Project title: Optimising algorithm performance on NISQ computers using machine learning

Supervisors: Stefan Scheel (German PI), Barry Sanders (Canadian PI), Christoph Simon (Canadian PI)

Background

Current state-of-the-art quantum computing processors typically employ between 50-250 qubits. This is far below what is needed for fault-tolerant error correction on the scale required for a useful quantum advantage; even with a gate error of only $\sim 0.01\%$, factoring a 1000-bit number using Shor's algorithm would still involve millions of physical qubits [1]. Instead, recent emphasis has been placed on finding algorithms which can run in the so-called noisy intermediate-scale quantum (NISQ) regime [2].

Typically, NISQ algorithms work by moving part of the computation to a classical computer, while leaving a reduced-depth quantum circuit to carry out the most demanding tasks. Early applications primarily focused on quantum simulations (most notably, computational quantum chemistry [3, 4, 5]), approximate optimization algorithms [6], and then later quantum machine learning (QML). A review of the current state of the art can be found in [7]. Perhaps the best known example of these is the variational quantum eigensolver (VQE) [3], used for finding ground-state energies of interacting quantum systems. This involves minimising a cost function, defined as $\langle H\rangle$, over a parameterised set of states using a classical optimiser. As the expectation values can be computed efficiently on a quantum computer, then for suitable parameterisations of the trial states, the VQE can iteratively find the ground state using a polynomial number of operations.

For NISQ systems, where the error rates are higher than the threshold for error correction, the priority is instead to focus on quantum error mitigation (QEM) rather than full-scale fault tolerance. Here, a number of techniques for short-depth circuits have been put forward, e.g. by artificially increasing noise in a sequence of runs and then extrapolating to the zero-noise limit (ZNE), or by decomposing an ideal circuit in a noisy basis [8, 9]. In particular, QEM should 1) not require too many additional quantum resources, and 2) not increase the circuit depth beyond the coherence time, i.e. inasmuch as possible, it should avoid any sources of unwanted noise. However, in most cases the circuit depth is not actively decreased, and there has been limited work towards developing general approaches in this direction.

Recently, we have demonstrated that reinforcement learning algorithms can find low gate circuit approximations for arbitrary state generation. In real NISQ computers, these turn out to perform better, on average, than exact circuits as the accuracy improvement of the ML approximation (as a function of circuit depth) quickly exceeds the typical gate error. Similar trade-offs also exist at the level of individual gate operations, where accepting a small, known, error can sometimes offset higher levels of unwanted noise. This suggests that ML-optimised circuits can provide a promising avenue for quantum error mitigation in NISQ computers.

Research goals and working programme

The goal of this project is to explore the use of hybrid (quantum-classical) machine learning approaches in real quantum processors. In particular, we aim

1. to investigate the performance limits which can be achieved, in principle, for both general and NISQ algorithms using approximate circuits on small qubit systems, and
2. to develop scalable approaches to error mitigation using hybrid machine learning.

In the first stage of the project, we will extend our existing algorithm (Fig. 1) to include a comprehensive noise model of a real-world system. A prime example would be the cloud-based quantum computers provided by IBM or Xanadu. Our goal is to investigate how well one can optimise the state-generation task by algorithm choice alone. Similar techniques can then be applied to the approximation of general circuits (i.e. of arbitrary unitary operators). Already, we have found evidence that some hardware restrictions (e.g. system architectures) can be better dealt with using this holistic approach rather than the available circuit transpilers. However, it seems clear that the noise characterisations typically published by system providers are insufficient to fully model performance. Thus, additional levels of sophistication may be required. Once a baseline has been achieved, we will then try to further limit noise through a variety of QEM techniques. Of particular interest is whether new approaches can be learnt using classical neural networks.

For small systems, both circuit approximation and quantum error mitigation can be modelled entirely using classical resources. In the second stage, we aim to find scalable implementations of the algorithm. The most obvious approach is to use some form of QML (in the hybrid sense). However, if noise mitigation is learnt on single qubits (or at least small subsets of qubits at a time), then these should be already inherently scalable, and one can directly apply them to existing NISQ algorithms.

The project supervisor in Rostock, Prof. Stefan Scheel, has a long-standing expertise in theoretical quantum optics and quantum information. The Canadian partner, Prof. Barry Sanders in Calgary, complements that with his expertise in classical and quantum machine learning. While the project is anchored in Rostock, the successful applicant is expected to spend between 6-12 months at the University of Calgary.

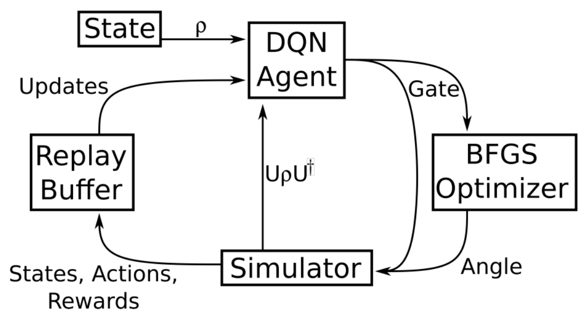

Figure 1: An example reinforcement algorithm approach to arbitrary state generation. Here a DQN agent is trained to choose a gate type (e.g. CNOT, or rotation), while a local optimiser finds the ideal parameter value.

[1] A. Fruchtman and I. Choi, *Technical roadmap for fault-tolerant quantum computing*, University of Oxford (2016).

[2] J. Preskill, *Quantum Computing in the NISQ era and beyond*, Quantum **2**, 79 (2018).

[3] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, *A variational eigenvalue solver on a photonic quantum processor*, Nature Communications **5**, 4213 (2014).

[4] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, *The theory of variational hybrid quantum-classical algorithms*, New J. Phys. **18**, 023023 (2016).

[5] D. Wecker, M. B. Hastings, and M. Troyer, *Progress towards practical quantum variational algorithms*, Phys. Rev. A **92**, 042303 (2015).

[6] E. Farhi, J. Goldstone, and S. Gutmann, *A quantum approximate optimization algorithm*, arXiv:1411.4028 (2014).

[7] 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. Kwek, and A. Aspuru-Guzik, *Noisy intermediate-scale quantum algorithms*, Rev. Mod. Phys. **94**, 015004 (2022).

[8] K. Temme, S. Bravyi, and J. M. Gambetta, *Error mitigation for short-depth quantum circuits*, Phys. Rev. Lett. **119**, 180509 (2017).

[9] Y. Li and S. C. Benjamin, *Efficient variational quantum simulator incorporating active error minimization*, Phys. Rev. X **7**, 021050 (2017).