Using Reinforcement Learning to Optimize Quantum Circuits in the Presence of Noise

FERMILAB-POSTER-20-105-QIS

Khalil Guy, Fisk University; Gabriel Purdue, Fermi Quantum Institute

Motivation and Objective

As we move towards devices which utilize more qubits, it becomes increasing more important to map quantum circuits in a way that uses resources efficiently as well as maximizes the reliability of the results of that circuit. To this end, we will need to rely on heuristic algorithms, specifically reinforcement learning (RL) as a method of building quantum circuits based on observations of the noise characteristics in its environment.

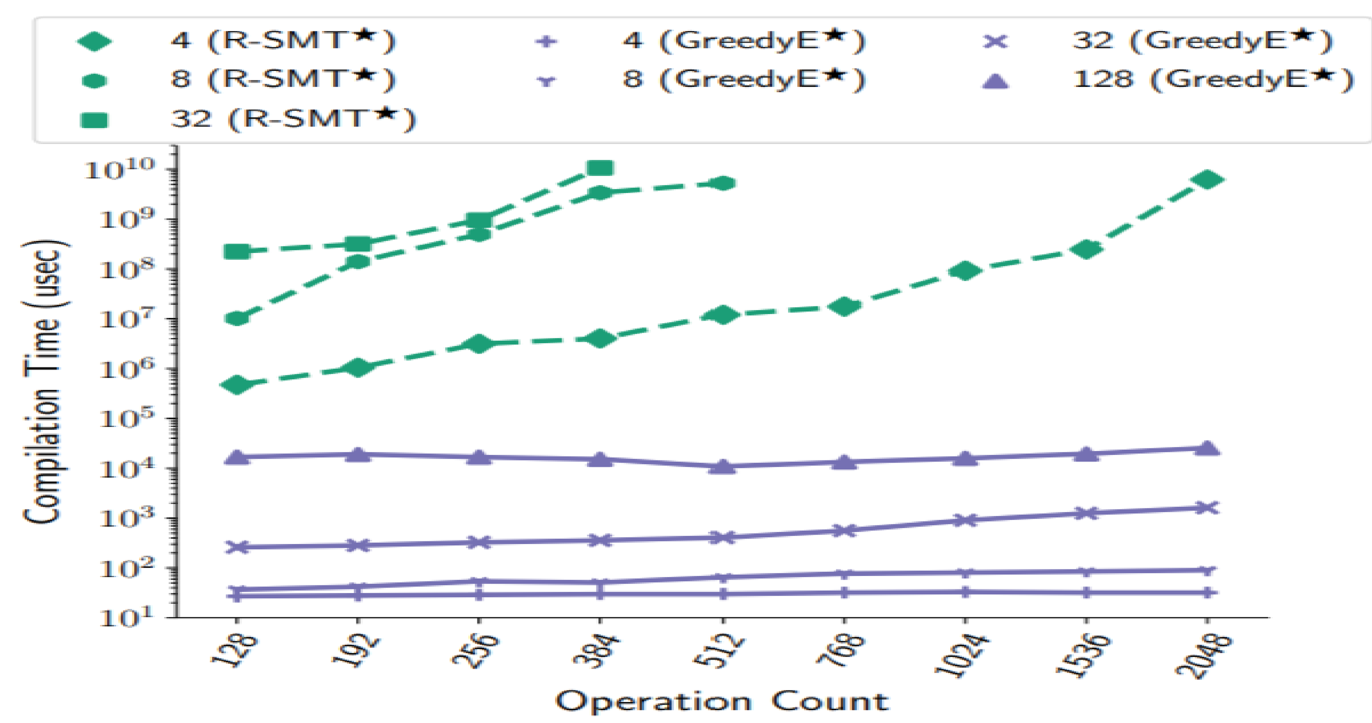

Fig 1: This figure shows how compilation time changes for different algorithms as the circuit depth increases (Murali, et el., 2019)

Overview: Quantum Circuits

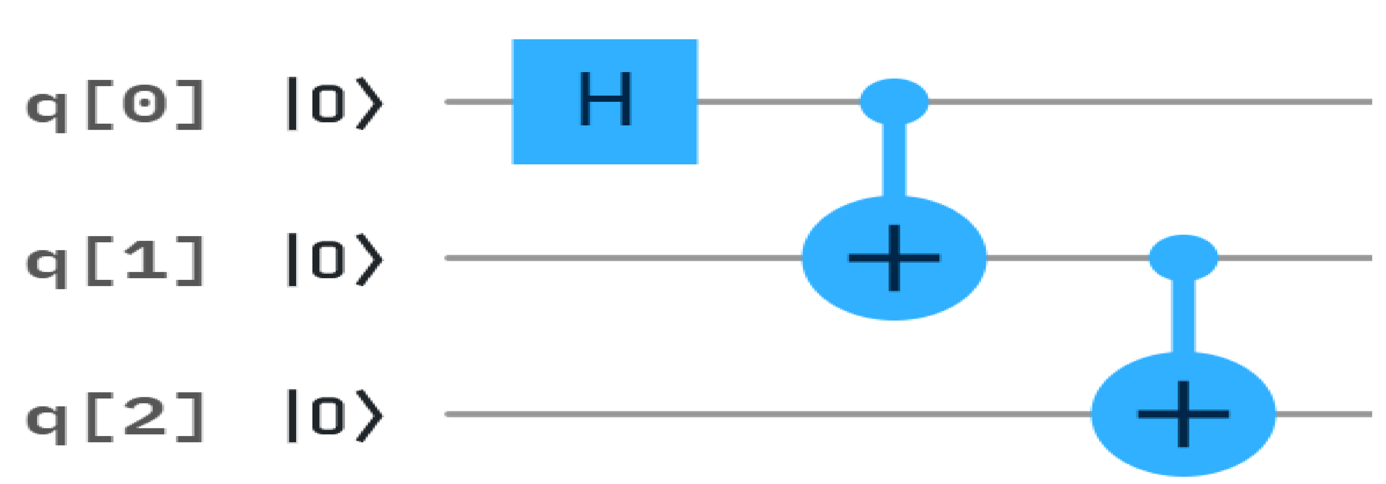

Fig 2.: A virtual representation of a quantum circuit implemented in Qiskit. The figure shows the circuit creates a three qubit GHZ state.

Acknowledgements

This manuscript has been authored by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics.

Reinforcement Learning

Reinforcement learning is a type of machine learning which uses an agent to choose from a certain set of actions based on observations from an environment to complete a task or maximize some reward.

Description of Reinforcement Learning Environment

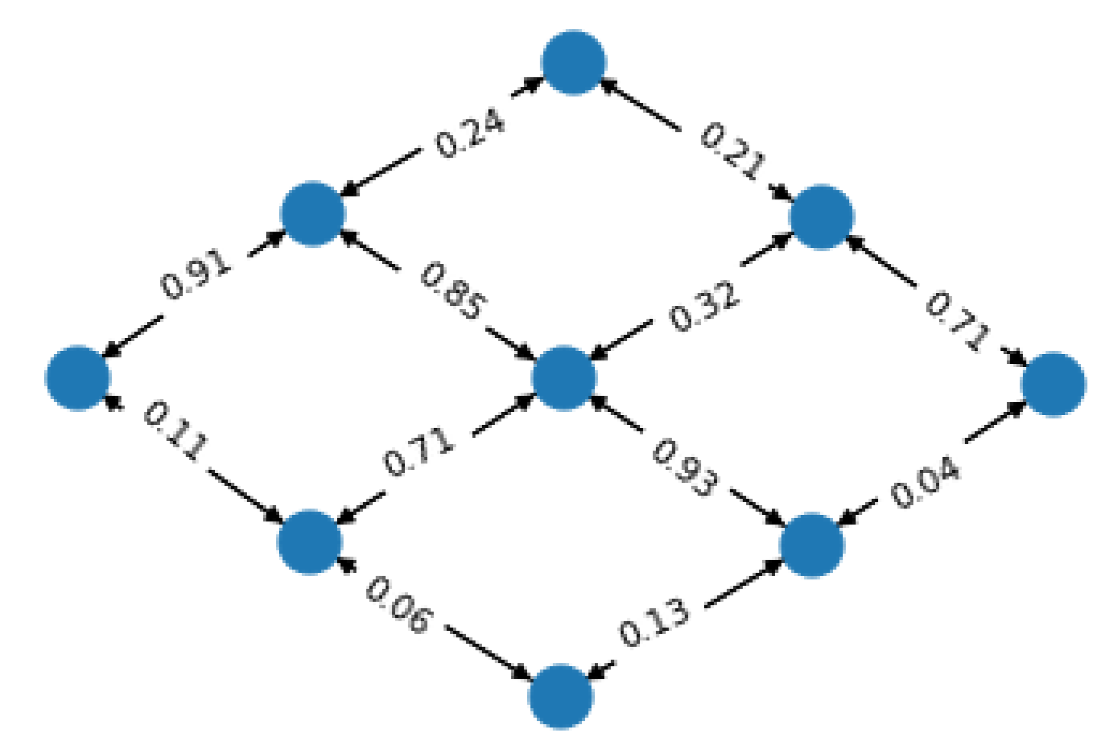

Fig 2.: A representation of the quantum hardware where each node represents a qubit and each edge represents and available connection to apply a two-qubit gate operation. The edge value represents the amount of noise along each edge.

The environment is designed with a specific class of circuits in mind, specifically the circuit which creates the GHZ state. This circuit takes a strictly line topology on a graph, making it relatively simple target for an RL agent.

We use a noisy unitary matrix to simulate coherent noise in our quantum hardware. Initially, we start with a matrix $\widehat{N}$ such that

$$\widehat{N} = aU + bT$$

Where U and T are two-qubit gate transformations and a and b are constants. $\widehat{N}$ is transformed into a unitary matrix using the Grahm-Shmidt process.

Using Random Tree Search to Determine Intermediate Rewards

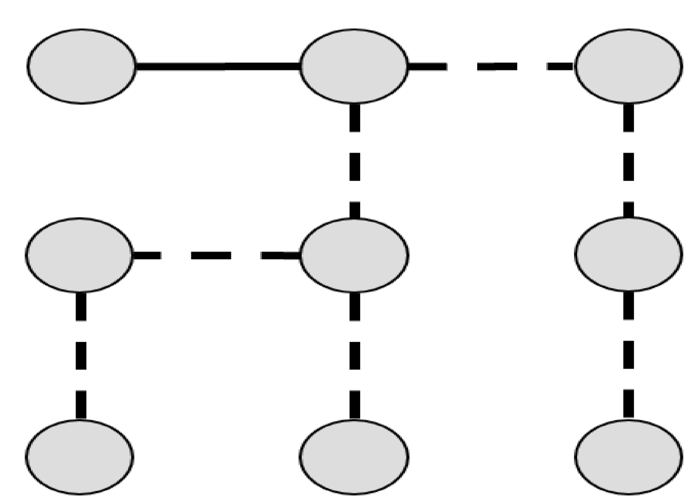

Fig 3: A representation of a random tree search on the quantum hardware. At each time step, we need to have some measure of goodness of the potential completed circuits that can be created from the agent's current position in the hardware at it current stage in completing the desired circuit.

Results and Future Works

From this study, we build a functional reinforcement learning environment in which a reinforcement learning agent can perform actions and receive rewards. Going forward, we want to train a reinforcement learning agent in the environment with the following questions in mind:

- What is the best state representation of the environment for the RL agent?
- Which RL model will have the best performance in the proposed environment?
- How do we generalize the environment to more classes of circuits?

References

Murali, Prakash & Baker, Jonathan & Javadi-Abhari, Ali & Chong, Frederic & Martonosi, Margaret. (2019). Noise-Adaptive Compiler Mappings for Noisy Intermediate-Scale Quantum Computers. 1015-1029. 10.1145/3297858.3304075

This manuscript has been authored by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics.