

Efficient Quantum Circuit Cutting by Neglecting Basis Elements

Daniel T. Chen*, Ethan H. Hansen*, Xinpeng Li*, Vinooth Kulkarni*, Vipin Chaudhary*, Bin Ren[†], Qiang Guan[†], Sanmukh Kuppannagari*, Ji Liu[§], Shuai Xu*

* Department of Computer and Data Science

Case Western Reserve University, Cleveland, OH 44106

Email: {txc461, eh50, xxl1337, vxk285, vxc204, sxx1942, sxx214}@case.edu

[†] Department of Computer Science

Kent State University, Kent, OH 44240

Email: qguan@kent.edu

[‡] College of William & Mary, Williamsburg, VA 23185

Email: bren@wm.edu

[§] Argonne National Laboratory, IL 60439

Email: ji.liu@anl.gov

Abstract—Quantum circuit cutting has been proposed to help execute large quantum circuits using only small and noisy machines. Intuitively, cutting a qubit wire can be thought of as classically passing information of a quantum state along each element in a basis set. As the number of cuts increase, the number of quantum degrees of freedom needed to be passed through scales exponentially. We propose a simple reduction scheme to lower the classical and quantum resources required to perform a cut. Particularly, we recognize that for some cuts, certain basis element might pass “no information” through the qubit wire and can effectively be neglected. We empirically demonstrate our method on circuit simulators as well as IBM quantum hardware, and we observed up to 33 percent reduction in wall time without loss of accuracy.

I. INTRODUCTION

While quantum computers can provide exponential speed-up for certain algorithms, unveiling potential diverse and disruptive applications [1], [2], [3], current quantum hardware lacks the scale and reliability to execute these quantum circuits. Preskill described such hardware as Noisy Intermediate-Scale Quantum (NISQ) computers, and much effort has been focused on increasing qubit fidelity and qubit count [4], [5], [6], [7]. Meanwhile, algorithmic developments such as variational quantum algorithms [8], [9], [10] have caught substantial attention for having the potential to realize quantum advantage utilizing only NISQ machines. However, there are few results demonstrating provable advantage for using variational quantum algorithms over known classical algorithms [11].

A natural question to ask when given limited qubits is: “Are there ways of simulating large quantum circuits

using small quantum devices?” Inspired by fragmentation methods in molecular simulations [12], [13], [14], [15], Peng *et al.* has demonstrated that it is possible to divide quantum circuit into subcircuits, or *fragments*, that can be independently executed on small quantum computers and then recombined [16]. This method does not impose assumptions on circuits but requires a large amount of classical computing resources—exponential in the number of cuts—in order to reconstruct the expected output of the respective uncut circuit.

Nonetheless, quantum circuit cutting still holds great promise as an effective tool for realizing quantum algorithms. Most notably, it has been empirically shown that reducing the circuit size, even when sufficiently large machines are available, can improve fidelity [17], [18], [19]. Moreover, tailoring circuit cutting to specific applications, *e.g.* combinatorial optimization [20] and error mitigation [21], can effectively avoid the exponential classical overhead. Alternatively, stochastic methods through randomized measurement [22] and sampling [23] have also been developed to reduce the postprocessing cost. In recent, state-of-the-art research, it has been empirically shown that optimization can be employed to replace traditional tomographic techniques and significantly reduce the classical resources needed for reconstruction, though at the cost of additional circuit evaluations [24]. Circuit cutting algorithms have also undergone improvement and variations in terms of performance. Maximum likelihood methods have been introduced to circuit cutting to formally account for finite-shot error [19]. Shadow tomography—an efficient classical representation of quantum states for predicting expectations of observables [25]—can also be applied to

This research was supported in part by NSF Award 2216923.

individual circuit fragments [26].

Our work contributes in the direction of reducing the computational resources needed. In this paper, we introduce the idea of a *golden cutting point*: a cut location on a quantum circuit where a certain measurement basis can be neglected, *e.g.*, when the quantum state operates in a lower-dimensional subspace prior to the cut. Such reduction lowers the runtime of reconstructing measurement statistics from fragments by approximately 33% for a single cut and avoids excessive circuit executions. We experimentally verify our method on both the Qiskit Aer simulator [27] and superconducting quantum devices from IBM [28].

The outline of this paper is as follows. Section II presents the mathematical formulation behind quantum circuit cutting and the idea of golden cutting points. We begin with an elaborate three-qubit example in Section II-A and generalize it to arbitrary circuit bipartitions in Section II-B. Then, we proceed to conduct numerical experiments in Section III. Finally, we provide insight into future directions in Section IV.

II. QUANTUM CIRCUIT CUTTING

Quantum circuit cutting seeks a qubit wire (or multiple qubit wires) such that, upon removing, the circuit decomposes into several independent fragments. This is accomplished by performing quantum tomography—learning a quantum state—on the qubit upstream of the cut and recreating the inferred state in the downstream qubit. One can think of tomography as estimating the coefficient after expanding a quantum state with respect to some basis set, which we take to be the Pauli basis

$$\left\{ I = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \right. \\ \left. Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \right\} \quad (1)$$

for concreteness. By measuring the quantum state with respect to each element in the basis, one could piece together the state prior to the cut. Then, by initializing the downstream qubit incident to the cut into the eigenstate of each element in the basis, one could learn the behavior of the circuit given the initial state. Lastly, reweighting the results obtained through multiple initializations yields the behavior of an uncut circuit.

If a circuit undergoes K cuts, tomography has to be performed on all K qubits at the location of the cuts. The set of length- K Pauli strings is then an appropriate basis for a K -qubit system and the measurement-preparation scheme described above must be done for all elements in the set. Thus, the cost of performing circuit cutting scales exponentially with the number of cuts. In this section, we introduce a special class of circuits that contains a *golden*

cutting point and that enjoys a reduction in measurement complexity and classical postprocessing cost. We will demonstrate this via a three-qubit example (Section II-A) and proceed to formalize the phenomenon for arbitrary circuit bipartitions (Section II-B).

A. A Three-qubit Example

Consider the following three-qubit state

$$\rho = U_{23}U_{12}|000\rangle\langle 000|U_{12}^\dagger U_{23}^\dagger \quad (2)$$

where unitaries U_{12} and U_{23} can be arbitrary quantum gates (Figure 1). Notice that the qubit wire between the two gates can be *cut*, in the sense that we can rewrite the above state as

$$\rho = \frac{1}{2} \sum_{M_2 \in \mathcal{B}} \rho_{f_1}(M_2) \otimes \rho_{f_2}(M_2) \quad (3)$$

where the set $\mathcal{B} = \{I, X, Y, Z\}$ is the Pauli basis, and

$$\rho_{f_1}(M_2) = \text{tr}_2(M_2 U_{12}|00\rangle\langle 00|U_{12}^\dagger), \quad (4)$$

$$\rho_{f_2}(M_2) = U_{23}(M_2 \otimes |0\rangle\langle 0|)U_{23}^\dagger. \quad (5)$$

Here, M_2 refers to applying the operator M onto to the second qubit. This convention will be held throughout the manuscript where numerical subscripts denote the qubit(s) an operator is acting on.

However, as currently presented, the fragments are not quantum states as Pauli matrices are traceless operators. To obtain a physical interpretation, consider the eigendecomposition $M = rM^r + sM^s$ for each $M \in \mathcal{B}$ where $r, s \in \{+1, -1\}$ are the eigenvalues and M^r, M^s are the eigenstates. Then, we can rewrite Equation 3 as

$$\rho = \frac{1}{2} \sum_{\substack{M_2 \in \mathcal{B} \\ r, s = \pm 1}} r s \rho_{f_1}(M_2^r) \otimes \rho_{f_2}(M_2^s). \quad (6)$$

The eigenstates of $M \in \mathcal{B}$ have unit trace and can be regarded as quantum states.

Intuitively, one can understand circuit cutting as classically passing information through the qubit wire that is cut. Inside the summation (fix an M_2 of choice), the information stored in the second qubit in the direction M_2 is measured—taking the partial trace in ρ_{f_1} —and passed onto the remaining circuit—initializing into state M_2^s . Since the *circuit fragments* ρ_{f_1} and ρ_{f_2} can be simulated independently, through repeating the measurement-initialization scheme, we can effectively cut the original circuit, run fragments in parallel, and combine the results classically.

We are often interested in finding the expectation of the quantum state with respect to a quantum observable, O . Without loss of generality, assume an observable can

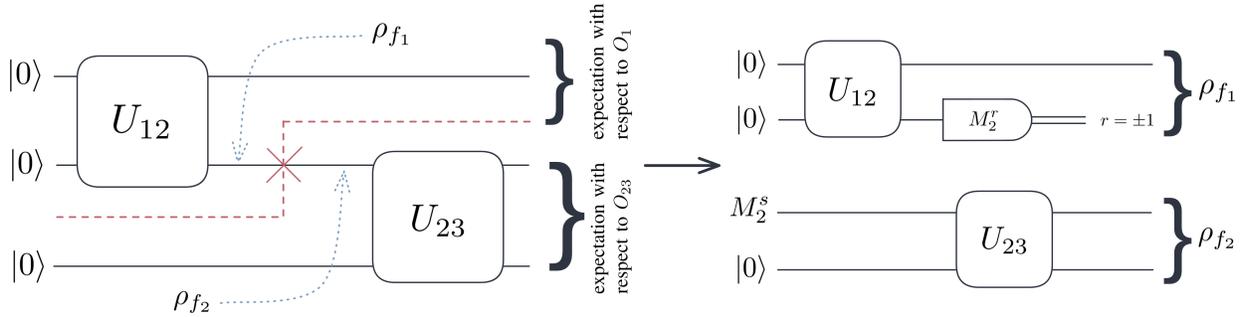


Fig. 1: Circuit diagram corresponding to the 3-qubit example

be decomposed as $O = O_1 \otimes O_{23}$. Then, we can rewrite the expectation in terms of ρ_{f_1} and ρ_{f_2} :

$$\text{tr}(O\rho) = \frac{1}{2} \sum_{M_2, r, s} r s \text{tr}(O_1 \rho_{f_1}(M_2^r)) \text{tr}(O_{23} \rho_{f_2}(M_2^s)) \quad (7)$$

$$= \frac{1}{2} \sum_{M_2, s} s \text{tr}(O_{23} \rho_{f_2}(M_2^s)) \sum_r r \text{tr}(O_1 \rho_{f_1}(M_2^r)). \quad (8)$$

The above summation involves 16 terms. Experimentally, the summation requires measuring the second qubit in the fragment upstream for each non-identity Pauli basis and initializing the first qubit in the downstream fragment to the respective eigenstates.

However, the computation above can potentially be reduced. Suppose that there exists an $M_{2,*} \in \mathcal{B}$ such that the last summation evaluates to zero, that is,

$$\sum_{r=\pm 1} r \text{tr}(O_1 \rho_{f_1}(M_{2,*}^r)) = 0. \quad (9)$$

Note that this can happen one of two ways:

- (i) Operator O_1 is orthogonal to the conditional state $\rho_{f_1}(M_{2,*}^r)$, i.e., $\text{tr}(O_1 \rho_{f_1}(M_{2,*}^r)) = 0$ for both $r = \pm 1$. For example, $O_1 = X$ and $U_{12}|00\rangle = (|00\rangle + |11\rangle)/\sqrt{2}$, the Bell state.
- (ii) The conditional state is in a subspace that ‘‘conveys no information’’ about the observable, specifically, $\text{tr}(O_1 \rho_{f_1}(M_{2,*}^r)) \neq 0$ but the weighted sum in Equation 9 leads to systematic cancellations. For example, $O_1 = |+\rangle\langle +|$ is a projector to the plus-state, and $U_{12}|00\rangle$ is again the Bell state.

Upon noting the reduction, the number of terms in the summation drops from 16 to 12. Moreover, since any term involving $M_{2,*}$ is neglected, there is no need to estimate the expectation in the downstream fragment ($\text{tr}(O_{23} \rho_{f_2}(M_{2,*}^s))$), which saves circuit evaluations that use the initial state $M_{2,*}^s$. We say that the cut is a *golden cutting point* if such reduction occurs.

B. Generalization

We formally extend the above case to more general bipartitions, but refrain from considering cutting schemes that result in more than three fragments (c.f. [26] for general expression of expectations with respect to circuit fragments). Suppose a N -qubit quantum circuit induces a state ρ . A set of K cuts, with injective function c mapping the i -th cut onto qubit $c(i) \in [N] = \{1, 2, \dots, N\}$. By cutting the K wires, the circuit can be bipartitioned into fragments f_1 and f_2 . Now, cutting requires passing information from K sets of basis \mathcal{B} through the cutting point, namely, measuring upstream qubits and preparing downstream qubits with respect to operators

$$\mathbf{M} = (M_{c(1)} \quad M_{c(2)} \quad \dots \quad M_{c(K)}) \in \mathcal{B}^K. \quad (10)$$

Furthermore, each operator M admits spectral decomposition. Letting

$$\mathbf{r} = (r_{c(1)} \quad r_{c(2)} \quad \dots \quad r_{c(K)}) \in \{-1, +1\}^K \quad (11)$$

be a tuple of eigenvalues, we define

$$\mathbf{M}^{\mathbf{r}} = (M_{c(1)}^{r(1)} \quad M_{c(2)}^{r(2)} \quad \dots \quad M_{c(K)}^{r(K)}) \quad (12)$$

to be the \mathbf{r} -th eigenstate of operator \mathbf{M} .

Using the above notation, we can compactly write the uncut state using the fragment-induced states $\rho_{f_i}(\mathbf{M})$ for $i = 1, 2$. Following Equation 6 gives the reconstruction formula in this bipartition case:

$$\rho = \frac{1}{2^K} \sum_{\substack{\mathbf{M} \in \mathcal{B}^K, \\ \mathbf{r}, \mathbf{s} \in \{-1, +1\}^K}} \prod_{i \in [K]} r_i s_i \rho_{f_1}(\mathbf{M}^{\mathbf{r}}) \otimes \rho_{f_2}(\mathbf{M}^{\mathbf{s}}). \quad (13)$$

Alternatively, for any desired quantum observable O , suppose the operator can be decomposed to accommodate the two fragments, i.e., $O = O_{f_1} \otimes O_{f_2}$ up to appropriate permutation of qubit indices. This is without loss of generality as expansions using Pauli strings would yield a linear combination of operators that are qubit-wise separable. Then, we can arrive at an expression

analogous to Equation 7 in terms of the fragments ρ_{f_i} and their respective observables O_{f_i} :

$$\begin{aligned} \text{tr}(O\rho) &= \\ \frac{1}{2^K} \sum_{\mathbf{M}, \mathbf{r}, \mathbf{s}} \prod_{i \in [K]} r_i s_i \text{tr}(O_{f_1} \rho_{f_1}(\mathbf{M}^{\mathbf{r}})) \text{tr}(O_{f_2} \rho_{f_2}(\mathbf{M}^{\mathbf{s}})). \end{aligned} \quad (14)$$

We now formally define the *golden circuit cutting point*.

Definition 1. A N -qubit circuit amenable to bipartition with K cuts has a golden cutting point if there exists $k^* \in [K]$ such that

$$\sum_{r_c(k^*)} r_c(k^*) \text{tr}(O_{f_1} \rho_{f_1}(\mathbf{M}^{\mathbf{r}})) = 0. \quad (15)$$

Note that golden cutting points are not necessarily unique. There can be multiple cuts with negligible bases and/or multiple negligible bases in one cut. Suppose there are K_g many golden cutting points and $K_r = K - K_g$ regular cutting points, the run time complexity of reconstructing the expectation in this bipartite case scales as $\mathcal{O}(4^{K_r} 3^{K_g})$. Moreover, there is no need of preparing the downstream fragments into the respective eigenstates, reducing the number of circuit evaluations from $\mathcal{O}(6^K)$ to $\mathcal{O}(6^{K_r} 4^{K_g})$. It is worth noting that the eigenstate preparation scheme is overcomplete and alternative bases, *e.g.* the symmetric informationally-complete (SICC) basis, can be used to achieve $\mathcal{O}(4^K)$ circuit evaluations without invoking golden circuit cutting formalism. However, employing the SICC basis would require more involved implementation, namely, solving linear systems, in order to construct the appropriate tensor during reconstruction.

We emphasize that golden cutting points as its current form is strictly a product of circuit design. That is, we can design circuits that admits golden cutting points at known locations (c.f. Section III). In reality, the existence of golden cutting points depends on the choice of the observable O , how it can be decomposed into form amenable to cutting, and the the state of the qubit prior to the cut. It is unlikely that an arbitrary circuit would exhibit such property without intricate designs, and we defer further discussion on finding golden cutting points to Section IV.

III. EXPERIMENTS

We demonstrate the capability of golden cutting points numerically. Specifically, experiments were designed to verify two aspects of our work: first, that our method does not sacrifice the correctness of reconstruction and second, that our method reduces the overall runtime of the algorithm. The circuit used in our experiment takes the form showed in Figure 2 where the sizes of the circuits were tailored to the size of the device it was

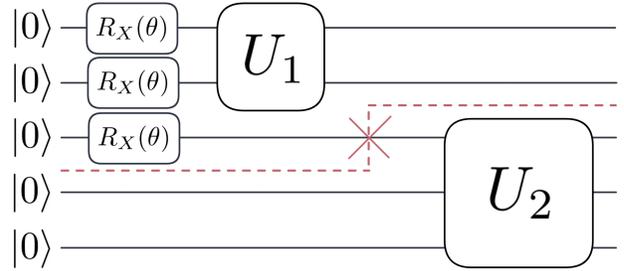


Fig. 2: Example 5-qubit circuit with golden cutting point. U_1 and U_2 are randomized circuits. Note that other odd-number circuit widths were used too, mainly 7-qubits split into subcircuits of 4 qubits each.

run on. Most commonly, we split 5- and 7-qubit circuits into fragments of 3 and 4 qubits respectively. In addition to quantum hardware, we also tested our method on the Aer simulator from Qiskit [27].

The generated circuits featured collections of RX gates with the rotation angle θ was chosen uniformly at random from the interval $[0, 6.28]$, as well as random gates generated using the `random_circuit()` function in Qiskit, denoted by U_1 and U_2 in Figure 2. In our experiments, we want to acquire the bitstring probability distribution generated by repeated measurements in the computational (Z) basis. Alternatively, to match the formalism of Section II-B, for all bitstrings $\hat{b} \in \{0, 1\}^n$, one can interpret the above as estimating the expectation of the projector observable $\Pi^{\hat{b}} = |\hat{b}\rangle\langle\hat{b}|$ where $n = 5, 7$ is the number of qubits in a circuit. This observable decomposes straightforwardly upon cutting:

$$\Pi^{\hat{b}} = \Pi^{\hat{b}_1} \otimes \Pi^{\hat{b}_2} \quad (16)$$

where \hat{b}_1 and \hat{b}_2 are bitstrings of length $\lfloor n/2 \rfloor$ and $\lceil n/2 \rceil$ respectively such that the concatenation recovers \hat{b} . By repeating many measurements in the computational basis and obtaining a bitstring distribution, we can then estimate the expectation of the projector observables. The restrictions imposed on our circuit ansatz creates a golden cutting point. In particular, the contribution of first fragment to the total expectation (with respect to the projector operator above, which has only diagonal components) conditioned on observing each eigenstate of the Pauli Y operator leads to components of equal magnitudes. Once weighed by the respective eigenvalues, the trace terms systematically cancels and induces a golden cutting point as defined in 15.

A. Verifying Reconstruction Accuracy

We first proceed to verify the correctness of our golden cutting point method. We performed noiseless (in the sense of hardware noise) simulations of the

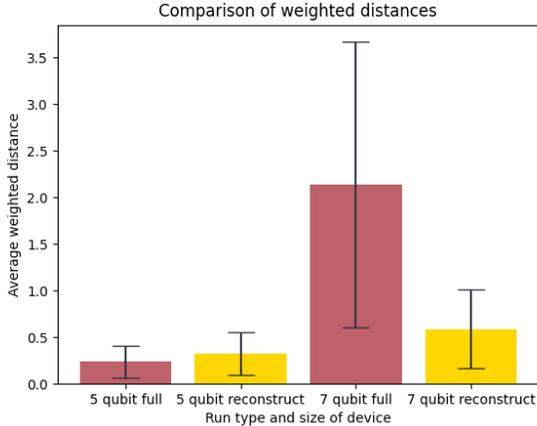


Fig. 3: Comparison of weighted distances for uncut circuit evaluations and reconstructed measurements from fragment data. Each bar was averaged over 10 independent trials. Each trial consisted of 10,000 shots per fragment. Error bars represent 95% confidence intervals.

full uncut circuit to obtain the ground truth distribution and compared our golden cutting point method to it. In addition, we ran the same circuits, both uncut and fragmented, on superconducting IBM Quantum devices [28]. To compare distributions, we now introduce a *weighted distance function* d_w :

$$d_w(p; q) = \sum_{x \in \mathcal{X}} \frac{(p(x) - q(x))^2}{q(x)}. \quad (17)$$

for distributions p and q with support \mathcal{X} . In the experiment, p was our “test” distribution and q was the “ground truth.” This weighted distance function penalizes large percentage deviations more than other metrics such as the total variational distance.

Using this weighted distance function, we determined the distance between the ground truth bitstring distribution from the Aer simulator and the distribution obtained by running the full circuit on quantum hardware. We then also determined the distance between the ground truth distribution and the distribution obtained by running our golden cutting method on quantum hardware. We repeated the experiment for machines of two different sizes—a 5-qubit device would run a 5-qubit circuit that is split into two 3-qubit subcircuits and a 7-qubit device would run a 7-qubit circuit split into two 4-qubit subcircuits. Each circuit configuration was repeated for 10 trials with each trial consisting of 10,000 shots per (sub)circuit. The results of the experiments are shown in Figure 3.

Because previous work has found that quantum circuit cutting has yielded a benefit in terms of fidelity, we expected similar results. We were surprised to find that

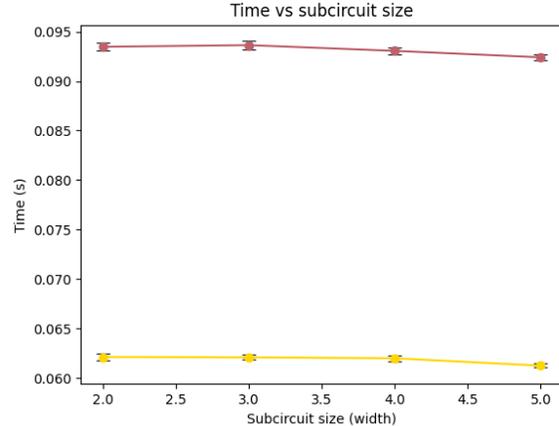


Fig. 4: Runtime comparison of circuits with (in gold) and without (in red) golden circuit cutting optimization. Each configuration was repeated for 1000 trials and each trial involved 1000 shots for each (sub)circuit. Error bars represent 95% confidence intervals.

such a benefit is non-existent on 5-qubit devices and not detectable within 95% confidence intervals on a 7-qubit device. This is likely due to the fact that our circuits were not particularly deep—only a few gates in each. As the circuit depth increases we should expect to see more of a benefit from circuit cutting. Additionally, an average of only 10 trials was performed, leaving our uncertainties relatively large. In our future work we will run more trials, potentially on larger devices, or with greater depths, to more precisely determine the true distance for this method. Regardless, in this work, we have verified that (within error) our method performs as well as full circuit execution on real hardware in terms of outputting the correct bitstring distribution.

B. Algorithm runtime

The presented golden cutting point method predicts a reduction in the number of measurements, thereby lowering the runtime of the procedure. In this experiment, we recorded the time taken for gathering fragment data and reconstructing them on a randomly generated circuit. We assumed the golden cutting point was known a priori (see Section IV for determining the existence of golden cutting points), and studied the runtime reduction from exploiting this knowledge. The results are illustrated in Figure 4.

We then repeated the same experiment using real devices available through the IBM Quantum Experience [28]. We report a speedup of 33 percent using our method as compared to the standard method [18], as shown in Figure 5. Specifically, we find that the average time for execution using the standard (without golden

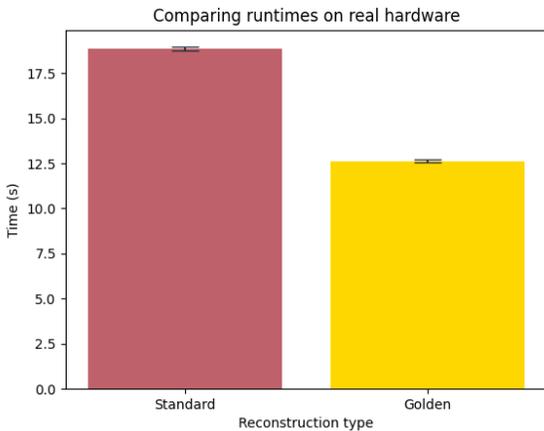


Fig. 5: Circuit cutting runtime with and without golden cutting point on quantum devices from IBM. Each bar represents the average runtime of 50 trials and each trial comprised of 1,000 shots for each (sub)circuit.

cutting point consideration) reconstruction method was 18.84 seconds whereas the golden cutting point method had a mean of 12.61 seconds. This reduction in run time is largely attributed to the reduced number of circuit evaluations. Particularly, we avoided having to execute a third of the total shots by neglecting one basis element, bringing the total number of circuit executions down from 4.5×10^5 to 3.0×10^5 . This demonstrates the applicability of our method: it is simple to implement and produces a considerable reduction in wall time needed to perform circuit reconstruction.

IV. CONCLUSION

In this paper, we introduced the idea of a golden cutting point where elements of a basis set can be neglected (Section II-B). Neglecting these basis elements allows for a considerable decrease in the runtime. The reduction can be attributed to (1) fewer terms are involved when combining measurement results from fragments, and (2) fewer circuit evaluations are needed in the fragment downstream of the golden cut. We verified the correctness and demonstrated the lowered runtime in both simulators and real quantum hardware (Section III). We observed that golden cutting points do not come at the cost of accuracy, and a statistically significant 33% decrease in runtime persists in both a simulated environment and quantum devices.

The introduction of the golden cutting point begs the question of applicable algorithms that have built-in golden cutting points. Variational circuits are a potential candidate due to their flexibility in the circuit ansatz. Although variational algorithms for combinatorial optimization or chemical simulations have great structural

assumptions given by the Hamiltonian, quantum machine learning circuits typically do not possess these constraints and therefore are likely more amenable to exploiting the golden cutting point formalism.

Moreover, detection of the existence of a golden cutting point is also an interesting one. In this work, we assumed the golden cutting point was known a priori. However, it is not obvious whether a basis can be neglected without simulating the circuit. Thus, we suspect there are methods for detecting golden cutting points “online” during the execution of the circuit cutting procedure through sequential empirical measurements. Note that this does not necessarily sacrifice the parallelism of circuit cutting as golden cutting points only affect the fragments directly incident to the cut. However, performing such a method would require further statistical analysis of acceptable error and the amplification of error through tensor contraction.

ACKNOWLEDGEMENTS

We acknowledge the use of IBM Quantum services for this work. The views expressed are those of the authors, and do not reflect the official policy or position of IBM or the IBM Quantum team. This material is based upon work supported by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers. The colors used in figures were based on Nord Theme [29]. Daniel T. Chen and Ethan H. Hansen contributed equally to this work.

REFERENCES

- [1] L. K. Grover, “A fast quantum mechanical algorithm for database search,” in *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, ser. STOC ’96. Association for Computing Machinery, 1996, pp. 212–219. [Online]. Available: <https://doi.org/10.1145/237814.237866>
- [2] P. W. Shor, “Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer,” *SIAM Journal on Computing*, vol. 26, no. 5, pp. 1484–1509, oct 1997. [Online]. Available: <https://doi.org/10.1137%2Fs0097539795293172>
- [3] A. Robert, P. K. Barkoutsos, S. Woerner, and I. Tavernelli, “Resource-efficient quantum algorithm for protein folding,” *npj Quantum Information*, vol. 7, no. 1, feb 2021. [Online]. Available: <https://doi.org/10.1038%2Fs41534-021-00368-4>
- [4] R. LaRose, A. Mari, S. Kaiser, P. J. Karalekas, A. A. Alves, P. Czarnik, M. E. Mandouh, M. H. Gordon, Y. Hindy, A. Robertson, P. Thakre, M. Wahl, D. Samuel, R. Mistri, M. Tremblay, N. Gardner, N. T. Stemen, N. Shammah, and W. J. Zeng, “Mitiq: A software package for error mitigation on noisy quantum computers,” *Quantum*, vol. 6, p. 774, aug 2022. [Online]. Available: <https://doi.org/10.22331%2Fq-2022-08-11-774>
- [5] M. H. Abobeih, Y. Wang, J. Randall, S. J. H. Loenen, C. E. Bradley, M. Markham, D. J. Twitchen, B. M. Terhal, and T. H. Taminiau, “Fault-tolerant operation of a logical qubit in a diamond quantum processor,” *Nature*, vol. 606, no. 7916, pp. 884–889, may 2022. [Online]. Available: <https://doi.org/10.1038%2Fs41586-022-04819-6>
- [6] A. A. Saki, A. Katarbarwa, S. Resch, and G. Umbrascu, “Hypothesis testing for error mitigation: How to evaluate error mitigation,” 2023. [Online]. Available: <https://arxiv.org/abs/2301.02690>

- [7] J. Gambetta, “Quantum-centric supercomputing: The next wave of computing,” Dec 2022. [Online]. Available: <https://research.ibm.com/blog/next-wave-quantum-centric-supercomputing>
- [8] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles, “Variational quantum algorithms,” *Nature Reviews Physics*, vol. 3, no. 9, pp. 625–644, aug 2021. [Online]. Available: <https://doi.org/10.1038/s42254-021-00348-9>
- [9] E. Farhi, J. Goldstone, and S. Gutmann, “A quantum approximate optimization algorithm,” *arXiv preprint arXiv:1411.4028*, 2014. [Online]. Available: <https://arxiv.org/abs/1411.4028>
- [10] D. A. Fedorov, B. Peng, N. Govind, and Y. Alexeev, “VQE method: a short survey and recent developments,” *Materials Theory*, vol. 6, no. 1, pp. 1–21, 2022. [Online]. Available: <https://doi.org/10.1186/s41313-021-00032-6>
- [11] B. Barak and K. Marwaha, “Classical Algorithms and Quantum Limitations for Maximum Cut on High-Girth Graphs,” in *13th Innovations in Theoretical Computer Science Conference (ITCS 2022)*, ser. Leibniz International Proceedings in Informatics (LIPIcs), M. Braverman, Ed., vol. 215. Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022, pp. 14:1–14:21. [Online]. Available: <https://drops.dagstuhl.de/opus/volltexte/2022/15610>
- [12] A. Warshel and M. Levitt, “Theoretical studies of enzymic reactions: dielectric, electrostatic and steric stabilization of the carbonium ion in the reaction of lysozyme,” *Journal of molecular biology*, vol. 103, no. 2, pp. 227–249, 1976. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/0022283676903119>
- [13] M. S. Gordon, D. G. Fedorov, S. R. Pruitt, and L. V. Slipchenko, “Fragmentation methods: A route to accurate calculations on large systems,” *Chemical reviews*, vol. 112, no. 1, pp. 632–672, 2012. [Online]. Available: <https://doi.org/10.1021/cr200093j>
- [14] W. Li, S. Li, and Y. Jiang, “Generalized energy-based fragmentation approach for computing the ground-state energies and properties of large molecules,” *The Journal of Physical Chemistry A*, vol. 111, no. 11, pp. 2193–2199, 2007. [Online]. Available: <https://doi.org/10.1021/jp067721q>
- [15] H. Li, W. Li, S. Li, and J. Ma, “Fragmentation-based QM/MM simulations: Length dependence of chain dynamics and hydrogen bonding of polyethylene oxide and polyethylene in aqueous solutions,” *The Journal of Physical Chemistry B*, vol. 112, no. 23, pp. 7061–7070, 2008. [Online]. Available: <https://doi.org/10.1021/jp800777e>
- [16] T. Peng, A. W. Harrow, M. Ozols, and X. Wu, “Simulating large quantum circuits on a small quantum computer,” *Physical Review Letters*, vol. 125, no. 15, oct 2020. [Online]. Available: <https://doi.org/10.1103/PhysRevLett.125.150504>
- [17] T. Ayrál, F.-M. Le Régent, Z. Saleem, Y. Alexeev, and M. Suchara, “Quantum divide and compute: Hardware demonstrations and noisy simulations,” in *2020 IEEE Computer Society Annual Symposium on VLSI (ISVLSI)*. IEEE, 2020, pp. 138–140. [Online]. Available: <https://ieeexplore.ieee.org/document/9155024>
- [18] T. Ayrál, F.-M. L. Régent, Z. Saleem, Y. Alexeev, and M. Suchara, “Quantum divide and compute: exploring the effect of different noise sources,” *SN Computer Science*, vol. 2, no. 3, pp. 1–14, 2021. [Online]. Available: <https://doi.org/10.1007/s42979-021-00508-9>
- [19] M. A. Perlin, Z. H. Saleem, M. Suchara, and J. C. Osborn, “Quantum circuit cutting with maximum-likelihood tomography,” *npj Quantum Information*, vol. 7, no. 1, p. 64, 2021. [Online]. Available: <https://doi.org/10.1038/s41534-021-00390-6>
- [20] Z. H. Saleem, T. Tomesh, M. A. Perlin, P. Gokhale, and M. Suchara, “Quantum Divide and Conquer for Combinatorial Optimization and Distributed Computing,” *arXiv preprint*, Jul. 2021. [Online]. Available: <https://arxiv.org/abs/2107.07532>
- [21] J. Liu, A. Gonzales, and Z. H. Saleem, “Classical simulators as quantum error mitigators via circuit cutting,” *arXiv preprint arXiv:2212.07335*, 2022. [Online]. Available: <https://arxiv.org/abs/2212.07335>
- [22] A. Lowe, M. Medvidović, A. Hayes, L. J. O’Riordan, T. R. Bromley, J. M. Arrazola, and N. Killoran, “Fast quantum circuit cutting with randomized measurements,” 2022. [Online]. Available: <https://arxiv.org/abs/2207.14734>
- [23] D. Chen, B. Baheri, V. Chaudhary, Q. Guan, N. Xie, and S. Xu, “Approximate quantum circuit reconstruction,” in *2022 IEEE International Conference on Quantum Computing and Engineering (QCE)*. IEEE, 2022, pp. 509–515.
- [24] G. Uchihara, T. M. Aamodt, and O. Di Matteo, “Rotation-inspired circuit cut optimization,” *arXiv preprint arXiv:2211.07358*, 2022. [Online]. Available: <https://arxiv.org/abs/2211.07358>
- [25] H.-Y. Huang, R. Kueng, and J. Preskill, “Predicting many properties of a quantum system from very few measurements,” *Nature Physics*, vol. 16, no. 10, pp. 1050–1057, 2020. [Online]. Available: <https://doi.org/10.1038/s41567-020-0932-7>
- [26] D. T. Chen, Z. H. Saleem, and M. A. Perlin, “Quantum divide and conquer for classical shadows,” *arXiv preprint arXiv:2212.00761*, 2022. [Online]. Available: <https://arxiv.org/abs/2212.00761>
- [27] A. tA-v *et al.*, “Qiskit: An open-source framework for quantum computing,” 2021.
- [28] “IBM Quantum,” 2023. [Online]. Available: <https://quantum-computing.ibm.com/>
- [29] S. Greb, “Nord theme.” [Online]. Available: <https://www.nordtheme.com/>