

PAPER • OPEN ACCESS

Software mitigation of coherent two-qubit gate errors

To cite this article: Lingling Lao *et al* 2022 *Quantum Sci. Technol.* **7** 025021

View the [article online](#) for updates and enhancements.

You may also like

- [Two-qubit gate operations in superconducting circuits with strong coupling and weak anharmonicity](#)
Xin-You Lü, S Ashhab, Wei Cui et al.
- [Escaping Detrimental Interactions with Microwave-Dressed Transmon Qubits](#)
Z. T. Wang, , Peng Zhao et al.
- [Simulating the performance of a distance-3 surface code in a linear ion trap](#)
Colin J Trout, Muyuan Li, Mauricio Gutiérrez et al.



Easy-to-use and Helium-3 free
cryogenics solutions



LEARN MORE

Quantum Science and Technology



PAPER

Software mitigation of coherent two-qubit gate errors

OPEN ACCESS

RECEIVED
25 November 2021

REVISED
11 February 2022

ACCEPTED FOR PUBLICATION
23 February 2022

PUBLISHED
15 March 2022

Original content from
this work may be used
under the terms of the
[Creative Commons
Attribution 4.0 licence](#).

Any further distribution
of this work must
maintain attribution to
the author(s) and the
title of the work, journal
citation and DOI.



Lingling Lao^{1,*} , Alexander Korotkov², Zhang Jiang², Wojciech Mruczkiewicz² ,
Thomas E O'Brien² and Dan E Browne¹

¹ Department of Physics and Astronomy, University College London, United Kingdom

² Google AI Quantum, United States of America

* Author to whom any correspondence should be addressed.

E-mail: laolinglingrolls@gmail.com

Keywords: quantum error mitigation, quantum compilation, gate decomposition, coherent errors

Abstract

Two-qubit gates are important components of quantum computing. However, unwanted interactions between qubits (so-called parasitic gates) can be particularly problematic and degrade the performance of quantum applications. In this work, we present two software methods to mitigate parasitic two-qubit gate errors. The first approach is built upon the Cartan's KAK decomposition and keeps the original unitary decomposition for the error-free native two-qubit gate. It counteracts a parasitic two-qubit gate by only applying single-qubit rotations and therefore has no two-qubit gate overhead. We show the optimal choice of single-qubit mitigation gates. The second approach applies a numerical optimisation algorithm to re-compile a target unitary into the error-parasitic two-qubit gate plus single-qubit gates. We demonstrate these approaches on the CPhase-parasitic iSWAP-like gates. The KAK-based approach helps decrease unitary infidelity by a factor of 3 compared to the noisy implementation without error mitigation. When arbitrary single-qubit rotations are allowed, recompilation could completely mitigate the effect of parasitic errors but may require more native gates than the KAK-based approach. We also compare their average gate fidelity under realistic noise models, including relaxation and depolarising errors. Numerical results suggest that different approaches are advantageous in different error regimes, providing error mitigation guidance for near-term quantum computers.

1. Introduction

Quantum computers can tackle problems that are intractable by classical computers. A key challenge in quantum computing is to implement high-fidelity building blocks, including single-qubit and two-qubit gates, qubit initialisation and readout. Two-qubit gates generate entanglement and are particularly important. Although tremendous progress has been made, the error rates of two-qubit gates remain high and limit the performance of quantum computers [1–6].

In superconducting circuits, two-qubit gates can be realised by resonantly coupling two two-qubit states. We consider the iSWAP-like gate family ($i\text{SWAP}(\theta)$), which can be realised by tuning the states $|01\rangle$ and $|10\rangle$ into resonance [1, 2]. During an iSWAP-like gate, the repulsion of state $|11\rangle$ from states $|02\rangle$ and $|20\rangle$ causes an extra phase for state $|11\rangle$, so that the actual gate implemented on hardware has an unwanted CPhase component (i.e., $C\text{Phase}(\psi)i\text{SWAP}(\theta)$). We refer to an unwanted two-qubit interaction on a hardware two-qubit gate as a *parasitic gate* error. In [2], parasitic $C\text{Phase}(\psi)$ errors are shown to be associated in general with the implementation of $i\text{SWAP}(\theta)$ on this hardware, with $\psi \propto \theta^2$ for a fixed-duration gate. The parasitic gate errors, if left unmitigated, can have a negative effect on application performance. These errors can be substantially suppressed on *hardware* by increasing the gate duration [1, 2]. However, longer gate implementation may introduce more noises because of the limited coherence time.

Besides the parasitic CPhase errors, iSWAP-like gates also have other coherent errors such as single-qubit Z rotations and offsets on iSWAP angles. A single-qubit Z rotation error can be cancelled out by simply applying its Hermitian conjugation. The iSWAP offset angles are normally small [7] and will be considered

in future work. This work focuses on the parasitic CPhase errors since they are more detrimental in current quantum devices [2, 5, 7]. All these error parameters can be characterised by gate calibration tools such as randomised benchmarking [8], gate set tomography [9], and cross-entropy benchmarking [10]. If these parameters drift and fluctuate quickly, one can use a fast Floquet calibration to learn their real-time values [7].

In this work, we demonstrate two *software* approaches for mitigating the effect of parasitic CPhase errors. The first relies on approximating the parasitic CPhase gate via single-qubit Z rotations. It reduces unitary infidelity by a factor of 3 and has been experimentally demonstrated in [5]. We derive this approach from the KAK decomposition [11–13] and generalise it for an arbitrary target two-qubit gate and arbitrary parasitic errors. We show that the optimal single-qubit mitigation gates depend only on the parasitic gate.

In the second, we do not attempt to correct the parasitic gate at all, but treat the whole hardware gate, $i\text{SWAP}(\theta)\text{CPhase}(\psi)$ as the native gate for the computation. We use a numerical decomposition approach to recompile target unitary gates directly into a gate set consisting of arbitrary single-qubit gates and the native two-qubit gate. The recompilation approach can give a decomposition with perfect fidelity and therefore completely mitigate the effect of parasitic CPhase errors. However, it may require more native gates than the KAK approximation for some target unitaries and the final unitary fidelity could be decreased by other hardware errors such as qubit relaxation. We compare these methods by implementing arbitrary $\text{SU}(4)$ gates and a set of excitation number-preserving two-qubit gates in different strength of parasitic errors and extra hardware errors including relaxation and depolarising errors. Our evaluation results provide suggestions on how to choose the best mitigation approach for a target unitary under realistic noise models.

This paper is organised as follows. We first introduce the background information in section 2. Then we present the error mitigation approach based on KAK approximation in section 3 and the recompilation approach in section 4. We compare different methods with other hardware errors in section 5 and conclude the paper in section 6.

2. Background


In this section, we introduce the background on two-qubit gates and unitary fidelity. Both the $i\text{SWAP}$ -like gate family and CPhase gate family are excitation number-preserving gates and can allow short-depth circuit implementation for quantum simulation [14] and the quantum approximate optimisation algorithm [15]. The matrix representation of $i\text{SWAP}(\theta)$ is defined as

$$i\text{SWAP}(\theta) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\theta) & -i \sin(\theta) & 0 \\ 0 & -i \sin(\theta) & \cos(\theta) & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}.$$

The matrix representation of $\text{CPhase}(\phi)$ is

$$\text{CPhase}(\phi) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & e^{-i\phi} \end{pmatrix}.$$

In principle, one can realise a continuous set of $i\text{SWAP}$ -like gates [2] or CPhase gates [16] to minimise circuit depth. However, it is challenging to calibrate and benchmark a continuous gate set on multiple qubits. Current quantum processors typically calibrate one two-qubit gate for high-fidelity implementation. For example, the sole native two-qubit gate in [7] is $\sqrt{i\text{SWAP}}^\dagger = i\text{SWAP}(\pi/4)$. The $\sqrt{i\text{SWAP}}^\dagger$ gate has powerful capabilities to express other two-qubit gates. It has been proven that any two-qubit gate can be expressed by at most three $\sqrt{i\text{SWAP}}^\dagger$ gates [17]. For example, a general $i\text{SWAP}(\theta)$ unitary can be implemented using six single-qubit Z rotations ($R_Z(\theta) = \exp(-i\theta Z/2)$) and two $\sqrt{i\text{SWAP}}^\dagger$ gates [7] as shown in figure 1. For some special angles such as $i\text{SWAP}(-\pi/4)$ and $i\text{SWAP}(\pm 3\pi/4)$, one can decompose these unitaries with only one $\sqrt{i\text{SWAP}}^\dagger$ gate. Unless otherwise stated, we will use the decomposition in figure 1 throughout the paper.



To evaluate the performance of error mitigation approaches, we calculate the average gate fidelity between the target unitary U (d -dimension) and its noisy implementation $\mathcal{E}$ [18], defined by

$$\begin{aligned}
 F(U, \mathcal{E}) &= \int d\psi \langle \psi | U^\dagger \mathcal{E}(|\psi\rangle\langle\psi|) U | \psi \rangle \\
 &= \frac{dF_{\text{pro}}(\mathcal{E}, U) + 1}{d + 1} \\
 &= \frac{\text{Tr}(S_U^\dagger S_{\mathcal{E}})/d + 1}{d + 1}.
 \end{aligned} \tag{1}$$

The integral is performed over the uniform distribution of all pure states. $F_{\text{pro}}(\mathcal{E}, U)$ is the process fidelity. S_U and $S_{\mathcal{E}}$ are the superoperator representation of U and $\mathcal{E}$. When only considering unitary errors, the fidelity calculation can be simplified to

$$F(U, U_H) = \frac{|\text{Tr}(U^\dagger U_H)|^2/d + 1}{d + 1}. \tag{2}$$

U_H is the unitary that is actually implemented on hardware.

3. KAK-based approximation

For a parasitic two-qubit gate error U_E , one could apply $U_E^\dagger$ to completely mitigate this error. Nevertheless, $U_E^\dagger$ may require several applications of native two-qubit gates, introducing more hardware errors in the computation. In this section, we present a mitigation method based on KAK decomposition. The method finds an optimal choice of single-qubit unitaries to counteract the parasitic two-qubit gate such that the unitary fidelity is maximised after mitigation.

3.1. KAK decomposition

By applying the KAK decomposition [11–13], any arbitrary two-qubit unitary U can be written as

$$U = (K_1 \otimes K_2) U_A(\alpha, \beta, \gamma) (K_3 \otimes K_4), \tag{3}$$

where K_i is a single-qubit unitary and

$$U_A(\alpha, \beta, \gamma) = \exp[i(\alpha X \otimes X + \beta Y \otimes Y + \gamma Z \otimes Z)], \tag{4}$$

with $\alpha, \beta, \gamma \in \mathbb{R}$. We define $K_l = K_1 \otimes K_2$ and $K_r = K_3 \otimes K_4$. Two unitaries are equivalent under local operations if they have the same $U_A(\alpha, \beta, \gamma)$. In this work, we will restrict to the Weyl chamber [13, 19]

$$\{\pi/4 \geq \alpha \geq \beta \geq |\gamma| \text{ and } \gamma \geq 0 \text{ if } \alpha = \pi/4 | (\alpha, \beta, \gamma) \in \mathbb{R}^3\}. \tag{5}$$

The unitary $U_A(\alpha, \beta, \gamma)$ may need to be further decomposed into several applications of a native two-qubit gate. For a target two-qubit unitary $U = K_l U_A(\alpha, \beta, \gamma) K_r$ and an implementable unitary $V = U_A(\alpha', \beta', \gamma')$, the fidelity $F(U, K_l' V K_r')$ after optimisation over single-qubit gates is maximised when taking $K_l' = K_l$ and $K_r' = K_r$ (section 3.1 in [20], lemma 66 in [21]). When V is the identity gate, one has

$$F(U, K_l K_r) = \max_{K_l', K_r'} F(U, K_l' K_r'). \tag{6}$$

3.2. KAK approximation for general unitary errors

We now show how to use the KAK decomposition to mitigate parasitic gate errors. Let us assume that the target two-qubit gate U_T has a two-qubit unitary error U_E , that is, the actual unitary implemented on quantum hardware is $U_H = U_E U_T$. The unitary error U_E can be decomposed as

$$U_E = K_{El} U_A(\alpha_E, \beta_E, \gamma_E) K_{Er}. \quad (7)$$

We then compute the unitary fidelity $F(U_T, U_H)$ based on equation (2),

$$\begin{aligned} F(U_T, U_H) &= \frac{|\text{Tr}(U_T^\dagger U_H)|^2 / 4 + 1}{5} \\ &= \frac{|\text{Tr}(U_E)|^2 / 4 + 1}{5} \\ &= F(U_E, I). \end{aligned} \quad (8)$$

Since two-qubit gates have higher error rates, we consider minimising this unitary error by only applying single-qubit rotations after the hardware gate. Assume the single-qubit mitigation gate is K_{EM} , then the unitary fidelity after mitigation becomes

$$\begin{aligned} F_{EM}(U_T, K_{EM} U_H) &= \frac{|\text{Tr}(U_T^\dagger K_{EM} U_H)|^2 / 4 + 1}{5} \\ &= \frac{|\text{Tr}(K_{EM} U_E)|^2 / 4 + 1}{5} \\ &= F(U_E, K_{EM}^\dagger). \end{aligned} \quad (9)$$

$F(U_E, K_{EM}^\dagger)$ is maximised when

$$K_{EM} = K_{Er}^\dagger K_{El}^\dagger, \quad (10)$$

which can be proved by substituting the unitary U in equation (6) with U_E in equation (7), that is,

$$F_{EM}(U_E, K_{El} K_{Er}) = \max_{K'_l, K'_r} F(U_E, K'_l K'_r). \quad (11)$$

Applying the result in equation (11) to equations (8) and (9), we can prove that performing the mitigation gate $K_{EM} = K_{Er}^\dagger K_{El}^\dagger$ improves the unitary fidelity (i.e., $F(U_E, K_{EM}^\dagger) \geq F(U_T, U_H)$) and this gate is optimal among all single-qubit rotations. We call this error mitigation approach **KAK-approx**. The maximally achievable unitary fidelity by KAK-approx only depends on the parasitic two-qubit gate,

$$\begin{aligned} F_{EM}^{\max}(U_T, K_{EM} U_H) &= \frac{|\text{Tr}(U_A(\alpha_E, \beta_E, \gamma_E))|^2 / 4 + 1}{5} \\ &= [1 + 4 \cos^2(\alpha_E) \cos^2(\beta_E) \cos^2(\gamma_E) + 4 \sin^2(\alpha_E) \sin^2(\beta_E) \sin^2(\gamma_E)] / 5. \end{aligned} \quad (12)$$

3.3. KAK approximation for excitation-preserving unitary errors

In this section, we consider a special class of parasitic two-qubit gate errors, which is, a general excitation number-preserving two-qubit gate with the following form

$$U_{NP}(\theta, \xi, \chi, \eta, \phi) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & e^{-i(\eta+\xi)} \cos(\theta) & -i e^{-i(\eta-\chi)} \sin(\theta) & 0 \\ 0 & -i e^{-i(\eta+\chi)} \sin(\theta) & e^{-i(\eta-\xi)} \cos(\theta) & 0 \\ 0 & 0 & 0 & e^{-i(2\eta+\phi)} \end{pmatrix}, \quad (13)$$

where θ is the iSWAP angle, ϕ is the CPhase angle, ξ, χ, η are single-qubit phase angles. This set of gates has been termed as the fermionic simulation gate set due to its natural representation in simulating fermionic operators [14] and can be decomposed into

$$U_{\text{NP}}(\theta, \xi, \chi, \eta, \phi) = R_Z(-\eta, -\eta) R_Z\left(\frac{-\xi + \chi}{2}, \frac{\xi - \chi}{2}\right) U_{\text{NP}}(\theta, 0, 0, 0, \phi) R_Z\left(\frac{\xi + \chi}{2}, \frac{-\xi - \chi}{2}\right), \quad (14)$$

where $R_Z(\phi_1, \phi_2) = \exp[i(\phi_1 + \phi_2)/2] R_Z(\phi_1) \otimes R_Z(\phi_2)$ [7]. Applying the KAK decomposition on equation (3), we get

$$U_{\text{NP}}(\theta, 0, 0, 0, \phi) = [R_Z(-\phi/2) \otimes R_Z(-\phi/2)] U_A(-\theta/2, -\theta/2, -\phi/4). \quad (15)$$

If the parasitic gate U_E on a target two-qubit unitary U_T is a general excitation-preserving two-qubit unitary, i.e., $U_H = U_E U_T = U_{\text{NP}}(\theta, \xi, \chi, \eta, \phi) U_T$, then the maximal fidelity that can be achieved by applying single-qubit gate mitigation (based on equations (9)–(12)) is

$$F_{\text{EM}}^{\text{max}}(U_T, K_{\text{EM}} U_{\text{NP}}(\theta, \xi, \chi, \eta, \phi) U_T) = [1 + 4 \cos^2(\theta/2) \cos^2(\theta/2) \cos^2(\phi/4) + 4 \sin^2(\theta/2) \sin^2(\theta/2) \sin^2(\phi/4)] / 5. \quad (16)$$

An iSWAP-like gate can be expressed as

$$\text{iSWAP}(\theta) = U_{\text{NP}}(\theta, 0, 0, 0) = U_A(-\theta/2, -\theta/2, 0). \quad (17)$$

For the target gate with a parasitic iSWAP error (i.e., $U_E = \text{iSWAP}(\delta)$), one cannot improve its unitary fidelity by only applying single-qubit gates. In comparison, for a parasitic error in the form of

$$\begin{aligned} \text{CPhase}(\phi) &= U_{\text{NP}}(0, 0, 0, \phi) \\ &= [R_Z(-\phi/2) \otimes R_Z(-\phi/2)] U_A(0, 0, -\phi/4), \end{aligned} \quad (18)$$

the unitary fidelity without error mitigation is

$$F(U_T, U_H) = \frac{3 \cos(\phi) + 7}{10}. \quad (19)$$

One can improve the fidelity by performing the single-qubit gate correction $K_{\text{EM}} = R_{Z_1}(\phi/2) R_{Z_2}(\phi/2)$. Afterwards, the local components of the parasitic CPhase are exactly cancelled, and only the entangling part remains. The fidelity after mitigation is

$$F_{\text{EM}} = F(U_T, K_{\text{EM}} U_H) = \frac{2 \cos(\phi/2) + 3}{5}. \quad (20)$$

When angle ϕ is small, the infidelities of the unmitigated and mitigated unitary can be approximated to

$$1 - F \approx \frac{3\phi^2}{20} \quad \text{and} \quad 1 - F_{\text{EM}} \approx \frac{\phi^2}{20}.$$

We have shown how to mitigate parasitic errors on a native two-qubit gate. We now apply this approach to a composite unitary that needs to be decomposed into several applications of native gates and evaluate its fidelity improvements. We consider a general iSWAP-like gate that requires two applications of $\sqrt{\text{iSWAP}}^\dagger$ gates as shown in figure 1. If each $\sqrt{\text{iSWAP}}^\dagger$ gate has a parasitic CPhase(ψ) error, the implemented unitary will be $\text{iSWAP}(\theta) \text{CPhase}(2\psi)$. After applying the KAK-approx mitigation, the unitary becomes $\text{iSWAP}(\theta) \exp(-iZZ\psi/2)$. Based on equations (19) and (20), the approximate infidelities of the implemented $\text{iSWAP}(\theta)$ gates without and with error mitigation are $3\psi^2/5$ and $\psi^2/5$, respectively.

In summary, KAK-approx reduces the unitary infidelity by a factor of 3 for both the single hardware two-qubit gate $\sqrt{\text{iSWAP}}^\dagger$ and the composite gate $\text{iSWAP}(\theta)$. Compared to the unmitigated implementation for CPhase angle $\psi = 9$ degrees (the largest CPhase error angle mentioned in [7]), the KAK-approx mitigation approach reduces the infidelities of $\text{iSWAP}(\theta)$ and $\sqrt{\text{iSWAP}}^\dagger$ gates from 1.5% to 0.5% and from 0.37% to 0.123%, respectively. We note that, since the circuit for implementing $\text{iSWAP}(\theta)$ already contains single-qubit R_Z rotations on either side of each $\sqrt{\text{iSWAP}}^\dagger$ gate (figure 1), the R_Z gates which implement this




Figure 2. Decomposition of a target unitary into several applications of hardware two-qubit gate U_H . R_i is a single-qubit gate and may need be decomposed into a sequence of R_X and R_Z rotations.

KAK-approx mitigation approach can be introduced by merely modifying the R_Z -rotation angles which already appear in the circuit, and therefore no additional gates are required. Mitigating parasitic CPhase errors on the $\sqrt{i\text{SWAP}}^\dagger$ gate by adding single-qubit rotations has been experimentally demonstrated in [5]. Here we show that it is a special application of the general KAK-approx approach.

4. Gate recompilation

In the KAK-approx approach, we assume a target unitary has been decomposed into several applications of two-qubit gate U and we apply single-qubit rotations to mitigate the effect of the parasitic gate U_E on U . In this section, we present a recompilation approach that directly recompiles a target unitary into the hardware two-qubit gate $U_H = U_E U$ interleaved with single-qubit gates (figure 2) by using a numerical decomposition method [22]. It has been shown that any arbitrary two-qubit gate can be constructed by several six applications of an entangling gate³ plus single-qubit rotations [23]. Therefore, if the hardware two-qubit gate is entangling (which is typically true), then the circuit in figure 2 is universal for two-qubit gates.

We choose the numerical optimisation technique because it has the flexibility to decompose any target unitary into any native gate and can achieve comparable performance as analytical decomposition methods [22, 24]. The decomposition performance is measured by the native gate count required for achieving an accuracy. We set the accuracy tolerance of the numerical decomposition method to be 10^{-8} (which is much higher than state-of-the-art gate fidelity). Once the infidelity of a decomposition reaches this threshold, the numerical optimisation will terminate. Higher-fidelity decomposition could be found if the accuracy tolerance is set to be a lower value, but longer optimisation time may be required.

We verify the performance of the numerical decomposition by evaluating the expressivity of the $\sqrt{i\text{SWAP}}^\dagger$ gate for arbitrary two-qubit gates. Figure 3 shows that around 53% of the two-qubit unitaries in the Weyl chamber can be implemented using two perfect $\sqrt{i\text{SWAP}}^\dagger$ with infidelity below 10^{-8} . All two-qubit unitaries can be composed by three applications of $\sqrt{i\text{SWAP}}^\dagger$ with nearly perfect fidelity. These evaluation results are similar to the results by using the analytical decomposition method in [17], demonstrating the good performance of the numerical decomposition method. Moreover, figure 3 also shows that the $\text{CPhase}(\pi/20)\sqrt{i\text{SWAP}}^\dagger$ gate has similar expressivity power as the $\sqrt{i\text{SWAP}}^\dagger$ gate.

We note that arbitrary single-qubit gates are typically required to find an exact decomposition⁴ for an arbitrary target unitary. Single-qubit R_Z rotations may be enough for decomposing a special class of two-qubit target unitaries with specific native two-qubit gate (see examples in figure 1). It may be beneficial to minimise the number of R_X gates because they could have higher error rates than R_Z rotations [7]. In this work, we evaluate two numerical mitigation approaches, one named **Recompile** uses arbitrary single-qubit gates and one named **Recompile-RZ** only allows single-qubit R_Z rotations. **Recompile-mG** has at most m hardware two-qubit gates.

Figure 4 shows the average infidelity and two-qubit gate count for implementing the $i\text{SWAP}(\theta)$ gates when using different error mitigation approaches and only considering parasitic CPhase errors. The baseline implementation does not apply any mitigation (NoMitigate) and directly uses the decomposition in figure 1 which requires two $\sqrt{i\text{SWAP}}^\dagger$ gates and six R_Z gates. The KAK-approx mitigation approach applies single-qubit Z rotations after each $\sqrt{i\text{SWAP}}^\dagger$. These rotations are combined with existing R_Z gates and therefore there is no extra gate overhead. Instead of using the above standard decomposition, the Recompile mitigation approach decomposes each target unitary into the actual hardware gate

³ Any nonlocal two-qubit gate that is not locally equivalent to the SWAP gate is an entangling gate [13].

⁴ In this work, we assume a decomposition is exact if its unitary infidelity $\leq 10^{-8}$.






Figure 3. The fidelity distribution for implementing SU(4) unitary gates that are uniformly chosen from the Weyl chamber (equation (5)) with step $\pi/80$. Either $\sqrt{i\text{SWAP}}^\dagger$ or $\text{CPhase}(\pi/20)\sqrt{i\text{SWAP}}^\dagger$ is used as native gate. The vertical dashed line marks the unitary fidelity at 0.995. Decomposition with at most three native gates can achieve near-perfect unitary fidelity (yellow and green solid lines overlap). Around 70% (53%) of the unitaries can be implemented using two native gates with infidelity below 5×10^{-3} (10^{-8}).



(a) The average unitary infidelity



(b) The average two-qubit gate count

Figure 4. Comparison of different error mitigation approaches when only considering parasitic CPhase errors. The error bars represent the standard deviation of the mean over 1000 $i\text{SWAP}(\theta)$ unitaries, the angles are evenly chosen from $(0, \pi]$. NoMitigate is the baseline implementation without any error mitigation, i.e., each $\sqrt{i\text{SWAP}}^\dagger$ in the decomposition in figure 1 is experimentally realised with a parasitic $\text{CPhase}(\psi)$. The KAK-approx approach applies single-qubit Z rotations after each noisy two-qubit gate to partially mitigate the effect of a CPhase error. Recompile represents the numerical decomposition approach that uses the hardware two-qubit gate $\text{CPhase}(\psi)\sqrt{i\text{SWAP}}^\dagger$ as native gate. Recompile allows arbitrary single-qubit rotations and Recompile-RZ uses only Z rotations. Recompile-mG allows at most m hardware two-qubit gates.

$\text{CPhase}(\psi)\sqrt{i\text{SWAP}}^\dagger$ plus single-qubit rotations. As shown all mitigation approaches can decrease the unitary infidelity compared to NoMitigate.

We note that Recompile-RZ-2G does not perform the same as KAK-approx. KAK-approx always uses two $\sqrt{i\text{SWAP}}^\dagger$ gates for $i\text{SWAP}(\theta)$ unitaries of which angles are not $\pi/4$ or $3\pi/4$ (figure 1) and the unitary infidelity is the same for these target unitaries (equation (20)). In comparison, around nine percent of $i\text{SWAP}(\theta)$ unitaries will be constructed by only one $\text{CPhase}(\psi)\sqrt{i\text{SWAP}}^\dagger$ gate when using Recompile-RZ-2G (two applications of hardware two-qubit gates will not improve unitary fidelity). These $i\text{SWAP}(\theta)$ unitaries are the ones close to $i\text{SWAP}(\frac{\pi}{4})$ or $i\text{SWAP}(\frac{3\pi}{4})$. The achieved unitary fidelity by Recompile-RZ-2G may vary across target unitaries, causing a large fidelity variance in figure 1. Recompile-RZ-3G has the same performance as Recompile-RZ-2G for $i\text{SWAP}(\theta)$ unitaries and are therefore not presented in figure 1. Since KAK-approx achieves similar mean fidelity as Recompile-RZ-2G but has a faster implementation, we will only consider KAK-approx in the later evaluation.

The recompiling approach with arbitrary single-qubit gates (Recompile-3G) can find an exact decomposition for each unitary. That is, Recompile-3G can completely mitigate the unitary errors

introduced by parasitic CPhase gates. Yet, it requires more native two-qubit gates for implementing $i\text{SWAP}(\theta)$ unitaries (around three per unitary) than other mitigation approaches (around two per unitary in KAK-approx). Interestingly, if the maximum number of two-qubit native gates is limited to 2 (Recompile-2G), the unitary infidelity of $i\text{SWAP}(\theta)$ can be reduced to around 0.1% (0.17%) when the CPhase angle is 7(9) degrees. In quantum systems with high gate error rates, it may be beneficial to use an approximate decomposition that uses fewer hardware gates (e.g., KAK-approx), i.e., improving the overall gate fidelity by trading off the hardware errors with decomposition inaccuracy.

5. Fidelity with hardware errors

In the previous section, we have shown that an $i\text{SWAP}(\theta)$ unitary can be implemented with at most two $\sqrt{i\text{SWAP}}^\dagger$ gates and an arbitrary $\text{SU}(4)$ gate requires at most three $\sqrt{i\text{SWAP}}^\dagger$ gates. We have also compared different mitigation approaches for implementing $i\text{SWAP}(\theta)$ unitaries when the native gate $\sqrt{i\text{SWAP}}^\dagger$ has a parasitic CPhase error. In this section, we perform similar evaluations on arbitrary $\text{SU}(4)$ gates. Specifically, we evaluate 3309 $\text{SU}(4)$ unitaries by uniformly discretising the Weyl chamber (equation (5)) with step $\pi/80$ and 20 of them are $i\text{SWAP}(\theta)$ unitaries.

Besides the parasitic-CPhase errors, we take other hardware errors into account and compare the average gate fidelity of the KAK-approx mitigation approach and the recompilation approach with m noisy two-qubit gates and arbitrary single-qubit gates (Recompile-mG). Specifically, we apply relaxation errors after each gate (including idling gate) based on T_1/T_2 times and its gate duration t . We assume $T_2 = 2T_1$ and model the relaxation noise as an amplitude damping channel with $p = 1 - e^{-t/T_1}$. We also apply single-qubit and two-qubit depolarising noises based on single-qubit and two-qubit gate error rates. In this evaluation, we use noise parameters which are similar to the values presented in [2, 4] as shown in table 1. We evaluate two sets of depolarising error rates. For the first set, we choose the error rates of the single-qubit R_X rotations and the two-qubit gates $\sqrt{i\text{SWAP}}^\dagger$ to be 0.0003 ($p_X^{(1)}$) and 0.0048 ($p_S^{(1)}$) such that when adding them together with relaxation errors the measured total gate error rates go to 0.001 and 0.005, respectively. For the second set, we set the depolarising error rates of single-qubit R_X gates and two-qubit gates to be 0.001 ($p_X^{(2)}$) and 0.005 ($p_S^{(2)}$). We do not apply depolarising errors on the single-qubit R_Z gates.


Furthermore, we compare the proposed software mitigation methods with a hardware method that can suppress the parasitic CPhase errors by increasing gate duration as introduced in [1, 2]. We assume the hardware two-qubit gate $\sqrt{i\text{SWAP}}^\dagger$ is CPhase-free if the gate duration is two times (2XLong-mG) or four times (4XLong-mG) long. We also assume the hardware mitigation methods suffer from the same relaxation and depolarising errors as other methods.

We estimate the average fidelity (equation (1)) of $i\text{SWAP}(\theta)$ gates and $\text{SU}(4)$ gates as shown in figures 5 and 6, respectively. For both $i\text{SWAP}(\theta)$ and $\text{SU}(4)$ gates, applying the KAK-approx error mitigation is always beneficial compared to the non-mitigated ones. The unitary fidelity achieved by KAK-approx or NoMitigate decreases quadratically as the parasitic CPhase angle increases and has a constant offset in the presence of relaxation and depolarising errors. The recompilation and long-duration approaches can achieve (nearly) perfect fidelity when only considering coherent errors. However, they have higher offsets than KAK-approx or NoMitigate when adding relaxation and depolarising errors, that is, their high decomposition fidelity is compromised by hardware errors. Specifically, KAK-approx can achieve higher fidelity than 4XLong-4G(-2G) when the parasitic CPhase angle is smaller than around 6(4) degrees for $i\text{SWAP}(\theta)$ gates and 8(5) degrees for $\text{SU}(4)$ gates. Therefore, whether it takes $2\times$ or $4\times$ as long time to implement a CPhase-free gate is critical in determining whether this hardware mitigation approach is worthwhile.

Furthermore, when relaxation and depolarising errors are dominating in quantum systems, unitary implementation with fewer gates and shorter circuit duration is more beneficial. For example, an $i\text{SWAP}(\theta)$ gate requires two $\sqrt{i\text{SWAP}}^\dagger$ and six R_Z rotations as shown in figure 1. This decomposition is used in the KAK-approx and the 2X(4X)-long hardware method. Compared to KAK-approx, Recompile uses arbitrary single-qubit rotations (figure 2) and may have more two-qubit gates. An arbitrary single-qubit rotation will be decomposed into several R_X and R_Z rotations [25]. Therefore, for implementing $i\text{SWAP}(\theta)$ gates under realistic noise models, Recompile could introduce more hardware errors than KAK-approx and achieves lower fidelity ((c) and (d) in figure 5). In contrast, most of $\text{SU}(4)$ unitaries typically requires three $\sqrt{i\text{SWAP}}^\dagger$ gates and eight arbitrary single-qubit rotations (only 0.6% of the evaluated $\text{SU}(4)$ unitaries are $i\text{SWAP}(\theta)$ gates). Recompile uses similar number of gates as a standard decomposition. This means Recompile can completely mitigate coherent errors without introducing extra incoherent errors and

Table 1. Hardware parameters used for calculating noise channels in section 5. p_X and p_S are the depolarising error rates of the single-qubit X rotation and the two-qubit gate $\sqrt{\text{iSWAP}}^\dagger$, respectively.


T_1	R_X	R_Z	$\sqrt{\text{iSWAP}}^\dagger$	$p_X^{(1)}$	$p_S^{(1)}$	$p_X^{(2)}$	$p_S^{(2)}$
25 μs	25 ns	10 ns	12 ns	0.0003	0.0048	0.001	0.005



therefore achieves higher fidelity than other methods ((c) and (d) in figure 6). In appendix B, we observe similar results for implementing $\text{CPhase}(\phi)$ and $\text{SU}(4)$ gates when $\text{CZ} = \text{CPhase}(\pi)$ is used as native gate and has an over-rotation angle ψ .

Based on these evaluations, we summarise the following mitigation strategy to improve the implementation fidelity of a target two-qubit unitary U_T :

- If U_T requires three applications of $\sqrt{\text{iSWAP}}^\dagger$ gates for an exact decomposition, then Recompile-3G is the best mitigation method.
- If only two $\sqrt{\text{iSWAP}}^\dagger$ gates are needed for U_T , then one can compare KAK-approx with the long-duration hardware mitigation method (if this method is available). The best approach may vary, depending on the parasitic CPhase angle and the duration of CPhase-free $\sqrt{\text{iSWAP}}^\dagger$ gate.




6. Conclusion and discussion

We have presented two software approaches, KAK-approx and Recompile, to mitigate the parasitic CPhase errors on two-qubit gates. We compared them with a hardware mitigation method, namely, the long gate implementation, under various hardware errors. The evaluation results imply that for different target unitaries, each mitigation approach has the best performance in different error regimes. Our work can provide guidance for efficient error mitigation on near-term quantum computers. That is, one can apply the most appropriate approach based on the calibration data to achieve the highest application fidelity. We have also shown that the proposed mitigation methods can be generalised to other unitary errors and other hardware two-qubit gates.

We note that the KAK-approx approach decreases the unitary infidelity of a CPhase-parasitic gate on average by a factor of 3, but it may not be effective for some applications. This is because the impact of coherent errors is state dependent. States that are close to the eigenstates of these errors will experience only a minor effect. For instance, the parasitic CPhase(ψ) gate causes a phase shift ($e^{-i\psi}$) on $|11\rangle$ of a superposition state $a_1|00\rangle + a_2|01\rangle + a_3|10\rangle + a_4|11\rangle$ (error-sensitive) and leaves state $a_1|00\rangle + a_2|01\rangle + a_3|10\rangle$ unchanged (error-insensitive). Applying KAK-approx on error-sensitive states will improve state fidelity. Applying this mitigation approach on error-insensitive states will introduce effective errors and decrease state fidelity. Further investigation may be required to understand what applications this approach is useful for. Nonetheless, since we expect generic highly entangled states generated in quantum computation to be far from an eigenstate of two-qubit gates, in typical cases we expect such error mitigation strategies to be advantageous.

Furthermore, it has been shown that a FSim gate ($\text{FSim}(\theta, \phi) = i\text{SWAP}(\theta)\text{CPhase}(\phi)$) and its ‘neighboring’ gate $\text{FSim}(\theta + \epsilon_1, \phi + \epsilon_2)$ have similar capability of expressing SU(4) gates (figure 8 in [22]). We conjecture that a general two-qubit gate $\exp(iH_1 t)$ in the Weyl chamber and its neighboring gate $\exp(iH_2 \epsilon)\exp(iH_1 t)$ (for small ϵ) would have similar expressivity. For mitigating a general two-qubit error by Recompile, we then would observe results similar to the parasitic CPhase error as presented in this work. We leave the numerical analysis on general coherent errors for future work. In addition, the single-qubit



rotations in the numerical decomposition approach are not unique, future work can optimise these rotations in terms of circuit duration or gate error rates.

We have only implemented the numerical decomposition for two-qubit gates in this work. One can generalise this approach to multiple-qubit gates by using the A* search algorithm [24]. Recompiling the whole circuit could be advantageous compared to recompiling each two-qubit gate in the circuit individually for some applications (see an example for the three-qubit Toffoli gate in appendix C where each two-qubit gate in the circuit is equivalent to the noiseless hardware gate up to single-qubit rotations). However, this may not be true for other circuits where the two-qubit gates in the circuit require a similar number of noiseless and noisy hardware gates (appendix C shows an example of decomposing the CPhase gates in a QFT circuit into the noiseless or noisy CZ gate). Moreover, the fidelity calculation in equation (2) scales exponentially with the number of qubits and the number of optimisation parameters increases linearly as the number of qubits [24]. For example, numerically decomposing the four-qubit and the three-qubit QFT circuits already takes around 400k s and 700 s respectively, compared to 3 s for decomposing a two-qubit QFT circuit [24]. Therefore, recompiling the whole circuit using numerical optimisation approaches may not be feasible for large circuits.

In addition, we have focused on mitigating parasitic two-qubit gate errors. Coherent errors can also be caused by crosstalk [26, 27]. For example, the dispersive coupling on transmon qubits can cause ZZ-type errors between gate qubits and between a gate qubit and a spectator qubit. Both KAK-approx and Recompile work for arbitrary two-qubit coherent errors and therefore could mitigate crosstalk errors in this form. In appendix B, we have evaluated both approaches for mitigating phase errors on CPhase gates. Nevertheless, it may not be feasible to use the recompilation approach for mitigating crosstalk errors that involve multiple qubits due to its limited scalability. Alternatively, it may be more efficient to use recompilation to mitigate the ZZ-type errors on gate qubits and use dynamical decoupling [28] to mitigate the errors between a gate qubit and a spectator qubit. Furthermore, both KAK-approx and Recompile require precise information of the coherent errors but characterising crosstalk errors is hard [27, 29]. One may convert crosstalk errors into Pauli errors and then apply generic error mitigation techniques such as probabilistic error cancellation [30].





Figure 9. Decomposition of $\text{CPhase}(\phi)$ into CZ and single-qubit gates. $R_X(\theta) = \exp(-i\theta X/2)$. $H = R_Z(\pi/2)R_X(\pi/2)R_Z(\pi/2)$.

Acknowledgments

LL thanks Yang Wang for useful feedback on the manuscript. LL and DEB acknowledge funding from the EPSRC Prosperity Partnership in Quantum Software for Modelling and Simulation (Grant No. EP/S005021/1).

Data availability statement

The data that support the findings of this study are available upon reasonable request from the authors.

Appendix A. Error bar plots for iSWAP-like gates

In this section, we show the error bars of the average gate fidelity when the native two-qubit gate is $\sqrt{\text{iSWAP}}$ (figures 7 and 8). We note that the required native two-qubit gate counts for implementing SU(4) unitaries vary from 1 to 3 while almost all $\text{iSWAP}(\theta)$ unitaries require two native two-qubit gates. This larger variance in gate count causes a larger variance in the amount of hardware errors. Therefore, the fidelity of SU(4) unitaries has a larger variance than the fidelity of $\text{iSWAP}(\theta)$ unitaries.






Figure 10. The fidelity distribution for implementing SU(4) unitaries that are uniformly chosen from the Weyl chamber (equation (5)) with step $\pi/80$. The native two-qubit gates are CZ and CPhase($\pi/20$)CZ. The vertical dashed line marks the unitary fidelity (infidelity) at 0.995 (0.005).

Table 2. Hardware parameters used for calculating noise channels in appendix B. p_X and p_s are the depolarising error rates of R_X and CZ, respectively.


T1	R_X	R_Z	CZ	$p_X^{(1)}$	$p_s^{(1)}$	$p_X^{(2)}$	$p_s^{(2)}$
25 μ s	25 ns	0 ns	15 ns	0.0003	0.0047	0.001	0.005




(a) with only coherent (CPhase offset) errors



(b) With both coherent and relaxation errors.




(c) With coherent, relaxation, and depolarising errors $(p_X^{(1)}, p_{CZ}^{(1)})$.



(d) With coherent, relaxation, and depolarising errors $(p_X^{(2)}, p_{CZ}^{(2)})$.

Figure 11. The average gate fidelity of 80 different CPhase(ϕ) gates, where ϕ is evenly chosen from $(0, \pi]$. The native two-qubit gate is CZ and has an over-rotation CPhase angle. Lines represent the mean values and error bars represent the standard deviation.



Appendix B. Phase errors on CPhase gates

The native two-qubit gates vary across quantum computers. For example, the quantum processor presented in [26] have $\text{CPhase}(\phi)$ as native gates and these gates may acquire conditional phase errors due to dispersive coupling. If the CPhase errors are not mitigated on the hardware and can be characterized, then one can mitigate these errors using the proposed software approaches in sections 3 and 4.

In this section, we evaluate the average unitary fidelity of the $\text{CPhase}(\phi)$ and SU(4) gates when $\text{CZ} = \text{CPhase}(\pi)$ is used as native gate and has an over-rotation angle ψ , i.e., the actual gate unitary is $\text{CPhase}(\pi + \psi)$. We use an analytical method to decompose a $\text{CPhase}(\phi)$ gate into CZ as shown in figure 9. When CZ gates have these systematic errors, KAK-approx will apply R_Z corrections after each CZ in this circuit. For decomposing arbitrary SU(4) gates with CZ or any decomposition based on the numerical approach, the circuit structure is the same as figure 2 which alternates arbitrary single-qubit gate layers and two-qubit gate layers.

Similar to figure 3, we first evaluate the expressivity of CZ gate to demonstrate the good performance of the numerical decomposition approach. Figure 10 shows that all the SU(4) gates can be constructed by three applications of perfect CZ, which is consistent with the theoretical results in [31, 32]. However, only around 13% (36%) of SU(4) gates can be composed by two CZ gates with infidelity below 10^{-8} (5×10^{-3}). The noisy CZ gate ($\text{CPhase}(\pi/20)\text{CZ}$) has similar expressivity power as the perfect CZ.

We then compare these error mitigation methods under different hardware errors. The noise parameters used in this evaluation are presented in table 2. We evaluate two sets of depolarising error rates. For the first set, the depolarising error rates of the single-qubit gate R_X and the two-qubit gate CZ are chosen to be 0.0003 and 0.0047 such that when adding them with relaxation errors the measured gate error rates go to 0.001 and 0.005, respectively. R_Z gates are assumed to be performed virtually and are error-free in this evaluation.

Figures 11 and 12 show similar results to section 5 that has $\sqrt{\text{iSWAP}}^\dagger$ as native two-qubit gate. For the $\text{CPhase}(\phi)$ gates in figure 11, Recompile can completely mitigate the effect of coherent errors. However, its




Figure 13. (a) The three-qubit Toffoli circuit. (b) Recompiled each two-qubit gate in (a) into the noisy CZ ($\text{CPhase}(\phi) = \text{CPhase}(\pi/20)\text{CZ}$) plus single-qubit rotations separately. In total, 12 noisy CZ gates are required for implementing Toffoli. (c) Recompiling the unitary of the whole circuit (a) only requires nine noisy CZ gates plus single-qubit rotations.




Figure 14. (a) The three-qubit QFT circuit. (b) Recompiled each two-qubit gate in (a) into the noisy CZ ($\text{CPhase}(\phi) = \text{CPhase}(\pi/20)\text{CZ}$) plus single-qubit rotations separately. In total, six noisy CZ gates are required. (c) Recompiling the unitary of the whole circuit in (a) could require more noisy CZ gates.

benefits are reduced when considering relaxation and depolarising errors due to its higher number of R_X gates compared to KAK-approx. For example, with relaxation errors, KAK-approx can achieve higher fidelity than Recompile when the CPhase over-rotation angle is smaller than around 7 degrees. In contrast, both Recompile and KAK-approx have similar circuit structures for implementing $\text{SU}(4)$ unitaries. Therefore, Recompile always outperforms KAK-approx for the $\text{SU}(4)$ gates as shown in figure 12. The average gate fidelity of $\text{SU}(4)$ unitaries has a larger variance than the fidelity of $\text{CPhase}(\phi)$ unitaries because of its larger variance in required native two-qubit gate count.

Appendix C. Recompilation of three-qubit unitary gates

To compare the performance of recompiling the unitary of a whole circuit with the performance of recompiling each two-qubit gate in the circuit separately, we have extended the numerical decomposition approach to three-qubit gates. Our implementation assumes specific circuit structures and more optimised circuits could be found by using the A* algorithm in [24]. We assume the native two-qubit gate is CZ and has a CPhase error $\text{CPhase}(\pi/20)$. Figures 13 and 14 show the examples of recompiling the Toffoli gate and the three-qubit QFT circuit, respectively. The CNOT gate in the original circuit implementation of the

Toffoli gate is equivalent to CZ up to single-qubit rotations but it requires two applications of the noisy CZ (CPhase($\pi/20$)CZ). Recompiling the Toffoli unitary achieves better performance than recompiling each CNOT gate separately. In comparison, the CPhase gates in the QFT circuit require two applications of CZ or noisy CZ gates. Recompiling the unitary of the entire QFT circuit does not outperform the recompilation of individual two-qubit gates. More optimised decomposition results may be found by improving the recompilation approach.

ORCID iDs

Lingling Lao  <https://orcid.org/0000-0001-6870-5670>

Wojciech Mruczkiewicz  <https://orcid.org/0000-0002-8497-6363>

Dan E Browne  <https://orcid.org/0000-0003-3001-158X>

References

- [1] Yan F, Krantz P, Sung Y, Kjaergaard M, Campbell D L, Orlando T P, Gustavsson S and Oliver W D 2018 Tunable coupling scheme for implementing high-fidelity two-qubit gates *Phys. Rev. Appl.* **10** 054062
- [2] Foxen B *et al* 2020 Demonstrating a continuous set of two-qubit gates for near-term quantum algorithms *Phys. Rev. Lett.* **125** 120504
- [3] McKay D C, Sheldon S, Smolin J A, Chow J M and Gambetta J M 2019 Three-qubit randomized benchmarking *Phys. Rev. Lett.* **122** 200502
- [4] Arute F *et al* 2019 Quantum supremacy using a programmable superconducting processor *Nature* **574** 505–10
- [5] Arute F *et al* 2020 Hartree–Fock on a superconducting qubit quantum computer *Science* **369** 1084–9
- [6] Collodo M C *et al* 2020 Implementation of conditional phase gates based on tunable ZZ interactions *Phys. Rev. Lett.* **125** 240502
- [7] Arute F *et al* 2020 Observation of separated dynamics of charge and spin in the Fermi–Hubbard model (arXiv:2010.07965)
- [8] Knill E *et al* 2008 Randomized benchmarking of quantum gates *Phys. Rev. A* **77** 012307
- [9] Greenbaum D 2015 Introduction to quantum gate set tomography (arXiv:1509.02921)
- [10] Boixo S, Isakov S V, Smelyanskiy V N, Babbush R, Ding N, Jiang Z, Bremner M J, Martinis J M and Neven H 2018 Characterizing quantum supremacy in near-term devices *Nat. Phys.* **14** 595–600
- [11] Khaneja N, Brockett R and Glaser S J 2001 Time optimal control in spin systems *Phys. Rev. A* **63** 032308
- [12] Kraus B and Cirac J I 2001 Optimal creation of entanglement using a two-qubit gate *Phys. Rev. A* **63** 062309
- [13] Zhang J, Vala J, Sastry S and Whaley K B 2003 Geometric theory of nonlocal two-qubit operations *Phys. Rev. A* **67** 042313
- [14] Kivlichan I D, McClean J, Wiebe N, Gidney C, Aspuru-Guzik A, Chan G K-L and Babbush R 2018 Quantum simulation of electronic structure with linear depth and connectivity *Phys. Rev. Lett.* **120** 110501
- [15] Farhi E, Goldstone J and Gutmann S 2014 A quantum approximate optimization algorithm (arXiv:1411.4028)
- [16] Lacroix N *et al* 2020 Improving the performance of deep quantum optimization algorithms with continuous gate sets *PRX Quantum* **1** 110304
- [17] Huang C, Ding D, Wu F, Kong L, Zhang F, Ni X, Shi Y, Zhao H-h and Chen J 2021 Towards ultra-high fidelity quantum operations: SQiSW gate as a native two-qubit gate (arXiv:2105.06074)
- [18] Horodecki M, Horodecki P and Horodecki R 1999 General teleportation channel, singlet fraction, and quasidistillation *Phys. Rev. A* **60** 1888–98
- [19] Cross A W, Bishop L S, Sheldon S, Nation P D and Gambetta J M 2019 Validating quantum computers using randomized model circuits *Phys. Rev. A* **100** 032328
- [20] Watts P, Vala J, Müller M M, Calarco T, Birgitta Whaley K, Reich D M, Goerz M H and Koch C P 2015 Optimizing for an arbitrary perfect entangler. I. Functionals *Phys. Rev. A* **91** 062306
- [21] Peterson E C, Crooks G E and Smith R S 2020 Two-qubit circuit depth and the monodromy polytope *Quantum* **4** 247
- [22] Lao L, Murali P, Martonosi M and Browne D 2021 Designing calibration and expressivity-efficient instruction sets for quantum computing *ACM/IEEE 48th Annual Int. Symp. on Computer Architecture (ISCA)* pp 846–59
- [23] Zhang J, Vala J, Sastry S and Whaley K B 2003 Exact two-qubit universal quantum circuit *Phys. Rev. Lett.* **91** 027903
- [24] Davis M G, Smith E, Tudor A, Sen K, Siddiqi I and Iancu C 2019 Heuristics for quantum compiling with a continuous gate set (arXiv:1912.02727)
- [25] Barenco A, Bennett C H, Cleve R, DiVincenzo D P, Margolus N, Shor P, Sleator T, Smolin J A and Weinfurter H 1995 Elementary gates for quantum computation *Phys. Rev. A* **52** 3457–67
- [26] Krinner S *et al* 2020 Benchmarking coherent errors in controlled-phase gates due to spectator qubits *Phys. Rev. Appl.* **14** 024042
- [27] Sarovar M, Proctor T, Rudinger K, Young K, Nielsen E and Blume-Kohout R 2020 Detecting crosstalk errors in quantum information processors *Quantum* **4** 321
- [28] Tripathi V, Chen H, Khezri M, Yip K-W, Levenson-Falk E M and Lidar D A 2021 Suppression of crosstalk in superconducting qubits using dynamical decoupling (arXiv:2108.04530)
- [29] Harper R, Flammia S T and Wallman J J 2020 Efficient learning of quantum noise *Nat. Phys.* **16** 1184–8
- [30] van den Berg E, Mineev Z K, Kandala A and Temme K 2022 Probabilistic error cancellation with sparse Pauli–Lindblad models on noisy quantum processors (arXiv:2201.09866)
- [31] Vatan F and Williams C 2004 Optimal quantum circuits for general two-qubit gates *Phys. Rev. A* **69** 032315
- [32] Vidal G and Dawson C M 2004 Universal quantum circuit for two-qubit transformations with three controlled-NOT gates *Phys. Rev. A* **69** 010301