Quantum circuit noise simulation with reinforcement learning

Simone Bordoni[1,2,*], Andrea Papaluca[2,3], Piergiorgio Buttarini[1,2], Alejandro Sopena[4], Stefano Carrazza[2] and Stefano Giagu[1]

[1] *La Sapienza University of Rome, Dep. of Physics, Rome, Italy.*
[2] *Technology Innovation Institute, Abu Dhabi, UAE.*
[3] *School of Computing, The Australian National University, Canberra, ACT, Australia*
[4] *Instituto de Física Teórica, UAM/CSIC, Universidad Autónoma de Madrid, Madrid, Spain.*

Abstract

Quantum computing in the NISQ era requires powerful tools to reduce the gap between simulations and quantum hardware execution. In this work, we present a machine learning approach for reproducing the noise of a specific quantum device during simulations. The proposed algorithm is meant to be more flexible, in reproducing different noise conditions, than standard techniques like randomized benchmarking or heuristic noise models. This model has been tested both with simulation and on real superconducting qubits.

1. Introduction

Noise Intermediate Scale Quantum (NISQ) [1] devices are limited in usability and reliability mainly because of errors due to: interactions with the environment, thermal relaxation, measurement errors and cross-talk [2, 3, 4]. Hence, it is widely regarded that near-term quantum advantage will only be achieved through advanced error mitigation techniques [5, 6, 7, 8] or only with the future generations of fault-tolerant quantum devices [9, 10, 11, 12]. Even if no advantage has been proven on NISQ devices, many algorithms have been developed and deployed on this hardware. In particular, machine learning inspired models have shown encouraging results in the last few years [13, 14, 15, 16]. To study this kind of algorithms it is important to be able to emulate their behavior when executed on imperfect quantum chips. In this work we propose to train a machine learning model for learning a hardware specific noise and use it to replicate NISQ's behaviour during circuit simulations. This objective is further motivated by the fact that very few techniques of noise modeling or noise prediction are available to this date [17, 18, 19]. In our approach we train a reinforcement learning (RL) agent [20, 21, 22] to add noise channels to different moments in the circuit and reproduce the noise pattern of a specific quantum chip. In this way we reduce as much as possible the heuristic assumption on the noise model, thus, increasing the adaptability and generalization properties of the algorithm.

AIQxQIA, International Workshop on AI for Quantum and Quantum for AI, Nov. 2023.
This work is supported by ICSC – Centro Nazionale di Ricerca in High Performance Computing, Big Data and Quantum Computing, funded by European Union – NextGenerationEU.
simone.bordoni@uniroma1.it (S. Bordoni); andrea.papaluca@anu.edu.au (A. Papaluca)

2. Background

Reinforcement Learning (RL) is a powerful paradigm in machine learning that involves the training of an agent to make optimal decisions in a dynamic environment. It relies on the fundamental concepts of policy and reward functions. The policy determines the behavior of the agent, mapping the states of the environment to actions. The reward function assigns a numeric value to state-action pairs, indicating the immediate desirability or cost associated with them. Training the algorithm involves finding an optimal policy that maximizes the expected long-term cumulative reward. During training, different episodes of agent-environment interaction are executed. At the end of each episode the reward is used to update the weights of the policy, which is often approximated by a Neural Network (NN). In recent years different optimization methods have been developed to improve RL convergence and stability during training [23]. In our work we have obtained the best results using the proximal policy optimization (PPO) [24].

Noise is a crucial challenge faced in quantum circuits as qubits are susceptible to environment interaction, spontaneous emission or calibration errors. Many mathematical tools have been developed in order to describe quantum noise [25, 26, 27]. In this work we will model quantum noise using the error channels illustrated in the following.
The local depolarizing channel is a simple model for incoherent noise. It is characterized by a parameter λ that tends to bring the state to the maximally mixed state:

$$Dep_{\lambda}(\rho) = (1-\lambda)\rho + \lambda \cdot \mathbb{I}/2 \tag{1}$$

The amplitude damping channel models physical processes occurring on the qubits that involve energy dissipation, such us spontaneous emission. Amplitude damping describes a decay process where the state $|0\rangle$ is conserved, while $|1\rangle$ decays to $|0\rangle$ with probability γ:

$$Damp(\gamma)\,|1\rangle = (1-\gamma)\,|1\rangle + \gamma\,|0\rangle \tag{2}$$

Single-qubit coherent errors can be represented with rotation gates (R_x, R_y, R_z). Although these errors can be corrected once identified, in the NISQ era they are more challenging to address since they accumulate after successive gate executions, leading to a bias in the output of the quantum circuit. In order to obtain a simple modeling of the noise, it is possible to use a technique called randomized benchmarking (RB) [28, 29, 30]. RB allows to estimate the magnitude of the average error of a set of quantum gates. The noise model obtained in this way is unrealistic as all noise sources are projected on the depolarizing channel. However, this technique allows for estimating the average gate fidelity and can be used as a basic benchmark for other, more sophisticated, noise characterization techniques.

3. Methodology

All the datasets used in this work are composed of ensembles of random quantum circuits with their relative density matrices (DM). DMs are used as ground truth labels during the training of the algorithm. They are analytically computed in simulation, whereas they are extracted using quantum state tomography [31] when circuits are run on hardware. All the circuits in the

datasets have been generated by extracting random gates among a set of native gates. For this work, we have used the native gates set of the Technology Innovation Institute (TII) quantum hardware, composed of the R_x, R_z and CZ gates. In order to train the RL agent it is necessary to translate the quantum circuit formalism to a tensor representation that can be fed to the policy neural network. We will call this array the quantum circuit representation (QCR). In the general case the shape of the QCR is: $[qubits, depth, encoding\_dim]$. The first entrance of the QCR identifies the circuit qubit, while the second specifies the circuit moment. The third entrance, *encoding_dim*, defines the dimension of the latent space used to embed the information about gates and noise channels that act on a specific qubit at a circuit moment.
For the training the following steps are performed in this order. A noiseless, quantum circuit representation is given to the agent at the beginning of each episode. For every circuit moment the agent performs an observation of the circuit and takes an action that consists in inserting any combination of the aforementioned noise channels with a chosen parameter. At the end of an episode, the DM of the circuit obtained with this process is computed (generated DM). The generated DM is compared with the real noisy DM, produced by the noise we want to model, and their fidelity is used to compute the final reward. After many episodes the agent should learn, in principle, where to insert noise channels in a non-noisy circuit in order to reconstruct the DM of the real noisy circuit. Once trained, the agent should be able to generalize to new unseen circuits and could be used to perform realistic noisy simulations.

4. Results

We tested the proposed algorithm for learning both a simulated noise and the real noise of a quantum superconducting chip. In the first case, we have tested the model on simulated circuits of one and three qubits with a custom noise model. This noise model uses all the errors introduced in section 2. In order to train the RL agent, 500 random circuits of depth 7 have been generated: 400 circuits for the training set and 100 for the test set. Figure 1 reports the DM fidelity and trace distance obtained during training for both a single qubit and three qubits circuits. The standard deviation over our train and test sets is represented by error bars. In both cases, the agent is able to learn the simulated noise, no overfit is observed. Convergence is reached after $\sim 4 \cdot 10^5$ and $\sim 1.5 \cdot 10^6$ time steps respectively, where the standard deviation of the fidelity begins to shrink. In order to test the generalization properties of the model, we tested our agent on circuits with depth spanning from 3 to 30. Figure 2 reports the comparison with the RB model under different metrics. The agent is capable of generalizing to both longer and shorter circuits and always provides more accurate results compared to RB. This means that, as RB considers all the noise sources as depolarizing, our algorithm is able to identify the specific features of the noise.
To test the algorithm on real quantum hardware, we have used a single qubit superconducting transmon [32] chip available at the Technology Innovation Institute (TII) of Abu Dhabi. Figure 3 reports the average DM fidelity and trace distance obtained during the training. Even though the noise of a real quantum chip is far more complex, the agent is still able to learn the noise. We are currently working on testing the algorithm on real three qubits chips.

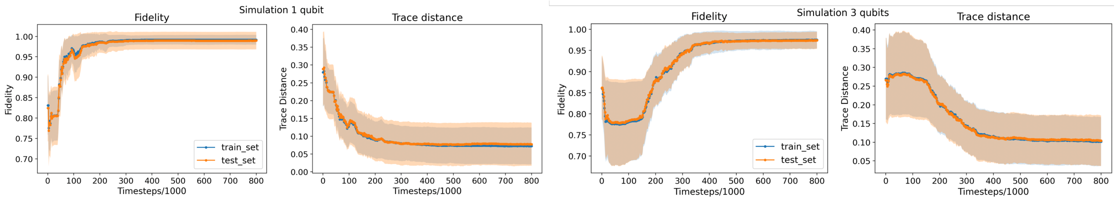

Figure 1: Average DM fidelity and trace distance, obtained during training for single qubit circuits (left) and three qubit circuits (right) with a custom noise model in simulations.

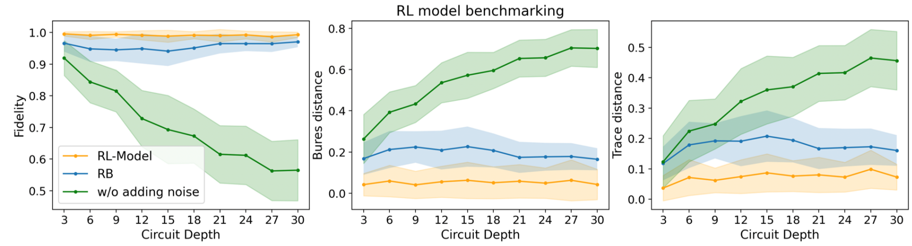

Figure 2: Performance comparison of the single qubit RL agent (yellow) with respect to the RB (blue) and the result obtained without adding any noise channel (green). The performance have been computed using different metrics and for circuits of different depth.

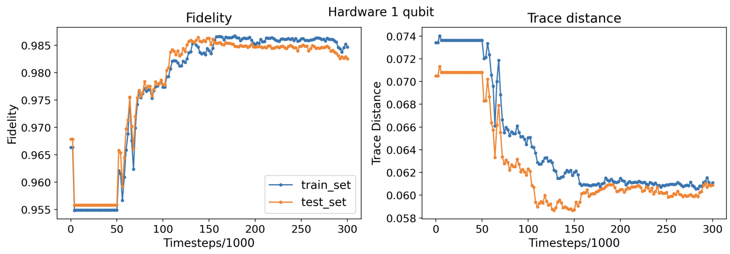

Figure 3: Training metrics of the RL algorithm for the single qubit TII quantum chip.

5. Conclusions and future work

We presented a RL algorithm to replicate, both, a simulated and a real quantum chip noise, which provided better results than RB. The applications of the proposed algorithm extends beyond noise characterization. By learning the error pattern associated to specific qubit-gate combinations, the model could be used to optimize the transpiling process [33], preferring the execution of gates on less noisy qubits during circuit routing. Moreover, we aim at investigating if a similar, but inverse, strategy could be used to perform error mitigation. The scaling of the model to circuits with many qubits is the current biggest limitation as state tomography is unfeasible for many qubits. A potential solution could be training the model to approximate the statistics of measurement outcomes, or using classical shadow state reconstruction [34, 35]. The neural networks used in this work do not require high computational power, a single policy NN has about 10^4 parameters. This opens the possibility of training an ensemble of networks that operate in parallel on subsets of qubits from a larger chip.

References

[1] J. Preskill, Quantum computing in the NISQ era and beyond, Quantum 2 (2018) 79. URL: https://doi.org/10.22331%2Fq-2018-08-06-79. doi:10.22331/q-2018-08-06-79.

[2] P. V. Klimov, J. Kelly, Z. Chen, M. Neeley, A. Megrant, B. Burkett, R. Barends, K. Arya, B. Chiaro, Y. Chen, A. Dunsworth, A. Fowler, B. Foxen, C. Gidney, M. Giustina, R. Graff, T. Huang, E. Jeffrey, E. Lucero, J. Y. Mutus, O. Naaman, C. Neill, C. Quintana, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. C. White, S. Boixo, R. Babbush, V. N. Smelyanskiy, H. Neven, J. M. Martinis, Fluctuations of energy-relaxation times in superconducting qubits, Phys. Rev. Lett. 121 (2018) 090502. URL: https://link.aps.org/doi/10.1103/PhysRevLett.121.090502. doi:10.1103/PhysRevLett.121.090502.

[3] M. Kjaergaard, M. E. Schwartz, J. Braumüller, P. Krantz, J. I.-J. Wang, S. Gustavsson, W. D. Oliver, Superconducting qubits: Current state of play, Annual Review of Condensed Matter Physics 11 (2020) 369–395. URL: https://doi.org/10.1146/annurev-conmatphys-031119-050605. doi:10.1146/annurev-conmatphys-031119-050605. arXiv:https://doi.org/10.1146/annurev-conmatphys-031119-050605.

[4] P. Zhao, K. Linghu, Z. Li, P. Xu, R. Wang, G. Xue, Y. Jin, H. Yu, Quantum crosstalk analysis for simultaneous gate operations on superconducting qubits, PRX Quantum 3 (2022) 020301. URL: https://link.aps.org/doi/10.1103/PRXQuantum.3.020301. doi:10.1103/PRXQuantum.3.020301.

[5] P. D. Nation, H. Kang, N. Sundaresan, J. M. Gambetta, Scalable mitigation of measurement errors on quantum computers, PRX Quantum 2 (2021) 040326. URL: https://link.aps.org/doi/10.1103/PRXQuantum.2.040326. doi:10.1103/PRXQuantum.2.040326.

[6] L. Funcke, T. Hartung, K. Jansen, S. Kühn, P. Stornati, X. Wang, Measurement error mitigation in quantum computers through classical bit-flip correction, Phys. Rev. A 105 (2022) 062404. URL: https://link.aps.org/doi/10.1103/PhysRevA.105.062404. doi:10.1103/PhysRevA.105.062404.

[7] S. S. Tannu, M. K. Qureshi, Mitigating measurement errors in quantum computers by exploiting state-dependent bias, in: Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture, MICRO '52, Association for Computing Machinery, New York, NY, USA, 2019, p. 279–290. URL: https://doi.org/10.1145/3352460.3358265. doi:10.1145/3352460.3358265.

[8] A. Sopena, M. H. Gordon, G. Sierra, E. López, Simulating quench dynamics on a digital quantum computer with data-driven error mitigation, Quantum Science and Technology 6 (2021) 045003. URL: https://dx.doi.org/10.1088/2058-9565/ac0e7a. doi:10.1088/2058-9565/ac0e7a.

[9] E. Knill, Quantum computing with realistically noisy devices, Nature 434 (2005) 39–44. doi:10.1038/nature03350.

[10] K. Fukui, A. Tomita, A. Okamoto, K. Fujii, High-threshold fault-tolerant quantum computation with analog quantum error correction, Phys. Rev. X 8 (2018) 021054. URL: https://link.aps.org/doi/10.1103/PhysRevX.8.021054. doi:10.1103/PhysRevX.8.021054.

[11] A. G. Fowler, M. Mariantoni, J. M. Martinis, A. N. Cleland, Surface codes: Towards practical large-scale quantum computation, Phys. Rev. A 86 (2012) 032324. URL: https:

//link.aps.org/doi/10.1103/PhysRevA.86.032324. doi:10.1103/PhysRevA.86.032324.
[12] B. Simone, G. Stefano, Convolutional neural network based decoders for surface codes, Quantum Information Processing 22 (2023) 151. URL: https://doi.org/10.1007/s11128-023-03898-2. doi:10.1007/s11128-023-03898-2.
[13] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, S. Lloyd, Quantum machine learning, Nature 549 (2017) 195–202. URL: https://doi.org/10.1038%2Fnature23474. doi:10.1038/nature23474.
[14] J. Choi, J. Kim, A tutorial on quantum approximate optimization algorithm (qaoa): Fundamentals and applications, in: 2019 International Conference on Information and Communication Technology Convergence (ICTC), 2019, pp. 138–142. doi:10.1109/ICTC46691.2019.8939749.
[15] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, P. J. Coles, Variational quantum algorithms, Nature Reviews Physics 3 (2021) 625–644. URL: https://doi.org/10.1038%2Fs42254-021-00348-9. doi:10.1038/s42254-021-00348-9.
[16] S. Bordoni, D. Stanev, T. Santantonio, S. Giagu, Long-lived particles anomaly detection with parametrized quantum circuits, Particles 6 (2023) 297–311. URL: https://www.mdpi.com/2571-712X/6/1/16. doi:10.3390/particles6010016.
[17] A. Zlokapa, A. Gheorghiu, A deep learning model for noise prediction on near-term quantum devices, 2020. arXiv:2005.10811.
[18] K. Georgopoulos, C. Emary, P. Zuliani, Modeling and simulating the noisy behavior of near-term quantum computers, Phys. Rev. A 104 (2021) 062432. URL: https://link.aps.org/doi/10.1103/PhysRevA.104.062432. doi:10.1103/PhysRevA.104.062432.
[19] R. Harper, S. T. Flammia, J. J. Wallman, Efficient learning of quantum noise, Nature Physics 16 (2020) 1184–1188. URL: https://doi.org/10.1038%2Fs41567-020-0992-8. doi:10.1038/s41567-020-0992-8.
[20] R. S. Sutton, A. G. Barto, Reinforcement learning: An introduction, MIT press, 2018.
[21] L. P. Kaelbling, M. L. Littman, A. W. Moore, Reinforcement learning: A survey, Journal of artificial intelligence research 4 (1996) 237–285.
[22] M. A. Wiering, M. Van Otterlo, Reinforcement learning, Adaptation, learning, and optimization 12 (2012) 729.
[23] R. Özalp, N. K. Varol, B. Taşci, A. Uçar, A Review of Deep Reinforcement Learning Algorithms and Comparative Results on Inverted Pendulum System, Springer International Publishing, Cham, 2020, pp. 237–256. URL: https://doi.org/10.1007/978-3-030-49724-8_10. doi:10.1007/978-3-030-49724-8_10.
[24] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, O. Klimov, Proximal policy optimization algorithms, 2017. arXiv:1707.06347.
[25] E. C. G. Sudarshan, P. M. Mathews, J. Rau, Stochastic dynamics of quantum-mechanical systems, Phys. Rev. 121 (1961) 920–924. URL: https://link.aps.org/doi/10.1103/PhysRev.121.920. doi:10.1103/PhysRev.121.920.
[26] M. A. Nielsen, I. L. Chuang, Quantum Computation and Quantum Information: 10th Anniversary Edition, Cambridge University Press, 2010. doi:10.1017/CBO9780511976667.
[27] Y. I. Bogdanov, A. Y. Chernyavskiy, A. S. Holevo, V. F. Lukichev, A. A. Orlikovsky, Mathematical models of quantum noise, in: A. A. Orlikovsky (Ed.), International Conference Micro-

and Nano-Electronics 2012, volume 8700, International Society for Optics and Photonics, SPIE, 2013, p. 870019. URL: https://doi.org/10.1117/12.2017396. doi:10.1117/12.2017396.

[28] J. Emerson, R. Alicki, K. Życzkowski, Scalable noise estimation with random unitary operators, Journal of Optics B: Quantum and Semiclassical Optics 7 (2005) S347. URL: https://dx.doi.org/10.1088/1464-4266/7/10/021. doi:10.1088/1464-4266/7/10/021.

[29] M. Heinrich, M. Kliesch, I. Roth, Randomized benchmarking with random quantum circuits, 2023. arXiv:2212.06181.

[30] J. Helsen, X. Xue, L. M. K. Vandersypen, S. Wehner, A new class of efficient randomized benchmarking protocols, npj Quantum Information 5 (2019) 71. doi:10.1038/s41534-019-0182-7. arXiv:1806.02048.

[31] M. Christandl, R. Renner, Reliable quantum state tomography, Phys. Rev. Lett. 109 (2012) 120403. URL: https://link.aps.org/doi/10.1103/PhysRevLett.109.120403. doi:10.1103/PhysRevLett.109.120403.

[32] M. H. Devoret, R. J. Schoelkopf, Superconducting circuits for quantum information: An outlook, Science 339 (2013) 1169–1174. URL: https://www.science.org/doi/abs/10.1126/science.1231930. doi:10.1126/science.1231930. arXiv:https://www.science.org/doi/pdf/10.1126/science.1231930.

[33] E. Wilson, S. Singh, F. Mueller, Just-in-time quantum circuit transpilation reduces noise, in: 2020 IEEE International Conference on Quantum Computing and Engineering (QCE), 2020, pp. 345–355. doi:10.1109/QCE49297.2020.00050.

[34] G. Struchalin, Y. A. Zagorovskii, E. Kovlakov, S. Straupe, S. Kulik, Experimental estimation of quantum state properties from classical shadows, PRX Quantum 2 (2021) 010307. URL: https://link.aps.org/doi/10.1103/PRXQuantum.2.010307. doi:10.1103/PRXQuantum.2.010307.

[35] J. Eisert, D. Hangleiter, N. Walk, I. Roth, D. Markham, R. Parekh, U. Chabaud, E. Kashefi, Quantum certification and benchmarking, Nature Reviews Physics 2 (2020) 382–390. URL: https://doi.org/10.1038%2Fs42254-020-0186-4. doi:10.1038/s42254-020-0186-4.