Standardized Low-Level Noise Characterization to Differentiate DOE Quantum Testbeds

Ruslan Shaydulin* Martin Suchara* Jeffrey Larson*
*Mathematics and Computer Science Division, Argonne National Laboratory, Lemont, IL 60439, USA

As the available quantum hardware becomes increasingly capable, opportunities arise to leverage these systems for nontrivial computation. However, implementing practically relevant algorithms on quantum hardware requires at least partial algorithmic error correction or error mitigation, an observation that is reflected in the roadmaps of all major quantum hardware vendors. The development of low-cost partial error mitigation techniques is therefore crucial to the success of quantum computing in the near term before the hardware advances sufficiently to support full fault tolerance. Moreover, it is widely believed that quantum advantage in applications on most quantum hardware platforms is unlikely without at least some error mitigation [3, 8].

To be useful in practice, error mitigation techniques must not introduce high overhead, in terms of the number of additional operations, qubits, or samples required. To keep the overhead low, the developed techniques must be tailored to the noise profile of the particular hardware platform. Information about the noise present on the device can be used to reduce the overhead: for example, if bit-flip errors are exponentially suppressed as compared with phase-flip errors, simpler codes can be used to correct the errors [5]. But the development of such techniques requires precise information about the error sources of the hardware.

Unfortunately, commercial vendors do not provide detailed characterizations of the rates and sources of the noise that is present on their hardware at a sufficiently low level. While some vendors do provide high-level statistics obtained from randomized benchmarking, some vendors do not release even this limited data. This reticence to release data may be due to a perceived need to protect commercially sensitive information or concerns about accidentally exposing proprietary know-how or simply because such activities are given a low priority. The opacity of noise levels and sources on commercially available systems severely hampers open research into error mitigation techniques by researchers who are not affiliated with a given vendor. Fortunately, federally funded quantum computing testbeds present a unique opportunity to fill this gap in the research ecosystem. Specifically, the U.S. Department of Energy's open quantum testbeds are in the lone position to address the needs of the community, including characterization (of quantum systems), comparability (between different systems), and verification (of results obtained on hardware).

Characterization Transparent and actionable noise characterization information is essential to executing quantum algorithms on near-term hardware, as well as to designing error mitigation and error correction schemes. Such information allows researchers to understand hardware capability and helps design experiments that produce meaningful outcomes. However, quantum hardware vendors typically release only T_1 and T_2 relaxation times and one- and two-qubit error rates for their systems. These statistics are useful in some cases, but they neither accurately capture the capability of the device nor are sufficiently detailed to provide an understanding of the sources of noise in the system. While any desired holistic or granular statistic could in principle be extracted from the device by running appropriate circuits, in practice the corresponding burden of designing and implementing such a scheme is high, especially for researchers who are not well versed in the intricacies of the hardware. Moreover, having each researcher perform the requisite characterization is wasteful, since the same characterization data can be used in many applications. This is especially true in the cases where access is provided as a fixed number of "credits" to be used in a cloud service.

Federally funded quantum testbeds should provide a daily-updated dashboard and API with both low-level and high-level noise characterization information. Such a database containing the characterization information should be centrally hosted and contain information from all federally funded testbeds. Designing the particular characterization protocols is an active area of research [1, 6, 7, 10] and presents an opportunity for collaboration between the scientific community and testbed scientists. Beyond merely providing the statistics obtained from characterization protocols, a simulator with a noise model that mimics the testbed—a "digital twin"—can be developed by using the granular characterization information. Such a digital twin for a given testbed allows for the use of classical HPC (hosted by the user) to perform preliminary analyses of proposed techniques. This reduces the demand to run redundant circuits when submitted by independent researchers, thereby ensuring more efficient use of valuable testbed resources.

Comparability Of course, simply providing characterization data is insufficient in many cases. To be useful, the noise data must be easily interpretable. For example, a compilation technique that uses device noise properties to decide what error mitigation strategy to insert into a given circuit would benefit from a standardized characterization data API. Moreover, detailed and standardized characterization data allows researchers to understand the advantages and disadvantages of each hardware platform, as well as to investigate the different sources of noise in them. While great advances have been made in designing standardized holistic performance metrics such as quantum volume (QV) [2] and CLOPS [9], no such metrics are available for comparing the *sources* of noise on different hardware platforms. Breaking down the error by sources of noise (e.g., leakage or cross-talk) in a standardized way will allow insights into the comparative advantages and disadvantages of different hardware platforms. As an analogy, quantum volume is similar to the overall profits of a company: a single number to describe its overall capability. Breaking down the error by components is analogous to separately describing revenues, losses due to depreciation, investment, and operating expenses. Understanding the different sources of noise will enable a more nuanced understanding of the potential of a given system as compared with just considering a single number such as QV. Similar to the use of generally accepted accounting principles in financial statements, standardized characterization of testbed error components will allow for a more fine-grained comparison of different platforms and error budgeting when implementing algorithms in a portable way. Proprietary versions of such characterization protocols have been shown to provide crucial insight into device behavior and helped drive improvements in hardware performance [4].

Verification In addition to comparing different testbeds, a database of historical characterization data will allow verification of results obtained on the testbeds beyond viewing the quantum device as a "black box". Concretely, the historical parameters corresponding to the time when a given experiment was performed can be used to approximately simulate (using a digital twin or simpler model) the effect of the noise on the experiment and validate the accuracy of reported results. As the number and quality of the qubits grow beyond the classically simulatable regime, a coarse-grained noise profile can be used to estimate the expected quality of a result and provide sanity checks for too-good-to-be-true results or unexpectedly poor performance.

A Scientific Asset Realizing the outlined goals will require overcoming several challenges. First, certain aspects of the quantum testbeds will have to be prioritized. Specifically, to enable daily execution of the characterization pipeline, the testbeds must support a sufficiently high throughput rate (i.e., they must execute circuits quickly enough). Second, collaborations between external researchers and testbed scientists will be needed to develop scalable yet fine-grained characterization techniques, and arrive at the consensus required for standardization. We are confident that if the proposed capabilities are implemented, they will be a great asset to the scientific community and enable the development of error mitigation techniques with broad applicability, including to commercial quantum computing systems.

References

[1] A. Cornelissen, J. Bausch, and A. Gilyén, *Scalable benchmarks for gate-based quantum computers*, arXiv:2104.10698, (2021).

[2] A. W. Cross, L. S. Bishop, S. Sheldon, P. D. Nation, and J. M. Gambetta, *Validating quantum computers using randomized model circuits*, Physical Review A, 100 (2019).

[3] D. S. França and R. García-Patrón, *Limitations of optimization algorithms on noisy quantum devices*, Nature Physics, 17 (2021), pp. 1221–1227.

[4] Google Quantum AI, *Exponential suppression of bit or phase errors with cyclic error correction*, Nature, 595 (2021), pp. 383–387.

[5] J. Guillaud and M. Mirrahimi, *Error rates and resource overheads of repetition cat qubits*, Physical Review A, 103 (2021).

[6] A. J. McCaskey, Z. P. Parks, J. Jakowski, S. V. Moore, T. D. Morris, T. S. Humble, and R. C. Pooser, *Quantum chemistry as a benchmark for near-term quantum computers*, npj Quantum Information, 5 (2019).

[7] T. Proctor, K. Rudinger, K. Young, E. Nielsen, and R. Blume-Kohout, *Measuring the capabilities of quantum computers*, arXiv:2008.11294, (2020).

[8] R. Shaydulin and A. Galda, *Error mitigation for deep quantum optimization circuits by leveraging problem symmetries*, in IEEE International Conference on Quantum Computing and Engineering (QCE), IEEE, Oct. 2021.

[9] A. Wack, H. Paik, A. Javadi-Abhari, P. Jurcevic, I. Faro, J. M. Gambetta, and B. R. Johnson, *Quality, speed, and scale: three key attributes to measure the performance of near-term quantum computers*, arXiv:2110.14108, (2021).

[10] K. Yeter-Aydeniz, E. F. Dumitrescu, A. J. McCaskey, R. S. Bennink, R. C. Pooser, and G. Siopsis, *Scalar quantum field theories as a benchmark for near-term quantum computers*, Physical Review A, 99 (2019).