

COMPUTER SCIENCE

Qubit readout error mitigation with bit-flip averaging

Alistair W. R. Smith*, Kiran E. Khosla, Chris N. Self, M. S. Kim

Quantum computers are becoming increasingly accessible and may soon outperform classical computers for useful tasks. However, qubit readout errors remain a substantial hurdle to running quantum algorithms on current devices. We present a scheme to more efficiently mitigate these errors on quantum hardware and numerically show that our method consistently gives advantage over previous mitigation schemes. Our scheme removes biases in the readout errors, allowing a general error model to be built with far fewer calibration measurements. Specifically, for reading out n -qubits, we show a factor of 2^n reduction in the number of calibration measurements without sacrificing the ability to compensate for correlated errors. Our approach can be combined with, and simplify, other mitigation methods, allowing tractable mitigation even for large numbers of qubits.

Copyright © 2021
The Authors, some
rights reserved;
exclusive licensee
American Association
for the Advancement
of Science. No claim to
original U.S. Government
Works. Distributed
under a Creative
Commons Attribution
License 4.0 (CC BY).

INTRODUCTION

Noisy-intermediate scale quantum (NISQ) computers (1) are running increasingly complicated algorithms on small to intermediate numbers of qubits (2–10). However, their usefulness continues to be limited by noise, leading to unreliable outputs. Error mitigation schemes (11, 12) compensate for errors through a combination of calibration measurements and postprocessing and have been applied to the benchmarking of NISQ hardware (13–16), quantum chemistry and solid-state physics problems (17–19), dynamical quantum simulations (10, 14), and demonstrations of quantum supremacy (4). They have been proposed to bridge the gap between current devices and future fault-tolerant error correction (20), which actively corrects errors in the quantum state. Error mitigation has already proven to be an important tool to reach previously unachieved benchmarks on existing hardware (2, 21).

Qubit readout is a significant source of error in quantum computing experiments. This is particularly true for the popular superconducting qubit architectures, which typically have per-qubit readout error probabilities of a few percent (8) [detailed information about readout error probabilities on current devices can be found through IBM's Qiskit platform (22)]. In practice, the measurement errors on transmon-based devices are additionally complicated by effects such as bias toward certain states and cross-talk-induced correlations (23–25). Furthermore, quantum experiments often involve measuring many qubits at a time (26, 27), compounding the impact of readout errors. Together, these effects make readout errors a significant hurdle to scaling up NISQ computation.

Readout error mitigation schemes combine an error model with calibration measurements. The calibrated model is then used to infer the “error-free” result of an experiment (11, 24, 28–41). The quality of the mitigation strongly depends on the choice of error model; however, there is a trade-off between model complexity and calibration cost. Simple models, for example, those assuming qubit-wise-independent errors (2, 8), require fewer calibration measurements but may not capture the true error process. In contrast, using fewer assumptions leads to a more general error model but at the cost of potentially requiring a prohibitive number of calibration measurements (22). Here, we present a scheme that addresses both these problems, giving a lossless reduction in error model complexity and introducing a single, model agnostic, calibration step. This allows the most suitable model to be chosen a posteriori.

Here, we introduce bit-flip averaging (BFA), a scheme that uses random bit-flips to simplify the effective error process. We analytically show that averaging over these random bit-flips allows one to more efficiently parameterize and estimate readout errors. The error process under BFA admits convenient mathematical symmetries that greatly simplify the inference of error-free experimental results. We compare our approach to full mitigation and tensor product noise (TPN) models and show that BFA outperforms both. The bit-flips introduced by our method can be uniquely inverted, allowing for mid-circuit measurement and feed-forward algorithms (22, 42) experimental overhead, requiring only a layer of single qubit gates and classical postprocessing.

Imperfect multi-qubit measurements can be effectively modeled as a classical process (16, 28, 43–45). This can be understood as a probabilistic corruption of the error-free result. Assuming that the measurements (in the computational basis) will be performed across a constant number of qubits, this model is expressed in terms of a response matrix $\mathbf{M}$ such that $M_{\sigma\sigma'} = p(\sigma | \sigma')$, is the probability of reading out σ given that the error-free outcome should have been σ' . The observed outcome probabilities $\mathbf{p}_{\text{obs}}$ are given by the action of the response matrix on the error-free probabilities $\mathbf{p}_{\text{true}}$

$$\mathbf{p}_{\text{obs}} = \mathbf{M}\mathbf{p}_{\text{true}} \quad (1)$$

In general, the matrix $\mathbf{M}$ is not symmetric as readout on many devices is biased toward some states (23). Our protocol uses random bit-flips to symmetrize the response matrix, averaging out the biases. This markedly reduces the number of parameters required to define this matrix; this reduction is $\mathcal{O}(2^{2n}) \rightarrow \mathcal{O}(2^n)$ for n read-out qubits. It also simplifies the matrix inversion task required to find $\mathbf{p}_{\text{true}}$.

The calibration step involves estimating $\mathbf{M}$ by preparing and measuring each of the computational basis states. The k th column of $\mathbf{M}$ is the vector of measurement outcome probabilities given an input computational basis state $|k\rangle$. This requires enough calibration shots to sufficiently determine 2^n (potentially) unique probabilities for each of the 2^n different $|k\rangle$, which is especially problematic if time-drifting errors necessitate frequent recalibrations. At worst, calibration costs scale as $\mathcal{O}(2^{2n})$; however, in practice, many of the error probabilities will be negligibly small (i.e., those for simultaneous errors on many qubits) and can be safely approximated as zero. We show in sample complexity and scaling of full BFA calibration that the number of calibration measurements needed to estimate a single distribution (column of $\mathbf{M}$) typically scales at a rate $\ll \mathcal{O}(2^n)$ (although

QOLS, Blackett Laboratory, Imperial College London, London, SW7 2AZ, UK.

*Corresponding author. Email: alistair.smith18@imperial.ac.uk

still exponentially in n). Nevertheless, even if each distribution can be described with a marked reduced set of probabilities, there are still exponentially many distributions (input states) to estimate.

Once $\mathbf{M}$ is estimated, readout errors are typically mitigated by either inverting $\mathbf{M}$ or by solving a constrained linear optimisation problem [minimizing $(\mathbf{p}_{\text{obs}} - \mathbf{M}\mathbf{p})^2$ over $\mathbf{p}$, subject to physical probabilities]. We note that both problems quickly become intractable with increasing numbers of qubits.

In practice, there will be some underlying structure to the readout error distributions. Several proposals have taken advantage of this by making assumptions about the error process (16, 28–30). A common and effective choice of simplified model assumes that the readout errors for each qubit are independent, yielding the so-called TPN model (2, 16, 46). This simplification allows the response matrix to be given in terms of $2n$ single qubit error probabilities $\{p_{(1|0)}^i\}$ and $\{p_{(0|1)}^i\}$. For TPN, the response matrix $\mathbf{M}_{\text{TPN}}$ is the tensor product of single-qubit response matrices

$$\mathbf{M}_{\text{TPN}} = \bigotimes_i \begin{pmatrix} p_{(0|0)}^i & p_{(0|1)}^i \\ p_{(1|0)}^i & p_{(1|1)}^i \end{pmatrix} \quad (2)$$

where $p_{(s|s')}^i$ is the probability that the i th qubit reads out s given the error-free readout should have been s' [and $p_{(0|0)}^i = 1 - p_{(1|0)}^i$ etc.]. The TPN model can be calibrated more efficiently than the full scheme as $\{p_{(1|0)}^i\}$ and $\{p_{(0|1)}^i\}$ can be found by sampling only the input states $|0\dots 0\rangle$ and $|1\dots 1\rangle$, respectively. The inverse of $\mathbf{M}_{\text{TPN}}$ is now tractable and is simply the tensor product of inverse single-qubit response matrices.

On real devices multi-qubit readout errors can be correlated (16, 23) (through cross-talk effects), limiting the accuracy of many simplified models. Alternative approaches have been proposed to deal with correlated errors in a scalable way, e.g., using continuous Markov processes (16) or via cumulant expansion (44). These methods extend the TPN approximation by characterizing the readout errors in terms of single qubit and pairwise (between physically/frequency close qubits) correlated error terms. Although we do not consider these models directly here, our BFA proposal naturally extends to these correlation-extended models. Furthermore, as our scheme eliminates the bias in the readout errors toward certain

states, it allows for these models to be expressed in terms of fewer parameters. This simplification allows these models to be calibrated with fewer measurements and thereby mitigate readout errors more efficiently.

RESULTS

Averaging out readout errors with bit-flips

Our BFA method yields a greatly simplified and more easily measurable response matrix without sacrificing the ability to capture correlated readout errors. By applying random pre- and postmeasurement bit-flips, we completely symmetrize the response matrix and remove readout biases. The process is qualitatively similar to the randomized benchmarking techniques (47) that are often used to efficiently quantify gate errors. Methods that tackle state-dependent bias have been proposed, e.g., the “Static Invert-and-Measure” scheme (30); however, the scheme introduced here allows for more active and efficient response matrix–based mitigation to be used. Our scheme also provides computational advantages in applying the mitigation; we give an analytic formula for the inverse of the simplified response matrix under BFA that can be calculated using only vector-matrix multiplication.

Following the standard response matrix approach, we assume that every measurement will be performed across a fixed number of qubits. Mitigation happens in two stages: a calibration stage where the response matrix is measured, and an experimental stage in which readout errors are mitigated using this response matrix. In each shot of the experiment, we bit-flip random qubits before measuring them and then invert the bit-flip in the (classical) readout for the corresponding qubits (Fig. 1). We repeat and average over this process, randomly selecting different qubits to bit-flip each shot. As we will show, this simplifies the measured effective response matrix (measurement of this is shown in Fig. 1A). The bit-flipped qubits are chosen uniformly at random per shot, and the bit-flips are implemented with an $\hat{X}$ gate. Here and in the following, we assume the gate errors introduced by $\hat{X}$ are negligible. By continuing the random bit-flips and classical correction when performing a quantum experiment, the readout errors continue to be simplified, allowing for easier mitigation (Fig. 1B). In the absence of any readout errors,




Fig. 1. Example of BFA for four qubits. (A) Determining the response matrix requires applying $\hat{X}$ gates to randomly chosen qubits (red squares) while leaving others unflipped (cyan squares). The bit-flipping is inverted in the classical readout result (red and cyan circles), leaving the resulting measurement invariant under the full process. Averaging over many random choices of bit-flipped qubits means that the effective readout error process is uniquely determined by considering only the logical zero state. (B) To use the BFA response matrix to mitigate errors in an experiment, it is crucial the experiment sees the same effective measurement process; this is simply achieved by applying BFA to the experimental measurement process. If required, the $\hat{X}$ gates can be inverted (transparent squares), leaving the quantum state invariant under BFA.

BFA has no effect on the readout results. Last, if necessary, we can include another set of $\hat{X}$ gates to the postmeasurement flipped qubits, ensuring a consistent a posteriori state.

Without loss of generality, we describe the measurement operation and corresponding response matrix using Kraus operators. For n qubits, we do this in terms of a set of 2^n measurement Kraus operators $\{A_\sigma\}$ where

$$A_\sigma = \sum_{\sigma'} \sqrt{p(\sigma|\sigma')} |\sigma'\rangle\langle\sigma'| \quad (3)$$

which was chosen in such a way to recover Eq. 1. This operator A_σ corresponds to an uncertain measurement (with $\{|\sigma'\rangle\}$ being computational basis vectors), yielding a classical readout bit-string σ . We note that A_σ corresponds to a quantum noise-limited measurement operator. An additional incoherent classical assignment error can be included; however, this also gives a response matrix of the form in Eq. 1. Our BFA scheme yields the same result for any combination of these two processes.

Eq. 1 is recovered by considering the probability $p_{\text{obs}}(\sigma) = \text{Tr}\{A_\sigma^\dagger A_\sigma \rho\}$ that a measurement of a quantum state ρ yields a readout σ

$$p_{\text{obs}}(\sigma) = \sum_{\sigma'} p(\sigma|\sigma') p_{\text{true}}(\sigma') \quad (4)$$

where $p_{\text{true}}(\sigma') = \langle\sigma'| \rho |\sigma'\rangle$ is the probability of an error-free measurement of ρ to yield the state $|\sigma'\rangle$. We identify this sum as the matrix equation $p_{\text{obs}} = \mathbf{M} p_{\text{true}}$.

BFA can be taken into account by applying the relevant bit-flipping operations and adjusted readout results directly to the measurement operators. We quantify a bit-flip in terms of a binary string s such that the unitary operator applied to the qubits is $\hat{X}^{(s)} = \bigotimes_i \hat{X}^{s_i}$, a tensor product of Pauli $\hat{X}$ operators and identity operators where s_i is the i th bit of s , e.g., $\hat{X}^{(01)} = 1 \otimes \hat{X}$. For an n -qubit measurement, we choose a random bit-string s with probability $1/2^n$, and given a bit-string readout σ , the corresponding measurement operator is

$$A_{s,\sigma} = \frac{1}{\sqrt{2^n}} \sum_{\sigma'} \sqrt{p(\sigma \oplus s|\sigma')} |\sigma' \oplus s\rangle\langle\sigma'| \quad (5)$$

where $|\sigma' \oplus s\rangle = \hat{X}^{(s)} |\sigma'\rangle$ and $\sigma' \oplus s$ is the bit-wise addition of s and σ' . We now consider how averaging over s changes the readout error process. From here, averaging over s is denoted by a tilde. The (s -averaged) probability of observing σ is

$$\tilde{p}_{\text{obs}}(\sigma) = \frac{1}{2^n} \sum_{s,\sigma'} p(\sigma \oplus s|\sigma') p_{\text{true}}(\sigma') \quad (6)$$

allowing us to identify a new response matrix $\tilde{\mathbf{M}}$ (such that $\tilde{p}_{\text{obs}} = \tilde{\mathbf{M}} p_{\text{true}}$) that describes the readout errors under BFA. The elements of this new matrix are simply the s -averaged conditional probabilities $\tilde{p}(\sigma|\sigma')$ under BFA and are given by

$$\tilde{M}_{\sigma\sigma'} = \frac{1}{2^n} \sum_s p(\sigma \oplus s|\sigma' \oplus s) \quad (7)$$

From this equation, we see that $\tilde{p}(\sigma \oplus s|\sigma' \oplus s) = \tilde{p}(\sigma|\sigma')$, and so, we have arrived at a far simpler, symmetrized error model with $2^n - 1$ parameters, instead of the $2^n(2^n - 1)$ in $\mathbf{M}$.

Thanks to this symmetry, we can express the whole response matrix under this bit-flipping protocol $\tilde{\mathbf{M}}$ in terms of just the parameters in its first column. The conditional index in the response matrix can now be dropped as

$$\tilde{M}_{\sigma\sigma'} = \tilde{p}_{(\sigma \oplus \sigma')|0} \equiv \tilde{p}_{(\sigma \oplus \sigma')} \quad (8)$$

Physically, this is because any computational basis state is equally likely to be bit-flipped into any other basis state. As it is the bit-flipped state that is measured by the physical apparatus, the effective error probability is averaged across all inputs, removing any bias toward certain states. This removal of bias gives a huge practical advantage of BFA over normal response matrix error mitigation; calibrating the entire symmetrized error matrix only requires measurement of the probabilities in its first column, which is done with just one input state $|0\dots 0\rangle$. As no assumptions are made about whether the errors are correlated, these correlations can be effectively dealt with using this scheme. In particular, one can infer correlations by comparing the relative probabilities of different errors in the calibration data.

For practical reasons, readout errors are often biased toward certain states. With superconducting qubits, the readout of qubits in the state $|0\rangle$ is typically more reliable than for the state $|1\rangle$. Adaptive mitigation schemes (30, 37) have been proposed to more effectively compensate for readout errors by exploiting this bias. These involve applying $\hat{X}$ gates to certain qubits before measurement to maximize the proportion that are measured in the state $|0\rangle$, reducing the probability of readout errors and allowing the error-corrupted measurement distributions to be estimated more easily. However, these schemes only yield improvements for certain states and require the circuits for suitable states to be run twice (without and with adaptive $\hat{X}$ gates) while also requiring the full response matrix to be found. By symmetrizing the readout errors, BFA increases the probability of some states (those with a high proportion of $|0\rangle$'s) being read out incorrectly, but this is balanced by effective error probability for other states (with a high proportion of $|1\rangle$'s) being reduced. The balancing of readout errors coupled with the factor of 2^n reduction in the cost of estimating the response matrix and the lack of circuit-specific calibration measurements allow BFA to compete with these adaptive schemes while being applicable to a broader set of states.

The probabilities $\tilde{p}_{(s)}$ have a convenient physical interpretation as the average probability that an error with syndrome S (the bit-string identifying which qubits are read out incorrectly) occurs, e.g., $\tilde{p}_{(1011)}$ is the probability that readout errors occur simultaneously on the 0th, 1st and 3rd qubits. The error matrix $\tilde{\mathbf{M}}$ is symmetric about both its diagonal and antidiagonal, allowing it to be decomposed into a compact form

$$\tilde{\mathbf{M}} = \sum_s \tilde{p}_{(s)} X^{(s)} \quad (9)$$

where we have used the same notation for the matrix $X^{(s)}$ as for the operator $\hat{X}^{(s)}$. This is due to Eq. 7 being invariant under $\sigma, \sigma' \rightarrow \sigma \oplus s, \sigma' \oplus s$ (i.e., invariant under $\tilde{\mathbf{M}} \rightarrow X^{(s)} \tilde{\mathbf{M}} X^{(s)}$). To take advantage of this sparse representation, one must continue to perform the bit-flipping and classical correction during experiments. This requirement adds negligible overhead as single qubit bit-flips can typically be performed with very high fidelity or are completely free if combined with an existing gate. The corresponding bit-flip of the measurement output requires only classical Boolean logic.

The decomposition in Eq. 9 gives us an advantage in both mitigation strategies (response matrix inversion and constrained least squares minimization). For the least-squares method, Eq. 9 tells us with what probability $\tilde{p}_{(s)}$ we expect a given readout string to be

corrupted by the binary addition of S . If many of these probabilities are zero (or negligibly small and so can be set to zero), then this would allow us to use a sparse matrix representation of $\tilde{\mathbf{M}}$, allowing the optimization problem to be solved more easily.

The matrix inverse mitigation strategy requires $\tilde{\mathbf{M}}^{-1}$ to be found. As $\tilde{\mathbf{M}}$ contains only tensor products of the identity and Pauli X matrix, all these terms can be simultaneously diagonalized by the application of the Hadamard matrix H . As we show in the section on “Derivation of inverse BFA-simplified response matrix,” this gives the vector of eigenvalues

$$\lambda = \sqrt{2}^n H^{\otimes n} \tilde{p} \quad (10)$$

where $\tilde{p}$ is a vector of the probabilities $(\tilde{p})_s \equiv \tilde{p}_{(s)}$ (corresponding to the first column of $\tilde{\mathbf{M}}$). Like $\tilde{\mathbf{M}}$, the inverse must be symmetric about both its diagonal and antidiagonal, meaning that it can also be decomposed onto Pauli X matrices, i.e., the form given in Eq. 9. As shown in the section on “Derivation of inverse BFA-simplified response matrix,” the inverse BFA-simplified response matrix is given in terms of the vector of reciprocal eigenvalues $\lambda^{-1} \equiv (1/\lambda_0, \dots, 1/\lambda_{n-1})$ by

$$\tilde{\mathbf{M}}^{-1} = \sum_s \tilde{q}_{(s)} X^{(s)}, \quad \tilde{q} = \frac{1}{\sqrt{2}^n} H^{\otimes n} \lambda^{-1} \quad (11)$$

This shows another clear advantage to bit-flipping over the full mitigation approach. Like with $\tilde{\mathbf{M}}$, we only need to find the elements of the inverse’s first column, and this can be done with simple matrix multiplication (as opposed to a computationally costly general matrix inverse).

As the primary function of BFA is to average out bias in the readout errors toward certain measurement outcomes, we can also use it to further simplify other simplified error models. As an example of how BFA can simplify other approximate measurement error mitigation protocols, we can consider how the TPN model transforms under bit-flipping. Under bit-flipping, the biases of the qubit-wise readout errors are averaged out, meaning that the BFA-symmetrized TPN matrix $\tilde{\mathbf{M}}_{\text{TPN}}$ is given by

$$\tilde{\mathbf{M}}_{\text{TPN}} = \bigotimes_i \begin{pmatrix} 1 - \tilde{p}_i & \tilde{p}_i \\ \tilde{p}_i & 1 - \tilde{p}_i \end{pmatrix}, \quad \tilde{p}_i = \frac{p_{(1|0)}^i + p_{(0|1)}^i}{2} \quad (12)$$

Combining the TPN model with BFA provides two main advantages: the first being that the number of parameters to estimate for the combined model on n qubits is n instead of $2n$. The second comes in estimating $\{\tilde{p}_i\}$; these probabilities can be measured by preparing the state $|0\dots 0\rangle$. This is the same experimental procedure as is required for calibration of the full BFA matrix $\tilde{\mathbf{M}}$, and so, a single set of calibration results can be used for both models. In this example, a TPN + BFA model could be calibrated first and its predictions for the different error probabilities checked against the calibration data. If this proves unsuitable, then a larger more general model could be used without requiring any further calibration measurements.

The information contained in the BFA calibration measurements of $|0\dots 0\rangle$ fully describes the response matrix, and so, it can be used to calibrate any response matrix-based approach. This means that one is not forced to make any assumptions about the model (e.g., independent errors, pairwise correlations, or a full model) before calibration.

This flexibility potentially allows for readout error mitigation to be performed even for large numbers of qubits, provided any

correlations in the readout errors have some degree of locality. If the qubits can be grouped into disjoint sets such that there are no intergroup correlations (for example, if correlations only occur between qubits coupled to the same readout cavity), then an expanded TPN-like model could be used in which each group has its own full response matrix. The response matrix for a measurement of all the qubits would then be given by a tensor product of those for each group. BFA would allow this model to be calibrated using only a single measured input state at a cost scaling at worst as $\mathcal{O}(2^k)$, where $k < n$ is the number of qubits in the largest grouping. While the groupings could initially be chosen on the basis of some knowledge of the device (e.g., by readout cavity, operating frequency, or some spatial consideration), BFA would allow this grouping to be changed retroactively to match the calibration data. In the Supplementary Text, we show an example calibration process for a sparse response matrix that exhibits correlations between only a restricted number of qubits. We demonstrate how an appropriate choice of model can make both the calibration and subsequent mitigation more accurate.

Sample complexity and scaling of full BFA calibration

To use BFA most effectively, it should be combined with an error model that best balances the trade-off between model expressibility and calibration cost. However, it is helpful to estimate the worst-case cost to calibrate a BFA model using the full symmetrized response matrix as in Eq. 9, under some physically motivated assumptions. On real devices, we expect that even when correlations are taken into account, the probabilities of errors occurring on many qubits simultaneously are negligible and can be neglected. This effectively reduces the number of parameters that must be estimated to find $\tilde{p}$ and, thereby, $\tilde{\mathbf{M}}$. The number of probabilities that give nonnegligible contributions to $\tilde{p}$ will provide an indication of the number of calibration measurements required to estimate $\tilde{\mathbf{M}}$, and equivalently the cost to estimate a single column of $\tilde{\mathbf{M}}$.

To give an idea of how many parameters must be retained in $\tilde{p}$, we consider a TPN model for an n -qubit readout with constant single-qubit readout error probabilities p_e for all qubits. While this neglects any correlations in the errors, we expect that these would act as a modest correction to the TPN model and so would not greatly affect the calibration cost. For a conservative scaling estimate, p_e could be the largest measured single-qubit readout error probability for the device in question. We then calculate the number N of error probabilities that must be retained so that their cumulative probability reaches above a threshold $\sum_{i=0}^{N-1} \tilde{p}_{(i)}^\downarrow > 1 - \epsilon$ (where $\tilde{p}_{(i)}^\downarrow$ is $\tilde{p}$ sorted in descending order).

Under these assumptions, the number of qubits Q that experience a readout error is binomially distributed $Q \sim B(n, p_e)$. From this, we can calculate the highest weight of error k (i.e., largest number of qubits that are simultaneously read out incorrectly) that must be retained such that the cumulative probability up to k is greater than $1 - \epsilon$. This is given by $k = S_B^{-1}(\epsilon; n, p_e)$, where $S_B(k; n, p_e) = 1 - \Pr(Q \leq k)$. Thus, N is the number of possible readout errors with weight less than or equal to k

$$N = \sum_{i=0}^k \binom{n}{i} \quad (13)$$

While this sum does not have a closed form solution, for $k/n \leq 1/2$ (expected for large n and $p_e \ll 1$), it is bounded (48) by

$$\frac{1}{\sqrt{8k(1-k/n)}} 2^{nH(k/n)} \leq N \leq 2^{nH(k/n)} \quad (14)$$

where $H(p) = -p \log_2(p) - (1-p) \log_2(1-p)$ is the binary entropy function. This means the number of measurements needed to calibrate a completely general $\tilde{\mathbf{M}}$ is expected to scale at worst as $\mathcal{O}(2^{nH(k/n)})$. Although the number of outcomes that significantly contribute to the total probability is typically much less than the 2^n worst case, the required number of terms still scales exponentially with n .

For large n , we can examine the scaling of N by approximating Q with a normal distribution $Q \sim \mathcal{N}(np_e, np_e(1-p_e))$ giving

$$k \approx np_e + \sqrt{np_e(1-p_e)} S_N^{-1}(\epsilon) + 0.5 \quad (15)$$

where $S_N(k) = 1 - \text{Erf}(k)$, and we have applied a continuity correction of 0.5. Using this approximation, we see that in the limit of large n , the lower bound tends to $2^{nH(p_e)}/\sqrt{8np_e(1-p_e)}$, and so, we require at least $\mathcal{O}(2^{nH(p_e)})$ samples to calibrate $\tilde{\mathbf{M}}$. As the upper bound on N tends to $2^{nH(p_e)}$, we can also identify a very rough rule of thumb for the error rates under which full mitigation is tractable, e.g., $nH(p_e) < 10$.

We can further quantify the sample complexity of estimating a typical $\tilde{\mathbf{M}}$ by applying bounds for the estimation of arbitrary discrete distributions. A useful bound for this is given in (49) and states that for an N -outcome discrete distribution with true probabilities $\mathbf{p}$, the expected total variation distance $\delta(\mathbf{p}, \hat{\mathbf{p}}^{(m)}) = \sum_i |p_i - \hat{p}_i^{(m)}|$ between $\mathbf{p}$ and the empirical distribution after m samples $\hat{\mathbf{p}}^{(m)}$ is bounded by

$$\mathbb{E}[\delta(\mathbf{p}, \hat{\mathbf{p}}^{(m)})] \leq \sqrt{\frac{N}{m}} \quad (16)$$

This shows that we can estimate $\mathbf{p}$ to an expected accuracy ϵ using $m = N/\epsilon^2$ shots. From (49), we also have a concentration bound for this inequality. The probability $P[\delta(\mathbf{p}, \hat{\mathbf{p}}^{(m)}) > \epsilon]$ that we observe a distribution after m shots that differs from the true distribution with $\delta > \epsilon$ is bounded for $\epsilon \geq \sqrt{N/m}$ by

$$P[\delta(\mathbf{p}, \hat{\mathbf{p}}^{(m)}) > \epsilon] \leq \exp\left(-\frac{m}{2}\left(\epsilon - \sqrt{\frac{N}{m}}\right)^2\right) \quad (17)$$

This implies that to be sure that we have estimated $\mathbf{p}$ to within an accuracy ϵ with a failure probability less than γ , we need at least $m > \left(\sqrt{N} + \sqrt{2\ln \frac{1}{\gamma}}\right)^2 / \epsilon^2$ shots.

Therefore, estimating an arbitrary n -qubit response matrix using BFA to an accuracy ϵ with failure probability less than γ requires at least $\tilde{m}(n, \epsilon, \gamma) \equiv (2^{n/2} + \sqrt{2\ln \frac{1}{\gamma}})^2 / \epsilon^2$ or $\mathcal{O}(2^n/\epsilon^2)$ shots. With full mitigation, all 2^n columns must be estimated independently, meaning a worst-case sample cost of $2^n \times \tilde{m}(n, \epsilon, \gamma)$ and an $\mathcal{O}(2^{2n}/\epsilon^2)$ sample complexity.

These worst-case sample complexities correspond to a completely general $\mathbf{M}$ in which all 2^n outcomes in each column are significant. From Eq. 14, the typical number of significant probabilities N given a representative single-qubit readout error probability p_e is bounded by $N \leq 2^{nH(k/n)}$ (where k is defined as before). This yields to give a more typical estimate for the sampling cost under BFA of $\tilde{m}(nH(k/n), \epsilon, \gamma)$ shots and, for large n , a sampling complexity of $\mathcal{O}(2^{nH(p_e)}/\epsilon^2)$ shots.

The exact sampling cost depends on the values of the various error probabilities. In the above analysis, we assume that the sampling cost under BFA can be reasonably well approximated by the cost of naively sampling a symmetric TPN response matrix with

unbiased and equal per-qubit readout errors. While this is a reasonable approximation for BFA, it ignores the readout error biases encountered in full mitigation schemes. With the typical biases observed for superconducting qubits, the columns (output distributions) corresponding to input states with large numbers of qubits in the state $|0\rangle$ are more sharply peaked than for states with many qubits in $|1\rangle$ (as the former will have lower readout error probabilities). Sharper distributions [e.g., $p_{(\sigma|0\dots 0)}$] can be estimated with fewer samples than flatter ones [e.g., $p_{(\sigma|1\dots 1)}$], so the required number of samples to reach a given accuracy will vary across the columns of $\mathbf{M}$. Because of the averaging that occurs, the sample complexity of estimating the single distribution $\tilde{p}_{(\sigma)}$ required for BFA will usually lie somewhere between the best-case cost of estimating $p_{(\sigma|0\dots 0)}$ and worst case of $p_{(\sigma|1\dots 1)}$.

While the cost of calibrating $\tilde{\mathbf{M}}$ is lower bounded by the cost of estimating the cheapest column of $\mathbf{M}$, the calibration will be significantly cheaper than for $\mathbf{M}$'s most expensive column. Because full mitigation requires estimating $\sim 2^{n-1}$ columns that are more sampling expensive than the single column required for BFA and $\sim 2^{n-1}$ that are cheaper, we expect the total reduction in calibration cost brought by BFA to average out to a factor of $\sim 2^n$, despite the variable column cost in $\mathbf{M}$ caused by readout error biases.

This analysis considers a naive estimation of $\mathbf{M}$ and $\tilde{\mathbf{M}}$ using a full dense matrix model rather than making any assumptions about correlations present. Simplified models (e.g., TPN) that require estimation of fewer parameters will incur a smaller sampling cost than the general model considered here. However, because BFA ensures symmetry in the effective response matrix thereby reducing the number of parameters needed to describe it, we expect that it will yield a reduction in calibration cost in practically all realistic cases. The bounds we use for the analysis of a generic $\mathbf{M}$ are valid for the estimation of arbitrary discrete probability distributions. Therefore, they are also applicable to the sampling cost of estimating the noisy (premitigation) measurement distribution for an arbitrary quantum circuit—i.e., this requires at most $\mathcal{O}(2^n/\epsilon^2)$ shots with or without BFA.

Simulated measurement of response matrices

To obtain realistic readout error models for our simulations, we measured full response matrices on *ibmq_manhattan* (50) for one to eight qubit readouts (data taken on 1 December 2020). To minimize sampling error, we used 2^{16} shots per computational basis state (per column of $\mathbf{M}$). Taking these measured response matrices as “exact,” we used $\mathbf{M}$ to simulate the readout error process as described in Eq. 1. This effectively simulates a full on-device readout process from which we benchmark various BFA strategies. The different schemes used for our simulations are summarized in Table 1.

Figure 2 demonstrates the advantage of using BFA, comparing the exact and calibrated 4-qubit response matrices. For both schemes, a budget of 100×2^4 shots were used to estimate the response matrix. Here, the BFA advantage is immediately obvious: For full mitigation, this budget must be divided between the 2^4 input states that are measured, while for BFA, all 100×2^4 shots are used to measure $|0000\rangle$ (with shot-by-shot bit-flipping). For this budget, BFA produces an accurate estimate of its target response matrix $\tilde{\mathbf{M}}$, in contrast to full mitigation's poor estimate of $\mathbf{M}$. A more accurate response matrix allows for more effective error mitigation in the final experiment.

We stress that the calibration shown in Fig. 2 assumes a full dense response matrix rather than a TPN model. It is likely that the

response matrix used for Fig. 2 (measured on-device) can be well approximated using a TPN model, allowing it to be sampled more efficiently (as only marginal distributions per qubit rather than full 2^n -outcome distributions must be estimated). However, to showcase the advantage BFA brings over the most general approach for generic response matrices, we have used the full dense model here that requires the 2^n outcome measurement distributions for all 2^n computational basis state inputs to be measured.

Figure 2 shows BFA greatly reduces the amount of calibration data required to obtain a faithful description of the response matrix. We can quantify this reduction by examining the number of shots required for the estimated response matrix to converge to that of the underlying error model. To measure the “closeness” of the estimated response matrix $\mathbf{M}_{\text{obs}}$ to the true response matrix $\mathbf{M}$, we use the average of the column-wise classical fidelities (as each column is a distinct probability distribution), which for distributions $\mathbf{p}$ and $\mathbf{q}$ is defined as $\sum_i \sqrt{p_i q_i}$. Our figure of merit $\mathcal{F}_{\mathcal{M}}$ (which we will refer to as the “response matrix fidelity”) for the closeness of two response matrices $\mathbf{M}$ and $\mathbf{N}$ is then given by

$$\mathcal{F}_{\mathcal{M}}(\mathbf{M}, \mathbf{N}) = \frac{1}{2^n} \sum_{ij} \sqrt{M_{ij} N_{ij}} \quad (18)$$

Figure 3A compares the fidelity between the exact and estimated response matrices for the different models considered. For a fair comparison, the same total number of shots are used to estimate each model, and the fidelities are averaged over 50 trials. The exact response matrix used here is one we directly measured for 5-qubits on *ibmq_manhattan* and so provides a realistic picture of how the schemes fare on current devices. For full mitigation and the TPN model, the target is the exact response matrix used to simulate the error process, while for the BFA and the combination of BFA + TPN, the target is the symmetrized version of the exact matrix (as in Eq. 9).

In comparison to the simplified schemes, full mitigation requires a far greater number of measurements to converge to a good recreation of the true response matrix. BFA very quickly converges to the maximum response matrix fidelity, requiring around two orders of magnitude fewer calibration shots than full mitigation to reach comparable fidelities. The schemes using TPN converge to $\mathcal{F}_{\mathcal{M}} \rightarrow 1$, indicating the TPN assumption (i.e., independent readout errors) is justified for this particular experimentally measured response matrix. The combination of TPN and BFA yields the best fidelity with the fewest calibration shots as it is the most parameter-efficient model that sufficiently captures the true readout error process.

As discussed in the section on “Averaging-out readout errors with bit-flips,” the same calibration measurements are required to

infer the BFA and BFA + TPN response matrices. In this instance, the combination of TPN + BFA was the fastest model to yield a useful description of the error process. However, in situations where significant cross-talk leads to correlated readout errors, the TPN approximation of independent errors becomes invalid. It is therefore helpful to also consider the scaling of the different schemes in a case where the TPN assumption manifestly fails.

Response matrix measurement with cross-talk

To provide a toy model for cross-talk-induced correlated errors, we change the experimentally measured response matrix $\mathbf{M}$ to artificially amplify correlations between readout errors on adjacent qubits. Specifically, we boost the probability of particular syndromes to get a new response matrix $\mathbf{M}_\gamma$. The strength of the amplification is parameterized by a boosting factor γ (see section on “Boosting correlations in readout errors” for details). Figure 3B shows how $\mathcal{F}_{\mathcal{M}}$ scales with the number of calibration shots. The 5-qubit response matrix used here is that of Fig. 3A, but correlation boosted with $\gamma = 20$. Here, $\gamma = 20$ was chosen to highlight the contrast between the TPN-based schemes and those capable of dealing with correlated errors.

We see that in the presence of correlated errors, the BFA provides a clear advantage, converging to the optimal fidelity with far fewer shots (around a factor $2^n = 32$) than required for full mitigation. The two TPN-based models saturate at fidelities well below the




Fig. 2. Illustrating the sampling advantage of BFA over full mitigation. Example response matrix plots from four qubits showing the exact matrix (left) and a finite shot (100×2^4 shots) estimation (right). For full measurement (top), the bias is manifested by lack of diagonal symmetry, and the many independent parameters increase the sampling error of its estimate. The BFA response matrix (bottom) exhibits many symmetries, and sampling error is nearly imperceptible. The exact response matrix for the error process (top left) is measured from *ibmq_manhattan* using 2^{20} total shots (2^{16} per input state), and the exact BFA response matrix (bottom left) is calculated using Eq. 8.

Table 1. A summary of the mitigation schemes considered here, highlighting key characteristics.

Error mitigation scheme summary			
Scheme	Num. measured states	Num. free parameters	Correlated errors
Full mitigation	2^n	$2^n(2^n - 1)$	Yes
TPN	2	$2n$	No
BFA	1	$2^n - 1$	Yes
BFA + TPN	1	n	No




Fig. 3. Scaling of the response matrix fidelities with total calibration shots. (A) Readout errors simulated using M measured for $n = 5$ qubits on *ibmq_manhattan*. The fidelity, Eq. 18, is calculated relative to the exact response matrix M (Full Mitigation and TPN) or $\tilde{M}$ (BFA and BFA + TPN). (B) As in (A), using the same response matrix but with artificially boosted correlations (with boosting factor $\gamma = 20$; see section on “Boosting correlations in readout errors”). For both plots, the fidelities are averaged over 50 repeats, with error bars showing the middle 95% percentile of data.

optimal value of 1 as they cannot account for correlated errors by construction. The fidelities obtained take longer to saturate than in Fig. 3 because the underlying readout error process with boosted correlations is more complicated than the original response matrix (which can be accurately described with TPN).

The advantage of BFA over full mitigation will become increasingly apparent as more qubits are measured, as demonstrated in Fig. 4. The plot shows the fidelity of ($\gamma = 20$) boosted response matrices taken on *ibmq_manhattan* estimated for different numbers of qubits. At each (n -qubit) point, $2^n \times 100$ simulated shots were used to estimate the response matrix.

Again, the full mitigation scheme has to share the $2^n \times 100$ budget among the 2^n input basis states leading to increasingly severe sampling errors. While the TPN (and BFA + TPN) schemes have much lower sampling error, they both suffer by their inability to express correlated errors. This limit is less of a problem for BFA + TPN, as the BFA symmetrization helps reduce the effect of correlations, and the fewer free parameters further reduce sampling errors. In particular, we note the response matrix estimated by TPN + BFA is a better approximation to the γ -boosted response matrix than is managed by TPN alone. This hints at a further advantage of the BFA scheme; averaging over different error probabilities damps biased correlations in the error model, yielding an effective error model that is closer to TPN. Last, we note that full BFA obtains by far the best fidelity as it most effectively balances sampling errors and (correlation) model expressibility.

Simulating graph state fidelity measurement

For direct comparison with previous work, we consider the example given in (16). We demonstrate our BFA scheme in a practical context by considering the problem of measuring the fidelity of a linear graph state of varying numbers of qubits. Again, we compare combinations of full mitigation and TPN models with BFA on simulated measurements using the experimentally measured response matrices from *ibmq_manhattan*. For a linear array of n qubits with initial state $|+\rangle^{\otimes n}$, a linear graph state $|g_n\rangle$ is created by applying controlled- $\hat{Z}$ gates to adjacent qubits. This graph state has a stabilizer group S^n generated by the set of Pauli operators $G_i^n = Z_{i-1} X_i Z_{i+1}$ (dropping the Z_{-1} and Z_n operators for $i = 0, n$ respectively). The state fidelity can be measured averaging the expectation value of




Fig. 4. Demonstrating the advantage of BFA with increasing numbers of qubits n . For each n , the response matrix was measured on *ibmq_manhattan*, and the correlations between neighboring qubits boosted with $\gamma = 20$. This matrix was then used to sample $2n \times 100$ shots and estimate the response matrix fidelities. The error bars (sampling error) show the middle 95% percentile of 50 fidelity estimates.

elements in S^n . As a simplification, and to ensure we are making consistent comparisons, we only measure the generators G_i^n of the stabilizer group themselves, providing an approximation to the fidelity.

Figure 5A shows this approximation $\mathcal{F}_n = \sum_i \langle G_i^n \rangle / n$ of the fidelity found for simulations in which graph states of different sizes are prepared noiselessly and then measured with readout errors. Readout error mitigation is then applied using the four schemes. The response matrices used to simulate the readout errors were those experimentally measured on the *ibmq_manhattan*. For comparison, the fidelities with no mitigation applied and those for mitigation using the exact response matrix of the noise process are also shown. For the simulation of an n -qubit graph state, a budget of 100×2^n calibration shots were allowed for measurement of the model's response matrix, while 10^5 shots were used for each of the two circuits used to measure $\{G_i^k\}$.

The two schemes incorporating BFA perform very well, particularly with increasing system size, giving fidelities close to that achieved when mitigation is performed using the exact response matrix (i.e., perfect knowledge of the error process). However, given this calibration




Fig. 5. Demonstrating the effectiveness of readout error mitigation schemes for observable (graph state fidelity) estimation. (A) Simulated estimates of graph state fidelity for varying numbers of qubits. For an n qubit graph state, 100×2^n calibration shots were used to estimate the response matrix, and 10^5 shots were used for the circuits to measure the fidelities. The readout errors were generated from response matrices directly measured on *ibmq_manhattan*. For comparison, the mitigation-free and exactly mitigated (using the *ibmq_manhattan* response matrices) fidelities are also shown. (B) Root-mean-square fidelity error $\epsilon_n^{(\text{model})} = (\mathbb{E}[(\mathcal{F}_n^{(\text{model})} - \mathbb{E}[\mathcal{F}_n^{(\text{exact})}])^2])^{1/2}$ with respect to mean of exact mitigation. For both plots, the data are averaged over 50 repeats, with error bars indicating the middle 95% of the population.

budget, full mitigation performs significantly worse than the other schemes. As in previous simulations, the full scheme only gets 100 calibration shots for each of the 2^n input states, resulting in high sampling errors.

It is tempting to quantify the performance of the mitigation schemes in terms of how close their fidelity is to the optimal value of 1 as we are performing simulations without gate errors. However, because we are calibrating the response matrix and calculating expectation values on the basis of finite shots, noncorrectable sampling noise remains the limiting factor. Therefore, a scheme's performance should be compared against exact mitigation with the response matrix used to simulate the error process. This gives the fairest estimate of the maximum improvement possible with this classical postprocessing-based mitigation strategies.

Figure 5B shows the average root-mean-square fidelity error of the different models presented with respect to the mean “exact mitigation” results, $\epsilon_n^{(\text{model})} = (\mathbb{E}[(\mathcal{F}_n^{(\text{model})} - \mathbb{E}[\mathcal{F}_n^{(\text{exact})}])^2])^{1/2}$, where the expectation $\mathbb{E}[\cdot]$ is taken over the 50 trials. Again, we see that the full mitigation scheme consistently underperforms the others, suffering from far greater errors relative to exact mitigation. The TPN model performs similarly to the two bit-flipping schemes until $n > 5$, where TPN + BFA outperforms TPN alone, indicating biased correlations that BFA averages out (cf. Figs. 3B and 4). For larger system sizes, both BFA-based schemes manage to replicate the exact mitigation results with remarkable accuracy. The simplified schemes benefit from receiving exponentially many more measurement shots per calibration circuit than for full mitigation, markedly reducing sampling error. Having a more accurate response matrix means the mitigation yields results closer to what would be found if using the response matrix that exactly generates the error process.

DISCUSSION

We have presented BFA, an effective scheme for readout error mitigation on near-term qubit devices. We demonstrate that BFA can augment, and consistently outperform, other measurement error mitigation strategies as it always simplifies the underlying error model. This simplification allows the response matrix to be measured using far fewer resources than would otherwise be required. Furthermore, all

BFA-augmented error models are calibrated from the same measurements, allowing these to inform the choice of model. In particular, BFA + TPN, BFA + full matrix mitigation, and all other combined schemes only require measurement of the state $|0\dots 0\rangle$.

BFA works by applying bit-flips to random qubits (premeasurement) and subsequently undoing these bit-flips in the classical readout result. This greatly simplifies the observed measurement response matrix, removing all biases toward particular input states. This bias is separate from readout error correlations, and so, BFA does not impose any assumptions about the error process. The resulting response matrix admits a highly symmetric form. We derive a general analytic expression for its inverse and show that it can be calculated with vector-matrix multiplication.

We benchmark BFA using numerical simulations estimating response matrix and quantum state fidelities and examine the role of correlated readout errors. The simulations are based on the empirical response matrices measured precisely on an IBM quantum device. Our results show that BFA can accurately estimate the response matrix with as many or fewer measurement shots than required by other schemes. Furthermore, when readout error correlations are artificially boosted, we show BFA requires orders of magnitude fewer calibration shots to find an optimal error description. Last, we test the performance of BFA in a realistic task of measuring the (simulated) fidelity of a linear graph state in the presence of readout errors. In each case, BFA results in a more accurate fidelity estimate compared to non-bit-flipped counterparts.

Implementing our scheme on real devices is technically simple but is prohibitively impractical in the current version of Qiskit. Previous works have demonstrated classical error models accurately describe on-device readout errors (16, 28, 45), which encouragingly suggests that BFA will continue to surpass other mitigation strategies on physical devices. We stress that practical implementation only requires minor changes to device access, and the BFA method itself only adds (effectively free) quantum bit-flips and classical postprocessing.

BFA is a useful tool for NISQ era quantum computing, allowing noisy measurements to be mitigated even in the presence of significant readout error correlations. This provides more freedom in the fabrication of solid-state quantum devices, allowing more compact qubit layouts and greater connectivity. Efficiently mitigating

14. Y. Li, S. C. Benjamin, Efficient variational quantum simulator incorporating active error minimization. *Phys. Rev. X* **7**, 021050 (2017).
15. S. Zhang, Y. Lu, K. Zhang, W. Chen, Y. Li, J.-N. Zhang, K. Kim, Error-mitigated quantum gates exceeding physical fidelities in a trapped-ion system. *Nat. Commun.* **11**, 587 (2020).
16. S. Bravyi, S. Sheldon, A. Kandala, D. C. McKay, J. M. Gambetta, Mitigating measurement errors in multiqubit experiments. *Phys. Rev. A* **103**, 042605 (2021).
17. A. J. McCaskey, Z. P. Parks, J. Jakowski, S. V. Moore, T. D. Morris, T. S. Humble, R. C. Pooser, Quantum chemistry as a benchmark for near-term quantum computers. *npj Quantum Inf.* **5**, 99 (2019).
18. A. Kandala, K. Temme, A. D. Córcoles, A. Mezzacapo, J. M. Chow, J. M. Gambetta, Error mitigation extends the computational reach of a noisy quantum processor. *Nature* **567**, 491–495 (2019).
19. C. N. Self, K. E. Khosla, A. W. R. Smith, F. Sauvage, P. D. Haynes, J. Knolle, F. Mintert, M. S. Kim, Variational quantum algorithm with information sharing. *npj Quantum Inf.* **7**, 116 (2021).
20. E. T. Campbell, B. M. Terhal, C. Vuillot, Roads towards fault-tolerant universal quantum computation. *Nature* **549**, 172–179 (2017).
21. A. J. Zalcman, A. Derk, A. Ho, A. Opremcak, A. Korotkov, A. G. Petukhov, A. Bengtsson, A. Dunsworth, A. Megrant, A. Fowler, B. Pató, B. Chiaro, B. Villalonga, B. B. Buckley, B. Burkett, B. R. Foxen, C. Erickson, C. Neill, C. Quintana, C. Jones, C. M. Gidney, D. Eppens, D. Sank, D. Landhuis, D. A. Buell, D. Strain, D. Kafri, E. Farhi, E. Ostby, E. Lucero, E. Jeffrey, F. Kostritsa, F. C. Arute, H. Neven, I. Aleiner, J. Yao, J. R. McClean, J. P. Hilton, J. Chen, J. A. Gross, J. Bardin, J. Mutus, J. Atalaya, J. Campero, J. Kelly, K. Miao, K. Satzinger, K. Kechedzhi, K. Arya, L. Ioffe, M. Szalay, M. Giustina, M. Mohseni, M. Jacob-Mitos, M. M. Ewen, M. Trevithick, M. Neeley, M. P. Harrigan, M. B. Broughton, M. Newman, M. Y. Niu, N. Bushnell, N. Redd, N. Rubin, O. Naaman, O. Martin, P. V. Klimov, P. Laptev, P. Roushan, P. Yeh, R. Barends, R. Collins, R. Babbush, S. Hong, S. Demura, S. Harrington, S. Kim, S. Isakov, S. Boixo, T. White, T. E. O'Brien, T. Huang, T. McCourt, V. Smelyanskiy, V. Shvarts, W. Courtney, W. J. Huggins, W. M. Muczkiewicz, X. Mi, Y. Chen, Z. Jiang, Accurately computing electronic properties of materials using eigenenergies, arXiv:2012.00921 (2020).
22. G. Aleksandrowicz, T. Alexander, P. Barkoutsos, L. Bello, Y. Ben-Haim, D. Bucher, Francisco Jose Cabrera-Hernández, J. Carballo-Franquis, A. Chen, C.-F. Chen, J. M. Chow, A. D. Córcoles-Gonzales, A. J. Cross, A. Cross, J. Cruz-Benito, C. Culver, Salvador De La Puente González, Enrique De La Torre, D. Ding, E. Dumitrescu, I. Duran, P. Eendebak, M. Everitt, I. F. Sertage, A. Frisch, A. Fuhrer, J. Gambetta, B. G. Gago, J. Gomez-Mosquera, D. Greenberg, I. Hamamura, V. Havlicek, J. Hellmers, L. Herok, H. Horii, S. Hu, T. Imamichi, T. Itoko, A. Javadi-Abhari, N. Kanazawa, A. Karazeev, K. Krsulich, P. Liu, Y. Luh, Y. Maeng, M. Marques, Francisco Jose Martin-Fernández, Douglas T. McClure, D. M. Kay, S. Meesala, A. Mezzacapo, N. Moll, D. M. Rodríguez, G. Nannicini, P. Nation, P. Ollitrault, Lee James O'Riordan, H. Paik, J. Pérez, A. Phan, M. Pistoia, V. Prutyay, M. Reuter, J. Rice, A. R. Davila, Raymond Harry Putra Rudy, M. Ryu, N. Sathaye, C. Schnabel, E. Schoute, K. Setia, Y. Shi, A. Silva, Y. Siraichi, S. Sivarajah, J. A. Smolin, M. Soeken, H. Takahashi, I. Tavernelli, C. Taylor, P. Taylour, K. Trabing, M. Treinish, W. Turner, D. Vogt-Lee, C. Vuillot, J. A. Wildstrom, J. Wilson, E. Winston, C. Wood, S. Wood, S. Wörner, I. Y. Akhalwaya, C. Zoufal, Qiskit: An open-source framework for quantum computing, (2019); 10.5281/zenodo.2562110.
23. T. Alexander, N. Kanazawa, D. J. Egger, L. Capelluto, C. J. Wood, A. Javadi-Abhari, D. C. McKay, Qiskit pulse: Programming quantum computers through the cloud with pulses. *Quant. Sci. Technol.* **5**, 044006 (2020).
24. Y. Chen, M. Farahzad, S. Yoo, T.-C. Wei, Detector tomography on ibm quantum computers and mitigation of an imperfect measurement. *Phys. Rev. A* **100**, 052315 (2019).
25. J. Heinsoo, C. K. Andersen, A. Remm, S. Krinner, T. Walter, Y. Salathé, S. Gasparinetti, J.-C. Besse, A. Potočnik, A. Wallraff, C. Eichler, Rapid high-fidelity multiplexed readout of superconducting qubits. *Phys. Rev. Appl.* **10**, 034040 (2018).
26. S. McArdle, S. Endo, A. Aspuru-Guzik, S. C. Benjamin, X. Yuan, Quantum computational chemistry. *Rev. Mod. Phys.* **92**, 015003 (2020).
27. R. Cleve, A. Ekert, C. Macciavella, M. Mosca, Quantum algorithms revisited. *Proc. R. Soc. Lond. A* **454**, 339–354 (1998).
28. F. B. Maciejewski, Z. Zimborás, M. Oszmaniec, Mitigation of readout noise in near-term quantum devices by classical post-processing based on detector tomography. *Quantum* **4**, 257 (2020).
29. L. Funcke, T. Hartung, K. Jansen, S. Kühn, P. Stornati, X. Wang, Measurement error mitigation in quantum computers through classical bit-flip correction. *arXiv* (2020); arxiv.org/abs/2007.03663.
30. S. S. Tannu and Moinuddin K. Qureshi, Mitigating measurement errors in quantum computers by exploiting state-dependent bias, in *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO '52)* (Association for Computing Machinery, New York, NY, USA, 2019) p. 279–290.
31. H. Kwon, J. Bae, A hybrid quantum-classical approach to mitigating measurement errors in quantum algorithms. *IEEE Trans. Comput.* **70**, 1401–1411 (2020).
32. C. Kim, K. D. Park, J.-K. Rhee, Quantum error mitigation with artificial neural network. *IEEE Access* **8**, 188853–188860 (2020).
33. A. Zlokap, A. Gheorghiu, A deep learning model for noise prediction on near-term quantum devices, *arXiv* (2020); arxiv.org/abs/2005.10811.
34. P. Czarnik, A. Arrasmith, P. J. Coles, L. Cincio, Error mitigation with Clifford quantum-circuit data, *arXiv* (2020); arxiv.org/abs/2005.10189.
35. M. P. Harrigan, K. J. Sung, M. Neeley, K. J. Satzinger, F. Arute, K. Arya, J. Atalaya, J. C. Bardin, R. Barends, S. Boixo, M. Broughton, B. B. Buckley, D. A. Buell, B. Burkett, N. Bushnell, Y. Chen, Z. Chen, Ben Chiaro, R. Collins, W. Courtney, S. Demura, A. Dunsworth, D. Eppens, A. Fowler, B. Foxen, C. Gidney, M. Giustina, R. Graff, S. Habegger, A. Ho, S. Hong, T. Huang, L. B. Ioffe, S. V. Isakov, E. Jeffrey, Z. Jiang, C. Jones, D. Kafri, K. Kechedzhi, J. Kelly, S. Kim, P. V. Klimov, A. N. Korotkov, F. Kostritsa, D. Landhuis, P. Laptev, M. Lindmark, M. Leib, O. Martin, J. M. Martinis, J. R. McClean, M. McEwen, A. Megrant, X. Mi, M. Mohseni, W. Muczkiewicz, J. Mutus, O. Naaman, C. Neill, F. Neukart, M. Y. Niu, T. E. O'Brien, B. O'Gorman, E. Ostby, A. Petukhov, H. Putterman, C. Quintana, P. Roushan, N. C. Rubin, D. Sank, A. Skolik, V. Smelyanskiy, D. Strain, M. Streif, M. Szalay, A. Vainsencher, T. White, Z. J. Yao, P. Yeh, A. Zalcman, L. Zhou, H. Neven, D. Bacon, E. Lucero, E. Farhi, R. Babbush, Quantum approximate optimization of non-planar graph problems on a planar superconducting processor. *Nat. Phys.* **17**, 332–336 (2021).
36. P. Czarnik, A. Arrasmith, P. J. Coles, L. Cincio, Error mitigation with Clifford quantum-circuit data, (2021); arxiv.org/abs/2005.10189.
37. R. Hicks, C. W. Bauer, B. Nachman, Readout rebalancing for near-term quantum computers. *Phys. Rev. A* **103**, 022407 (2021).
38. E. van den Berg, Z. K. Mineev, K. Temme, Model-free readout-error mitigation for quantum expectation values, (2021); arxiv.org/abs/2012.09738.
39. K. Wang, Y.-A. Chen, Xin Wang, Measurement error mitigation via truncated neumann series (2021); arxiv.org/abs/2103.13856.
40. E. Peters, A. C. Y. Li, G. N. Perdue, Perturbative readout error mitigation for near term quantum computers (2021); arxiv.org/abs/2105.08161.
41. F. B. Maciejewski, F. Baccari, Z. Zimborás, M. Å. Oszmaniec, Modeling and mitigation of cross-talk effects in readout noise with applications to the quantum approximate optimization algorithm. *Quantum* **5**, 464 (2021).
42. I. Cong, S. Choi, M. D. Lukin, Quantum convolutional neural networks. *Nat. Phys.* **15**, 1273–1278 (2019).
43. M. R. Geller, M. Sun, Efficient correction of multiqubit measurement errors (2020); arxiv.org/abs/2001.09980.
44. K. E. Hamilton, T. Kharazi, T. Morris, A. J. McCaskey, R. S. Bennink, R. C. Pooser, Scalable quantum processor noise characterization, in *Proceedings of the 2020 IEEE International Conference on Quantum Computing and Engineering (QCE)* (IEEE, 2020) pp. 430–440.
45. B. Nachman, M. Urbanek, W. A. de Jong, C. W. Bauer, Unfolding quantum computer readout noise. *npj Quantum Inf.* **6**, 84 (2020).
46. M. R. Geller, Rigorous measurement error correction. *Quant. Sci. Technol.* **5**, 03LT01 (2020).
47. E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, D. J. Wineland, Randomized benchmarking of quantum gates. *Phys. Rev. A* **77**, 012307 (2008).
48. Robert Ash, *Information Theory* (Interscience Tracts in Pure and Applied Mathematics, Interscience, 1965) p. 121.
49. D. Berend, A. Kontorovich, On the convergence of the empirical distribution (2012); arxiv.org/abs/1205.6711.
50. IBM Quantum team, "ibmq_manhattan v1.6.2," (2020); quantum-computing.ibm.com.

Acknowledgments: We acknowledge the use of IBM Quantum services for this work. The views expressed are those of the authors and do not reflect the official policy or position of IBM or the IBM Quantum team. **Funding:** This work has been supported by the U.K. EPSRC (EP/P510257/1), the EPSRC Hub in Quantum Computing and Simulation (EP/T001062/1), the Royal Society, and the Samsung GRP grant. **Author contributions:** A.W.R.S. and M.S.K. planned the research. A.W.R.S., K.E.K., and C.N.S. discussed the initial idea that was then developed and analyzed by A.W.R.S. All authors discussed the results and contributed to the final manuscript. M.S.K. supervised the project. **Competing interests:** The authors are inventors on a patent application related to this work filed by Samsung Electronics Co. Ltd. (US 63/147388). The authors declare no additional competing interests. **Data and materials availability:** All data needed to evaluate the conclusions in the paper are present in the paper and/or the Supplementary Materials. The raw data used to generate the plots (Figs. 3 to 5) and the raw response matrix calibration data from *ibmq_manhattan* can be found at <https://doi.org/10.5281/zenodo.5267397>.

Submitted 1 April 2021

Accepted 28 September 2021

Published 17 November 2021

10.1126/sciadv.abi8009

Qubit readout error mitigation with bit-flip averaging

Alistair W. R. SmithKiran E. KhoslaChris N. SelfM. S. Kim

Sci. Adv., 7 (47), eabi8009. • DOI: 10.1126/sciadv.abi8009

View the article online

<https://www.science.org/doi/10.1126/sciadv.abi8009>

Permissions

<https://www.science.org/help/reprints-and-permissions>

Use of this article is subject to the [Terms of service](#)