Characterization of Coherent Errors in Noisy Quantum Devices

Noah Kaufmann,* Ivan Rojkov, and Florentin Reiter†
Institute for Quantum Electronics, ETH Zürich, Otto-Stern-Weg 1, 8093 Zürich, Switzerland
(Dated: July 19, 2023)

Characterization of quantum devices generates insights into their sources of disturbances. State-of-the-art characterization protocols often focus on incoherent noise and eliminate coherent errors when using Pauli or Clifford twirling techniques. This approach biases the structure of the effective noise and adds a circuit and sampling overhead. We motivate the extension of an incoherent local Pauli noise model to coherent errors and present a practical characterization protocol for an arbitrary gate layer. We demonstrate our protocol on a superconducting hardware platform and identify the leading coherent errors. To verify the characterized noise structure, we mitigate its coherent and incoherent components using a gate-level coherent noise mitigation scheme in conjunction with probabilistic error cancellation. The proposed characterization procedure opens up possibilities for device calibration, hardware development, and improvement of error mitigation and correction techniques.

I. INTRODUCTION

Current quantum computers are highly affected by noise [1]. The errors originate from interactions of the devices with their environment [2], unwanted dynamics between the qubits [3], or imperfect control signals [4]. Characterization aims to find error sources and quantify their effect on the performance of a computing processor [5]. The information obtained from the characterization of a device facilitates the development of a less error-prone hardware platform [6], the implementation of schemes for correcting specific errors [7], and the application of noise mitigation protocols [8]. Therefore, noise characterization stands at the core of future technological progress.

Demarcating noise along the line of purity conservation leads to the distinction between coherent and incoherent noise [9]. *Coherent* noise occurs due to systematic, reversible perturbations such as imperfect calibrations, imprecise control signals, or couplings to low-frequency external fields [10–13]. Conversely, *incoherent* noise, e.g., bit-flip errors, is stochastic and typically arises from insufficient isolation of the system from its environment [12], causing a non-reversible loss of information. Since controllability and isolation of a system are generally opposing goals in hardware design [4, 14], distinguishing the structure of the limiting noise is critical to advance quantum computing platforms.

Numerous characterization techniques have been developed to identify and quantify coherent and incoherent noise in quantum processes. Full process tomography [15, 16] and gate-set tomography [17, 18] are common techniques for obtaining a general representation of a physical operation. Although they provide a full representation, they lack in discriminating the noise processes from the actual operation. Furthermore, as they involve the reconstruction of density matrices, they are experimentally intractable on devices with more than a few qubits due to exponential scaling of the number of operations with the size of the system's Hilbert space and the resource-intensive post-processing [19]. While some protocols reach scalable noise characterization using principles of randomized benchmarking [20–22], they are often limited to a few heuristic metrics that do not provide enough insights for existing noise calibration, mitigation, or correction techniques [23]. Characterization approaches that exceed this limitation while still being scalable exist for incoherent noise [7, 24]. These *noise reconstruction* protocols [25] introduce a model and estimate its parameters by fitting it to the real process. They aim to quantify noise of a gate layer, where a restricted noise correlation length is the central structural assumption. However, those models do not portray the actual noise process completely, as the protocols require the application of probabilistic *twirling* schemes [20, 26–30] that convert coherent errors into incoherent ones. Currently, no similar approach to the noise reconstruction protocols exists that is sensitive to coherent errors.

A useful benchmark to evaluate the effectiveness of a noise characterization protocol is to apply its results to *error mitigation* techniques, i.e., methods aiming to remove, on average, noise-induced bias from experimental results. Unitary errors can, for instance, be suppressed with dynamical decoupling [31–33], dynamically corrected operations [34, 35], or hidden inverses [36]. Despite their effectiveness in certain situations [3, 37–41], coherent error mitigation techniques are often tailored to specific systems or require precise knowledge of unitary errors. Incoherent noise can be mitigated using probabilistic error cancellation or zero-noise extrapolation techniques [8, 42]. These have been proven to be efficient against stochastic errors using the information from noise reconstruction methods [24, 43–45], but only twirling of the intrinsic noise processes allows to apply those techniques for coherent noise mitigation. Consequently, it is essential to develop scalable noise characterization protocols that provide information simultaneously for coherent and incoherent noise mitigation techniques.

* knoah@ethz.ch
† freiter@phys.ethz.ch

In this article, we propose a characterization technique that accounts for unitary errors and show how it can be used to estimate the model parameters. To characterize the coherent part of the model, we devise a protocol that is compatible with the noise reconstruction method. The proposed model and protocol are hardware-agnostic and scalable to larger systems, assuming a low noise-correlation length. To demonstrate our method, we apply it to characterize a gate layer on a 7-qubit superconducting circuit device provided by IBM Quantum (IBMq) [46]. This allows us to determine the most significant coherent single- and two-qubit errors and their behavior over time. Finally, we leverage the gained information from the protocol and mitigate the coherent errors through a gate-level approach. We thereby show the validity of our model and the accuracy of our protocol. Our work paves the way for advancements in hardware platform development, quantum error correction and mitigation schemes, and the design of noise-aware algorithms.

II. MODEL

An ideal noiseless operation on a quantum computer can be expressed by a *unitary channel* $\mathcal{U}_I(\rho) = U_I \rho U_I^\dagger$, with U_I being some unitary matrix. A simple model for the noisy implementation of this operation on real hardware $\mathcal{E}_\mathcal{P}$ consists of the concatenation of a noise modeling *Pauli channel* $\mathcal{P}$ and the ideal unitary,

$$\mathcal{E}_\mathcal{P}(\rho) = \mathcal{P}(\mathcal{U}_I(\rho)) = \sum_i p(i) P_i U_I \rho U_I^\dagger P_i, \qquad (1)$$

where the Pauli error rates $p(i)$ form a probability distribution over the n-qubit Pauli operators $P_i \in \mathbb{P}^{\otimes n} = \{I, X, Y, Z\}^{\otimes n}$. Compared to general completely positive trace-preserving maps, Pauli channels offer a concise description of noisy operations by restricting the model to only Pauli operators. Despite those substantial expressibility restrictions of the model, a wide range of incoherent noise, including dephasing and depolarization, can be described by Pauli channels. In this framework, a n-qubit noisy operation is modeled at most by 4^n Pauli error rates and operators. To reduce this scaling from exponential to polynomial, we follow the widely used and experimentally motivated approach of only regarding *locally* correlated noise, i.e., considering n-qubit Pauli operators P_i that act nontrivially on l qubits only [7, 24]. These l qubits are selected with respect to the qubits' connectivity in the quantum computing architecture. This work focuses on a correlation length $l = 2$.

Pauli noise models are appealing for different physical and practical reasons. Not only is the interaction of a qubit system with its environment often dominated by depolarization and dephasing [47, 48], but quantum error correction also leads to noise better approximated by Pauli channels [49–51]. Additionally, existing mitigation schemes applicable to Pauli noise [8, 24] are generally implementable with low-depth circuits on near-term devices. Lastly, Pauli channels are interesting for theoretical work, as they can be efficiently simulated classically [52].

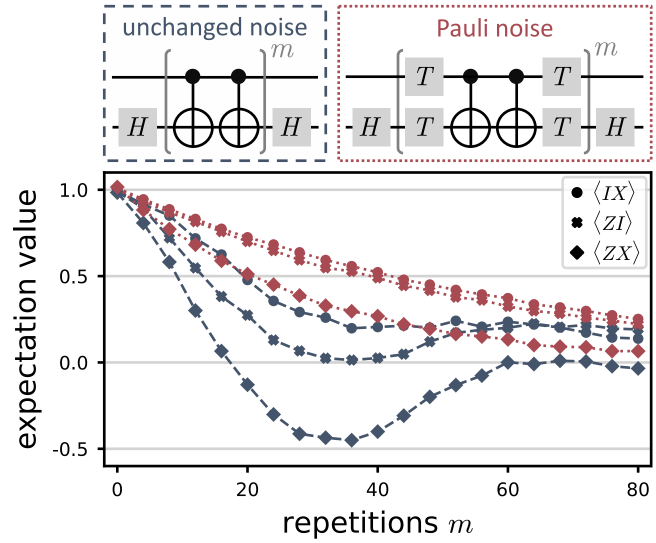

FIG. 1. Noise examination on identity circuit. The data is obtained on `ibm_lagos` and plots the expectation value of the three two-qubit Pauli operators P_{IX}, P_{ZI}, P_{ZX} against the number of noisy identities m implemented by two CX gates. In the experiment, we prepare and measure the system in the state $|0+\rangle$ using one Hadamard gate H. The blue dashed curves correspond to the unchanged noise. The dotted red curves are obtained using Pauli twirling scheme that projects the noise to a Pauli noise (cf. Eq. (1)). Each circuit is measured 4096 times and the twirled results are averaged over 32 different twirling gates T. A readout mitigation scheme [53] is applied in post-processing based on the reported readout error rates of IBM Quantum.

To illustrate the limitations of the Pauli noise model and motivate its extension, we study noise introduced by identity circuits on IBMq processors. We perform an *echo-experiment* (similar to Ref. [54]) where we prepare a two-qubit system in the Pauli basis state $|0+\rangle$, apply an even number of CX gates, and measure the evolved state in the same Pauli basis as it was prepared. This measurement allows us to estimate the expectation value of all two-qubit Pauli operators of which the prepared state is an eigenstate, specifically P_{IX}, P_{ZI}, and P_{ZX}. In Fig. 1, we diagram these expectation values against the number of noisy identities m obtained from the `ibm_lagos` processor. The results for different input states and superconducting hardware platforms available on IBMq are similar. In the noiseless case, all of those expectation values are independent of the number of repetitions of the identity and equal to 1. In the exclusive presence of Pauli noise, a purely exponential decay of the expectation values, characteristic of incoherent error processes, is expected [20–22, 27, 55]. Consequently, the Pauli noise model of Eq. (1) cannot explain the oscillations observed in the experiment of Fig. 1 (blue dashed curves).

A well-established technique to address this discrepancy between model and reality is Pauli twirling [20, 26–30], a procedure which involves stochastically introducing Pauli gates before and after a Clifford operation. Since the Pauli group is fixed under Clifford conjugation, carefully chosen sets of Pauli gates can effectively cancel each other out and enforce the errors to follow a Pauli noise model given in Eq. (1) [27, 55]. The Pauli twirled version of the echo-experiment presented above is illustrated in Fig. 1 with red dotted curves. Notably, we observe that the expectation values follow a desired exponential trend. However, when examining the physical noise sources of a system, altering the noise structure is often not desired, as it obscures the distinction between coherent and incoherent noise. Furthermore, the probabilistic nature of twirling creates a circuit and sampling overhead [24, 26, 27].

Instead of altering the noise, we modify the noise model and introduce a coherently rotated Pauli noise model. We extend the Pauli noise model in Eq. (1) with a unitary channel $\mathcal{U}_\theta$ (cf. Fig. 2) intended to explain the oscillatory behavior observed in Fig. 1

$$\mathcal{E}_{\mathcal{P},\theta}\left(\rho\right) = \mathcal{U}_\theta(\, \mathcal{E}_{\mathcal{P}}(\rho)) = U_\theta \sum_i p(i) P_i U_I \rho U_I^\dagger P_i U_\theta^\dagger \quad (2)$$

with U_θ being a unitary operator parameterized by a Hermitian matrix H_θ as $U_\theta = e^{-iH_\theta}$. This matrix can be expressed in the Pauli basis as $H_\theta = \sum_k \theta_k P_k$, with real coefficients θ_k and Pauli matrices P_k. Assuming small noise, i.e., U_θ is close to identity, we can rewrite the coherent noise as $U_\theta \approx \prod_k \exp(-i\theta_k P_k)$, which is a valid approximation up to $\mathcal{O}(\theta^2)$. Each product term reads

$$\exp(-i\theta_k P_k) = \cos\left(\theta_k\right) \mathbb{I} - i \sin\left(\theta_k\right) P_k \approx \mathbb{I} - i\theta_k P_k\,, \quad (3)$$

and represents in the Bloch sphere picture a rotation around the k-axis by the angle $2\,\theta_k$, cf. Fig. 2. To reach a scalable model, we consider, as for the Pauli noise model in Eq. (1), locally correlated noise operators, namely single- and two-qubit rotations among neighboring qubits according to their connectivity in the hardware. Therefore, the *coherently rotated Pauli noise model* defined in Eq. (2) consists of nine two-qubit rotations for each neighboring qubit pair and three single-qubit rotations for each qubit, leading to a magnitude of these rotations of $|\theta| = 15$.

Note that this model is not universal. Nonunital noise [56], such as decay, surpasses its expressibility. In the discourse regarding the influence of coherent noise on error correcting schemes and the fault tolerance threshold, similar noise models to Eq. (2) are utilized in Refs. [57–59]. However, those models are limited to specific types of Pauli and coherent noise, and the works do not focus on the characterization of a noisy process but on analyzing the properties of a process with the respective noise model.

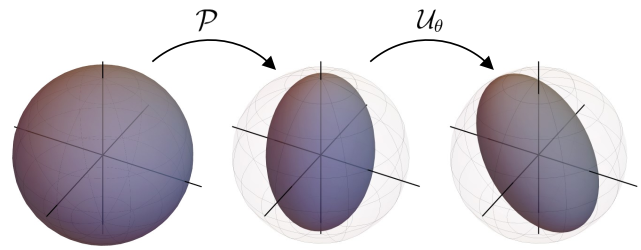

FIG. 2. Bloch sphere transformations. Image of the Bloch sphere on the left under the action of a Pauli channel $\mathcal{P}$ (middle), and an additional unitary transformation $\mathcal{U}_\theta$ (right).

III. CHARACTERIZATION

We now present our novel protocol for characterizing all the parameters of the coherently rotated Pauli noise model introduced in Eq. (2). The protocol relies on the assumption that the amplitude of both coherent and incoherent noise is small, allowing us to utilize the small angle approximation in Eq. (3).

We express the coherent noise characterization protocol in the *Pauli transfer matrix* (PTM) formalism [60]. As the Pauli matrices build a complete basis of the operator space, a n-qubit state ρ can be vectorized $|\rho) \in \mathbb{R}^{4^n}$ in the Pauli basis with the components being their expectation values $|\rho)_i = \frac{1}{2^n}\mathrm{Tr}\left[P_i\rho\right]$. These 4^n expectation values can be estimated from 3^n measurements of ρ in different Pauli bases [7, 61–63]. For a single qubit, this state representation is known as the Bloch vector [64]. In this formalism, a channel $\mathcal{E}$ is represented by a matrix $T \in \mathbb{R}^{4^n} \times \mathbb{R}^{4^n}$ that acts on the vectorized states according to $|\mathcal{E}(\rho)) = T\,|\rho)$. The matrix T is called the PTM, with elements $T_{ij} = \mathrm{Tr}\left[P_i\, \mathcal{E}(P_j)\right]$.

Before addressing the characterization protocol applicable to a quantum process modeled by the coherently rotated Pauli noise model of Eq. (2), we present the case of estimating U_θ for a 2-qubit process described by $\mathcal{E}_\theta(\rho) = U_\theta \rho U_\theta^\dagger$, meaning that we disregard both the Pauli noise $\mathcal{N}$ and the ideal unitary $\mathcal{U}_I$ channels. The Pauli transfer matrix T_θ corresponding to $\mathcal{E}_\theta$ is a square matrix of order 16, parameterized by 15 single- and two-qubit rotation angles. While the first row and column are trivial and independent of θ, in the small angle approximation (cf. Eq. (3)), the other elements of this matrix are given by

$$\left(T_\theta\right)_{ij} \approx \delta_{ij} - \frac{1}{2}\theta_k \mathrm{Tr}\left[P_i P_k P_j\right] = \delta_{ij} + \theta_k\, C_{kij}\,, \qquad (4)$$

where we use Einstein's notation for the summation over k. Applying the channel to an arbitrary state $|\rho)$ leads to the equation $|\mathcal{E}_\theta(\rho))_j = |\rho)_j + \theta_k\, C_{kij}\, |\rho)_j$. Defining the matrix $(B_{|\rho)})_{ik} = C_{kij}\,|\rho)_j$, the estimation problem boils down to a linear inverse problem of form $B_{|\rho)}\,\theta = |\mathcal{E}_\theta(\rho)) - |\rho)$ where the two-dimensional matrix $B_{|\rho)}$ and the vectors $|\rho)$ and $|\mathcal{E}_\theta(\rho))$ can be obtained using state tomography. $B_{|\rho)}$ is singular and cannot be inverted. Therefore, experimentally estimating $|\rho)$ and

$|\mathcal{E}_\theta(\rho))$ for a single state ρ does not allow to estimate θ. However, by estimating the vector representation of k different states before and after the evolution by $\mathcal{E}_\theta$, we can create an over-determined system of $14N$ independent equations for the 15 unknown parameters. We concatenate the matrices $B_{|\rho)}$ for N different states into a single matrix $\tilde{B}$ of size $15N \times 15$ and combine N vectors $|\mathcal{E}_\theta(\rho)) - |\rho)$ to one vector y of length $15N$. Finally, we estimate θ using the least square method,

$$\hat{\theta} = \left(\tilde{B}^T \tilde{B}\right)^{-1} \tilde{B}^T y\,. \tag{5}$$

The strategy of the characterization protocol for the entire n-qubit system is to apply this estimation procedure to every involved two-qubit subsystem according to their connectivity in the hardware. Three points must be considered when moving from the discussed simple case to the model given in Eq. (2): first, the action of the ideal unitary $\mathcal{U}_I$; second, the presence of the Pauli noise $\mathcal{N}$; and lastly, the impact of two-qubit noise on qubits that are outside of the considered subsystem.

The effect of the ideal unitary channel $\mathcal{U}_I$ can be incorporated in post-processing after estimating the state $|\rho)$. For an arbitrary U_I, this procedure requires a full-state tomography of the system [52]. To preserve scalability, we restrict to operations, separable into multiple single- and two-qubit unitaries $U_I = U_1 \otimes U_2 \otimes ... \otimes U_l$. This requirement is fulfilled, e.g., in the characterization of so-called *gate layers*, i.e., n-qubit operations consisting of single- and two-qubit gates that can be executed simultaneously in experiments.

We address the presence of the Pauli noise channel $\mathcal{N}$ by assuming that both the incoherent and coherent noises are relatively small. In the scenario that $\max_{ijk} |p(i) - p(j)|\, |\theta_k| \ll 1$ where $p(i)$ and θ_k are the Pauli, respective coherent-error rates of Eq. (2), we can consider the Pauli channel to approximately commute with the coherent rotations. Then we can show that the presented estimation procedure is not disturbed by the presence of an additional Pauli channel (cf. App. A). Intuitively this is visible when illustrating the action of the two types of noise for a single qubit on a Bloch sphere. While coherent noise corresponds to a rotation, Pauli noise leads to a shrinkage of the sphere (cf. Fig. 2) which does not alter the orientation of the sphere poles and does, therefore, not interfere with the estimation of unitary rotations.

Lastly, when focusing exclusively on two-qubit subsystems, one has to address the correlated multi-qubit noise between a qubit in the regarded subsystem and the external ones. To handle this problem in the estimation of $|\rho)$ and $|\mathcal{E}_\theta(\rho))$, we prepare the surrounding qubits in different basis states for every single experiment run. This technique isolates the two desired qubits and randomizes the effect of the noise introduced outside the regarded subsystem, similar to twirling schemes.

We now use the presented protocol to characterize the coherent noise induced by an arbitrary gate layer executed on the IBMq Falcon Processor `ibm_lagos`. The result is summarized in Fig. 3, where we show the absolute values of the angles of the characterized single- and two-qubit coherent errors for an arbitrary gate layer. The layer is composed of two CX gates and three single-qubit gates, namely H, $\sqrt{X}$, and X. Considering the connectivity of the 7 qubits in the hardware and our assumption that the noise is only locally correlated, we prepare 216 initial states, all being product states of $|0\rangle$, $|1\rangle$, $|+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$, $|-\rangle = (|0\rangle - |1\rangle)/\sqrt{2}$, $|i\rangle = (|0\rangle + i\,|1\rangle)/\sqrt{2}$ and $|-i\rangle = (|0\rangle - i\,|1\rangle)/\sqrt{2}$. Each two-qubit subsystem is prepared in 36 states, and for each of those states, the environment is randomized with 6 additional states. The dimension of the y-vector in the estimation problem given in Eq. (5) is thus equal to 36 for each subsystem. To further increase the value of N, we execute 0 to 3 repetitions of the layer, amounting to 4 circuits per preparation. Finally, bypassing the effect of the unitary gate layer $\mathcal{U}_I$ in the post-processing necessitates performing state tomography on up to three-qubit subsystems, leading to 3^3 Pauli measurements per circuit. Overall, we run $216 \cdot 4 \cdot 27 = 23'328$ different circuits with 128 shots each, leading to a total running time in the system of about 30 minutes.

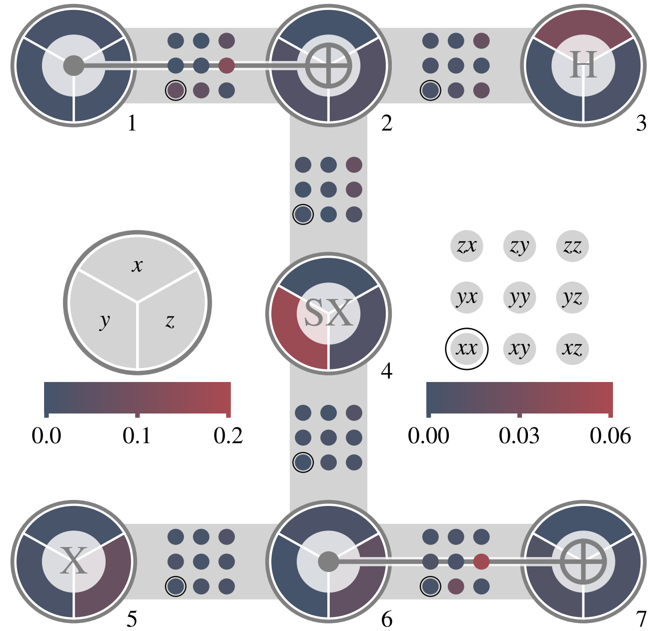

FIG. 3. Characterization of the coherent noise introduced by executing a gate layer on the 7-qubit IBMq Processor `ibm_lagos` [46]. The gate layer is drawn on top of the connectivity tree (qubit connectivity in the quantum computing architecture), where each big circles marks a different qubit. The gate layer consists of two CX gates between qubits 1 and 2, and qubits 6 and 7, a Hadamard (H) gate on qubit 3, a $\sqrt{X}$ (SX) gate on qubit 4 and a X gate on qubit 5. The three pieces inside the qubit representing circles represent the single-qubit coherent errors, and the nine small circles between two qubits display the two-qubit coherent cross-talk errors. All errors are measured in radians and the absolute value is plotted.

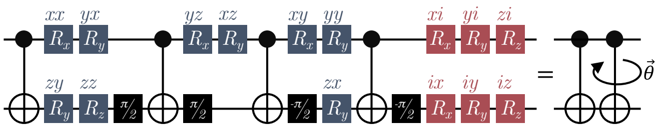

FIG. 4. Coherent noise mitigating circuit element. The circuit element comprises 4 CX gates, 4 fixed z-rotations by $\pi/2$, respectively $-\pi/2$, and 15 parameterized single-qubit rotations. The first 9 rotations (blue) account for different two-qubit coherent errors, and the final 6 rotations (red) mitigate single-qubit errors. Without the first CX gate, the circuit element represents a corrected CX gate. We represent the element by the symbol on the right-hand side, where $\vec{\theta}$ corresponds to the angles of the 15 parameterized rotations.

The used IBMq platforms report readout errors on the order of 1%, caused by relaxation, imperfect coupling to the readout resonator, and signal amplification errors [65]. To counter the influence of these errors, we apply a post-processing readout mitigation scheme in all executed experiments [53].

From Fig. 3, we conclude that the most significant two-qubit coherent errors occur between qubits 1 and 2 and qubits 6 and 7. Those are associated with the Pauli P_{YZ} error. This result is expected, as those are the pairs of qubits on which the CX gates are acting. This Pauli error has already been noted in the effective analysis of the echoed cross-resonance gate used to realize the CX gate [66]. The predominant coherent errors are the single-qubit ones. We note no strong bias towards a specific single-qubit rotation error, and those errors could be related to the hardware implementation of each single-qubit gate.

We conduct the characterization of the same device and gate layer as in Fig. 3 three times, with time differences between the runs being 6 hours and 19 days. The results available in App. B show that the leading coherent errors in the system do not change over time. The most significant single-qubit coherent errors change by no more than 10% and the two-qubit coherent rotation errors shift maximally by 0.018 radians. This result suggests that the platform suffers from systematic coherent errors not influenced by the daily re-calibration of the system.

IV. MITIGATION

In principle, it is possible that the characterization results in Fig. 3 originate from overfitting the non-ideal process to the introduced model. Indeed, it is unknown whether the coherently rotated Pauli noise model of Eq. (2) is expressive enough to resemble the main characteristics of the probed IBMq process. To verify that the model captures the dominant noise contributions of a process on the probed device, we examine whether the output of the presented protocol facilitates significant noise mitigation.

Returning to the unmitigated circuits in Fig. 1, our goal is to suppress the observed oscillatory behavior in the echo-experiment by correcting the characterized coherent errors at the circuit level. In case of a successful correction, according to the introduced model, one would expect a purely exponential decay of the expectation values with the number of process repetitions due to the Pauli noise. From this exponential decay, we can then characterize the Pauli error rates [7, 24] and mitigate these noise components using the probabilistic error cancellation method (PEC) [8]. If the final measurement statistics after coherent error mitigation and PEC closely follow the ideal statistics, we demonstrated that the proposed model is physical, and the characterization protocol provides valuable insights.

Any unitary channel is *reversible*. Therefore, in our model, the effect of the coherent noise can be canceled by single- and two-qubit rotations with opposite rotation angles than the characterized ones. Single-qubit rotations are readily available in most of the leading quantum computing architectures. To invert the two-qubit rotations, an entangling gate is necessary. We use the CX gate as it is a native gate of the chosen platform. By interleaving 15 single-qubit rotations into a structure of four CX gates and four fixed single-qubit rotations, as shown in Fig. 4, a circuit element with 15 parameters is created. These parameters, which we will denote as $\vec{\theta}$, are directly obtained from the characterization procedure presented in the previous section. The element corresponds to the identity if all parameters are set to zero.

We run the echo-experiments for three different situations: unmitigated; coherent-error mitigated; and coherent-error mitigated combined with PEC. The results are presented in Fig. 5. Fig. 5(a) and (b) were obtained with 2048 shots per circuit, whereas for Fig. 5(c), an ensemble of 280 altered circuits with 512 shots each was executed per experiment.

For the unmitigated case in Fig. 5(a), we see two main characteristics: Firstly, the expectation values of the 12 Pauli operators that should be ideally 0 are spreading. Secondly, the expectation values that ideally equal 1, meaning that the states are eigenstates of the corresponding Pauli, decay. Note that the expectation value of the identity in the plot is not based on actual experiments and is always 1 due to the trace constraints on density matrices.

Applying the presented coherent noise characterization protocol on the data of Fig. 5(a) leads to a first estimate of the rotation angles of the coherent errors $\hat{\theta}_0$. To further improve the characterization, we run the characterization scheme with $\vec{\theta} = \hat{\theta}_0$ and obtain $\hat{\theta}_1$. This iterative approach is expected to improve the characterization because it compensates for violations of the assumed commutation of single- and two-qubit rotation errors. Fig. 5(b) shows the result of the final mitigation iteration with the estimate $\vec{\theta} = \hat{\theta}_1$ using the gate-based mitigation circuit from Fig. 4. The observed changes in the angles between the first and the last iteration were about an order of magnitude smaller than the biggest

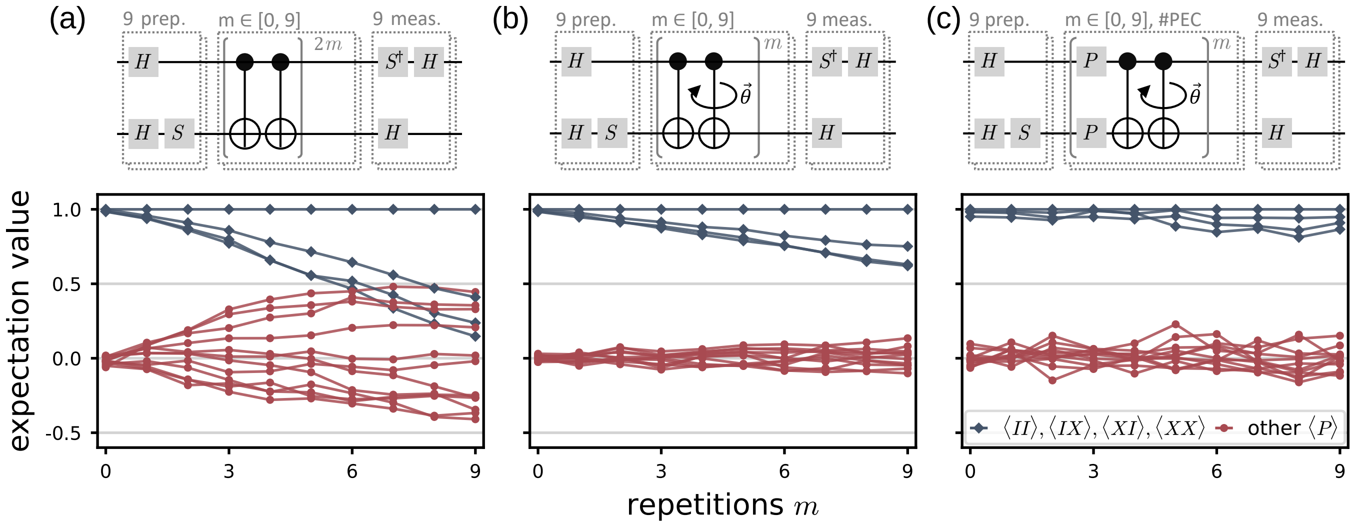

FIG. 5. Noise mitigation according to the coherently rotated Pauli noise model in Eq. (2) on `ibm_lagos`. The plot shows the estimated expectation values of all 9 two-qubit Pauli operators for the initial state $|{++}\rangle$ after m repetitions of a noisy identity (a) without mitigation, (b) with coherent error mitigation based on data of (a), and (c) with probabilistic error cancellation based on data of (b) and coherent error mitigation. $|{++}\rangle$ is an eigenstate of the operators P_{IX}, P_{XI}, and P_{XX} (blue), as well as the identity P_{II} (purple), in the noiseless case, their expectation values are 1. The expectation values of the other 12 Pauli operators (red) are 0. The circuit diagrams contain the preparation of each qubit in one of the states $|0\rangle, |+\rangle, |i\rangle$, by applying Hadamard ($H$) and/or phase ($S$) gates. The evolved states are measured in 9 bases to estimate all 16 Pauli expectation values. In (c), specific Pauli operators (P) are inserted between the circuit elements according to the PEC method.

coherent errors and thus supporting the validity of the taken approximation. In Fig. 5(b), the mitigation effect is clearly visible, as the spreading behavior of the expectation values that are ideally 0 is largely suppressed. The expectation values of the Pauli operators, of which the prepared state is a +1 eigenstate, are higher in the mitigated case compared to the unmitigated estimation. However, the coherent noise mitigation does not retrieve their ideal-case expectation value of 1, as according to our model, the statistic is still affected by Pauli noise.

To show that the examined noise follows the coherently rotated Pauli noise model in Eq. (2), we estimate the Pauli error rates from the mitigated curves of Fig. 5(b) and apply PEC to mitigate the Pauli noise. Note that different from the above presented coherent noise correction, PEC is a mitigation technique that does not correct deterministically individual circuits (as Pauli channels are not reversible) but allows noise-free estimation of expectation values from circuit ensembles. It is important to emphasize that compared to previous works demonstrating PEC [24, 43], we do not use Pauli twirling and, thus, avoid the associated circuit and sampling overhead.

The Pauli eigenvalues of the Pauli noise channel are estimated by fitting exponential curves to the expectation values that are supposed to be one in the ideal case. The Pauli error rates can then be found by a Walsh-Hadamard transform [7]. We follow the PEC protocol of van den Berg *et al.* [24] to mitigate the incoherent noise. It consists of applying the non-physical inverse of a Pauli channel by running an ensemble of unitary processes composing this inverse. Those results are weighted according to the Pauli error rates and combined in post-processing. The results of the echo-experiments with coherent error mitigation and PEC are displayed in Fig. 5(c). Compared to Fig. 5(b), the data are noisier, which we suspect to be due to the PEC sampling overhead. The obtained measurement statistics clearly resemble the main features of the ideal measurement statistics. Therefore, we can conclude that the introduced noise model explains the leading noise terms in the examined system and that the presented characterization protocol generates valid physical insights into the noise. Similar characterization and mitigation results were obtained on `ibm_jakarta`.

Based on the data from Fig 5(a) and (b), we can reconstruct the density matrices for the 9 prepared initial states at every repetition of the noisy identity. We then estimate their fidelity with their ideal initial state. The standard state fidelity formula is used $F(\sigma,\rho) = \mathrm{Tr}\left[\sqrt{\sqrt{\rho}\sigma\sqrt{\rho}}\right]^2$. We observe in Fig. 6 that the mitigation decreases the average infidelity by about a factor of two.

V. GENERALIZATION OF THE PROTOCOL

We motivated our coherently rotated Pauli noise model and devised the protocol for its characterization in a manner agnostic toward the choice of the qubit-based quantum computing architecture. The same procedure can be employed to characterize coherent errors in superconducting devices as we demonstrated in this work, but also

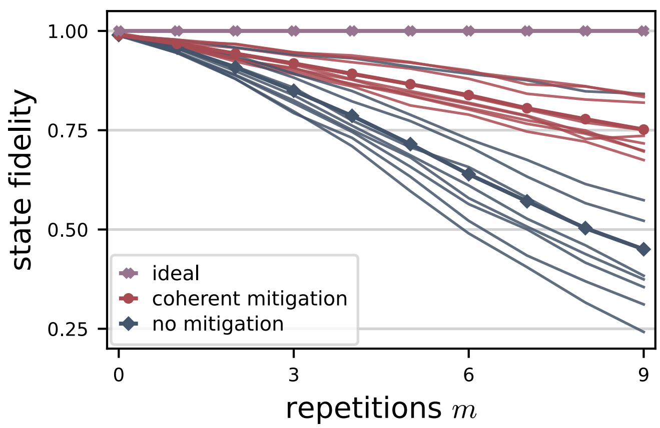

FIG. 6. Influence of the coherent error mitigation on fidelity measurement between the ideal state and states evolved by the noisy identity. From the data of Fig. 8, the density matrix of the evolved state is estimated for each number of repetitions of the noisy identity and each of the 9 prepared states. The plot shows the fidelity between these state estimates and the ideally prepared state. The thin lines represent the evolution of the fidelity with the number of repetitions of the noisy identity for each state, and the thick lines are the average of the 9 thin lines.

in trapped ions, neutral atoms, or any other platforms, as it only requires the preparation of elementary Pauli states without the need for entanglement.

Furthermore, we can readily extend the proposed model to qudit-based quantum processors, which are also susceptible to coherent errors [67–71]. In these systems, Pauli X and Z operators are generalized as $X_d |s\rangle = |s+1 \bmod d\rangle$ and $Z_d |s\rangle = \omega^s |s\rangle$, respectively. Here, $|s\rangle$ represents a qudit computational state with $s \in \{1, 2, \ldots, d\}$ and $\omega = \exp(i2\pi/d)$. Consequently, the Pauli noise model for a n-qudit system can be expressed similarly to Eq. (1) but with $P_i \in \{(X_d)^p (Z_d)^q \mid q, p \in \{1, 2, \ldots, d\}\}^{\otimes n}$ representing the generalized Pauli operators. Since the generalized Pauli operators exhibit similar properties to qubit Pauli gates (such as being unitary 1-designs and normalized by the qudit Clifford group), the Pauli transfer matrix formalism remains valid in the qudit framework. It thus enables the same twirling and characterization techniques employed in the qubit case to be utilized [67–71]. Moreover, this implies that the unitary errors can be effectively modeled by the coherently rotated Pauli noise model described in Eq. (2) and that the characterization protocol presented in Sec. III can be utilized.

Similarly, our characterization protocol can be expanded to encompass the logical level, as logical qubits can be susceptible to unitary noise often caused by coherent errors at the physical level [72, 73]. To extend our protocol, it becomes necessary to prepare and measure logical Pauli basis states. These operations are inherent to various error correcting codes, thereby enabling the generalization of our protocol.

VI. CONCLUSION AND OUTLOOK

In this paper, we introduced a characterization technique to estimate coherent errors in noisy quantum devices. An existing Pauli noise model is extended to incorporate coherent errors by adding unitary rotations to the model. We devise a procedure to learn the unknown parameters of this coherently rotated Pauli noise model. The scalability of our approach is affirmed by assuming a limited correlation length of the coherent noise. As the results of the characterization of a physical quantum processor showed, the protocol is feasible to be executed on an actual device. For the tested device, the characterization revealed the leading coherent errors introduced by the probed gate layer. Notably, these characterized coherent errors remained stable over multiple re-calibration cycles, suggesting a systematic nature. To verify the validity of our noise model, we mitigated the coherent errors and the characterized Pauli noise. The coherent error mitigation was achieved with a gate-level correction scheme for two-qubit coherent errors. To counter the Pauli noise, we applied probabilistic error cancellation. The combination of the two techniques resulted in almost ideal measurement statistics. This finding shows that the model is sufficiently expressible to explain the main noise contributions of a specific process on an actual processor. Moreover, the successful mitigation of the characterized errors confirms the applicability of the developed protocol. Our protocol is hardware agnostic and thus expected to be executable on a variety of hardware platforms.

Our work offers several promising avenues for future research. The method can enhance mitigation schemes that rely on detailed knowledge of the noise structure. One such approach is optimizing the resilience to coherent noise during circuit transpilation based on the noise insights for each gate or gate layer [36, 40, 41, 74, 75]. A clear example of this is the hidden inverses method [36, 41], which mitigates the impact of coherent errors by compiling circuits in a way that induces destructive interference among the coherent errors of the gates. However, implementing hidden inverses requires knowledge of the physical nature of these errors, which in Ref. [41] was obtained through gate-set tomography. Integrating our approach into the calibration process can provide valuable information to the compiler, enabling it to leverage the hidden inverses technique while actively inverting the non-corrected coherent errors. Similarly, our technique opens the door to studying and developing noise-aware versions of advanced mitigation techniques such as dynamical decoupling or composite pulses. Moreover, information on coherent errors facilitates Pauli conjugation [76]. An efficient noise twirling method that limits sampling overhead compared to common Pauli twirling and, therefore, diminishes the reduction of the threshold of error-correcting codes.

Finally, identifying the nature and strength of the errors is also crucial for the fault-tolerant operation of a quantum processor unit at the logical level [51, 72, 77, 78]. While quantum error correction protocols transform local unitary errors into correctable probabilistic ones [49, 50], coherent noise can lead to significantly larger worst-case logical errors than only incoherent noise [9, 48, 50, 51, 58, 72, 73, 78–80]. The knowledge gained from coherent errors, which effectively represents soft information, can then be leveraged to enhance the performance of quantum error-correcting codes at both the decoding and encoding stages. Indeed, while optimal decoders are notoriously hard to find [81], recent studies have demonstrated the benefit of soft information in improving existing decoders and overall quantum error correction performance [82–84]. At the encoding level, it is possible to design codes to be naturally robust against specific coherent noises, as demonstrated in Refs. [85, 86]. Thus, our protocol has the potential to be a diagnostic tool that could help to decide on the next higher-level encoding.

In the long run, the presented characterization protocol paves the way for the development of improved hardware platforms by identifying the limiting imperfections. Yet, already in the near term it can serve as a tool for various applications that rely on knowledge of the existing noise structure, such as quantum error correction or hardware-aware algorithm design.

ACKNOWLEDGMENTS

The authors thank Jonathan Home, Elias Zapusek, Roland Matt, Jeremy Flannery, and Luca Huber for helpful comments and discussions throughout the project. This work was supported by the Swiss National Science Foundation (SNSF) through the National Centre of Competence in Research - Quantum Science and Technology (NCCR QSIT) grant 51NF40–160591. I.R. and F.R. acknowledge financial support by the Swiss National Science Foundation (Ambizione grant no. PZ00P2_186040). We acknowledge the use of IBM Quantum services for this work. The views expressed are those of the authors, and do not reflect the official policy or position of IBM or the IBM Quantum team.

APPENDICES

Appendix A: Addressing Pauli noise neglect during coherent noise characterization

This section complements the intuitive reasoning of Sec. III, on why we can neglect the influence of the Pauli noise channel when characterizing the present coherent noise, with a mathematically more rigorous explanation.

The action of a noisy channel modeled by $\mathcal{E}_{\mathcal{P},\theta}(\rho)$ as defined in Eq. (2) can be represented in terms of the vectorized form of the density matrix of ρ and the PTM of the channel. It reads $|\mathcal{E}_{\mathcal{P},\theta}(\rho)) = T_{\mathcal{E}_{\mathcal{P},\theta}}\,|\rho) = T_{U_\theta} T_{\mathcal{P}} T_{U_I}\,|\rho)$ where T_{U_θ}, $T_{\mathcal{P}}$ and T_{U_I} are the transfer matrices of the coherent, Pauli and the ideal unitary channels, respectively. As stated in the main text, we discuss the case of T_{U_I} being the identity. We moreover assume that $|\mathcal{E}_{\mathcal{P},\theta}(\rho))$ and $|\rho)$ can be experimentally estimated.

The estimation of the 15 parameters of T_{U_θ} can be written as a minimization problem over the estimator $\hat{T}_{U_\theta}$.

Statement 1. *With the assumption $T_{\mathcal{P}}$ and T_{U_θ} being close to the identity. The minimization of the function*

$$\left\| |\mathcal{E}_{\mathcal{P},\theta}(\rho)) - \hat{T}_{U_\theta} T_{\mathcal{P}}\,|\rho) \right\|^2 \tag{A1}$$

over $\hat{T}_{U_\theta}$ is equivalent to the minimization of

$$\left\| |\mathcal{E}_{\mathcal{P},\theta}(\rho)) - \hat{T}_{U_\theta}\,|\rho) \right\|^2 \,. \tag{A2}$$

Proof. A noisy identity process described by the coherent noise model, $|\mathcal{E}_{\mathcal{P},\theta}(\rho))$ equals $T_{U_\theta} T_{\mathcal{P}}\,|\rho)$. We show that $\theta' = \theta$ minimizes the expression $\|T_{\mathcal{P}} T_{U_\theta}\,|\rho) - T_{U_{\theta'}}\,|\rho)\|$ independent of the choice of ρ.

$$\begin{aligned}
&\min_{\theta'} \left\| |\mathcal{E}_{\mathcal{P},\theta}(\rho)) - T_{U_{\theta'}}\,|\rho) \right\|^2 \\
&= \min_{\theta'} \left\| T_{U_\theta} T_{\mathcal{P}}\,|\rho) - T_{U_{\theta'}}\,|\rho) \right\|^2 \\
&= \min_{\theta'} |\rho)^\top \left(T_{\mathcal{P}}^\top T_{U_\theta}^\top - T_{U_{\theta'}}^\top \right) \left(T_{U_\theta} T_{\mathcal{P}} - T_{U_{\theta'}} \right) |\rho) \\
&= \min_{\theta'} (\rho| \left(T_{\mathcal{P}}^\top T_{U_\theta}^\top T_{U_\theta} T_{\mathcal{P}} - T_{U_{\theta'}}^\top T_{U_\theta} T_{\mathcal{P}} - \right. \\
&\qquad \left. - T_{\mathcal{P}} T_{U_\theta}^\top T_{U_{\theta'}} + T_{U_{\theta'}}^\top T_{U_{\theta'}} \right) |\rho) \,.
\end{aligned} \tag{A3}$$

Given that the noise channel is close to the identity, we take the approximation that the diagonal matrix $T_{\mathcal{P}}$ commutes with the matrices T_{U_θ} and $T_{U'_\theta}$. Furthermore, because of because of the unitarity of the coherent rotation channel, $T_{U_\theta}^\top T_{U_\theta} = \mathbb{I}$ is valid. Then, the minimization problem boils down to

$$\min_{\theta'} (\rho| \left(\mathbb{I} + T_{\mathcal{P}}^2 - T_{\mathcal{P}} \left(T_{U_\theta}^\top T_{U_{\theta'}} + \left(T_{U_\theta}^\top T_{U_{\theta'}} \right)^\top \right) \right) |\rho) \,. \tag{A4}$$

The off-diagonal part of $T_{U_\theta}^\top T_{U_{\theta'}}$ is anti-symmetric. Consequently, $T_{U_\theta}^\top T_{U_{\theta'}} + \left(T_{U_\theta}^\top T_{U_{\theta'}}\right)^\top$ is diagonal and smaller or equal than $2\,\mathbb{I}$. Thereby, all involved matrices are diagonal. From $1 + T_{\mathcal{P}}^2 \geq 2\,T_{\mathcal{P}}$, it follows that the expression in Eq. (A4) is minimal when $T_{U_\theta}^\top T_{U_{\theta'}} + \left(T_{U_\theta}^\top T_{U_{\theta'}}\right)^\top = 2\,\mathbb{I}$. This equation has a unique solution corresponding to $T_{U_\theta} = T_{U_{\theta'}}$, which finally implies $\theta = \theta'$.

□

This statement implies that we can optimize for the coherent noise term without considering the present Pauli noise, allowing to apply the closed form solution of Eq. (5) or conduct a least squares optimization method for $\| |\mathcal{E}_{\mathcal{P},\theta}(\rho)) - \hat{T}_{U_\theta}\,|\rho)\|^2$.

The approximation of the diagonal matrix $T_{\mathcal{P}}$ commuting with the matrices T_{U_θ} is limited by the term $\max_{ijk} |p(i) - p(j)|\,|\theta_k| \ll 1$. The first part of this expression quantifies how isotropic the action of the Pauli

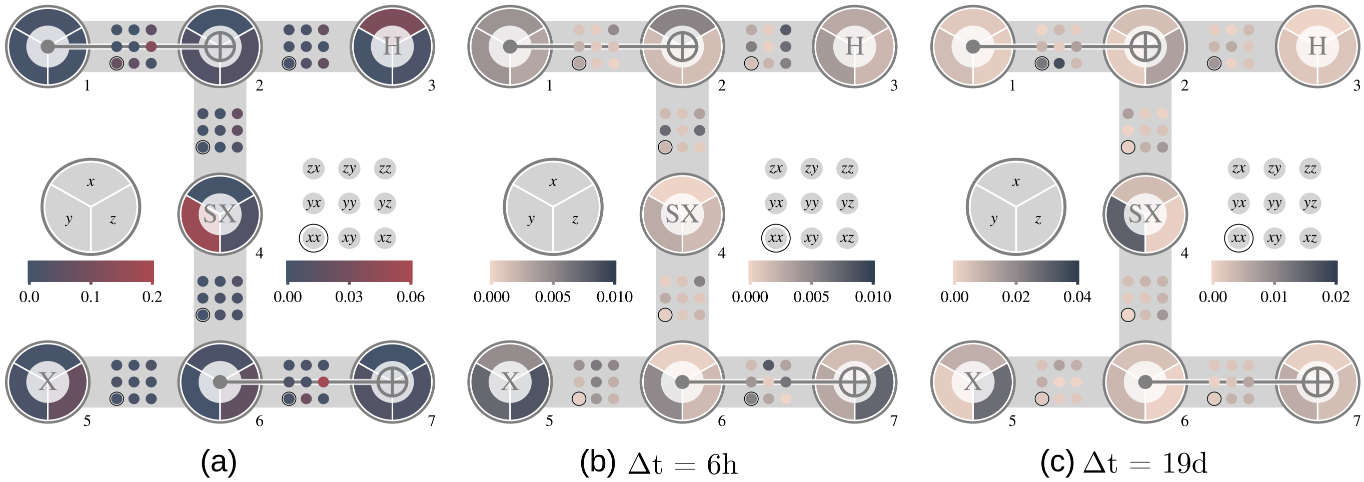

FIG. 7. Coherent noise drift. We repeated the characterization result of Fig. 3 at 3 different times. Subplot (a) shows the baseline data of Fig. 3. Subplot (b) displays the observed changes in the coherent error characterization over a time span of 19 days and subplot (c) exhibits the differences over (6) hours. The experiments were run on the 7-qubit IBM Quantum Processor `ibm_lagos` [46] on September 22, 2022, and October 11, 2022.

noise is. For the depolarizing channel which for a single qubit is illustrated as a uniformly contracting Bloch sphere, $\max_{ij}|p(i) - p(j)|$ is zero. In that case $T_{\mathcal{P}}$ can be written as $f \cdot \mathbb{I}$, with f being the depolarizing factor, and thus commutes with all other matrices. The second term of the expression, namely $\max_k|\theta_k|$, is a measure for the maximal off-diagonal terms in the coherent noise modeling PTM. As the Pauli noise modeling matrix is diagonal, the violation of the commutation is only facilitated by the off-diagonal of T_{U_θ}. The geometric pictures of scaling and rotation helps to understand this commutation relation better.

Appendix B: Coherent noise drift

We conducted the characterization of the gate layer represented in Fig. 3 again after 6 hours as well as 19 days. The shifts in the characterization results between those three experiments are shown in Fig. 7. First, we observe that the leading coherent errors in the system did not change over the 19 days. However, as expected, the changes observed for the two experiments within 6 hours are significantly smaller than the differences over the complete time span. After 6 hours, the maximum deviation for single-qubit coherent errors was 0.009 radians and 0.004 radians on average. The two-qubit errors changed maximally by 0.008 radians and on average by 0.003 radians. In the total period (i.e. after 19 days), the most significant single-qubit coherent errors have changed by maximally 10% (corresponding to about 0.03 radians). The two-qubit coherent rotation errors have changed by maximally 0.018 radians. The average deviation was 0.007 radians for the single-qubit errors and 0.003 radians for the two-qubit errors. These results suggest that the platform is affected by systematic coherent errors which are not influenced by the hourly and daily re-calibrations of the system during which qubits' frequency and readout angle, as well as pulses' amplitude and phase of basic single- and two-qubit gates, are calibrated [46].

Appendix C: Complementary mitigation data

Fig. 8 displays the complete data set used to create Fig. 6. The influence on the expectation values that ideally are 0 for all number of repetitions is clearly visible for all states, as they disperse much less in the mitigated case. Furthermore, the expectation values of the operators, of which the prepared state is an eigenstate, are closer to the ideal value of 1 in the mitigated case. While the mitigation did not work equally well for all 9 states, the whole data set gives a clear indication of the positive effect of the presented coherent noise mitigation scheme.

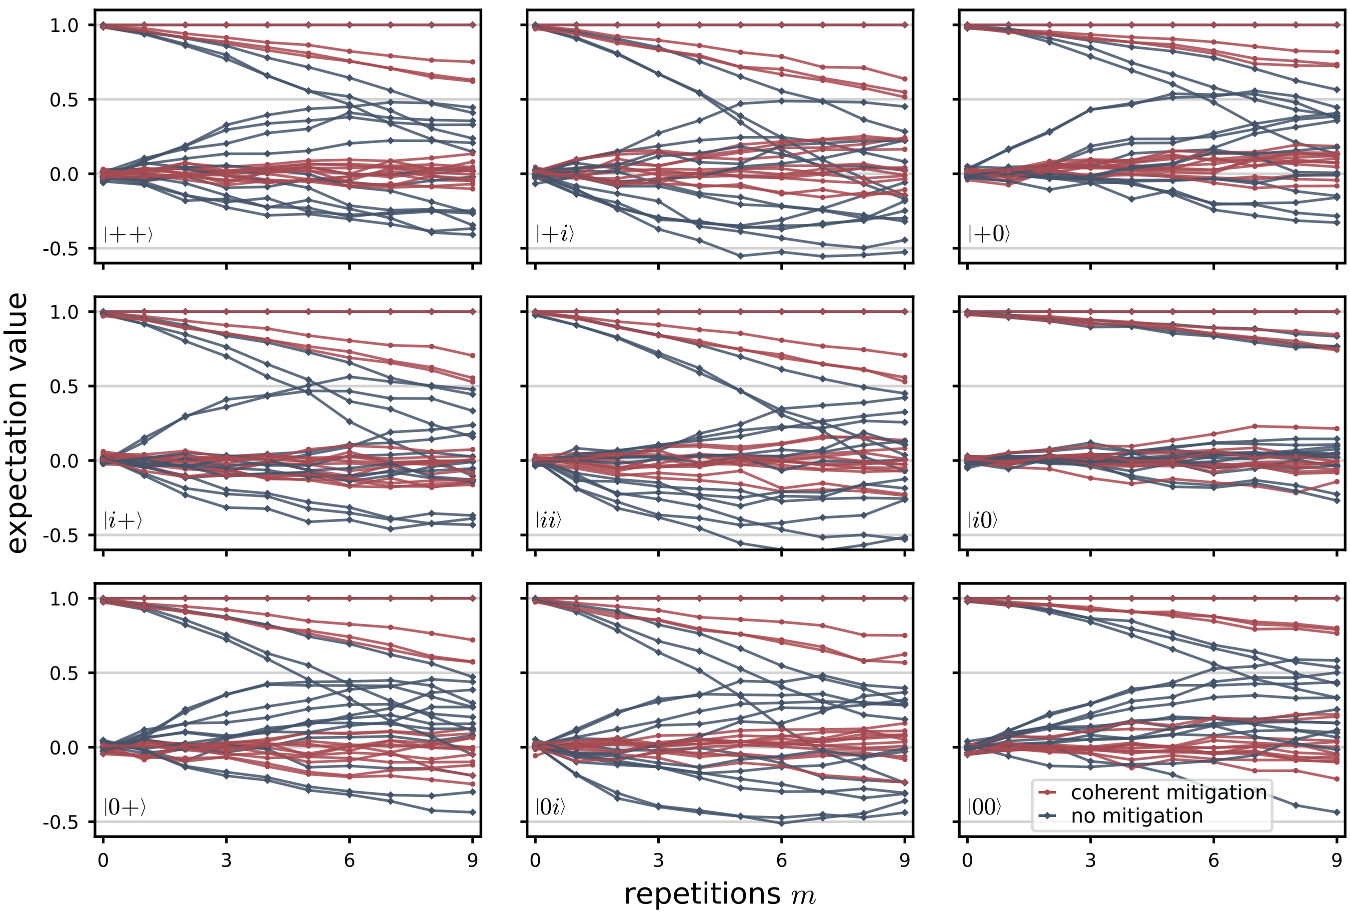

FIG. 8. Complete data set from coherent error mitigation. Collected data of the coherent error mitigated experiment of Fig. 3. The blue curves display the unmitigated results as shown in Fig. 3 (a). Red shows the curves of the experiment with coherent mitigation according to Fig. 3 (b). Each of the 9 different panels exhibits a different initial 2-qubit state.

[1] J. Preskill, Quantum computing in the nisq era and beyond, Quantum **2**, 79 (2018).

[2] C. D. Wilen, S. Abdullah, N. A. Kurinsky, C. Stanford, L. Cardani, G. D'Imperio, C. Tomei, L. Faoro, L. B. Ioffe, C. H. Liu, A. Opremcak, B. G. Christensen, J. L. DuBois, and R. McDermott, Correlated charge noise and relaxation errors in superconducting qubits, Nature **594**, 369 (2021).

[3] S. Krinner, S. Lazar, A. Remm, C. Andersen, N. Lacroix, G. Norris, C. Hellings, M. Gabureac, C. Eichler, and A. Wallraff, Benchmarking Coherent Errors in Controlled-Phase Gates due to Spectator Qubits, Phys. Rev. Appl. **14**, 024042 (2020).

[4] S. Resch and U. R. Karpuzcu, Benchmarking quantum computers and the impact of quantum noise, ACM Comput Surv (CSUR) **54**, 1 (2021).

[5] E. Nielsen, K. Rudinger, T. Proctor, A. Russo, K. Young, and R. Blume-Kohout, Probing quantum processor performance with pygsti, Quantum Sci. Technol. **5**, 044002 (2020).

[6] W. Jiang, J. Xiong, and Y. Shi, A co-design framework of neural networks and quantum circuits towards quantum advantage, Nat. Commun. **12**, 1 (2021).

[7] R. Harper, S. T. Flammia, and J. J. Wallman, Efficient learning of quantum noise, Nat. Phys. **16**, 1184 (2020).

[8] K. Temme, S. Bravyi, and J. M. Gambetta, Error mitigation for short-depth quantum circuits, Phys. Rev. Lett. **119**, 180509 (2017).

[9] J. Wallman, C. Granade, R. Harper, and S. T. Flammia, Estimating the coherence of noise, New J. Phys. **17**, 113020 (2015).

[10] J. T. Muhonen, J. P. Dehollain, A. Laucht, F. E. Hudson, R. Kalra, T. Sekiguchi, K. M. Itoh, D. N. Jamieson, J. C. McCallum, A. S. Dzurak, and A. Morello, Storing quantum information for 30 seconds in a nanoelectronic device, Nat. Nanotechnol. **9**, 986 (2014).

[11] M. A. Fogarty, M. Veldhorst, R. Harper, C. H. Yang, S. D. Bartlett, S. T. Flammia, and A. S. Dzurak, Nonexponential fidelity decay in randomized benchmarking with low-frequency noise, Phys. Rev. A **92**, 022326 (2015).

[12] D. Suter and G. A. Álvarez, Colloquium: Protecting quantum information against environmental noise, Rev. Mod. Phys. **88**, 041001 (2016).

[13] H. Ball, W. D. Oliver, and M. J. Biercuk, The role of master clock stability in quantum information process-

ing, npj Quantum Inf. **2**, 1 (2016).

[14] J. Yoneda, K. Takeda, T. Otsuka, T. Nakajima, M. R. Delbecq, G. Allison, T. Honda, T. Kodera, S. Oda, Y. Hoshi, *et al.*, A quantum-dot spin qubit with coherence limited by charge noise and fidelity higher than 99.9%, Nat. Nanotechnol. **13**, 102 (2018).

[15] I. Chuang and M. A. Nielsen, Prescription for experimental determination of the dynamics of a quantum black box, J. Mod. Opt. **44**, 2455 (1997).

[16] J. Poyatos, J. I. Cirac, and P. Zoller, Complete characterization of a quantum process: the two-bit quantum gate, Phys. Rev. Lett. **78**, 390 (1997).

[17] R. Blume-Kohout, J. K. Gamble, E. Nielsen, J. Mizrahi, J. D. Sterk, and P. Maunz, Robust, self-consistent, closed-form tomography of quantum logic gates on a trapped ion qubit, arXiv preprint arXiv:1310.4492 (2013).

[18] S. T. Merkel, J. M. Gambetta, J. A. Smolin, S. Poletto, A. D. Córcoles, B. R. Johnson, C. A. Ryan, and M. Steffen, Self-consistent quantum process tomography, Phys. Rev. A **87**, 062119 (2013).

[19] H. Häffner, W. Hänsel, C. F. Roos, J. Benhelm, D. Chekal kar, M. Chwalla, T. Körber, U. D. Rapol, M. Riebe, P. O. Schmidt, C. Becher, O. Gühne, W. Dür, and R. Blatt, Scalable multiparticle entanglement of trapped ions, Nature **438**, 643 (2005).

[20] J. Emerson, R. Alicki, and K. Życzkowski, Scalable noise estimation with random unitary operators, J. Opt. B: Quantum Semiclassical Opt. **7**, S347 (2005).

[21] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland, Randomized benchmarking of quantum gates, Phys. Rev. A **77**, 012307 (2008).

[22] E. Magesan, J. M. Gambetta, and J. Emerson, Scalable and robust randomized benchmarking of quantum processes, Phys. Rev. Lett. **106**, 180504 (2011).

[23] M. Sarovar, T. Proctor, K. Rudinger, K. Young, E. Nielsen, and R. Blume-Kohout, Detecting crosstalk errors in quantum information processors, Quantum **4**, 321 (2020).

[24] E. van den Berg, Z. K. Minev, A. Kandala, and K. Temme, Probabilistic error cancellation with sparse pauli–lindblad models on noisy quantum processors, Nat. Phys. , 1 (2023).

[25] M. L. Dahlhauser and T. S. Humble, Benchmarking characterization methods for noisy quantum circuits, arXiv preprint arXiv:2201.02243 (2022).

[26] O. Kern, G. Alber, and D. L. Shepelyansky, Quantum error correction of coherent errors by randomization, EPJ D **32**, 153 (2005).

[27] J. J. Wallman and J. Emerson, Noise tailoring for scalable quantum computation via randomized compiling, Phys. Rev. A **94**, 052325 (2016).

[28] J. Emerson, M. Silva, O. Moussa, C. Ryan, M. Laforest, J. Baugh, D. G. Cory, and R. Laflamme, Symmetrized Characterization of Noisy Quantum Processes, Science **317**, 1893 (2007).

[29] O. Moussa, M. P. da Silva, C. A. Ryan, and R. Laflamme, Practical Experimental Certification of Computational Quantum Gates Using a Twirling Procedure, Phys. Rev. Lett. **109**, 070504 (2012).

[30] M. R. Geller and Z. Zhou, Efficient error models for fault-tolerant architectures and the pauli twirling approximation, Phys. Rev. A **88**, 012314 (2013).

[31] L. Viola and S. Lloyd, Dynamical suppression of decoherence in two-state quantum systems, Phys. Rev. A **58**, 2733 (1998).

[32] L. Viola, E. Knill, and S. Lloyd, Dynamical Decoupling of Open Quantum Systems, Phys. Rev. Lett. **82**, 2417 (1999).

[33] K. Khodjasteh and D. A. Lidar, Fault-Tolerant Quantum Dynamical Decoupling, Phys. Rev. Lett. **95**, 180501 (2005).

[34] K. Khodjasteh and L. Viola, Dynamically Error-Corrected Gates for Universal Quantum Computation, Phys. Rev. Lett. **102**, 080501 (2009).

[35] D. Buterakos, S. Das Sarma, and E. Barnes, Geometrical Formalism for Dynamically Corrected Gates in Multiqubit Systems, PRX Quantum **2**, 010341 (2021).

[36] S. Majumder, C. G. Yale, T. D. Morris, D. S. Lobser, A. D. Burch, M. N. H. Chow, M. C. Revelle, S. M. Clark, and R. C. Pooser, Characterizing and mitigating coherent errors in a trapped ion quantum processor using hidden inverses, Quantum **7**, 1006 (2023).

[37] M. J. Biercuk, H. Uys, A. P. VanDevender, N. Shiga, W. M. Itano, and J. J. Bollinger, Optimized dynamical decoupling in a model quantum memory, Nature **458**, 996 (2009).

[38] T. van der Sar, Z. H. Wang, M. S. Blok, H. Bernien, T. H. Taminiau, D. M. Toyli, D. A. Lidar, D. D. Awschalom, R. Hanson, and V. V. Dobrovitski, Decoherence-protected quantum gates for a hybrid solid-state spin register, Nature **484**, 82 (2012).

[39] J. P. Kestner, X. Wang, L. S. Bishop, E. Barnes, and S. Das Sarma, Noise-Resistant Control for a Spin Qubit Array, Phys. Rev. Lett. **110**, 140502 (2013).

[40] D. C. Murphy and K. R. Brown, Controlling error orientation to improve quantum algorithm success rates, Phys. Rev. A **99**, 032318 (2019).

[41] B. Zhang, S. Majumder, P. H. Leung, S. Crain, Y. Wang, C. Fang, D. M. Debroy, J. Kim, and K. R. Brown, Hidden Inverses: Coherent Error Cancellation at the Circuit Level, Phys. Rev. Appl. **17**, 034074 (2022).

[42] Y. Li and S. C. Benjamin, Efficient Variational Quantum Simulator Incorporating Active Error Minimization, Phys. Rev. X **7**, 021050 (2017).

[43] S. Endo, S. C. Benjamin, and Y. Li, Practical Quantum Error Mitigation for Near-Future Applications, Phys. Rev. X **8**, 031027 (2018).

[44] A. Kandala, K. Temme, A. D. Córcoles, A. Mezzacapo, J. M. Chow, and J. M. Gambetta, Error mitigation extends the computational reach of a noisy quantum processor, Nature **567**, 491 (2019).

[45] Y. Kim, C. J. Wood, T. J. Yoder, S. T. Merkel, J. M. Gambetta, K. Temme, and A. Kandala, Scalable error mitigation for noisy quantum circuits produces competitive expectation values, Nat. Phys. , 1 (2023).

[46] IBM Quantum. https://quantum-computing.ibm.com/ (2021).

[47] E. Magesan, D. Puzzuoli, C. E. Granade, and D. G. Cory, Modeling quantum noise for efficient testing of fault-tolerant circuits, Phys. Rev. A **87**, 012324 (2013).

[48] Y. R. Sanders, J. J. Wallman, and B. C. Sanders, Bounding quantum gate error rate based on reported average fidelity, New J. Phys. **18**, 012002 (2015).

[49] S. J. Beale, J. J. Wallman, M. Gutiérrez, K. R. Brown, and R. Laflamme, Quantum error correction decoheres noise, Phys. Rev. Lett. **121**, 190501 (2018).

[50] J. K. Iverson and J. Preskill, Coherence in logical quantum channels, New J. Phys. **22**, 073066 (2020).
[51] E. Huang, A. C. Doherty, and S. Flammia, Performance of quantum error correction with coherent errors, Phys. Rev. A **99**, 022313 (2019).
[52] D. Gottesman, The heisenberg representation of quantum computers, arXiv preprint quant-ph/9807006 (1998).
[53] 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).
[54] G. Cenedese, G. Benenti, and M. Bondani, Correcting coherent errors by random operation on actual quantum hardware, Entropy **25**, 324 (2023).
[55] M. Ware, G. Ribeill, D. Riste, C. A. Ryan, B. Johnson, and M. P. Da Silva, Experimental pauli-frame randomization on a superconducting qubit, Phys. Rev. A **103**, 042604 (2021).
[56] M. Gutiérrez, L. Svec, A. Vargo, and K. R. Brown, Approximation of realistic errors by clifford channels and pauli measurements, Phys. Rev. A **87**, 030302 (2013).
[57] S. J. Devitt, W. J. Munro, and K. Nemoto, Quantum error correction for beginners, Rep. Prog. Phys. **76**, 076001 (2013).
[58] R. Kueng, D. M. Long, A. C. Doherty, and S. T. Flammia, Comparing experiments to the fault-tolerance threshold, Phys. Rev. Lett. **117**, 170502 (2016).
[59] Y. Suzuki, K. Fujii, and M. Koashi, Efficient simulation of quantum error correction under coherent error based on the nonunitary free-fermionic formalism, Phys. Rev. Lett. **119**, 190503 (2017).
[60] D. Greenbaum, Introduction to Quantum Gate Set Tomography, arXiv preprint arXiv:1509.02921 (2015).
[61] C. A. Riofrío, D. Gross, S. T. Flammia, T. Monz, D. Nigg, R. Blatt, and J. Eisert, Experimental quantum compressed sensing for a seven-qubit system, Nat. Commun. **8**, 15305 (2017).
[62] H.-Y. Huang, R. Kueng, and J. Preskill, Efficient Estimation of Pauli Observables by Derandomization, Phys. Rev. Lett. **127**, 030503 (2021).
[63] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, Nat. Phys. **16**, 1050 (2020).
[64] F. Bloch, Nuclear induction, Phys. Rev. **70**, 460 (1946).
[65] D. Bultrini, M. H. Gordon, E. López, and G. Sierra, Simple mitigation strategy for a systematic gate error in ibmq, arXiv preprint arXiv:2012.00831 (2020).
[66] 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).
[67] M. S. Blok, V. V. Ramasesh, T. Schuster, K. O'Brien, J. M. Kreikebaum, D. Dahlen, A. Morvan, B. Yoshida, N. Y. Yao, and I. Siddiqi, Quantum Information Scrambling on a Superconducting Qutrit Processor, Phys. Rev. X **11** (2021).
[68] A. Morvan, V. Ramasesh, M. Blok, J. Kreikebaum, K. O'Brien, L. Chen, B. Mitchell, R. Naik, D. Santiago, and I. Siddiqi, Qutrit Randomized Benchmarking, Phys. Rev. Lett. **126**, 210504 (2021).
[69] N. Goss, A. Morvan, B. Marinelli, B. K. Mitchell, L. B. Nguyen, R. K. Naik, L. Chen, C. Jünger, J. M. Kreikebaum, D. I. Santiago, J. J. Wallman, and I. Siddiqi, High-fidelity qutrit entangling gates for superconducting circuits, Nat. Commun. **13**, 10.1038/s41467-022-34851-z (2022).
[70] S. Cao, D. Lall, M. Bakr, G. Campanaro, S. Fasciati, J. Wills, V. Chidambaram, B. Shteynas, I. Rungger, and P. Leek, Efficient qutrit gate-set tomography on a transmon, arXiv preprint arXiv:2210.04857 (2022).
[71] N. Goss, S. Ferracin, A. Hashim, A. Carignan-Dugas, J. M. Kreikebaum, R. K. Naik, D. I. Santiago, and I. Siddiqi, Extending the Computational Reach of a Superconducting Qutrit Processor, arXiv preprint arXiv:2305.16507 (2023).
[72] M. Gutiérrez, C. Smith, L. Lulushi, S. Janardan, and K. R. Brown, Errors and pseudothresholds for incoherent and coherent noise, Phys. Rev. A **94**, 042338 (2016).
[73] D. Greenbaum and Z. Dutton, Modeling coherent errors in quantum error correction, Quantum Sci. Technol. **3**, 015007 (2017).
[74] D. M. Debroy, M. Li, M. Newman, and K. R. Brown, Stabilizer slicing: coherent error cancellations in low-density parity-check stabilizer codes, Phys. Rev. Lett. **121**, 250502 (2018).
[75] J. Berberich, D. Fink, and C. Holm, Robustness of quantum algorithms against coherent control errors, arXiv preprint arXiv:2303.00618 (2023).
[76] Z. Cai, X. Xu, and S. C. Benjamin, Mitigating coherent noise using Pauli conjugation, npj Quantum Inf. **6**, 1 (2020).
[77] J. P. Barnes, C. J. Trout, D. Lucarelli, and B. Clader, Quantum error-correction failure distributions: Comparison of coherent and stochastic error models, Phys. Rev. A **95**, 062338 (2017).
[78] P. Iyer and D. Poulin, A small quantum computer is needed to optimize fault-tolerant protocols, Quantum Sci. Technol. **3**, 030504 (2018).
[79] J. J. Wallman and S. T. Flammia, Randomized benchmarking with confidence, New J. Phys. **16**, 103032 (2014).
[80] S. Bravyi, M. Englbrecht, R. König, and N. Peard, Correcting coherent errors with surface codes, npj Quantum Inf. **4**, 1 (2018).
[81] P. Iyer and D. Poulin, Hardness of Decoding Quantum Stabilizer Codes, IEEE Trans. Inf. Theory **61**, 5209 (2015).
[82] X. Xue, B. D'Anjou, T. F. Watson, D. R. Ward, D. E. Savage, M. G. Lagally, M. Friesen, S. N. Coppersmith, M. A. Eriksson, W. A. Coish, and L. M. Vandersypen, Repetitive Quantum Nondemolition Measurement and Soft Decoding of a Silicon Spin Qubit, Phys. Rev. X **10**, 021006 (2020).
[83] C. A. Pattison, M. E. Beverland, M. P. da Silva, and N. Delfosse, Improved quantum error correction using soft information, arXiv preprint arXiv:2107.13589 (2021).
[84] N. Raveendran, N. Rengaswamy, A. K. Pradhan, and B. Vasić, Soft Syndrome Decoding of Quantum LDPC Codes for Joint Correction of Data and Syndrome Errors, in *2022 IEEE International Conference on Quantum Computing and Engineering (QCE)* (2022) pp. 275–281.
[85] D. M. Debroy, L. Egan, C. Noel, A. Risinger, D. Zhu, D. Biswas, M. Cetina, C. Monroe, and K. R. Brown, Optimizing Stabilizer Parities for Improved Logical Qubit Memories, Phys. Rev. Lett. **127**, 240501 (2021).
[86] Y. Ouyang, Avoiding coherent errors with rotated concatenated stabilizer codes, npj Quantum Inf. **7**, 1 (2021).