

Understanding the Effect of Transpilation in the Reliability of Quantum Circuits

Original

Understanding the Effect of Transpilation in the Reliability of Quantum Circuits / Dilillo, N., Giusto, E., Dri, E., Baheri, B., Guan, Q., Montrucchio, B., Rech, P.. - ELETTRONICO. - (2023), pp. 232-235. (IEEE International Conference on Quantum Computing and Engineering (QCE) Bellevue, WA (USA) 17-22 September 2023) [10.1109/QCE57702.2023.10220].

Availability:

This version is available at: 11583/2982673 since: 2023-12-12T10:35:49Z

Publisher:

IEEE

Published

DOI:10.1109/QCE57702.2023.10220

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2023 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

Understanding the Effect of Transpilation in the Reliability of Quantum Circuits

Nicola Dilillo^{✉*}, Edoardo Giusto^{✉*}, Emanuele Dri^{✉*}, Betis Baheri^{✉†},
Qiang Guan^{✉†}, Bartolomeo Montrucchio^{✉*}, Paolo Rech^{✉‡}

*DAUIN, Politecnico di Torino, Italy, {<name.surname>}@polito.it

†Kent State University, US, {bbaheri, qguan}@kent.edu

‡University of Trento, Italy, paolo.rech@unitn.it

Abstract—Transpiling is a necessary step to map a logical quantum algorithm to a circuit executed on a physical quantum machine, according to the available gate set and connectivity topology. Different transpiling approaches try to minimize the most critical parameters for the current transmon technology, such as Depth and CNOT number. Crucially, these approaches do not take into account the reliability of the circuit. In particular, transpilation can modify how radiation-induced transient faults propagate. In this paper, we aim at advancing the understanding of transpilation impact on fault propagation by investigating the low-level reliability of several transpiling approaches. We considered 4 quantum algorithms transpiled for 2 different architectures, increasing the number of qubits, and all possible logical-to-physical qubit mapping, adding to a total of 4,640 transpiled circuits. We inject a total of 202,124 faults and track their propagation. Our experiments show that by simply choosing the proper transpilation, the reliability of the circuit can improve by up to 14%.

Index Terms—Transpilation, Transient faults, Reliability, Co-Design

I. INTRODUCTION

The pursuit of quantum computing is akin to the race to the moon, with competitors globally investing heavily and taking pride in its potential for critical applications such as quantum simulation, combinatorial optimization, machine learning, and cryptography. Constructing a quantum computing system requires a collaborative, interdisciplinary effort from both theoretical and engineering perspectives.

Quantum computing became a viable computing paradigm when fault-tolerant qubits were successfully developed [1], [2], enabling the computation of small but fundamental circuits. This marked a significant turning point in the field. The challenge of ensuring qubit reliability stems from their sensitivity to external disturbances and the inherent unpredictability of quantum mechanics, as the state of qubits can change randomly. Pioneering research has revealed that, in addition to noise, it is crucial to consider the effect of natural radiation on superconducting qubits [3]–[8], as ionizing particles can significantly reduce their fault tolerance [9], [10]. Due to their higher sensitivity to external disturbances compared to CMOS transistors, quantum computers may actually be more vulnerable to ionizing radiation than classical computers. In fact, recent research has revealed that light particles, including muons [10] and even infrared light [11], can impact qubits,

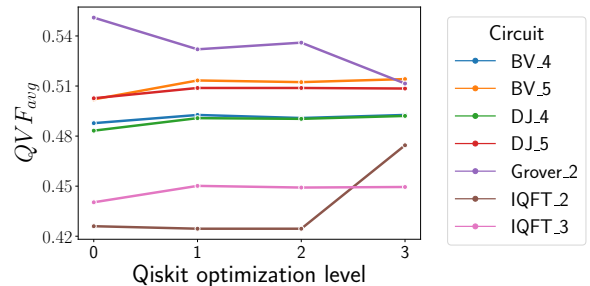


Fig. 1: Average QVF_{avg} (probability for a fault to propagate to the output) for various circuits transpiled with optimization levels from 0 to 3 for the IBM Santiago machine.

while not having enough energy to significantly affect CMOS behavior.

Since the fault is originated in the underlying hardware but its effect is observed in the software output, the way the (logical) algorithm is mapped in the available (hardware) resources can potentially significantly impact the probability for the fault to propagate. This process, called *transpilation*, can thus significantly modify fault propagation. With this paper, we intend to understand the impact of the transpilation process in the reliability of a quantum algorithm.

The objective of transpilation is to align a high-level quantum program with a specific quantum device, improving the program’s performance and ensuring that it meets the device’s physical constraints. Building on prior research that introduced the Quantum Vulnerability Factor (QVF) [12], [13] (i.e., the probability for a fault to propagate to the circuit output) as an indicator of a circuit’s susceptibility to transient faults, we identify the resources that, if compromised, are more likely to affect the accuracy of the output. This approach enables us to rapidly identify vulnerabilities in the system in order to then take appropriate measures to mitigate them.

As an example, Figure 1 shows the average QVF value for the 7 considered circuits, each transpiled with Qiskit¹ transpiler with optimization level from 0 to 3. These are only the average values, but already highlight a significant difference depending on the selected approach. The mean

¹<https://qiskit.org/>

difference in QVF_{avg} is 1.9%, the minimum difference is 0.5%, obtained by the Deutsch-Josza algorithm on 5 qubits, while the maximum difference is 5%, obtained by the Inverse Quantum Fourier Transform on 2 qubits.

The main contributions of this work are the following:

- Our design involves creating a quantum fault injector that is interoperable with both Qiskit and PennyLane² and can run on simulators or on actual quantum devices.
- We assess the impact of faults on output probabilities by employing the QVF metric at the transpiled circuit level.
- We perform an in-depth analysis of the reliability profile of the considered quantum algorithms, inspecting the contribution of different approaches of the Qiskit transpiler.

II. BACKGROUND AND RELATED WORKS

The capability of quantum computers to perform calculations is bounded by intrinsic machine noise, which greatly reduces the accuracy of quantum computation [14]. A qubit can properly maintain its state for a limited amount of time, which is made up by T1 and T2 times [15]. These two times are well understood and characterized by researchers and manufacturers, who are constantly trying to increase them with innovative machine design and isolation techniques.

Research is making progress not only in the hardware domain. On the software side, the quantum community effort runs on two parallel tracks: *compilation* and *Quantum Error Correction* (QEC). QEC techniques have taken a prominent role in increasing the resilience of quantum machines with respect to intrinsic noise. Despite the large cost for their implementation, standard QEC techniques are incapable of handling radiation-induced transient faults. Precursory research in this domain highlighted the impact of radiation-induced faults in transmon qubits [6], [7], [9], [10]. Within these phenomena, a particle strike deposits a certain amount of charge in the substrate of the qubit housing, generating a current that alters the qubit state, generating a fault. The influence of these phenomena is so significant that it could impede the large-scale adoption of such technology [3]–[5].

III. METHODOLOGY

This section describes methods adopted to simulate the behaviour of transient faults on a target transpiled circuit and the metrics used to evaluate its reliability. In particular, the focus here is on selecting a transpilation that is more resilient to faults than others.

The experiment was conducted on a server with the following specifications: 8 x Intel(R) Xeon(R) CPU E3-1245 v5 @ 3.50GHz processor; 32 GB @ 2400 Mhz RAM; 1 TB SSD storage; Ubuntu 18.04.6 LTS (Bionic Beaver) as Operating System.

A. Quantum Vulnerability Factor

The reliability assessment of the different transpilations is carried out by means of the Quantum Vulnerability Factor [12].

This metric measures the impact of a corruption in a qubit on the circuit output probability distribution.

A value of QVF near 0 means that the circuit output is same as we expected or very similar, thus no issues have arisen. Instead, a value near 1 means that the circuit output is far from the expected one and that one or more issues have arisen in the circuit.

B. QuFI

The tool used to carry out the fault injection campaign is QuFI, the Quantum Fault Injector [16]. QuFI introduces a fault in each possible injection point (i.e. before each gate) for each qubit in the circuit. The injected fault is modeled as a phase shift of the two angles (θ, ϕ) of the state of the qubit on the Bloch sphere. This shift is achieved with the use of a generic U gate. This gate is parameterized with θ values ranging from 0 to π and ϕ values ranging from 0 to 2π , both with a step of $\pi/12$.

QuFI enables us to easily assess how a fault propagates in a circuit and what impact it has on the circuit output. For our experiment, we simulate without noise, since the aim of the work is to focus on transient-fault-related issues, regardless of other influence.

C. Fault Injection Campaign

A fault injection campaign was performed on all possible Qiskit transpilations of a circuit, which means:

- all different optimization levels, from 0 to 3, where the larger is the optimization level, the fewer the number of CNOT gates and overall depth of the transpiled circuit will be;
- all possible initial layouts, used to map the logical qubits to the physical ones.

Tested circuits, alongside the total number of transpilations, are listed in Table I. All the experiments have been run for two different backends, both with 5 qubits: IBM Santiago and Belem machines. This way, the number of initial layouts is always the same. Each of these initial layouts has been simulated using the 4 different optimization levels available for the Qiskit transpiler. In total, 4640 transpiled circuits have been simulated and analysed.

D. Metric

Each transpiled circuit is coupled with a QVF heatmap, shown in Figure 4.b and 4.c is associated with the function $h(\theta, \phi)$, used to describe the impact of faults.

TABLE I: Fault Injection Campaign Detail

Circuit	# Initial layouts	# Transpilations
Bernstein-Vazirani 4	120	480
Bernstein-Vazirani 5	120	480
Deutsch-Josza 4	120	480
Deutsch-Josza 5	120	480
Grover 2	20	80
IQFT 2	20	80
IQFT 3	60	240

²<https://pennylane.ai/>

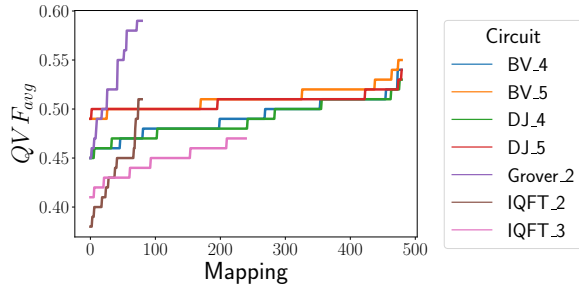


Fig. 2: Sorted QVF_{avg} average values for the 7 considered circuits, transpiled with Qiskit transpiler with all optimization levels and using all possible initial layouts (lowest QVF_{avg} on the left).

This heatmap cannot be directly used to make reliability evaluations because it is based on more than one value. To circumvent this, the metric in Formula 1 has been selected to enable absolute comparisons between different transpilation:

$$QVF_{avg} = \frac{\sum_{\theta=0}^{\pi} \sum_{\phi=0}^{2\pi} h(\theta, \phi)}{325} \quad (1)$$

Where 325 is the total number of heatmap squares - the number of angle pairs. This function basically describes the average value of QVF in the circuit heatmap.

IV. RESULTS AND DISCUSSION

Zooming out from Figure 1, it is possible to highlight maximum and minimum QVF_{avg} values for the considered circuits, regardless of the optimization level, as in Figure 2. This Figure shows, for each circuit, all the possible transpilation with all possible initial mapping layouts. These transpilation are sorted in ascending order for QVF_{avg} value. Hence, in this Figure, the x -axis is purely a list of different mappings. Since different algorithms have different pools of possible initial configurations, as detailed in Table I, the length of their lines is different. This plot is useful to highlight the variability in QVF_{avg} values. The quantum circuit showing the largest such variability is Grover’s algorithm, with a 14% difference, while the smallest is Deutsch-Josza’s algorithm on 5 qubits, with a difference of 5%. The average difference is 8.7%. The smallest value of QVF_{avg} is 0.38, in IQFT’s algorithm on 2 qubits, while the greatest is 0.59, in Grover’s algorithm.

Digging deeper in the behavior of a particular algorithm, Figure 3(a) shows the QVF_{avg} distribution for the Bernstein-Vazirani algorithm on 4 qubits, transpiled with all optimization levels for the IBM Santiago machine. We selected BV4 algorithm since it is the one with largest possible initial mapping layouts and largest variability in QVF_{avg} values. As the Figure shows, optimization level 0 is able to achieve generally lower QVF_{avg} values when compared to other optimization levels, having both maximum and minimum values lower than the other transpiling approaches. Furthermore, it is also possible to highlight that optimization level 3 has the largest

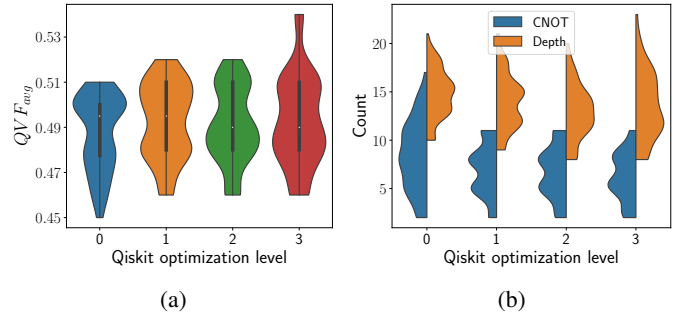


Fig. 3: (a) Distribution of QVF_{avg} and average circuit depth and (b) CNOT for Bernstein-Vazirani algorithm on 4 qubits (BV4), transpiled on the IBM Santiago architecture.

variability among all optimization levels. It is also interesting to compare these values against the obtained CNOT number and circuit depth depending on the optimization level, as in Figure 3(b). The number of CNOT gates and the circuit depth are, in fact, the parameters that the Qiskit transpiler tries to minimize. It is interesting to note that maximum and minimum circuit depths are close among all optimization levels, with optimization level 3 having a minimum depth of just 2 gates less than the optimization level 0 minimum. For what concerns CNOT count instead, optimization level 0 has a maximum count higher than the others, which are the same. However, all optimization levels report not only the same minimum count but also the distributions show that this minimum is achieved a comparable number of times for all optimization levels.

Figure 4 shows a direct comparison of the reliability profile for the best and worst transpilation considering QVF_{avg} value (i.e. lowest and highest). Figure 4(a) quickly compares the considered metrics - QVF_{avg} , circuit depth and CNOT count - normalized from 0 to 1, on the total number of gates. Figures 4(b,c) show the related circuit QVF heatmaps to closely highlight the differences between the two circuits. While the heatmap for the best transpilation shows significant vulnerability only for the $(\phi \simeq 0, \theta \simeq \pi)$ and $(\phi \simeq 2\pi, \theta \simeq \pi)$ injections, the heatmap for the worst transpilation boasts a wider range of angles that have a significant detrimental impact on the output on the circuit. Eventually, Figure 4(d) plots the Δ -heatmap, i.e. each square represents the difference between the heatmap of the best and the worst transpilation, and it is instrumental in understanding the benefit of the best one. This heatmap is composed only of values that are lower or equal to zero since only white and blue squares are present. This means the best transpilation has equal or lower QVF values for all pairs of angles injected.

V. CONCLUSIONS AND FUTURE PROSPECTS

In this paper, we have highlighted the importance of the transpilation of a logical quantum circuit for its execution on a real quantum device. We have also underlined the susceptibility of this technology to stochastic radiation events, developing a fault injector working at the transpiled circuit level. For our analysis, we considered 4 well-known quantum circuits, acting

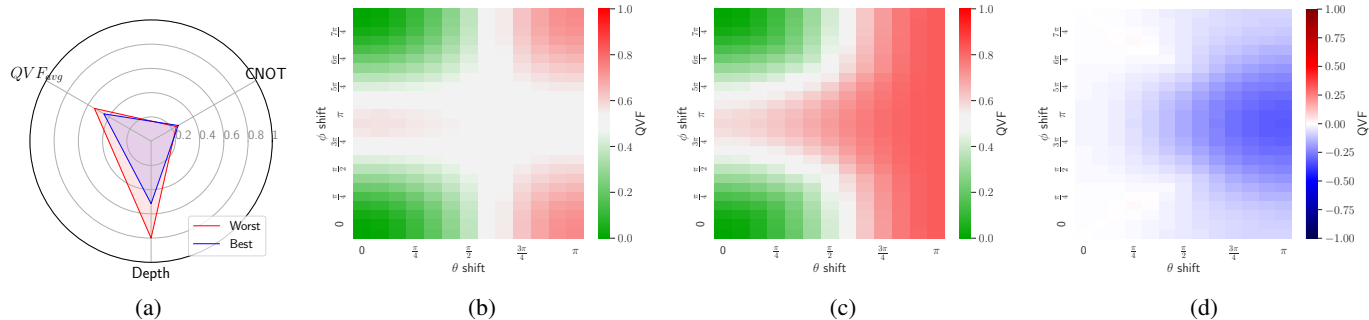


Fig. 4: Comparison of best and worst transpilation of BV4 algorithm on IBM Santiago architecture considering the QVF_{avg} value. (a) Radar graph for QVF_{avg} value, CNOT count and circuit depth for the two circuits. QVF heatmaps for best (b) and worst (c) transpilation. (d) Δ -heatmap between best and worst transpilation.

on increasing number of qubits, transpiled for two different architectures, with all 4 levels of optimization for the transpiler, and all possible initial logical-to-qubit mapping. Results have shown that, counterintuitively, the qiskit transpiler, set with the highest degree of optimization (level 3), produces a circuit that is shorter but more prone to transient fault errors than longer, less optimized circuits. Given the uneven reliability assessment of the transpiled circuits, we suggest including such analysis in the definition of the transpilation of choice for a given circuit, alongside the standard considered metrics, trading off between them. In fact, depending on the criticality of the calculation to be performed, a certain transpilation policy could prove more effective regarding transient-fault related issues. Future work will entail the broadening of current evaluation by analyzing the performance of other transpiling approaches, acting on more quantum algorithms, and other quantum physical architectures.

REFERENCES

- [1] S. Bravyi, M. Englbrecht, R. König, and N. Peard, “Correcting coherent errors with surface codes,” *npj Quantum Information*, vol. 4, oct 2018.
- [2] C. Chamberland, A. Kubica, T. J. Yoder, and G. Zhu, “Triangular color codes on trivalent graphs with flag qubits,” *New Journal of Physics*, vol. 22, p. 023019, feb 2020.
- [3] A. D. Corcoles, J. M. Chow, J. M. Gambetta, C. Rigetti, J. R. Rozen, G. A. Keefe, M. Beth Rothwell, M. B. Ketchen, and M. Steffen, “Protecting superconducting qubits from radiation,” *Applied Physics Letters*, vol. 99, no. 18, p. 181906, 2011.
- [4] L. Cardani, F. Valenti, N. Casali, G. Catelani, T. Charpentier, M. Clemenza, I. Colantoni, A. Cruciani, G. D’Imperio, L. Gironi, L. Grünhaupt, D. Gusenkova, F. Henriques, M. Lagoin, M. Martinez, G. Pettinari, C. Rusconi, O. Sander, C. Tomei, A. V. Ustinov, M. Weber, W. Wernsdorfer, M. Vignati, S. Pirro, and I. M. Pop, “Reducing the impact of radioactivity on quantum circuits in a deep-underground facility,” *Nature Communications*, vol. 12, no. 1, p. 2733, 2021.
- [5] J. M. Martinis, “Saving superconducting quantum processors from decay and correlated errors generated by gamma and cosmic rays,” *npj Quantum Information*, vol. 7, no. 1, p. 90, 2021.
- [6] Z. Chen, K. J. Satzinger, J. Atalaya, A. N. Korotkov, A. Dunsworth, D. Sank, C. Quintana, M. McEwen, R. Barends, P. V. Klimov, S. Hong, C. Jones, A. Petukhov, D. Kafri, S. Demura, B. Burkett, C. Gidney, A. G. Fowler, A. Paler, H. Putterman, I. Aleiner, F. Arute, K. Arya, R. Babbush, J. C. Bardin, A. Bengtsson, A. Bourassa, M. Broughton, B. B. Buckley, D. A. Buell, N. Bushnell, B. Chiaro, R. Collins, W. Courtney, A. R. Derk, D. Eppens, C. Erickson, E. Farhi, B. Foxen, M. Giustina, A. Greene, J. A. Gross, M. P. Harrigan, S. D. Harrington, J. Hilton, A. Ho, T. Huang, W. J. Huggins, L. B. Ioffe, S. V. Isakov, E. Jeffrey, Z. Jiang, K. Kechedzhi, S. Kim, A. Kitaev, F. Kostritsa, D. Landhuis, P. Laptev, E. Lucero, O. Martin, J. R. McClean, T. McCourt, X. Mi, K. C. Miao, M. Mohseni, S. Montazeri, W. Mruzckiewicz, J. Mutus, O. Naaman, M. Neeley, C. Neill, M. Newman, M. Y. Niu, T. E. O’Brien, A. Opremcak, E. Ostby, B. Pató, N. Redd, P. Roushan, N. C. Rubin, V. Shvarts, D. Strain, M. Szalay, M. D. Trevithick, B. Villalonga, T. White, Z. J. Yao, P. Yeh, J. Yoo, A. Zalcman, H. Neven, S. Boixo, V. Smelyanskiy, Y. Chen, A. Megrant, J. Kelly, and G. Q. AI, “Exponential suppression of bit or phase errors with cyclic error correction,” *Nature*, vol. 595, no. 7867, pp. 383–387, 2021.
- [7] L. Grünhaupt, N. Maleeva, S. T. Skacel, M. Calvo, F. Levy-Bertrand, A. V. Ustinov, H. Rotzinger, A. Monfardini, G. Catelani, and I. M. Pop, “Loss mechanisms and quasiparticle dynamics in superconducting microwave resonators made of thin-film granular aluminum,” *Phys. Rev. Lett.*, vol. 121, p. 117001, Sep 2018.
- [8] D. Oliveira, E. Auden, and P. Rech, “Atmospheric neutron-induced fault generation and propagation in quantum bits and quantum circuits,” *IEEE Transactions on Nuclear Science*, vol. 70, no. 4, pp. 345–353, 2023.
- [9] A. P. Vepsäläinen, A. H. Karamlou, J. L. Orrell, A. S. Dogra, B. Loer, F. Vasconcelos, D. K. Kim, A. J. Melville, B. M. Niedzielski, J. L. Yoder, S. Gustavsson, J. A. Formaggio, B. A. VanDevender, and W. D. Oliver, “Impact of ionizing radiation on superconducting qubit coherence,” *Nature*, vol. 584, no. 7822, pp. 551–556, 2020.
- [10] C. D. Wilen, S. Abdullah, N. A. Kurinsky, C. Stanford, L. Cardani, G. D’Imperio, C. Tomei, L. Faoro, L. B. Ioffe, C. H. Liu, A. Opremcak, B. G. Christensen, J. L. DuBois, and R. McDermott, “Correlated charge noise and relaxation errors in superconducting qubits,” *Nature*, vol. 594, no. 7863, pp. 369–373, 2021.
- [11] R. Barends, J. Wenner, M. Lenander, Y. Chen, R. C. Bialczak, J. Kelly, E. Lucero, P. O. Malley, M. Mariantoni, D. Sank, H. Wang, T. C. White, Y. Yin, J. Zhao, A. N. Cleland, J. M. Martinis, and J. J. A. Baselmans, “Minimizing quasiparticle generation from stray infrared light in superconducting quantum circuits,” *Applied Physics Letters*, vol. 99, no. 11, p. 113507, 2011.
- [12] D. Oliveira, E. Giusto, B. Baheri, Q. Guan, B. Montrucchio, and P. Rech, “A systematic methodology to compute the quantum vulnerability factors for quantum circuits,” *arXiv preprint arXiv:2111.07085*, 2021.
- [13] N. Casciola, E. Giusto, E. Dri, D. Oliveira, P. Rech, and B. Montrucchio, “Understanding the impact of cutting in quantum circuits reliability to transient faults,” in *2022 IEEE 28th International Symposium on On-Line Testing and Robust System Design (IOLTS)*, pp. 1–7, 2022.
- [14] R. Harper, S. T. Flammia, and J. J. Wallman, “Efficient learning of quantum noise,” *Nature Physics*, vol. 16, pp. 1184–1188, Dec 2020.
- [15] X. Hu, R. de Sousa, and S. D. Sarma, “Decoherence and Dephasing in Spin-Based Solid State Quantum Computers,” in *Foundations of Quantum Mechanics in the Light of New Technology ISQM-Tokyo ’01*, pp. 3–11, Oct. 2002.
- [16] D. Oliveira, E. Giusto, E. Dri, N. Casciola, B. Baheri, Q. Guan, B. Montrucchio, and P. Rech, “Qufi: a quantum fault injector to measure the reliability of qubits and quantum circuits,” in *2022 52nd Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, pp. 137–149, 2022.