

QRAFT: Reverse Your Quantum Circuit and Know the Correct Program Output

Extended Abstract

Tirthak Patel and Devesh Tiwari
Northeastern University

1. Motivation

Quantum computing has transitioned from theoretical promise to practical physical realization. However, one of the major bottlenecks toward wide-adoption of quantum computers is the high error rate in current Noisy Intermediate-Scale Quantum (NISQ) machines [13, 7]. On quantum computers, the program may complete successfully, but the the final program output may be incorrect – hence, the programmers do not know the correct program output (the output if the program was executed on an error-free quantum machine). While quantum computing promises orders of magnitude performance improvement for a class of algorithms, such improvements are not useful if the programmer cannot deduce the correct program output. Therefore, the goal of this paper is to *help quantum programmers automatically deduce the correct program output while running on error-prone NISQ machines.*

2. Limitations of the State of the Art

NISQ machines suffer from multiple types of errors which can make the final algorithm output erroneous, including qubit coherence errors, quantum gate operation errors, and state preparation and measurement errors. However, error rate on different qubits on a given machine may differ significantly depending upon the quantum operation type. Therefore, the impact of error experienced by a quantum program can be reduced by intelligently mapping a program’s logical operations to the least erroneous physical component on a given machine (e.g., avoid mapping the program on physical qubits that exhibit higher error rate for a certain type of quantum gate or operation).

Today, existing approaches reduce the impact of error experienced by a quantum program by intelligently mapping a program’s logical operations to the least erroneous physical component on a given quantum machine [16, 4, 23, 14, 20, 8, 15, 17, 2]. This approach is known as “**optimal circuit mapping**” – where the logical operations of a quantum program are intelligently *mapped* on a set of physical qubits considering multiple factors including the historical reliability characteristics of the physical qubits, the physical connection between qubits, and the sequence of operations in the quantum program. However, the current approaches have two major limitations:

I. Lack of knowledge about a program’s true output: Current approaches aim to minimize the error occurrence prob-

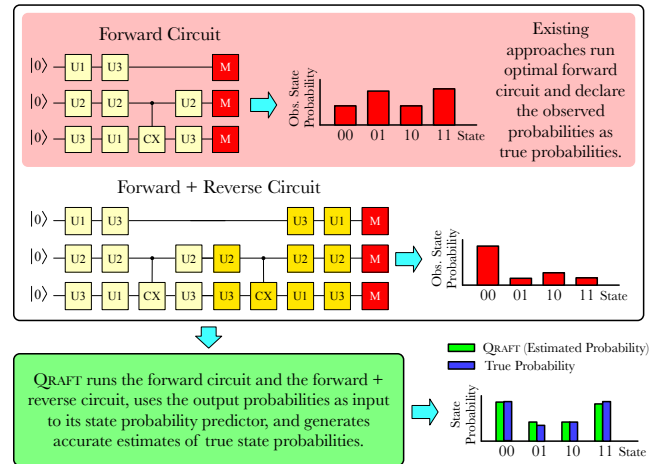


Figure 1: QRAFT leverages the forward + reverse circuit to generate accurate estimates of true state probabilities, as compared to existing optimal circuit mapping approaches.

ability and learn certain characteristics of the program under study (e.g., number of operations, type of operations, etc.) [21, 11, 1, 3, 23]. *However, these approaches do not (and cannot) know the “correct/true program output” of the program irrespective of how many times or on which qubits the program is executed. Hence, the observed output can only be used as the best guess of the correct program output.*

All prior studies assume that they know the correct program output apriori and report the difference between the observed output and the apriori-known correct program output as the “error in estimation of the program output” [17, 10, 18, 4, 22]. Unfortunately, this approach is not useful for programmers who may not always know the correct program output and may have multiple correct output states with probabilities of different magnitudes. Estimating the correct probability of these output states is critical, but current approaches are simply insufficient to tackle this.

II. Everyone has access to the “best” qubits and knows “which” qubits are the best: Existing approaches over-optimistically assume that each end user has access to all the qubits and their historical error rates [12, 17, 4, 8, 15, 9, 20], especially the most reliable ones to run their programs – an assumption that will be seriously challenged in future quantum computing systems shared among multiple users concurrently [2]. Moreover, historical error information may become business-sensitive, similar to classical computing systems.

3. Key Insights

The key insight behind QRAFT¹ is to *reverse the quantum circuit and repeatedly execute the full (forward + reverse) circuit* to deduce the correct program output of the original (or forward) circuit. Quantum operations, unlike classical operations, are reversible – that is, the original input can be restored by applying an inverse operation. QRAFT leverages this property and extends it to the whole quantum circuit. Since quantum states cannot be cloned or checkpointed, QRAFT appends the reverse circuit at the end of the original (forward) circuit and executes the full (forward + reverse) circuit.

Reversing a circuit enables QRAFT to partially know the inherent correct output of the full circuit – all output states must get reduced to original input states – a piece of “ground truth” information that is not available to us irrespective of how many times the original circuit is executed, or on whichever qubit it is executed. This allows QRAFT to have a better estimation of how much error might have occurred in the original circuit execution, compared to existing approaches. This is because existing approaches only focus on optimizing the circuit map. As shown in Fig. 1, the reversibility property of quantum operations that allows us to know the ground truth of states at the end of the execution of the full circuit has never been exploited to deduce the correct program output.

4. Main Artifacts

QRAFT is a new methodology for automatically deducing the correct program output, which is implemented as a software tool and evaluated on multiple IBM QX quantum machines. The workflow of QRAFT includes generating, reversing, and running quantum circuits to construct the training dataset, training a prediction model using the dataset, and using the prediction model for predicting true state probabilities of quantum algorithms. Artifacts of QRAFT include the over 1400 random quantum circuits and quantum algorithms run using the python-based Qiskit quantum programming framework on IBM QX cloud quantum computers and the model training tools of QRAFT implemented in MATLAB. QRAFT’s framework will be open-source available with the publication of the paper.

5. Key Results and Contributions

QRAFT is a novel, automated method to accurately estimate the true program output – QRAFT demonstrates for the first time that reversing the circuit can *reveal* what the program actually does. While this approach appears to be encouraging, a straight forward rule-based application of this approach does not yield the expected results due to the complex interactions of errors with the underlying original circuit. QRAFT designs and develops a learning-based prediction model that generates

¹QRAFT, pronounced as craft, conveniently stands for quantum circuit reversal for attaining the full truth (about the program output).

the true output state probabilities based on the observed state probabilities from the forward circuit and the observed errors from the forward + reverse circuit.

QRAFT, while being simple in hindsight, is surprisingly effective and improves upon the state-of-the-art approaches based on optimizing the circuit mappings to accurately estimate the magnitude of probabilities for programs with multiple states with non-zero output probabilities – departure from existing works which do not handle programs effectively with multiple states with non-zero output probabilities states and focus largely on dominant states [19, 15, 17, 1].

Our evaluation demonstrates that QRAFT reduces the median state error by up to 7%, dominant state error by up to 30%, and total program error by up to 20% across different algorithms. The state-of-the-art approach has only 20% of the circuit states with 0% error, while QRAFT ensures that over 70% of the states have 0% error when tested with 200 previously unseen and randomly generated circuits.

Furthermore, QRAFT demonstrates that it is possible to deduce the correct program output successfully, even in the absence of the optimal circuit map (relaxing a prerequisite of existing approaches which assume unrestricted access to the “best” qubits and know “which” qubits are the best based on the historical error information). As expected, using the optimal circuit map makes QRAFT deliver near-perfect deduction of correct program output in many cases, but QRAFT remains fairly effective even when a non-optimized circuit map is used. Our evaluation confirms that QRAFT scales to larger number of qubits, and its effectiveness is not sensitive to the choice of the platform or specific circuit characteristics such as circuit depth and number of operations.

6. Why ASPLOS

QRAFT provides a novel solution in the quantum computing architecture and systems domain – a non-traditional but emerging area (one of the foci of ASPLOS conference as per the call for papers). QRAFT attempts to bridge the gap between the traditional computer systems community and the quantum computing community. Recently, ASPLOS has published multiple quantum computing systems papers focusing on improving usability and resiliency of NISQ devices [9, 5, 8, 6, 19, 14, 4].

7. Citation for Most Influential Paper Award

QRAFT was the first work to leverage and demonstrate the reversibility property of quantum computing to deduce the correct output from the erroneous executions of quantum programs on then error-prone near-term intermediate scale quantum computers. QRAFT’s insights and open-source contributions enabled researchers to study and advance various other aspects of early quantum programs including resilience, debugging, correctness, and effective resource management.

References

- [1] Debjyoti Bhattacharjee, Abdullah Ash Saki, Mahabubul Alam, Anupam Chattopadhyay, and Swaroop Ghosh. MUQUT: Multi-Constraint Quantum Circuit Mapping on NISQ Computers. In *38th IEEE/ACM International Conference on Computer-Aided Design, ICCAD 2019*, page 8942132. Institute of Electrical and Electronics Engineers Inc., 2019.
- [2] Poulami Das, Swamit S Tannu, Prashant J Nair, and Moinuddin Qureshi. A Case for Multi-Programming Quantum Computers. In *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*, pages 291–303. ACM, 2019.
- [3] Pranav Gokhale, Yongshan Ding, Thomas Propson, Christopher Winkler, Nelson Leung, Yunong Shi, David I Schuster, Henry Hoffmann, and Frederic T Chong. Partial Compilation of Variational Algorithms for Noisy Intermediate-Scale Quantum Machines. In *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*, pages 266–278. ACM, 2019.
- [4] Gushu Li, Yufei Ding, and Yuan Xie. Tackling the Qubit Mapping Problem for NISQ-Era Quantum Devices. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 1001–1014. ACM, 2019.
- [5] Gushu Li, Yufei Ding, and Yuan Xie. Towards efficient superconducting quantum processor architecture design. In *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 1031–1045, 2020.
- [6] Ji Liu, Gregory T Byrd, and Huiyang Zhou. Quantum Circuits for Dynamic Runtime Assertions in Quantum Computation. In *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 1017–1030, 2020.
- [7] Margaret Martonosi and Martin Roetteler. Next Steps in Quantum Computing: Computer Science’s Role. *arXiv preprint arXiv:1903.10541*, 2019.
- [8] Prakash Murali, Jonathan M Baker, Ali Javadi-Abhari, Frederic T Chong, and Margaret Martonosi. Noise-Adaptive Compiler Mappings for Noisy Intermediate-Scale Quantum Computers. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 1015–1029. ACM, 2019.
- [9] Prakash Murali, David C McKay, Margaret Martonosi, and Ali Javadi-Abhari. Software Mitigation of Crosstalk on Noisy Intermediate-Scale Quantum Computers. In *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 1001–1016, 2020.
- [10] Daniel C Murphy and Kenneth R Brown. Controlling Error Orientation to Improve Quantum Algorithm Success Rates. *Physical Review A*, 99(3):032318, 2019.
- [11] Tirthak Patel, Baolin Li, Rohan Basu Roy, and Devesh Tiwari. {UREQA}: Leveraging operation-aware error rates for effective quantum circuit mapping on nisq-era quantum computers. In *2020 {USENIX} Annual Technical Conference ({USENIX}{ATC} 20)*, pages 705–711, 2020.
- [12] Tirthak Patel and Devesh Tiwari. Veritas: Accurately Estimating the Correct Output on Noisy Intermediate-Scale Quantum Computers. In *2020 SC20: International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*, pages 188–203. IEEE Computer Society, 2020.
- [13] John Preskill. Quantum Computing in the NISQ Era and Beyond. *Quantum*, 2:79, 2018.
- [14] Yunong Shi, Nelson Leung, Pranav Gokhale, Zane Rossi, David I Schuster, Henry Hoffmann, and Frederic T Chong. Optimized Compilation of Aggregated Instructions for Realistic Quantum Computers. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 1031–1044. ACM, 2019.
- [15] Kaitlin N Smith and Mitchell A Thornton. A Quantum Computational Compiler and Design Tool for Technology-Specific Targets. In *Proceedings of the 46th International Symposium on Computer Architecture*, pages 579–588. ACM, 2019.
- [16] Bochen Tan and Jason Cong. Optimality Study of Existing Quantum Computing Layout Synthesis Tools. *arXiv preprint arXiv:2002.09783*, 2020.
- [17] Swamit S Tannu and Moinuddin Qureshi. Ensemble of Diverse Mappings: Improving Reliability of Quantum Computers by Orchestrating Dissimilar Mistakes. In *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*, pages 253–265. ACM, 2019.
- [18] Swamit S Tannu and Moinuddin 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*, pages 279–290. ACM, 2019.
- [19] Swamit S Tannu and Moinuddin K Qureshi. Not All Qubits are Created Equal: A Case for Variability-Aware Policies for NISQ-Era Quantum Computers. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 987–999. ACM, 2019.
- [20] Robert Wille, Lukas Burgholzer, and Alwin Zulehner. Mapping Quantum Circuits to IBM QX Architectures Using the Minimal Number of SWAP and H Operations. In *Proceedings of the 56th Annual Design Automation Conference 2019*, page 142. ACM, 2019.
- [21] Ellis Wilson, Sudhakar Singh, and Frank Mueller. Just-in-time Quantum Circuit Transpilation Reduces Noise. *arXiv preprint arXiv:2005.12820*, 2020.
- [22] Alwin Zulehner, Alexandru Paler, and Robert Wille. An Efficient Methodology for Mapping Quantum Circuits to the IBM QX Architectures. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 2018.
- [23] Alwin Zulehner and Robert Wille. Compiling SU (4) Quantum Circuits to IBM QX Architectures. In *Proceedings of the 24th Asia and South Pacific Design Automation Conference*, pages 185–190. ACM, 2019.