



CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

Mitigating measurement errors in multiqubit experiments

Sergey Bravyi, Sarah Sheldon, Abhinav Kandala, David C. McKay, and Jay M. Gambetta

Phys. Rev. A **103**, 042605 — Published 9 April 2021

DOI: [10.1103/PhysRevA.103.042605](https://doi.org/10.1103/PhysRevA.103.042605)

Mitigating measurement errors in multi-qubit experiments

Sergey Bravyi,¹ Sarah Sheldon,² Abhinav Kandala,¹ David C. McKay,¹ and Jay M. Gambetta¹

¹IBM Quantum, T. J. Watson Research Center, Yorktown Heights, NY 10598

²IBM Quantum, Almaden Research Center, San Jose, CA 95120

Reducing measurement errors in multi-qubit quantum devices is critical for performing any quantum algorithm. Here we show how to mitigate measurement errors by a classical post-processing of the measured outcomes. Our techniques apply to any experiment where measurement outcomes are used for computing expected values of observables. Two error mitigation schemes are presented based on tensor product and correlated Markovian noise models. Error rates parameterizing these noise models can be extracted from the measurement calibration data using a simple formula. Error mitigation is achieved by applying the inverse noise matrix to a probability vector that represents the outcomes of a noisy measurement. The error mitigation overhead, including the the number of measurements and the cost of the classical post-processing, is exponential in ϵn , where ϵ is the maximum error rate and n is the number of qubits. We report experimental demonstration of our error mitigation methods on IBM Quantum devices using stabilizer measurements for graph states with $n \leq 12$ qubits and entangled 20-qubit states generated by low-depth random Clifford circuits.

I. INTRODUCTION

Quantum computing experiments are beginning a shift from few-qubit demonstrations of entangling gates and short quantum circuits to larger multi-qubit quantum algorithms attempting to address practically important computational problems. Although available devices are noisy, error mitigation methods have emerged as a possible near-term solution of the fault-tolerance problem [1–5]. Error mitigation schemes are particularly attractive because they introduce no overhead in terms of the number of qubits and gates. These methods have recently enabled a reliable simulation of shallow quantum circuits on a noisy hardware without resorting to quantum error correcting codes [6, 7]. The key idea behind error mitigation is to combine outcomes of multiple experiments in a way that cancels the contribution of noise to the quantity of interest, see [1–3] for details.

Readout errors introduced by imperfect qubit measurements are often the dominant factor limiting scalability of near-term devices. Here we introduce methods for mitigating readout errors for a class of quantum algorithms, where measurement outcomes are only used for computing mean values of observables. Notable examples of such algorithms are variational quantum eigensolvers [8, 9], quantum machine learning [7], and tomography of entangled states [10]. However, our techniques are not applicable to single-shot measurements. The latter are indispensable to many applications such as quantum teleportation, measurement-based computation, and quantum error correction.

Prior work [11–15] demonstrated that readout errors in quantum devices based on superconducting qubits can be well understood in terms of purely classical noise models. Such models describe a noisy n -qubit measurement by a matrix of transition probabilities A of size $2^n \times 2^n$ such that $A_{y,x}$ is the probability of observing a measurement outcome y provided that the true outcome is x . It is common to simplify the noise model further by assuming that the noise acts independently on each qubit [9, 16].

This defines a tensor product noise model such that A is a tensor product of 2×2 noise matrices. The model depends on $2n$ error rates describing single-qubit readout errors $0 \rightarrow 1$ and $1 \rightarrow 0$. The open-source software package, Qiskit [17], provides a toolkit for calibrating A and mitigating readout errors for either a tensor product or a general noise model.

Although the tensor product noise model is appealingly simple, it leaves aside cross-talk errors encountered in real-world setups. Cross-talk during readout can arise from the underlying qubit-qubit coupling, and spectral overlap of readout resonators with stray couplings or multiplexing [18]. Here we introduce a correlated noise model based on Continuous Time Markov Processes (CTMP). The corresponding noise matrix has the form $A = e^G$, where G is a sum of local operators generating single- and two-qubit readout errors such as $01 \rightarrow 10$, $11 \rightarrow 00$, etc. The model depends on $2n^2$ error rates.

Measurement calibration aims at learning the unknown parameters of a noise model. This is achieved by preparing a set of well-characterized input states (we use the standard basis vectors) and repeatedly performing n -qubit measurements on each input state. We show how to extract parameters of the considered noise models from the calibration data using simple analytic formulas. The number of input states required to calibrate the tensor product and CTMP noise models scales linearly with the number of qubits n .

Once parameters of the noise model are known, error mitigation proceeds by applying the inverse noise matrix to a probability vector that represents the noisy measurement outcomes. We show how to sidestep the explicit computation of the inverse noise matrix, which may be prohibitive for a large number of qubits. Instead, we follow ideas of [1] and cancel errors by sampling noise matrices from a suitable quasi-probability distribution. The total error mitigation overhead, including the number of measurements and the cost of classical post-processing, scales as $e^{4\gamma} \text{poly}(n)$, where γ is a parameter quantifying the noise strength, as formally defined below. For ex-

ample, if each qubit independently undergoes a bit-flip readout error with probability ϵ then $\gamma = n\epsilon$. Thus our methods are practical whenever γ is a small constant. For comparison, all previously studied readout error mitigation techniques require manipulations with probability vectors of size 2^n which limits their scalability. Importantly, both tensor product and CTMP error mitigation methods provide an unbiased estimator of the ideal mean value to be measured. For a fixed number of qubits, the cost of approximating the ideal mean value with a specified error tolerance δ scales as δ^{-2} .

The proposed error mitigation techniques are demonstrated using a 20-qubit IBM Quantum device. The device has a few-percent readout error rate. For small number of qubits ($n \leq 7$), a comparison is performed between the full noise matrix A , the tensor product, and the CTMP noise models. It is found that the CTMP model provides a considerably more accurate approximation of the readout noise. Error-mitigated stabilizer measurements of n -qubit graph states are reported for $n \leq 12$. Finally, we perform error-mitigated stabilizer measurements on 20-qubit states generated by low-depth random Clifford circuits. The classical post-processing associated with the error mitigation, including the measurement calibration and the noise inversion steps, can be done on a laptop in a few-second time frame for each of these experiments. Note that the error mitigation runtime is crucial for VQE-type algorithms [9] where quantum mean value estimation is repeatedly used in a feedback loop.

Let us briefly discuss the previous work. Seif et al [19] considered a system of trapped-ion qubits and the problem of discriminating between different basis states based on data collected by an array of photon detectors. It was shown that a neural network trained on the measurement calibration data can realize a high-fidelity state discrimination [19]. It remains to be seen whether similar methods can address the problem of high-precision mean value estimation considered here. Our techniques build off earlier work from Kandala et al [9] and are most closely related to the work of Sun and Geller [16, 20, 21] who showed how to characterize readout errors in multi-qubit devices. In particular, Ref. [16] introduced a noise model based on single- and two-qubit correlation functions which can be viewed as an analogue of the CTMP model considered here. Ref. [22] elucidated how performance of readout error mitigation methods is affected by the quality of the measurement calibration, number of measurement samples, and the presence of coherent (non-classical) measurement errors, see also [11, 23]. Ref. [24] pointed out a surprising connection between quantum readout error mitigation methods and unfolding algorithms used in high energy physics to cope with a finite resolution of particle detectors [25]. However the classical processing cost of unfolding methods may be prohibitive for multi-qubit experiments since they require manipulations with exponentially large probability vectors. Applications of the noise matrix inversion method in the context of variational quantum algorithms are discussed

in [26].

After completion of the present paper we became aware of a closely related work by Hamilton, Kharazi et al. [27] which developed a scalable characterization of readout errors based on a cumulant expansion. The latter can be viewed as a counterpart of the CTMP noise model considered here. Namely, both methods provide a recipe for combining single- and two-qubit noise matrices into the full noise matrix describing a multi-qubit device.

The paper is organized as follows. Section II defines an error-mitigated mean value of an observable and shows that it provides an unbiased estimator of the ideal mean value. The tensor product and CTMP noise models are discussed in Sections III,IV. The parameters of these noise models can be extracted from the measurement calibration data as described in Section V. The experimental demonstration of our error mitigation methods is reported in Section VI. Appendices A,B,C prove technical lemmas used in the main text and provide additional details on the implementation of our algorithms. The experimental hardware is described in Appendix D.

II. UNBIASED ERROR MITIGATION

Let ρ be the output state of a (noisy) quantum circuit acting on n qubits. Our goal is to measure the expected value of a given observable O on the state ρ within a specified precision δ . For simplicity, below we assume that the observable O is diagonal in the standard basis and takes values in the range $[-1, 1]$, that is,

$$O = \sum_{x \in \{0,1\}^n} O(x)|x\rangle\langle x|, \quad |O(x)| \leq 1.$$

Consider M independent experiments where each experiment prepares the state ρ and then measures each qubit in the standard basis. Let $s^i \in \{0, 1\}^n$ be the string of measurement outcomes observed in the i -th experiment. In the absence of measurement errors one can approximate the mean value $\text{Tr}(\rho O)$ by an empirical mean value

$$\mu = M^{-1} \sum_{i=1}^M O(s^i). \quad (1)$$

By Hoeffding's inequality, $|\mu - \text{Tr}(\rho O)| \leq \delta$ with high probability (at least $2/3$) if we choose $M = 4\delta^{-2}$. Furthermore, μ is an unbiased estimator of $\text{Tr}(\rho O)$.

Suppose now that measurements are noisy. The most general model of a noisy n -qubit measurement involves a POVM with 2^n elements $\{\Pi_x\}$ labeled by n -bit strings x . The probability of observing a measurement outcome x is given by $P(x) = \text{Tr}(\rho \Pi_x)$. In the ideal case one has $\Pi_x = |x\rangle\langle x|$ while in the presence of noise Π_x could be arbitrary positive semi-definite operators that obey the normalization condition $\sum_x \Pi_x = I$. To enable efficient error mitigation, we make a simplifying assumption that

all POVM elements Π_x are diagonal in the standard basis, that is,

$$\Pi_y = \sum_x \langle y|A|x\rangle |x\rangle\langle x|$$

for some stochastic matrix A of size 2^n . In other words, $\langle y|A|x\rangle$ is the probability of observing an outcome y provided that the true outcome is x . For example, suppose $n = 1$ and

$$A = \begin{bmatrix} 0.9 & 0.2 \\ 0.1 & 0.8 \end{bmatrix}.$$

Then the probability of measuring 1 on a qubit prepared in the state $|0\rangle$ is 10%. The probability of measuring 0 on a qubit prepared in the state $|1\rangle$ is 20%. For brevity, we shall refer to such readout errors as $0 \rightarrow 1$ and $1 \rightarrow 0$ respectively. In this section we assume that the matrix A is known (this assumption is relaxed in Section V). By analogy with Eq. (1), define an error-mitigated empirical mean value as

$$\xi = M^{-1} \sum_{i=1}^M \sum_x O(x) \langle x|A^{-1}|s^i\rangle. \quad (2)$$

Lemma 1. *The random variable ξ is an unbiased estimator of $\text{Tr}(\rho O)$ with the standard deviation $\sigma_\xi \leq \Gamma M^{-1/2}$, where*

$$\Gamma = \max_y \sum_x |\langle x|A^{-1}|y\rangle|. \quad (3)$$

A proof of the lemma is given in Appendix A. By Hoeffding's inequality, $|\xi - \text{Tr}(\rho O)| \leq \delta$ with high probability (at least $2/3$) if the number of measurements is

$$M = 4\delta^{-2}\Gamma^2. \quad (4)$$

Recall that $M \sim \delta^{-2}$ in the ideal case. Thus the quantity Γ^2 can be viewed as an error mitigation overhead. In other words, error mitigation increases the number of measurements required to achieve a given precision δ by a factor of Γ^2 compared with the case of ideal measurements.

III. TENSOR PRODUCT NOISE

Consider first a simple case when A is a tensor product of 2×2 stochastic matrices,

$$A = \begin{bmatrix} 1 - \epsilon_1 & \eta_1 \\ \epsilon_1 & 1 - \eta_1 \end{bmatrix} \otimes \cdots \otimes \begin{bmatrix} 1 - \epsilon_n & \eta_n \\ \epsilon_n & 1 - \eta_n \end{bmatrix}. \quad (5)$$

Here ϵ_j and η_j are error rates describing readout errors $0 \rightarrow 1$ and $1 \rightarrow 0$ respectively. The error-mitigated mean value ξ defined in Eq. (2) can be easily computed when the observable O has a tensor product form,

$$O = O_1 \otimes \cdots \otimes O_n.$$

For example, O could be a product of Pauli Z operators on some subset of qubits. Alternatively, O could project some subset of qubits onto 0 or 1 states. Define a single-qubit state $|e\rangle = |0\rangle + |1\rangle$. From Eqs. (2,5) one gets

$$\xi = M^{-1} \sum_{i=1}^M \prod_{j=1}^n \langle e|O_j \begin{bmatrix} 1 - \epsilon_j & \eta_j \\ \epsilon_j & 1 - \eta_j \end{bmatrix}^{-1} |s_j^i\rangle. \quad (6)$$

Here $s^i \in \{0, 1\}^n$ is the outcome observed in the i -th measurement and $s_j^i \in \{0, 1\}$ is j -th bit of s^i . The single-qubit matrix element that appears in Eq. (6) can be computed in constant time for any fixed pair i, j . Thus one can compute ξ in time $\approx nM$, that is, the classical post-processing runtime is linear in the number of qubits and the number of measurements. The above algorithm is generalized to arbitrary observables O in Appendix C. The algorithm has runtime $\approx nM\tau$, where τ is the cost of computing the observable $O(x)$ for a given x .

Suppose the noise is weak such that $\epsilon_j, \eta_j \leq 1/2$ for all j . Substituting the definition of A into Eq. (3) one gets

$$\Gamma = \prod_{j=1}^n \frac{1 + |\epsilon_j - \eta_j|}{1 - \epsilon_j - \eta_j}. \quad (7)$$

We show how to extract the error rates ϵ_j and η_j from the readout calibration data in Section V. Then Eqs. (4,7) determine the number of measurements M , as a function of the desired precision δ . Assuming that $\epsilon_j, \eta_j \ll 1$ one gets $\Gamma \approx e^{2\gamma}$, where γ is the *noise strength* defined as

$$\gamma = \sum_{j=1}^n \max\{\epsilon_j, \eta_j\}.$$

The error mitigation overhead thus scales as $\Gamma^2 \approx e^{4\gamma}$.

IV. CORRELATED MARKOVIAN NOISE

In this section we propose a noise model based on Continuous Time Markov Processes (CTMP) [29]. It naturally extends the tensor product noise model and accounts for correlated (cross-talk) errors. Define a CTMP noise model as a stochastic matrix A of size $2^n \times 2^n$ such that ¹

$$A = e^G, \quad G = \sum_{i=1}^{2n^2} r_i G_i \quad (8)$$

where $e^G = \sum_{p=0}^{\infty} G^p/p!$ is the matrix exponential, $r_i \geq 0$ are error rates, and G_i are single-qubit or two-qubit operators from the following list

¹ Although our definition makes no reference to time, one can express A as a solution of a differential equation $\dot{A}(t) = GA(t)$ with the initial condition $A(0) = I$. Thus G serves as a generator of a continuous time Markov process such that $\langle y|G|x\rangle$ is the rate of transitions from a state x to a state $y \neq x$.

Generator G_i	Readout error	Number of generators
$ 1\rangle\langle 0 - 0\rangle\langle 0 $	$0 \rightarrow 1$	n
$ 0\rangle\langle 1 - 1\rangle\langle 1 $	$1 \rightarrow 0$	n
$ 10\rangle\langle 01 - 01\rangle\langle 01 $	$01 \rightarrow 10$	$n(n-1)$
$ 11\rangle\langle 00 - 00\rangle\langle 00 $	$00 \rightarrow 11$	$n(n-1)/2$
$ 00\rangle\langle 11 - 11\rangle\langle 11 $	$11 \rightarrow 00$	$n(n-1)/2$

Each operator G_i generates a readout error on some bit or some pair of bits (errors $10 \rightarrow 01$ are equivalent to $01 \rightarrow 10$ with the reversed order of bits). The right column shows the number of ways to choose qubit(s) acted upon by a generator. The negative terms in G_i ensure that A is a stochastic matrix. The CTMP model depends on $2n^2$ parameters r_i , as can be seen by counting the number of generators G_i of each type. We show how to extract the error rates r_i from the readout calibration data in Section V. Furthermore, the tensor product model Eq. (5) is a special case of CTMP with $r_i = 0$ for all two-qubit errors. Indeed, in this case Eq. (8) defines a tensor product of single-qubit stochastic matrices

$$\begin{bmatrix} 1 - \epsilon & \eta \\ \epsilon & 1 - \eta \end{bmatrix} = e^G, \quad G = -\frac{\log(1 - \epsilon - \eta)}{\epsilon + \eta} \begin{bmatrix} -\epsilon & \eta \\ \epsilon & -\eta \end{bmatrix}.$$

Here we assume that $\epsilon + \eta < 1$.

Next let us show how to compute the error-mitigated mean value ξ defined in Eq. (2). Since ξ itself is only a δ -estimate of the true mean value $\text{Tr}(\rho O)$, it suffices to approximate ξ within an error $\approx \delta$. We use a version of the well-known Gillespie algorithm [29] for simulating continuous Markov processes combined with the following lemma proved in Appendix B.

Lemma 2. *Suppose A is an invertible stochastic matrix. There exist stochastic matrices S_α and real coefficients c_α such that $\|c\|_1 \equiv \sum_\alpha |c_\alpha| \geq \Gamma$ and*

$$A^{-1} = \sum_\alpha c_\alpha S_\alpha. \quad (9)$$

Furthermore, $\|c\|_1 = \Gamma$ for some decomposition as above.

Here Γ is the quantity defined in Eq. (3). Suppose we have found a decomposition Eq. (9) for the inverse CTMP noise matrix $A^{-1} = e^{-G}$. Then the error-mitigated mean value ξ defined in Eq. (2) can be approximated using the following algorithm.

Algorithm 1:

```

 $T \leftarrow 4\delta^{-2}\|c\|_1^2$ 
for  $t = 1$  to  $\lceil T \rceil$  do
  Sample  $i \in [M]$  uniformly at random
  Sample  $\alpha$  from the distribution  $q_\alpha = |c_\alpha|/\|c\|_1$ 
  Sample  $x$  from the distribution  $\langle x|S_\alpha|s^i \rangle$ 
   $\xi_t \leftarrow \text{sgn}(c_\alpha)O(x)$ 
end for
return  $\xi' = T^{-1}\|c\|_1 \sum_{t=1}^T \xi_t$ 

```

Here c_α, S_α are the same as in Lemma 2. We claim that the output of this algorithm satisfies $|\xi' - \xi| \leq \delta$ with high probability (at least 2/3). Indeed, substituting Eq. (9) into Eq. (2) one gets

$$\xi = M^{-1}\|c\|_1 \sum_{i=1}^M \sum_\alpha \sum_x q_\alpha \text{sgn}(c_\alpha) O(x) \langle x|S_\alpha|s^i \rangle. \quad (10)$$

that is,

$$\xi = \|c\|_1 \mathbb{E}_{i,\alpha,x} \text{sgn}(c_\alpha) O(x(i,\alpha)). \quad (11)$$

Here $i \in [M]$ is picked uniformly at random, α is sampled from q_α , and $x(i,\alpha)$ is sampled from the distribution $\langle x|S_\alpha|s^i \rangle$ with fixed i,α . Accordingly, a random variable $\|c\|_1 \text{sgn}(c_\alpha) O(x(i,\alpha))$ is an unbiased estimator of ξ with the variance at most $\|c\|_1^2$. By Hoeffding's inequality, $|\xi' - \xi| \leq \delta$ with high probability (at least 2/3).

In the special case of tensor product noise one can efficiently compute a decomposition $A^{-1} = \sum_\alpha c_\alpha S_\alpha$ defined in Lemma 2 with $\|c\|_1 = \Gamma$, see Appendix C. This yields an instantiation of Algorithm 1 that computes a δ -estimate of ξ in time roughly

$$n\Gamma^2\delta^{-2} \sim nM. \quad (12)$$

Here M is the number of measurements determined by Eq. (4). Further details on the implementation of Algorithm 1 for the tensor product noise can be found in Appendix C.

Next, let us show how to implement Algorithm 1 for the CTMP model $A = e^G$. Define a parameter

$$\gamma = \max_{x \in \{0,1\}^n} -\langle x|G|x \rangle. \quad (13)$$

We shall see that the error mitigation overhead introduced by the CTMP method scales as $e^{4\gamma}$. To avoid a confusion with individual error rates r_i we shall refer to the quantity γ as the CTMP noise strength. The function $-\langle x|G|x \rangle$ can be viewed as a classical Ising-like Hamiltonian with two-spin interactions. Although finding the maximum energy of such Hamiltonians is NP-hard in the worst case, this can be easily done for moderate system sizes, say $n \leq 50$, using heuristic optimizers such as simulated annealing. Below we assume that the noise strength γ has been already computed.

First we note that $\gamma \geq 0$ since G has non-positive diagonal elements. Define a matrix

$$B = I + \gamma^{-1}G. \quad (14)$$

Using Eq. (13) and the fact that G has zero column sums, one can check that B is a stochastic matrix. Furthermore,

$$A^{-1} = e^{-G} = e^\gamma \cdot e^{-\gamma B} = \sum_{\alpha=0}^{\infty} \frac{e^\gamma (-\gamma)^\alpha}{\alpha!} B^\alpha \equiv \sum_{\alpha=0}^{\infty} c_\alpha S_\alpha \quad (15)$$

where

$$S_\alpha = B^\alpha \quad \text{and} \quad c_\alpha = \frac{e^{\gamma(-\gamma)^\alpha}}{\alpha!}. \quad (16)$$

Clearly, S_α is a stochastic matrix for any integer $\alpha \geq 0$. Thus Eqs. (15,16) provide a stochastic decomposition of A^{-1} stated in Lemma 2 with the 1-norm

$$\|c\|_1 = \sum_{\alpha=0}^{\infty} \frac{e^{\gamma}\gamma^\alpha}{\alpha!} = e^{2\gamma}. \quad (17)$$

From Eq. (4) and Lemma 2 one infers that the number of measurements required for error mitigation scales as $M \sim e^{4\gamma}$. Thus the quantity $e^{4\gamma}$ determines the error mitigation overhead. The probability distribution

$$q_\alpha = |c_\alpha|/\|c\|_1 = \frac{e^{-\gamma}\gamma^\alpha}{\alpha!} \quad (18)$$

is the Poisson distribution with the mean γ . Substitute the decomposition of Eqs. (15,16) into Algorithm 1. The only non-trivial step of the algorithm is sampling x from the distribution

$$\langle x|S_\alpha|s^i\rangle = \langle x|B^\alpha|s^i\rangle.$$

This amounts to simulating α steps of a Markov Chain with the transition matrix B . Note that B is a sparse efficiently computable matrix. More precisely, each column of B has $\approx n^2$ non-zero elements. Thus one can simulate a single step of the Markov chain in time $\approx n^2$. Since $\mathbb{E}(\alpha) = \gamma$, simulating α steps of the Markov chain on average takes time $\approx \gamma n^2$. The overall runtime of Algorithm 1 becomes

$$n^2\gamma\|c\|_1^2\delta^{-2} \sim n^2\gamma e^{4\gamma}\delta^{-2}.$$

Here we used Eq. (17). We note that the above algorithm is well defined as long as $\gamma \geq -\langle x|G|x\rangle$ for all x . Indeed, this ensures that the matrix B in Eq. (14) is stochastic. Thus one can choose γ as an arbitrary upper bound on the optimal solution of the optimization problem Eq. (13).

Assuming that cross-talk errors result from short-range interactions between qubits, one should expect that each qubit participates in a constant number of generators G_i independent of n . If this is the case, each column of G has at most Cn non-zero elements for some constant C . Then $\gamma \leq C\epsilon n$, where $\epsilon = \max_i r_i$ is the maximum error rate. Thus the exponential term in the runtime of Algorithm 1 becomes $e^{4C\epsilon n}$. Experimental data presented in Section VI are consistent with a linear scaling $\gamma \sim n$, see Fig. 5.

V. MEASUREMENT CALIBRATION

Calibration aims at learning parameters of the noise model from experimental data. A single calibration

round initializes the n -qubit register in a basis state $|x\rangle$ and performs a noisy measurement of each qubit, keeping the record of the measured outcome y . Let $m(y, x)$ be the number of rounds with the input state x and the measured output state y . We fix some set of input states $\mathcal{C} \subseteq \{0, 1\}^n$ and perform N_{cal} calibration rounds for each input state $x \in \mathcal{C}$. Thus $\sum_y m(y, x) = N_{\text{cal}}$ if $x \in \mathcal{C}$ and $m(y, x) = 0$ if $x \notin \mathcal{C}$. The calibration requires $N_{\text{cal}} \cdot |\mathcal{C}|$ experiments in total.

Learning the full noise matrix requires calibration of each possible input state, i.e. $\mathcal{C} = \{0, 1\}^n$. Let A_{full} be the empirical estimate of the full noise matrix, i.e.

$$\langle y|A_{\text{full}}|x\rangle = \frac{m(y, x)}{N_{\text{cal}}} \quad (19)$$

where $x, y \in \{0, 1\}^n$.

To learn error rates of the tensor product or the CTMP model it suffices to calibrate a small subset of basis states. Consider first the tensor product model. To ensure that each possible single-qubit readout error is probed on some input state, we assume that for any qubit j the set \mathcal{C} contains at least one state x with $x_j = 0$ and at least one state with $x_j = 1$. Let ϵ_j and η_j be the rates of errors $0 \rightarrow 1$ and $1 \rightarrow 0$ on the j -th qubit, see Eq. (5). We set

$$\begin{aligned} \epsilon_j &= \frac{\sum_{x,y} m(y, x) \langle 1|y_j\rangle \langle x_j|0\rangle}{\sum_{x,y} m(y, x) \langle x_j|0\rangle}, \\ \eta_j &= \frac{\sum_{x,y} m(y, x) \langle 0|y_j\rangle \langle x_j|1\rangle}{\sum_{x,y} m(y, x) \langle x_j|1\rangle} \end{aligned} \quad (20)$$

For example, ϵ_j is the fraction of calibration rounds that resulted in a readout error $0 \rightarrow 1$ on the j -th qubit among all rounds with $x_j = 0$. Let A_{tp} be the tensor product noise matrix defined by the rates ϵ_j, η_j and Eq. (5).

Next consider the CTMP model. Recall that the model is defined by $2n^2$ error rates describing readout errors $0 \leftrightarrow 1$, $01 \leftrightarrow 10$, and $00 \leftrightarrow 11$. To ensure that each possible two-qubit readout error is probed on some input state, we assume that for any pair of qubits $j \neq k$ and for any $v_j, v_k \in \{0, 1\}$ the set \mathcal{C} contains at least one state x such that $x_j = v_j$ and $x_k = v_k$. For example, \mathcal{C} could include bit strings 0^n , 1^n , and all weight-1 strings. Fix a pair of qubits $j \neq k$. Given a bit string $x \in \{0, 1\}^n$, let $x_{\text{in}} \in \{0, 1\}^2$ be the restriction of x onto the bits j, k and $x_{\text{out}} \in \{0, 1\}^{n-2}$ be the restriction of x onto all remaining bits $i \notin \{j, k\}$. We choose CTMP error rates on qubits j, k by defining a local noise matrix $A(j, k)$ and fitting $A(j, k)$ with a two-qubit CTMP model. More formally, $A(j, k)$ is a stochastic matrix of size 4×4 with matrix elements

$$\langle w|A(j, k)|v\rangle = \frac{\sum_{x,y} m(y, x) \langle w|y_{\text{in}}\rangle \langle x_{\text{in}}|v\rangle \langle x_{\text{out}}|y_{\text{out}}\rangle}{\sum_{x,y} m(y, x) \langle x_{\text{in}}|v\rangle \langle x_{\text{out}}|y_{\text{out}}\rangle}.$$

Here $v, w \in \{0, 1\}^2$. In words, $\langle w|A(j, k)|v\rangle$ is the fraction of calibration rounds that resulted in a readout error $v \rightarrow w$ on qubits j, k and no errors on qubits $i \notin \{j, k\}$ among

all rounds with $x_{\text{in}} = v$ that resulted in no errors on qubits $i \notin \{j, k\}$. Let

$$G(j, k) = \log A(j, k).$$

be the matrix logarithm of $A(j, k)$. We choose the branch of the log function such that $G(j, k) = 0$ if $A(j, k) = I$. The rates of two-qubit errors occurring on the qubits j, k are chosen as shown in the following table.

Readout error	CTMP error rate
$01 \rightarrow 10$	$\langle 10 G'(j, k) 01 \rangle$
$10 \rightarrow 01$	$\langle 01 G'(j, k) 10 \rangle$
$00 \rightarrow 11$	$\langle 11 G'(j, k) 00 \rangle$
$11 \rightarrow 00$	$\langle 00 G'(j, k) 11 \rangle$

Here $G'(j, k)$ is a matrix obtained from $G(j, k)$ by setting to zero all negative off-diagonal elements. Let $r_j^{0 \rightarrow 1}$ and $r_j^{1 \rightarrow 0}$ be the rates of single-qubit readout errors $0 \rightarrow 1$ and $1 \rightarrow 0$ occurring on the j -th qubit. We set

$$r_j^{0 \rightarrow 1} = \frac{1}{2(n-1)} \sum_{k \neq j} \langle 10 | G'(j, k) | 00 \rangle + \langle 11 | G'(j, k) | 01 \rangle,$$

$$r_j^{1 \rightarrow 0} = \frac{1}{2(n-1)} \sum_{k \neq j} \langle 00 | G'(j, k) | 10 \rangle + \langle 01 | G'(j, k) | 11 \rangle.$$

Here we noted that two-qubit errors $01 \leftrightarrow 10$ or $00 \leftrightarrow 11$ on qubits j, k can only be generated by the local noise matrix $A(j, k)$. Meanwhile, single-qubit errors $0 \leftrightarrow 1$ on the j -th qubit can be generated by any local noise matrix $A(j, k)$ with $k \neq j$. The above definition of single-qubit error rates amounts to taking the average over all such possibilities. Let A_{ctmp} be the CTMP noise model defined by the above error rates and Eq. (8).

The calibration method described above is well-defined only if the set of input states \mathcal{C} is *complete* in the sense that any two-bit readout error can be probed on some input state $x \in \mathcal{C}$. More formally, let us say that a set of n -bit strings \mathcal{C} is complete if for any pair of bits $1 \leq a < b \leq n$ and for any bit values $\alpha, \beta \in \{0, 1\}$ there exists at least one string $x \in \mathcal{C}$ such that $x_a = \alpha$ and $x_b = \beta$, see Fig. 1 for an example. Below we describe complete sets of input states used in our experiments.

Weight-1 calibration: the set \mathcal{C} includes all n -bit strings with the Hamming weight 0, 1, n . For example, if $n = 4$ then \mathcal{C} includes 0000, 1111, and all permutations of 1000.

Weight-2 calibration: the set \mathcal{C} includes all n -bit strings with the Hamming weight 0, 1, 2. For example, if $n = 4$ then \mathcal{C} includes 0000, all permutations of 1000, and all permutations of 1100.

Hadamard calibration: Let p be the smallest integer such that $n < 2^p$. Given an integer a in the range $[0, 2^p - 1]$, let $a_1, \dots, a_p \in \{0, 1\}$ be the binary digits of a . We choose

1	1	1	1	0	0	0	0	0	0
1	0	0	0	1	1	1	0	0	0
0	1	0	0	1	0	0	0	1	1
0	0	1	0	0	1	0	1	0	1
0	0	0	1	0	0	1	1	1	0
1	1	1	1	1	1	1	1	1	1

FIG. 1. A complete set of input states for CTMP calibration on $n = 10$ qubits. Each row of the matrix defines a bit string $x \in \mathcal{C}$. One can check that any pair of columns contains 00, 01, 10, and 11 in at least one row. Thus any two-bit readout error can be probed on some input $x \in \mathcal{C}$. We numerically checked that $|\mathcal{C}| \geq 6$ for any complete set $\mathcal{C} \subseteq \{0, 1\}^{10}$.

$\mathcal{C} = \{x^0, \dots, x^{2^p-1}\}$, where $x^a \in \{0, 1\}^n$ is defined by

$$x_b^a = \sum_{i=1}^p a_i b_i \pmod{2}$$

for $b = 1, \dots, n$. For example, if $n = 4$ then \mathcal{C} includes all even-weight bit strings $x \in \{0, 1\}^4$. In general, $|\mathcal{C}| \leq 2n$. One can easily check that each two-bit error is probed on exactly 2^{p-2} input states $x \in \mathcal{C}$. Thus the calibration resources are allocated evenly among all possible errors.

VI. EXPERIMENTAL RESULTS

All experiments are performed using n -qubit registers of the 20-qubit IBM Quantum device called `ibmq_johannesburg`, see Appendix D for details. Only the chosen n -qubit register is used for readout calibration. The unused qubits are initialized in the $|0\rangle$ state in each calibration round. Measurement outcomes on the unused qubits are ignored. All experiments were performed with $N_{\text{cal}} = 8192$ calibration rounds.

First, let us discuss how well the tensor product (TP) and the CTMP models approximate the readout noise observed in the experiment. We quantify the difference between a pair of n -qubit noise matrices A and A' using the Total Variation Distance (TVD)

$$\text{TVD}(A, A') = \frac{1}{2} \max_x \sum_y |\langle y | A | x \rangle - \langle y | A' | x \rangle|$$

The TVD provides an upper bound on the probability of distinguishing output distributions $A|p\rangle$ and $A'|p\rangle$ for any input distribution p . This enables a comparison between noise models in the worst-case scenario rather than for a specific quantum state and a specific observable.

For small number of qubits, $n \leq 7$, all three calibration methods described in Section V have been performed. Let A_{full} be the full noise matrix. Let A_{tp} and A_{ctmp} be the noise matrices predicted by the TP and CTMP models with weight-2 calibration. Fig. 2 shows the observed distances $\text{TVD}(A_{\text{full}}, A_{\text{ctmp}})$ and $\text{TVD}(A_{\text{full}}, A_{\text{tp}})$ for $n = 4, 5, 6, 7$ qubits (with the qubits chosen for $n < 7$ consisting of subsets of the $n = 7$ case) and 16 independent

experiments per each number of qubits. It indicates that the CTMP model provides a more accurate approximation of the readout noise achieving 2X reduction in the TVD metric for $n = 6, 7$ qubits. The increased separation between TP and CTMP models for larger numbers of qubits indicates that readout crosstalk is larger in the additional qubits used for the circuits with more qubits. Note that this improvement comes at no additional cost since both models have access to exactly the same calibration data set.

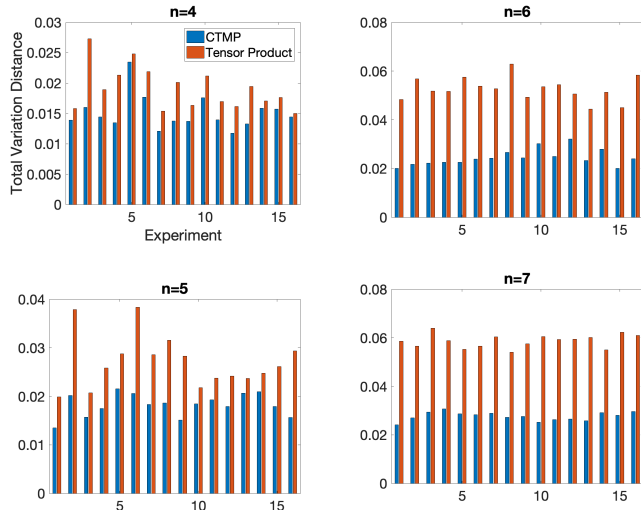


FIG. 2. *Distance from the full A -matrix.* A comparison is made between the full A -matrix, the tensor product, and the CTMP noise models. The plot shows the total variation distance $\text{TVD}(A_{\text{full}}, A_{\text{ctmp}})$ (blue bars) and $\text{TVD}(A_{\text{full}}, A_{\text{tp}})$ (red bars) for $n = 4, 5, 6, 7$ qubits. For each number of qubits we performed 16 independent experiments.

In order to compare the methods of measurement mitigation on relevant quantum circuits, we use graph states as an example of a highly entangled state that can serve as a benchmark for quantum devices. We estimate the fidelity with stabilizer measurements for n -qubit graph states [30] of the form

$$|G_n\rangle = \left(\prod_{j=1}^{n-1} CZ_{j,j+1} \right) H^{\otimes n} |0^n\rangle. \quad (21)$$

The corresponding graph is a path with n vertices. The state $|G_n\rangle$ is a stabilizer state with the stabilizer group \mathcal{S} generated by Pauli operators $S_1 = X_1 Z_2$, $S_n = Z_{n-1} X_n$, and $S_j = Z_{j-1} X_j Z_{j+1}$ for $1 < j < n$. Let ρ be an approximate version of $|G_n\rangle$ prepared in the lab by executing the quantum circuit Eq. (21). The fidelity $\langle G_n | \rho | G_n \rangle$ can be estimated by picking a random element of the stabilizer group $S \in \mathcal{S}$ and measuring its mean value on the state ρ ,

$$F = \langle G_n | \rho | G_n \rangle = 2^{-n} \sum_{S \in \mathcal{S}} \text{Tr}(\rho S) = \mathbb{E}_S \text{Tr}(\rho S). \quad (22)$$

A comparison between the error-mitigated fidelity F obtained using the full A -matrix, TP, and CTMP methods with weight-1 calibration is shown in Fig. 3(a). We also show raw values of F obtained without error mitigation. The full A -matrix model was only used for $n \leq 7$. The data suggests that CTMP provides much more accurate estimates of the fidelity, compared with the TP model which systematically over-estimates the fidelity. The mean value of each stabilizer was estimated by measuring $M = 8192$ copies of the graph state $|G_n\rangle$. The entire experiment (calibration and mean value measurements) was repeated 16 times in order to estimate error bars.

For the six-qubit graph state, there are few enough stabilizers that we can measure all of them. Let $\xi_{\text{full}}(S)$, $\xi_{\text{tp}}(S)$, and $\xi_{\text{ctmp}}(S)$ be the error-mitigated mean values $\text{Tr}(\rho S)$ of a stabilizer S obtained using the full A -matrix, the TP, and the CTMP method respectively. Fig. 3(b) shows the differences $\xi_{\text{full}}(S) - \xi_{\text{tp}}(S)$ and $\xi_{\text{full}}(S) - \xi_{\text{ctmp}}(S)$ averaged over stabilizers of the same weight, $|S|$. It can be seen that the TP method systematically over-estimates the mean values for stabilizers with the weight 3, 4, 5, 6, while CTMP shows a much better agreement with the full A -matrix method. The difference between CTMP and TP increases for higher weight stabilizers, which we expect to be more sensitive to correlated readout errors. Fig. 4 (a) shows the combined rate of correlated two-qubit errors $01 \leftrightarrow 10$ and $00 \leftrightarrow 11$, as described by the CTMP model, for each pair of qubits in the chosen 12-qubit register. It can be seen that some pairs of qubits, such as (7, 8), experience correlated errors with a few-percent rate.

Recall that the error mitigation overhead introduced by the CTMP method is roughly $e^{4\gamma}$, where γ is the noise strength defined in Eq. (13). Fig. 5 shows the noise strength γ observed in the graph state experiments as a function of the number of qubits. Error bars were estimated from 16 independent experiments. The data is consistent with a linear scaling, $\gamma \approx 0.05n$, and with single qubit readout fidelities of a few percent (see Appendix D for more details on the readout errors in the hardware). We observed $\gamma \leq 0.7$ for all experiments reported above. The error mitigation overhead is therefore at most $e^{4 \cdot 0.7} \leq 20$, indicated a factor of 20 increase in the number of measurements required.

Our last experiment demonstrates error mitigated stabilizer measurements on an entangled 20-qubit state $|\psi\rangle$ generated by a random depth-4 Clifford circuit. The circuit consists of two layers of CNOT gates on nearest-neighbor pairs of qubits and two layers of single-qubit Clifford operators, see Fig. 6. Such a state $|\psi\rangle$ can be specified by a group of Pauli-type stabilizers $S \in \pm\{I, X, Y, Z\}^{\otimes 20}$ such that $S|\psi\rangle = |\psi\rangle$. We have measured mean values $\langle \psi | S | \psi \rangle$ for ≈ 500 stabilizers S chosen randomly such that the weight of S varies between 1 and 20. In the absence of readout and gate errors one has $\langle \psi | S | \psi \rangle = 1$ for all stabilizers S . Our results for the error mitigated mean values $\langle \psi | S | \psi \rangle$ obtained with the TP

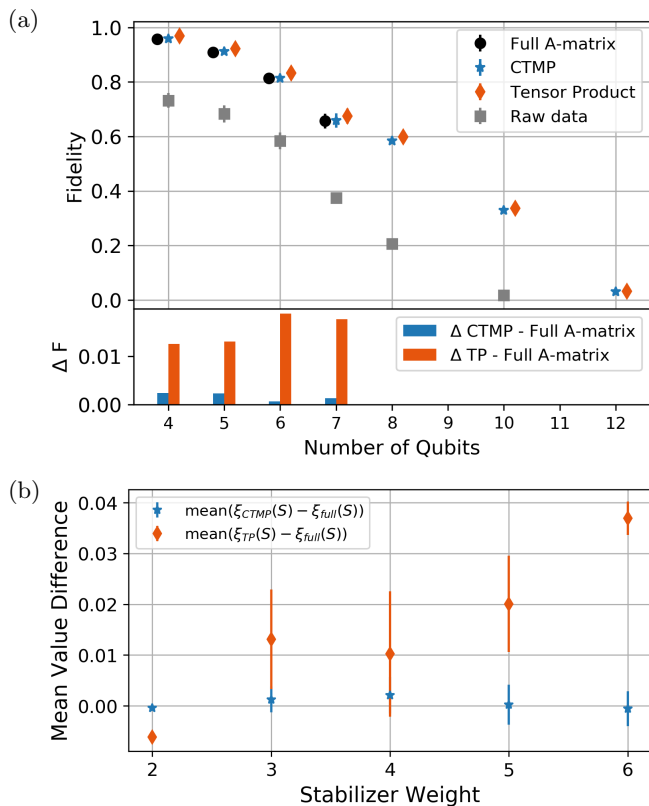


FIG. 3. (a) Average mean values of graph state stabilizers for $4 \leq n \leq 12$ qubits obtained using different error mitigation methods (Full A-matrix in black circles, CTMP in blue stars, TP in red diamonds) and without error mitigation (raw, grey squares) with the difference, ΔF , between CTMP and full A-matrix (blue) and TP and full A-matrix (red) below. To avoid clutter, data points on the plot (a) that represent the error mitigated results (black, blue, and red) are offset along the horizontal axis. The fidelity was estimated by averaging the stabilizer mean value over ≈ 100 random stabilizers. For each stabilizer 16 independent experiments were performed to estimate error bars. (b) Differences in the error-mitigated mean values $\xi_{\text{ctmp}}(S) - \xi_{\text{full}}(S)$ (blue) and $\xi_{\text{tp}}(S) - \xi_{\text{full}}(S)$ (red) for stabilizers S of the 6-qubit graph state from (a) obtained using the CTMP, the tensor product, and the full A-matrix noise models. Data sets are averaged over all stabilizers of the same weight, indicated on the horizontal axis. The graph state has no stabilizers with weight=1, two stabilizers with weight=2, eight stabilizers with weight=3, etc.

and CTMP methods and Hadamard-type calibration are shown on Fig. 6. In contrast to the graph state experiment, error rates associated with correlated two-qubit errors $01 \leftrightarrow 10$ and $00 \leftrightarrow 11$ observed in the 20-qubit experiment are less than 1% for all pairs of qubits, see Fig. 4(b). Accordingly, the difference between error mitigated mean values obtained using the TP and CTMP methods is less pronounced.

The 20-qubit experiment showcases the scalability of our method. For example, computing the error-mitigated mean value of a 20-qubit stabilizer using the CTMP

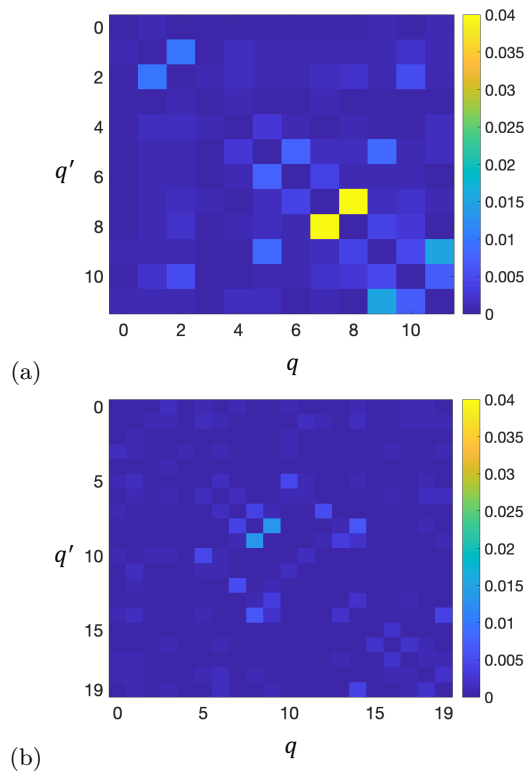


FIG. 4. Combined rate of correlated two-qubit readout errors for each pair of qubits $0 \leq q, q' < n$. (a) 12-qubit graph state experiment. (b) 20-qubit random Clifford circuit experiment.

method with $T = 10^6$ samples in Algorithm 1 takes about 2 seconds on a laptop computer (recall that T affects the cost of a classical post-processing but not the number of experiments). Extracting CTMP error rates from the 20-qubit calibration data takes about 1 second. Thus we expect that our methods can be scaled up to a larger number of qubits.

VII. CONCLUSIONS

We introduced scalable methods of mitigating readout errors in multi-qubit experiments. Our methods are capable of mitigating correlated cross-talk errors and enable efficient implementation of the calibration and the noise inversion steps. The proposed error mitigation methods are demonstrated experimentally for measurements of up to 20 qubits. We believe that our techniques will be useful for many near term quantum applications, and their scalability will be increasingly important as the problem sizes become larger.

ACKNOWLEDGMENTS

The authors thank Kristan Temme for helpful discussions and Neereja Sundaresan for experimental contribu-




FIG. 5. CTMP noise strength $\gamma = \max_x -\langle x|G|x\rangle$, where G is is generator matrix of the CTMP model such that $A_{\text{ctmp}} = e^G$. The error mitigation overhead introduced by the CTMP method scales as $e^{4\gamma}$. The data suggests a linear scaling, $\gamma \approx 0.05n$. We observed $\gamma \approx 1.1$ in the $n = 20$ experiment.

tions. This work was supported in part by ARO under Contract No. W911NF-14-1-0124 and by the IBM Research Frontiers Institute. The authors declare that they have no competing interests.

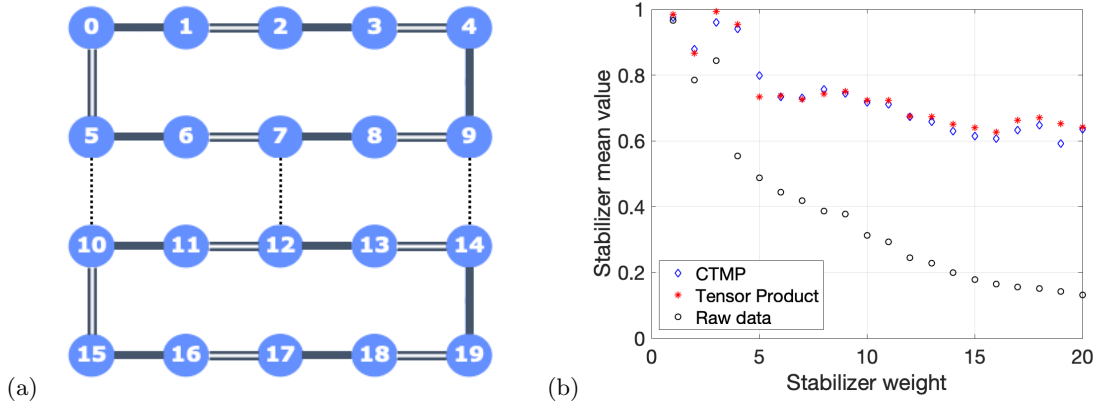


FIG. 6. Error-mitigated stabilizer measurements performed on a 20-qubit state $|\psi\rangle$ generated by a random depth-4 Clifford circuit. The circuit consists of alternating layers of single-qubit Clifford operators and entangling layers composed of nearest-neighbor CNOT gates. (a) the coupling map for the 20-qubit device, `ibmq_johannesburg`. The patterns of CNOTs in the two entangling layers are indicated by single and double solid lines, with unused physical connections in dotted lines. (b) Error mitigated and raw mean values $\langle\psi|S|\psi\rangle$ for ≈ 500 stabilizer operators S chosen such that the weight of S varies between 1 and 20. Each data point represents the average over all stabilizers S with a given weight. In the absence of gate and readout errors each stabilizer S has mean value +1. Error mitigation was performed using the CTMP and Tensor Product methods with the Hadamard-type calibration. Black circles indicate raw mean values measured without error mitigation.

-
- [1] Kristan Temme, Sergey Bravyi, and Jay M Gambetta. Error mitigation for short-depth quantum circuits. *Physical Review Letters*, 119(18):180509, 2017.
- [2] Ying Li and Simon C Benjamin. Efficient variational quantum simulator incorporating active error minimization. *Physical Review X*, 7(2):021050, 2017.
- [3] Suguru Endo, Simon Benjamin, and Ying Li. Practical quantum error mitigation for near-future applications. *Physical Review X*, 8:031027, 2018.
- [4] Matthew Otten and Stephen Gray. Recovering noise-free quantum observables. *arXiv preprint arXiv:1806.07860*, 2018.
- [5] X Bonet-Monroig, R Sagastizabal, M Singh, and TE O'Brien. Low-cost error mitigation by symmetry verification. *arXiv preprint arXiv:1807.10050*, 2018.
- [6] Abhinav Kandala, Kristan Temme, Antonio D Córcoles, Antonio Mezzacapo, Jerry M Chow, and Jay M Gambetta. Error mitigation extends the computational reach of a noisy quantum processor. *Nature*, 567(7749):491–495, 2019.
- [7] Vojtěch Havlíček, Antonio D Córcoles, Kristan Temme, Aram W Harrow, Abhinav Kandala, Jerry M Chow, and Jay M Gambetta. Supervised learning with quantum-enhanced feature spaces. *Nature*, 567(7747):209, 2019.
- [8] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J Love, and Jeremy L O'Brien. A variational eigenvalue solver on a photonic quantum processor. *Nature Communications*, 5:4213, 2014.
- [9] Abhinav Kandala, Antonio Mezzacapo, Kristan Temme, Maika Takita, Markus Brink, Jerry M. Chow, and Jay M. Gambetta. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature*, 549:242–246, 2017.
- [10] Leonardo DiCarlo, Matthew D Reed, Luyan Sun, Blake R Johnson, Jerry M Chow, Jay M Gambetta, Luigi Frunzio, Steven M Girvin, Michel H Devoret, and Robert J Schoelkopf. Preparation and measurement of three-qubit entanglement in a superconducting circuit. *Nature*, 467(7315):574, 2010.
- [11] Yanzhu Chen, Maziar Farahzad, Shinjae Yoo, and Tzu-Chieh Wei. Detector tomography on IBM quantum computers and mitigation of an imperfect measurement. *Physical Review A*, 100(5):052315, 2019.
- [12] J. M. Chow, L. DiCarlo, J. M. Gambetta, A. Nunnenkamp, Lev S. Bishop, L. Frunzio, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf. Detecting highly entangled states with a joint qubit readout. *Phys. Rev. A*, 81:062325, Jun 2010.
- [13] Colm A. Ryan, Blake R. Johnson, Jay M. Gambetta, Jerry M. Chow, Marcus P. da Silva, Oliver E. Dial, and Thomas A. Ohki. Tomography via correlation of noisy measurement records. *Phys. Rev. A*, 91:022118, Feb 2015.
- [14] Seth T. Merkel, Jay M. Gambetta, John A. Smolin, Stefano Poletto, Antonio D. Córcoles, Blake R. Johnson, Colm A. Ryan, and Matthias Steffen. Self-consistent quantum process tomography. *Phys. Rev. A*, 87:062119, Jun 2013.
- [15] Jerry M. Chow, Jay M. Gambetta, A. D. Córcoles, Seth T. Merkel, John A. Smolin, Chad Rigetti, S. Poletto, George A. Keefe, Mary B. Rothwell, J. R. Rozen, Mark B. Ketchen, and M. Steffen. Universal quantum gate set approaching fault-tolerant thresholds with superconducting qubits. *Phys. Rev. Lett.*, 109:060501, Aug 2012.
- [16] Mingyu Sun and Michael R Geller. Efficient characterization of correlated spam errors. *arXiv preprint arXiv:1810.10523*, 2018.
- [17] Gadi Aleksandrowicz, Thomas Alexander, Panagiotis Barkoutsos, Luciano Bello, Yael Ben-Haim, David Bucher, Francisco Jose Cabrera-Hernandez, Jorge Carballo-Franquis, Adrian Chen, Chun-Fu Chen, Jerry M. Chow, Antonio D. Córcoles-Gonzales, Abigail J. Cross, Andrew Cross, Juan Cruz-Benito, Chris Culver, Salvador De La Puente Gonzalez, Enrique De La Torre, Delton Ding, Eugene Dumitrescu, Ivan Duran, Pieter Eendebak, Mark Everitt, Ismael Faro Sertage, Albert Frisch, Andreas Fuhrer, Jay Gambetta, Borja Godoy Gago, Juan Gomez-Mosquera, Donny Greenberg, Ikko Hamamura, Vojtech Havlicek, Joe Hellmers, ukasz Herok, Hiroshi Horii, Shaohan Hu, Takashi Imamichi, Toshinari Itoko, Ali Javadi-Abhari, Naoki Kanazawa, Anton Karazeev, Kevin Krsulich, Peng Liu, Yang Luh, Yunho Maeng, Manoel Marques, Francisco Jose Martn-Fernandez, Douglas T. McClure, David McKay, Srujan Meesala, Antonio Mezzacapo, Nikolaj Moll, Diego Moreda Rodriguez, Giacomo Nannicini, Paul Nation, Pauline Ollitrault, Lee James O'Riordan, Han-hee Paik, Jess Prez, Anna Phan, Marco Pistoia, Viktor Prutyaynov, Max Reuter, Julia Rice, Abdn Rodriguez Davila, Raymond Harry Putra Rudy, Mingi Ryu, Ninad Sathaye, Chris Schnabel, Eddie Schoute, Kanav Setia, Yunong Shi, Adenilton Silva, Yukio Siraichi, Seyon Sivarajah, John A. Smolin, Mathias Soeken, Hitomi Takahashi, Ivano Tavernelli, Charles Taylor, Pete Taylor, Kenso Trabing, Matthew Treinish, Wes Turner, Desiree Vogt-Lee, Christophe Vuillot, Jonathan A. Wildstrom, Jessica Wilson, Erick Winston, Christopher Wood, Stephen Wood, Stefan Wrner, Ismail Yunus Akhalwaya, and Christa Zoufal. Qiskit: An Open-source Framework for Quantum Computing, January 2019.
- [18] Johannes Heinsoo, Christian Kraglund Andersen, Ants Remm, Sebastian Krinner, Theodore Walter, Yves Salathé, Simone Gasparinetti, Jean-Claude Besse, Anton Potočnik, Andreas Wallraff, and Christopher Eichler. Rapid high-fidelity multiplexed readout of superconducting qubits. *Phys. Rev. Applied*, 10:034040, Sep 2018.
- [19] Alireza Seif, Kevin A Landsman, Norbert M Linke, Caroline Figgatt, Christopher Monroe, and Mohammad Hafezi. Machine learning assisted readout of trapped-ion qubits. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 51(17):174006, 2018.
- [20] Michael R Geller. Rigorous measurement error correction. *arXiv preprint arXiv:2002.01471*, 2020.
- [21] Michael R Geller and Mingyu Sun. Efficient correction of multiqubit measurement errors. *arXiv preprint arXiv:2001.09980*, 2020.
- [22] Filip B Maciejewski, Zoltán Zimborás, and Michał Oszmaniec. Mitigation of readout noise in near-term quantum devices by classical post-processing based on detector tomography. *Quantum*, 4:257, 2020.

- [23] Jinzhao Sun, Xiao Yuan, Takahiro Tsunoda, Vlatko Vedral, Simon C Benjamin, and Suguru Endo. Mitigating realistic noise in practical noisy intermediate-scale quantum devices. *arXiv preprint arXiv:2001.04891*, 2020.
- [24] Benjamin Nachman, Miroslav Urbanek, Wibe A de Jong, and Christian W Bauer. Unfolding quantum computer readout noise. *arXiv preprint arXiv:1910.01969*, 2019.
- [25] Glen Cowan. A survey of unfolding methods for particle physics. *Proc. Advanced Statistical Techniques in Particle Physics, Durham*, 2002.
- [26] Kathleen E Hamilton and Raphael C Pooser. Error-mitigated data-driven circuit learning on noisy quantum hardware. *arXiv preprint arXiv:1911.13289*, 2019.
- [27] Kathleen E Hamilton, Tyler Kharazi, Titus Morris, Alexander J McCaskey, Ryan S Bennink, and Raphael C Pooser. Scalable quantum processor noise characterization. *arXiv preprint arXiv:2006.01805*, 2020.
- [28] Ulrich Schollwöck. The density-matrix renormalization group in the age of matrix product states. *Annals of Physics*, 326(1):96–192, 2011.
- [29] Daniel T Gillespie. Exact stochastic simulation of coupled chemical reactions. *The Journal of Physical Chemistry*, 81(25):2340–2361, 1977.
- [30] Robert Raussendorf and Hans J Briegel. A one-way quantum computer. *Physical Review Letters*, 86(22):5188, 2001.

Appendix A: Proof of Lemma 1

Represent the measurement outcomes s^1, \dots, s^M by a probability vector $|\psi\rangle = M^{-1} \sum_{i=1}^M |s^i\rangle$. Then

$$\xi = \sum_x O(x) \langle x|A^{-1}|\psi\rangle$$

and

$$\mathbb{E}(\xi^2) = \sum_{x,y} O(x)O(y) \langle x|A^{-1}\mathbb{E}(|\psi\rangle\langle\psi|)(A^{-1})^T|y\rangle. \quad (\text{A1})$$

Let $P(x) = \langle x|\rho|x\rangle$ and $|P\rangle = \sum_x P(x)|x\rangle$. Then

$$\mathbb{E}(|\psi\rangle\langle\psi|) = \frac{M-1}{M} A|P\rangle\langle P|A^T + \frac{1}{M} \sum_x |x\rangle\langle x|\langle x|A|P\rangle. \quad (\text{A2})$$

Substituting Eq. (A2) into Eq. (A1) gives

$$\mathbb{E}(\xi^2) \leq \mathbb{E}(\xi)^2 + \frac{1}{M} \sum_{x,y} O(x)O(y) \langle x|A^{-1}D(A^{-1})^T|y\rangle, \quad (\text{A3})$$

where D is a diagonal operator with entries $\langle x|D|x\rangle = \langle x|A|P\rangle$. Since $\langle x|A|P\rangle$ is a normalized probability distribution, one concludes that $\mathbb{E}(\xi^2) - \mathbb{E}(\xi)^2$ is upper bounded by

$$\frac{1}{M} \max_z \sum_{x,y} |O(x)O(y)| \cdot |\langle x|A^{-1}|z\rangle| \cdot |\langle y|A^{-1}|z\rangle|. \quad (\text{A4})$$

By assumption, $|O(x)| \leq 1$ for all x , which gives $\mathbb{E}(\xi^2) - \mathbb{E}(\xi)^2 \leq M^{-1}\Gamma^2$. Thus the standard deviation of ξ is at most $M^{-1/2}\Gamma$.

Appendix B: Proof of Lemma 2

Suppose M is a real $N \times N$ matrix (we shall be interested in the case $M = A^{-1}$). Let us state necessary and sufficient conditions under which M admits a decomposition

$$M = \sum_a c_a S_a \quad (\text{B1})$$

where S_a are stochastic matrices and c_a are real coefficients. Define the j -th column sum of M as $\sigma_j(M) = \sum_{i=1}^N M_{i,j}$.

Lemma 3. *A matrix M admits a decomposition Eq. (B1) iff all column sums of M are the same.*

Proof. From Eq. (B1) one gets $\sigma_j(M) = \sum_a c_a$ for all j .

Conversely, suppose $\sigma_j(M) = \gamma$ for all j . For any integers $i, j \in [N]$ define a matrix $T(i, j) = |1\rangle\langle j| - |i\rangle\langle j|$. Note that all column sums of $T(i, j)$ are zero. One can easily verify that $T(i, j)$ is a linear combination of two stochastic matrices. Let

$$M' = M + \sum_{i=2}^N \sum_{j=1}^N M_{i,j} T(i, j).$$

The matrix M' has zero rows $2, 3, \dots, N$ and $M'_{1,j} = \gamma$ for all j . Thus $M' = \gamma S$ for a stochastic matrix S . This shows that M is a linear combination of stochastic matrices. \square

As a corollary, the inverse of any stochastic matrix can be written as a linear combination of stochastic matrices. Indeed, if $M = A^{-1}$ for some stochastic matrix A then all column sums of M are equal to one.

Given a matrix M let $\Gamma(M)$ be the maximum 1-norm of its columns,

$$\Gamma(M) = \max_j \sum_{i=1}^N |M_{i,j}|.$$

Lemma 2 is a special case of the following.

Lemma 4. *Suppose all column sums of M are the same. Then any decomposition Eq. (B1) obeys $\|c\|_1 \geq \Gamma(M)$ and $\|c\|_1 = \Gamma(M)$ for some decomposition.*

Proof. Suppose M is written as in Eq. (B1). By triangle inequality,

$$\Gamma(M) = \sum_i |M_{i,j}| \leq \sum_a |c_a| \sum_i |\langle i|S_a|j\rangle| = \|c\|_1 \quad (\text{B2})$$

Conversely, let $\Lambda(M)$ be the minimum 1-norm $\|c\|_1$, where the minimum is over all decompositions Eq. (B1). From Eq. (B2) one infers that $\Lambda(M) \geq \Gamma(M)$. It suffices to prove that

$$\Lambda(M) \leq \Gamma(M). \quad (\text{B3})$$

Let k be the number of non-zero elements in M . We shall prove Eq. (B3) using induction in k . The base of induction is $k = 0$. In this case M is the all-zeros matrix and $\Lambda(M) = \Gamma(M) = 0$.

Suppose we have already proved Eq. (B3) for all matrices M with at most k non-zeros. Consider a matrix M with $k + 1$ non-zeros. Let $\gamma = \sigma_j(M)$ be the column sum of M (by assumption, it does not depend on j). We shall consider two cases.

Case 1: $\gamma \neq 0$. Then all columns of M are non-zero and each column contains a non-zero entry $M_{i,j}$ with the same sign as γ . Let ω be a smallest magnitude non-zero entry of M with the same sign as γ . Assume wlog that $\omega = M_{1,1}$ (otherwise, permute rows and/or columns of M). Choose any function $f : [N] \rightarrow [N]$ such that $f(1) = 1$ and $M_{f(j),j}$ is a largest magnitude entry in the j -th column with the same sign as γ . Define a new matrix

$$M' = M - \omega \sum_{j=1}^N |f(j)\rangle\langle j|. \quad (\text{B4})$$

Note that M' has at most k non-zeros since $M_{f(j),j} \neq 0$ for all j and $M'_{1,1} = 0$. Since $\sum_{j=1}^N |f(j)\rangle\langle j|$ is a stochastic matrix, Eq. (B4) gives

$$\Lambda(M) \leq \Lambda(M') + |\omega|. \quad (\text{B5})$$

We claim that

$$\Gamma(M') \leq \Gamma(M) - |\omega|. \quad (\text{B6})$$

Indeed, let $\Gamma_j(M)$ and $\Gamma_j(M')$ be the 1-norm of the j -th column of M and M' respectively. We have $\Gamma_1(M') = \Gamma_1(M) - |\omega|$ since the first column of M' is obtained from the one of M by setting to zero the entry $M_{1,1} = \omega$. If $j \geq 2$ then

$$\begin{aligned} \Gamma_j(M') &= |M_{f(j),j} - \omega| + \sum_{i \neq f(j)} |M_{i,j}| \\ &= \text{sgn}(\omega) \cdot (M_{f(j),j} - \omega) + \sum_{i \neq f(j)} |M_{i,j}| \\ &= \Gamma_j(M) - |\omega|. \end{aligned}$$

Here we used $\text{sgn}(\omega) = \text{sgn}(M_{f(j),j})$ and $|\omega| \leq |M_{f(j),j}|$. This proves Eq. (B6). By induction hypothesis, $\Lambda(M') \leq \Gamma(M')$. Combining this and Eqs. (B5,B6) one arrives at

$$\Lambda(M) \leq \Lambda(M') + |\omega| \leq \Gamma(M') + |\omega| \leq \Gamma(M). \quad (\text{B7})$$

Case 2: $\gamma = 0$. Then each non-zero column of M contains at least one positive and at least one negative entry. Let ω be a non-zero entry of M with the smallest magnitude. Assume wlog that $\omega = M_{1,1}$. Choose any function $f : [N] \rightarrow [N]$ such that $f(1) = 1$ and $M_{f(j),j}$ is the largest magnitude entry in the j -th column with the same sign as ω . If the j -th column is zero then set $f(j)$ arbitrarily. Define a new matrix M' using Eq. (B4). The same arguments as above show that $\Gamma_1(M') = \Gamma_1(M) - |\omega|$ and

$\Gamma_j(M') \leq \Gamma_j(M) - |\omega|$ whenever the j -th column of M is non-zero. Suppose now that the j -th column of M is zero. Then $\Gamma_j(M) = 0$ and $\Gamma_j(M') = |\omega|$. We claim that $\Gamma(M) \geq 2|\omega|$. Indeed, by assumption, M has at least one non-zero column. Such column must have at least two non-zero entries with the magnitude at least $|\omega|$, that is, $\Gamma(M) \geq 2|\omega|$. Thus, if the j -th column of M is zero then $\Gamma_j(M') = |\omega| \leq \Gamma(M) - |\omega|$. This proves Eq. (B6). The same arguments as above give $\Lambda(M) \leq \Gamma(M)$. \square

Appendix C: Error mitigation for product noise

Here we instantiate Algorithm 1 for the tensor product noise defined in Eq. (5). For each $j \in [n]$ define the following quantities

$$\begin{aligned} c_{j0} &= \frac{2 - \epsilon_j - \eta_j}{2(1 - \epsilon_j - \eta_j)}, & c_{j1} &= -\frac{\epsilon_j + \eta_j}{2(1 - \epsilon_j - \eta_j)}, \\ c_{j2} &= \frac{\epsilon_j - \eta_j}{2(1 - \epsilon_j - \eta_j)}, & c_{j3} &= \frac{\eta_j - \epsilon_j}{2(1 - \epsilon_j - \eta_j)}. \end{aligned}$$

Define single-bit boolean functions F_0, F_1, F_2, F_3 as

$$F_0(b) = b, \quad F_1(b) = 1 \oplus b, \quad F_2(b) = 0, \quad F_3(b) = 1.$$

Here $b \in \{0, 1\}$. Consider the following algorithm.

Algorithm 2

```

T ← 4δ-2Γ2
for t = 1 to ⌈T⌉ do
  ξt ← 1
  Sample i ∈ [M] uniformly at random.
  xt ← si
  for j = 1 to n do
    Sample α ∈ {0, 1, 2, 3} with probabilities
    qjα = |cjα|(|cj0| + |cj1| + |cj2| + |cj3|)-1
    xjt ← Fα(xjt)
    ξt ← ξt · sgn(cjα)
  end for
end for
return ξt = T-1Γ ∑t=1T ξtO(xt)

```

Here $x_j^t \in \{0, 1\}$ is the j -th bit of x^t . We claim that Algorithm 2 is a special case of Algorithm 1 stated in the main text. Indeed, define stochastic matrices

$$A_j = \begin{bmatrix} 1 - \epsilon_j & \eta_j \\ \epsilon_j & 1 - \eta_j \end{bmatrix} \quad \text{and} \quad V_i = \sum_{b=0,1} |F_i(b)\rangle\langle b|.$$

A simple calculation shows that

$$A_j^{-1} = c_{j0}V_0 + c_{j1}V_1 + c_{j2}V_2 + c_{j3}V_3. \quad (\text{C1})$$

and

$$|c_{j0}| + |c_{j1}| + |c_{j2}| + |c_{j3}| = \frac{1 + |\epsilon_j - \eta_j|}{1 - \epsilon_j - \eta_j}. \quad (\text{C2})$$

Taking n -fold tensor product of stochastic decompositions defined in Eq. (C1) one gets $A^{-1} = \sum_{\alpha} c_{\alpha} S_{\alpha}$, where S_{α} is a tensor product of stochastic matrices V_i and c_{α} are real coefficients such that $\|c\|_1 = \Gamma$. The inner loop of Algorithm 2 samples x^t from the distribution $S_{\alpha}|s^i\rangle$, where α is sampled from $q_{\alpha} = |c_{\alpha}|/\Gamma$. Thus Algorithm 2 is indeed a special case of Algorithm 1. We conclude that the output of Algorithm 2 satisfies $|\xi' - \xi| \leq \delta$ with high probability (at least $2/3$). The algorithm has runtime $\approx nT\tau \sim nM\tau$, where M is the number of measurements determined by Eqs. (4,7) and τ is the cost of computing the observable $O(x)$ for a given x . For simplicity, here we assume that each sampling step in Algorithm 2 can

be done in unit time.

Appendix D: Experimental Hardware

All experiments were performed on the 20-qubit IBM Quantum device called `ibmq_johannesburg`. The device coupling map is shown in Fig. 6(a) of the main text. Typical mean error rates are $(1.75 \pm 1.05) \times 10^{-2}$ for CX gates, $(4.35 \pm 1.70) \times 10^{-4}$ for single qubit gates, and $(3.44 \pm 1.72) \times 10^{-2}$ for readout errors; a more detailed breakdown of the error rates is given in the histograms in Fig. 7. The graph state experiments in Fig. 3 were performed on n -qubit subsets of qubits 0 through 11.

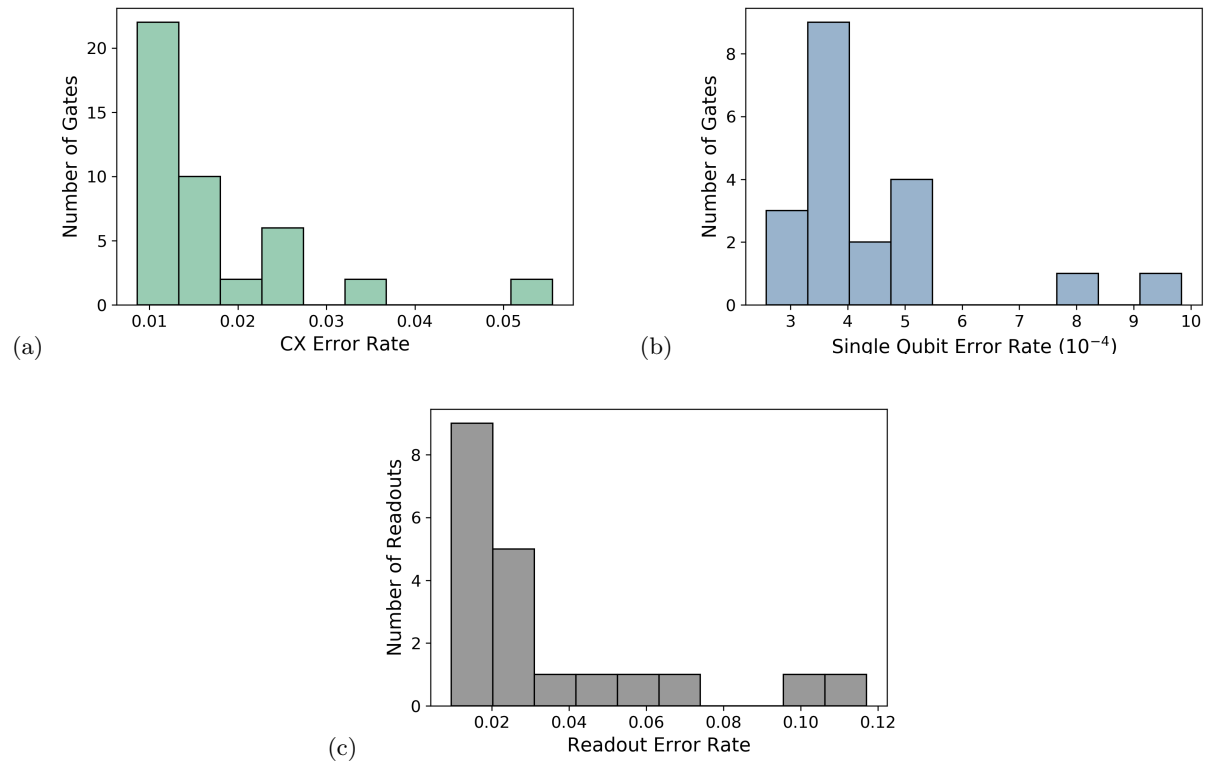


FIG. 7. Histograms of the (a) two-qubit CX (b) single-qubit and (c) readout error rates across the full 20-qubit device, `ibmq_johannesburg`, for a typical set of calibrations.