

Averaged Circuit Eigenvalue Sampling

Steven T. Flammia  

AWS Center for Quantum Computing, Pasadena, CA, USA
California Institute of Technology, Pasadena, CA, USA

Abstract

We introduce ACES, a method for scalable noise metrology of quantum circuits that stands for Averaged Circuit Eigenvalue Sampling. It simultaneously estimates the individual error rates of all the gates in collections of quantum circuits, and can even account for space and time correlations between these gates. ACES strictly generalizes randomized benchmarking (RB), interleaved RB, simultaneous RB, and several other related techniques. However, ACES provides much more information and provably works under strictly weaker assumptions than these techniques. Finally, ACES is extremely scalable: we demonstrate with numerical simulations that it simultaneously and precisely estimates all the Pauli error rates on every gate and measurement in a 100 qubit quantum device using fewer than 20 relatively shallow Clifford circuits and an experimentally feasible number of samples. By learning the detailed gate errors for large quantum devices, ACES opens new possibilities for error mitigation, bespoke quantum error correcting codes and decoders, customized compilers, and more.

2012 ACM Subject Classification Hardware → Quantum computation

Keywords and phrases Quantum noise estimation, quantum benchmarking, QCVV

Digital Object Identifier 10.4230/LIPIcs.TQC.2022.4

Related Version *Averaged circuit eigenvalue sampling*: <https://arxiv.org/abs/2108.05803>

Supplementary Material The source code used to perform the simulations and generate the figures is available on GitHub:

Software (Source Code): <https://github.com/sflammia/ACES>

Funding This work was supported by the ARO QCISS program grant W911NF2110001.

Acknowledgements We thank Laura DeLorenzo, Robin Harper, Robert Huang, Alex Kubica, Ryan O'Donnell, Colm Ryan, Prasahnt Sivarajah, and Giacomo Torlai for discussions.

1 Introduction

Estimating errors in quantum computers is essential for progress towards fault tolerant quantum computation (FTQC) [36]. An error is any undesired quantum evolution, and so errors can be as general as the set of allowed quantum dynamics, making them difficult to estimate and characterize. The most relevant errors in the context of FTQC can be broadly cast into the two archetypes of coherent and incoherent errors [31], though this is not an exclusive dichotomy.

Coherent errors are roughly those that we wish to reduce through improved calibration or eliminate via dynamical decoupling [42], though clever choices of quantum codes and circuits can also be tailored to handle coherent noise [9, 22, 46]. These methods reach natural limits when the coherent noise becomes too complex to efficiently describe. While in principle coherent errors can accumulate badly during a computation [31], quantum error correction itself seems to reduce the coherence of noise [23, 3, 26].

Incoherent noise, by contrast, can generally only be completely fixed by quantum error correction and fault tolerance, though near-term strategies for error mitigation could also help [38, 32, 11, 8, 34]. Optimizing the codes, decoders, and circuits for FTQC requires a



© Steven T. Flammia;

licensed under Creative Commons License CC-BY 4.0

17th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2022).

Editors: François Le Gall and Tomoyuki Morimae; Article No. 4; pp. 4:1–4:10

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

comprehensive understanding of the incoherent noise in a quantum device. Many techniques have been developed to estimate incoherent errors, including randomized benchmarking (RB) [10], interleaved RB [35], simultaneous RB [17], character RB [21], and Pauli noise estimation [14] among others. Each of these techniques has in common that a general quantum noise source (which may include coherent errors) is actively *averaged* to obtain an incoherent noise model with the same fidelity using randomized control techniques [41, 28, 29, 43, 44]. It is this averaged noise that RB-type methods seek to estimate.

In this paper, we show that incoherent noise, modeled as a Pauli channel, can be learned extremely efficiently using averaged circuit eigenvalue sampling, or ACES. It is already known that Pauli channels can be (individually) estimated efficiently and in a manner that is robust to state preparation and measurement (SPAM) errors [19, 20, 14, 16], and they are effective at modeling noise for FTQC [18, 27]. ACES goes beyond this prior work and *simultaneously* estimates a large collection of Pauli noise channels associated to a quantum device. Indeed, we give numerical simulations showing that ACES can characterize every error rate associated to the Clifford gates in a 100 qubit quantum device using fewer than 20 circuits and a reasonable number of samples. Moreover, it requires only very simple classical resources to process these data, unlike other characterization techniques based on simulating or implementing general quantum circuits, or using challenging tensor network simulations [4, 5, 37, 39, 33, 6].

2 The Pauli and Clifford groups

The n -qubit *Pauli group* P_n consists of n -fold tensor products of single-qubit Pauli operators labeled as follows. Let \mathbf{a} be a $2n$ -bit string $\mathbf{a} = a_1 a_2 \dots a_{2n}$ and write $P_{\mathbf{a}} = i^{\mathbf{a}^T \Upsilon \mathbf{a}} \prod_{j=1}^n X_j^{a_{2j-1}} Z_j^{a_{2j}}$, where X_j and Z_j are single-qubit Paulis acting on qubit j , and $\Upsilon = \bigoplus_{k=1}^n \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ is such that $P_{\mathbf{a}}$ is always hermitian. The group P_n contains these $P_{\mathbf{a}}$, together with the overall phases $\{\pm 1, \pm i\}$, composed under matrix multiplication. All elements of the Pauli group satisfy

$$P_{\mathbf{a}} P_{\mathbf{b}} = (-1)^{\langle \mathbf{a}, \mathbf{b} \rangle} P_{\mathbf{b}} P_{\mathbf{a}}, \quad (1)$$

where the sign is determined by the binary symplectic form $\langle \mathbf{a}, \mathbf{b} \rangle = \mathbf{a}^T (\Upsilon + \Upsilon^T) \mathbf{b} \bmod 2$.

The normalizer of the Pauli group inside the unitary group, modulo phases, is the *Clifford group* C_n , and it is generated by the controlled-NOT gate $CX_{i \rightarrow j}$ from control i to target j , the Hadamard gate H_j , and the phase gate S_j ¹.

Pauli channels are quantum channels of the form

$$\rho \mapsto \sum_{\mathbf{a}} p_{\mathbf{a}} P_{\mathbf{a}} \rho P_{\mathbf{a}}^{\dagger}, \quad (2)$$

where $p_{\mathbf{a}}$ is a (possibly subnormalized) probability distribution called the *Pauli error rates*. Leakage from the qubit space occurs when $\sum_{\mathbf{a}} p_{\mathbf{a}} < 1$. When a general quantum channel $\mathcal{E} = \sum_j K_j \cdot K_j^{\dagger}$ is *twirled* by the Pauli group, it becomes a Pauli channel denoted \mathcal{E}^P ,

$$\mathcal{E}^P(\rho) = \frac{1}{4^n} \sum_{\mathbf{a}} P_{\mathbf{a}}^{\dagger} \mathcal{E}(P_{\mathbf{a}} \rho P_{\mathbf{a}}^{\dagger}) P_{\mathbf{a}}. \quad (3)$$

¹ From this definition, elements of the Clifford group are actually equivalence classes up to an overall phase, but by a slight abuse of language we can speak about a ‘‘Clifford unitary’’ to mean any representative element up to a phase and refer to uniqueness of a Clifford unitary when we really mean uniqueness up to an overall phase.

If $K_j = \sum_{\mathbf{a}} \nu_{j,\mathbf{a}} P_{\mathbf{a}}$, then the Pauli error rates of \mathcal{E}^P are $p_{\mathbf{a}} = \sum_j |\nu_{j,\mathbf{a}}|^2$. Thus we can speak of the Pauli error rates of a general channel by considering its Pauli twirl. Note that we can interpret twirling as the mean of a random process where a Pauli is selected uniformly at random and applied both before and after the channel.

The eigenvectors of a Pauli channel \mathcal{E} are just the Pauli operators. Indeed, from Equation (1) we have $\mathcal{E}(P_{\mathbf{b}}) = \lambda_{\mathbf{b}} P_{\mathbf{b}}$ where the *Pauli eigenvalues* $\lambda_{\mathbf{b}}$ are,

$$\lambda_{\mathbf{b}} = \sum_{\mathbf{a}} (-1)^{\langle \mathbf{a}, \mathbf{b} \rangle} p_{\mathbf{a}}. \quad (4)$$

This equation can be inverted to express the error rates in terms of the eigenvalues ²,

$$p_{\mathbf{a}} = \frac{1}{2^n} \sum_{\mathbf{b}} (-1)^{\langle \mathbf{a}, \mathbf{b} \rangle} \lambda_{\mathbf{b}}. \quad (5)$$

We now introduce a “ \mathcal{G} -twisted” Pauli twirl. For a given Clifford \mathcal{G} and Pauli $P_{\mathbf{a}}$, let $P_{\mathbf{a}'} = \mathcal{G}(P_{\mathbf{a}})$. Note that since \mathcal{G} is unitary, the set of all \mathbf{a} and \mathbf{a}' are in one-to-one correspondence. We wish to expand a noisy gate as $\tilde{\mathcal{G}} = \mathcal{G}\mathcal{E}$ for some general noise channel $\mathcal{E} = \mathcal{G}^\dagger \tilde{\mathcal{G}}$. Intuitively, \mathcal{E} is close to the identity, though the definition doesn’t assume that. Then the \mathcal{G} -twisted twirl of $\tilde{\mathcal{G}}$ is

$$\tilde{\mathcal{G}}^{\mathcal{G}P}(\rho) = \frac{1}{4^n} \sum_{\mathbf{a}} P_{\mathbf{a}}^\dagger \tilde{\mathcal{G}}(P_{\mathbf{a}} \rho P_{\mathbf{a}}^\dagger) P_{\mathbf{a}'} = \mathcal{G}(\mathcal{E}^P(\rho)). \quad (6)$$

From the last equality, we see that the \mathcal{G} -twisted twirl isolates the Pauli noise around a given noisy implementation $\tilde{\mathcal{G}}$ of an ideal Clifford gate \mathcal{G} .

\mathcal{G} -twisted twirled channels have an analogous eigendecomposition to a Pauli twirled channel, but with the notion of *generalized eigenvector*. Given such a channel $\tilde{\mathcal{G}}^{\mathcal{G}P}$, the generalized eigenvectors with respect to \mathcal{G}_0 are vectors such that $\tilde{\mathcal{G}}^{\mathcal{G}P}(v) = \lambda \mathcal{G}_0(v)$. We see from Equation (6) that choosing $\mathcal{G}_0 = \mathcal{G}$ gives exactly the Paulis as the generalized eigenvectors with eigenvalues given by the Pauli eigenvalues of the noise map \mathcal{E}^P .

3 Averaged circuits

Let us consider a Clifford circuit (i.e., a circuit composed solely of CX , H , and S gates or an equivalent generating set), denoted \mathcal{C} . Any physical implementation of these circuits will be noisy, and we seek to characterize the incoherent Pauli-averaged noise in these circuits, specifically in the generators used to create the circuits. To that end, from the circuit \mathcal{C} we create a new ensemble of circuits \mathcal{C}^P by sampling a \mathcal{G} -twisted Pauli twirl across each Clifford circuit element and recompiling the Pauli gate. This approach to Pauli frame randomization is known as randomized compiling [43]. Each circuit in the ensemble implements the same unitary operation, but now the noise has been averaged over the Pauli group. In Ref. [43], it was proven that circuits sampled in this way yield on average a circuit that interleaves Pauli-averaged noise with ideal gates (except possibly in the final measurement step). This result rigorously holds whenever the noise on each Pauli gate is the same, and furthermore Ref. [43] provides some robustness guarantees in the event that this assumption is perturbatively violated ³.

² Note that $\lambda_{\mathbf{a}}$ and $p_{\mathbf{a}}$ are essentially Fourier transform pairs since the transform relating them is the Walsh-Hadamard transform (up to a permutation). A helpful intuition is that λ lives in the “time domain” where we can efficiently sample, and p lives in the “frequency domain” where we wish to reconstruct the signal.

³ While the gate-independent noise assumption may seem unrealistic, it should be noted that the successful and widely used method of standard RB makes the *much* stronger assumption that the noise is gate-independent across all n -qubit Clifford gates, whereas ACES weakens this substantially to just the Pauli gates.

4:4 Averaged Circuit Eigenvalue Sampling

These considerations motivate considering only *averaged circuits*, denoted $\mathcal{C}^{\mathbb{P}}$, so that the noisy physical implementations will have the form

$$\widetilde{\mathcal{C}}^{\mathbb{P}} = \widetilde{\mathcal{G}}_T^{\mathcal{G}_T^{\mathbb{P}}} \dots \widetilde{\mathcal{G}}_1^{\mathcal{G}_1^{\mathbb{P}}} = \mathcal{G}_T \mathcal{E}_{\mathcal{G}_T} \dots \mathcal{G}_1 \mathcal{E}_{\mathcal{G}_1}. \quad (7)$$

4 Eigenvalue sampling

Let us suppose for the moment that a given circuit \mathcal{C} ideally implements the identity unitary. Under the gate-independent noise assumption, it follows that the noisy implementation of the averaged circuit, $\widetilde{\mathcal{C}}^{\mathbb{P}}$, will be a Pauli channel. It therefore has Pauli eigenvalues, namely $\widetilde{\mathcal{C}}^{\mathbb{P}}(P_{\mathbf{a}}) = \Lambda_{\mathcal{C},\mathbf{a}} P_{\mathbf{a}}$, where we use capital Λ to denote this circuit-level eigenvalue. Because of the gate-independent noise assumption, this eigenvalue depends only on the eigenvector ($P_{\mathbf{a}}$) and on the circuit (\mathcal{C}), so it is labeled accordingly as $\Lambda_{\mathcal{C},\mathbf{a}}$.

If the circuit \mathcal{C} does not implement the identity unitary, but rather some net Clifford operation, something similar still holds. If the ideal circuit maps an input Pauli $P_{\mathbf{a}}$ to an output Pauli $\mathcal{C}(P_{\mathbf{a}}) = \pm P_{\mathbf{a}'}$, then the overall \pm sign and the value of \mathbf{a}' can be efficiently computed [1]. The noisy version of the circuit will give an averaged operator that satisfies the generalized eigenvalue equation

$$\widetilde{\mathcal{C}}^{\mathbb{P}}(P_{\mathbf{a}}) = \Lambda_{\mathcal{C},\mathbf{a}} \mathcal{C}(P_{\mathbf{a}}) = \pm \Lambda_{\mathcal{C},\mathbf{a}} P_{\mathbf{a}'}. \quad (8)$$

From the orthogonality of the Pauli basis, it follows that

$$\Lambda_{\mathcal{C},\mathbf{a}} = \pm \frac{1}{2^n} \text{Tr}(P_{\mathbf{a}'} \widetilde{\mathcal{C}}^{\mathbb{P}}(P_{\mathbf{a}})), \quad (9)$$

and this suggests a prescription for estimating the (generalized) eigenvalue $\Lambda_{\mathcal{C},\mathbf{a}}$ that we call eigenvalue sampling.

To estimate $\Lambda_{\mathcal{C},\mathbf{a}}$ via eigenvalue sampling, let us focus on the case where $P_{\mathbf{a}}$ is a single-qubit Pauli. We begin by selecting uniformly at random an eigenstate ψ_{\pm} on the support of $P_{\mathbf{a}}$ having eigenvalue ± 1 (ignoring the other registers). Then we send ψ_{\pm} into a randomly chosen element in the circuit ensemble $\mathcal{C}^{\mathbb{P}}$ and measure the output in the basis defined by $P_{\mathbf{a}'}$. Our overall estimate for $\Lambda_{\mathcal{C},\mathbf{a}}$ consists of measuring N independent experiments and taking the difference of the sample averages between the ψ_+ and ψ_- experiments. It is easy to check that this differencing trick makes Equation (9) hold in expectation, so this is an unbiased estimator of $\Lambda_{\mathcal{C},\mathbf{a}}$. This sampling strategy was first analyzed in Ref. [15], and it is straightforward to generalize to the n -qubit case. Note that it will be most efficient if the support of $P_{\mathbf{a}}$ and $P_{\mathbf{a}'}$ are relatively small, and also that Paulis with disjoint support can implement such measurements simultaneously.

5 Relating circuit and gate eigenvalues

We have seen how eigenvalue sampling on averaged circuits gives us access to the (generalized) Pauli eigenvalues $\Lambda_{\mathcal{C},\mathbf{a}}$ in the implemented circuit ensemble $\widetilde{\mathcal{C}}^{\mathbb{P}}$. This is already a useful method for estimating the quality of the circuit implementation $\widetilde{\mathcal{C}}$, since it can be interpreted as a fidelity-like measure for how faithfully the circuit executes given the input $P_{\mathbf{a}}$. However, we seek to characterize not just the global circuit noise, but the error rates associated to the constituent gates as well. How do the (generalized) eigenvalues of the individual gates relate to the eigenvalue of the total circuit $\mathcal{C} = \mathcal{G}_T \dots \mathcal{G}_1$?

Let us apply the generalized eigenvalue relation sequentially to the gates in a Clifford circuit. For the first gate we obtain $\tilde{\mathcal{G}}_1^{\mathcal{P}}(P_{\mathbf{a}_1}) = \lambda_{1,\mathbf{a}_1} \mathcal{G}_1(P_{\mathbf{a}_1}) = (\pm)_1 \lambda_{1,\mathbf{a}_1} P_{\mathbf{a}_2}$. Acting on this with $\tilde{\mathcal{G}}_2^{\mathcal{P}}$, we obtain

$$\begin{aligned} \tilde{\mathcal{G}}_2^{\mathcal{P}} \tilde{\mathcal{G}}_1^{\mathcal{P}}(P_{\mathbf{a}_1}) &= (\pm)_1 (\pm)_2 \lambda_{1,\mathbf{a}_1} \lambda_{2,\mathbf{a}_2} \mathcal{G}_2^{\mathcal{P}}(P_{\mathbf{a}_2}) \\ &= (\pm)_1 (\pm)_2 \lambda_{1,\mathbf{a}_1} \lambda_{2,\mathbf{a}_2} P_{\mathbf{a}_3}. \end{aligned}$$

Continuing in this fashion, we find that

$$\tilde{\mathcal{C}}^{\mathcal{P}}(P_{\mathbf{a}_1}) = (\pm) \prod_{k=1}^T \lambda_{k,\mathbf{a}_k} \mathcal{C}(P_{\mathbf{a}_1}), \quad (10)$$

where the overall sign and the individual \mathbf{a}_k can be computed efficiently [1]. Comparing with Equation (8), we see that $\Lambda_{\mathcal{C},\mathbf{a}_1} = (\pm) \prod_k \lambda_{k,\mathbf{a}_k}$. We will use the freedom to reinterpret the sign of the input Pauli $P_{\mathbf{a}_1}$ to ensure that we always have a + sign in this equation, and therefore we have the relation

$$\Lambda_{\mathcal{C},\mathbf{a}_1} = \prod_{k=1}^T \lambda_{k,\mathbf{a}_k}. \quad (11)$$

With this sign convention, in the regime of interest $\Lambda_{\mathcal{C},\mathbf{a}_1}$ is positive and not too small. We therefore focus on sets of circuits \mathcal{C}_k and input labels \mathbf{a}_{k_i} such that $\Lambda_{\mathcal{C}_k,\mathbf{a}_{k_i}}$ is always larger than, say, $1/2$, and gates where $\lambda_{k,\mathbf{a}_k} > 0$.

6 Estimating gate errors via ACES

We now consider a circuit \mathcal{C}_k and an input label \mathbf{a}_{k_i} ; we give this combination a composite index $\mu = (\mathcal{C}_k, \mathbf{a}_{k_i})$. From the above discussion, we can obtain an empirical estimate $\hat{\Lambda}_\mu$ of Λ_μ by eigenvalue sampling on the averaged circuit ensemble for the circuit/input label μ . Similarly, we assemble all gate-level eigenvalues under a single index to get λ_ν , where ν labels pairs of gates and Paulis whose noise we wish to model. Since all eigenvalue quantities are positive in the regime of interest, we can introduce new variables,

$$\Lambda_\mu = e^{-b_\mu}, \quad \lambda_\nu = e^{-x_\nu}. \quad (12)$$

The new variables are related by the *linear* equations

$$\sum_{\nu} A_{\mu\nu} x_\nu = b_\mu. \quad (13)$$

We refer to the matrix A as the *design matrix*. Once enough independent equations (μ) are obtained so that A is full rank, an estimate for \mathbf{x} can be obtained in any number of ways, most straightforwardly via least squares as $\hat{\mathbf{x}} = A^+ \hat{\mathbf{b}}$, where $\hat{\mathbf{b}}$ denotes an empirical estimate for \mathbf{b} and A^+ is the pseudoinverse of A . Inverting Equation (12) subsequently gives us estimates for λ_ν , and Pauli error rates can be obtained by using Equation (5).

The precision of our estimate depends in part on the choice of A , as well as the precision of the initial estimates of the Λ_μ . The estimates for λ_ν are always accurate in the sense that these are consistent estimators, however they will in general have some bias. In the numerical simulations below, no attempt was made to find optimal designs A , and only random choices were used. We leave open the question of finding optimal design matrices that maximize the precision and accuracy of these estimators.

4:6 Averaged Circuit Eigenvalue Sampling

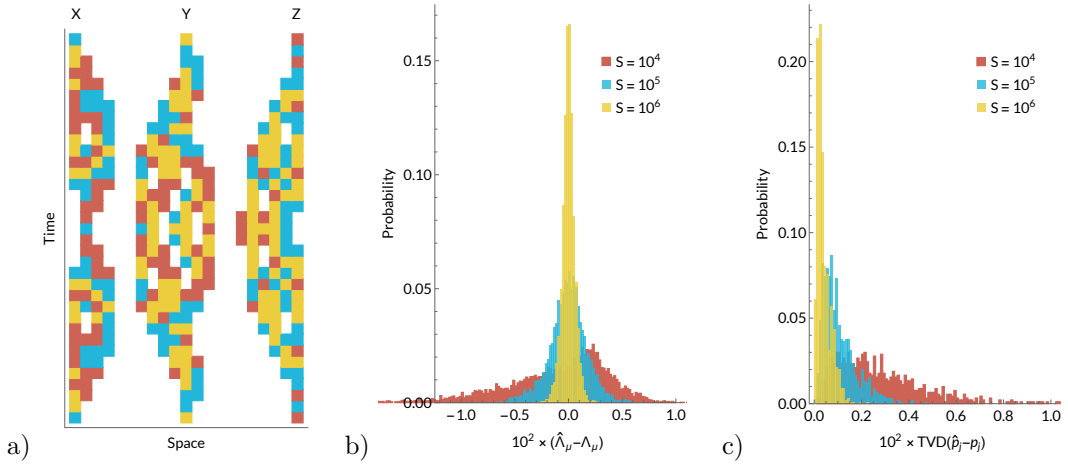


Figure 1 a) Sending X , Y , and Z Paulis (blue, yellow, and red, resp.) through a small “mirror circuit” (i.e., one of the form UU^\dagger) with $n = 21$ qubits, depth $d = 34$, and nearest-neighbor gates in 1D. Normalized histograms of b) the absolute error for the μ th estimated circuit eigenvalue $\hat{\Lambda}_\mu$, and c) the total variation distance (TVD) for the estimated Pauli error rates $\hat{\rho}_j$ of the noisy gate $\hat{\mathcal{G}}_j$ in a $n = 100$ qubit simulation. There are 898 gates (including measurements) in the model, and 10,155 estimated circuit fidelities (which are estimated in large batches due to the n -bit measurements) to estimate $N = 5070$ parameters. Plots are for a number of samples S per Λ_μ of 10^4 , 10^5 , and 10^6 .

7 Correlations and SPAM

The ACES methodology is flexible enough to allow independent estimation of SPAM errors as well as space- and/or certain time-correlated errors. To estimate measurement noise, we simply add a list of variables x_ν associated to each Pauli measurement error that we wish to model. We caution that separating preparation errors from measurement errors will not be possible if they are introduced into the model in a symmetric way (because then A will not be full rank); this problem is not unique to ACES however [4] and we do not attempt to resolve it here.

To handle space-correlated errors, we reinterpret the gates that generate our circuits to come in correlated groups. For example, if we want to model correlated noise between the Hadamard gates H_1 and H_2 , we could have separate variables for the gates H_1 , H_2 , and $H_1 \otimes H_2$. This is analogous to interleaved [35] and simultaneous RB [17], except that all of the data are used to fit all of the gates and correlations symmetrically and simultaneously.

Limited forms of time-correlated errors can be handled similarly by introducing variables for pairs of gates in time. For example, if the noise on H_1 depends on whether S_1 was applied or not right before, then we can introduce separate variables for these cases.

The only condition for a unique and consistent estimate in all of these scenarios is that the design matrix A is full rank. If A were random, then we only need as many equations as unknowns for this to hold with high probability. From this heuristic, we expect that the number of experiments should be about as large as, or a little larger than, the number of variables.

8 Numerical results

We now demonstrate the scalability of ACES via numerical simulations. Rigorous proofs of the consistency of ACES and bounds on the sample complexity will be presented elsewhere.

We consider the most general model of inhomogeneous but uncorrelated noise, plus readout errors⁴. In this model on n qubits there are $O(n)$ variables: CX gates acting between neighbors, together with six single-qubit Clifford gates (modulo the Paulis), and independent readout errors on each qubit in each Pauli basis.

We generated $C = 19$ random 1D Clifford circuits on $n = 100$ qubits of varying depths from $d = 2$ up to $d = 89$. The sum of all the circuit depths, including the measurement rounds, was 354. We then computed the circuit eigenvalues obtained from sending in all single-qubit Paulis and, on some circuits, two-qubit Paulis on nearest neighbors as well. We found it challenging to generate a full-rank design matrix A using the “mirror circuits” shown in Figure 1a, so we padded each mirror circuit with a depth 5 random circuit layer at the end. This means that the Paulis measured at the output had, in some cases, weight as high as 6, though most still had weight 1 or 2. Constant-weight Pauli operators can nonetheless be estimated efficiently from single-qubit Pauli measurements [12, 24, 25], and this only increases the sample complexity by a constant factor. We then generated a “true” noise model by assigning to each gate random Pauli error rates consistent with the estimates reported in the Arute et al. experiment [2]. The entire implementation can be found in the associated Mathematica notebook accompanying this manuscript [13].

Despite its seeming simplicity, this model still has $N = 5070$ parameters. Even under the simplifying assumptions of RB with Clifford averaging where the noise is depolarizing on each gate, there would still be 798 parameters (neglecting SPAM) to be estimated through interleaved RB, and even then the required Clifford randomizations would be prohibitively expensive.

ACES estimates all of these parameters with just these 19 random circuits (and their Pauli randomizations). This is possible because each measurement is an n -bit measurement, so many parameters are estimated in parallel.

In Figs. 1b and 1c we plot the convergence of the ACES estimate as a function of S , the number of samples per circuit eigenvalue estimate. Estimates \hat{x}_ν of the model parameters x_ν were obtained from the simulated data by solving Equation (13) with the simplest possible estimator, a truncated least-squares estimate (i.e., finding the least squares solution and truncating any negative values).

Counting an n -bit measurement as one sample, the total sample complexity is $O(SC)$ where C is the number of different averaged circuits used, in this case $C = 19$. Results are shown for $S = 10^4, 10^5, 10^6$. Even for $S = 10^4$, nearly all circuit eigenvalue estimates (1b) are within 1% of the true answer, and the total variation distance (TVD) between the estimated and true Pauli error rates on each gate are within .64% with 95% confidence. This latter figure improves to .1% with 95% confidence for $S = 10^6$, a remarkably precise estimate given how many parameters there are and that no regularization was used to avoid potential overfitting.

9 Discussion

There are many potential applications for ACES, and many avenues for improving and extending it as well. For example, in addition to the tailored codes and decoders mentioned above, error mitigation is one of the most natural applications of ACES [38, 32, 11, 8, 34], and it can also be used to debias estimates of classical shadows following the ideas in

⁴ We do not model state preparation errors, only readout errors, for the reason discussed in the main text that these are not separately identifiable.

Refs. [7, 30, 40]. Regarding extensions, while we have focused entirely on Clifford gates, it is easy to see that ACES can accommodate circuits with a constant number of T gates in specific configurations. However, extending beyond this to universal gate sets in general is an important question for future research. A differential analysis suggests that obtaining circuit eigenvalue estimates such that $\hat{\Lambda}_\mu = \Lambda_\mu \pm \epsilon \Lambda_\mu$ suffices to obtain gate-level eigenvalue estimates of order $\hat{\lambda}_\nu = \lambda_\nu \pm O(\|A^+\| \epsilon) \lambda_\nu$. Thus, finding circuits, Pauli inputs, and noise models whose associated design matrix minimizes $\|A^+\|$ could help optimize the sample efficiency of ACES. There are additional desiderata for the design matrix, such as requiring only few experiments and using circuits that map few-qubit Paulis to few-qubit Paulis. Finding a general understanding of which circuits behave best is an open question. While ACES can test for correlations in a given noise model, it would be more powerful to include a large model and then search for dominant correlations by enforcing sparsity. One way forward might be to test clusters of gates for inclusion using methods such as group LASSO [45]. Finally, the most obvious open problem is to implement ACES in a near-term experiment.

References

- 1 S. Aaronson and D. Gottesman. Improved simulation of stabilizer circuits. *Phys. Rev. A*, 70(5):052328, November 2004. doi:10.1103/PhysRevA.70.052328.
- 2 Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando G. S. L. Brandao, David A. Buell, Brian Burkett, Yu Chen, Zijun Chen, Ben Chiaro, Roberto Collins, William Courtney, Andrew Dunsworth, Edward Farhi, Brooks Foxen, Austin Fowler, Craig Gidney, Marissa Giustina, Rob Graff, Keith Guerin, Steve Habegger, Matthew P. Harrigan, Michael J. Hartmann, Alan Ho, Markus Hoffmann, Trent Huang, Travis S. Humble, Sergei V. Isakov, Evan Jeffrey, Zhang Jiang, Dvir Kafri, Kostyantyn Kechedzhi, Julian Kelly, Paul V. Klimov, Sergey Knysh, Alexander Korotkov, Fedor Kostritsa, David Landhuis, Mike Lindmark, Erik Lucero, Dmitry Lyakh, Salvatore Mandrà, Jarrod R. McClean, Matthew McEwen, Anthony Megrant, Xiao Mi, Kristel Michielsen, Masoud Mohseni, Josh Mutus, Ofer Naaman, Matthew Neeley, Charles Neill, Murphy Yuezhen Niu, Eric Ostby, Andre Petukhov, John C. Platt, Chris Quintana, Eleanor G. Rieffel, Pedram Roushan, Nicholas C. Rubin, Daniel Sank, Kevin J. Satzinger, Vadim Smelyanskiy, Kevin J. Sung, Matthew D. Trevithick, Amit Vainsencher, Benjamin Villalonga, Theodore White, Z. Jamie Yao, Ping Yeh, Adam Zalcman, Hartmut Neven, and John M. Martinis. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, October 2019. doi:10.1038/s41586-019-1666-5.
- 3 Stefanie J. Beale, Joel J. Wallman, Mauricio Gutiérrez, Kenneth R. Brown, and Raymond Laflamme. Coherence in quantum error-correcting codes. *Phys. Rev. Lett.*, 121:190501, 2018. doi:10.1103/PhysRevLett.121.190501.
- 4 Robin Blume-Kohout, John King Gamble, Erik Nielsen, Kenneth Rudinger, Jonathan Mizrahi, Kevin Fortier, and Peter Maunz. Demonstration of qubit operations below a rigorous fault tolerance threshold with gate set tomography. *Nature Communications*, 8:, 2016. doi:10.1038/ncomms14485.
- 5 Sergio Boixo, Sergei V. Isakov, Vadim N. Smelyanskiy, Ryan Babbush, Nan Ding, Zhang Jiang, Michael J. Bremner, John M. Martinis, and Hartmut Neven. Characterizing quantum supremacy in near-term devices. *Nature Physics*, 14(6):595–600, April 2018. doi:10.1038/s41567-018-0124-x.
- 6 Haoyuan Cai, Qi Ye, and Dong-Ling Deng. Sample complexity of learning quantum circuits, 2021. arXiv:2107.09078.
- 7 Senrui Chen, Wenjun Yu, Pei Zeng, and Steven T. Flammia. Robust shadow estimation, 2020. arXiv:2011.09636.

- 8 Piotr Czarnik, Andrew Arrasmith, Patrick J. Coles, and Lukasz Cincio. Error mitigation with Clifford quantum-circuit data, 2021. [arXiv:2005.10189](https://arxiv.org/abs/2005.10189).
- 9 Dripto M. Debroy, Muyuan Li, Michael Newman, and Kenneth R. Brown. Stabilizer slicing: Coherent error cancellations in low-density parity-check stabilizer codes. *Phys. Rev. Lett.*, 121(25):250502, December 2018. doi:10.1103/physrevlett.121.250502.
- 10 Joseph Emerson, Robert Alicki, and Karol Życzkowski. Scalable noise estimation with random unitary operators. *J. Opt. B*, 7(10):S347, 2005. doi:10.1088/1464-4266/7/10/021.
- 11 Suguru Endo, Simon C. Benjamin, and Ying Li. Practical quantum error mitigation for near-future applications. *Phys. Rev. X*, 8(3):031027, July 2018. doi:10.1103/physrevx.8.031027.
- 12 Tim J. Evans, Robin Harper, and Steven T. Flammia. Scalable Bayesian Hamiltonian learning, 2019. [arXiv:1912.07636](https://arxiv.org/abs/1912.07636).
- 13 S. T. Flammia. ACES. <https://github.com/sflammia/ACES>, 2021.
- 14 Steven Flammia and Joel Wallman. Efficient estimation of Pauli channels. *ACM Transactions on Quantum Computing*, 1(1):1–32, 2020. doi:10.1145/3408039.
- 15 Steven T. Flammia and Yi-Kai Liu. Direct fidelity estimation from few Pauli measurements. *Phys. Rev. Lett.*, 106(23):230501, June 2011. doi:10.1103/PhysRevLett.106.230501.
- 16 Steven T. Flammia and Ryan O’Donnell. Pauli error estimation via population recovery. *Quantum*, 5:549, September 2021. doi:10.22331/q-2021-09-23-549.
- 17 Jay M. Gambetta, A. D. Córcoles, S. T. Merkel, B. R. Johnson, John A. Smolin, Jerry M. Chow, Colm A. Ryan, Chad Rigetti, S. Poletto, Thomas A. Ohki, Mark B. Ketchen, and M. Steffen. Characterization of addressability by simultaneous randomized benchmarking. *Phys. Rev. Lett.*, 109:240504, December 2012. doi:10.1103/PhysRevLett.109.240504.
- 18 Michael R. Geller and Zhongyuan Zhou. Efficient error models for fault-tolerant architectures and the Pauli twirling approximation. *Phys. Rev. A*, 88(1):012314, July 2013. doi:10.1103/physreva.88.012314.
- 19 Robin Harper, Steven T. Flammia, and Joel J. Wallman. Efficient learning of quantum noise. *Nature Physics*, 16(12):1184–1188, August 2020. doi:10.1038/s41567-020-0992-8.
- 20 Robin Harper, Wenjun Yu, and Steven T. Flammia. Fast estimation of sparse quantum noise. *PRX Quantum*, 2(1):010322, February 2021. doi:10.1103/prxquantum.2.010322.
- 21 Jonas Helsen, Xiao Xue, Lieven MK Vandersypen, and Stephanie Wehner. A new class of efficient randomized benchmarking protocols. *npj Quantum Information*, 5(1):71, August 2019. doi:10.1038/s41534-019-0182-7.
- 22 Jingzhen Hu, Qingzhong Liang, Narayanan Rengaswamy, and Robert Calderbank. Mitigating coherent noise by balancing weight-2 Z -stabilizers, 2021. [arXiv:2011.00197](https://arxiv.org/abs/2011.00197).
- 23 Eric Huang, Andrew C. Doherty, and Steven Flammia. Performance of quantum error correction with coherent errors. *Phys. Rev. A*, 99:022313, 2019. doi:10.1103/PhysRevA.99.022313.
- 24 Hsin-Yuan Huang, Richard Kueng, and John Preskill. Predicting many properties of a quantum system from very few measurements. *Nature Physics*, 16(10):1050–1057, June 2020. doi:10.1038/s41567-020-0932-7.
- 25 Hsin-Yuan Huang, Richard Kueng, and John Preskill. Efficient estimation of Pauli observables by derandomization, 2021. [arXiv:2103.07510](https://arxiv.org/abs/2103.07510).
- 26 Joseph K Iverson and John Preskill. Coherence in logical quantum channels. *New Journal of Physics*, 22(7):073066, August 2020. doi:10.1088/1367-2630/ab8e5c.
- 27 Amara Katarbarwa and Michael R. Geller. Logical error rate in the Pauli twirling approximation. *Scientific Reports*, 5(1), September 2015. doi:10.1038/srep14670.
- 28 O. Kern, G. Alber, and D. L. Shepelyansky. Quantum error correction of coherent errors by randomization. *Euro. Phys. J. D*, 32(1):153–156, January 2005. doi:10.1140/epjd/e2004-00196-9.
- 29 E. Knill. Quantum computing with realistically noisy devices. *Nature*, 434(7029):39–44, March 2005. doi:10.1038/nature03350.
- 30 Dax Enshan Koh and Sabee Grewal. Classical shadows with noise, 2020. [arXiv:2011.11580](https://arxiv.org/abs/2011.11580).

- 31 Richard Kueng, David M. Long, Andrew C. Doherty, and Steven T. Flammia. Comparing experiments to the fault-tolerance threshold. *Phys. Rev. Lett.*, 117:170502, October 2016. doi:10.1103/PhysRevLett.117.170502.
- 32 Ying Li and Simon C. Benjamin. Efficient variational quantum simulator incorporating active error minimization. *Phys. Rev. X*, 7(2):021050, June 2017. doi:10.1103/physrevx.7.021050.
- 33 Yunchao Liu, Matthew Otten, Roozbeh Bassirianjahromi, Liang Jiang, and Bill Fefferman. Benchmarking near-term quantum computers via random circuit sampling, 2021. arXiv:2105.05232.
- 34 Angus Lowe, Max Hunter Gordon, Piotr Czarnik, Andrew Arrasmith, Patrick J. Coles, and Lukasz Cincio. Unified approach to data-driven quantum error mitigation, 2020. arXiv:2011.01157.
- 35 Easwar Magesan, Jay M. Gambetta, B. R. Johnson, Colm A. Ryan, Jerry M. Chow, Seth T. Merkel, Marcus P. da Silva, George A. Keefe, Mary B. Rothwell, Thomas A. Ohki, Mark B. Ketchen, and M. Steffen. Efficient measurement of quantum gate error by interleaved randomized benchmarking. *Phys. Rev. Lett.*, 109:080505, August 2012. doi:10.1103/PhysRevLett.109.080505.
- 36 John M. Martinis. Qubit metrology for building a fault-tolerant quantum computer. *npj Quantum Information*, 1:, October 2015. doi:10.1038/npjqi.2015.5.
- 37 Runzhou Tao, Yunong Shi, Jianan Yao, John Hui, Frederic T. Chong, and Ronghui Gu. Gleipnir: Toward practical error analysis for quantum programs. In *Proceedings of the 42nd ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI 2021)*, 2021. doi:10.1145/3453483.3454029.
- 38 Kristan Temme, Sergey Bravyi, and Jay M. Gambetta. Error mitigation for short-depth quantum circuits. *Phys. Rev. Lett.*, 119(18):180509, November 2017. doi:10.1103/physrevlett.119.180509.
- 39 Giacomo Torlai, Christopher J. Wood, Atithi Acharya, Giuseppe Carleo, Juan Carrasquilla, and Leandro Aolita. Quantum process tomography with unsupervised learning and tensor networks, 2020. arXiv:2006.02424.
- 40 Ewout van den Berg, Zlatko K. Mineev, and Kristan Temme. Model-free readout-error mitigation for quantum expectation values, 2021. arXiv:2012.09738.
- 41 Lorenza Viola and Emanuel Knill. Random decoupling schemes for quantum dynamical control and error suppression. *Phys. Rev. Lett.*, 94:060502, February 2005. doi:10.1103/PhysRevLett.94.060502.
- 42 Lorenza Viola, Emanuel Knill, and Seth Lloyd. Dynamical decoupling of open quantum systems. *Phys. Rev. Lett.*, 82:2417–2421, March 1999. doi:10.1103/PhysRevLett.82.2417.
- 43 Joel J. Wallman and Joseph Emerson. Noise tailoring for scalable quantum computation via randomized compiling. *Phys. Rev. A*, 94:052325, November 2016. doi:10.1103/PhysRevA.94.052325.
- 44 Matthew Ware, Guilhem Ribeill, Diego Ristè, Colm A. Ryan, Blake Johnson, and Marcus P. da Silva. Experimental Pauli-frame randomization on a superconducting qubit. *Phys. Rev. A*, 103(4):042604, April 2021. doi:10.1103/physreva.103.042604.
- 45 Ming Yuan and Yi Lin. Model selection and estimation in regression with grouped variables. *J. R. Statist. Soc. B*, 68(1):49–67, 2006. doi:10.1111/j.1467-9868.2005.00532.x.
- 46 Bichen Zhang, Swarnadeep Majumder, Pak Hong Leung, Stephen Crain, Ye Wang, Chao Fang, Dripto M. Debroy, Jungsang Kim, and Kenneth R. Brown. Hidden inverses: Coherent error cancellation at the circuit level, 2021. arXiv:2104.01119.