Reinforcement learning for quantum compilation

Supervisor(s):
Roberto Bondesan (Department of Computing, r.bondesan@imperial.ac.uk)

Project description:
Quantum compilation is the process of implementing a high-level quantum algorithm on hardware. It is a challenging optimisation problem that needs to be adapted to different hardware characteristics. By posing the problem as a reinforcement learning problem, we aim to develop an agent that can compile a quantum circuit taking into account the noise present in quantum computers. The AI challenge is that observation and action spaces of the agent are high dimensional and characterising the noise in a quantum computer is itself a challenging problem. We shall tackle the problem of parametrising the agent using graph neural networks adapted to model quantum circuits and develop a machine learning model for the noise to obtain compiled circuits that can be deployed on real use cases.

Timeline (tentative):
Oct-Dec: background reading and development of idea.
Dec-March: experiment with synthetic data.
March-June: deployment and benchmarking.
July/Aug: thesis writing.

Minimum viable thesis:
Implementing previous literature on the problem (refs [1,2] below) add graph neural network part and benchmark using existing libraries (ref [3] below).

Required background & skills:
Maths, programming, basics of quantum computing or quantum physics.

Representative References:
[1] https://arxiv.org/abs/2103.07585
[2] https://www.nature.com/articles/s42005-021-00684-3
[3] https://arxiv.org/abs/2308.02536

Additional notes:
To benchmark progress we will use the data publicly available under Apache 2.0 license at https://github.com/Quantomatic/pyzx/tree/master/circuits