

## Probabilistic Interpolation of Quantum Rotation Angles

Bálint Koczor<sup>1,2,3,\*</sup> John J. L. Morton<sup>1,4</sup> and Simon C. Benjamin<sup>1,3,†</sup>

<sup>1</sup>Quantum Motion, 9 Sterling Way, London N7 9HJ, United Kingdom

<sup>2</sup>Mathematical Institute, University of Oxford, Woodstock Road, Oxford OX2 6GG, United Kingdom

<sup>3</sup>Department of Materials, University of Oxford, Parks Road, Oxford OX1 3PH, United Kingdom

<sup>4</sup>London Centre for Nanotechnology, UCL, 17-19 Gordon St, London WC1H 0AH, United Kingdom



(Received 6 July 2023; accepted 15 February 2024; published 27 March 2024)

Quantum computing requires a universal set of gate operations; regarding gates as rotations, any rotation angle must be possible. However a real device may only be capable of  $B$  bits of resolution, i.e., it might support only  $2^B$  possible variants of a given physical gate. Naive discretization of an algorithm's gates to the nearest available options causes coherent errors, while decomposing an impermissible gate into several allowed operations increases circuit depth. Conversely, demanding higher  $B$  can greatly complexify hardware. Here, we explore an alternative: probabilistic angle interpolation (PAI). This effectively implements any desired, continuously parametrized rotation by randomly choosing one of three discretized gate settings and postprocessing individual circuit outputs. The approach is particularly relevant for near-term applications where one would in any case average over many runs of circuit executions to estimate expected values. While PAI increases that sampling cost, we prove that (a) the approach is optimal in the sense that PAI achieves the least possible overhead and (b) the overhead is remarkably modest even with thousands of parametrized gates and only seven bits of resolution available. This is a profound relaxation of engineering requirements for first generation quantum computers where even 5–6 bits of resolution may suffice and, as we demonstrate, the approach is many orders of magnitude more efficient than prior techniques. Moreover we conclude that, even for more mature late noisy intermediate-scale quantum era hardware, no more than nine bits will be necessary.

DOI: [10.1103/PhysRevLett.132.130602](https://doi.org/10.1103/PhysRevLett.132.130602)

*Introduction.*—Producing quantum computers of the scale and fidelity needed to solve practically useful problems requires development not just of the quantum processor itself, but of the analog and digital electronics used for the control and readout of qubits. These electronics may include field-programmable gate array systems [1,2] and customised integrated circuits [3–6], typically cooled to improve performance and integration with the qubits. Understanding and minimizing the required specifications of the electronics supporting the quantum computer is an essential step in developing scalable systems in general. These issues, however, become even more acute when considering cryogenic electronics with limited power budgets [7], and/or quantum processor architectures targeting close integration of quantum systems with classical systems, such as for silicon-based spin qubits [8].

Consider, for example, the instruction to implement a parametrized Pauli gate in which a user has specified the

kind of gate they want to implement and to which qubits. As the gate angles are defined and implemented using digital electronics, they must be discretized into  $B$  bits of resolution. The choice of  $B$  has significant impacts on elements such as the bandwidth of communication channels between different elements of the control stack and the memory requirements of any gate instruction cache, and in the digital-to-analog converters used ultimately to produce the driving fields acting on the qubits. There is therefore a strong benefit in minimizing  $B$  to the point where it is just sufficient to provide the required gate fidelities for a given application or circuit. Most of the currently leading qubit hardware platforms operate optimally at cryogenic temperatures, including superconducting qubits [5,9], trapped ions [10,11], semiconductor spin qubits [12,13], and photonic qubits [14]. This has motivated significant effort on developing control systems that can also operate adjacent to the qubits, at low temperatures, where the motivation to minimize power consumption becomes even greater.

In principle even three bits of angular resolution would already guarantee a universal computing machine via sequences of discrete Clifford and T gates [15]. However, realizing the desired gate by a sequence of discrete options, one would significantly deepen the overall circuit; this is undesirable in general and particularly so for near-term

---

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International license](https://creativecommons.org/licenses/by/4.0/). Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

quantum computers where one always prefers shallow circuits. Furthermore, focusing on the impact of control limitations on a single gate operation [7], rather than on the average output of a quantum calculation, risks significantly overspecifying the control hardware requirements. Our approach is to consider the output state of the quantum device, and examine its sensitivity to the number of bits used in the angular discretization of the gates (which can ultimately be related to parameters such as the bit resolution and/or sampling rate of a qubit control digital-to-analog converter).

In this Letter, we develop the method of probabilistic angle interpolation (PAI) and show that we can effectively upgrade the capabilities of a physical device with only a set of discrete angles such that on average it can implement any continuous rotation gate—and this is achieved without increasing circuit depths. We do so by randomly instructing the control infrastructure to perform one of the discretized rotation angles, and we subsequently combine the individual outputs of the device such that on average we obtain the same output as the ideal device with infinite resolution. The PAI approach thus allows one to fully exploit the power of algorithms that (nominally) require continuously parametrized gates. It does so with only a marginal increase in repetition (sampling) cost for any reasonable number of parameters as long as the aim is to estimate expected values of observables as relevant in most near-term quantum algorithms [16–18]. As opposed to comparable error mitigation techniques addressing the issue of gate infidelity due to random and unknown variations [19,20] our approach addresses the challenge of discretization. We thus first build a model using unitary gates and later discuss combinations with error mitigation.

*Probabilistic angle interpolation.*—Summary of protocol: Focusing on quantum systems of  $N$  qubits, we consider parametrized quantum gates  $R(\theta) = e^{-i\theta G/2}$  with gate generators  $G$  of eigenvalues  $\pm 1$ . These include parametrized SWAP gates and Pauli gates for any Pauli string  $G \in \{\mathbb{1}, X, Y, Z\}^{\otimes N}$ . These gates encompass most gate sets developed for quantum technologies, such as single qubit  $X$ ,  $Y$ , or  $Z$  rotations or two-qubit  $XX$  entangling gates—and we discuss below that the PAI method can also be extended to any other physical gate sets. We denote the superoperator of our parametrized gates as  $\mathcal{R}(\theta)$ , which acts by conjugation as  $\mathcal{R}(\theta)\rho := e^{-i\theta P_k/2}\rho e^{i\theta P_k/2}$ .

Figure 1 illustrates a physical device that can perfectly perform parametrized gates  $\mathcal{R}(\Theta_k)$  but only with a finite angular resolution of  $B$  bits as

$$\Theta_k = k\Delta \quad \text{with} \quad \Delta = \frac{2\pi}{2^B}, \quad k \in \{0, 1, \dots, 2^B - 1\}, \quad (1)$$

and we detail below the generalization to nonuniformly distributed (nonlinear) set of angles. We define any continuous rotation angle as an over-rotation of one of




FIG. 1. Continuously parametrized Pauli rotation gates  $\mathcal{R}(\theta)$  encompass most typical gates developed for quantum technologies. Because of the use of digital electronics, the rotation angles are divided into  $2^B$  equiangular segments  $\Theta_k$ . In order to reduce engineering complexity, the number  $B$  of bits is chosen as small as possible. PAI realizes an arbitrary, continuous rotation  $\mathcal{R}(\Theta_k + \theta)$  by randomly instructing the quantum hardware to apply one of the two nearest notch settings  $\mathcal{R}(\Theta_k)$  and  $\mathcal{R}(\Theta_{k+1})$  or with a small probability to apply the antipolar rotation  $\mathcal{R}(\Theta_k + \pi)$ .

the discrete settings  $\mathcal{R}(\Theta_k + \theta)$  by an angle  $0 \leq \theta < \Delta$ . Given a relative position  $\lambda = \theta/\Delta$  between two discrete settings, the most simple solution would be to round to the nearest notch; however, this leads to systematic coherent errors that we demonstrate below in numerics can be remarkably severe.

PAI randomly chooses one of three allowed notch settings (Fig. 1) for each parametrized gate in a circuit and exactly implements the desired continuous rotation angle by postprocessing measurement outcomes. In particular, in each circuit execution we randomly choose either the nearest two notch settings  $\Theta_k$  and  $\Theta_{k+1}$  or with a small probability we choose the antipolar angle setting  $\Theta_k + \pi$  as illustrated in Fig. 1. When estimating expected values with PAI, the individual circuit outputs are multiplied by a sign  $-1$  whenever the third rotation angle was chosen.

Thus, the expected value estimation yields a probability distribution in Fig. 3 (left, gray histogram) that is centered around the same mean value that one would obtain via an infinite angular resolution (blue). However, the ability to exactly implement continuous rotations while having access to only discrete rotation angles comes at the price of an increased number of circuit repetitions that scales in the worst case exponentially as  $e^{\nu\Delta^2/4}$  with the number of gates  $\nu$ . We find in Fig. 2, however, that at  $B = 7$  bits of resolution this overhead is still reasonable when the number of parametrized gates in the circuit is not more than a few thousand (of course, there can be arbitrarily many additional nonparameterized gates that align to the available, discrete rotations).

PAI of a single rotation gate: Introducing the notation for the aforementioned discrete notch settings as

$$\mathcal{R}_1 := \mathcal{R}(\Theta_k), \quad \mathcal{R}_2 := \mathcal{R}(\Theta_{k+1}), \quad \mathcal{R}_3 := \mathcal{R}(\Theta_k + \pi), \quad (2)$$

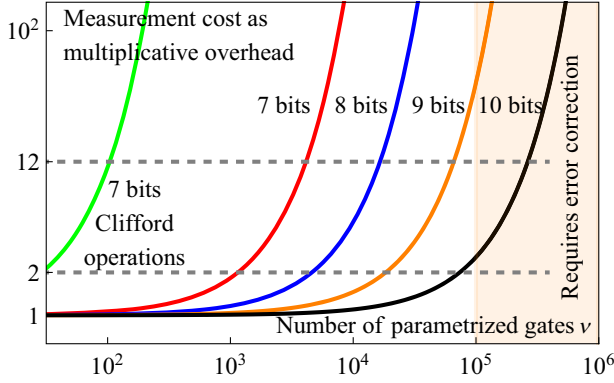


FIG. 2. The measurement cost of PAI is increased compared to the case when one has access to continuous rotation angles; see Eq. (6). Solid lines: worst-case measurement overhead  $\|g\|_1^2$  of PAI as a function of the number of parametrized gates in the quantum circuit. The number of gates one can reasonably (with an overhead at most 12) implement with PAI is approximately  $2^{2(B_{\min}-1)}$  where  $B_{\min}$  is the number of bits used to digitize the rotation angle in Fig. 1. As these estimates rely on worst-case bounds, we observe in numerical simulations that the actual number of gates can be significantly larger. Red vs green lines: our optimal scheme can achieve many orders of magnitude smaller overheads than prior techniques based on Clifford operations [21]. Yellow region: very deep circuits will require quantum error correction and thus even for late noisy intermediate-scale quantum era devices no more than nine bits will be necessary.

the main observation we build on is that we can exactly express any over-rotation  $\mathcal{R}(\Theta_k + \theta)$  as a linear combination of the discrete gates as

$$\mathcal{R}(\Theta_k + \theta) = \gamma_1(\theta)\mathcal{R}_1 + \gamma_2(\theta)\mathcal{R}_2 + \gamma_3(\theta)\mathcal{R}_3. \quad (3)$$

By solving a system of trigonometric equations, we obtain the analytic form of the coefficients  $\gamma_l(\theta)$  as a function of the continuous angle  $\theta$  in the Supplemental Material [22]. In a fashion analogous to quasiprobability sampling methods [19,20,36], which mitigate nonunitary error effects, our angular decomposition leads to the following implementation.

*Statement 1.*—We define a sampling scheme whereby we randomly choose one of the three discrete gate variants  $\{\mathcal{R}_l\}_{l=1}^3$  from Eq. (2) according to the probabilities  $p_l(\theta) = |\gamma_l(\theta)|/\|\gamma(\theta)\|_1$ , where  $\|v\|_1 = \sum_l |v_l|$  is the usual  $\ell_1$  norm. This yields the unbiased estimator of the rotation gate as

$$\hat{\mathcal{R}}(\Theta_k + \theta) = \|\gamma(\theta)\|_1 \text{sign}[\gamma_l(\theta)]\mathcal{R}_l, \quad (4)$$

such that  $\mathbb{E}[\hat{\mathcal{R}}(\Theta_k + \theta)] = \mathcal{R}(\Theta_k + \theta)$ .

To intuitively understand the above approach, we expand the trigonometric probabilities to leading order in the small  $\Delta$ . With probability  $p_1(\theta) = (1 - \lambda) + \mathcal{O}(\Delta^2)$  we apply

the gate at the notch setting  $\Theta_k$ , and with probability  $p_2(\theta) = \lambda + \mathcal{O}(\Delta^2)$  we apply the gate at the next notch setting  $\Theta_{k+1}$ . In leading order this would be equivalent to a naive, non-unitary approximation that we detail in the Supplementary Materials. In contrast, we obtain the desired, unitary operation by applying the antipolar rotation  $\mathcal{R}(\Theta_k + \pi)$  with a small probability  $p_3(\theta) = \frac{1}{4}\lambda(1 - \lambda)\Delta^2 + \mathcal{O}(\Delta^4)$ . Additionally, we multiply any observable measurement-outcome with the factor  $\|\gamma(\theta)\|_1 \text{sign}[\gamma_l(\theta)]$ .

In Theorem 1 we explicitly prove the above approach is optimal in the sense that it yields a minimal  $\|\gamma\|_1$  and present a general solution that can be applied to nonuniform notch settings, too.

*Theorem 1 (informal summary).*—Given any set  $\{\mathcal{R}(\Theta_q)\}$  of discrete (possibly nonuniform) notch settings that a machine can realize, the optimal protocol that minimizes  $\|\gamma\|_1$  uses  $\Theta_k$  and  $\Theta_{k+1}$  as the two nearest notch settings to  $\theta$  and we choose the third gate to be the notch setting nearest to  $\Theta_k + \pi + (\Delta/2)$ , where we defined the distance  $\Delta := \Theta_{k+1} - \Theta_k$ .

PAI of parametrized circuits: We now consider a quantum circuit  $\mathcal{U}_{\text{circ}}$  that contains  $\nu$  continuously parametrized gates and additionally may also contain other nonparameterized gates. We apply Statement 1 to each of the continuously parametrized gates. Given that each parametrized gate has a desired continuous rotation angle  $\Theta_{k_j} + \theta_j$ , we first determine the corresponding notch settings  $(\Theta_{k_1}, \Theta_{k_2}, \dots, \Theta_{k_\nu})$  and corresponding over-rotation angles  $(\theta_1, \theta_2, \dots, \theta_\nu)$  in Eq. (3). At each execution of the circuit we randomly replace a parametrized gate with the corresponding discrete gate variant, i.e., the  $j$ th parametrized gate is replaced by one of the discrete gate variants  $\mathcal{R}_{l_j}^{(j)}$  from Eq. (2) according to the probability distribution  $p_{l_j}(\theta_j)$  from Statement 1.

The result is a set of circuit variants  $\mathcal{U}_{\underline{j}}$  that contain only discrete notch settings according to the multi index  $\underline{j} = (j_1, j_2, \dots, j_\nu) \in 3^\nu$ .

*Statement 2.*—Given a circuit  $\mathcal{U}_{\text{circ}}$  of  $\nu$  parametrized gates we choose a multi index  $\underline{l} \in 3^\nu$  according to the probability distribution  $p(\underline{l}) = |g_{\underline{l}}|/\|g\|_1$  where  $g_{\underline{l}}$  are simply products of the single-gate factors from Statement 1. We obtain an unbiased estimator of the ideal circuit as

$$\hat{\mathcal{U}}_{\text{circ}} = \|g\|_1 \text{sign}(g_{\underline{l}})\mathcal{U}_{\underline{l}} \quad (5)$$

by executing the circuit variants  $\mathcal{U}_{\underline{l}}$  in which all continuously parametrized gates are replaced by the discrete ones according to the multi index  $\underline{l}$ . Thus,  $\mathbb{E}[\hat{\mathcal{U}}_{\text{circ}}] = \mathcal{U}_{\text{circ}}$ .

The above scheme can be compared to probabilistic error cancellation [19,20]. Despite the close formal connection, PAI is quite different conceptually, e.g., all gates involved in PAI are unitary, and PAI does not apply gate insertions but rather applies the same gate at different angle settings.




FIG. 3. Left: distribution of estimated expected values  $\langle Z_0 \rangle$  using 1000 repetitions (shots) in a 12-qubit,  $l = 50$  layer trotter circuit of  $\nu = 1786$  parametrized gates. Left, red: rounding to nearest notch settings at seven bits resolution results in a shifted mean (black vertical line) due to over- or under-rotations. Experimentally estimated histogram (gray) of PAI at seven bits resolution is centered around the same mean as the ideal one (blue, infinite resolution) but its distribution width is slightly increased. Right: energy distance  $\Delta E$  from the ground state during a gradient descent search of a 12-qubit spin-ring problem (energy shown assuming infinite resolution to inform about the optimizer’s progress) using  $\nu = 540$  parametrized gates. Gradient estimation is performed with (blue) infinite resolution and infinite number of shots; (red) rounding to nearest notch settings at seven bits of resolution and using  $10^6$  shots; (green) using PAI with  $10^4$  shots at only 100 different circuit configurations. PAI (green) significantly outperforms the naive approach (red) despite it using the same amount of quantum resources and essentially recovers the performance of the ideal optimizer (blue). Additionally shown is the energy (dashed gray) at the notch settings nearest to the ground-state parameters.

Furthermore, the quasiprobability decomposition in PAI in Eq. (3) is known by construction as opposed to the experimentally learned approximate models in probabilistic error cancellation [36–40].

Estimating expected values: Typical near-term and early fault-tolerant quantum algorithms use quantum computers for estimating expected values  $o = \text{Tr}[OU_{\text{circ}}|0\rangle\langle 0|]$  of an observable  $O$  [16–18]. Thus, one applies a state-preparation circuit to a fixed reference state as  $U_{\text{circ}}|0\rangle\langle 0|$ . One then performs a measurement whose outcome is generally a random variable. By averaging over many repeated measurement outcomes, one obtains an empirical estimate of the expected value  $o$ . Without loss of generality we assume a normalized observable  $\|O\|_{\infty} = 1$  and thus the number of repetitions required to determine the expected value to a precision  $\epsilon$  is upper bounded as  $N_s \leq \epsilon^{-2}$  given the single-shot variance is less than one.

Since we assume access to only discrete rotations, whenever the hardware is instructed to execute the state-preparation circuit  $U_{\text{circ}}$  the parametrized gates are replaced by one of the discrete circuit variants  $U_l$ . After performing a measurement, one multiplies the random outcome with a factor  $\|g\|_1 \text{sign}(g_l)$  that can have negative signs. As a consequence, the variance of the estimator is magnified, which implies an increased number of circuit repetitions.

*Statement 3.*—Applying PAI to the estimation of an expected value results in an unbiased estimator  $\hat{o}$  of the expected value of an observable as  $\mathbb{E}[\hat{o}] = \text{Tr}[OU_{\text{circ}}|0\rangle\langle 0|] = o$ . The number of repetitions required to determine the expected value  $o$  to accuracy  $\epsilon$  is upper bounded as

$$N_s \leq \epsilon^{-2} \|g\|_1^2 = \epsilon^{-2} \prod_{j=1}^{\nu} \|\gamma^{(j)}(\theta_j)\|_1^2, \quad (6)$$

where  $\|g\|_1$  is simply a product of the single-gate norms  $\|\gamma^{(j)}(\theta_j)\|_1$  from Statement 1.

Indeed, to achieve the same precision, PAI has an increased measurement cost compared to having physical access to continuously parametrized gates. In the worst case, when all gate angles are exactly halfway between two notches as  $\theta_j = \Delta_j/2$ , this overhead scales as  $\propto e^{\nu \Delta_{\text{max}}^2/4}$  where  $\Delta_{\text{max}}$  is the largest discretization across the different parametrized gates. The overhead is actually quite reasonable as long as the exponent does not significantly exceed one, as illustrated in Fig. 2. Thus, in order for the circuit repetitions to not exceed a 12-fold increase (gray dashed line in Fig. 2), the number  $\nu$  of parametrized gates in a circuit that can be implemented with PAI is limited by the lowest resolution  $B_{\text{min}}$  of the gate discretizations as  $\nu \leq 2^{2(B_{\text{min}}-1)}$ .

For example, at  $B = 7$  bits resolution 4096 gates can still reasonably be implemented, while 10 bits resolution allows over a quarter of a million gates, which is certainly sufficient for most near-term applications [16–18]. Furthermore, one can significantly reduce these costs by “turning off” PAI for the gates that are not contained in the light cone of the analyzed observable  $O$  [41,42].

*Numerical simulations.*—We consider a typical practical benchmarking task of simulating the spin-ring Hamiltonian

$$\mathcal{H} = \sum_{k \in \text{ring}(N)} \omega_k Z_k + J \vec{\sigma}_k \cdot \vec{\sigma}_{k+1}, \quad (7)$$

with coupling  $J = 0.3$  and uniformly random  $-1 \leq \omega_k \leq 1$ . This spin problem is relevant in condensed-matter physics in understanding many-body localisation in which early quantum computers might be very useful [43–45].

**Time evolution:** We first consider simulating time evolution, which is one of the most natural applications of quantum computers [46–49] and focus on Trotterization that is a commonly applied simulation technique [50]; it approximates the time evolution  $e^{-i\mathcal{H}t}$  as repeated layers of evolutions under the individual Hamiltonian terms for small times  $\delta t$ . Since the evolution under each Hamiltonian term in Eq. (7) is a Pauli rotation gate, e.g.,  $\mathcal{R}(\omega_k \delta t)$ , a layer of the time evolution circuit is just a series of rotation gates, each tuned to its relevant small rotation angle. This layer is then repeated a large number  $t/\delta t$  of times.

Rounding the rotation angles to nearest notch settings, e.g.,  $\omega_k \delta t \mapsto \Theta_1$ , leads to a significant coherent error as it implements incorrect evolution times and/or incorrect interaction terms  $\omega_k$ . Thus, near phase transitions the discrepancy might be radical. Thus, measuring an expected value with a fixed number of shots leads to a biased probability distribution as we illustrate in Fig. 3 (left, red). In contrast, PAI results in a probability distribution that is centered around the exact mean in Fig. 3 (left, blue and gray), while its distribution width is slightly increased. The increase in width is actually significantly lower than our worst-case estimates in Statement 3, which we analyze in the Supplemental Material [22]—and we also discuss that PAI is highly enabling for randomized compilers, such as qDRIFT [23].

**Finding eigenstates:** We next consider finding eigenstates of the Hamiltonian in Eq. (7). A broad range of techniques are available in the literature, including ones that target near-term and early fault-tolerant quantum computers [16–18,33,49,51,52]. We use the same Trotterized circuit structure as we used for time evolution but we optimize the angles of the rotation gates so that the energy  $\text{Tr}[\mathcal{H}\mathcal{U}_{\text{circ}}|0\rangle\langle 0|]$  of the emerging state is minimal—this variant of the variational quantum eigensolver uses the Hamiltonian variational ansatz as in case of the quantum approximate optimization algorithm [16–18,53].

Figure 3 (right, red lines) illustrates that a gradient descent optimizer does not manage to meaningfully lower the energy when the gradient is calculated using only nearest notch settings due to the coherent discrepancy in the output state. In contrast, using the same quantum resources (same number of circuit repetitions and discretized gates) but estimating the gradient using PAI matches the performance of an ideal quantum circuit that has infinite angular resolution in Fig. 3 (right, green and blue lines). We also note that formally our PAI protocol applies a different, randomly chosen circuit variant at each circuit repetition. However, reconfiguring circuits will likely be a bottleneck for some quantum hardware platforms and thus it is desired

to run the same circuit variant multiple times. Indeed, Fig. 3 (right, green) only uses 100 different circuit variants—each of which is repeated  $10^4$  times—which demonstrably does not compromise the optimizer’s performance.

*Discussion and conclusion.*—We present PAI, which effectively upgrades the capabilities of a quantum hardware that can only realize discrete rotation angles to a device that can perform arbitrary, continuous rotation angles.

As we detail in the Supplemental Material [22], a number of generalizations and further applications of our approach are apparent, including its combination with well-established error mitigation techniques. We thus conclude that the present technique will be an important and useful tool in designing optimal quantum hardware. Our analysis suggests that early quantum hardware, being practically limited to only a few thousand gate operations (due to limited coherence times), will need no more than seven bits of resolution in the control systems. As the technology matures, future generations of hardware are expected to be able to execute tens of thousands of quantum gates without error correction, which still, however, requires no more than nine bits of angular resolution.

B. K. thanks the University of Oxford for a Glasstone Research Fellowship and Lady Margaret Hall, Oxford for a Research Fellowship. The numerical modeling involved in this study made use of the Quantum Exact Simulation Toolkit (QuEST), and the recent development QuESTlink [54], which permits the user to use *Mathematica* as the integrated front end, and pyQuEST [55], which allows access to QuEST from PYTHON. We are grateful to those who have contributed to all of these valuable tools. The authors would like to acknowledge the use of the University of Oxford Advanced Research Computing (ARC) facility [56] in carrying out this work and specifically the facilities made available from the EPSRC QCS Hub Grant (Agreement No. EP/T001062/1). The authors also acknowledge funding from the EPSRC projects Robust and Reliable Quantum Computing (RoRQ, EP/W032635/1) and Software Enabling Early Quantum Advantage (SEEQA, EP/Y004655/1).

\*Corresponding author: koczor@maths.ox.ac.uk

†Corresponding author: simon.benjamin@materials.ox.ac.uk

- [1] Y. Xu, G. Huang, J. Balewski, R. Naik, A. Morvan, B. Mitchell, K. Nowrouzi, D. I. Santiago, and I. Siddiqi, *IEEE Trans. Quantum Eng.* **2**, 1 (2021).
- [2] E. Mount, D. Gaultney, G. Vrijsen, M. Adams, S.-Y. Baek, K. Hudek, L. Isabella, S. Crain, A. van Rynbach, P. Maunz, and J. Kim, *Quantum Inf. Process.* **15**, 5281 (2016).
- [3] K. Das and T. Lehmann, in *Proceedings of the 2011 IEEE 54th International Midwest Symposium on Circuits and Systems (MWSCAS)* (IEEE, New York, 2011).
- [4] X. Xue *et al.*, *Nature (London)* **593**, 205 (2021).

- [5] J. C. Bardin *et al.*, *IEEE J. Solid-State Circuits* **54**, 3043 (2019).
- [6] S. J. Pauka, K. Das, R. Kalra, A. Moini, Y. Yang, M. Trainer, A. Bousquet, C. Cantaloube, N. Dick, G. Gardner, M. J. Manfra, and D. Reilly, *Nat. Electron.* **4**, 64 (2021).
- [7] J. P. G. van Dijk, E. Kawakami, R. N. Schouten, M. Veldhorst, L. M. K. Vandersypen, M. Babaie, E. Charbon, and F. Sebastiano, *Phys. Rev. Appl.* **12**, 044054 (2019).
- [8] M. F. Gonzalez-Zalba, S. de Franceschi, E. Charbon, T. Meunier, M. Vinet, and A. S. Dzurak, *Nat. Electron.* **4**, 872 (2021).
- [9] R. Acharya *et al.*, *Nature (London)* **614**, 676 (2023).
- [10] G. Pagano, P. W. Hess, H. B. Kaplan, W. L. Tan, P. Richerme, P. Becker, A. Kyprianidis, J. Zhang, E. Birkelbaw, M. R. Hernandez, Y. Wu, and C. Monroe, *Quantum Sci. Technol.* **4**, 014004 (2018).
- [11] M. F. Brandl, M. W. van Mourik, L. Postler, A. Nolf, K. Lakhmanskii, R. R. Paiva, S. Möller, N. Daniilidis, H. Häffner, V. Kaushal, T. Ruster, C. Warschburger, H. Kaufmann, U. G. Poschinger, F. Schmidt-Kaler, P. Schindler, T. Monz, and R. Blatt, *Rev. Sci. Instrum.* **87**, 113103 (2016).
- [12] S. Schaal, A. Rossi, V. N. Ciriano-Tejel, T.-Y. Yang, S. Barraud, J. J. L. Morton, and M. F. Gonzalez-Zalba, *Nat. Electron.* **2**, 236 (2019).
- [13] S. G. J. Philips, M. T. Madzik, S. V. Amitonov, S. L. de Snoo, M. Russ, N. Kalhor, C. Volk, W. I. L. Lawrie, D. Brousse, L. Tryputen, B. P. Wuetz, A. Sammak, M. Veldhorst, G. Scappucci, and L. M. K. Vandersypen, *Nature (London)* **609**, 919 (2022).
- [14] M. Dong, G. Clark, A. J. Leenheer, M. Zimmermann, D. Dominguez, A. J. Menssen, D. Heim, G. Gilbert, D. Englund, and M. Eichenfield, *Nat. Photonics* **16**, 59 (2022).
- [15] C. M. Dawson and M. A. Nielsen, *arXiv:quant-ph/0505030*.
- [16] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles, *Nat. Rev. Phys.* **3**, 625 (2021).
- [17] S. Endo, Z. Cai, S. C. Benjamin, and X. Yuan, *J. Phys. Soc. Jpn.* **90**, 032001 (2021).
- [18] K. Bharti, A. Cervera-Lierta, T. H. Kyaw, T. Haug, S. Alperin-Lea, A. Anand, M. Degroote, H. Heimonen, J. S. Kottmann, T. Menke, W.-K. Mok, S. Sim, L.-C. Kwak, and A. Aspuru-Guzik, *Rev. Mod. Phys.* **94**, 015004 (2022).
- [19] K. Temme, S. Bravyi, and J. M. Gambetta, *Phys. Rev. Lett.* **119**, 180509 (2017).
- [20] S. Endo, S. C. Benjamin, and Y. Li, *Phys. Rev. X* **8**, 031027 (2018).
- [21] Y. Suzuki, S. Endo, K. Fujii, and Y. Tokunaga, *PRX Quantum* **3**, 010345 (2022).
- [22] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevLett.132.130602>, which includes Refs. [23–35], for further details and for proofs of the mathematical statements.
- [23] E. Campbell, *Phys. Rev. Lett.* **123**, 070503 (2019).
- [24] M. Braun and S. J. Glaser, *J. Magn. Reson.* **207**, 114 (2010).
- [25] O. R. Meitei, B. T. Gard, G. S. Barron, D. P. Pappas, S. E. Economou, E. Barnes, and N. J. Mayhall, *npj Quantum Inf.* **7**, 155 (2021).
- [26] R. De Keijzer, O. Tse, and S. Kokkelmans, *Quantum* **7**, 908 (2023).
- [27] J. Foldager and B. Koczor, *J. Phys. A: Math. Theor.* 015306 (2023).
- [28] H. Y. Huang, R. Kueng, and J. Preskill, *Nat. Phys.* **16**, 1050 (2020).
- [29] K. Mitarai and K. Fujii, *New J. Phys.* **23**, 023021 (2021).
- [30] C. Piveteau and D. Sutter, *arXiv:2205.00016*.
- [31] F. Martins, F. K. Malinowski, P. D. Nissen, E. Barnes, S. Fallahi, G. C. Gardner, M. J. Manfra, C. M. Marcus, and F. Kuehnmeth, *Phys. Rev. Lett.* **116**, 116801 (2016).
- [32] G. Wolfowicz, M. Urdampilleta, M. L. W. Thewalt, H. Riemann, N. V. Abrosimov, P. Becker, H.-J. Pohl, and J. J. L. Morton, *Phys. Rev. Lett.* **113**, 157601 (2014).
- [33] B. Koczor and S. C. Benjamin, *Phys. Rev. Res.* **4**, 023017 (2022).
- [34] Y. Kim, A. Eddins, S. Anand, K. X. Wei, E. Van Den Berg, S. Rosenblatt, H. Nayfeh, Y. Wu, M. Zaletel, K. Temme *et al.*, *Nature (London)* **618**, 500 (2023).
- [35] B. Koczor, *Phys. Rev. X* **11**, 031057 (2021).
- [36] Z. Cai, R. Babbush, S. C. Benjamin, S. Endo, W. J. Huggins, Y. Li, J. R. McClean, and T. E. O'Brien, *arXiv:2210.00921*.
- [37] A. Strikis, D. Qin, Y. Chen, S. C. Benjamin, and Y. Li, *PRX Quantum* **2**, 040330 (2021).
- [38] A. Lowe, M. H. Gordon, P. Czarnik, A. Arrasmith, P. J. Coles, and L. Cincio, *Phys. Rev. Res.* **3**, 033098 (2021).
- [39] E. Van Den Berg, Z. K. Mineev, A. Kandala, and K. Temme, *Nat. Phys.* **19**, 1116 (2023).
- [40] A. Montanaro and S. Stanisic, *arXiv:2102.02120*.
- [41] H. Jnane, J. Steinberg, Z. Cai, H. C. Nguyen, and B. Koczor, *PRX Quantum* **5**, 010324 (2024).
- [42] M. C. Tran, K. Sharma, and K. Temme, *arXiv:2303.06496*.
- [43] R. Nandkishore and D. A. Huse, *Annu. Rev. Condens. Matter Phys.* **6**, 15 (2015).
- [44] A. M. Childs, D. Maslov, Y. Nam, N. J. Ross, and Y. Su, *Proc. Natl. Acad. Sci. U.S.A.* **115**, 9456 (2018).
- [45] D. J. Luitz, N. Laflorencie, and F. Alet, *Phys. Rev. B* **91**, 081103(R) (2015).
- [46] S. McArdle, S. Endo, A. Aspuru-Guzik, S. C. Benjamin, and X. Yuan, *Rev. Mod. Phys.* **92**, 015003 (2020).
- [47] Y. Cao, J. Romero, J. P. Olson, M. Degroote, P. D. Johnson, M. Kieferová, I. D. Kivlichan, T. Menke, B. Peropadre, N. P. Sawaya *et al.*, *Chem. Rev.* **119**, 10856 (2019).
- [48] H. H. S. Chan, R. Meister, M. L. Goh, and B. Koczor, *arXiv:2212.11036*.
- [49] G. Boyd and B. Koczor, *Phys. Rev. X* **12**, 041022 (2022).
- [50] D. W. Berry, G. Ahokas, R. Cleve, and B. C. Sanders, *Commun. Math. Phys.* **270**, 359 (2007).
- [51] B. van Straaten and B. Koczor, *PRX Quantum* **2**, 030324 (2021).
- [52] B. Koczor and S. C. Benjamin, *Phys. Rev. A* **106**, 062416 (2022).
- [53] E. Farhi, J. Goldstone, and S. Gutmann, *arXiv:1411.4028*.
- [54] T. Jones and S. Benjamin, *Quantum Sci. Technol.* **5**, 034012 (2020).
- [55] R. Meister, pyQuEST—A Python interface for the Quantum Exact Simulation Toolkit (2022), <https://github.com/rmeister/pyQuEST>.
- [56] A. Richards, University of Oxford Advanced Research Computing (2015), 10.5281/zenodo.22558.