A Day In the Life of a Quantum Error

Salonik Resch, Swamit Tannu,
Ulya R. Karpuzcu, and Moinuddin Qureshi

Abstract—When a fault occurs in a computational system, it is of interest what effects it has. If it is known what can go wrong, it may also be known how to mitigate or correct for it. While complex, it is possible to obtain this information with rigorous fault injection in classical systems. This is also desirable for quantum systems, unfortunately it is much more difficult. The exponential information content of quantum systems makes this likely impossible for larger systems. However, analyses on smaller systems may provide valuable insight into the behavior of larger systems. This letter analyzes the effect of error on different parts of quantum programs.

Index Terms—Quantum computing, fault injection

1 Introduction

In the future, quantum computers will have sufficient qubit counts and fidelity in order to perform quantum error correction, and hence will be capable of computations of arbitrary size. This is unfortunately not true for current machines, which have both limited qubit counts and fidelity. Although small in size and noisy, the existing quantum computers with fifty-plus qubits are capable of powerful computations. To leverage existing quantum computers, researchers are developing software techniques to mitigate hardware errors. To that end, recent proposals use machine-specific noise characteristics to increase the likelihood of measuring the correct output on Noisy Intermediate Scale Quantum (NISQ) computers [9], [11], [12]. However, there is a lack of studies on understanding how an error affects the later instructions and the output of quantum programs. Such analysis can be beneficial as it can allow the execution of of noise tolerant parts of a program on less reliable qubits. By exploiting the unique properties of specific programs, potentially larger applications can be run reliably.

Our approach to sensitivity analysis corresponds to the quantum equivalent of well-explored, classical statistical fault injection, where each experiment entails (i) injecting a fault in space or time at a specific point in computation (e.g., by corrupting temporal state, using a representative noise model to best capture the actual manifestation of a given type of fault); (ii) tracking the fault propagation all the way to the end results; and (iii) repeating this procedure for a statistically significant number of times to be able to draw a meaningful conclusion. Similar to this procedure, our experiments (in simulation) involve corruption of one qubit (to capture spatial noise tolerance) and one gate operation (to capture temporal noise tolerance) at a time, in performing a quantum algorithm, and simulating the respective experiment with density matrices to achieve a statistically significant result.

Of interest, is how errors on qubits temporally and spatially propagate when performing a quantum algorithm. If this can be properly characterized, such information can enable novel optimization opportunities in mapping an application on to quantum hardware, to best match the spatio-temporal variation in quantum noise tolerance of the underlying quantum substrate. The idea is mapping noise-sensitive computations to the least noisy components of the quantum computer (spatial scheduling). Additionally, performing gates earlier or later can reduce sensitivity to noise (temporal scheduling). Going even further, characterizing the noise tolerance of specific algorithms could aid in the design of application-specific quantum hardware [8].

In many ways, quantum computing significantly increases the complexity and difficulty of statistical fault injection. As accurate noise models are difficult to develop, simplifying assumptions must be made. Additionally, noise rates are much higher in quantum systems, which makes validation by physical experiment difficult, as errors are not manifested as single isolated events. This leaves as an open question if classically-inspired fault-injection analysis can be effectively used for quantum computers.

2 Background

2.1 Noise Types

Quantum noise is complex, and often simplifying assumptions are made to ease simulation and analysis [6]. Stochastic Pauli Noise is commonly used, which corresponds to the random insertion of extra X or Z gates. However, this is not a complete representation [3]. More realistic noise models capture Amplitude Damping (AD) and coherent errors. AD is the collapse of a qubit into its low energy state, and is the result of imperfect isolation from the environment. The likelihood of this occurrence can be calculated as a function of experimentally determined coherence times [13]. Notably this can cause loss of entanglement between qubits. For example, if two qubits are in the maximally entangled bell state, $|00\rangle + |11\rangle$, both qubits are in a superposition and their states depend on each other. If AD occurs on the second qubit, the state will transform to $|00\rangle + |10\rangle$, where the second qubit is now guaranteed to be in state $|0\rangle$. Additionally, the state of the second qubit no longer depends on the first qubit, and they are no longer entangled.

Coherent noise can be slight over- or under-rotation of qubits due to imperfect control hardware [1], [3], [5], [7], [14], and is the dominant noise source in modern quantum computing systems [6]. Unlike AD, this noise source is unitary. Hence it does not destroy the quantum state, but will move it into a different, and undesired, coherent state [2]. A phase error (Z-rotation) will change the phase between the $|0\rangle$ and $|1\rangle$ states. A Z-rotation by angle θ will transform the quantum state from $\alpha|0\rangle + \beta|1\rangle$ to $\alpha|0\rangle + e^{i\theta}\beta|1\rangle$. A rotation by π will effectively be a Z gate.

We model noise by causing AD with 100 percent chance of collapse and by performing Z-rotations by π.[1] This is not to suggest that physical quantum noise manifests in this manner. Rather, we test the spatial and temporal susceptibility of quantum programs to noise of this nature.

In practice, it may be beneficial to use machine-specific noise models. The approach developed here still applies, and can be made machine-specific by simply swapping out the specific noise model. Such machine-specific noise models will change over time and across re-calibration cycles. Regardless, AD and coherent rotations are representative of real-world processes and can provide insight into quantum program sensitivity.

1. For the case studies used in this work, different AD probabilities and Z-rotation angles affect the severity of the impact, but not do not change the nature of the impact.

- *Salonik Resch and Ulya R. Karpuzcu are with the Department of Electrical and Computer Engineering, University of Minnesota, Twin Cities, Minneapolis, MN 55455 USA. E-mail: {resc0059, ukarpuzc}@umn.edu.*
- *Swamit Tannu is with the Computer Science Department, University of Wisconsin-Madison, Madison, WI 53706 USA. E-mail: swamit@gatech.edu.*
- *Moinuddin Qureshi is with the School of Computer Science, Georgia Institute of Technology, Atlanta, GA 30332 USA. E-mail: moin@ece.gatech.edu.*

Manuscript received 22 Oct. 2020; revised 11 Dec. 2020; accepted 11 Dec. 2020. Date of publication 17 Dec. 2020; date of current version 8 Jan. 2021.
(Corresponding author: Salonik Resch.)
Digital Object Identifier no. 10.1109/LCA.2020.3045628

2.2 Metrics

Evaluating the effects of quantum noise also depends on the metrics of interest. *Process Fidelity* [4] is an intuitive metric as this effectively measures the *distance* between the actual quantum state and the ideal. However, depending on context, this could be misleading. Phase errors will cause degradation of the process fidelity, but will not directly affect measurement outcomes (in the Z-basis). Hence, this may be overly pessimistic. A similar metric, the *Hellinger Fidelity*, focuses exclusively on measurement probabilities. It is a measure of how similar samples from a probability distribution are. This can easily be determined by performing a quantum program multiple times and comparing the measured results to the expected. This metric is useful as it roughly corresponds to how likely one is to get results from the correct set of possible answers. However, it has its own drawbacks. The Hellinger fidelity can vary widely and report pessimistic results if there are only a few possible measurement outcomes in the ideal case. For example, say the ideal case has a 100 percent chance of measuring $|0010010\rangle$. If there is an X error causing a flip on the last qubit, there will be a 100 percent chance of measuring $|0010011\rangle$. While these outputs are highly similar, and it may be possible to find the correct answer with post processing, the Hellinger fidelity will be 0 due to no overlap in the output measurements.

Probability of Successful Trial (PST) is another widely used metric. It is the probability of measuring the correct output. In some cases it may also be overly pessimistic. In the above example, the PST would also be 0. PST finds greater applicability if there is no classical means to extrapolate the correct output from a noisy measurement.

3 Case Studies

Different quantum circuits (programs) can vary in how they are affected by noise. We use circuits with starkly different characteristics to understand how faults evolve in time and space. Moreover, we use quantum benchmarks, which can be parametrized and scaled without significantly changing the workload structure.

Quantum Adder (ADDER). Quantum adder circuits perform addition on two quantum states. We use quantum adders because (1) Output of quantum adder is trivial to verify using conventional computers (2) Adders have a rich structure that uses Toffoli gates, which are sensitive to both amplitude and phase errors. Moreover, quantum adders are used in Shor's prime factorization, and they are a basic building block of many quantum arithmetic functions.

Bernstein-Vazirani (BV). is a quantum algorithm that demonstrates quantum speedup by querying the oracle only once to find the encoded secrete. On execution, BV outputs a binary string corresponding to the secret key. We use an oracle function that uses CNOTs and leverage phase kickback. The sensitivity to phase errors but tolerance against bit-flip errors make BV an interesting case study.

Quantum Approximate Optimization Algorithm. The Quantum Approximation Optimization Algorithm(QAOA) can solve combinatorial optimization problems such as MAX-CUT. Certain characteristics of QAOA make it tolerant to noise. For example, QAOA can be run in a variational mode where it is run multiple times with classical parameter updates in order to achieve an optimal output. One of the unique structural traits of QAOA is that all qubits have an equal number of connections (participate in 2-qubit gates), and the circuit is typically chosen to have a shallow depth. We use QAOA implementation that requires only nearest-neighbor connections. This greatly eases execution on a physical machine as it does not require any qubit swap operations. The input to QAOA is the $|00..0\rangle$ state.

Quantum Fourier Transform (QFT). The Quantum Fourier Transform performs the Discrete-Time Fourier Transform on the amplitudes of a quantum state. It is the core of Shor's algorithm for factorization [10]. The QFT circuit is asymmetrical, for an n-qubit QFT, the least-significant qubit is the target of $n-1$ controlled-rotations and the most-significant qubit is not the target of any controlled-rotations. However, it is the control of $n-1$ such operations.

The input to the QFT circuit is typically a highly entangled state. In the case of Shor's algorithm, the amplitudes are the outputs of the modular exponentiation function. We use an entangled state where the amplitudes form a periodic function, representative of input to the QFT in Shor's algorithm.

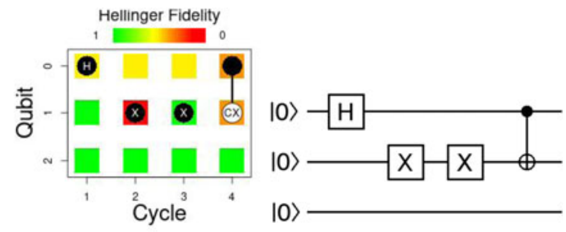

(a) Heatmap Circuit (b) Circuit

Fig. 1. Heat map circuit alongside equivalent circuit diagram. Heatmap shows hellinger fidelity of output if AD (with 100 percent chance of collapse) were to occur at each location.

4 Fault-Injection Analysis

To determine the sensitivity of quantum circuits we need to insert noise at specific locations and times. This can be done by inserting a noise event on a single qubit during a single cycle of the circuit. While this noise event will have an immediate impact on the quality of the quantum state, and analyzing this immediate impact can build intuition, it is of more interest on how it will impact the state once the entire circuit has been executed. This is for two reasons. One is that only effects visible at the output will degrade the measurement results. Hence, if a noise event induces an error which becomes obscured, we don't care. The second reason is that, unless performing quantum error correction, we will not be able to detect, mitigate, or correct the error in real time during execution. The effects remain invisible until a measurement is performed, which is at the end of the circuit for these near term applications.

For our simulations, we insert a single noise event on a single qubit in a single cycle and view the impact on the output of the circuit. We repeat this for every qubit and every cycle in each circuit. To accentuate the impact of the noise event, all other operations are noiseless, including gates, idling, and measurement. To visualize the impact, we plot the circuit over the heatmap of the output quality. This way, it is easy to spot the sensitive regions of each circuit.

An example is shown in Fig. 1. The figure follows the same conventions as quantum circuits, the qubits are stacked vertically and the time flows from left to right. Gates are represented by black circles. White circles are the targets of 2-qubit gates. For clarity, we omit the traditional horizontal lines tracing a qubit in traditional circuit diagrams. The heatmap shows the hellinger fidelity of the *output* if a *single* error occurred at that location in space and time. Hence, each noise location represents an individual experiment, where noise occurred at that location and nowhere else. The color on a gate represents an error which occurs immediately after the gate. The heatmap scales from perfect output, green meaning the hellinger fidelity is 1, to yellow at 0.5, and finally red at 0.

The ideal output of the circuit in Fig. 1 is $|000\rangle + |110\rangle$. If AD occurs on qubit 0 on cycle 1, after the H gate, the output will be $|000\rangle$. This has considerable overlap with the desired output, and the hellinger fidelity is 0.46 and the location on the heatmap circuit is indicated as such. If AD occurs after the first X gate on qubit 1, the final output will be $|010\rangle + |100\rangle$. This has no overlap with

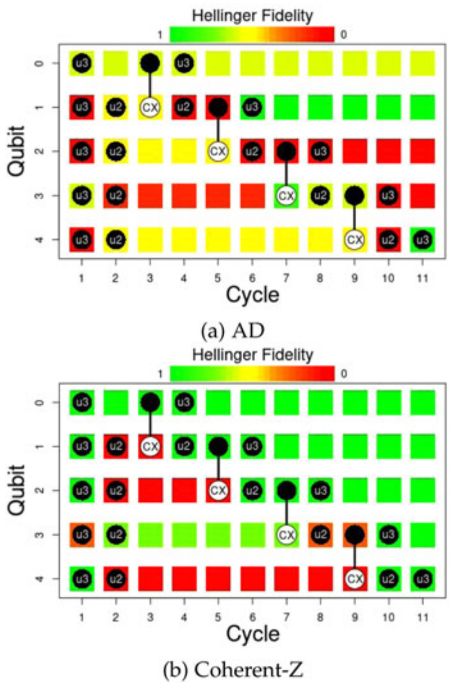

(a) AD
(b) Coherent-Z

Fig. 2. QAOA.

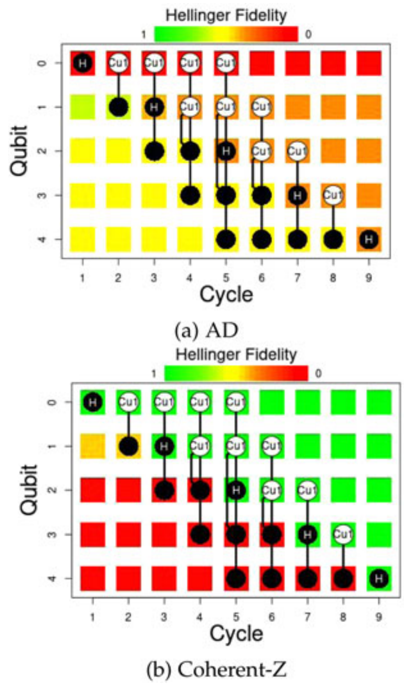

(a) AD
(b) Coherent-Z

Fig. 3. QFT.

ideal output, hence the hellinger fidelity is 0. If AD occurs on qubit 1 after the final CX (CNOT) gate, this will cause a loss of entanglement between qubits 0 and 1. The output state will be $|000\rangle + |100\rangle$, and the hellinger fidelity is 0.29. As qubit 2 remains at $|0\rangle$ the entire time, and is not entangled with the other qubits, AD on qubit 2 has no effect on output.

For noise events, we use both Amplitude Damping, which has a 100 percent chance of collapse, and coherent Z-rotations by π. For metrics, we show the hellinger fidelity.

As can be seen, the effect of noise highly depends on the circuit. For circuits which create highly entangled output states, AD tends to have more destructive results later in the program, whereas coherent Z-rotation tends to have more impact earlier in the program. This is intuitive, as AD causes a loss of information and entanglement. The further along the program is, the more entangled the quantum state is and hence a greater loss of information with the non-unitary collapse in AD. Coherent noise will not cause loss of entanglement, and as it is unitary (reversible), does not destroy information in the quantum state. Effectively, it changes the program that is being performed. The earlier this occurs in the circuit, the more opportunity it has to spread to other qubits and further corrupt the state. This trend is observable even in these small scale simulations, most notably for the QFT in Fig. 3. It is expected to have an even more significant impact on larger circuits, where the increased circuit depth provides more opportunities for spreading error and entangled states exist for longer periods. Notably phase errors do not cause a change in the measurement outcomes if they occur immediately before measurement. Hence, Z-rotation errors show no degradation if occurring after the last gate on a qubit.

If spatial sensitivity is observed (some qubits/gates in a cycle have worse errors than others), the sensitive gates/qubits can be scheduled on a machine's most reliable physical qubits (spatial gate scheduling). This will only help if these qubits can be mapped correctly to the physical machine without introducing additional swaps, or if the swaps introduce less error than is saved. For example, qubit 0 of QAOA (Fig. 2) and qubit 2 of the adder (Fig. 4) show low sensitivity to both AD and coherent noise, relative to other qubits in the circuit. Hence, some robustness may be gained by mapping these logical qubits to less reliable physical qubits in physical experiment. This technique can be added on top of well-studied methods of mapping quantum circuits to the most reliable physical qubits [11]. Effectively, in addition to using the overall most reliable qubits, the reliability of the physical qubits used are matched to sensitivity of the logical circuit.

Temporal gate scheduling also plays an important role. Exemplified by BV (Fig. 5), performing gates as late as possible can reduce sensitivity. If all qubits are initiated and put into a superposition with H gates early, this leaves them vulnerable to both AD and coherent errors for extended periods of time. In this case, scheduling the initial H gates later clearly provides superior performance, as the qubits are more resilient in the $|0\rangle$ state. However, this is not universally true. For QAOA (Fig. 2), performing the second u2 gate on qubit 3 later (immediately prior to following CX

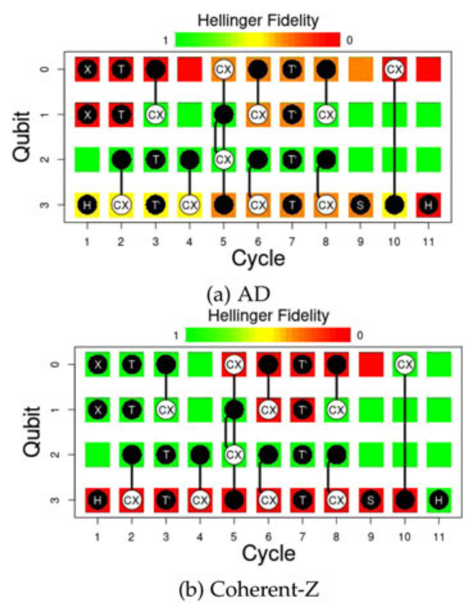

(a) AD
(b) Coherent-Z

Fig. 4. Adder.

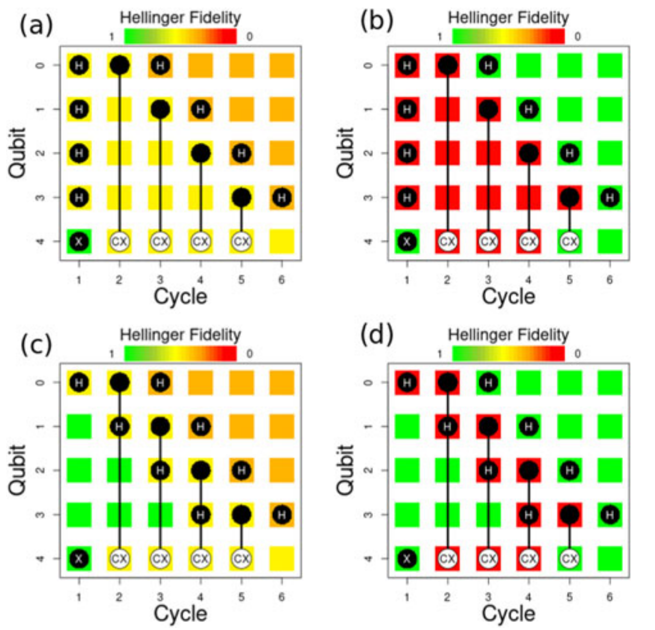

Fig. 5. *Eager Bernstein-Vazirani:* (a) with AD noise (b) Coherent-Z noise, *Late Bernstein-Vazirani:* (c) with AD noise (b) Coherent-Z noise.

TABLE 1
PST of BV-5 on IBM Hardware for Eager and Delayed Scheduling

Schedule	Valencia	Casablanca	Bagota
Eager	0.61	0.74	0.45
Delayed	0.73	0.82	0.65

gate) would reduce its susceptibility to AD noise. However, this would have the opposite effect if the noise source is coherent, where the delay in u2 would leave the state more susceptible. This condition is also seen for qubit 4, but with the reverse conclusion. Hence, gate scheduling not only critically depends on the circuit but the noise source. Both must be considered for optimal scheduling.

The high-level insights from fault-injection studies hold on a real hardware. For instance, we observe substantial reliability improvements by using a delayed schedule, similar to what we observe for the fault-injection studies with coherent-Z and AD noise in the Fig. 5. Table 1 reports a Probability of Successful Trial (PST) for five qubit Bernstein-Vazirani(BV) for two scheduling policies - Eager and Delayed when executed on IBM quantum computers with 5 to 7 qubits. Note that both instance of circuits with eager and delayed schedule have identical gate count and physical to logical mapping, the only difference is the instruction order.

5 Conclusion

The well studied classical error injection process is also applicable to quantum computers. Finding critical and noise-tolerant regions will enable an application to be adapted to both mitigate the effect of quantum noise and save precious resources. The errors that arise from noise are highly spatially and temporally dependent and are determined by the circuit that is being performed. Unfortunately, quantum computers have additional complications. The errors are also highly depended on the nature of the quantum noise, which is not well understood. Assumptions made about the noise become critical to finding meaningful results and techniques which will remain valid in experimentation.

Acknowledgments

This work was supported in part by the Office of the Vice Chancellor for Research and Graduate Education at the University of Wisconsin–Madison with funding from the Wisconsin Alumni Research Foundation.

References

[1] H. Ball *et al.*, "Effect of noise correlations on randomized benchmarking," *Physical Rev. A*, vol. 93, 2016, Art. no. 022303.
[2] J. P. Barnes *et al.*, "Quantum error-correction failure distributions: Comparison of coherent and stochastic error models," *Physical Rev. A*, vol. 95, 2017, Art. no. 062338.
[3] S. Bravyi *et al.*, "Correcting coherent errors with surface codes," *npj Quantum Inf.*, vol. 4, 2018, Art. no. 55.
[4] S. T. Flammia *et al.*, "Direct fidelity estimation from few pauli measurements," *Physical Rev. Lett.*, vol. 106, 2011, Art. no. 230501.
[5] D. Greenbaum *et al.*, "Modeling coherent errors in quantum error correction," *Quantum Sci. Technol.*, vol. 3, 2017, Art. no. 015007.
[6] J. J. Wallman *et al.*, "Noise tailoring for scalable quantum computation via randomized compiling," *Physical Rev. A*, vol. 94, 2016, Art. no. 052325.
[7] R. Kueng *et al.*, "Comparing experiments to the fault-tolerance threshold," *Physical Rev. Lett.*, vol. 117, 2016, Art. no. 170502.
[8] G. Li *et al.*, "Towards efficient superconducting quantum processor architecture design," in *Proc. 25th Int. Conf. Archit. Support Program. Lang. Operating Syst.*, 2020, pp. 1031–1045.
[9] P. Murali *et al.*, "Software mitigation of crosstalk on noisy intermediate-scale quantum computers," in *Proc. 25th Int. Conf. Archit. Support Program. Lang. Operating Syst.*, 2020, pp. 1001–1016.
[10] P. W. Shor, "Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer," *SIAM J. Comput.*, vol. 26, pp. 1484–1509, 1999.
[11] S. S. Tannu *et al.*, "Not all qubits are created equal: A case for variability-aware policies for NISQ-era quantum computers," in *Proc. 24th Int. Conf. Archit. Support Program. Lang. Operating Syst.*, 2019, pp. 987–999.
[12] S. S. Tannu and M. K. Qureshi, "Mitigating measurement errors in quantum computers by exploiting state-dependent bias," in *Proc. 52nd Annu. IEEE/ACM Int. Symp. Microarchit.*, 2019, pp. 279–290.
[13] Y. Tomita and K. M. Svore, "Low-distance surface codes under realistic quantum noise," *Phys. Rev. A*, vol. 90, p. 062320, Dec. 2014. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevA.90.062320
[14] J. Wallman *et al.*, "Estimating the coherence of noise," *New J. Phys.*, vol. 17, 2015, Art. no. 113020.