Demonstration of Algorithmic Quantum Speedup

Bibek Pokharel[1,*] and Daniel A. Lidar[2,†]
[1]Department of Physics & Astronomy and Center for Quantum Information Science & Technology, University of Southern California, Los Angeles, California 90089, USA
[2]Departments of Electrical & Computer Engineering, Chemistry, and Physics & Astronomy, and Center for Quantum Information Science & Technology, University of Southern California, Los Angeles, California 90089, USA

(Received 19 November 2022; accepted 20 April 2023; published 26 May 2023)

Despite the development of increasingly capable quantum computers, an experimental demonstration of a provable algorithmic quantum speedup employing today's non-fault-tolerant devices has remained elusive. Here, we unequivocally demonstrate such a speedup within the oracular model, quantified in terms of the scaling with the problem size of the time-to-solution metric. We implement the single-shot Bernstein-Vazirani algorithm, which solves the problem of identifying a hidden bitstring that changes after every oracle query, using two different 27-qubit IBM Quantum superconducting processors. The speedup is observed on only one of the two processors when the quantum computation is protected by dynamical decoupling but not without it. The quantum speedup reported here does not rely on any additional assumptions or complexity-theoretic conjectures and solves a bona fide computational problem in the setting of a game with an oracle and a verifier.

DOI: 10.1103/PhysRevLett.130.210602

The quest to demonstrate a quantum speedup using physical hardware for a computational problem over a range of increasing problem sizes—an algorithmic quantum speedup—has motivated the field of quantum computing from its inception [1]. Many quantum algorithms are now known that theoretically outperform classical algorithms in solving problems of increasing size [2–14]. However, computational errors must be suppressed to realize this potential, especially in the current noisy intermediate-scale quantum (NISQ) era [15]. Better-than-classical algorithmic performance has been demonstrated a number of times, e.g., on ion-trap [16,17], superconducting [11,18–23], photonic [24–28], and Rydberg atom [29] quantum processors. In most cases, this was done by exceeding the corresponding classical algorithmic success probability at a fixed or small set of problem sizes [16,17,21,25,28], by outperforming a limited set of classical algorithms [18,19,29], or under complexity-theoretic assumptions [20,22,24,26]. See [30] for a survey of existing experimental demonstrations of better-than-classical algorithmic results.

To qualify as a provable, unqualified algorithmic quantum speedup, we stipulate that the speedup—quantified in terms of the scaling with problem size of the time-to-solution (TTS) metric [56]—is (i) relative to the best possible classical algorithm ("provable") and (ii) free of any assumptions or conjectures ("unqualified"). Moreover, to avoid detectable finite-size effects, we stipulate that (iii) the speedup is exhibited up to the largest experimentally accessible problem size. Similar criteria were established in Ref. [56]. However, no provable, unqualified algorithmic quantum speedup that satisfies these stringent criteria has been empirically demonstrated so far. Here, we provide the first such demonstration. To achieve this, we revisit the Bernstein-Vazirani (BV) algorithm, which was one of the very first theoretical examples of a quantum vs classical complexity class separation [3]. In the original BV problem, an oracle outputs $f_b(x) = b \cdot x (\mathrm{mod}\ 2) \in \{0, 1\}$, where x and b are both length-n bitstrings. Here, x is a guess provided by the user, and b is a secret bitstring the user is trying to learn in as few oracle queries as possible. The best classical algorithm requires n queries, since each query can only provide one new bit of information about b. By solving the problem with certainty in a *single* query, the BV algorithm provides a linear speedup over the best-classical algorithm.

Here, we consider a modified, *single-shot* version of BV, denoted ssBV-n, where the hidden bitstring b changes after every query. We colloquially refer to this as the "BV guessing game": after one query of the single-shot oracle, the player is allowed one guess of the bitstring b. If the verifier confirms that the guess is correct, the player wins; if the guess is wrong, the game continues with a new bitstring. In this setting, the optimal classical algorithm is to query the oracle with $x = 0 \ldots 01_i 0 \ldots 0$ (i is arbitrary), which reveals b_i, and then guess the remaining $n - 1$ bits. This yields classical success probability $p_s = 2^{1-n}$, only twice better than a random guess (this cannot be improved [57]). In stark contrast, a player with access to a quantum computer (QC) running the original BV algorithm has success probability $p_s = 1$ after each query, which

becomes an exponential advantage in the speedup ratio (defined below) over the classical setting.

Quantum speedup quantified.—In a head-to-head comparison of success probabilities, $p_s > 2^{1-n_0}$, for a fixed problem size n_0, implies a better-than-classical result. This is the context in which better-than-classical results have been achieved for the Grover and BV algorithms [16,17,21]. However, the success probability at a fixed problem size is not a reliable measure of quantum speedup, as detecting an algorithmic speedup requires computing the scaling with problem size. Moreover, p_s is itself a function of the time $t_r(n)$ taken to run the calculation, i.e., the time required to run the complete quantum or classical circuit once. Instead, we quantify quantum speedup in terms of the scaling with the problem size n of the speedup ratio of the classical and quantum total run times: $S(n) = [\mathrm{TTS}_C(n)/\mathrm{TTS}_Q(n)]$, where the total runtime is quantified using the well-established TTS metric [56]:

$$\mathrm{TTS}(n) = t_r(n)R(n), \qquad R(n) = \frac{\log(1-p_d)}{\log(1-p_s[t_r(n)])}. \quad (1)$$

$\lceil R(n) \rceil$ is the number of repetitions—oracle calls in the present context—needed to find a solution at least once with desired probability p_d, given that a single repetition succeeds with probability $p_s[t_r(n)]$; we set $p_d = 0.99$ henceforth. Thus, the TTS quantifies the total time it takes to win the BV guessing game, whether classically or with access to a QC.

We choose to measure $t_r(n)$ in terms of the circuit execution time and the readout duration and ignore the postprocessing overhead, as the latter is a constant in our experiments, set by the time it takes to count the number of times the secret bitstring b appears out of the constant total number S of "shots" (circuit runs); we use the maximum allowed number S for all n. Also, as detailed below, the circuit unitaries are specified in terms of the native gates of the device and do not incur a compilation overhead. It follows from the BV circuit structure (Fig. 1) that $t_r(n) = c\tau_{2q}n + \tau_0$, where $1 \le c \le 2$ depends on the qubit connectivity graph, with the two limits corresponding to all-to-all connectivity ($c = 1$) and a chain ($c = 2$). For our IBM Quantum implementation, we found $c \approx 1.76$ [30]. The two-qubit gate time, τ_{2q}, and the sum of the single qubit and readout times, τ_0, depend on the specific QC and can vary by orders of magnitude across platforms.

When accounting for gate and measurement imperfections, we expect $\mathrm{TTS}_Q(n)$ in the ssBV-n case to scale as $n2^{\lambda n}$ (with $\lambda > 0$, an effective noise parameter), instead of as $t_r \sim n$, as would be the case for a noiseless QC. The factor $2^{\lambda n}$ arises from a (naive) noise model wherein gate and measurement fidelities multiply in a circuit of depth $O(n)$, yielding a TTS denominator that scales as $\log(1-2^{-\lambda n}) \approx 2^{-\lambda n}$. When ssBV-$n$ is solved classically, computing $f_b(x) = b \cdot x (\mathrm{mod}\ 2)$ also takes

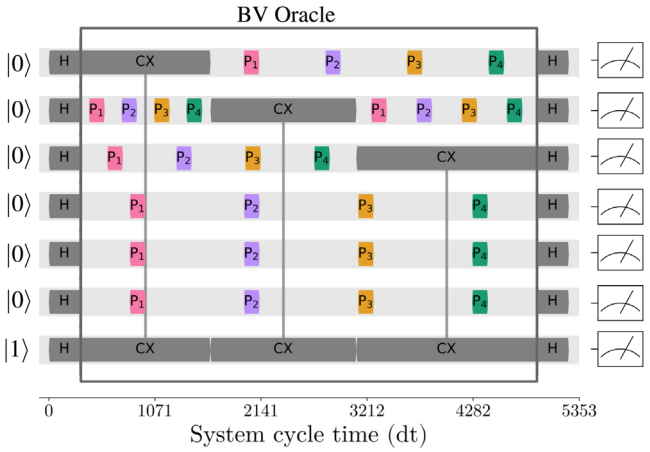

FIG. 1. Circuit for the BV algorithm, including dynamical decoupling (DD) pulses. The oracle shown encodes the unknown bitstring $b = 111\,000$ for the ssBV-6 problem. A controlled-NOT (CNOT, or CX) or identity gate is performed from qubit i to the ancilla qubit if $b_i = 1$ or 0, respectively. Note that the quantum and classical oracles are identical in the ssBV-n problem, and so both take time $t_r \propto |b|$ to run, where $|b|$ is the Hamming weight of b. Each BV-n circuit requires $n+1$ qubits. A Hadamard gate (H) is applied to each qubit before and after the oracle, and each qubit is measured in the computational basis for a total circuit depth $\ge |b| + 3$ (with equality only for fully connected architectures). DD pulses (P_i) are turned on during idle times. Pulse placement is schematic but illustrates the principles we used in practice: (i) DD fills all available idle times; (ii) pulse intervals are varied depending on the available idle time per qubit. The actual timeline is shown in units of $dt = 2/9$ ns—the inverse sampling rate of the backend's arbitrary waveform generators.

time $\propto n$ (the cost of adding n bits), so we obtain $\mathrm{TTS}_C(n) \propto n/\log(1-2^{1-n}) \approx n2^{n-1}$. We thus expect

$$S(n) \sim 2^{(1-\lambda)n}, \qquad n \in [n_{\min}, n_{\max}], \quad (2)$$

where $n_{\max}$ denotes the largest number of data qubits available to the quantum algorithm and $n_{\min}$ is identified empirically by excluding small-size effects. *We will declare a quantum speedup if the speedup exponent* $\lambda < 1$. It is important to emphasize that the speedup exponent must be extracted using n reaching up to and including $n_{\max}$, since otherwise, one cannot hope to draw conclusions that reflect asymptotic scaling behavior. Using this criterion, we demonstrate below that a statistically significant quantum speedup is achieved for DD-protected ssBV-n quantum circuits, but no speedup is obtained for "bare" quantum circuits implemented without DD.

Dynamical decoupling.—DD protection has a long history of experimental demonstrations on various quantum devices (see Ref. [58] for a review), and has also been shown to improve various performance metrics, such as qubit memory fidelity [59,60], crosstalk mitigation [61,62], quantum volume [63], and algorithmic fidelity [64].

However, we are unaware of prior experimental demonstrations of the use of DD to directly improve quantum algorithmic scaling.

A major challenge in using DD is that pulse imperfections can significantly deteriorate performance, necessitating a careful choice of DD sequence. Building on a survey of numerous known sequences [65], we selected the universally robust (UR_p) sequence family [66] as the top performance enhancer. This sequence was designed to suppress pulse axis and angle errors and has been shown previously to enhance performance in superconducting-qubit-based NISQ devices [60,67]. It suppresses pulse errors up to $O(\varepsilon^{p/2})$ using p pulses, where ε is the nominal infidelity [66].

Experimental implementation.—We implemented ssBV-n on two different 27-qubit QCs: Montreal and Cairo. While similar in their connectivity, they have different quantum volumes, qubit generations, and gate fidelities [30]. Since the BV algorithm uses one ancilla, we have $n_{\max} = 26$.

Given the unknown string b, the BV oracle is implemented by performing CNOTs from a subset of the first n qubits to the ancilla qubit (numbered $n+1$), and the number of CNOTs is the Hamming weight $k = |b|$ (see Fig. 1). There are 2^n different oracles, and in the ssBV-n problem with n fixed, one is selected at random in each round. However, to increase our confidence in the results we exploited the fact that given n and k, the circuits for all $\binom{n}{k}$ distinct bitstrings are identical up to qubit permutation; we used this symmetry and tested the $n+1$ permutationally inequivalent strings $b = 1^k 0^{n-k}$ with $0 \le k \le n$ for each n.

For DD, we employed a "decouple then compute" strategy [68,69], whereby pulses constituting short but complete DD sequences are inserted into the idle intervals of the quantum circuit. These idle intervals arise because the algorithm specifies them or due to limited connectivity of the underlying architecture, which requires information swapping between some qubits while others are idle. We implemented one repetition of UR_{14} and UR_{18} per idle interval on Montreal and Cairo, respectively; see Fig. 1.

We took $S = 100$ K (32 K) shots using Cairo (Montreal) for each unique circuit. We then sampled the corresponding results for all BV-n oracles using bootstrapping [70] and report the mean TTS for BV-n along with error bars corresponding to $\pm 5\sigma$ for the bootstrapped distribution. See [30] for more experimental implementation details.

Results.—The Cairo results for BV-6, both with and without DD, are shown in Fig. 2. The oracles and outputs bitstrings are sorted by increasing Hamming weight. It is clear from these results that a higher Hamming weight results in a decreasing success probability without DD; this is consistent with our expectation that deeper circuits have a lower overall fidelity. With DD, this problem is significantly mitigated, which already suggests that error suppression through DD will be central to our quantum speedup demonstration. In fact, with DD the single-shot output success probability exceeds $1/2$ for all oracles, which allows reaching the bounded-error quantum polynomial (BQP) threshold of $2/3$ for all possible inputs by classical majority vote on multiple repetitions [3,17]. Without DD, the single-shot output success probability is below $1/2$ for $7/64$ of the inputs, so for these inputs, the BQP threshold cannot be reached. With (without) DD, the average single-shot success probability is 80.3% (73.2%). While this is much higher than the classical single-shot probability of $2^{-5} \sim 3\%$, it does not suffice for claiming a quantum speedup, as this requires that we demonstrate a scaling advantage as a function of the problem size n. Moreover, even demonstrating such an advantage just for $p_s(n)$ is insufficient [30].

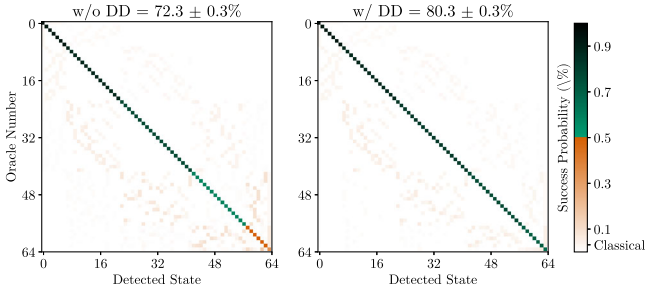

FIG. 2. Full output distribution for BV-6 from Cairo. Oracles f_b are numbered from 0 to 63, corresponding to $b \in \{0^6, \ldots, 1^6\}$, sorted by increasing Hamming weight. Ideally, the output state for oracle f_b (vertical axis) is b, but in reality, other bitstrings (horizontal axis) are observed as well. Green dots on the diagonal correspond to $p_s > 1/2$, where p_s is the empirical frequency (success probability) with which b was output for oracle f_b. Success probabilities are reported with 5σ confidence intervals.

Our main result is presented in Fig. 3, which shows the TTS vs the problem size n for both Montreal and Cairo. White grid lines show the classical TTS (scaling as $n2^{n-1}$), and the ideal quantum TTS (equal to $t_r \sim n$) is shown for reference by the two dashed lines—one each for Montreal and Cairo. As is apparent, the scaling without DD (empty symbols) for both devices is worse than the classical scaling at large n. We attribute this, beyond the aforementioned exponential fidelity loss with circuit depth, to the fact that transmon-based devices suffer from spontaneous emission errors, as a result of which they preferentially generate bitstrings with low Hamming weight, which is worse than a uniformly random guess. This is also consistent with the result shown in Fig. 2 (left).

With DD, this problem is mitigated, so that $p_s > 0$ is extended for Cairo (blue) to $n = n_{\max} = 23$ (excessive readout noise required us to treat Cairo as a device with $n_{\text{tot}} = 24$ [30]). Most notably, it is clear that with DD the Montreal scaling (orange) is better than classical and extends to $n = n_{\max} = 26$, suggesting a quantum speedup.

To quantify this and extract the speedup exponent λ as conservatively as possible, we compute the worst-case

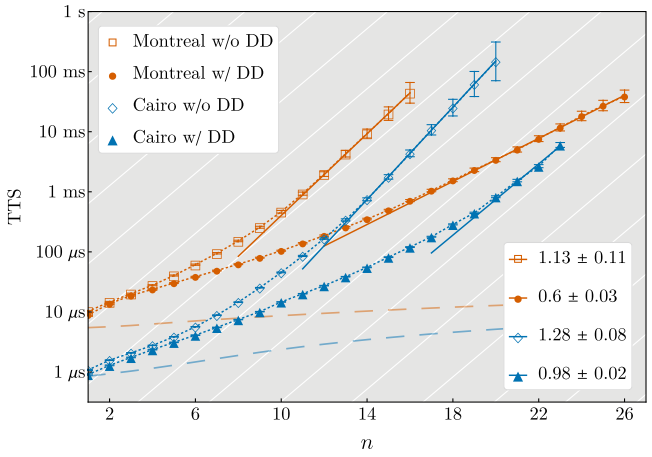

FIG. 3. Time to solution (TTS) as a function of problem size or number of data qubits n. We report $\mathrm{TTS}(n) = (1/2^n)\sum_b \mathrm{TTS}(n,b)$, where $\mathrm{TTS}(n,b)$ is given by Eq. (1), with $p_d = 0.99$ and $t_r(n)$ replaced by $t_r(n,b)$, since each oracle (labeled by $b \in \{0,1\}^n$) takes a different time to run. Results for Montreal and Cairo are shown by the orange and blue symbols, respectively, and filled (empty) symbols represent results with (without) DD; dotted lines are guides to the eye. The asymptotic classical scaling $\mathrm{TTS}_C(n) \sim 2^n$ is shown as white grid lines, and the hypothetical, ideal quantum scaling $\mathrm{TTS}_Q(n) \propto n$ of each QC is indicated by the dashed lines (for QC-specific parameter values see [30]). The solid lines give the worst-case scaling fit for each curve, whose slopes λ are reported in the bottom legend, with uncertainties representing 95% confidence intervals. Without DD, the TTS curves terminate at $n'_{\max} = 16$ ($n'_{\max} = 20$) for Montreal (Cairo), since we find $p_s = 0$ for $n > n'_{\max}$. Moreover, $\lambda > 1$ without DD, indicating a worse-than-classical scaling. With DD protection, on Cairo, the $p_s > 0$ range is extended to $n = 23$, and λ is just below the breakeven point of 1, but the uncertainty is too large to conclude that quantum speedup has occurred. In contrast, the Montreal scaling with DD does exhibit quantum speedup. Since two-qubit operations and readout durations are shorter for Cairo, it exhibits a consistently lower absolute TTS than Montreal. We report 5σ confidence intervals from bootstrapping for each data point; error bars are mostly covered by the symbols.

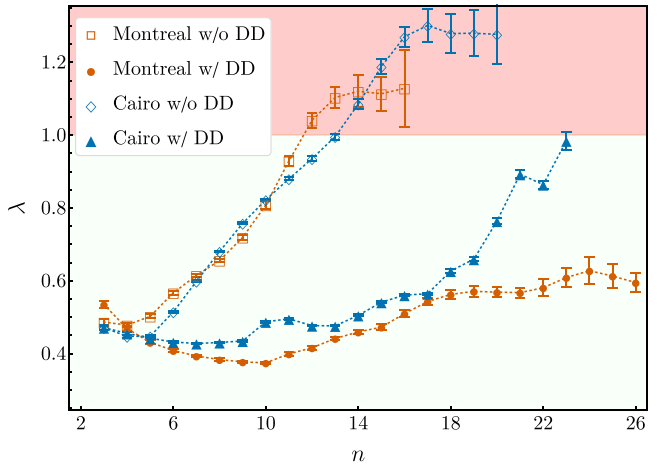

FIG. 4. Results for $\lambda_{h_{\max}}$, the maximum local slope of each of the curves in Fig. 3 for $n \le h_{\max}$, i.e., the worst-case scaling when Fig. 3 is restricted to $h_{\max}+1$ qubits. Only Montreal with DD exhibits an unambiguous quantum speedup, with $\lambda_{h_{\max}}$ well below 1 for all $n \le h_{\max}$. Error bars represent 2σ confidence intervals.

scaling from our data [30]. The results are shown as the straight blue and orange lines in Fig. 3, along with the numerical values of λ in the legend. Without DD, we obtain $\lambda = 1.13 \pm 0.11$ and 1.28 ± 0.08 for Montreal and Cairo, respectively, meaning a quantum slowdown. For Cairo, the scaling with DD is $\lambda = 0.98 \pm 0.02$, not a statistically significant difference from the classical scaling. However, the fit confirms that *Montreal with DD exhibits an algorithmic quantum speedup*: $\lambda = 0.60 \pm 0.03$. All the reported uncertainties represent 2σ symmetric confidence intervals [30]. The difference between Cairo and Montreal agrees with the reported larger quantum volume (128 vs 64) of Montreal [71], and suggests that the latter is a relevant performance metric also in the present context of algorithmic speedups.

All the slopes vary with n in Fig. 3. One might thus ask what the scaling would appear to be for a hypothetical QC with fewer qubits ($h_{\max}$) than the actual $n_{\max} = 26$; we address this in Fig. 4. This figure shows the maximum local slope of each of the curves in Fig. 3 for $n \le h_{\max}$ [30]. The results clearly show the growth of the speedup exponents $\lambda_{h_{\max}}$ for Cairo with and without DD, and for Montreal without DD, to the point $\lambda > 1$ or beyond, where no quantum speedup survives. In contrast, the speedup exponent for Montreal with DD is well within the quantum speedup region of $\lambda < 1$ for all values of $h_{\max}$.

Discussion and conclusions.—The ssBV-n problem has a provable, conjecture-free exponential speedup over the best possible classical algorithm in the setting of a game involving an oracle and a verifier. The main weakness of this setting is its oracular nature: we are forced to hide the internal structure of the circuit from the players since the BV circuit can be efficiently simulated classically by virtue of the fact that it uses only Clifford gates [72]. In contrast, quantum supremacy is not subject to oracular restrictions and is in this sense a more interesting type of quantum advantage. However, this advantage only holds under certain conjectures. Another interesting class of non-oracular speedups is quantum constant depth circuits vs classical logarithmic depth circuits [9,10] and quantum limited-space advantage [11]. Here, the assumption is a classical resource constraint. Some sort of tradeoff between computational complexity assumptions, constraints, and oracularity thus appears to be inevitable.

To test for a quantum speedup, we compared the asymptotic scaling of the TTS metric with problem size for both classical and quantum algorithms. We demonstrated a statistically significant algorithmic quantum

speedup on Montreal using this metric. A crucial feature in our implementation was error suppression through DD, without which the speedup was not exhibited.

It is natural to question whether this speedup can be expected to continue indefinitely. Given the flatness of $\lambda_{h_{\max}}$ seen in Fig. 4, extrapolating the quantum speedup result for Montreal to n slightly larger than $n_{\max}$ appears reasonable. However, for $n \gg n_{\max}$ the DD-enabled speedup cannot survive, since in the absence of a mechanism for entropy removal, such as fault-tolerant quantum error correction [73], decoherence always eventually dominates. Thus, one should expect the DD-enabled quantum speedup to disappear at some finite upper limit on n. The fact that this upper limit is not observed in our experiments satisfies a key goal of implementing a quantum algorithm on a NISQ device, namely to check whether a quantum advantage is already accessible even before the advent of fault tolerance, up to the largest problem sizes supported by the device. We have shown here that, with the help of error suppression via DD, this is indeed the case.

Another natural question is to what extent the speedup reported here can be further improved. We certainly expect that methods such as measurement error mitigation [74] and further DD sequence optimization [64,75–77] will have such an effect, though TTS_Q should then account for the additional classical computation time they incur. For measurement error mitigation, this cost overwhelms the quantum speedup we have observed [78]. Device-tailored optimization of DD sequences with advanced low-level pulse control is an exciting frontier that remains largely unexplored and appears particularly promising. While we focused on superconducting-qubit devices, DD protection can be beneficial across platforms, as all NISQ devices are affected by computational errors such as decoherence and crosstalk.

An ideal quantum computer would yield an exponential TTS speedup for the ssBV-n problem. Our results are comparatively less impressive: we demonstrated what amounts to a polynomial quantum speedup, by reducing the exponent of the TTS scaling to below its classical minimum. Our work provides a path to testing such speedups across platforms and algorithms in the NISQ era.

Access to the IBM Quantum Network was obtained through the membership of UNM in the IBM Quantum Hub at NC State. We are grateful to UNM for generously providing us access through their membership. This material is based in part upon work supported by the National Science Foundation the Quantum Leap Big Idea under Grant No. OMA-1936388. This work is also partially supported by a DOE/HEP QuantISED program grant, QCCFP/Quantum Machine Learning and Quantum Computation Frameworks (QCCFP-QMLQCF) for HEP, Grant No. DE-SC0019219. We are grateful to Dr. Namit Anand, Dr. Victor Kasatkin, and Dr. Evgeny Mozgunov for their useful comments on the manuscript.

*pokharel@usc.edu
†lidar@usc.edu

[1] R. P. Feynman, Int. J. Theor. Phys. **21**, 467 (1982).
[2] D. Deutsch and R. Jozsa, Proc. R. Soc. A **439**, 553 (1992).
[3] E. Bernstein and U. Vazirani, SIAM J. Comput. **26**, 1411 (1997).
[4] L. K. Grover, Phys. Rev. Lett. **79**, 325 (1997).
[5] P. W. Shor, SIAM Rev. **41**, 303 (1999).
[6] A. M. Childs, R. Cleve, E. Deotto, E. Farhi, S. Gutmann, and D. A. Spielman, in *STOC '03: Proceedings of the Thirty-Fifth Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery, New York, 2003), pp. 59–68, 10.1145/780542.780552.
[7] W. Van Dam, S. Hallgren, and L. Ip, SIAM J. Comput. **36**, 763 (2006).
[8] A. W. Harrow, A. Hassidim, and S. Lloyd, Phys. Rev. Lett. **103**, 150502 (2009).
[9] S. Bravyi, D. Gosset, and R. König, Science **362**, 308 (2018).
[10] S. Bravyi, D. Gosset, R. König, and M. Tomamichel, Nat. Phys. **16**, 1040 (2020).
[11] D. Maslov, J.-S. Kim, S. Bravyi, T. J. Yoder, and S. Sheldon, Nat. Phys. **17**, 894 (2021).
[12] A. Montanaro, npj Quantum Inf. **2**, 15023 (2016).
[13] K. Bharti, A. Cervera-Lierta, T. H. Kyaw, T. Haug, S. Alperin-Lea, A. Anand, M. Degroote, H. Heimonen, J. S. Kottmann, T. Menke *et al.*, Rev. Mod. Phys. **94**, 015004 (2022).
[14] A. J. Daley, I. Bloch, C. Kokail, S. Flannigan, N. Pearson, M. Troyer, and P. Zoller, Nature (London) **607**, 667 (2022).
[15] J. Preskill, Quantum **2**, 79 (2018).
[16] C. Figgatt, D. Maslov, K. A. Landsman, N. M. Linke, S. Debnath, and C. Monroe, Nat. Commun. **8**, 1 (2017).
[17] K. Wright, K. M. Beck, S. Debnath, J. M. Amini, Y. Nam, N. Grzesiak, J.-S. Chen, N. C. Pisenti, M. Chmielewski, C. Collins *et al.*, Nat. Commun. **10**, 5464 (2019).
[18] T. Albash and D. A. Lidar, Phys. Rev. X **8**, 031016 (2018).
[19] A. D. King, J. Raymond, T. Lanting, S. V. Isakov, M. Mohseni, G. Poulin-Lamarre, S. Ejtemaee, W. Bernoudy, I. Ozfidan, A. Y. Smirnov *et al.*, Nat. Commun. **12**, 1113 (2021).
[20] F. Arute *et al.*, Nature (London) **574**, 505 (2019).
[21] T. Roy, S. Hazra, S. Kundu, M. Chand, M. P. Patankar, and R. Vijay, Phys. Rev. Appl. **14**, 014072 (2020).
[22] Y. Wu, W.-S. Bao, S. Cao, F. Chen, M.-C. Chen, X. Chen, T.-H. Chung, H. Deng, Y. Du, D. Fan *et al.*, Phys. Rev. Lett. **127**, 180501 (2021).
[23] H.-Y. Huang, M. Broughton, J. Cotler, S. Chen, J. Li, M. Mohseni, H. Neven, R. Babbush, R. Kueng, J. Preskill *et al.*, Science **376**, 1182 (2022).
[24] H.-S. Zhong, Y.-H. Deng, J. Qin, H. Wang, M.-C. Chen, L.-C. Peng, Y.-H. Luo, D. Wu, S.-Q. Gong, H. Su, Y. Hu *et al.*, Phys. Rev. Lett. **127**, 180502 (2021).
[25] V. Saggio, B. E. Asenbeck, A. Hamann, T. Strömberg, P. Schiansky, V. Dunjko, N. Friis, N. C. Harris, M. Hochberg, D. Englund *et al.*, Nature (London) **591**, 229 (2021).
[26] F. Centrone, N. Kumar, E. Diamanti, and I. Kerenidis, Nat. Commun. **12**, 850 (2021).
[27] Y. Xia, W. Li, Q. Zhuang, and Z. Zhang, Phys. Rev. X **11**, 021047 (2021).

[28] M.-G. Zhou, X.-Y. Cao, Y.-S. Lu, Y. Wang, Y. Bao, Z.-Y. Jia, Y. Fu, H.-L. Yin, and Z.-B. Chen, Research **2022**, 9798679 (2022).

[29] S. Ebadi, A. Keesling, M. Cain, T. T. Wang, H. Levine, D. Bluvstein, G. Semeghini, A. Omran, J.-G. Liu, R. Samajdar *et al.*, Science **376**, 1209 (2022).

[30] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevLett.130.210602 for further details, which includes Refs. [31–55].

[31] J. Preskill, arXiv:1203.5813.

[32] S. Aaronson and L. Chen, in *Proceedings of the 32nd Computational Complexity Conference* (Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, DEU, 2017), Vol. 79, pp. 22:1–22:67, 10.4230/LIPIcs.CCC.2017.22.

[33] A. W. Harrow and A. Montanaro, Nature (London) **549**, 203 (2017).

[34] A. P. Lund, M. J. Bremner, and T. C. Ralph, npj Quantum Inf. **3**, 15 (2017).

[35] A. Zlokapa, S. Boixo, and D. Lidar, npj Quantum Inf. **9**, 36 (2023).

[36] F. Pan, K. Chen, and P. Zhang, Phys. Rev. Lett. **129**, 090502 (2022).

[37] X. Gao and L. Duan, arXiv:1810.03176.

[38] B. Barak, C.-N. Chou, and X. Gao, in *12th Innovations in Theoretical Computer Science Conference (ITCS 2021)* (2021), Vol. 185, pp. 30:1–30:20, 10.4230/LIPIcs.ITCS.2021.30.

[39] D. Aharonov, X. Gao, Z. Landau, Y. Liu, and U. Vazirani, arXiv:2211.03999.

[40] Z. Zhu, A. J. Ochoa, and H. G. Katzgraber, Phys. Rev. Lett. **115**, 077201 (2015).

[41] S. Lee, J. Lee, H. Zhai, Y. Tong, A. M. Dalzell, A. Kumar, P. Helms, J. Gray, Z.-H. Cui, W. Liu *et al.*, Nat. Commun. **14**, 1952 (2023).

[42] A. D. King, J. Raymond, T. Lanting, R. Harris, A. Zucca, F. Altomare, A. J. Berkley, K. Boothby, S. Ejtemaee, C. Enderud *et al.*, arXiv:2207.13800.

[43] T. Lubinski, S. Johri, P. Varosy, J. Coleman, L. Zhao, J. Necaise, C. H. Baldwin, K. Mayer, and T. Proctor, arXiv:2110.03137.

[44] K. Zhang, P. Rao, K. Yu, H. Lim, and V. Korepin, Quantum Inf. Process. **20**, 233 (2021).

[45] James King, Sheir Yarkoni, Mayssam M. Nevisi, Jeremy P. Hilton, and Catherine C. McGeoch, arXiv:1508.05087.

[46] W. Vinci and D. A. Lidar, Phys. Rev. Appl. **6**, 054016 (2016).

[47] S. Mandrà, Z. Zhu, W. Wang, A. Perdomo-Ortiz, and H. G. Katzgraber, Phys. Rev. A **94**, 022337 (2016).

[48] L. Viola, S. Lloyd, and E. Knill, Phys. Rev. Lett. **83**, 4888 (1999).

[49] K. Khodjasteh and D. A. Lidar, Phys. Rev. A **78**, 012355 (2008).

[50] K. Khodjasteh and L. Viola, Phys. Rev. Lett. **102**, 080501 (2009).

[51] L. Viola, E. Knill, and S. Lloyd, Phys. Rev. Lett. **82**, 2417 (1999).

[52] P. D. Nation, H. Kang, N. Sundaresan, and J. M. Gambetta, PRX Quantum **2**, 040326 (2021).

[53] S. Bravyi, S. Sheldon, A. Kandala, D. C. Mckay, and J. M. Gambetta, Phys. Rev. A **103**, 042605 (2021).

[54] F. B. Maciejewski, Z. Zimborás, and M. Oszmaniec, Quantum **4**, 257 (2020).

[55] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven, Nat. Phys. **14**, 595 (2018).

[56] T. F. Rønnow, Z. Wang, J. Job, S. Boixo, S. V. Isakov, D. Wecker, J. M. Martinis, D. A. Lidar, and M. Troyer, Science **345**, 420 (2014).

[57] M. Naseri, T. V. Kondra, S. Goswami, M. Fellous-Asiani, and A. Streltsov, Phys. Rev. A **106**, 062429 (2022).

[58] D. Suter and G. A. Álvarez, Rev. Mod. Phys. **88**, 041001 (2016).

[59] B. Pokharel, N. Anand, B. Fortman, and D. A. Lidar, Phys. Rev. Lett. **121**, 220502 (2018).

[60] A. M. Souza, Quantum Inf. Process. **20**, 237 (2021).

[61] V. Tripathi, H. Chen, M. Khezri, K.-W. Yip, E. M. Levenson-Falk, and D. A. Lidar, Phys. Rev. Appl. **18**, 024068 (2022).

[62] Z. Zhou, R. Sitler, Y. Oda, K. Schultz, and G. Quiroz, arXiv:2208.05978.

[63] P. Jurcevic, A. Javadi-Abhari, L. S. Bishop, I. Lauer, D. F. Bogorin, M. Brink, L. Capelluto, O. Günlük, T. Itoko, N. Kanazawa *et al.*, Quantum Sci. Technol. **6**, 025020 (2021).

[64] G. S. Ravi, K. N. Smith, P. Gokhale, A. Mari, N. Earnest, A. Javadi-Abhari, and F. T. Chong, arXiv:2112.05821.

[65] N. Ezzell, B. Pokharel, L. Tewala, G. Quiroz, and D. A. Lidar, arXiv:2207.03670.

[66] G. T. Genov, D. Schraft, N. V. Vitanov, and T. Halfmann, Phys. Rev. Lett. **118**, 133202 (2017).

[67] A. Gautam, Arvind, and K. Dorai, arXiv:2112.05821.

[68] J. R. West, D. A. Lidar, B. H. Fong, and M. F. Gyure, Phys. Rev. Lett. **105**, 230503 (2010).

[69] H. K. Ng, D. A. Lidar, and J. Preskill, Phys. Rev. A **84**, 012305 (2011).

[70] B. Efron, in *Breakthroughs in Statistics: Methodology and Distribution* (Springer, New York, NY, 1992), pp. 569–593.

[71] E. Pelofske, A. Bärtschi, and S. Eidenbenz, arXiv:2112.05821.

[72] D. Gottesman, *Group22: Proceedings of the XXII International Colloquium on Group Theoretical Methods in Physics* (International Press, Cambridge, MA, 1999).

[73] E. T. Campbell, B. M. Terhal, and C. Vuillot, Nature (London) **549**, 172 (2017).

[74] K. Temme, S. Bravyi, and J. M. Gambetta, Phys. Rev. Lett. **119**, 180509 (2017).

[75] A. Zlokapa and A. Gheorghiu, arXiv:2005.10811.

[76] P. Das, S. Tannu, S. Dangwal, and M. Qureshi, in *MICRO-54: 54th Annual IEEE/ACM International Symposium on Microarchitecture* (Association for Computing Machinery, New York, 2021), pp. 950–962, 10.1145/3466752.3480059.

[77] S. Niu and A. Todri-Sanial, arXiv:2204.14251.

[78] S. Srinivasan, B. Pokharel, G. Quiroz, and B. Boots, arXiv:2210.12284.