


Performance Optimization for Drift-Robust Fidelity Improvement of Two-Qubit Gates

G.A.L. White,¹ C.D. Hill,^{1,2} and L.C.L. Hollenberg^{1,*}

¹*School of Physics, University of Melbourne, Parkville, Victoria 3010, Australia*

²*School of Mathematics and Statistics, University of Melbourne, Parkville, Victoria 3010, Australia*

 (Received 23 January 2020; revised 22 September 2020; accepted 6 November 2020; published 13 January 2021)

Quantum-system characterization techniques represent the front line in the identification and mitigation of noise in quantum computing but can be expensive in terms of quantum resources and time to repeatedly employ. Another challenging aspect is that the parameters governing the performance of various operations tend to drift over time and monitoring these is hence a difficult task. One of the most promising characterization techniques, gate-set tomography (GST), provides a self-consistent estimate of the completely positive trace-preserving (CPTP) maps for a complete set of gates, as well as preparation and measurement operators. We develop a method for performance optimization seeded by tomography (POST), which couples the power of GST with a classical optimization routine to achieve a consistent gate improvement in just a short number of steps within a given calibration cycle. By construction, the POST procedure finds the best available gate operation given the hardware and is therefore robust to the effects of drift. Further, in comparison to other quantum error-mitigation techniques, it builds upon a one-time application of GST. To demonstrate the performance of this method on a real quantum computer, we map out the operations of six qubit pairs on the superconducting *ibmq_poughkeepsie* quantum device. Under the restriction of logical-only control, we monitor the performance of the POST approach on a chosen controlled-NOT (CNOT) gate over a period of six weeks. In this time, we achieve a consistent improvement in gate fidelity, averaging a fidelity increase of 21.1% as measured by randomized benchmarking. The POST approach should find wide applicability as it is hardware agnostic and can be applied at the upper logical level or at a deeper pulse-control level.

DOI: [10.1103/PhysRevApplied.15.014023](https://doi.org/10.1103/PhysRevApplied.15.014023)

I. INTRODUCTION

The nascent field of quantum computing has seen an emergence of many experimentally realized small-scale devices in recent years, most notably in superconducting qubit systems and trapped atomic spins [1–5]. Different architectures have achieved high-fidelity one- and two-qubit gates, as well as the construction of multiqubit entangled states [2,6–12]. Despite this progress, current hardware cannot yet demonstrate large-scale topological quantum error correction below threshold and there are many significant obstacles to overcome before qubit numbers can be scaled up to useful levels. Quantum computers presently face the challenge of imperfections in state preparation, measurement errors, and erroneous logical gates. Before improvement can be achieved, comprehensive characterization techniques are essential in mapping where deficiencies lie.

Noise on real quantum devices is challenging to understand quantitatively. In particular, it is difficult to isolate

device behavior given the tendency of noise and system parameters to *drift* [13–15]. This is one of the many barriers facing the improvement of quantum hardware. Common characterization techniques such as quantum process tomography (QPT) [16] and randomized benchmarking (RB) [17] offer insights into the quality of a qubit but suffer from respective self-consistency and limited information issues. Gate-set tomography (GST), introduced in Refs. [18,19], provides a relatively novel method in which the preparation, gate, and measurement operations can be implemented in conjunction with each other and separately characterized. The results can be highly accurate, but with the trade-off that a large number of experiments are required to provide the data. The analysis itself is also computationally demanding. As a consequence, there are relatively few examples of two-qubit GST carried out experimentally in the literature [20–22]. Importantly, GST puts a lens on otherwise inaccessible quantities such as the average gate infidelity, dominant noise channels, and its diamond norm. Here, we demonstrate how the comprehensive information provided by GST can be leveraged to reduce the coherent noise in quantum operations. We

*lloydch@unimelb.edu.au

consider both the diamond norm of the gate, which may be extracted from the initial estimate, and the randomized benchmarking infidelity, which may be efficiently estimated in the tune-up procedure.

Two-qubit gates are the most significant source of error in many quantum circuits and so minimization of their infidelity is critical to the performance of quantum algorithms. In this paper, we develop a method for performance optimization seeded by tomography (POST) to consistently improve two-qubit CNOT gates based on a hybrid quantum-classical approach. We characterize the bare two-qubit logic gate using GST, find the optimal corrective parameters for bookend single-qubit unitaries, and then use these as a seed to the Nelder-Mead algorithm in order to find the best improvement for a given calibration cycle. We consider two regimes of control: corrective gates acting solely on the control qubit and corrective gates acting both on the control and target qubits. Following the one-time overhead of GST, each daily optimization is performed in a small (<150) number of single-length RB experiments to overcome any drift that has occurred. This is similar to the use of RB as an objective function in Ref. [23] but with the incorporation of the more detailed information provided by GST. We discuss later the practical advantages to seeding this procedure with the initial GST estimates. We test the method on the *ibmq_poughkeepsie* quantum device, a 20-qubit transmon device with a quantum volume of 16, CNOT error rates typically ranging between 1% and 5%, and calibrations usually performed once per day. In particular, we measure the CNOT gate under consideration to have an RB infidelity r_{CNOT} of 0.0337, a diamond norm $\mathcal{E}_{\diamond\text{CNOT}}$ of 0.0683, and a unitarity (as defined in Ref. [24]) of 0.955. The computational costliness of computing confidence factories on two-qubit GST means that we do not have error bars on these values. We discuss in Sec. II the reliability of the GST estimates and how the procedure is relatively insensitive to these estimates. The coupling of the GST seed with classical optimization is successful at improving the gate. When tested on an experimental device, we find that the POST approach is effective even weeks after the initial characterization. At approximately 0.5 ms per shot, this procedure requires about 5 s per shot of device time for the initial GST experiment and less than 0.2 s per shot for the subsequent tune-ups. As such, this procedure could be incorporated into the daily calibration routines of these devices. The hybrid technique brings the gate as close as possible to its target, up to the hardware limit, but the actual effectiveness depends on the level of control afforded. Basic calibration routines tend to be amplitude optimizations in Rabi experiments—however, the calibration routine in this particular instance is opaque to us. In expanding the calibration model to a full process matrix, it appears as though a logical tune-up to bring a quantum gate as close as possible to its target can be an effective and efficient calibration

technique. Further, this procedure can detect and correct for unitary errors that do not add coherently, which can be a limitation of some detection schemes [25]. Although GST has previously been proposed as part of a quantum error-mitigation protocol in Refs. [21,26], we emphasize the need to avoid repeated application of GST in any gate-improvement techniques, owing to its extremely high experimental and computational overheads.

We perform our experiments on an IBM Cloud-based quantum computer with only logical-level control. As a consequence, corrections to the CNOT gate can only be made through single-qubit gate corrections, which themselves are erroneous. With control of the CNOT pulse scheme, we anticipate that the corrections are likely to be more effective.

In addition to the testing of the POST gate-improvement scheme on a specific two-qubit case, we also conduct two-qubit GST experiments on six separate CNOT gates as an investigation into the performance and types of noise that occur on real superconducting devices. Understanding the real noise that occurs on devices is important for several reasons: it can help to inform future characterization, which results in these procedures being less computationally expensive; it can help us to understand noise channels, which is important for quantum error correction [27]; and it can help to identify hardware issues up on a real machine for better future implementation [1]. We present these results, as well as a theoretical evaluation of the effectiveness of our technique on the additional qubit pairs.

II. AN OVERVIEW OF GATE-SET TOMOGRAPHY

GST is a hardware-agnostic method of characterizing quantum operations. This section provides a brief overview of its methodology—for a comprehensive guide to the techniques involved, see Ref. [1,28]. In this work, we operate in the Pauli-transfer-matrix (PTM) representation of quantum channels. The matrices in this representation are mappings of the Stokes vector of a given density matrix. For some channel Λ , the action of which on a density matrix ρ has a Kraus decomposition $\Lambda(\rho) = \sum_i K_i \rho K_i^\dagger$, its matrix representation R_Λ is given by

$$(R_\Lambda)_{ij} = \text{Tr}[P_i \Lambda(P_j)], \quad (1)$$

where the $P_{i,j}$ refer to the normalized ordered set of Pauli matrices spanning the d^2 -dimensional space of bounded linear operators on some Hilbert space \mathcal{H}_d . That is, the n -qubit basis is the set

$$\mathcal{P}_n = \{I/\sqrt{2}, X/\sqrt{2}, Y/\sqrt{2}, Z/\sqrt{2}\}^{\otimes n}. \quad (2)$$

The PTM representation is therefore a mapping of eigenvectors of the Pauli matrices. Application of the superoperator formalism makes this picture more convenient. Here,

density matrices are represented as d^2 -dimensional vectors $|\rho\rangle\rangle$ on a Hilbert-Schmidt space with inner product $\langle\langle\rho_1|\rho_2\rangle\rangle = \text{Tr}(\rho_1^\dagger\rho_2)$. This allows the action of quantum channels to be given by ordinary matrix multiplication. The k th component of these vectors is equal to $\text{Tr}(P_k\rho)$. The operational action is given by $|\Lambda(\rho)\rangle\rangle = R_\Lambda|\rho\rangle\rangle$ and the map composition by $R_{\Lambda_2\circ\Lambda_1} = R_{\Lambda_2}R_{\Lambda_1}$.

The only experimentally accessible quantities in a laboratory are probabilities of measurement outcomes. For example, after a sequence of quantum operations G , many measurements are taken in order to form an estimate of

$$\langle\langle E|G|\rho\rangle\rangle \quad (3)$$

for some preparation Hilbert-Schmidt vector $|\rho\rangle\rangle$ and some measurement effect $\langle\langle E|$. QPT is a technique that provides an estimate for G by the operator acting on a complete set of preparations, followed by a complete set of basis measurements. This style of characterization assumes known state preparation and measurement (SPAM). When the SPAM errors are not negligible, however, QPT can produce gate estimates considerably far away from the true maps [28]. This is particularly an issue since the primary source of error in current quantum computers is SPAM errors. RB curves provide a metric for the quality of a gate operation but this number alone does not elucidate how that gate might be improved. Furthermore, since it is insensitive to SPAM errors, there is little or no information produced about the character of SPAM on a device. The ideal characterization should produce an accurate picture of all quantum channels, including projection operations.

GST aims to fully characterize a complete set of gates. The self-consistency in this method is achieved by including the preparation and measurement operations within the gate set $\{|\rho\rangle\rangle, \langle\langle E|, G_0, G_1, \dots\}$. A set of gates is chosen first as the object of characterization. The only requirement is that these operations (or their compositions) generate an informationally complete set of preparations and measurements. That is, they form a complete basis of the Hilbert-Schmidt space. These SPAM operations are known as *fiducials* and are denoted by $\mathcal{F} = \{F_0, F_1, \dots, F_n\}$. They are optimally selected to form the most mutually distinguishable informationally complete set. Further, a set of gate compositions $\mathcal{G} = \{g_0, g_1, \dots, g_n\}$ is generated. The elements of this set are termed *germs* and each comprises a sequence of operations from the gate set. Gates contain a large number of free parameters and the emergence of errors in these depend on the input state, the sequence of operations, and the basis in which a measurement is made. Germs are chosen from an extensive search such that a possible error in each gate parameter may be amplified and made detectable by the repetition of at least one germ. In order to reduce the statistical error in detecting noise, each germ is repeated L times for many different values of L .

When every possible noisy parameter is made detectable, the germs set is termed *amplificationally complete* [1]. For all values of i, j, k , and L , the experimental data are then collected in the object

$$p_{ijk}^L = \langle\langle E|F_i^{(\text{prep})}g_j^L F_k^{(\text{meas})}|\rho\rangle\rangle. \quad (4)$$

Linear inversion can provide an estimate of each gate at this point but there is no natural way to include physicality constraints or to select the correct gauge. A maximum-likelihood estimate (MLE) is therefore performed to provide the best estimate for the experimental gate set, consistent with Eq. (4), which obeys the conditions of complete positivity and trace preservation. Following this is a final gauge optimization, selecting the gauge that most closely resembles the target gate set.

GST provides a device-independent estimate of the quantum channel. Like all characterization processes, it assumes a model of the physical process taking place. The model assumption is purely Markovian, including zero leakage and weak environmental couplings—that maps are composable and without context dependence. When violations of these model assumptions show up in the data, the non-Markovianity is then flagged. It is for this reason that the GST estimates must be treated with care: in physical terms, a quantum operation cannot necessarily be treated in isolation and a circuit cannot be decomposed into its constituent maps.

III. IMPROVEMENT OF GATE-SET PARAMETERS

The CPTP map of quantum gates given by GST highlights all Markovian errors in the operation. Examples of errors of this nature include control errors, such as axis tilt or errors in pulse shaping, or erroneous coherent rotation of the qubits due to external couplings. Our general goal in this work is to examine how best to feed this characterization back into a quantum device with less noise. We consider only the estimate determined for the CNOT gate. The best method with which to address these noisy parameters depends on the level of control afforded in the device. Given the recent advent of Cloud-based NISQ computers, where users only have restricted control of the device, the POST protocol introduced here makes use only of additional logical operations on the qubits (although we note in the conclusion that the extension to high-level pulse control is possible). At the time of the experiments, this was the only control available to the authors and so the only scheme examined. The natural CNOT gate on these devices is a cross-resonance ZX interaction bookended by local rotations. With pulse-level control, provided by IBM through OpenPulse [29], we speculate that the proposed blueprint could be modified to absorb the corrections into the local pulses—rather than applying logical corrections

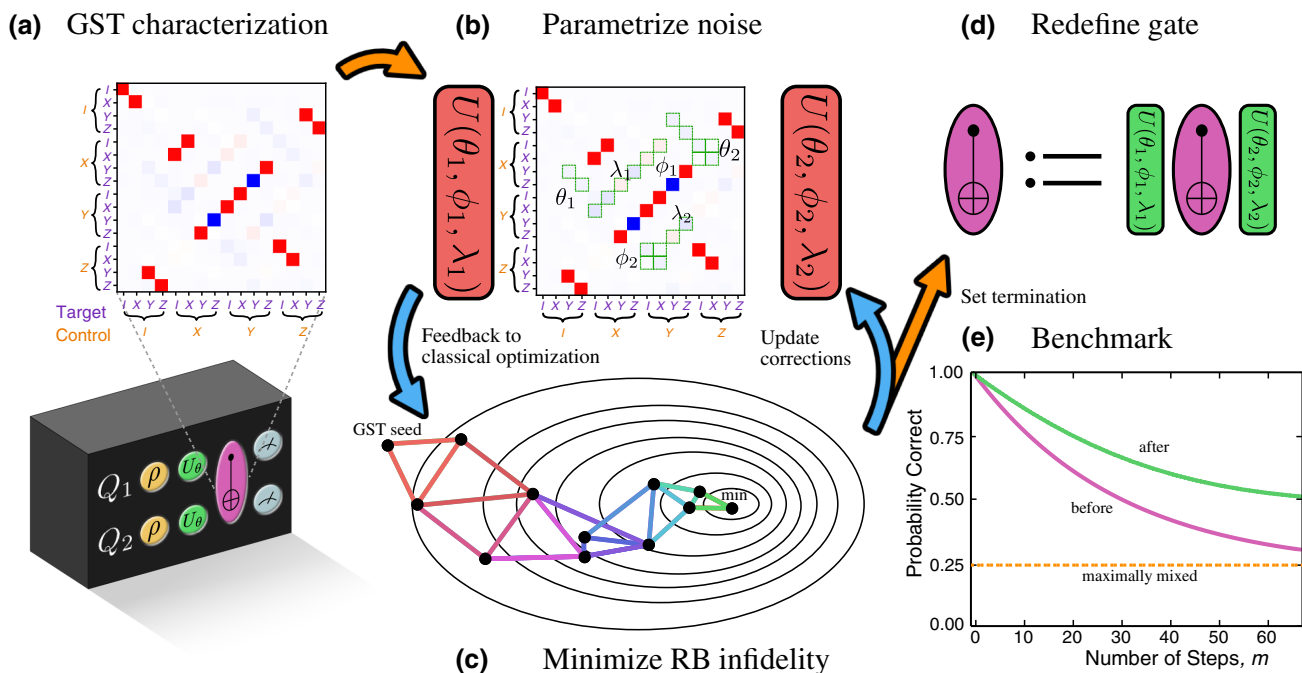


FIG. 1. An overview of the POST approach for CNOT gate-fidelity improvement. (a) The initial characterization of the CNOT Pauli transfer matrix (PTM) using GST. (b) The gate noise is parametrized in the action of arbitrary single-qubit unitaries acting before and after the CNOT gate. The parameter values that minimize the Frobenius distance between the noisy and ideal CNOT gates are then found. (c) These values then seed a Nelder-Mead optimization, where the objective function is the experimental infidelity of a chosen-length randomized benchmarking experiment. With each iteration of POST, a new simplex is chosen and the randomized benchmarking experiment is performed for each vertex, with the single-qubit unitaries taking on the parameters at that point. (d) When the minimum infidelity is found, the CNOT gate is redefined to include the corrective unitaries. (e) The new gate is then fully benchmarked against the native gate, showing a revived fidelity for a much larger number of applications.

around it. We summarize the overall POST procedure in Fig. 1 and describe it in further detail here.

GST provides a means by which errors can be identified but it is not necessarily straightforward to then mitigate their effects. Errors occurring on two-qubit gates such as a CNOT gate tend to be an order of magnitude greater than those of single qubits. From the perspective of logical corrections, it is therefore optimal to address two-qubit noise with the application of single-qubit gates.

Consider a quantum device with an informationally complete set of two-qubit controls. This can be used to conduct a GST experiment in the standard way on a qubit pair. The GST analysis of a CNOT gate produces an estimate for the CPTP map of the gate, designated by \bar{G}_{CX} . This can be decomposed into the product of an ideal CNOT gate, G_{CX} , and some residual noise channel G_{Λ} —the inverse of which is generally unphysical [30]. Previous approaches have typically treated this noise with quasiprobability decompositions [31]. In the case of solely logical control, the nearest physical corrective G_{Λ}^{-1} map may not be within the user’s control set. Importantly, given gate calibration and general hardware drift, a GST estimate’s accuracy quickly expires over time. The sole utilization of GST will require a significant overhead every

time the corrections are implemented. From this, it is clear that GST on its own faces limitations as a practical method of improving gates.

Without direct control of the hardware, correction of all the coherent two-qubit errors will not be possible, since these corrections will contain associated errors equal to or greater in magnitude than the existing ones. Further, cross-resonance errors are not controllable at this level. In this control regime, we propose placing single-qubit corrective gates before and after the native CNOT gate in order to correct as much of the local noise as possible and then optimizing over their parameters. Using a unitary parametrization for the four correction gates U_i ($i \in \{1, 2, 3, 4\}$),

$$U_i(\theta_i, \phi_i, \lambda_i) = \begin{pmatrix} \cos(\theta_i/2) & -e^{i\lambda_i} \sin(\theta_i/2) \\ e^{i\phi_i} \sin(\theta_i/2) & e^{i\lambda_i + i\phi_i} \cos(\theta_i/2) \end{pmatrix}, \quad (5)$$

we propose a superlogical CNOT gate structured as

$$(U_1 \otimes U_2) \cdot \bar{G}_{CX} \cdot (U_3 \otimes U_4), \quad (6)$$

provided that the cumulative error of four single-qubit gates is not greater than one two-qubit gate. In the case of

high single-qubit error rates (or crosstalk between simultaneous gates), a similar approach can be made by applying local corrections exclusively on the control qubit (this allows for broader manipulation than corrections solely on the target qubit), at the expense of more limited noise targeting. That is,

$$(U_1 \otimes \mathbb{I}) \cdot \bar{G}_{CX} \cdot (U_2 \otimes \mathbb{I}). \quad (7)$$

The distinction between these two levels of correction is not immediately obvious. In order to clearly see the difference between corrections on both qubits versus corrections acting solely on the control qubit, we illustrate the addressable parts of the CNOT matrix in green in Figs. 2(c) and 2(d). That is, this indicates the extra accessible dimensions of control offered in either case. Whether that control is necessary in practice will depend on where the noise

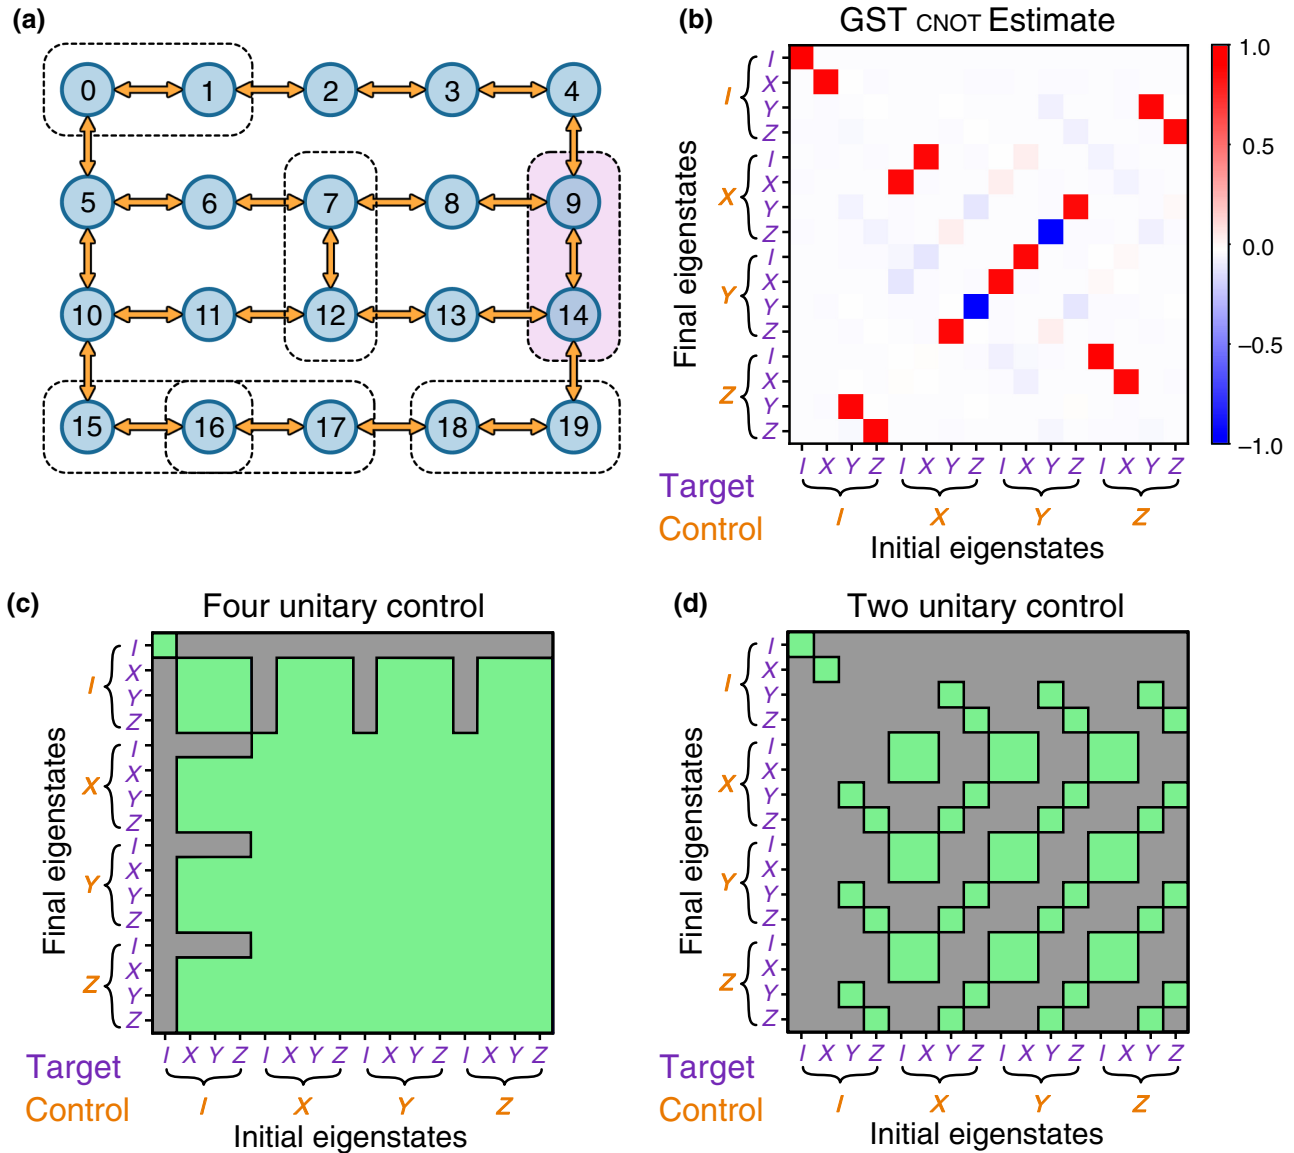


FIG. 2. (a) The *ibmq_poughkeepsie* device layout, showing both the geometry and the connectivity. The dotted lines indicate the six pairs of qubits that are characterized under GST. Qubits 9 and 14 (highlighted in pink) are chosen as the subject for testing POST. (b) The GST-reconstructed estimate of the PTM of the 14–9 CNOT gate under investigation. (c) The green parts of this matrix indicate the controllable parts of the overall map when combining single-qubit operations on the control and target qubits both before and after the native CNOT gate. Because not all elements of the PTM can be *independently* controlled, this does not imply that all errors falling in the green region can be corrected. Under this regime, almost all coherent noise is addressable in principle. (d) The green parts of this matrix indicate the controllable parts of the map when restricted to single-qubit operations on the control qubit before and after the native CNOT gate. This regime is far more restrictive than the previous one due to the trade-off of introducing approximately half of the amount of single-qubit error. These may also be viewed as indicators of the propagation of unitary error.

manifests and the trade-off in introducing further single-qubit error. Note that not all of the elements in the green region can necessarily be independently configured. We call a matrix element “controllable” under a particular set of control operations if it is possible to change that matrix element by varying the parameters of the control. Matrix elements that are uncontrollable under a set of operations cannot be changed at all and are marked in gray. Not all errors that fall in the green region will necessarily be completely correctable.

The ϕ_i , θ_i , λ_i are first selected with a simple optimization to minimize the distance between the corrected gate and the ideal map. At this stage, any norm could be selected. For computational convenience, we choose the Frobenius norm, which also minimizes the average gate infidelity,

$$\|U_1 \otimes U_2 \cdot \bar{G}_{CX} \cdot U_3 \otimes U_4 - G_{CX}\|_F. \quad (8)$$

If the GST estimates of a quantum process are perfect and static with time, then this will be sufficient to yield an improved CNOT gate. However, because GST is only an estimate of a Markovian map within a (generally) non-Markovian system, a simple mathematical minimization will not necessarily result in a physical optimization. This could be either if the noisy parameter values change over time or if the GST estimate is slightly incorrect. Instead, what we propose is a tune-up procedure that optimizes the *performance* of the gate by using GST to identify the neighborhood in which certain parameter values are able to be modified to improve the accuracy of the gate. As such, it is robust to the drift of different noisy parameters and does not rely on the absolute accuracy of the GST estimate. The only assumption is that the noisy channel will remain structurally similar enough to those of the GST estimate that an optimization seeded by GST will bring us back to a better gate than the native operation in few iterations. We define our objective function as the RB infidelity of the gate, in order to make use of the most general metric of performance. The algorithm to implement the POST procedure is as follows:

1. Conduct a series of experiments given by the requirements of GST. Use these to produce an output estimate of the PTM of the gate.
2. Using a classical minimization technique, find the six or 12 parameter values that numerically minimize the Frobenius distance between the superlogical and the ideal CNOT gates, given in Eq. (8). These will be the seed parameter values. The choice of the Frobenius distance is not necessarily special but we elect to use it to make the resulting matrices as similar as possible.
3. Define the objective function to be the average survival probability over an ensemble of circuits for an RB experiment of some length m . Taking step 2 as a newly

defined CNOT gate, compute the objective function for both the native and new CNOT gates as a point of comparison. Using the parameters obtained from GST as a seed, perform an optimization of the CNOT gate by feeding the parameter values into the Nelder-Mead algorithm, where for each vertex of the simplex, the m -length RB infidelity is computed as the objective function.

4. Converge at some predefined level of change. The newly improved CNOT gate is then defined by the composition of the final single-qubit unitary gates on either side of the native CNOT gate.
5. Conduct a full RB experiment to compare the new gate fidelity to the original.

The overview of the procedure is to use the GST estimate as a seed for the Nelder-Mead optimization algorithm, which then works to minimize the infidelity of the CNOT gate by varying the parameters given in Eq. (6) or Eq. (7). We take the infidelity measure to be a set of fixed m -length randomized benchmarking experiments. In a short number of iterations, this locates the optimal corrective rotations to make for a given day. This process is summarized in Fig. 1. The flexibility of the procedure is not only its robustness to drift but generic steps (the optimization procedure, noise parametrization, and objective function) can all be chosen at the user’s discretion. Part of the rationale in performing GST is the ability to specifically determine the level of coherent noise in the gate from the outset, motivating whether it is possible to improve the gate at all. Further, it provides useful information through the determination of the diamond distance between the noisy gate and the ideal. The diamond distance is a means of assessing the distinguishability of two quantum channels. It is a worst-case error rate, taking the largest output trace distance over all possible input matrices, with the identity on an equal-sized ancilla system. That is, for two quantum channels Φ_1 and Φ_2 ,

$$\|\Phi_1 - \Phi_2\|_{\diamond} := \sup_{\rho} \frac{1}{2} \|\Phi_1 \otimes \mathbb{I}_{d^2}(\rho) - \Phi_2 \otimes \mathbb{I}_{d^2}(\rho)\|_1, \quad (9)$$

where $\|\cdot\|_1$ is the trace distance, a common measure of distinguishability between two density matrices. This metric between channels is commonly used in fault-tolerance calculations for quantum error-correcting codes. It is also much more sensitive to coherent error in the gate than the average gate infidelity, which is typically the figure provided from RB curves. By performing GST, we are able to find the global minimum for the gate infidelity, which also globally minimizes the diamond distance. In the next section, we discuss, using our data, how small differences in local minima of the infidelity can correspond to large changes in the diamond norm.

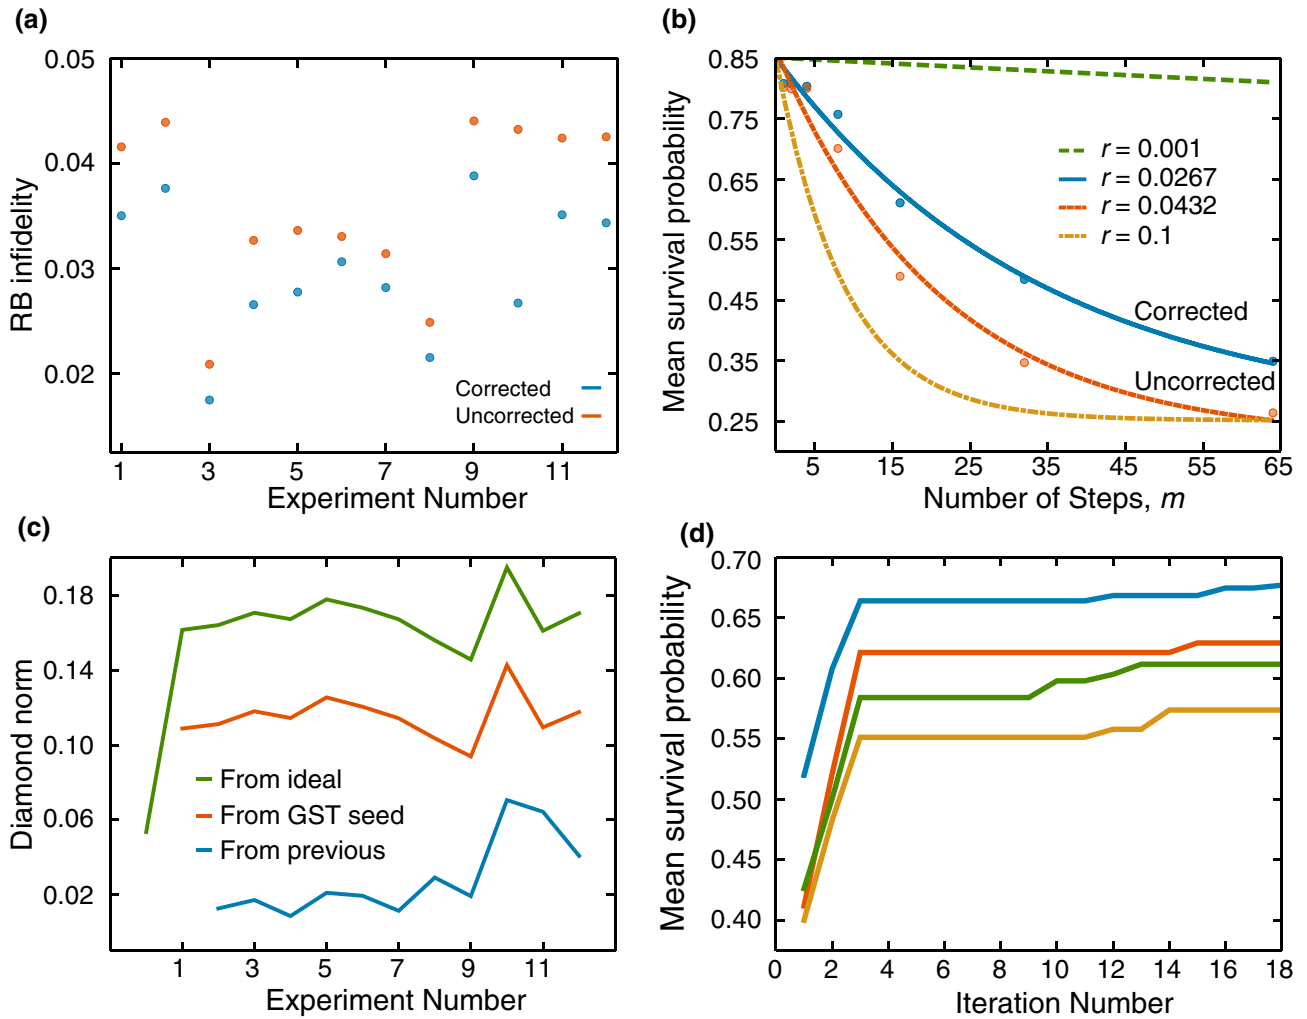


FIG. 3. The efficacy and robustness of the POST protocol—an experimental test on the 14 – 9 CNOT gate. (a) RB values for each experiment taken over the six-week period, both for the corrected and the uncorrected gates. (b) An example RB curve showing our best improved CNOT gate decay versus the native-gate RB experiment for experiment 10 (our most emphatic example of gate improvement). Also shown for comparison are theoretical curves corresponding to 10% and 0.1% error rates. (c) A comparison of the diamond norm of the corrective channel over time: first, with respect to the ideal CNOT gate; second, with respect to the initial GST corrections; and, finally, with respect to each previous experiment. (d) An indication of the convergence speed of the Nelder-Mead gate optimizer for four example runs. This figure shows how the survival probability of the RB experiment increases with each iteration. We observe an initially large increase in the fidelity of the experiment, followed by only small changes thereafter, and zero changes after 18 iterations. The sequence very quickly finds its best gate in a small number of steps.

IV. EXPERIMENTAL IMPLEMENTATION

We test the POST framework for CNOT gate characterization and improvement on the 20-qubit *ibmq_poughkeepsie* superconducting quantum device. Two-qubit GST is performed on six pairs of qubits with the gate set $\{\mathbb{I}, G_{XI}(\pi/2), G_{IX}(\pi/2), G_{YI}(\pi/2), G_{IY}(\pi/2), G_{CX}\}$ up to a germ repetition of $L = 4$ for a total of 10 500 circuits at 8190 shots each. The layout and connectivity of this device are shown in Fig. 2(a). We also indicate the qubits on which experiments are performed. Using the notation “control target” to indicate the physical qubit pair

used, respectively, as the control and target of a CNOT gate, we characterize the gates of qubits 0–1, 12–7, 14–9, 15–16, 16–17, and 18–19. In the estimation procedure, each gate is constrained to be CPTP. We then elect to test the POST procedure on the gate that has most recently been characterized, for which the control is qubit 14 and target qubit 9. For the germ generation and MLE steps, we use the comprehensive open-source PYTHON package pyGSTi, introduced in Ref. [32]. With the tools available, we generate the required germs from our target gate sets and conduct the analysis of our experimental data. The GST estimate from the 14 – 9 qubits, used here in the POST

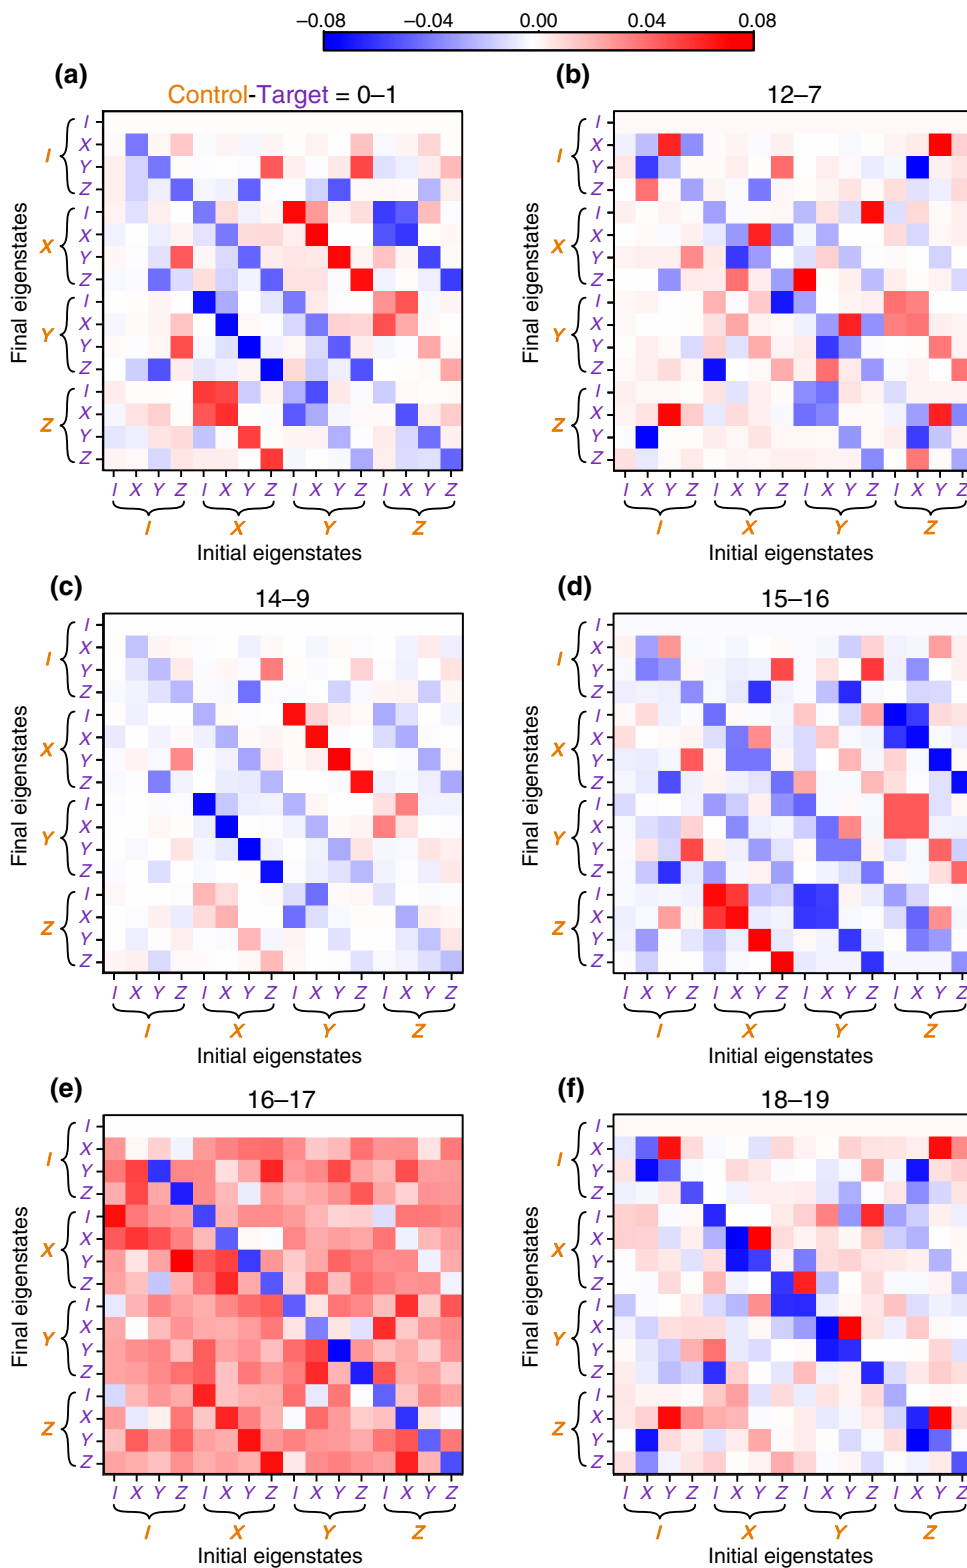


FIG. 4. The matrices here show the GST-estimated PTM representations of the CNOT gate error generator between six pairs of different qubits. The control and target designations correspond to the device layout given in Fig. 2.

tests, is shown in Fig. 2(b). The resulting noise maps for each additional CNOT gate can be found later in this paper, in Fig. 4. In the gate-set estimation procedure, *pyGSTi* flags violation of the Markovian process matrix model, typically indicating some form of non-Markovianity in the

device. Model violation is usually summarized in the form of difference between the observed and expected maximum log-likelihood $\log(\mathcal{L})$ for a given k degrees of freedom. A model-abiding $2\Delta \log(\mathcal{L})$ is expected to be distributed with mean k and standard deviation $\sqrt{2k}$. For length-1

sequences where $k = 737$, our data show $2\Delta \log(\mathcal{L}) = 2861$, for length-2 where $k = 6006$, $2\Delta \log(\mathcal{L}) = 15\,938$ and for length-4 where $k = 26\,901$, $2\Delta \log(\mathcal{L}) = 107\,889$. This suggests that the device is not totally Markovian and that the estimation not totally perfect.

Results Summary—The initial GST analysis of the 14 – 9 CNOT gate took place on March 31, 2019 and its corrective parameters are used as the base vertex for the Nelder-Mead simplex method. The procedure was implemented a total of 12 times over a period of approximately six weeks, corresponding to overlap with approximately 40 different calibration cycles. Figure 3(a) displays a summary of the improvement shown over the native gate with each experiment run for bare RB infidelity r_u and corrected RB infidelity r_c . In each case, both the corrected and uncorrected benchmarking experiments were conducted in the same job submission to avoid any bias in gate drift throughout the day and the tune-up was often conducted shortly after the calibration of the machines. The total average improvement, which we define as $r_u/r_c - 1$ is 21.1%, with a notable outlier of 61.8% in experiment 10. The median observed improvement is 19.1%. In the next section, we discuss how this compares to theoretical figures based on the GST estimates. Figure 3(b) is a comparison RB curve showing the decay of an example improved gate over the native fidelity. For clarity and comparison, we also plot example curves with 10% and 0.1% error rates. Note that this RB number is from the overall curve, which is composed of single and two-qubit gates. For these experiments, this partitions into $r = 3/4 \cdot r_{\text{CNOT}} + 1/4 \cdot r_{\text{single}}$. To reduce the total number of experiments per day, we do not compute multiple curves with different fractions of CNOT and single-qubit gates. Consequently, $r_u/r_c - 1$ is really $(3/4 \cdot r_{u,\text{CNOT}} + 1/4 \cdot r_{\text{single}})/(3/4 \cdot r_{c,\text{CNOT}} + 1/4 \cdot r_{\text{single}}) - 1$, which is a lower bound for the improvement of the CNOT gate. Given that $r_{\text{single}} < r_{\text{CNOT}}$ by about an order of magnitude, we do not expect that the figure differs substantially.

The minimized objective function is the average infidelity of 20 randomly sampled RB circuits, consisting of 16 circuit layers in addition to the preparation and measurement layers. Each circuit is run at 8190 shots to minimize statistical error in the optimization. The use of an RB experiment as the objective function is a flexible metric and can be chosen as the user desires. In principle, the context dependence of a gate may affect the versatility of the improved gate in the sense that gates may affect their system differently depending on the circuit; however, at this stage, RB curves are the most robust assessment of a gate performance of a gate and require the fewest assumptions.

Although the insertion of four local unitaries provides more coherent control, we find empirically on this device that this introduces more error than it eliminates. Consequently, we choose to apply corrections only on the

control qubit, before and after the CNOT gate. This reduces the optimization to six parameters. In the Nelder-Mead method, a dimension-6 simplex with seven vertices is constructed, with the base vertex given by the GST parameters. The objective function is then evaluated for each point. In order to save on computation, we elect to omit the shrink step. After five iterations of no further improvement, we then terminate the algorithm and redefine our gate with the best point.

We use a relatively new form of RB known as *direct randomized benchmarking* (DRB) [33]. In a single circuit, DRB prepares stabilizer states, followed by m randomly selected layers of gates native to that stabilizer, before finally performing a stabilizer measurement to give the success probability. A number of randomly generated circuits can then be used to provide an overall average. The utility of this over Clifford RB is the ability to specify the occurrence of given gates. It also makes an accommodation for future instances of the protocol where the tune-up might be efficiently computed over a larger number of qubits. Here, we randomly generate 20 RB circuits, with CNOT gates composing on average 3/4 of the total circuit. The average probability of success at length m , P_m is then plotted over a series of values for m . These points are then fitted to $P_m = A + Bp^m$ for the fit parameters A , B , and p . The RB number $r = (15/16)(1 - p)$ is then the probability of an error occurring under a stochastic model. We remark here that the RB number is not necessarily a true estimate of the average gate infidelity and can differ under, for example, gate-dependent noise. Here, our characterization produces an initial infidelity estimate of 1.191×10^{-2} for the CNOT gate, whereas the estimated RB number is 3.369×10^{-2} . RB is currently the most efficient method for estimating the quality of a gate and the number most often cited in reporting the performance of hardware. For this reason, we use this as indicative of the gate infidelity, even though the true figures may differ. Indeed, this suggests that IBM may have better gates than their RB curves suggest (for further details, see Refs. [24,34]).

As previously stated, the utility of the GST experiments is both in determining the coherence of the gate noise and in effecting a search that is placed near the global minimum of both the gate infidelity and the diamond norm. For example, using the GST estimate, we simulate the randomized benchmarking tune-up when seeded from the zero vector (that is, the case in which we optimize the RB number without the initial GST seed). In this case, the Nelder-Mead search terminates at a point where $r_{\text{CNOT}} = 0.0254$ and $\mathcal{E}_\diamond = 0.0426$ —whereas the numerical search for the global minimum yields $r_{\text{CNOT}} = 0.0210$ and $\mathcal{E}_\diamond = 0.0286$. Randomly seeded starts are able to find the global minimum but only with the equivalent of between 10 000 and 20 000 RB experiments, which is comparable to

TABLE I. The values of the corrective parameters obtained after a full improvement cycle for each day of experiments, given to the third decimal place. These correspond to the unitary parametrization given in Eq. (5).

Date	θ_1	ϕ_1	λ_1	θ_2	ϕ_2	λ_2
March 31, 2019 (initial GST)	0.046	-1.271	0.480	0.029	0.480	0.393
April 2, 2019	0.116	-1.234	0.536	0.116	0.545	0.403
April 9, 2019	0.130	-1.218	0.461	0.084	0.565	0.475
April 24, 2019	0.116	-1.234	0.552	0.119	0.558	0.415
April 27, 2019	0.148	-1.208	0.502	0.010	0.530	0.427
April 29, 2019	0.089	-1.228	0.523	0.071	0.523	0.436
May 1, 2019	0.218	-1.199	0.552	0.010	0.552	0.365
May 3, 2019	0.089	-1.228	0.523	0.071	0.523	0.436
May 10, 2019	0.096	-1.221	0.530	0.079	0.530	0.443
May 13, 2019	0.141	-1.310	0.497	0.123	0.575	0.488

the initial number of circuits required for GST. Crucially, this approach comes with no evidence of having reached the global minimum, which has large implications for \mathcal{E}_\diamond .

Tracking the Tune-Up—It is prudent to ask how much the corrective values change over time. We stop short of claiming that these values directly quantify drift, since a flat objective function could plausibly entail corrections the value of which changes without much impact on the RB number. However, empirically we find that the previous day’s optimal corrections typically reduce the performance of the gate, rather than improving it. It is necessary, therefore, to perform the optimization, rather than simply applying the results of previous days. Operating on the assumption that this method finds the most appropriate values for the corrective parameters on a particular day, we can use these data to loosely quantify the change in the coherent noise on a real quantum information processor (QIP). The values of the corrective parameters for each experiment are provided in Table I. Moreover, the variables do not all independently affect the final map, meaning that change in the parameter values themselves might obfuscate the fact that the channel overall has not varied much. In this table, the θ values appear to change the most. These values are equivalent to Y rotations, which change the populations of the control qubit. The reason for this is not entirely clear; however, inspection of the error channels in Fig. 4(c) shows that some of the most dominant noise occurs in the X and Z control qubit blocks. This could mean that they are more likely to change with time.

In order to paint a more concrete picture of the effects of the parameters, we study the case of the improved gate as though it comprises three perfect gates. We then compare the diamond distance of this channel from the ideal CNOT gate, from the initial GST seed, and from the previous experiment. Our results are summarized in Fig. 3(c). One of the key assumptions in this method is that the system and its gate noise never change too much from the initial GST seed that the optimizer cannot easily find the best gate

for the day. The distance of the corrective channel on a given day from the GST seed supports this stance: over a six-week period, the mean diamond distance is 0.115, with a standard deviation of 1.2×10^{-2} .

Convergence Speed—Any error-mitigation protocol designed to address drift in a QIP will need to be regularly implemented as part of a tune-up procedure. It is therefore ideal that it should require as few experiments as possible. In Fig. 3(d), we present how each iteration of the Nelder-Mead optimization increases the survival probability of the RB experiment for four sample experiments. In each case, a substantial improvement is found in the first three iterations, beyond which we observe only small fluctuations, and zero change after 18 iterations. Each iteration requires 11 circuits to run and so most of the improvement is found in 22 circuits, with the worst case requiring approximately 200—depending on when one accepts that the algorithm should terminate. Assuming 0.5 ms per shot, POST would take approximately 9 min to converge in the worst-case scenario. Depending on the crosstalk limitations of the device, it could then be run in parallel across all nonoverlapping qubit pairs.

A. Tomography on other qubit pairs

In addition to gate improvement on qubits 14 and 9 on the *ibmq_poughkeepsie*, we also perform GST experiments on five other qubit CNOT-gate pairs on the quantum device. This incorporates a further five experiments of between 9367 and 20 530 circuits at 8190 shots each. We present this data as a case study for the noise that occurs in a real quantum device and we theoretically assess the effectiveness of the POST technique at mitigating the noise. Here, the fidelities of these gates and their structure indicate whether noise can be addressed by the single-qubit unitaries used. Figure 4 shows the noise PTMs for each of the six CNOT gates investigated. Specifically, we present this in the form of error generators \mathcal{L} such that $G_{CX} = e^{\mathcal{L}}G_{CX}$. This presents the noise as though it occurs after the ideal gate, which allows for a helpful—though not

TABLE II. For each GST analysis of the qubit pairs, we present the RB number of the CNOT gate, as well as its diamond distance from the ideal CNOT gate. Further, we show the same figures in the idealized scenario where perfect single-qubit corrective unitaries are employed on the control qubit to remove coherent error. In the final column, we include the predicted RB number for each gate with the complete removal of coherent noise. The diamond norm of almost every entry stands to be significantly improved with the removal of coherent error.

Qubit pair (control target)	r_{CNOT}	$\mathcal{E}_{\diamond\text{CNOT}}$	POST minimum r_{CNOT}	POST minimum $\mathcal{E}_{\diamond\text{CNOT}}$	Stochastic r_{CNOT}
0–1	2.28×10^{-2}	4.74×10^{-2}	1.79×10^{-2}	2.44×10^{-2}	1.78×10^{-2}
12–7	4.52×10^{-2}	9.02×10^{-2}	4.16×10^{-2}	7.77×10^{-2}	2.88×10^{-2}
14–9	3.37×10^{-2}	6.83×10^{-2}	2.10×10^{-2}	2.86×10^{-2}	2.07×10^{-2}
15–16	3.55×10^{-2}	8.10×10^{-2}	2.29×10^{-2}	3.79×10^{-2}	2.21×10^{-2}
16–17	5.56×10^{-2}	8.09×10^{-2}	5.26×10^{-2}	7.70×10^{-2}	5.09×10^{-2}
18–19	2.20×10^{-2}	4.39×10^{-2}	2.16×10^{-2}	4.21×10^{-2}	1.94×10^{-2}

necessarily physically accurate—picture. The control and target numbers given refer, respectively, to the qubits of Fig. 2(a) acting as the control and target of the CNOT gate under characterization. We elect to map out these qubits in order to obtain a relatively uniform sample of the full device geometry. It is instructive to compare these matrix plots with the schematics given in Figs. 2(c) and 2(d), which, respectively, indicate local target or control and sole control rotations. PTM noise, the locations of which are correspondingly indicated in green in the schematics, can be explained as a local rotation occurring either before or after the CNOT gate. For example, the blocklike features prominent in 0–1 and 15–16 make up the landscape of Fig. 2(d), suggesting a rotation of the Z eigenstates of the control qubit into X and Y eigenstates both before and after the CNOT gate. Any noise that falls outside the green region of either schematic can be attributed either to decoherence or cross-resonance errors.

We also use these data to estimate the effectiveness of the POST procedure on the other qubit pairs. Using the GST estimates for each CNOT gate, we compute its RB number as well as its diamond distance to the ideal case. We then investigate how these numbers are ideally minimized both in the case with corrective unitaries on the target qubit and where all coherent noise is removed. This provides an insight into the location of most of the noise in these gates, as well as a partition of the gate infidelity into coherent and stochastic error. Interestingly, in most cases (barring the 12–7 pair), there is little difference between the best ideal gate using two single-qubit corrective unitaries and where a perfect $U(4)$ gate is used to remove all coherent noise. A summary of the data is provided in Table II.

Pulse sequences for the implementation of a CNOT gate typically consist of a local pulse to each qubit, as well as an additional cross-resonance pulse coupling the two. We would expect that implementation of POST with absorbed corrective rotations into the native CNOT

pulse sequence would see a much larger increase in fidelity.

V. DISCUSSION AND CONCLUSIONS

The transition from mathematical maps to physical operations is not always a seamless one. Besides errors in the GST characterization, absent a good method of characterizing non-Markovian behavior, assumptions must be made of weak system-environment correlations, composability of operations, and minimal crosstalk between qubits. The emergence of unexpected behavior from quantum systems means that in-principle operational improvements, such as the direct application of corrections from GST estimates, cannot always be relied upon. We present a general quantum-classical hybrid method that uses the real-life performance of the gate as the feed-forward for corrective updates. The success of the procedure is therefore self-fulfilling.

Randomized benchmarking is a robust method of measuring the Markovian fidelity of a given operation. However, in a system with environmental back-action or context-dependent gates, it is not clear whether the situation will always be as simple as transplanting a redefined gate into a quantum circuit and seeing an increased fidelity in this new context. The markedly better performance of quantum algorithms consisting of these redefined gates remains to be demonstrated and will be the subject of future work. In particular, the POST algorithm would be easily adapted to any characterization technique more inclusive of non-Markovian behavior.

The development of high-fidelity gate hardware is imperative for the field of quantum computing to achieve its ambitious aims. An underrated measure of device quality, however, is *consistency*—the ability to achieve reported minimal error rates again and again despite gradual changes in device parameters and system-environment correlations. In this work, we present a consistent method that combines an initial overhead of GST with a classical

optimization algorithm that delivers an improved two-qubit gate in relatively few experiments. We emphasize that although POST is tested on an IBM Quantum device, it is applicable to any hardware with logical-level control. Furthermore, the method is adaptable to any level of control. The key aspect is to identify the noisy parameters from the GST estimate using the afforded set of device controls. In particular, we would expect to see significantly better results with pulse-level control, wherein instead of separate implementation of the corrective unitaries, they would be absorbed into modifying the CNOT pulse. This would introduce minimal additional gate errors, without any increase in depth.

ACKNOWLEDGMENTS

We are grateful to K. Modi and F. Pollock for valuable conversations, and to D. Broadway for figure advice. This work was supported by the University of Melbourne through the establishment of an IBM Quantum Network Hub at the University. G.A.L.W. is supported by an Australian Government Research Training Program Scholarship. C.D.H. is supported through a Laby Foundation grant at The University of Melbourne. Computational resources are acknowledged from the National Computational Infrastructure (NCI) and the Pawsey Supercomputer Center under the National Computational Merit Allocation Scheme (NCMAS).

-
- [1] R. Blume-Kohout, J. K. Gamble, E. Nielsen, K. Rudinger, J. Mizrahi, K. Fortier, and P. Maunz, Demonstration of qubit operations below a rigorous fault tolerance threshold with gate set tomography, *Nat. Commun.* **8**, 1 (2017).
 - [2] F. Arute, *et al.*, Quantum supremacy using a programmable superconducting processor, *Nature* **574**, 505 (2019).
 - [3] C. Nay, IBM Opens Quantum Computation Center in New York; Brings World's Largest Fleet of Quantum Computing Systems Online, Unveils New 53-Qubit Quantum System for Broad Use, (2019).
 - [4] J. Hsu, Intel's 49-Qubit Chip Shoots for Quantum Supremacy, (2018).
 - [5] W. Zeng, Unsupervised Machine Learning on Rigetti 19Q with Forest 1.2, (2017).
 - [6] R. Barends, J. Kelly, A. Megrant, A. Veitia, D. Sank, E. Jeffrey, T. C. White, J. Mutus, A. G. Fowler, and B. Campbell, *et al.*, Superconducting quantum circuits at the surface code threshold for fault tolerance, *Nature* **508**, 500 (2014).
 - [7] J. Zeuner, A. N. Sharma, M. Tillmann, R. Heilmann, M. Gräfe, A. Moqanaki, A. Szameit, and P. Walther, Integrated-optics heralded controlled-NOT gate for polarization-encoded qubits, *Npj Quantum Inf.* **4**, 13 (2018).
 - [8] T. P. Harty, M. A. Sepiol, D. T. C. Allcock, C. J. Ballance, J. E. Tarlton, and D. M. Lucas, High-Fidelity Trapped-Ion Quantum Logic Using Near-Field Microwaves, *Phys. Rev. Lett.* **117**, 140501 (2016).
 - [9] Y. He, S. K. Gorman, D. Keith, L. Kranz, J. G. Keizer, and M. Y. Simmons, A two-qubit gate between phosphorus donor electrons in silicon, *Nature* **571**, 371 (2019).
 - [10] S. S. Hong, A. T. Papageorge, P. Sivarajah, G. Crossman, N. Didier, A. M. Polloreno, E. A. Sete, S. W. Turkowski, M. P. da Silva, and B. R. Johnson, Demonstration of a parametrically activated entangling gate protected from flux noise, *Phys. Rev. A* **101**, 012302 (2020).
 - [11] M. Kjaergaard, M. E. Schwartz, J. Braumüller, P. Krantz, J. I.-J. Wang, S. Gustavsson, and W. D. Oliver, Superconducting qubits: Current state of play, *Annu. Rev. Condens. Matter Phys.* **11**, 369 (2020).
 - [12] C. E. Bradley, J. Randall, M. H. Abobeih, R. C. Berrevoets, M. J. Degen, M. A. Bakker, M. Markham, D. J. Twitchen, and T. H. Taminiau, A Ten-Qubit Solid-State Spin Register with Quantum Memory up to One Minute, *Phys. Rev. X* **9**, 31045 (2019).
 - [13] P. V. Klimov, J. Kelly, Z. Chen, M. Neeley, A. Megrant, B. Burkett, R. Barends, K. Arya, B. Chiaro, and Yu Chen, *et al.*, Fluctuations of Energy-Relaxation Times in Superconducting Qubits, *Phys. Rev. Lett.* **121**, 090502 (2018).
 - [14] 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—Atomic, Molecular, Opt. Phys.* **92**, 1 (2015).
 - [15] J. M. Chow, J. M. Gambetta, L. Tornberg, J. Koch, L. S. Bishop, A. A. Houck, B. R. Johnson, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf, Randomized Benchmarking and Process Tomography for Gate Errors in a Solid-State Qubit, *Phys. Rev. Lett.* **102**, 090502 (2009).
 - [16] J. F. 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] 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 - Atomic, Molecular, Opt. Phys.* **77**, 12307 (2008).
 - [18] R. Blume-Kohout, J. Gamble, E. Nielsen, J. Mizrahi, J. Sterk, and P. Maunz, Robust, self-consistent, closed-form tomography of quantum logic gates on a trapped ion qubit, arXiv:1310.4492 (2013).
 - [19] 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**, 62119 (2013).
 - [20] P. L. W. Maunz, Characterization of two-qubit quantum gates in Sandia's high optical access surface ion trap (2016).
 - [21] S. Zhang, Y. Lu, K. Zhang, W. Chen, Y. Li, J. N. Zhang, and K. Kim, Error-mitigated quantum gates exceeding physical fidelities in a trapped-ion system, *Nat. Commun.* **11**, 587 (2020).
 - [22] C. Song, J. Cui, H. Wang, J. Hao, H. Feng, and Y. Li, Quantum computation with universal error mitigation on a superconducting quantum processor, *Sci. Adv.* **5**, eaaw5686 (2019).
 - [23] J. Kelly, R. Barends, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. G. Fowler, I. C. Hoi, and

- E. Jeffrey, *et al.*, Optimal Quantum Control Using Randomized Benchmarking, *Phys. Rev. Lett.* **112**, 1 (2014).
- [24] J. Wallman, C. Granade, R. Harper, and S. T. Flammia, Estimating the coherence of noise, *New J. Phys.* **17**, 1 (2015).
- [25] S. Sheldon, L. S. Bishop, E. Magesan, S. Filipp, J. M. Chow, and J. M. Gambetta, Characterizing errors on qubit operations via iterative randomized benchmarking, *Phys. Rev. A* **93**, 012301 (2016).
- [26] S. Endo, S. C. Benjamin, and Y. Li, Practical Quantum Error Mitigation for Near-Future Applications, *Phys. Rev. X* **8**, 031027 (2018).
- [27] T. E. O'Brien, B. Tarasinski, and L. DiCarlo, Density-matrix simulation of small surface codes under current and projected experimental noise, *Npj Quantum Inf.* **3**, 39 (2017).
- [28] D. Greenbaum, Introduction to quantum gate set tomography, arXiv:1509.02921v1 (2015).
- [29] D. C. McKay, T. Alexander, L. Bello, M. J. Biercuk, L. Bishop, J. Chen, J. M. Chow, A. D. Córcoles, D. Egger, and S. Filipp, *et al.*, Qiskit backend specifications for openqasm and openpulse experiments, arXiv:1809.03452 (2018).
- [30] T. F. Jordan, Maps and inverse maps in open quantum dynamics, *Ann. Phys.* **325**, 2075 (2010).
- [31] K. Temme, S. Bravyi, and J. M. Gambetta, Error Mitigation for Short-Depth Quantum Circuits, *Phys. Rev. Lett.* **119**, 1 (2017).
- [32] E. Nielsen, K. Rudinger, T. Proctor, A. Russo, K. Young, and R. Blume-Kohout, Probing quantum processor performance with pyGSTi, arXiv:2002.12476 (2020).
- [33] T. J. Proctor, A. Carignan-Dugas, K. Rudinger, E. Nielsen, R. Blume-Kohout, and K. Young, Direct Randomized Benchmarking for Multiqubit Devices, *Phys. Rev. Lett.* **123**, 30503 (2019).
- [34] T. Proctor, K. Rudinger, K. Young, M. Sarovar, and R. Blume-Kohout, What Randomized Benchmarking Actually Measures, *Phys. Rev. Lett.* **119**, 130502 (2017).