  $R_{ZZ}$  gates. Next, in Sec. III we leverage this pulse-efficient gate generation to create arbitrary SU(4) gates, which we benchmark with quantum process tomography [44,45]. In Sec. IV we show how pulse-efficient gates can be included in automated circuit transpiler passes. Finally, in Sec. V we demonstrate the advantage of our pulse-efficient transpilation by applying it to the quantum approximate optimization algorithm (QAOA) [4].

\*deg@zurich.ibm.com

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

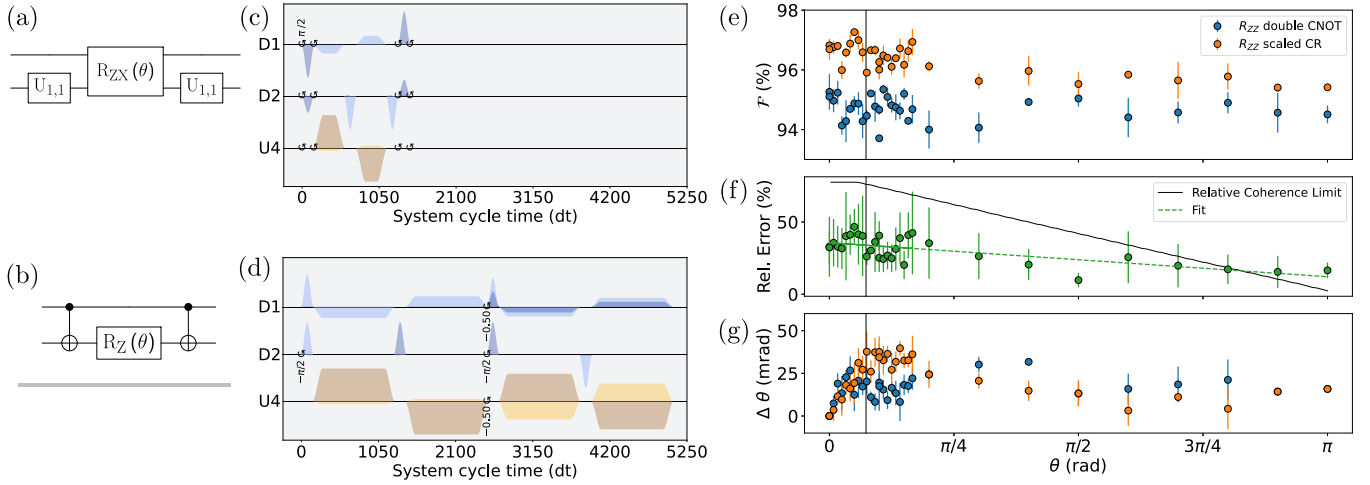


FIG. 1.  $R_{ZZ}(\theta)$  characterization for qubits one and two on *ibmq\_paris*. (a) Continuous gate implementation where  $U_{1,1} = R_Z(\pi/2)\sqrt{X}R_Z(\pi/2)$  and its corresponding schedule in (c) for  $\theta = 0.5$ . Here  $R_{ZX}(\theta)$  is a scaled cross-resonance pulse with a built-in echo. The virtual Z gates, indicated by the  $\odot$  symbols, all have an angle  $\pi/2$  in (c). (b) Double-CNOT benchmark and its corresponding schedule in (d) for  $\theta = 0.5$ . (e) Gate fidelity  $\mathcal{F}[U_{\text{meas}}, R_{ZZ}(\theta)]$  of the double-CNOT implementation (blue) and the scaled cross-resonance pulses (orange). The vertical line indicates the angle at which  $w = 0$ . (f) The relative error between the two implementations (green dots), and the theoretical expectations for a coherence limited gate (solid black line). (g) The deviation angle  $\Delta\theta = \theta - \theta_{\text{max}}$  corresponding to the data in (e) that achieves the maximum gate fidelity  $\mathcal{F}[U_{\text{meas}}, R_{ZZ}(\theta_{\text{max}})]$ .

## II. SCALING HARDWARE-NATIVE CROSS-RESONANCE GATES

We consider an all-microwave fixed-frequency transmon architecture that implements the echoed cross-resonance gate [36]. A two-qubit system in which a control qubit is driven at the frequency of a target qubit and evolves under the time-dependent cross-resonance Hamiltonian  $H_{cr}(t)$ . The time-independent approximation of  $H_{cr}(t)$  is

$$\bar{H}_{cr} = \frac{1}{2}(Z \otimes B + I \otimes C), \quad (1)$$

where  $B = \omega_{ZI}I + \omega_{ZX}X + \omega_{ZY}Y + \omega_{ZZ}Z$  and  $C = \omega_{IX}X + \omega_{IY}Y + \omega_{IZ}Z$ . Here  $X$ ,  $Y$ , and  $Z$  are Pauli matrices,  $I$  is the identity, and  $\omega_{ij}$  are drive strengths. An echo sequence [36] and rotary tones [37] isolate the  $ZX$  interaction, which ideally results in the unitary  $R_{ZX}(\theta) = \exp(-i\theta ZX/2)$ . The rotation angle  $\theta$  is  $t_{cr}\omega_{ZX}(\bar{A})$  where  $t_{cr}$  is the duration of the cross-resonance drive. The drive strength  $\omega_{ZX}$  has a nonlinear dependency on the average drive amplitude  $\bar{A}$  as shown by a third-order approximation of the cross-resonance Hamiltonian [39,46].

IBM Quantum systems expose to their users a calibrated CNOT gate built from  $R_{ZX}(\pi/2)$  rotations implemented by the echoed cross-resonance gate. The pulse sequence of  $R_{ZX}(\pi/2)$  on the control qubit is  $\text{CR}(\pi/4)X\text{CR}(-\pi/4)X$ . Here  $\text{CR}(\pm\pi/4)$  are flat-top pulses of amplitude  $A^*$ , width  $w^*$ , and Gaussian flanks with standard deviations  $\sigma$ , truncated after  $n_\sigma$  times  $\sigma$ . Their area is  $\alpha^* = |A^*|[w^* + \sqrt{2\pi}\sigma\text{erf}(n_\sigma)]$  where the star superscript refers to the parameter values of the calibrated pulses in the CNOT gate. During each CR pulse rotary tones are applied to the target qubit to help reduce the magnitude of the undesired  $\omega_{IY}$  interaction. We can create  $R_{ZX}(\theta)$  rotations by scaling the area of the CR and rotary pulses following  $\alpha(\theta) = 2\theta\alpha^*/\pi$  as done in Ref. [43]. To create a target area  $\alpha(\theta)$  we first scale  $w$  to minimize the

effect of the nonlinearity between the drive strength  $\omega_{ZX}(\bar{A})$  and the pulse amplitude. When  $\alpha(\theta) < |A^*|\sigma\sqrt{2\pi}\text{erf}(n_\sigma)$  we set  $w = 0$  and scale the pulse amplitude such that  $|A(\theta)| = \alpha(\theta)/[\sigma\sqrt{2\pi}\text{erf}(n_\sigma)]$ .

We investigate the effect of the pulse scaling methodology with quantum process tomography by carefully benchmarking scaled  $R_{ZZ}(\theta)$  gates [see Fig. 1(a)] with respect to the double-CNOT decomposition; see Fig. 1(b). The schedule of the scaled gate is significantly shorter than the double-CNOT benchmark for small angles; compare Fig. 1(c) with Fig. 1(d). We measure the process fidelity  $\mathcal{F}[U_{\text{meas}}, R_{ZZ}(\theta)]$  between the target gate  $R_{ZZ}(\theta)$  and the measured gate  $U_{\text{meas}}$ . To determine  $U_{\text{meas}}$  we prepare each qubit in  $|0\rangle$ ,  $|1\rangle$ ,  $(|0\rangle + |1\rangle)/\sqrt{2}$ , and  $(|0\rangle + i|1\rangle)/\sqrt{2}$  and measure in the  $X$ ,  $Y$ , and  $Z$  bases. Two-qubit process tomography therefore requires a total of 148 circuits for each angle of interest, which includes four circuits needed to mitigate readout errors [15,16]. The scaled pulses consistently have a better fidelity than the double CNOT benchmark as demonstrated by the data gathered on *ibmq\_paris* with qubits one and two; see Fig. 1(e). Appendix B shows key device parameters and additional data taken on other devices, which illustrate the reliability of the methodology. The relative error reduction of the measured gate fidelity correlates well to the relative error reduction of the coherence limited average gate fidelity [37,47,48]; see Fig. 1(f) and details in Appendix C. We therefore attribute the error reduction to the shorter schedules as they use less coherence time.

In addition to the gate fidelity, we compare the deviation  $\Delta\theta$  from the target angle of both implementations of the  $R_{ZZ}(\theta)$  rotation. The deviation  $\Delta\theta$  is the difference between the target rotation angle  $\theta$  and the angle  $\theta_{\text{max}}$ , which satisfies  $\mathcal{F}[U_{\text{meas}}, R_{ZZ}(\theta_{\text{max}})] \geq \mathcal{F}[U_{\text{meas}}, R_{ZZ}(\theta')]$   $\forall \theta'$ . Since the  $R_Z(\theta)$  is virtual [49] the implementation with two CNOT gates does not depend on the desired target angle; see Fig. 1(g).

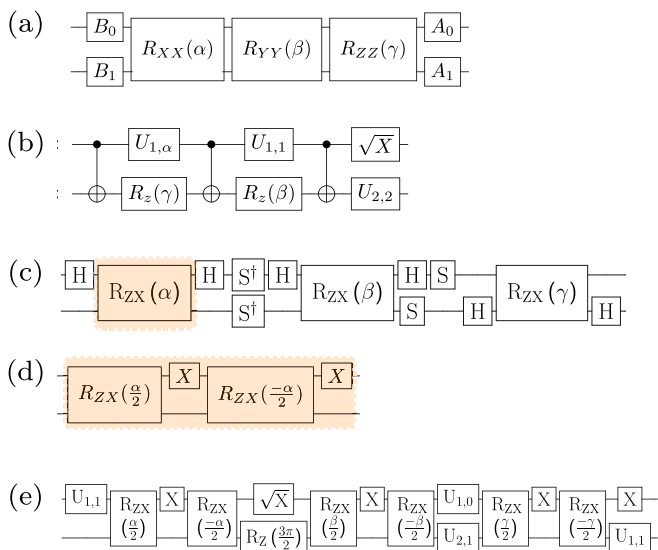


FIG. 2. Cartan’s  $KAK$  decomposition. (a) Circuit representation of the  $KAK$  decomposition of a two-qubit gate  $U \in SU(4)$  with  $k_1 = (A_1 \otimes A_0)$  and  $k_2 = (B_1 \otimes B_0)$ . (b) Circuit in (a) without  $k_{1,2}$  and decomposed into three CNOT gates and transpiled to the basis gates ( $R_Z(\theta)$ ,  $\sqrt{X}$ , CNOT). (c) Circuit in (a) decomposed into the hardware-native  $R_{ZX}$  gates. Here each  $R_{ZX}$  gate has a built-in echo as shown in (d). Transpiling circuit (c) to the basis ( $R_Z(\theta)$ ,  $\sqrt{X}$ ,  $X$ ,  $R_{ZX}(\theta)$ ) with the echoes exposed to the transpiler results in the pulse-efficient circuit shown in (e) where the scaled  $R_{ZX}$  gates do not have an echo. We replaced  $R_Z(n\pi/2)\sqrt{X}R_Z(m\pi/2)$  with  $U_{n,m}$  and  $U_{1,\alpha} = R_Z(\pi/2)\sqrt{X}R_Z(\alpha)$  to shorten the notation.

However, the scaled gate has two competing nonlinearities: an expected nonlinearity from the amplitude scaling and an unexpected one from scaling the width. As the width is scaled down, the angle deviation increases from  $\sim 10$  mrad to  $\sim 35$  mrad. Once the amplitude scaling begins, a nonlinearity arises which reduces the deviation angle of the scaled gates. At  $\alpha(\theta) \approx |A^*|\sigma\sqrt{2\pi}\text{erf}(n\sigma)/2$  the angle deviation of the scaled gates once again matches the deviation of the benchmark within the measured standard deviation.

### III. CREATING ARBITRARY SU(4) GATES

We now generalize the results from Sec. II. Cartan’s decomposition of an arbitrary two-qubit gate  $U \in SU(4)$  is  $U = k_1 A k_2$ , which we refer to as Cartan’s  $KAK$  decomposition [50]. Here  $k_1$  and  $k_2$  are local operations, i.e.,  $k_{1,2} \in SU(2) \otimes SU(2)$ , and  $A = e^{ik^T \cdot \Sigma/2} \in SU(4) \setminus SU(2) \otimes SU(2)$  is a nonlocal operation with  $\Sigma^T = (XX, YY, ZZ)$  [51–53]; see Fig. 2(a). The nonlocal term is defined by the three angles  $\mathbf{k}^T = (\alpha, \beta, \gamma) \in \mathbb{R}^3$  satisfying  $\alpha + \beta + \gamma \leq 3\pi/2$  and  $\pi \geq \alpha \geq \beta \geq \gamma \geq 0$ . Geometrically, the  $KAK$  decomposition is represented in a tetrahedron known as the Weyl chamber in the three-dimensional space; see Fig. 3. Every point  $(\alpha, \beta, \gamma)$  in the Weyl chamber (except in the base) defines a continuous set of two-qubit gates equivalent up to single-qubit rotations [51]. For instance, the point  $(\frac{\pi}{2}, 0, 0)$ , labeled as  $C$  in Fig. 3, corresponds to the local equivalence class of the CNOT gate,




FIG. 3. Weyl chamber of  $SU(4)$ . The coordinates of the chamber are  $O = (0, 0, 0)$ ,  $A_1 = (\pi, 0, 0)$ ,  $A_2 = (\frac{\pi}{2}, \frac{\pi}{2}, 0)$ , and  $A_3 = (\frac{\pi}{2}, \frac{\pi}{2}, \frac{\pi}{2})$ .  $C$  corresponds to the CNOT gate. The blue dots represent the data from Fig. 4 taken on *ibmq\_mumbai*.

and the point  $(\frac{\pi}{2}, \frac{\pi}{2}, \frac{\pi}{2})$ , labeled as  $A_3$ , represents the SWAP gate.

Since the rotations generated by  $XX$ ,  $YY$ , and  $ZZ$  are locally equivalent to rotations generated by  $ZX$  we decompose the nonlocal  $e^{ik^T \cdot \Sigma/2}$  term into a circuit with three  $R_{ZX}$  rotations; see Fig. 2(c). We shorten the total duration of the circuit by exposing the echo in the cross-resonance gate [see Fig. 2(d)] to the transpiler. This ensures that at most one single-qubit pulse is needed on each qubit between each nonechoed cross-resonance  $R_{ZX}$  gate. By scaling the cross-resonance pulses we create the  $R_{ZX}$  gates for arbitrary angles and therefore generalize the methods of Sec. II to arbitrary gates in  $SU(4)$ .

We generate  $R_{ZX}$ -based circuits as shown in Fig. 2(e) for  $(\alpha, \beta, \gamma)$  angles chosen at random from the Weyl chamber and measure their fidelity using process tomography with read-out error mitigation. Each  $R_{ZX}$ -based circuit is benchmarked against its equivalent three CNOT decomposition presented in Ref. [54] and shown in Fig. 2(b). The experiments are run on *ibmq\_dublin* and *ibmq\_mumbai* with 2048 shots for each circuit, which we measure three times to gain statistics. The pulse-efficient  $R_{ZX}$ -based decomposition of the circuits results in a significant fidelity increase for almost all angles; see Fig. 4. A subset of the data is also shown in the Weyl chamber in Fig. 3. We observe that not all data points taken on *ibmq\_dublin* consistently improve the fidelity, see also Appendix B. We attribute this to either noise in *ibmq\_dublin* or a nonlinear response in the pulse transfer function. The correlation between the relative error reduction and the relative schedule duration indicates that the gains in fidelity come from a better usage of the finite coherence time as the scaled cross-resonance pulses achieve the same unitary in less time. Remarkably, these results were achieved without recalibrating any pulses.

### IV. PULSE-EFFICIENT TRANSPILER PASSES

The quantum circuits of an algorithm are typically expressed using generic gates such as the CNOT or controlled-

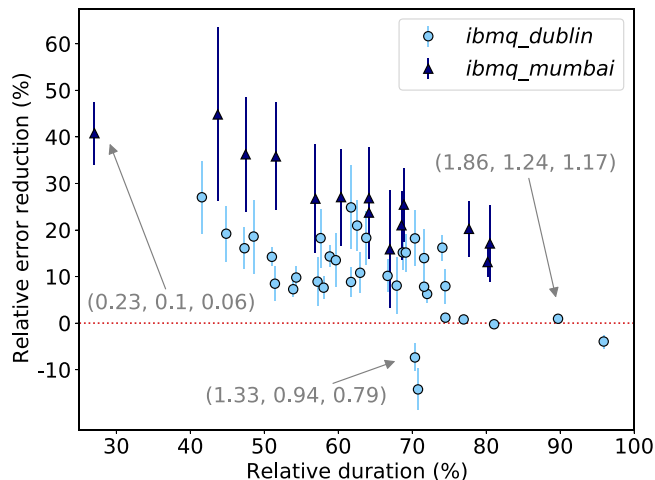


FIG. 4. Gate error reduction of the pulse-efficient SU(4) gates relative to the three CNOT benchmark for random angles in the Weyl chamber measured on *ibmq\_dublin*, qubits one and two (light blue circles), and *ibmq\_mumbai*, qubits 19 and 16 (dark blue triangles). The x-axis is the duration of the pulse-efficient SU(4) gates relative to the three CNOT benchmark. The angles of three gates are indicated in parentheses as example.

phase gate and then transpiled to the hardware on which they are run [55]. Quantum algorithms can benefit from the continuous family of gates presented in Secs. II and III if the underlying quantum circuit is either directly built from, or transpiled to, the hardware native  $R_{ZX}(\theta)$  gate. We now show how to transpile quantum circuits to a  $R_{ZX}(\theta)$ -based circuit with template substitution [56].

A template is a quantum circuit made of  $|T|$  gates acting on  $n_T$  qubits that compose to the identity  $U_1 \cdots U_{|T|} = \mathbb{1}$ ; see, e.g., Figs. 5(b) and 5(c). In a template substitution transpilation pass we identify a subset of the gates in the template  $U_a \cdots U_b$  that match those in a given quantum circuit. Next, if a cost of the matched gates is higher than the cost of the unmatched gates in the template we replace  $U_{\text{match}} = U_a \cdots U_b$  with  $U_{\text{match}} = U_{a-1}^\dagger \cdots U_1^\dagger U_{|T|}^\dagger \cdots U_{b+1}^\dagger$ . As cost we use a heuristic that sums the cost of each gate defined as an integer weight which is higher for two-qubit gates; details

are provided in Appendix A. The complexity of the template matching algorithm on a circuit with  $|C|$  gates and  $n_C$  qubits is

$$O(|C|^{|T|+3} |T|^{|T|+4} n_C^{n_T-1}), \quad (2)$$

i.e., exponential in the template length [56]. We therefore create short templates where the inverse of the intended match, i.e.,  $U_{\text{match}}^\dagger$ , is specified as a single gate with rules to further decompose it into  $R_{ZX}$  and single-qubit gates in a subsequent transpilation pass. In these decompositions we expose the echoed cross-resonance implementation of  $R_{ZX}$  to the transpiler by writing  $R_{ZX}(\theta) = X R_{ZX}(-\theta/2) X R_{ZX}(\theta/2)$ . This allows the transpiler to further simplify the single-qubit gates that would otherwise be hidden in the schedules of the two-qubit gates, as exemplified in the circuit in Fig. 5(e). Finally, once the  $R_{ZX}(\theta)$  gates are introduced into the quantum circuit we run a third transpilation pass to attach pulse schedules to each  $R_{ZX}(\theta)$  gate built from the backend’s calibrated CNOT gates following the procedure in Sec. II. The attached schedules consist of the scaled cross-resonance pulse and rotary tone without any echo. Details on the Qiskit implementation are given in Appendix A.

### V. IMPROVING QAOA WITH CARTAN’S DECOMPOSITION

We use the QAOA [4,57,58], applied to MAXCUT, to demonstrate gains of a pulse-efficient circuit transpilation on noisy hardware. QAOA maps a quadratic binary optimization problem with  $n$  decision variables to a cost function Hamiltonian  $\hat{H}_C = \sum_{i,j} \alpha_{i,j} Z_i Z_j$  where  $\alpha_{ij} \in \mathbb{R}$  are problem dependent and  $Z_i$  are Pauli  $Z$  operators. The ground state of  $\hat{H}_C$  encodes the solution to the problem. Next, a classical solver minimizes the energy  $\langle \psi(\boldsymbol{\beta}, \boldsymbol{\gamma}) | \hat{H}_C | \psi(\boldsymbol{\beta}, \boldsymbol{\gamma}) \rangle$  of a trial state  $|\psi(\boldsymbol{\beta}, \boldsymbol{\gamma})\rangle$  created by applying  $p$ -layers of the operator  $\exp(-i\beta_k \sum_{i=0}^{j-n} X_j) \exp(-i\gamma_k \hat{H}_C)$  where  $k = 1, \dots, p$  to the equal superposition of all states.

Implementing the operator  $\exp(-i\gamma_k \hat{H}_C)$  requires applying the  $\exp(-i\gamma_k \alpha_{i,j} ZZ)$  gate on pairs of qubits. However, to overcome the limited connectivity of superconducting qubit chips [59], several  $R_{ZZ}$  gates are followed or preceded by a

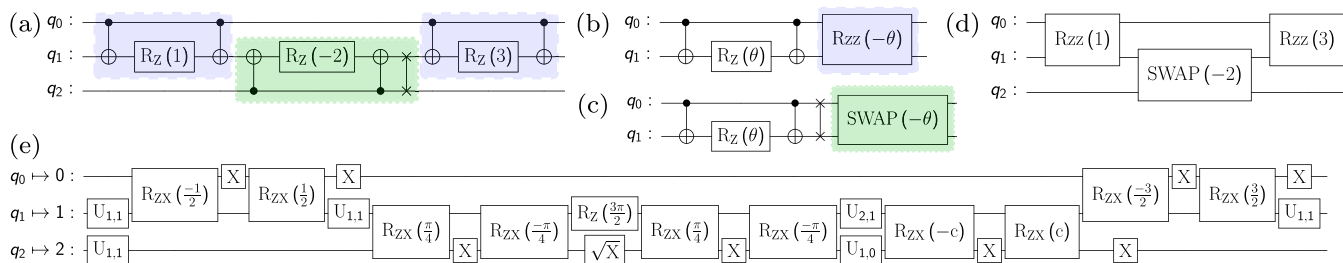


FIG. 5. Pulse-efficient transpilation example. (a) Circuit of the cost operator for a QAOA circuit implemented on three qubits connected in a line. (b), (c) Templates of the  $R_{ZZ}$  and phase-swap gates, respectively. Here  $R_{ZZ}(\theta)$  and  $\text{SWAP}(\theta)$  hold the rules with which to decompose them into the hardware-native  $R_{ZX}$  gates. (d) Circuit resulting from the template matching of (b) and (c) performed on circuit (a). (e) Circuit resulting from a transpilation of (d), which uses the decomposition rules of  $R_{ZZ}(\theta)$  and  $\text{SWAP}(\theta)$  into  $R_{ZX}$ . To shorten the circuit figure we replaced  $R_Z(n\pi/2)\sqrt{X}R_Z(m\pi/2)$  with  $U_{n,m}$  and  $c = 0.215$ .




FIG. 6. Depth-one QAOA energy landscape. (a) Noiseless simulation of the cut value, averaged over all 4096 bit-strings sampled from  $|\psi(\beta, \gamma)\rangle$ , obtained using the QASM simulator for the weighted graph shown in (b). The maximum cut, with value 28, is indicated by the color of the nodes in (b). Panels (c) and (e) show hardware results obtained by transpiling to CNOT gates and by using the  $R_{ZX}$  pulse-efficient methodology, respectively. Panels (d) and (f) share the same color scale and show the absolute deviation from the ideal averaged cut values in panels (c) and (e), respectively.

SWAP resulting in the unitary operator

$$\text{SWAP}(\theta) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & e^{i\theta} & 0 \\ 0 & e^{i\theta} & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad (3)$$

up to a global phase. When mapped to the  $KAK$  decomposition  $\text{SWAP}(\theta)$  corresponds to  $\mathbf{k}^T = (\eta\pi/2, \eta\pi/2, \theta + \eta\pi/2)$  where  $\eta = -1$  if  $\theta > 0$  and 1 otherwise. This allows us to reduce the total cross-resonance duration using the methodology presented in Sec. III.

We perform a depth-one QAOA circuit for an 11-node graph, shown in Fig. 6(b), built from CNOT gates. We map the decision variables zero to 10 to qubits 7, 10, 12, 15, 18, 13, 8, 11, 14, 16, 19 on *ibmq\_mumbai*, respectively. Since the graph is non-hardware-native eight SWAP gates are needed to implement the circuits. In QAOA the optimal values of  $(\beta, \gamma)$  are found with a classical optimizer [60]. Here we scan  $\beta$  and  $\gamma$  from  $\pm 2$  rad and  $\pm 1$  rad, respectively, as we submit jobs through the queue of the cloud-based quantum computers. For each  $(\beta, \gamma)$  pair we run the circuits with the noiseless QASM simulator in Qiskit [see Fig. 6(a)] and twice on the hardware. The first hardware run is done using a CNOT decomposition with the Qiskit transpiler on optimization level three; see Fig. 6(c) for results. The second run is done with the pulse-efficient circuit transpilation; see Fig. 6(e) for results. Here we first perform the template substitution with the

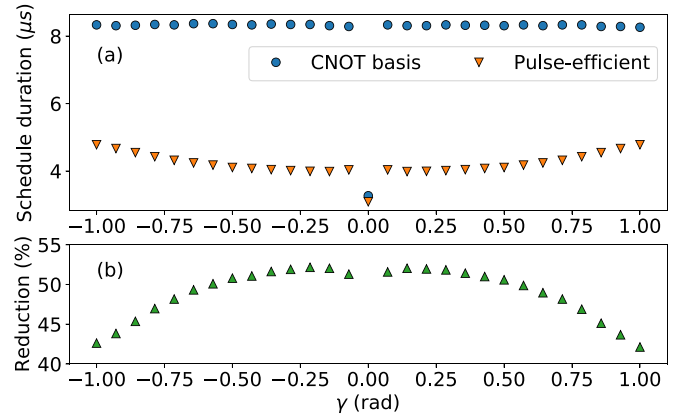


FIG. 7. QAOA schedule durations. (a) Duration of the scheduled quantum circuits transpiled to CNOTs with optimization level three (blue circles) and with the pulse-efficient methodology (orange triangles). In both cases we removed the final measurements from the quantum circuits. The angle  $\gamma = 0$  corresponds to the trivial operation and therefore results in the same schedule duration for both methods. (b) Length of the pulse efficient schedules relative to the CNOT-based schedules.

$R_{ZZ}(\theta)$  and  $\text{SWAP}(\theta)$  templates, shown in Figs. 5(b) and 5(c); see Appendix A for further details. A second transpilation pass then exposes the  $R_{ZX}(\theta)$  gates to which we attach pulse schedules in a third transpilation pass following Secs. II–IV. In each case we measure 4096 shots. The pulse-efficient circuits produce less noisy average cut values [compare Fig. 6(c) with Fig. 6(e)] and have a lower absolute deviation from the noiseless simulation than the circuits transpiled to CNOT gates; compare Fig. 6(d) with Fig. 6(f). The maximum error in the cut value averaged over the sampled bit strings is reduced by 38% from 3.65 to 2.26. We attribute the increased quality of the results to the decrease in total cross-resonance time and the fact that the pulse-efficient transpilation keeps the number of single-qubit pulses to a minimum. In total, we observe a reduction in total schedule duration ranging from 42% to 52% depending on  $\gamma$  when using the pulse efficient transpilation methodology; see Fig. 7. Since the schedule duration of  $R_{ZZ}(\gamma\alpha_{i,j})$  and  $\text{SWAP}(\gamma\alpha_{i,j})$  decreases and increases as  $\gamma$  decreases, respectively, we observe a nonmonotonous reduction in the schedule duration of the QAOA circuit as a function of  $\gamma$ .

## VI. DISCUSSION AND CONCLUSION

The results in Secs. II and III showed that by scaling cross-resonance gates we can automatically create a continuous family of gates which implements  $\text{SU}(4)$ . These scaled gates typically have shorter pulse schedules and higher fidelities than the digital CNOT implementation. This fidelity is limited by coherence, imperfections in the initial calibration, and nonlinear effects showing that it is easier to engineer pulses for systems with a linear response. Crucially, the resulting gate-tailored pulse schedules do not require additional calibration experiments and can therefore be automatically generated by the transpiler. In the future, if quantum systems become stable enough and models accurate enough so that

closed-loop optimal control is not necessary [26,31], it may be possible to embed open-loop optimal control [24,25,61] in transpiler passes such as template matching. Transpilation passes, as discussed in Sec. IV, can be leveraged to identify and attach the scaled pulse schedules to the gates in a quantum circuit. Here the template-matching algorithm of Ref. [56] can accommodate multiqubit gates, such as the iToffoli gate [62], but can be computationally costly. Alternatively, transpiler passes that decompose a series of two-qubit gates into  $R_{XX}(\alpha)$ ,  $R_{YY}(\beta)$ , and  $R_{ZZ}(\gamma)$  may run faster but are limited to hardware with two-qubit interactions. Furthermore, exposing the echo in the cross-resonance gate to the transpiler allows further simplifications of the single-qubit gates. We used this pulse-efficient transpilation methodology to reduce errors in an 11-qubit depth-one QAOA.

Scaled gates are particularly appealing for Trotter-based applications, as shown in Ref. [43], and could therefore benefit quantum simulations [63]. Future work may also include scaling direct cross-resonance gates [9] and benchmarking their impact on quantum volume [8]. Methods to interpolate pulse parameters based on a set of reference  $R_{ZX}(\theta)$  gates, calibrated at a few reference angles  $\theta$ , might also improve the gate fidelity and help deal with nonlinearities between the rotation angle  $\theta$  and pulse parameters. For variational algorithms, such as the variational quantum eigensolver, the scaled SU(4) gates may allow for better results due to the shorter schedules while still being robust to some unitary errors such as angle errors [64,65].

We believe that the methods presented in our work will help users of noisy quantum hardware to reap the benefits of pulse-level control without having to know its intricacies. This can improve the quality of a broad class of quantum applications running on noisy quantum hardware.

## ACKNOWLEDGMENTS

The authors acknowledge use of the IBM Quantum devices for this work. The authors also thank L. Capelluto, N. Kanazawa, N. Bronn, T. Itoko, and E. Pritchett for insightful discussions and S. Woerner for a careful reading of the manuscript.

## APPENDIX A: QISKIT IMPLEMENTATION

Quantum circuits often have repeating subcircuits with different parameters. For instance, QAOA circuits include many  $R_{ZZ}(\gamma\alpha_{ij})$  and SWAP( $\gamma\alpha_{ij}$ ) gates where  $\gamma$  is one of the variational parameters and the  $\{\alpha_{ij}\}$  depend on the problem instance. We therefore need parametric templates when running the template substitution algorithm.

We extended the Qiskit implementation of Ref. [56] to parametric templates. To avoid a symbolic description of the unitary matrix of each gate we first match gates by qubits and name. This is, however, not sufficient to create a valid match since, for example, the parametric template in Fig. 8(a) produces two tentative matches on the circuit in Fig. 8(b). We therefore form a system of equations based on the tentative match. If this system of equations accepts a solution the match is valid. For example, the tentative match in Fig. 8(b), indicated by the dashed blue box, results in the system of




FIG. 8. (a) Parametric template of a controlled-Z gate. (b) Circuit on which the template matching is run. The dashed blue and dotted purple boxes indicate potential matches based on circuit instruction names and qubits.

equations

$$\begin{cases} -\theta = -2 \\ -\theta = -2 \\ \theta = 2 \end{cases}, \quad (\text{A1})$$

which accepts the solution  $\theta = 2$  and is therefore valid. However, the second tentative match, highlighted by the dotted purple box, results in the system of equations

$$\begin{cases} -\theta = 3 \\ -\theta = 3 \\ \theta = 3 \end{cases}, \quad (\text{A2})$$

which has no solution and is therefore not valid.

We achieve a pulse-efficient circuit transpilation with Qiskit by using three transpilation steps shown in Fig. 9. First, the TemplateOptimization transpilation pass is applied with the SWAP and rzz templates as shown in Figs. 5(b) and 5(c). The next step, a standard transpiler pass with a low optimization level, i.e. one, exposes the rzx definition of the gates in the matched templates. Finally, the RZXCalibrationBuilderNoEcho class scales the pulses of the cross-resonance gates and attaches them to the  $R_{ZX}(\theta)$

```
# 1. Template substitution: introduces custom
# gates with an rzx decomposition
pass_ = TemplateOptimization(
    template_list=[swap(), rzz()],
    user_cost_dict=cost_dict)

qct = PassManager(pass_).run(qc)

# 2. Unroll to rzx basis and simplify gates
gates = ['rz', 'sx', 'rzx', 'cx', 'x']
qct_rzx = transpile(qct,
                    backend,
                    optimization_level=1,
                    basis_gates=gates)

# 3. Add custom rzx pulse schedules
pass_ = RZXCalibrationBuilderNoEcho(backend)
cal_qc = PassManager(pass_).run(qct_rzx)
```

FIG. 9. Example of a circuit transpilation that achieves a pulse-efficient circuit transpilation.

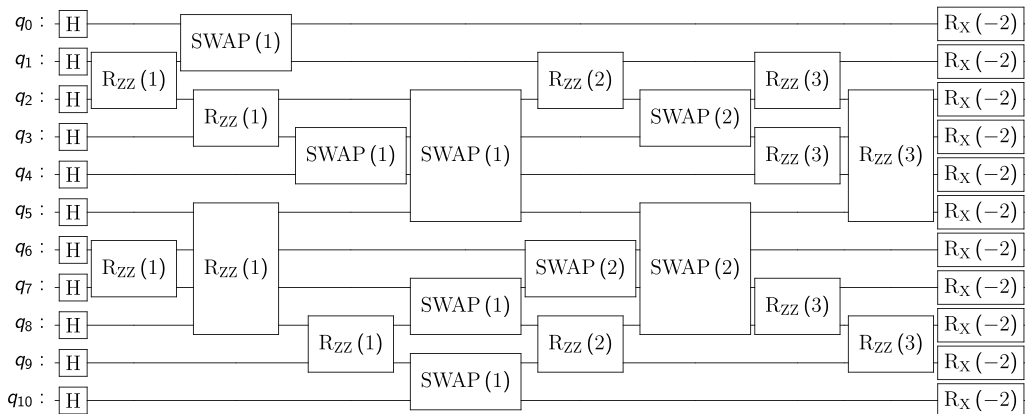


FIG. 10. 11 qubit QAOA circuit with  $\gamma = 1$  and  $\beta = -2$  after the template substitution. The final measurement instructions have been omitted.

gates in the circuit. Figure 10 exemplifies the result of the first transpilation pass applied to the QAOA circuit in Sec. V.

The template optimization pass requires a cost dictionary to determine if it is favorable to replace the matched gates  $U_{\text{match}} = U_a \cdots U_b$  from a template with the Hermitian conjugate of the remaining part of the template  $U_{a-1}^\dagger \cdots U_1^\dagger U_{|T|}^\dagger \cdots U_{b+1}^\dagger$ . The cost dictionary has gates as keys and their cost as value. The cost of  $U_a \cdots U_b$  is the sum of the costs of each individual gate  $U_a$  to  $U_b$ . We used the cost dictionary  $\{ \text{'sx': 1, 'x': 1, 'rz': 0, 'cx': 2, 'rzz': 0, 'swap': 6, 'phase\_swap': 0} \}$  which assigns a zero cost to the rzz and phase\_swap gates, which correspond to the pulse-efficient implementation of  $R_{ZZ}(\theta)$  and  $\text{SWAP}(\theta)$ . Single-qubit gates have unit cost except for rz, which is implemented with virtual Z-rotations. The CNOT gate, i.e., cx, and the standard SWAP gate, i.e., swap have costs two and six, respectively. This cost dictionary ensures that the template substitution will include  $R_{ZZ}(\theta)$  and  $\text{SWAP}(\theta)$  in the case of a match. Future work could improve this heuristic cost dictionary either by using the fidelity of the gates (if this metric is available) or the duration of the underlying pulse schedules as cost.

**APPENDIX B: PROPERTIES OF THE QUANTUM DEVICES AND ADDITIONAL DATA**

Since the qubit coherence times as well as the CNOT gate duration and error mainly limit the fidelity of the scaled cross-resonance gates, we list their values for the qubits and devices we experimented with in Table I. To illustrate that scaling imperfect cross-resonance gates improves the gate fidelity we measured the process fidelity on several IBM Quantum devices and qubit pairs. In almost all measurements the scaled gates have a higher fidelity than the double CNOT benchmark, and the relative error reduction increases as the schedule duration decreases; see Fig. 11.

Figure 12 shows additional quantum process tomography results for Cartan-decomposed circuits chosen at random in the Weyl chamber. The experiments were performed on *ibmq\_dublin* [see Fig. 12(a)] and *ibmq\_paris* using different qubit pairs [see Figs. 12(b)–12(d)]. For almost all angles the relative error reduction is positive which demonstrates the

advantage of a hardware-native, scaled cross-resonance gate-based circuit implementation.

**APPENDIX C: THEORETICAL COHERENCE LIMIT**

In Sec. II we compared the relative decrease in gate error to the coherence limit on the average gate error  $\mathcal{E}$ . This limit is implemented in Qiskit Ignis for two qubits  $a$  and  $b$  as  $\mathcal{E} = \frac{3}{4}(1 - u_1 - u_2)$  where

$$u_1 = \frac{1}{15}(e^{-t/T_{1,a}} + e^{-t/T_{1,b}} + e^{-t/T_{1,a}-t/T_{1,b}}), \tag{C1}$$

$$u_2 = \frac{2}{15}(e^{-t/T_{2,b}} + e^{-t/T_{2,b}-t/T_{1,a}} + e^{-t/T_{2,a}} + e^{-t/T_{2,a}-t/T_{1,b}} + 2e^{-t/T_{2,a}-t/T_{2,b}}). \tag{C2}$$

The derivation of this limit is discussed in more detail in Appendix G of Ref. [37]. Here  $t$  is the gate duration, while  $T_{1,a}$  and  $T_{2,a}$  are the  $T_1$  and  $T_2$  times for qubit  $a$ . Since the

TABLE I. Summary of the properties of the CNOT gates and coherence times for the qubits used to benchmark the performance of the scaled cross-resonance gates.

Device	CNOT			
	Error (%)	Duration (ns)	$T_1$ times ( $\mu\text{s}$ )	$T_2$ times ( $\mu\text{s}$ )
<i>ibmq_mumbai</i>				
q1, q2	1.27	739	102, 157	34, 228
q16, q19	0.84	754	84, 141	105, 132
<i>ibmq_paris</i>				
q1, q2	1.70	597	66, 92	82, 128
q13, q14	1.28	434	100, 23	27, 33
q18, q15	5.36	448	86, 74	103, 50
q18, q17	1.76	725	41, 71	94, 157
<i>ibmq_dublin</i>				
q1, q2	0.76	540	110, 103	174, 89
q3, q2	0.83	370	78, 96	100, 83
<i>ibmq_montreal</i>				
q14, q16	0.88	356	97, 87	97, 52
<i>ibmq_guadalupe</i>				
q7, q10	0.61	299	99, 68	153, 90

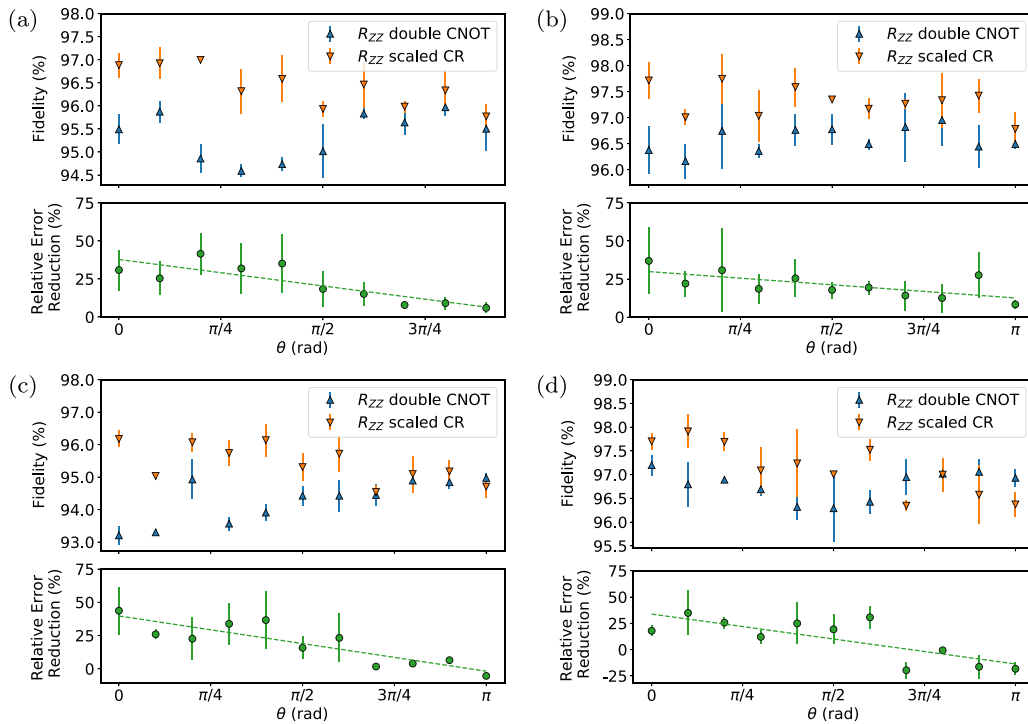


FIG. 11. Gate fidelity measured with quantum process tomography and relative error reduction for the ZZ-gate as a function of  $\theta$  on *ibmq\_mumbai* (qubits 16 and 19), *ibmq\_montreal* (qubits 14 and 16), *ibmq\_paris* (qubits 1 and 2), and *ibmq\_guadalupe* (qubits 7 and 10). (a)–(d), (top) Process fidelities for the  $R_{ZZ}$  double CNOT (blue up-triangles) and scaled CR (orange down-triangles) circuit implementation. (a)–(d), (bottom) Relative error reduction calculated from the fidelities.

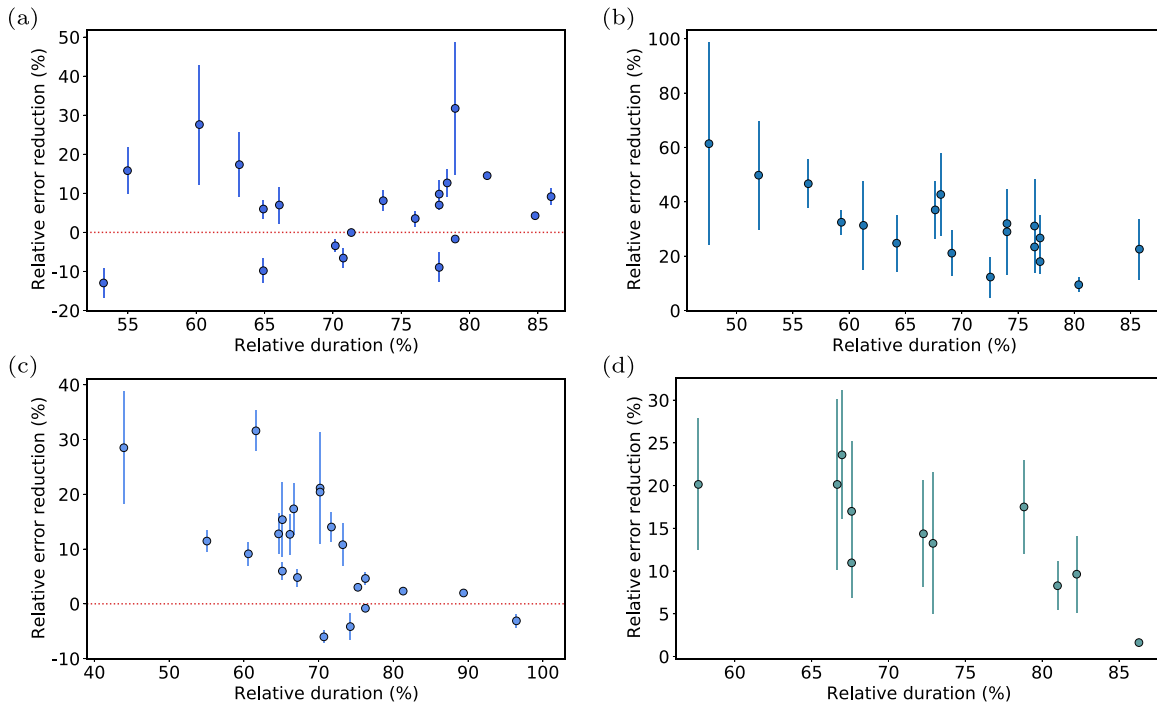


FIG. 12. Quantum process tomography results for random angles in the Weyl chamber on (a) *ibmq\_dublin* (qubits 3 and 2), (b) *ibmq\_paris* (qubits 18 and 15), (c) *ibmq\_paris* (qubits 13 and 14), and (d) *ibmq\_paris* (qubits 18 and 17).



process fidelity and the average gate fidelity are linearly related [47,48], we compare the relative error reduction in the

measured process fidelity with the theoretical relative error reduction in  $\mathcal{E}$ .

- 
- [1] N. Moll, P. Barkoutsos, L. S. Bishop, J. M. Chow, A. Cross, D. J. Egger, S. Filipp, A. Fuhrer, J. M. Gambetta, M. Ganzhorn *et al.*, Quantum optimization using variational algorithms on near-term quantum devices, *Quantum Sci. Technol.* **3**, 030503 (2018).
- [2] R. Orús, S. Mugel, and E. Lizaso, Quantum computing for finance: Overview and prospects, *Rev. Phys.* **4**, 100028 (2019).
- [3] 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).
- [4] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, [arXiv:1411.4028](https://arxiv.org/abs/1411.4028) [quant-ph].
- [5] D. J. Egger, J. Mareček, and S. Woerner, Warm-starting quantum optimization, *Quantum* **5**, 479 (2021).
- [6] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, Quantum machine learning, *Nature (London)* **549**, 195 (2017).
- [7] V. Havlicek, A. D. Corcoles, K. Temme, A. W. Harrow, A. Kandala, J. M. Chow, and J. M. Gambetta, Supervised learning with quantum-enhanced feature spaces, *Nature (London)* **567**, 209 (2019).
- [8] 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* **100**, 032328 (2019).
- [9] 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).
- [10] P. Krantz, M. Kjaergaard, F. Yan, T. P. Orlando, S. Gustavsson, and W. D. Oliver, A quantum engineer's guide to superconducting qubits, *Appl. Phys. Rev.* **6**, 021318 (2019).
- [11] M. Kjaergaard, M. E. Schwartz, J. Braumüller, P. Krantz, J. I.-J. Wang, S. Gustavsson, and W. D. Oliver, Superconducting qubits: Current state of play, *Annu. Rev. Condens. Matter Phys.* **11**, 369 (2020).
- [12] J. Koch, T. M. Yu, J. M. Gambetta, A. A. Houck, D. I. Schuster, J. Majer, A. Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, Charge-insensitive qubit design derived from the Cooper pair box, *Phys. Rev. A* **76**, 042319 (2007).
- [13] C. Rigetti, J. M. Gambetta, S. Poletto, B. L. T. Plourde, J. M. Chow, A. D. Corcoles, J. A. Smolin, S. T. Merkel, J. R. Rozen, G. A. Keefe *et al.*, Superconducting qubit in a waveguide cavity with a coherence time approaching 0.1 ms, *Phys. Rev. B* **86**, 100506(R) (2012).
- [14] IBM Quantum, <https://quantum-computing.ibm.com/> (2021).
- [15] S. Bravyi, S. Sheldon, A. Kandala, D. C. McKay, and J. M. Gambetta, Mitigating measurement errors in multiqubit experiments, *Phys. Rev. A* **103**, 042605 (2021).
- [16] G. S. Barron and C. J. Wood, Measurement error mitigation for variational quantum algorithms, [arXiv:2010.08520](https://arxiv.org/abs/2010.08520) [quant-ph].
- [17] K. Temme, S. Bravyi, and J. M. Gambetta, Error Mitigation for Short-Depth Quantum Circuits, *Phys. Rev. Lett.* **119**, 180509 (2017).
- [18] A. Kandala, K. Temme, A. D. Corcoles, A. Mezzacapo, J. M. Chow, and J. M. Gambetta, Error mitigation extends the computational reach of a noisy quantum processor, *Nature (London)* **567**, 491 (2018).
- [19] N. Lacroix, C. Hellings, C. K. Andersen, A. Di Paolo, A. Remm, S. Lazar, S. Krinner, G. J. Norris, M. Gabureac, J. Heinsoo *et al.*, Improving the performance of deep quantum optimization algorithms with continuous gate sets, *PRX Quantum* **1**, 110304 (2020).
- [20] B. Foxen, C. Neill, A. Dunsworth, P. Roushan, B. Chiaro, A. Megrant, J. Kelly, Z. Chen, K. J. Satzinger, R. Barends *et al.* (Google AI Quantum), Demonstrating a Continuous Set of Two-Qubit Gates for Near-Term Quantum Algorithms, *Phys. Rev. Lett.* **125**, 120504 (2020).
- [21] P. Gokhale, A. Javadi-Abhari, N. Earnest, Y. Shi, and F. T. Chong, Optimized quantum compilation for near-term algorithms with openpulse, in *2020 53rd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)* (IEEE, New York, 2020), pp. 186–200.
- [22] M. C. Collodo, J. Herrmann, N. Lacroix, C. K. Andersen, A. Remm, S. Lazar, J.-C. Besse, T. Walter, A. Wallraff, and C. Eichler, Implementation of Conditional Phase Gates Based on Tunable ZZ interactions, *Phys. Rev. Lett.* **125**, 240502 (2020).
- [23] A. D. Hill, M. J. Hodson, N. Didier, and M. J. Reagor, Realization of arbitrary doubly-controlled quantum phase gates, [arXiv:2108.01652](https://arxiv.org/abs/2108.01652) [quant-ph].
- [24] 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).
- [25] 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 Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*, ASPLOS '19 (Association for Computing Machinery, New York, 2019), pp. 1031–1044.
- [26] D. J. Egger and F. K. Wilhelm, Adaptive Hybrid Optimal Quantum Control for Imprecisely Characterized Systems, *Phys. Rev. Lett.* **112**, 240503 (2014).
- [27] N. Wittler, F. Roy, K. Pack, M. Werninghaus, A. S. Roy, D. J. Egger, S. Filipp, F. K. Wilhelm, and S. Machnes, Integrated Tool Set for Control, Calibration, and Characterization of Quantum Devices Applied to Superconducting Qubits, *Phys. Rev. Appl.* **15**, 034080 (2021).
- [28] J. Kelly, R. Barends, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. G. Fowler, I.-C. Hoi, E. Jeffrey *et al.*, Optimal Quantum Control Using Randomized Benchmarking, *Phys. Rev. Lett.* **112**, 240504 (2014).
- [29] M. Werninghaus, D. J. Egger, F. Roy, S. Machnes, F. K. Wilhelm, and S. Filipp, Leakage reduction in fast

- superconducting qubit gates via optimal control, *npj Quantum Inf.* **7**, 14 (2021).
- [30] A. M. Rol, C. C. Bultink, T. E. O'Brien, S. R. de Jong, L. S. Theis, X. Fu, F. Luthi, R. F. L. Vermeulen, J. C. de Sterke, A. Bruno *et al.*, Restless Tuneup of High-Fidelity Qubit Gates, *Phys. Rev. Appl.* **7**, 041001(R) (2017).
- [31] M. Werninghaus, D. J. Egger, and S. Filipp, High-speed calibration and characterization of superconducting quantum processors without qubit reset, *PRX Quantum* **2**, 020324 (2021).
- [32] S. Machnes, E. Assémat, D. Tannor, and F. K. Wilhelm, Tunable, Flexible, and Efficient Optimization of Control Pulses for Practical Qubits, *Phys. Rev. Lett.* **120**, 150401 (2018).
- [33] Y. Baum, M. Amico, S. Howell, M. Hush, M. Liuzzi, P. Mundada, T. Merkh, A. R. R. Carvalho, and M. J. Biercuk, Experimental deep reinforcement learning for error-robust gateset design on a superconducting quantum computer, [arXiv:2105.01079](https://arxiv.org/abs/2105.01079) [quant-ph].
- [34] A. R. R. Carvalho, H. Ball, M. J. Biercuk, M. R. Hush, and F. Thomsen, Error-Robust Quantum Logic Optimization Using a Cloud Quantum Computer Interface, *Phys. Rev. Appl.* **15**, 064054 (2021).
- [35] J. M. Chow, A. D. Córcoles, J. M. Gambetta, C. Rigetti, B. R. Johnson, J. A. Smolin, J. R. Rozen, G. A. Keefe, M. B. Rothwell, M. B. Ketchen, and M. Steffen, Simple All-Microwave Entangling Gate for Fixed-Frequency Superconducting Qubits, *Phys. Rev. Lett.* **107**, 080502 (2011).
- [36] S. Sheldon, E. Magesan, J. M. Chow, and J. M. Gambetta, Procedure for systematically tuning up cross-talk in the cross-resonance gate, *Phys. Rev. A* **93**, 060302(R) (2016).
- [37] N. Sundaresan, I. Lauer, E. Pritchett, E. Magesan, P. Jurcevic, and J. M. Gambetta, Reducing unitary and spectator errors in cross resonance with optimized rotary echoes, *PRX Quantum* **1**, 020318 (2020).
- [38] D. C. McKay, T. Alexander, L. Bello, M. J. Biercuk, L. Bishop, J. Chen, J. M. Chow, A. D. Córcoles, D. J. Egger, S. Filipp *et al.*, Qiskit backend specifications for OpenQASM and OpenPulse experiments, [arXiv:1809.03452](https://arxiv.org/abs/1809.03452) [quant-ph].
- [39] T. Alexander, N. Kanazawa, D. J. Egger, L. Capelluto, C. J. Wood, A. Javadi-Abhari, and D. C. McKay, Qiskit pulse: Programming quantum computers through the cloud with pulses, *Quantum Sci. Technol.* **5**, 044006 (2020).
- [40] S. Garion, N. Kanazawa, H. Landa, D. C. McKay, S. Sheldon, A. W. Cross, and C. J. Wood, Experimental implementation of non-Clifford interleaved randomized benchmarking with a controlled- $S$  gate, *Phys. Rev. Res.* **3**, 013204 (2021).
- [41] S. Oomura, T. Satoh, M. Sugawara, and N. Yamamoto, Design and application of high-speed and high-precision CV gate on IBM Q OpenPulse system, [arXiv:2102.06117](https://arxiv.org/abs/2102.06117) [quant-ph].
- [42] K. Heya and N. Kanazawa, Cross cross resonance gate, [arXiv:2103.00024](https://arxiv.org/abs/2103.00024) [quant-ph].
- [43] J. P. T. Stenger, N. T. Bronn, D. J. Egger, and D. Pekker, Simulating the dynamics of braiding of Majorana zero modes using an IBM quantum computer, *Phys. Rev. Res.* **3**, 033171 (2021).
- [44] M. Mohseni, A. T. Rezakhani, and D. A. Lidar, Quantum-process tomography: Resource analysis of different strategies, *Phys. Rev. A* **77**, 032322 (2008).
- [45] R. C. Bialczak, M. Ansmann, M. Hofheinz, E. Lucero, M. Neeley, A. D. O'Connell, D. Sank, H. Wang, J. Wenner, M. Steffen *et al.*, Quantum process tomography of a universal entangling gate implemented with Josephson phase qubits, *Nat. Phys.* **6**, 409 (2010).
- [46] E. Magesan and J. M. Gambetta, Effective hamiltonian models of the cross-resonance gate, *Phys. Rev. A* **101**, 052308 (2020).
- [47] M. Horodecki, P. Horodecki, and R. Horodecki, General teleportation channel, singlet fraction, and quasidistillation, *Phys. Rev. A* **60**, 1888 (1999).
- [48] E. Magesan, R. Blume-Kohout, and J. Emerson, Gate fidelity fluctuations and quantum process invariants, *Phys. Rev. A* **84**, 012309 (2011).
- [49] D. C. McKay, C. J. Wood, S. Sheldon, J. M. Chow, and J. M. Gambetta, Efficient  $z$  gates for quantum computing, *Phys. Rev. A* **96**, 022330 (2017).
- [50] N. Khaneja and S. J. Glaser, Cartan decomposition of  $SU(2^n)$  and control of spin systems, *Chem. Phys.* **267**, 11 (2001).
- [51] J. Zhang, J. Vala, S. Sastry, and B. K. Whaley, Geometric theory of nonlocal two-qubit operations, *Phys. Rev. A* **67**, 042313 (2003).
- [52] R. R. Tucci, An introduction to Cartan's KAK decomposition for QC programmers, [arXiv:quant-ph/0507171](https://arxiv.org/abs/quant-ph/0507171).
- [53] B. Drury and P. Love, Constructive quantum Shannon decomposition from Cartan involutions, *J. Phys. A: Math. Theor.* **41**, 395305 (2008).
- [54] G. Vidal and C. M. Dawson, Universal quantum circuit for two-qubit transformations with three controlled-not gates, *Phys. Rev. A* **69**, 010301(R) (2004).
- [55] M. Treinish, J. Gambetta, P. Nation, P. Kassebaum, qiskit-bot, D. M. Rodriguez, S. de la Puente González, S. Hu, K. Krsulich, L. Zdanski *et al.*, Qiskit: An open-source framework for quantum computing (2021), doi: [10.5281/zenodo.2573505](https://doi.org/10.5281/zenodo.2573505).
- [56] R. Iten, R. Moyard, T. Metger, D. Sutter, and S. Woerner, Exact and practical pattern matching for quantum circuit optimization, [arXiv:1909.05270](https://arxiv.org/abs/1909.05270) [quant-ph].
- [57] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm applied to a bounded occurrence constraint problem, [arXiv:1412.6062](https://arxiv.org/abs/1412.6062) [quant-ph].
- [58] Z.-C. Yang, A. Rahmani, A. Shabani, H. Neven, and C. Chamon, Optimizing Variational Quantum Algorithms Using Pontryagin's Minimum Principle, *Phys. Rev. X* **7**, 021027 (2017).
- [59] 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).
- [60] A. Bengtsson, P. Vikstål, C. Warren, M. Svensson, X. Gu, A. F. Kockum, P. Krantz, C. Križan, D. Shiri, I.-M. Svensson *et al.*, Improved Success Probability with Greater Circuit Depth for the Quantum Approximate Optimization Algorithm, *Phys. Rev. Appl.* **14**, 034010 (2020).
- [61] D. J. Egger and F. K. Wilhelm, Optimized controlled- $Z$  gates for two superconducting qubits coupled through a resonator, *Supercond. Sci. Technol.* **27**, 014001 (2013).
- [62] Y. Kim, A. Morvan, L. B. Nguyen, R. K. Naik, C. Jünger, L. Chen, J. M. Kreikebaum, D. I. Santiago, and I. Siddiqi, High-fidelity Toffoli gate for fixed-frequency superconducting qubits, [arXiv:2108.10288](https://arxiv.org/abs/2108.10288) [quant-ph].

- [63] S. Tornow, W. Gehrke, and U. Helmbrecht, Non-equilibrium dynamics of a dissipative two-site Hubbard model simulated on the IBM quantum computer, [arXiv:2011.11059](https://arxiv.org/abs/2011.11059) [quant-ph].
- [64] J. I. Colless, V. V. Ramasesh, D. Dahlen, M. S. Blok, M. E. Kimchi-Schwartz, J. R. McClean, J. Carter, W. A. de Jong, and I. Siddiqi, Computation of Molecular Spectra on a Quantum Processor with an Error-Resilient Algorithm, *Phys. Rev. X* **8**, 011021 (2018).
- [65] D. J. Egger, M. Ganzhorn, G. Salis, A. Fuhrer, P. Müller, P. Kl. Barkoutsos, N. Moll, I. Tavernelli, and S. Filipp, Entanglement Generation in Superconducting Qubits Using Holonomic Operations, *Phys. Rev. Appl.* **11**, 014017 (2019).