

Quancorde: Boosting fidelity with Quantum Canary Ordered Diverse Ensembles

Gokul Subramanian Ravi^{*1}, Jonathan M. Baker¹, Kaitlin N. Smith^{1,2}, Nathan Earnest³,
Ali Javadi-Abhari³, and Frederic T. Chong^{1,2}

¹University of Chicago

²Super.tech (a division of ColdQuanta)

³IBM Quantum, IBM T. J. Watson Research Center

Abstract—On today’s noisy imperfect quantum devices, execution fidelity tends to collapse dramatically for most applications beyond a handful of qubits. It is therefore imperative to employ novel techniques that can boost quantum fidelity in new ways.

This paper aims to boost quantum fidelity with Clifford canary circuits by proposing *Quancorde: Quantum Canary Ordered Diverse Ensembles*, a fundamentally new approach to identifying the correct outcomes of extremely low-fidelity quantum applications. It is based on the key idea of diversity in quantum devices - variations in noise sources, make each (portion of a) device unique, and therefore, their impact on an application’s fidelity, also unique.

Quancorde utilizes Clifford canary circuits (which are classically simulable, but also resemble the target application structure and thus suffer similar structural noise impact) to order a diverse ensemble of devices or qubits/mappings approximately along the direction of increasing fidelity of the target application. Quancorde then estimates the correlation of the ensemble-wide probabilities of each output string of the application, with the canary ensemble ordering, and uses this correlation to weight the application’s noisy probability distribution. The correct application outcomes are expected to have higher correlation with the canary ensemble order, and thus their probabilities are boosted in this process.

Doing so, Quancorde improves the fidelity of evaluated quantum applications by a mean of 8.9x/4.2x (wrt. different baselines) and up to a maximum of 34x.

I. INTRODUCTION

Quantum computing is a revolutionary computational model that takes advantage of quantum mechanical phenomena to solve intractable problems. Quantum computers can potentially leverage superposition, interference, and entanglement to give significant computing advantage in chemistry [22], optimization [27], machine learning [6] etc.

In near-term quantum computing, we expect to work with machines which comprise 100-1000s of imperfect qubits [35]. These machines suffer from high error rates in the form of state preparation and measurement (SPAM) errors, gate errors, qubit decoherence, crosstalk, etc. While, in the near future, it is clear that we will be unable to execute large-scale quantum algorithms like Shor’s Factoring [41] and Grover Search [20], which would require error correction comprised of millions




Fig. 1. Quancorde intuition for a trivial example: A target quantum circuit is executed on a diverse ensemble (eg., on different machines) as shown by the different colors. Then, for each output bitstring produced in noisy circuit execution (from ‘00’ - ‘11’), all machines are ordered by that string’s occurrence probability - these are the first 4 columns. Next, Quancorde runs a ‘canary’ circuit (which resembles the target circuit structure and is classically simulable) on the ensemble to obtain a canary ordering, i.e., the machine ordering that increases the occurrence of the correct canary output. The canary ordering is expected to be a close match to the ordering created by the correct outcome for the original circuit - this is because the two circuits have the same structure and suffer fairly similar noise effects. Thus, the output string from the original circuit that produces a machine ordering closely correlated with the canary ordering is likely to be the correct original outcome. This is ‘11’ in the figure. A full-fledged design illustration is shown in Fig.2.

of qubits to create fault-tolerant quantum systems [32], it is important to uncover some potential for quantum advantage in the near-term.

Classical computing today is able to perfectly simulate quantum systems up to around 50 qubits (with the help of supercomputers, if required) [7], [14], [21], [26], [49]. Thus, advancing today’s quantum computing to the doorstep of quantum advantage would require us to solve quantum applications of around 50 qubits with reasonable fidelity¹. However, on today’s quantum devices, execution fidelity tends to collapse dramatically for most circuits beyond a handful of qubits.

Multiple error mitigation techniques have been explored recently. These techniques reduce the effect of noise on circuit execution on the quantum machine. Several promising strate-

¹fidelity: the probability of finding the correct outcome or set of outcomes from the output distribution produced by a noisy quantum circuit.

* Correspondence: gravi@uchicago.edu.




Fig. 2. Quancorde overview. (1) A canary circuit for the target application is constructed by replacing the non-Clifford gates in the target (orange) with the nearest Clifford gates (green). (2) The target circuit is executed on the diverse noisy quantum ensemble R1-R5 (e.g. different machines or qubits/mappings). (3) In parallel, the canary is also executed on the same noisy ensemble. (4) The correct output of the canary is obtained by running it ideally (noise-free) on a classical machine - possible since Clifford circuits are efficiently classically simulable. (5) Since the correct canary outcome is known, the ensemble is ordered based on the noisy execution fidelity of the canary - baseline - the correct answer '11' has low probability. (6) The noisy distribution on any machine is selected as the baseline. (7) Ensemble orderings are produced for the different noisy outputs (of non-negligible probability) of the original circuit and are compared with the canary ordering (like in Fig.1) to estimate a correlation value for each output 0 '11' has a high correlation. (8) The correlations are then used to weight the baseline distribution to produce a new distribution in which the '11' probability is boosted to become the winner.

gies have recently emerged [4], [8], [9], [12], [15], [15], [18], [24], [29], [30], [34], [39], [42]–[44], [46]–[48], [50], [51]. However, the resulting execution fidelity is still minuscule for most circuits beyond 10 qubits, let alone real-world use cases. Thus, while we must continue to innovate across the hardware and software stack to improve the fidelity we can get out of any given quantum device, it is also imperative to ponder entirely different techniques that are able to boost quantum application fidelity in new ways, beyond the fundamental limitations of today's noisy quantum devices. Only then can we advance quantum frontiers in the near future.

This work proposes one such fundamentally new approach to identifying the correct outcomes of extremely low fidelity quantum applications. It is based on the key idea of diversity in quantum devices - no two quantum devices are identical. Today's device errors stem from multiple noise sources such as the imperfect classical control of the device, thermal fluctuations, destructive qubit coupling, imperfect insulation of the qubits, quasi-particles, and other external stimuli [10], [23], [25], [28], [40]. Because current fabrication techniques lack the precision to make homogeneous batches of quantum devices, the noise properties are distinct for each device. Furthermore, the dynamic nature of quantum systems causes these noise sources to suffer from spatial and temporal variation. Thus, each (portion of a) device is unique, and its impact on an application's execution / fidelity, is also unique.

Simply put, a device which is known to be 'better' for the target application would produce higher application fidelity (compared to other devices). In other words, it would produce a higher probability of occurrence of the correct application outcomes in its output distribution. Extending this notion

further, every output string that an application might produce in noisy execution, will have a unique ordering of devices that increases the string's probability of occurrence from low to high. If one of these orderings is somehow known to be the 'correct' device ordering that improves the fidelity of the application, then the bitstrings that produce this (or similar) device orders would most likely be among the correct outcomes of the application. Of course, knowing this correct ordering is not trivial, but in this work we show that this is very feasible. Note: this diversity could also be different qubit mappings within a single device.

We propose *Quancorde: Boosting fidelity with Quantum Canary Ordered Diverse Ensembles*. Quancorde exploits diversity and ordering, as expressed above, to identify the correct outcome of a target application. It is introduced in Fig.1. In Quancorde, a target quantum circuit is executed on a diverse ensemble of resources, (eg., on different machines) as shown by the different colored blocks in the figure. Then, for each output bitstring of the circuit (of non-negligible occurrence), machines are ordered by the string's probability of occurrence. Each bitstring can produce a unique ordering as shown for 4 different output strings ('00', '01', '10' and '11') for a 2-qubit circuit in the figure. Now, the critical step is to identify the (nearly) correct ensemble ordering that improves the application fidelity.

To do so, Quancorde constructs a quantum 'canary' circuit (inspired in name by prior classical canary circuits [16]). The canary circuit is important in two ways: ① First, the canary circuit is made up of only Clifford gates and is constructed by replacing all non-Clifford gates in the target circuit with the 'nearest' Clifford gates. A circuit made up

of only Clifford gates is efficiently classically simulable [19], and thus the correct output of the canary circuit is obtained via ideal classical simulation. ② Second, the canary circuit maintains the exact device-mapped circuit structure of the original circuit - it has the same circuit critical depth/paths, the same 2-qubit CNOT gates, and the same measurement bits. Thus, it suffers from the same noise sources, and their relative impact on its circuit fidelity is similar to that on the original circuit, even if the quantum states explored by the two circuits are different. Next, Quancorde runs this canary circuit on the ensemble to obtain a canary ordering, i.e. the ensemble ordering for the correct canary output (which is known from ①) - this is shown to the right of the figure. The canary ordering is expected to be a close match to the ordering created by the correct output for the original circuit (as discussed in ②). Next, the ensemble order produced by each noisy output string of the application is correlated against the canary ordering. A correlation distribution is obtained in which the correct outcome of the application is likely to have a strong correlation to the canary order (again, from ②). In the figure, this corresponds to the ‘11’ string, whose ordering is almost identical to the canary ordering. Finally, this correlation distribution is used to weight the original application noisy distribution. Since the likely correct outcomes are weighted higher, this tremendously boosts application fidelity. More the diversity of the ensemble, greater is the uniqueness of the ensemble order. This creates a more well defined correlation distribution, leading to a larger boost to application fidelity via the Quancorde approach. A step-by-step breakdown of Quancorde is shown in Fig.2.

Overall, we make the following contributions:

- ① We propose a fundamentally new approach to identifying the correct outcomes of extremely low fidelity quantum applications by exploiting quantum device diversity.
- ② We propose Clifford canary circuits to order a diverse ensemble of devices or qubits/mappings along the direction of increasing fidelity of the target application.
- ③ We then identify the correlation of the application outputs’ ensemble orderings with the canary order, and use this to weight the application output distribution, resulting in boosted application fidelity.
- ④ The above ideas are presented as Quancorde, which improves the fidelity of our evaluated quantum applications (with baseline fidelity as low as 0.1%) by a mean of 8.9x/4.2x (wrt. different baselines) and a maximum of 34x.
- ⑤ Importantly, Quancorde can surpass fundamental limitations of any particular noisy quantum device and thus improve application fidelity beyond the capability of any single device.

II. BACKGROUND AND MOTIVATION

A. Fidelity in the noisy quantum era

Today’s quantum devices are error-prone and up to around 100 qubits in size [35]. Multiple error mitigation strategies have been proposed to correct different forms of quantum errors. These include, but are not limited to, noise-aware compilation [29], [47], scheduling for crosstalk [15], [30],




Fig. 3. Sorted probability of occurrence for output strings produced by a 9-bit circuit executed on IBMQ Montreal. The correct outcome is indicated in red.




Fig. 4. Diversity in IBM Quantum machine characteristics, obtained from [11] on 04/17/2022. T1 times across the qubits on IBM Quantum Hanoi and Montreal respectively. The bottom 2 figures show CNOT and Readout errors for each qubit on Montreal.

1Q gate scheduling in idle windows [42], dynamical decoupling [34], [43], [50], zero-noise extrapolation [18], [24], [48], readout error mitigation [9], [46], exploiting quantum reversibility [33], [42] and many more [4], [8], [12], [18], [24], [39], [44], [48], [51]. In addition, some of these can be used in conjunction to achieve better fidelity [38]. However, the resulting fidelity is still insufficient for most circuits beyond 10 qubits, with the correct circuit outcome(s) often not being among the most dominant in the output distribution.

Fig.3 shows the sorted probability of the output bitstrings obtained from a 9-qubit circuit executed on IBM Quantum Montreal with the highest noise-aware compiler optimizations. The correct outcome is shown in red. Not only is the probability of the correct output (i.e., fidelity) very low at 4%, it is also 2x lower than the probability of the highest occurring output. Furthermore, the output distribution has nearly 40 bitstrings with occurrence probabilities higher than that of the correct outcome. It is evident that for circuits such as this, and those of greater complexity, today’s machines will be very limited in the fidelity they achieve. While error mitigation techniques improve fidelity and will continue to do so as more novel solutions are proposed, it is clear that noisy quantum machines will be fundamentally limited, and we require entirely novel approaches to identify the correct outcomes of applications of real-world criticality. In this work, we turn to diversity.

B. Diversity in Quantum Devices

Even if QCs are manufactured in a highly controlled setting, unavoidable variation results in intrinsic properties that impact performance. Fig.4 shows data on IBM Quantum machine




Fig. 5. Simulating Clifford circuits on the IBMQ Stabilizer Simulator [1] run on a laptop. The linear scaling with number of gates is clearly visible. Circuits with 300 qubits / 50,000 gates can be simulated in 30 seconds.

characteristics obtained from [1] on 04/17/2022. The figures show T1 coherence times across the 27 qubits on IBM Quantum Hanoi and Montreal, respectively. These coherence times are indicative of the circuit durations at which amplitude damping occurs. In the data shown, the average coherence time of Hanoi is lower than that of Montreal, indicating a higher likelihood of damping. Furthermore, some qubits are much worse than others. On Montreal, a 4-qubit circuit mapped to qubits (2, 3, 5, 6) would decohere faster than one mapped to (9, 10, 12, 13). Thus, clearly the circuit fidelity (especially for deeper circuits) could be very different across these two mappings and devices. Similar diversity is present in other forms of error such as gate errors and readout errors [37].

From the above discussion, it is intuitive that the same application executed on different machines or on different qubits/mappings within a machine, will experience different noise characteristics and therefore produce different output probability distributions. Finally, it is important to recognize that for circuits of even reasonable complexity, the fidelity impact cannot be directly inferred simply by observing error rates and decoherence times. We discuss this further in Fig.10 and Section IV-B, but this is intuitive, since if this were possible, then noise-aware quantum simulators would be perfect at mimicking real device fidelity (which they are not).

C. Clifford Circuits

Classical simulation of quantum problems usually requires exponential resources (otherwise the need for quantum computers is obviated). Even using high-performance supercomputers, the simulation is restricted to around 50 qubits [7], [14], [21], [26], [49]. An exception to the above is the classical simulation of the Clifford space which is a subset of the total quantum Hilbert space. Circuits made up of only Clifford operations can be exactly simulated in polynomial time (almost linear in the number of gates and qubits) as stated by the Gottesman-Knill theorem [19].

We quantitatively showcase this in Fig.5. The figure shows simulation time for purely Clifford circuits of up to 50,000 gates and 300 qubits. These are simulated on the IBM Stabilizer Simulator [1] run on a laptop computer. Clearly the simulation time is seen to scale fairly linearly in the number of gates. Further, even the largest circuits are simulated in roughly 30 seconds.

While the Clifford group operations do not provide a universal set of quantum gates, they have many benefits. In this work, Clifford circuits are used to construct canary circuits that closely resemble the structure of the target application. By being classically simulable, the correct outcomes of the canary circuits can be estimated on a classical computer in an efficient manner. This is then utilized towards ordering the diverse ensemble.

III. LEVERAGING ORDERED ENSEMBLES

In this section, we discuss how ensemble ordering enables the boosting of noisy application fidelity. This section assumes that an ensemble order that is well correlated to the fidelity of the target application is known; the practical implementation of this, with canary circuits, is discussed in Section IV.

Fig.6.a shows a 2-qubit circuit that exploits phase kickback (integral to applications like Bernstein-Vazirani [5]) to produce ‘11’ as the correct outcome. When executing on real machines, such a circuit would suffer from noise and thus, the output distribution produced is a set of many bitstrings, which in this case are ‘00’, ‘01’, ‘10’ and ‘11’. Our goal is to identify that ‘11’ is the correct outcome. Note that in the presence of significant noise, ‘11’ is not necessarily the most probable output (similar to Fig.3), so simply capturing the output with highest probability is often insufficient.

In this experiment, we execute this 2-qubit circuit on a diverse ensemble of 20 diverse simulated ‘machines’ - 20 simulations with varying error characteristics. For this experiment, we assume that the ensemble ordering that is well correlated with the circuit’s fidelity is known (i.e., we have an intuition for the ordering of machines that improves the fidelity of the application, even if we don’t know the correct outcome of the application). All errors increase uniformly from machines 1 to 20, so we assume that the reverse ordering will be well correlated with application fidelity.

Now, recall that we still do not know the correct outcome of our target circuit. Next, we plot the probability of occurrence of each output string against the error-ordered machine ensemble. This is shown in Fig.6.b. Clearly the probability of the ‘11’ outcome (green) falls almost monotonically as we go from best (1) to worst (20) machines, showing high (nearly perfect) positive correlation with the ‘correct’ machine order. On the other hand, the ‘00’ outcome probability (blue) increases as machine quality worsens (negative correlation) and the ‘10’ and ‘01’ outcomes (yellow / plum) have no discernible correlation.

Next, we use these correlations to weight the original application’s output probability distribution (this weighting is similar to Bayesian Reconstruction [13]). Outputs with high positive correlation are boosted, those with low positive correlation are diminished, and those with negative correlation are simply removed. This can be done for the application’s output distribution on any machine or every machine or a known fairly good machine. We show this for every machine in Fig.6.c. Clearly, the ‘11’ outcome probability is tremendously improved due to its high positive correlation. Now, from this




Fig. 6. Leveraging an ordered diverse ensemble to boost application fidelity. (a) 2-qubit circuit with a ‘11’ correct output. (b) The 2-qubit circuit is executed on 20 simulated machines with increasing noise, and the output bitstring probabilities are plotted. The probability of the correct bitstring ‘11’ decreases in high positive correlation with decreasing machine quality, whereas other bitstrings have low / negative correlation. (c) The correlations are used to weight the output probabilities. The correct bitstring is boosted due to high correlation, while others are diminished. Thus, it stands out as the clear estimated outcome. (d) Flow chart of the Quancorde proposal assuming a given canary-ordered ensemble.




Fig. 7. Generating Clifford canary circuits. The top circuit shows a target 4-qubit circuit after mapping to an IBM Quantum device. The bottom circuit shows a Clifford canary generated for the target circuit.

figure, it is easy to identify that ‘11’ is the correct application outcome since it clearly stands apart from other outputs. The resulting fidelity boost is significant for any machine in the ensemble, and clearly showcases the capability of Quancorde. A flow chart of the Quancorde proposal, generalizing the steps to exploit a diverse ensemble is shown in Fig.6.d and Fig.2.

Greater diversity and size of the ensemble are likely to imply more comprehensible identification of the correct outcome or a smaller subset of outcomes. Note that the ensemble does not have to be unique devices, it could also be different sets of qubits or different mapping of qubits with a single device or both - thus ensemble size can be increased.

IV. CANARIES FOR ENSEMBLE ORDERING

In Section III, we discussed how knowing an ensemble ordering that is well correlated with the fidelity of the target application can help identify the correct application outcome. While a ‘good’ ensemble ordering is fairly trivial to estimate for very simple circuits (roughly proportional to a product of different noise impacts), this is not the case for more complex circuits, for which fidelity is dependent on more complex interactions among the qubits and among the different noise sources. Thus, to help find an appropriate ensemble ordering we propose ‘canary’ circuits, which are designed with two insights.

① First, we observe that the impact of noise on the fidelity of a circuit is closely tied to the device-mapped structure of the circuit: how many 2-qubit CNOT gates are present and on which qubits they are placed; which qubits in the circuit are

involved in critically long, potentially decohering paths; which qubits contribute to worse readout errors, etc. Thus, a canary circuit that closely mimics the device-mapped structure of the original circuit, plus has as much overlap as possible with the original output state, could be useful for understanding the ensemble-wide impact of noise on the original circuit. But for this to be feasible, the correct outcome of the canary circuit should be known.

② Thus, second, we construct these canary circuits with only Clifford gates, by replacing all non-Clifford gates in the target circuit with the ‘nearest’ Clifford gates. A purely Clifford circuit is efficiently classically simulable (Section II-C), thus the correct output of the canary circuit is obtained via ideal classical simulation. Note that the correct output of the canary is almost always different from that of the original circuit (however, choosing the nearest Cliffords helps maintain state overlap to the best extent possible). But importantly, it still maintains the structure of the original circuit since only specific 1-qubit gates in the circuit have to be replaced by Clifford gates. Since the correct canary output is known, the ensemble can be ordered by the improving fidelity order of the canary circuit on the ensemble. This is estimated by running the canary circuit over the noisy ensemble and evaluating the probability of the correct canary output that they each produce. Then, as discussed above, this ordering is considered to be well suited to the original circuit (and we empirically show this to be the case). This ordering is then utilized as described in Section III to deduce the correct outcome of the target circuit. Construction and use of canaries is also illustrated in Fig.2.

A. Designing clifford canary circuits

Fig.7 shows a circuit example of Clifford canary generation. The top circuit shows a target 4-qubit circuit after mapping to a IBM Quantum device. On these devices, the basis gates are CX, ID, RZ, SX, X [1]. These basis gates are such the only non-Clifford gates are rotational gates about the Z-axis (RZ) with angles that are not multiples of $\pi/4$. Any number of such gates can be present in a circuit. The bottom circuit shows a Clifford canary generated for the target circuit. It is generated by ‘rounding’ the non-Clifford RZ gates to the nearest Clifford gates (i.e., to nearest multiple of $\pi/4$). Since all gates in the canary are Clifford, it can be efficiently classically simulated (Section II-C), and thus the canary correct outcome is known.




Fig. 8. Fidelity over a noisy ensemble for a 10-qubit circuit. Figure shows - a) ‘Original Farthest’: circuit with all 1-qubit gates midway between Cliffords, b) ‘Original Near’: same circuit structure, with all 1-qubit gates as near-Cliffords, c) ‘Closest Clifford’: with all 1-qubit gates as the closest Cliffords to the prior, d) ‘Random Clifford’: with all 1-qubits as random Cliffords.

Also, it is evident from the figure that the circuit structure is maintained in the canary circuit. As discussed earlier, the structural similarity leads to similar fidelity trends across the ensemble for both the original and the canary.

We quantify these observations in Fig.8. All the lines represent a 10-qubit circuit with 50-CX gates and 50 1-qubit gates, simulated on 10 different levels of noise. The blue line ‘Original Farthest’ represents a circuit in which all the 1-qubit gate angles are midway between Cliffords i.e., they have gate rotations: $\pi/4$ — hence these angles are farthest away from Cliffords. The red line ‘Original Near’ has the same circuit structure as above, but all the 1-qubit gates have only a small deviation from Cliffords: $\pi/8$. The green line ‘Closest Clifford’ circuit has the same structure, but with all angles as the nearest Clifford to the above: $\pi/2$. The yellow line ‘Random Clifford’ circuit has the same structure, but with all angles as random Cliffords: $\pi/2$.

The first point to note is that the ‘Random Clifford’ only captures the Original circuit’s fidelity trend very loosely — this alone is insufficient for our goals — this circuit is far from the Original, with insignificant state overlap. On the other hand, the ‘Closest Clifford’ is excellent at capturing the fidelity trend, closely mimicking the two Original circuits. Unsurprisingly, it almost overlaps the red line, since the angle difference (and hence the output state difference) between the two is marginal, but most importantly it also closely captures the trend of the blue line, which is the state that is farthest away from it. This is clearly indicative that even though the states of the two circuits are not the same, the impact of structural noise is significant enough to create similar fidelity trends as long as there is some non-trivial state overlap.

B. Quantifying canary ordering capabilities

Now, we show the capability of the canary circuits to capture fidelity trends of real applications on real devices. Fig.9 shows the high correlation in ensemble-wide fidelity between a 10-qubit adder circuit (orange) and its corresponding canary circuit (green) generated by the methodology described earlier. The comparison is performed across an ensemble of 50 different mappings on IBMQ Montreal. Note that, while this figure assumes that the output of the original adder circuit (and




Fig. 9. High correlation in ensemble-wide fidelity between a 10q adder (orange) and its canary (green), across an ensemble of 50 mappings on IBM Quantum Montreal.




Fig. 10. A 4-qubit QAOA circuit’s fidelity across an ensemble of 3 mappings each on 7 IBM Quantum machines is shown (orange). Fidelity of the corresponding canary is also shown (green) and achieves near-perfect correlation with the original. Also shown is the poorly predicted fidelity by a simulated noise model (blue) that uses the day’s machine calibration data.

therefore its fidelity) is known, this would not be the case for real-world quantum applications.

First, observe that the outputs produced by the two circuits are different (as shown in the legend) - this is expected from prior discussion. Second, note that the fidelity of both the original and the canary circuits are very low, under 4% - these are hard to execute circuits on today’s quantum devices. Third, within this ensemble, it is unclear which mapping will produce the highest fidelity. It is not the ‘optimal’ mapping chosen by the Qiskit transpiler, which is the first instance shown in the figure, and is clearly lower in fidelity compared to a couple of other instances. Though the output produced by the circuits are different and at low fidelity, maintaining the same circuit structure (as seen in Fig.7) enables the circuits to experience similar noise effects across the ensemble. The two trends are fairly similar across the ensemble, even if there are some instances of non-negligible deviation. The fidelity correlation between the two circuits is 87%, indicating that the canary is highly accurate in predicting a good ensemble order for this relatively low-fidelity original circuit.

Although it is evident from Fig.9 that the canary circuit can produce a good ensemble ordering for the target circuit, it would be fair to ask if an equally good ordering can be produced by simpler techniques, such as simply using the static noise information about the circuit. After all, the mapped qubits are known and the corresponding coherence times and gate errors are available from the device’s most recent calibration (roughly performed once a day for IBMQ machines). Perhaps this noise data can be used to construct some fidelity heuristic? We argue that this is insufficient for fairly complex




Fig. 11. Spearman correlation (varies from -1 to 1) between the canary-produced ensemble ordering and the ideal ordering for the target circuit, shown for 40 circuits (from real machine execution). The circuits are ordered by increasing size from 4-14 qubits. Green (fidelity > 10%) / Blue (2 – 10%) / Orange (< 2%). Circuits always show (very) strong correlation.

circuits and/or real device execution, for two reasons. First, the noise data obtained from the devices are stale - the devices are characterized only at coarse granularities of time (usually at the time of calibration) because per-qubit characterization is very time-consuming, and the noise characteristics of the qubits tend to drift over time [36]. Second, even if noise data were accurate, it is non-trivial to actually obtain the impact of these noise characteristics on the fidelity of a circuit of any reasonable complexity. This would require understanding every stage of the circuit’s execution, which is as complex as the circuit execution itself (and thus classically inefficient).

To quantify this argument, we perform an experiment with a fairly simple 4-qubit QAOA circuit in Fig.10. The fidelity of the circuit across an ensemble of 3 mappings, each on 7 IBM Quantum machines, is shown in orange. The fidelity of the corresponding canary circuit across the ensemble is shown in green and is seen to achieve near-perfect correlation with the original circuit. This circuit is of low width and depth, and this is a fairly easy task for the canary. In addition, the fidelity predicted by the Qiskit simulation noise model is shown in blue. This noise model uses the day’s calibration data from the machines’ qubits and gates. Clearly, the fidelity trend from the noise model is poorly correlated (only 24%) with the original, even for a simple circuit. The peaks and valleys of the original circuit’s fidelity across the ensemble are not captured and, thus, this is poor for ensemble ordering. Further, we argue that this noise model based simulation is a loose upper bound on the accuracy that any noise based fidelity heuristic can produce. A more simple heuristic might multiply different error rates and coherence times, whereas the noise model is actually simulating the circuit with the noise data. Note that this noise model simulation is itself not a scalable approach. Thus, it is clear that noise data based heuristic predictions of fidelity are inaccurate even for simple circuits, clearly motivating the canary approach to order the ensemble.

Next, Fig.11 quantifies the high-accuracy ensemble orderings that our canary circuits are able to produce for a variety of applications. The figure shows the Spearman correlation between the canary-produced ensemble order and the ideal ensemble order for the target circuit (which would be unknown for complex non-simulable circuits), for 40 circuits. The circuits are ordered by increasing size from 4-14 qubits. Green (fidelity > 10%) / Blue (1 – 10%) / Orange (< 1%).




Fig. 12. Spearman correlation (varies from -1 to 1) between the canary-produced ensemble ordering and the ideal ordering for the target circuit, shown for 8 deep 12q circuits with many gates (in simulation).

Spearman correlation measures the strength and direction of monotonic association between two variables and varies from -1 to 1. Positive values, especially those closer to one, indicate a very strong positive correlation. Correlations for all circuits are in the strong to very strong range, clearly indicating the effectiveness of Clifford canaries on real machines, even for circuits which are close to random output distributions.

C. Scaling to more complex circuits

To evaluate the capabilities of the Clifford canaries on more complex circuits, we perform experiments in noisy simulation on very deep, reasonably wide, circuits. In Fig.12, we show 12-qubit circuits with circuit depth ranging from 195-1,300 CX gates and with the total number of non-Clifford gates ranging from 96-656. It can be observed that the correlation between the Clifford canaries and the original circuits is always in the very strongly correlated range. This is a clear indicator of the scalability of this approach to more complex circuits, as and when they can be run on quantum devices with non-negligible fidelity. Note that the correlations seen here are generally higher than those in Fig.11 - this is because real machine noise is more erratic than noisy simulation. Nevertheless, the utility of the Clifford canaries is evident.

V. METHODOLOGY

A. Applications

Our benchmarks are representative of real-world usecases, and are described below and detailed in Table I. All applications have average baseline fidelity at 10% or lower - i.e., we are primarily interested in applications for which correct outcomes are relatively hard to identify.

Ripple Carry Adder: Adders are a critical building block for quantum logic. We implement linear-depth ripple-carry adder quantum circuits for 2-6 bits that utilize 6-14 qubits respectively [11].

Quantum Fourier Transform: QFT is a circuit used as a building block for applications such as Shor’s factoring and phase estimation. It converts a quantum state from the computational basis to the Fourier basis through the introduction of phase. QFT was constructed for 6-8 qubits [31].

Quantum Approximate Optimization Algorithm: QAOA [17] is a variational quantum-classical algorithm to solve combinatorial optimization problems. We use QAOA ansatz constructed for 6 qubits for 4 different input graphs.

App	Q	Output	Depth	MaxFid	AvgFid
QAOA6_1	6	101011+011000	41	29%	11.0%
ADD6_1	6	110000	29	32%	10.0%
ADD6_2	6	111100	29	25%	9.0%
QFT6	6	Equal superposn.	36	42%	7.0%
QAOA6_2	6	011001+001000	93	10%	5.0%
ADD8_1	8	00111100	41	13%	3.7%
QAOA6_3	6	001001+001000	243	5%	3.7%
QAOA6_4	6	011001+001000	283	6%	3.7%
ADD8_2	8	11111100	41	11%	2.6%
ADD10_1	10	0011000000	53	3.5%	1.0%
QFT8	8	Equal superposn.	52	1.7%	0.9%
ADD10_2	10	1111110000	53	2%	0.4%
ADD12_1	12	000011001100	65	0.7%	0.5%
ADD14_1	14	00000011110000	77	0.5%	0.4%
ADD12_2	12	111111111100	65	0.5%	0.35%

TABLE I

EVALUATED ADDER, QAOA, AND QFT BENCHMARKS. CIRCUIT OUTPUTS, PRE-MAPPING CX DEPTH, MAXIMUM FIDELITY, AND AVERAGE FIDELITY ACROSS ALL ENSEMBLES ARE SHOWN.

B. Infrastructure / Overheads

Quancorde is implemented to interface with Qiskit [3] and is evaluated on IBMQ [2] devices. Canary circuits are efficiently constructed on classical compute.

Intra-device evaluations: Our intra-device evaluations target the 27-qubit IBM Quantum Montreal and we run 50 different circuit instances across the device. One instance is that which is mapped and routed to the device by the Qiskit transpiler when run with highest optimizations enabled. The other 49 are random mappings.

Inter-device evaluations: Our inter-device evaluations target 7 IBM Quantum devices: Montreal (27q), Toronto (27q), Guadalupe (16q), Hanoi (27q), Cairo (27q), Mumbai (27q) and Brooklyn (65q). Only one circuit is run per device, mapped and routed to the device by the Qiskit transpiler when run with highest optimizations enabled.

Correlation post-processing: The correlation processing step involves calculating the correlation of the ensemble ordering for each output string in the baseline probability distribution with the canary ordering. At first glance, it may seem that output strings are possible, but this is *not* the case and the cost is not exponential. There are two upper bounds that can be established on the number of strings which have to be analyzed. The trivial upper-bound is $\frac{2^N}{N}$ i.e., the number of unique strings produced cannot be more than number of shots run over all the devices - this number would be reached if each device produced a purely random distribution, and each machines distribution was mutually exclusive. The number of shots is not exponential (if it were, quantum computing would be infeasible) and can be derived from the target decipherable fidelity on the quantum devices. There is an even lower upper-bound that can be derived directly from a minimum baseline fidelity target. If we establish that the minimum application fidelity we are interested in boosting is, say, f (since in this region, the output distribution starts becoming purely random), then we would at most have to analyze only the top $\frac{1}{f}$ ($= 1000$ for $f = 0.001$) unique strings in terms of their probability of occurrence. This is because there cannot be more than

999 strings which have probability of occurrence greater than or equal to than the correct outcome, because if it were so, that would mean that the fidelity would have to be less than 0.1%. This can be trivially reasoned mathematically. Thus, the correct outcome will have to be among the top 1000 strings. Or more generally, among the top $\frac{1}{f}$ strings for a target lower bound baseline fidelity f . This reasoning can similarly be extended to multi bit-string scenarios as well, but again, clearly it is non exponential.

C. Evaluation Metrics

We evaluate Quancorde benefits through 2 metrics:

① **Rank of correct outcome:** This metric indicates the ranking of the correct outcome among the correlations of all possible outcomes with the canary-produced ensemble order. We refer to this as ‘Rank’. If Quancorde works perfectly, the rank of the correct outcome is 1. But for more challenging circuits, the rank is not always 1. Note that while a rank nearer to 1 is always more beneficial, it is not required for Quancorde’s success - even lower ranks allow significant fidelity improvements over the baseline.

② **Fidelity:** This metric is a standard evaluation metric that indicates the likelihood of finding the correct outcome of a quantum application. We show the benefits of Quancorde against: a) the maximum baseline fidelity obtained for a target application across all mappings and machines, and b) the average benefits across the entire ensemble. Practical benefits can lie somewhere in between.

VI. EVALUATION

A. Single-application 10-qubit adder analysis

Fig.13 is a real machine / application example of the benefits produced by Quancorde, with an intra-machine ensemble. It shows the analysis for a 10-qubit adder which has a baseline machine fidelity of under 1% - the correct answer is a single bitstring ‘1111110000’. In all 3 graphs, the x-axis shows all the different 10-bit strings. The top figure shows the output probability distribution for the baseline. The correct outcome is indicated. Note that the probability of its occurrence (i.e., fidelity) is very low, only 1.7%. Further, it is not the output string with the highest probability. Clearly, the correct outcome cannot be identified. The middle figure shows Quancorde’s correlation analysis - the correlations vary from -1 to 1, with near 1 indicating high positive correlation. The correlation of the correct outcome is indicated and it has the highest correlation of nearly 0.9. Next, the bottom figure shows the weighted output distribution by incorporating the correlation into the baseline distribution. Only a few strings remain, and the correct outcome’s probability is boosted by 10x, to 20%. Further, it is a clear winner among all the strings.

B. Outcome Ranks and Correlation Strings

Figure 14 shows two Quancorde results for 15 applications. The results are averaged across 3 trials of intra-machine and inter-machine ensembles.




Fig. 13. Bitstring distribution produced by the baseline for a 10-q adder is shown in red. The probability of the correct outcome is low in magnitude (2%) and also lower than other strings. The Correlation distribution produced by Quancorde (blue), with an intra-machine ensemble, shows the correct outcome to have the highest correlation with the canary ensemble. When this is used to weight the original distribution, the Output distribution produced (green) has the correct outcome at 20% probability and is a clear winner.




Fig. 14. Left axis / bars: Rank produced by Quancorde for 15 applications. Right axis / line: Number of output strings (with occurrence probability greater than 0.1%) analyzed for correlation.

First, we look at the left axis / green columns, which shows the rank of the correct outcome in the correlation distribution. As discussed in Section V-C, a rank closer to 1 can improve fidelity further, since this would imply the correct outcome having a higher correlation compared to other outputs, thus boosting it further in the final output distribution. But lower ranks can still produce substantial benefits. It is evident that many of the relatively smaller circuits (6 qubits) have perfect Rank. The larger more complicated circuits (8 qubits and greater) have Ranks in the range of 1-30. Even the worst-case Rank 30 occurs for a distribution which has 4096 potential outputs, clearly showcasing the potential for harnessing correlation. Mean ranks across the applications show that the correct outcome usually occurs in the top 2-3 outputs in the correlation distribution. Also, we noted

that Quancorde performs marginally better on inter-machine diversity compared to intra-machine diversity. The difference is small enough to suggest that any form of diversity is useful.

Next, we look at the right axis / orange line, which shows the number of strings analyzed in correlation processing. As discussed in Section V, with a target baseline fidelity lower bound of 0.1%, the maximum strings analyzed would be 1000. From the figure, we see that the actual number only goes up to 316. The smaller circuits have nearly all their bitstrings with a fidelity greater than 0.1%, whereas the large circuits, like the 14-qubit adder, have less than 2% of their strings worthy of correlation processing. This clearly shows that the cost of correlation processing is very small, making Quancorde a computationally scalable approach.

C. Fidelity benefits




Fig. 15. Fidelity improvements of Quancorde compared to the baseline. Results are averaged over inter- and intra-machine ensembles. Comparison to EDM is also shown.

Next, we show the fidelity improvements of Quancorde compared to the baseline. Results are averaged over inter- and intra-machine ensembles. Two sets of fidelity improvements are shown — ‘Mean’ and ‘vs Best’ — these were discussed in Section V. Quancorde achieves considerable fidelity boosting, especially for the larger applications. For smaller applications, the baseline mean fidelity is already around 10% and as high as 42%, so the boosting potential is limited. Nevertheless, Quancorde achieves very high fidelity after boosting. For larger applications, Quancorde’s benefits are very attractive. The improvements are as high as 15x / 34x (‘Mean’ and ‘vs Best’ respectively) for the 12-qubit Adder, which had a correct outcome rank of under 5, out of 4k bitstrings. On average, fidelity boosting of 8.9x / 4.2x is achieved, which is higher than typical error mitigation techniques, clearly highlighting Quancorde’s novel capability of harness a diverse ensemble. Note that considerable improvements are achieved even on uniform multi-output distributions like QFT. Non-uniform distributions are discussed separately in Section VI-D.

We also show comparison against EDM [45], another ensembling approach, albeit for a different goal of avoiding specific occurrences of correlated errors. We use an ensemble size of 4 best mappings (which is the default setting). We observe fidelity improvements from EDM to be substantially lower than that obtained from Quancorde - 1.1x on average




Fig. 16. Quancorde benefits in non-uniformly weighted multi-outcome applications, shown for a 7-qubit circuit with 17 outputs. The Quancorde output distribution closely matches the Ideal distribution, while the baseline distribution does not. Some correct / incorrect peaks are indicated.




Fig. 17. Rank of correct outcome improves as size of intra-machine ensemble increases, for a 10-qubit adder on IBM Quantum Montreal.

and a maximum of 1.6x. EDM is seen to provide benefits primarily in scenarios of reasonably high baseline fidelity - the top 5-6 benchmarks in Table I. To our knowledge, this is likely because EDM is suited to reducing the probability of occurrence of very specific competing incorrect bitstrings which have uncharacteristically high occurrence due to unique correlated errors. Such scenarios are unlikely near the threshold of device capability, at which point, the impact of all qubits, their gate errors and coherence times are all significant.

D. Non-uniform Multi-output usecase

Quancorde can also achieve substantial benefits for applications with non-uniformly weighted multi-outcomes. Fig.16 shows Quancorde benefits for a 7-qubit circuit with a correct outcome distribution comprising of a superposition of 17 7-bit strings. The experiment is run with an intra-machine ensemble on Montreal. The baseline distribution (orange) is far off from the ideal distribution in purple. For example, peak (2) in the baseline is absent in the ideal, and other peaks can be similarly inferred. Quancorde's correlation processing identified all 17 strings amongst its top 24 Ranks. The weighted output is shown in green. Clearly it bears good resemblance to the Ideal with peaks such as (1). While still imperfect, a fidelity boosting of 4.6x / 2.1x ('Mean' / 'vs Best', respectively) is achieved over the baseline.

E. Sensitivity to ensemble size

Greater diversity and larger ensemble sizes could increase the potential for correct outcome identification. Fig.17 shows how the Rank of the correct outcome for a 10-qubit adder circuit improves as we increase the size of its intra-machine Quancorde ensemble. The experiment is performed on IBM Q Montreal. Even a small ensemble of just five mappings is fairly successful in achieving a top-30 Rank. As is evident, increasing the ensemble size improves the Rank. This is because a large ensemble has more diverse mappings or machines. This means that there are more diverse noise effects over which the occurrence probabilities of the different output strings are evaluated. Consequently, the correlation of the unique ensemble order for the correct outcome, with the canary-produced ensemble order, is relatively magnified, thus boosting fidelity further. The benefits reach the maximum (Rank = 1) at an ensemble size of 30 mappings. The optimal size of the ensemble is not directly related to the size of the circuit and is more tied to the baseline fidelity.

VII. CONCLUSION

Quancorde is a novel approach to identifying the correct outcomes of extremely low fidelity quantum applications by exploiting quantum device diversity. It uses Clifford canary circuits to order a diverse ensemble of devices or qubits/mappings along the direction of increasing fidelity of the target application. It then identifies the correlation of the application outputs' ensemble ordering with the canary order, and uses this to weight the application output distribution, resulting in boosted application fidelity. Quancorde is especially useful when diversity is significant, and applications' ideal outcomes are hard to produce. Quancorde can have a revolutionary impact on the way noisy quantum resources are leveraged to boost the fidelity of real-world quantum use cases.

ACKNOWLEDGEMENT

This work is funded in part by EPIQC, an NSF Expedition in Computing, under award CCF-1730449; in part by STAQ under award NSF Phy-1818914; in part by NSF award 2110860; in part by the US Department of Energy Office of Advanced Scientific Computing Research, Accelerated Research for Quantum Computing Program; and in part by the NSF Quantum Leap Challenge Institute for Hybrid Quantum Architectures and Networks (NSF Award 2016136) and in part based upon work supported by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers. This work was completed in part with resources provided by the University of Chicago's Research Computing Center. GSR is supported as a Computing Innovation Fellow at the University of Chicago. This material is based upon work supported by the National Science Foundation under Grant # 2030859 to the Computing Research Association for the CIFellows Project. KNS is supported by IBM as a Postdoctoral Scholar at the University of Chicago and the Chicago Quantum Exchange. FTC is Chief Scientist for Quantum Software at ColdQuanta and an advisor to Quantum Circuits, Inc.

REFERENCES

- [1] "Ibm quantum experience," <https://quantum-computing.ibm.com>, accessed: 2020-11-18.
- [2] "Ibm quantum systems," <https://quantum-computing.ibm.com/services?systems=all>, accessed: 2021-12-01.
- [3] H. Abraham, AduOffei, R. Agarwal, I. Y. Akhalwaya, G. Aleksandrowicz, T. Alexander, M. Amy, E. Arbel, Arijit02, A. Asfaw, A. Avkhadiiev, C. Azaustre, AzizNgoueya, A. Banerjee, A. Bansal, P. Barkoutsos, G. Barron, G. S. Barron, L. Bello, Y. Ben-Haim, D. Bevenius, A. Bhohe, L. S. Bishop, C. Blank, S. Bolos, S. Bosch, Brandon, S. Bravyi, Bryce-Fuller, D. Bucher, A. Burov, F. Cabrera, P. Calpin, L. Capelluto, J. Carballo, G. Carrascal, A. Chen, C.-F. Chen, E. Chen, J. C. Chen, R. Chen, J. M. Chow, S. Churchill, C. Claus, C. Clauss, R. Cocking, F. Correa, A. J. Cross, A. W. Cross, S. Cross, J. Cruz-Benito, C. Culver, A. D. Córcoles-Gonzales, S. Dague, T. E. Dandachi, M. Daniels, M. Dartiailh, DavideFrr, A. R. Davila, A. Dekusar, D. Ding, J. Doi, E. Drechsler, Drew, E. Dumitrescu, K. Dumon, I. Duran, K. EL-Safty, E. Eastman, G. Eberle, P. Eendebak, D. Egger, M. Everitt, P. M. Fernández, A. H. Ferrera, R. Fouilland, FranckChevallier, A. Frisch, A. Fuhrer, B. Fuller, M. GEORGE, J. Gacon, B. G. Gago, C. Gambella, J. M. Gambetta, A. Gammanpila, L. Garcia, T. Garg, S. Garion, A. Gilliam, A. Giridharan, J. Gomez-Mosquera, S. de la Puente González, J. Gorzinski, I. Gould, D. Greenberg, D. Grinko, W. Guan, J. A. Gunnels, M. Haglund, I. Haide, I. Hamamura, O. C. Hamido, F. Harkins, V. Havlicek, J. Hellmers, L. Herok, S. Hillmich, H. Horii, C. Howington, S. Hu, W. Hu, J. Huang, R. Huisman, H. Imai, T. Imamichi, K. Ishizaki, R. Iten, T. Itoko, JamesSeaward, A. Javadi, A. Javadi-Abhari, Jessica, M. Jivrajani, K. Johns, S. Johnstun, Jonathan-Shoemaker, V. K, T. Kachmann, N. Kanazawa, Kang-Bae, A. Karazeev, P. Kassebaum, J. Kelso, S. King, Knabberjoe, Y. Kobayashi, A. Kovyrshin, R. Krishnakumar, V. Krishnan, K. Krsulich, P. Kumkar, G. Kus, R. LaRose, E. Lacal, R. Lambert, J. Lapeyre, J. Latone, S. Lawrence, C. Lee, G. Li, D. Liu, P. Liu, Y. Maeng, K. Majmudar, A. Malyshev, J. Manela, J. Marecek, M. Marques, D. Maslov, D. Mathews, A. Matsuo, D. T. McClure, C. McGarry, D. McKay, D. McPherson, S. Meesala, T. Metcalfe, M. Mevissen, A. Meyer, A. Mezzacapo, R. Midha, Z. Minev, A. Mitchell, N. Moll, J. Montanez, M. D. Mooring, R. Morales, N. Moran, M. Motta, MrF, P. Murali, J. Müggenburg, D. Nadlinger, K. Nakanishi, G. Nannicini, P. Nasion, E. Navarro, Y. Naveh, S. W. Neagle, P. Neuweiler, J. Nicander, P. Niroula, H. Norlen, NuoWenLei, L. J. O'Riordan, O. Ogunbayo, P. Ollitrault, R. Otaolea, S. Oud, D. Padilha, H. Paik, S. Pal, Y. Pang, S. Perriello, A. Phan, F. Piro, M. Pistoia, C. Piveteau, P. Poceau, A. Pozas-iKerstjens, V. Prutyaynov, D. Puzzuoli, J. Pérez, Quintii, R. I. Rahman, A. Raja, N. Ramagiri, A. Rao, R. Raymond, R. M.-C. Redondo, M. Reuter, J. Rice, M. L. Rocca, D. M. Rodríguez, RohithKarur, M. Rossmannek, M. Ryu, T. SAPV, SamFerracin, M. B. Sletfjerding, H. Sandesara, R. Sapra, H. Sargsyan, A. Sarkar, N. Sathaye, B. Schmitt, C. Schnabel, Z. Schoenfeld, T. L. Scholten, E. Schoute, J. Schwarm, I. F. Sertage, K. Setia, N. Shammah, Y. Shi, A. Silva, A. Simonetto, N. Singstock, Y. Siraichi, I. Sidiokov, S. Sivarajah, M. B. Sletfjerding, J. A. Smolin, M. Soeken, I. O. Sokolov, I. Sokolov, SooluThomas, Starfish, D. Steenken, M. Stypulkoski, S. Sun, K. J. Sung, H. Takahashi, T. Takawale, I. Tavernelli, C. Taylor, P. Taylour, S. Thomas, M. Tillet, M. Tod, M. Tomasik, E. de la Torre, K. Trabing, M. Treinish, TrishaPe, D. Tuls, W. Turner, Y. Vaknin, C. R. Valcarce, F. Varchon, A. C. Vazquez, V. Villar, D. Vogt-Lee, C. Vuillot, J. Weaver, J. Weidenfeller, R. Wiecek, J. A. Wildstrom, E. Winston, J. J. Woehr, S. Woerner, R. Woo, C. J. Wood, R. Wood, S. Wood, S. Wootton, D. Yeralin, D. Yonge-Mallo, R. Young, J. Yu, C. Zachow, L. Zdanski, H. Zhang, C. Zoufal, Zoufal, a kapila, a matsuo, bcammorrison, brandhsn, nick bronn, chlorophyll zz, dekel.meirom, dekelmeirom, dekok, dime10, dholmie, dtrenev, ehchen, elfrocampeador, faisaldebouni, fanizamarco, gabrieleagl, gadial, galeinston, georgios ts, gruu, hhorii, hykavitha, jagunther, jliu45, jscott2, kanejess, klinvill, krutik2966, kurarr, lerongil, ma5x, merav aharoni, michelle4654, ordmoj, sagar pahwa, rmoyard, saswati qiskit, scottkelso, stfmerkel, strickroman, sumitpur, tigerjack, toulal, tsura crinaldo, vvilpas, welien, willhbang, yang.luh, yotam-vakninibm, and M. Čepulkovskis, "Qiskit: An open-source framework for quantum computing," 2019.
- [4] G. S. Barron and C. J. Wood, "Measurement error mitigation for variational quantum algorithms," 2020.
- [5] E. Bernstein and U. Vazirani, "Quantum complexity theory," *SIAM J. Comput.*, vol. 26, no. 5, p. 1411–1473, oct 1997. [Online]. Available: <https://doi.org/10.1137/S0097539796300921>
- [6] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, "Quantum machine learning," *Nature*, vol. 549, no. 7671, pp. 195–202, 2017.
- [7] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven, "Characterizing quantum supremacy in near-term devices," *Nature Physics*, vol. 14, no. 6, p. 595–600, Apr 2018. [Online]. Available: <http://dx.doi.org/10.1038/s41567-018-0124-x>
- [8] L. Botelho, A. Glos, A. Kundu, J. A. Miszczyk, Özlem Salehi, and Z. Zimborás, "Error mitigation for variational quantum algorithms through mid-circuit measurements," 2021.
- [9] S. Bravyi, S. Sheldon, A. Kandala, D. C. McKay, and J. M. Gambetta, "Mitigating measurement errors in multiqubit experiments," *Physical Review A*, vol. 103, no. 4, p. 042605, 2021.
- [10] J. Burnett, A. Bengtsson, M. Scigliuzzo, D. Niepce, M. Kudra, P. Delsing, and J. Bylander, "Decoherence benchmarking of superconducting qubits. npj quantum inf. 5," 2019.
- [11] S. A. Cuccaro, T. G. Draper, S. A. Kutin, and D. P. Moulton, "A new quantum ripple-carry addition circuit," *arXiv preprint quant-ph/0410184*, 2004.
- [12] P. Czarnik, A. Arrasmith, P. J. Coles, and L. Cincio, "Error mitigation with clifford quantum-circuit data," 2020.
- [13] P. Das, S. Tannu, and M. Qureshi, "Jigsaw: Boosting fidelity of nisq programs via measurement subsetting," in *MICRO-54: 54th Annual IEEE/ACM International Symposium on Microarchitecture*, ser. MICRO '21. New York, NY, USA: Association for Computing Machinery, 2021, p. 937–949. [Online]. Available: <https://doi.org/10.1145/3466752.3480044>
- [14] H. De Raedt, F. Jin, D. Willsch, M. Willsch, N. Yoshioka, N. Ito, S. Yuan, and K. Michielsen, "Massively parallel quantum computer simulator, eleven years later," *Computer Physics Communications*, vol. 237, p. 47–61, Apr 2019. [Online]. Available: <http://dx.doi.org/10.1016/j.cpc.2018.11.005>
- [15] Y. Ding, P. Gokhale, S. F. Lin, R. Rines, T. Propson, and F. T. Chong, "Systematic crosstalk mitigation for superconducting qubits via frequency-aware compilation," *arXiv preprint arXiv:2008.09503*, 2020.
- [16] D. Ernst, Nam Sung Kim, S. Das, S. Pant, R. Rao, Toan Pham, C. Ziesler, D. Blaauw, T. Austin, K. Flautner, and T. Mudge, "Razor: a low-power pipeline based on circuit-level timing speculation," in *Proceedings. 36th Annual IEEE/ACM International Symposium on Microarchitecture, 2003. MICRO-36., 2003*, pp. 7–18.
- [17] E. Farhi, J. Goldstone, and S. Gutmann, "A quantum approximate optimization algorithm," 2014.
- [18] T. Giurgica-Tiron, Y. Hindy, R. LaRose, A. Mari, and W. J. Zeng, "Digital zero noise extrapolation for quantum error mitigation," in *2020 IEEE International Conference on Quantum Computing and Engineering (QCE)*. IEEE, 2020, pp. 306–316.
- [19] D. Gottesman, "The heisenberg representation of quantum computers," *arXiv preprint quant-ph/9807006*, 1998.
- [20] L. K. Grover, "A fast quantum mechanical algorithm for database search," in *ANNUAL ACM SYMPOSIUM ON THEORY OF COMPUTING*. ACM, 1996, pp. 212–219.
- [21] T. Häner and D. S. Steiger, "0.5 petabyte simulation of a 45-qubit quantum circuit," *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, Nov 2017. [Online]. Available: <http://dx.doi.org/10.1145/3126908.3126947>
- [22] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, "Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets," *Nature*, vol. 549, no. 7671, pp. 242–246, 2017.
- [23] P. 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. Mutus, O. Naaman, C. Neill, C. Quintana, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. White, S. Boixo, R. Babbush, V. Smelyanskiy, H. Neven, and J. Martinis, "Fluctuations of energy-relaxation times in superconducting qubits," *Physical Review Letters*, vol. 121, no. 9, aug 2018. [Online]. Available: <https://doi.org/10.1103/PhysRevLett.121.090502>
- [24] Y. Li and S. C. Benjamin, "Efficient variational quantum simulator incorporating active error minimization," *Phys. Rev. X*, vol. 7, p.

- 021050, Jun 2017. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevX.7.021050>
- [25] J. M. Martinis, K. B. Cooper, R. McDermott, M. Steffen, M. Ansmann, K. Osborn, K. Cicak, S. Oh, D. P. Pappas, R. W. Simmonds *et al.*, “Decoherence in josephson qubits from dielectric loss,” *Physical review letters*, vol. 95, no. 21, p. 210503, 2005.
 - [26] M. Mohseni, P. Read, H. Neven, S. Boixo, V. Denchev, R. Babbush, A. Fowler, V. Smelyanskiy, and J. Martinis, “Commercialize quantum technologies in five years,” *Nature*, vol. 543, p. 171–174, 2017. [Online]. Available: <http://www.nature.com/news/commercialize-quantum-technologies-in-five-years-1.21583>
 - [27] N. Moll, P. Barkoutsos, L. S. Bishop, J. M. Chow, A. Cross, D. J. Egger, S. Filipp, A. Fuhrer, J. M. Gambetta, M. Ganzhorn *et al.*, “Quantum optimization using variational algorithms on near-term quantum devices,” *Quantum Science and Technology*, vol. 3, no. 3, p. 030503, 2018.
 - [28] C. Müller, J. H. Cole, and J. Lisenfeld, “Towards understanding two-level-systems in amorphous solids: insights from quantum circuits,” *Reports on Progress in Physics*, vol. 82, no. 12, p. 124501, 2019.
 - [29] P. Murali, J. M. Baker, A. Javadi-Abhari, F. T. Chong, and M. 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*, 2019, pp. 1015–1029.
 - [30] P. Murali, D. C. McKay, M. Martonosi, and A. 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*, 2020, pp. 1001–1016.
 - [31] M. A. Nielsen and I. Chuang, *Quantum computation and quantum information*. Cambridge University Press, 2010.
 - [32] J. O’Gorman and E. T. Campbell, “Quantum computation with realistic magic-state factories,” *Physical Review A*, vol. 95, no. 3, Mar 2017. [Online]. Available: <http://dx.doi.org/10.1103/PhysRevA.95.032338>
 - [33] T. Patel and D. Tiwari, “Qraft: reverse your quantum circuit and know the correct program output,” in *Proceedings of the 26th ACM International Conference on Architectural Support for Programming Languages and Operating Systems*, 2021, pp. 443–455.
 - [34] B. Pokharel, N. Anand, B. Fortman, and D. A. Lidar, “Demonstration of fidelity improvement using dynamical decoupling with superconducting qubits,” *Physical review letters*, vol. 121, no. 22, p. 220502, 2018.
 - [35] J. Preskill, “Quantum computing in the nisq era and beyond,” *Quantum*, vol. 2, p. 79, 2018.
 - [36] T. Proctor, M. Reville, E. Nielsen, K. Rudinger, D. Lobser, P. Maunz, R. Blume-Kohout, and K. Young, “Detecting and tracking drift in quantum information processors,” *Nature Communications*, vol. 11, no. 1, oct 2020.
 - [37] G. S. Ravi, K. N. Smith, P. Gokhale, and F. T. Chong, “Quantum computing in the cloud: Analyzing job and machine characteristics,” in *2021 IEEE International Symposium on Workload Characterization (IISWC)*, 2021, pp. 39–50.
 - [38] G. S. Ravi, K. N. Smith, P. Gokhale, A. Mari, N. Earnest, A. Javadi-Abhari, and F. T. Chong, “Vaqem: A variational approach to quantum error mitigation,” 2021.
 - [39] E. Rosenberg, P. Ginsparg, and P. L. McMahon, “Experimental error mitigation using linear rescaling for variational quantum eigensolving with up to 20 qubits,” *Quantum Science and Technology*, Nov 2021. [Online]. Available: <http://dx.doi.org/10.1088/2058-9565/ac3b37>
 - [40] S. Schlör, J. Lisenfeld, C. Müller, A. Bilmes, A. Schneider, D. P. Pappas, A. V. Ustinov, and M. Weides, “Correlating decoherence in transmon qubits: Low frequency noise by single fluctuators,” *Physical review letters*, vol. 123, no. 19, p. 190502, 2019.
 - [41] P. W. Shor, “Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer,” *SIAM Journal on Computing*, vol. 26, no. 5, p. 1484–1509, Oct 1997. [Online]. Available: <http://dx.doi.org/10.1137/S0097539795293172>
 - [42] K. N. Smith, G. S. Ravi, P. Murali, J. M. Baker, N. Earnest, A. Javadi-Abhari, and F. T. Chong, “Error mitigation in quantum computers through instruction scheduling,” *arXiv preprint arXiv:2105.01760*, 2021.
 - [43] A. M. Souza, G. A. Álvarez, and D. Suter, “Robust dynamical decoupling,” *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 370, no. 1976, pp. 4748–4769, 2012.
 - [44] R. Takagi, S. Endo, S. Minagawa, and M. Gu, “Fundamental limits of quantum error mitigation,” 2021.
 - [45] S. S. Tannu and M. 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*, ser. MICRO ’52. New York, NY, USA: Association for Computing Machinery, 2019, p. 253–265. [Online]. Available: <https://doi.org/10.1145/3352460.3358257>
 - [46] S. S. Tannu and 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*, 2019, pp. 279–290.
 - [47] —, “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*, 2019, pp. 987–999.
 - [48] K. Temme, S. Bravyi, and J. M. Gambetta, “Error mitigation for short-depth quantum circuits,” *Physical review letters*, vol. 119, no. 18, p. 180509, 2017.
 - [49] J. Tilly, H. Chen, S. Cao, D. Picozzi, K. Setia, Y. Li, E. Grant, L. Wossnig, I. Rungger, G. H. Booth, and J. Tennyson, “The variational quantum eigensolver: a review of methods and best practices,” 2021.
 - [50] L. Viola, E. Knill, and S. Lloyd, “Dynamical decoupling of open quantum systems,” *Physical Review Letters*, vol. 82, no. 12, p. 2417, 1999.
 - [51] S. Wang, P. Czarnik, A. Arrasmith, M. Cerezo, L. Cincio, and P. J. Coles, “Can error mitigation improve trainability of noisy variational quantum algorithms?” 2021.