

Quantum Error Mitigation

External Public

June 2025

Dell Technologies team

Miguel Paredes Quiñones

Rômulo Pinho

Micael Araujo

Smriti Bajaj

João Pinto

IQM team

Alessio Calzona

Amin Hosseinkhani

Adrian Auer

Table of Contents

Introduction..... 2

 Sources of errors in quantum computers 2

 Quantum error mitigation and the path to quantum advantage..... 3

 Challenges of this technology 3

Using AI to mitigate errors in quantum machines 4

 Generate synthetic data for AI training..... 4

 How is it possible to correct quantum computing’s output using AI? 5

 What other problems in quantum computing can be solved using AI? 9

Actions to Take 10

 Long term usage of error mitigation 10

 How Dell & IQM are working together on hybrid quantum-classical computing 10

 Further Reading..... 10

 Contact us 10

 References 11

The information in this publication is provided “as is.” Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

This document may contain certain words that are not consistent with Dell's current language guidelines. Dell plans to update the document over subsequent future releases to revise these words accordingly.

This document may contain language from third party content that is not under Dell's control and is not consistent with Dell's current guidelines for Dell's own content. When such third party content is updated by the relevant third parties, this document will be revised accordingly.

Copyright © 2025 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. [6/2/2025] [OFFICE OF THE CTO]

Introduction

Sources of errors in quantum computers

The execution of a quantum circuit is typically comprised of three stages: (i) initialization of the qubit register, (ii) application of single- and multi-qubit quantum gates, and (iii) the measurement of the qubits. The accurate and reliable execution of quantum algorithms relies on the precise implementation of these fundamental quantum operations. However, these operations are inherently susceptible to errors due to imperfect control and unwanted interaction with the environment. The accumulation of these errors leads to biased and incorrect results, compromising the fidelity of the computation. While the physical origins of these errors vary across hardware platforms used for realizing quantum processing units (QPUs), their effects on basic quantum operations can be characterized more generally.

One important source of errors that affects quantum gates, the fundamental building block of quantum circuits, is the imperfect control over the physical systems that make up the QPU. For example, in superconducting quantum processors, as developed by IQM, microwave and magnetic flux pulses are used to control the qubits and their interaction.¹ Imperfect calibration of these pulses leads to coherent errors in gate implementation. For example, an unintended extra rotation can lead to the implemented unitary operation deviating slightly from the intended one.

Other sources of errors that affect gate operations include leakage and crosstalk. Leakage errors are inherent in many physical realizations of qubits, which often utilize a two-dimensional subspace of a multi-dimensional quantum system. In this case, leakage results in a transition out of this intended computational subspace to other states (for example, due to insufficiently calibrated control pulses), thereby altering the calculation results. Crosstalk, on the other hand, is particularly relevant for larger QPUs where qubits interact with many neighboring qubits. The primary concern here is unwanted interaction processes that can lead to the execution of unintended operations on one or more neighboring qubits.

In general, even when qubits are idle, quantum systems constantly interact with their environment in uncontrolled ways, leading to decoherence and information loss. This can occur through energy transfer between a qubit and its surroundings causing a relaxation process that brings the qubit to its ground state. Additionally, stochastic fluctuations in the qubit frequency, driven by the environment, can lead to a loss of phase coherence, even without triggering transitions between energy levels.

Errors that affect the initialization and measurement processes are typically referred to as State Preparation and Measurement (SPAM) errors. Qubit initialization can be compromised by finite-temperature effects, leading to unwanted thermal population of higher-energy states in the quantum system. Measurement errors occur when the two computational states of a qubit cannot be fully distinguished during the readout process, resulting in incorrect state assignments. This can happen, for example, due to an unwanted relaxation process during the readout time. The relative impact of gate errors versus SPAM errors depends on the specific QPU and the quantum circuit being used. For shallow circuits that require measuring multiple qubits, SPAM errors are likely to be significant. Additionally, for deep circuits, the cumulative effects of gate errors within the circuit are likely to be more critical.

Quantum error mitigation and the path to quantum advantage

Quantum advantage refers to the point when quantum computers can solve real world problems faster or more efficiently than classical computers. However, as explained in the previous section, quantum systems are highly sensitive to their environment. Even slight interactions can cause loss of information, leading to errors in quantum computations.ⁱⁱ This makes maintaining the integrity of quantum information over time very difficult. Quantum advantage thus depends on achieving fault-tolerant quantum computing, but it remains uncertain when fault-tolerant computing will be available.

Quantum error mitigation (QEM) has been hailed as the path to “quantum utility,” when quantum computing may be able to show advantages over classical even under the effects of noise. Mitigation is typically a classical post-processing step on the readouts of quantum circuits. AI-driven error mitigation has proven to be an important tool to help achieve the quantum utility milestone.ⁱⁱⁱ AI applications’ main task is to identify patterns that generalize well to unseen conditions. As a result, AI-driven error mitigation can streamline processes, making it easier to manage and correct errors without continual human intervention.

Challenges of this technology

Error mitigation in quantum computing is a crucial area of research, but it comes with several challenges:

- **Resource demands:** Error mitigation often involves statistical inference, which becomes increasingly difficult as the system size grows. For larger quantum systems, an extremely large number of measurements are required to accurately estimate results, making effective noise mitigation infeasible.ⁱⁱ
- **Variable landscape:** The way noise affects quantum algorithms depends on the quantum technology/configuration in place. Mitigation methods need to be tailored to the specific type of noise present, which adds complexity to the process.^{iv}
- **Classical scalability:** As quantum systems scale up, the classical methods used for error mitigation may also need to scale. This can lead to exponential increases in the classical computation overhead for error mitigation.^{iv}

Some of these challenges will limit the usage of error mitigation when the size of quantum systems grows. Consequently, we believe that error mitigation is a transitory technology that will enable the first use cases for showing the potential of quantum computing and will later be complemented by quantum error correction (QEC) to continue improving quantum computing results.

Using AI to mitigate errors in quantum machines

Generate synthetic data for AI training

The Dell Technologies team utilized a framework to generate synthetic data for training our AI models. This involved designing diverse circuit types and simulating their execution, both with and without error considerations. To optimize resource utilization within Dell's classical infrastructure, the team leveraged CPUs and GPUs, deploying multiple instances of IQM's Garnet simulators in parallel. Approximately 32,000 circuits with different sizes and structures were executed, aiming to create a heterogeneous dataset representative of a wide range of possible circuit configurations that are most common in quantum computing algorithms.

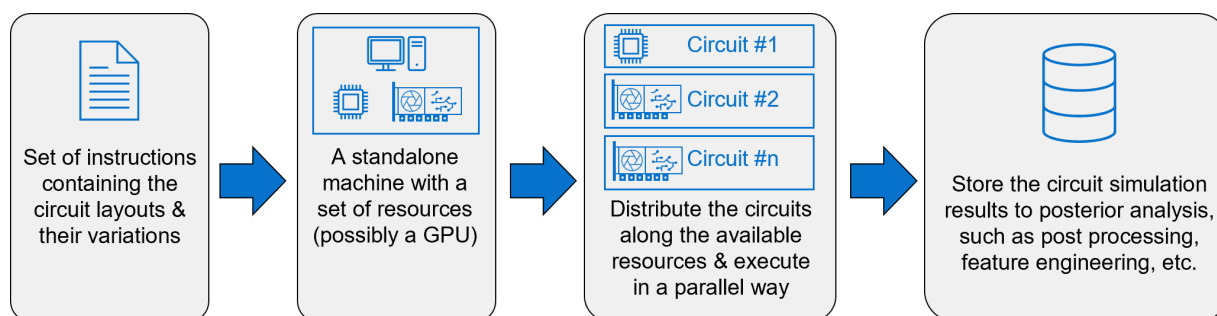


Figure 1 Framework to efficiently generate synthetic data for error mitigation models.

This framework can be reused in other quantum computing projects that need large amounts of data generated from real or simulated quantum computers. Additionally, we needed efficient methods to encode circuits into a more AI readable format. This generation of data can be done using both classical and quantum resources, making it a task-intensive and costly process compared to other AI-based steps executed for error mitigation, but fortunately, it only needs to be done once for training.

How is it possible to correct quantum computing's output using AI?

The Dell team developed an AI-based error mitigation method to predict the noise-free version of the noisy state vector measured as the output of a given quantum circuit. A high-level view of the model used for the method is presented in Figure 2. It uses information from the quantum computing hardware (IQM's Garnet, in this case) and the circuit, combining such information with the state vector measured as the output of executing the circuit n times (or shots) on the quantum hardware laid out as a probability distribution. The estimated noise-free state vector is also presented in the form of a probability distribution.

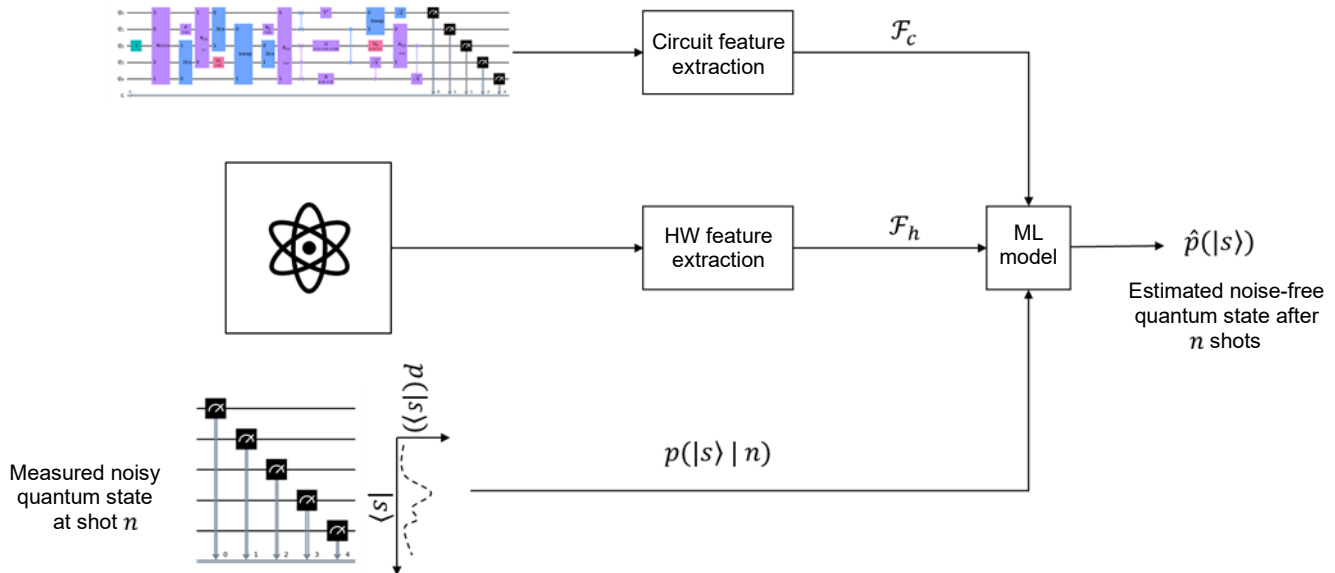


Figure 2 Architecture of the error mitigation ML model using inputs from the circuit and the empirical probabilities from Garnet that will be corrected by output error mitigated probabilities.

We incorporated key methods such as a training step, validation step, forward pass, and optimizer configuration, ensuring a robust approach to model development. To train the model, we utilized 80% of the synthetic data, reserving the remaining 20% for validation to monitor performance effectively. A notable highlight is the seamless integration of our ML models with PyTorch Lightning – a PyTorch framework that enabled us to implement advanced parallelization techniques. By combining the Distributed Data Parallel (DDP) technique with model parallelism, we achieved a transformative upgrade in training efficiency. This optimization significantly boosted training efficiency, achieving a 5x speedup that reduced training time from 3-4 days to just 7-8 hours. In addition to the initial upgrades, we implemented further modifications to enhance the training process, including the early stopping and checkpointing techniques. That enabled the model to terminate training when optimal performance was achieved and to periodically save model weights, accordingly. Additionally, we employed regularization techniques, which helped guide the model towards more optimal configurations considering the lower fidelity values.

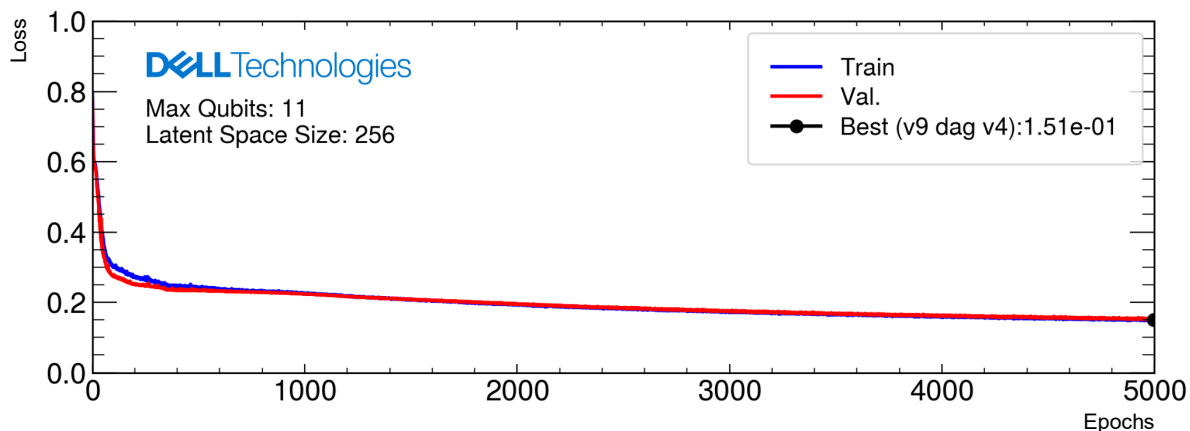


Figure 3 Model performance without regularization. The loss is 0.151 achieved within 5000+ epochs.

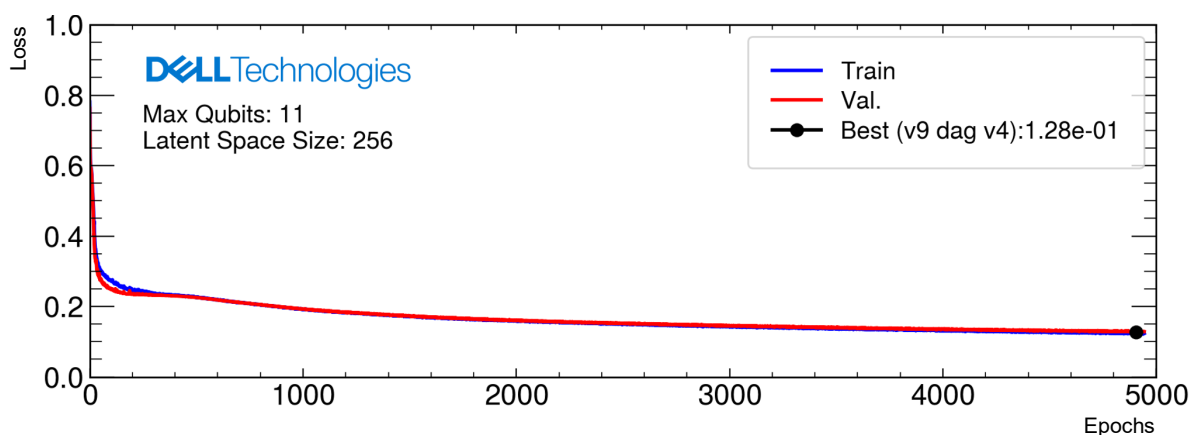


Figure 4 Model performance with L2 regularization. The loss is 0.128 achieved within approx. 4900 epochs

We used data coming from IQM's Garnet to test our AI-based method for error mitigation.ⁱ The data consists of 1,485 random circuits with the number of qubits ranging from 9 to 11 and depth from 4 to 512, run on Garnet and an ideal simulator. Garnet's capacity is up to 20 qubits, but we chose to use up to 11 to have more control over the hardware's intrinsic fidelity and more robustly assess the performance of the AI-based error mitigation. Each circuit was then executed 1,000 times on Garnet to yield a noisy state vector measured as output. Each circuit was also executed on an ideal simulator to obtain a ground truth of the noise-free state vector, which was then used to train the AI model described above.

Figure 5 depicts the results of applying the model on the validation data. The quality measure employed was the *fidelity* between the estimated noise-free state vector and its ground truth obtained with the ideal simulator, shown on the vertical axis. The fidelity is a measure of similarity between quantum states, quantified in the range [0, 1]. On the x axis, we show the fidelity between the state vectors measured from Garnet and their ground-truths using the ideal simulator. The desired result is that the AI-based error mitigation approach increases the fidelity value relative to the fidelity observed with Garnet. In other words, the points on the scatter plot should move to the left of the diagonal. Indeed, this is what occurred for all qubit counts. It should be noted that Garnet's fidelity (like any quantum hardware) decreases with more qubits. Still, the AI-based method consistently increased fidelity even in the most difficult cases.




Figure 5 Scatter plot of the relation between fidelity resulted from the ML model and the fidelity of the noisy probability grouped by number of qubits.

Figure 6 breaks down the results into a visualization of the fidelity as a function of the number of operations (layers, or depth) of the circuits. The effects of noise tend to accumulate with more operations on the circuits, reducing the fidelity. Indeed, Garnet results (in blue) show that the mean fidelity relative to the ideal state vectors monotonically decrease as more operations are added to the circuits. Although a small decrease can also be observed with the predictions of the model (in yellow), the noise-free estimates consistently improve fidelity across all circuit depths present in the validation dataset. Overall, this set of experiments demonstrates that using AI for error mitigation is feasible, although comparisons should be made with other state-of-the-art techniques.

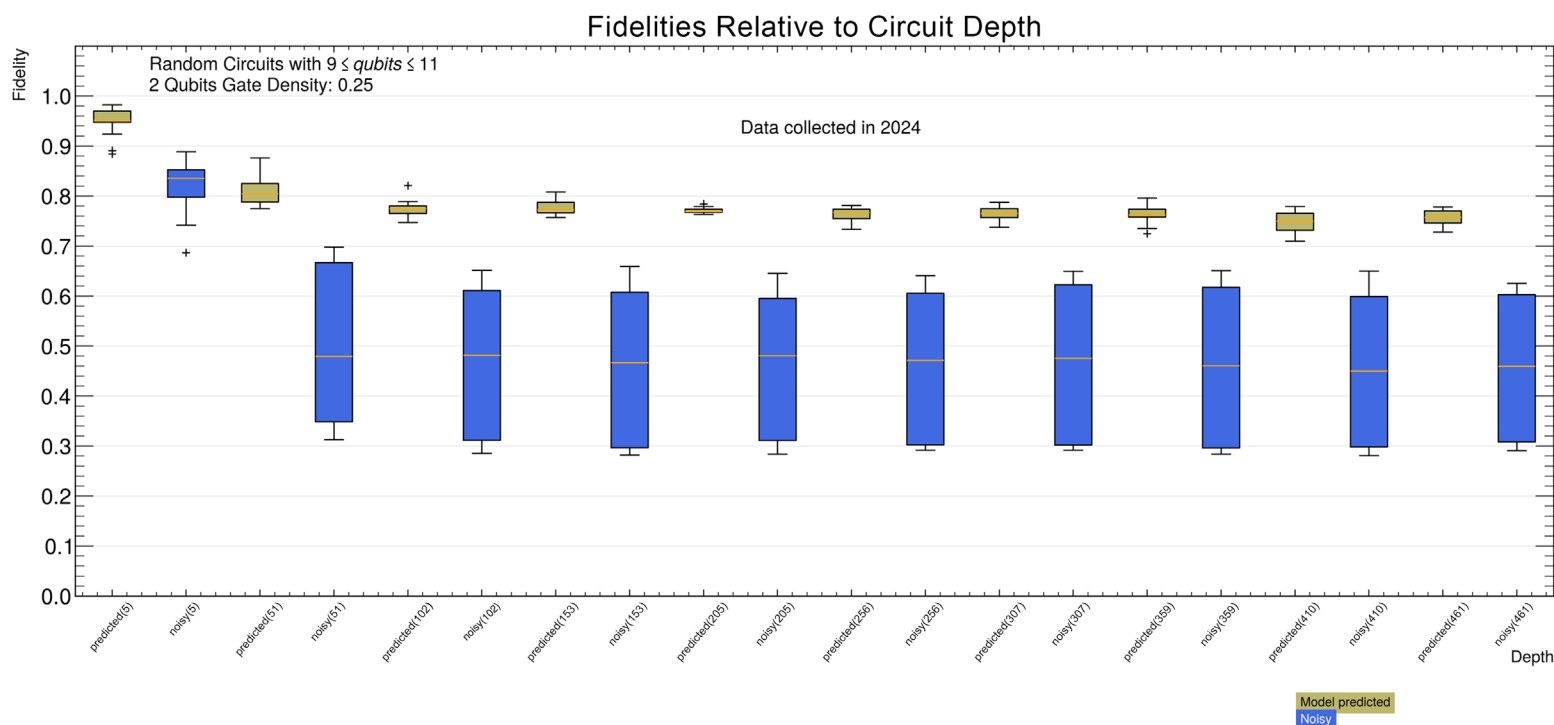


Figure 6 Box plots of fidelity distribution grouped by the number of layers in a circuit for the prediction of the model (yellow) and the noisy output of Garnet (blue).

We note that researchers at IQM recently introduced Noise-Robust Estimation (NRE) [Hosseinkhani et al., 2025], a novel, noise-agnostic, and general-purpose error mitigation method. NRE uncovers a robust bias-dispersion correlation in noisy quantum data and applies bias-aware regression to recover ideal values of quantum observables. While evaluated on a different set of problems than those analyzed in this study, NRE demonstrated significantly better performance than many traditional error mitigation techniques on the same Garnet QPU. Furthermore, NRE can potentially be enhanced through AI-based strategies, offering a promising direction for hybrid approaches.

What other problems in quantum computing can be solved using AI?

Classical AI can significantly complement the execution of quantum computing workloads in several ways, addressing complex problems and improving efficiency. Drawing upon the experience the Dell team has in AI and hybrid quantum-classical computing, we predict that AI will also positively impact other key areas:

- **Quantum simulation:** AI can help simulate quantum systems more efficiently. Machine learning algorithms can recognize patterns in data that traditional methods might miss, aiding in reproducing the particularity of some quantum systems with very non-linear behavior.
- **Orchestration problems:** Quantum computing excels at solving optimization problems and orchestration has an intrinsic optimization problem to be solved. AI algorithms can help identify optimal dispatch solutions faster and more accurately, which is beneficial in fields like logistics, finance, and supply chain management.
- **Error correction:** AI can improve error correction techniques by learning patterns of the recurrence and correction of errors. One of the challenges of this topic is the scalability of the techniques depending on the size of the quantum systems. AI models can approximate error correction techniques that scale better than the current ones. Another issue when implementing AI-based error correction is the latency between certain quantum computer modalities and classical computers; their short coherence times make it difficult to execute classical AI and apply error correction sequentially.
- **Quantum algorithm development:** AI, especially GenAI, can assist in the development of new quantum algorithms by exploring vast search spaces and identifying promising algorithmic approaches that might be too complex for human researchers to find on their own.

These applications demonstrate the powerful synergy between AI and quantum computing that will predictably drive further innovation.

Actions to Take

Long term usage of error mitigation

In the long term, the goal of quantum computation is to develop fully scalable, fault-tolerant systems that can reduce logical error rates to arbitrarily low levels. However, before reaching this stage, there will likely be an interim period where quantum systems can implement QEC but still produce results deeply affected by non-negligible logical error rates. For instance, residual logical errors can arise from constraints in system size, which limit achievable code distances, insufficient resources for magic state distillation, and/or challenges in large-scale decoding. Additionally, a key role will also be played by algorithmic errors, stemming from the need to implement arbitrary unitary operations using only a small set of discrete, fault-tolerant operations while maintaining a fixed total circuit depth to prevent an exponential increase in the logical circuit error rate.

In this initial scenario, also referred to as early fault-tolerant quantum computing, it is crucial to develop strategies to mitigate the effects of these residual errors on the desired output. Fortunately, the quantum error mitigation framework, although originally developed for NISQ (Noisy Intermediate-Scale Quantum) hardware, can also be applied at the logical level.ⁱⁱ This means that post-processing the noisy output of logical circuits can deliver results with improved precision, albeit at the cost of increased runtime. QEM thus emerges as a promising way to compensate for the scarcity of quantum resources during the upcoming early fault-tolerant era, while still providing high-quality results.

Importantly, this interplay between QEM and QEC remains an active area of research (e.g., with open questions concerning the types of algorithms implemented during this early fault-tolerant era). While QEM is typically effective at mitigating the expectation values of observables, its implementation on sampling algorithms may face additional challenges, especially with scalability. The synergic development of capable and flexible QEM techniques, efficient early fault-tolerant platforms and quantum algorithms will be key to unlocking more useful applications of quantum computation on the path toward fully scalable, fault-tolerant systems where logical errors can be suppressed to arbitrary levels.

How Dell & IQM are working together on hybrid quantum-classical computing

Dell and IQM have been partnering on the advancement of quantum error mitigation (QEM) techniques that are especially applicable to IQM's quantum computing technologies. While Dell has been leading the development of new AI-driven QEM algorithms, IQM provides their advanced quantum hardware for experimentation and additional technical support for its usage. Together, experts from both companies analyze the results and discuss the possible improvements, both on the classical software part as well as on the quantum hardware. Through this collaboration, Dell and IQM exercise many of the potentials of hybrid quantum-classical computation.

Further Reading

Dell Technologies Quantum Computing: dell.com/quantum-computing

Contact us

Dell Technologies team: quantum.assist@dell.com

IQM team: info@meetiqm.com

References

- ⁱ L. Abdurakhimov *et al.*, “Technology and Performance Benchmarks of IQM’s 20-Qubit Quantum Computer,” Aug. 22, 2024, *arXiv*: arXiv:2408.12433. doi: 10.48550/arXiv.2408.12433.
- ⁱⁱ M. Swayne, “Quantum Error Mitigation May Face Hard Limits,” The Quantum Insider. Accessed: Jan. 06, 2025. [Online]. Available: <https://thequantuminsider.com/2024/07/26/quantum-error-mitigation-may-face-hard-limits/>
- ⁱⁱⁱ H. Liao, D. S. Wang, I. Sitdikov, C. Salcedo, A. Seif, and Z. K. Minev, “Machine learning for practical quantum error mitigation,” *Nat Mach Intell*, vol. 6, no. 12, pp. 1478–1486, Dec. 2024, doi: 10.1038/s42256-024-00927-2.
- ^{iv} Z. Cai *et al.*, “Quantum Error Mitigation,” *Rev. Mod. Phys.*, vol. 95, no. 4, p. 045005, Dec. 2023, doi: 10.1103/RevModPhys.95.045005.