



Quantum Divide and Compute: Exploring the Effect of Different Noise Sources

Thomas Ayril¹ · François-Marie Le Régent^{1,2} · Zain Saleem³ · Yuri Alexeev³ · Martin Suchara³

Received: 7 December 2020 / Accepted: 6 February 2021 / Published online: 10 March 2021
© The Author(s) 2021

Abstract

Our recent work (Ayril et al. in Proceedings of IEEE computer society annual symposium on VLSI, ISVLSI, pp 138–140, 2020. <https://doi.org/10.1109/ISVLSI49217.2020.00034>) showed the first implementation of the Quantum Divide and Compute (QDC) method, which allows to break quantum circuits into smaller fragments with fewer qubits and shallower depth. This accommodates the limited number of qubits and short coherence times of quantum processors. This article investigates the impact of different noise sources—readout error, gate error and decoherence—on the success probability of the QDC procedure. We perform detailed noise modeling on the Atos Quantum Learning Machine, allowing us to understand tradeoffs and formulate recommendations about which hardware noise sources should be preferentially optimized. We also describe in detail the noise models we used to reproduce experimental runs on IBM’s Johannesburg processor. This article also includes a detailed derivation of the equations used in the QDC procedure to compute the output distribution of the original quantum circuit from the output distribution of its fragments. Finally, we analyze the computational complexity of the QDC method for the circuit under study via tensor-network considerations, and elaborate on the relation the QDC method with tensor-network simulation methods.

Keywords Quantum circuit compilation · Noise modeling · Simulation · NISQ

Introduction

The advent of Noisy Intermediate Scale Quantum (NISQ) technologies [2] makes multiqubit processors with modest but increasing numbers of qubits available. Google, IBM, and Intel have recently announced quantum computers with 72, 65, and 49 qubits, respectively [3–5]; and new systems with 50–200 qubits are expected to be commercially available in the next few years. However, our ability to use the hardware to solve interesting problems is lagging. Solving practical computational problems typically requires evaluating quantum circuits with many hundreds or even thousands

of qubits, exceeding the size of the devices. In addition, large gate errors and short qubit coherence times prevent accurate evaluations of deep circuits.

Despite the remarkable progress in manufacturing and controlling these small multiqubit systems, building hardware with a sufficiently high number of high-fidelity qubits remains an extremely challenging task. Engineering challenges worsen as the systems scale and are inherent for all major qubit technologies, including superconducting qubits (errors due to Josephson junction defects and spurious microwave resonances [6]), ion traps (susceptibility to noise and difficulty to address individual ions [7]), neutral atoms (motion of the atoms inside the lattice [8]), and quantum dots (difficulty to entangle multiple qubits [9, 10]).

Successfully solving practical computational problems can be achieved only by developing techniques that can simultaneously map large problems onto small qubit systems and mitigate the effects of noise. The Quantum Divide and Compute (QDC) approach is one such technique. In this approach, we divide a large and potentially deep quantum circuit to suit the number of qubits and coherence times available in current quantum hardware. We then perform

This article is part of the topical collection “Quantum Computing: Circuits Systems Automation and Applications” guest edited by Himanshu Thapliyal and Travis S. Humble.

✉ Thomas Ayril
thomas.ayral@atos.net

¹ Atos Quantum Laboratory, Les Clayes-sous-Bois, France

² Ecole Polytechnique, Palaiseau, France

³ Argonne National Laboratory, Lemont, IL, USA

the computations on the subcircuits obtained by this division on a quantum processor, and we finally recombine our output results to obtain the output of the original circuit. This allows us to compute the outputs of quantum circuits that are too deep or too wide to be run on existing small-scale quantum processors.

There has been some previous work related to this approach. Bravyi et al. [11] showed that a quantum circuit on $n + k$ qubits can be simulated by sparse circuits on n qubits and exponential classical processing that takes time $2^{O(k)}\text{poly}(n)$. A more general approach that allows fragmenting larger quantum circuits into smaller subcircuits was introduced in [12]. In this work, tensor-network techniques were used to show how to decompose a circuit with a large quantum volume [13] into smaller subcircuits with quantum volumes compatible with NISQ devices. The classical computing overhead of the circuit fragmenting techniques was reduced in [14], and maximum likelihood tomography was applied on top of the circuit fragmentation to ensure that the reconstructed probability distributions are strictly non-negative and normalized. This work also showed, with the help of classical simulations, that the QDC strategy, when combined with maximum likelihood tomography, can estimate the output of a clustered circuit with higher fidelity than the full circuit execution. In [15], a method was introduced to locate the optimal location of the cut (the location where the circuit should be fragmented). The QDC strategy was applied to commonly known circuits in quantum computing such as supremacy circuits, Grover and Bernstein-Vazirani circuits, and was shown to achieve a high quantum circuit evaluation fidelity.

The ultimate test for the quantum computing field—the ability to use controlled quantum systems to perform tasks surpassing what can be done using classical computers, also called quantum supremacy [16]—has received considerable attention from both the scientific community and the general public. The largest classical supercomputers are capable of reliably simulating quantum systems with approximately 50 qubits [17, 18], and there is evidence that devices with more than 50 qubits may be able to demonstrate quantum supremacy even in the presence of noise [19]. While quantum supremacy is not one of the goals of this work, the developed techniques will allow increasing the size of circuits that can be evaluated on quantum hardware as well as on quantum simulators run on classical hardware [20–23] by a constant factor. Consequently, it will be possible to evaluate quantum circuits with hundreds of qubits and use quantum algorithms to solve problems larger than ever before.

Circuit cutting naturally complements variational quantum-classical algorithms such as the Variational Quantum Eigensolver (VQE) [24, 25] and the Quantum Approximate Optimization Algorithm (QAOA) [26]. These approaches have successfully produced suitable quantum

circuits for optimization problems by combining shallow quantum circuits with classical processing; and they allow some control over the width, depth, and connectivity of the circuits. However, the quality of the approximate solutions produced by VQE and QAOA decreases as the width and depth of their circuits decreases, and solving most interesting problems still requires hundreds of qubits [27, 28].

Circuit cutting offers numerous benefits. First, the technique does not compromise the quality of the solution as the size of the subcircuits decreases (overhead may scale exponentially with the number of cuts, however). Second, the technique can be applied to any sparsely connected quantum circuit, irrespective of the structure of the problem. Third, circuit cutting has a close relationship with tensor network quantum simulation techniques that are used to address scalability limitations due to memory requirements that grow exponentially with the size of the simulated systems. Fourth, circuit cutting can enhance the performance of existing quantum-classical variational approaches because it can increase the size of the subproblems tackled by the variational quantum eigensolver.

In this article, we follow up on our previous work on the topic [1]: we start by giving a detailed derivation of the formula for the output reconstruction of the original circuit from the outputs of its fragments, and a description of the noise models we chose to reproduce the experimental results (“[Methods](#)”). We quantify the performance of the QDC method by recalling our previous results [1] on a 20-qubit IBM processor for different qubit counts and fragment sizes (“[Summary of Previous Results](#)”). Then in “[Results](#)”, based on noisy simulations, we quantify the differential influence of various noise sources such as readout error, gate error and decoherence on the success probability of the algorithm for different qubit counts and fragment sizes. Finally, we discuss the classical complexity of the method, its relation to tensor-network simulation approaches, and its implications for homogeneous and heterogeneous quantum computing.

Methods

Circuit Cutting

An algorithm that allows circuit cutting was first described in [12]. In this section, we provide a self-contained derivation that allows to compute the probability distribution of a circuit that has been fragmented into several smaller disconnected pieces. We first derive a formula that uses probability distributions of two fragments to obtain the probability distribution of the original circuit. We then generalize the formula for cases with more than two fragments.

Two-Fragment Case: Definitions

Let us consider a m -qubit circuit described as the following composition of operations:

$$\mathcal{O} = \mathcal{O}_A^a \circ \mathcal{O}_B^a \circ \mathcal{O}_A^b \circ \mathcal{O}_B^b,$$

where the support of superoperators \mathcal{O}_A^b and \mathcal{O}_B^b is a bipartition of the qubits; similarly, the support of \mathcal{O}_A^a and \mathcal{O}_B^a is a bipartition such that the two “a” (for “after”) sets differ from the “b” (for “before”) sets by one qubit. Without loss of generality, one can assume that up to a relabeling, the support of \mathcal{O}_A^b is q_0, \dots, q_n and that of \mathcal{O}_B^b is q_{n+1}, \dots, q_{m-1} , and the “a” supports, q_0, \dots, q_{n-1} and q_n, \dots, q_{m-1} (see Fig. 1a).

The final state of the circuit is given by the density matrix:

$$\rho = \mathcal{O}(\rho_0) = \mathcal{O}_A^a \circ \mathcal{O}_B^a \circ \mathcal{O}_A^b \circ \mathcal{O}_B^b(\rho_0)$$

where ρ_0 is the initial density matrix. The probability of measuring a state i with binary representation $i = (\hat{b}_0(i), \dots, \hat{b}_{m-1}(i))$ is given by

$$p(i) = \text{Tr}[\Pi_i \cdot \rho], \tag{1}$$

where Π_i is the projector on state i ($i = 0 \dots 2^m$). It can be expressed as $\Pi_i = |i\rangle\langle i| = \otimes_{k=0}^{m-1} |\hat{b}_k(i)\rangle\langle \hat{b}_k(i)|$, where $\hat{b}_k(i)$ is the value of the k th bit of i . We note that $\Pi_i^\dagger = \Pi_i$, and $\sum_i \Pi_i = \otimes_k \sum_{\hat{b}_k=0}^1 |\hat{b}_k\rangle\langle \hat{b}_k| = I$. Thus:

$$p(i) = \text{Tr}[\Pi_i^\dagger \cdot \mathcal{O}_A^a \circ \mathcal{O}_B^a \circ \mathcal{O}_A^b \circ \mathcal{O}_B^b(\rho_0)]. \tag{2}$$

We now switch to a Pauli-basis representation (see “Appendix A” for a reminder). Using Eq. (16), we get

$$p(i) = 2^m \langle \langle \Pi_i | \mathcal{R}_A^a \mathcal{R}_B^a \mathcal{R}_A^b \mathcal{R}_B^b | \rho_0 \rangle \rangle \tag{3}$$

where $\mathcal{R}_{A/B}^{a/b}$ is the Pauli transfer matrix (PTM) representation of superoperator $\mathcal{O}_{A/B}^{a/b}$.

Bipartite Splitting Formula

Basic formula We now derive the splitting formula. Let us decompose the one-qubit PTM representation of the identity superoperator as

$$\mathcal{R}_I = \sum_{\alpha=X,Y,Z} \sum_{bb' \in \{0,1\}} \tilde{\gamma}_\alpha^{bb'} |\sigma_\alpha^b\rangle\langle \sigma_\alpha^{b'}|, \tag{4}$$

where $|\sigma_\alpha^b\rangle\rangle$ are the (real) coordinates in the Pauli basis of the density matrix corresponding to the b th eigenvector $|\psi_\alpha^b\rangle$ of Pauli matrix σ_α . The $\tilde{\gamma}$ tensor is given by $\tilde{\gamma}_X^{bb'} = \tilde{\gamma}_Y^{bb'} = 2\delta_{bb'} - 1$ and $\tilde{\gamma}_Z^{bb'} = 2\delta_{bb'}$.

Inserting \mathcal{R}_I (acting on qubit q_n) in the expression for the probability, Eq. (3), we obtain

$$\begin{aligned} p(i) &= 2^m \langle \langle \Pi_i | \underbrace{\mathcal{R}_A^a}_{q_0 \dots q_{n-1}} \underbrace{\mathcal{R}_B^a}_{q_n \dots q_{m-1}} \underbrace{\mathcal{R}_I}_{q_n} \underbrace{\mathcal{R}_A^b}_{q_0 \dots q_n} \underbrace{\mathcal{R}_B^b}_{q_{n+1} \dots q_{m-1}} | \rho_0 \rangle \rangle \\ &= 2^m \sum_{\alpha=X,Y,Z} \sum_{bb' \in \{0,1\}} \tilde{\gamma}_\alpha^{bb'} \\ &\quad \times \langle \langle \Pi_i |_{q_0 \dots q_{n-1}} \langle \langle \Pi_i |_{q_n \dots q_{m-1}} \underbrace{\mathcal{R}_A^a}_{q_0 \dots q_{n-1}} \underbrace{\mathcal{R}_B^a}_{q_n \dots q_{m-1}} | \sigma_\alpha^b \rangle \rangle_{q_n} \\ &\quad \times \langle \langle \sigma_\alpha^{b'} |_{q_n} \underbrace{\mathcal{R}_A^b}_{q_0 \dots q_{n-1}} \underbrace{\mathcal{R}_B^b}_{q_{n+1} \dots q_{m-1}} | \rho_0 \rangle \rangle_{q_0 \dots q_n} | \rho_0 \rangle \rangle_{q_{n+1} \dots q_{m-1}} \\ &= 2^m \sum_{\alpha=X,Y,Z} \sum_{bb' \in \{0,1\}} \tilde{\gamma}_\alpha^{bb'} 2^{-n-1} p_A^\alpha(i_{|0 \dots n-1}; b') 2^{-m+n} \\ &\quad \times p_B^{ab}(i_{|n \dots m-1}). \end{aligned}$$

We thus obtain the final formula (with $i = (\hat{b}_0 \dots \hat{b}_{m-1})$):

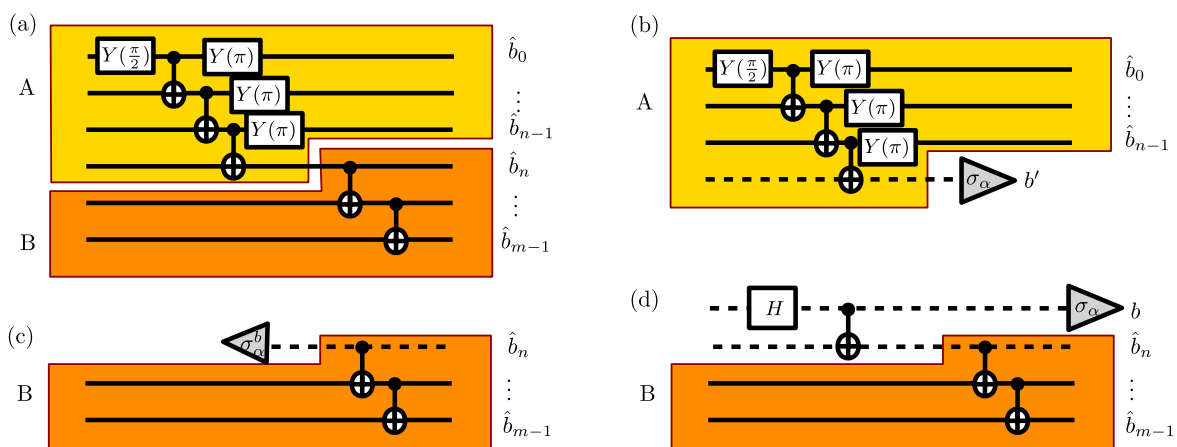


Fig. 1 Cutting sketch in the two-fragment case. **a** Original circuit, **b** Upper fragment of the circuit, **c** lower fragment of the circuit, **d** lower fragment in its Bell-state variant

$$p(\hat{b}_0 \dots \hat{b}_{m-1}) = \frac{1}{2} \sum_{\alpha=X,Y,Z} \sum_{bb' \in \{0,1\}} \tilde{\gamma}_\alpha^{bb'} p_A^\alpha(\hat{b}_0 \dots \hat{b}_{n-1}; b') \times p_B^{ab}(\hat{b}_n \dots \hat{b}_{m-1}) \tag{5}$$

with

$$p_A^\alpha(\hat{b}_0 \dots \hat{b}_{n-1}; b') \equiv 2^{n+1} \langle \langle \Pi_{\hat{b}_0 \dots \hat{b}_{n-1}} | \langle \langle \sigma_\alpha^{b'} |_{q_n} \mathcal{R}_A | \rho_0 \rangle \rangle_{q_0 \dots q_n} \rangle \tag{6}$$

$$p_B^{ab}(\hat{b}_n \dots \hat{b}_{m-1}) \equiv 2^{m-n} \langle \langle \Pi_{\hat{b}_n \dots \hat{b}_{m-1}} | \mathcal{R}_B | \sigma_\alpha^b \rangle \rangle_{q_{n+1} \dots q_{m-1}}, \tag{7}$$

where we have regrouped $\mathcal{R}_A \equiv \mathcal{R}_A^a \mathcal{R}_A^b$ and $\mathcal{R}_B \equiv \mathcal{R}_B^a \mathcal{R}_B^b$. In other words, $p_A^\alpha(\hat{b}_0 \dots \hat{b}_{n-1}; b')$ is the probability of measuring bitstring $\hat{b}_0 \dots \hat{b}_{n-1}, b'$ when measuring the final state of fragment A with a measurement on axis α for qubit q_n (see Fig. 1b), and $p_B^{ab}(\hat{b}_n \dots \hat{b}_{m-1})$ is the probability of measuring bitstring $\hat{b}_n \dots \hat{b}_{m-1}$ when measuring the final state of fragment B with qubit q_n initially prepared in the b th eigenstate of Pauli matrix σ_α (see Fig. 1c).

Variation using Bell pair We now derive a different expression based on the following idea: instead of preparing both eigenstates of σ_α , one can use an ancilla qubit, prepare a Bell state, and measure the value of the ancilla along measurement axis α and obtain an equivalent result, with a slightly different expression.

Switching from the Pauli-basis expression back to the original representation, Eq. (7) is equivalent to

$$p_B^{ab}(i) = \text{Tr}[\Pi_i \mathcal{O}_B(\sigma_\alpha^b \otimes \rho_0)]$$

where $\sigma_\alpha^b = |\psi_\alpha^b\rangle\langle\psi_\alpha^b|$. Let us decompose

$$|\psi_\alpha^b\rangle = \sum_{k \in \{0,1\}} \langle k | \psi_\alpha^b \rangle |k\rangle$$

then

$$\begin{aligned} p_B^{ab}(i) &= \sum_{kk'} \langle k | \psi_\alpha^b \rangle \langle \psi_\alpha^b | k' \rangle \text{Tr}[\Pi_i \cdot \mathcal{O}_B(|k\rangle\langle k'| \otimes \rho_0)] \\ &= \sum_{kk'} \langle \psi_\alpha^{b*} | k \rangle \langle k' | \psi_\alpha^{b*} \rangle \\ &\quad \times \text{Tr}[(I \otimes \Pi_i) \cdot (I \otimes \mathcal{O}_B)(I \otimes |k\rangle\langle k'| \otimes \rho_0)] \\ &= \text{Tr}\left[(|\psi_\alpha^{b*}\rangle\langle\psi_\alpha^{b*}| \otimes \Pi_i) \right. \\ &\quad \left. \times (I \otimes \mathcal{O}_B) \left(\sum_{kk'} |k\rangle\langle k'| \otimes |k\rangle\langle k'| \otimes \rho_0 \right) \right] \\ &= 2\text{Tr}[(\Pi_\alpha^{b*} \otimes \Pi_i) \cdot (I \otimes \mathcal{O}_B)(\rho_{\Phi^+} \otimes \rho_0)] \end{aligned}$$

where $\Pi_\alpha^{b*} = |\psi_\alpha^{b*}\rangle\langle\psi_\alpha^{b*}|$ is the projector onto the complex conjugate of the b th eigenstate of the σ_α Pauli matrix, and ρ_{Φ^+} is the density matrix of the Bell state

$$|\Phi^+\rangle \equiv \frac{1}{\sqrt{2}} \sum_{k=0,1} |kk\rangle. \tag{8}$$

In the second line, we have added an ancilla qubit. Now, let us note that for $\alpha = X, Z$, $|\psi_\alpha^b\rangle = |\psi_\alpha^{b*}\rangle$ (the eigenvector is real-valued), while $|\psi_Y^{b*}\rangle = |\psi_Y^{1-b}\rangle$, and let us define

$$\hat{p}_B^\alpha(b; i) \equiv \text{Tr}[\Pi_\alpha^b \otimes \Pi_i (I \otimes \mathcal{O}_B)(\rho_{\Phi^+} \otimes \rho_0)]. \tag{9}$$

Then

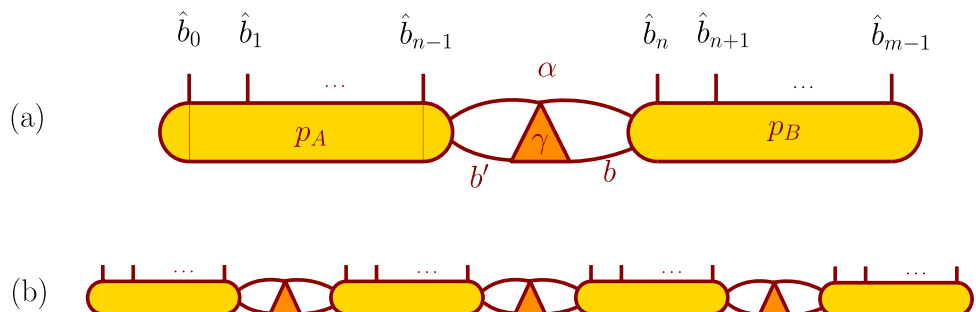
$$p_B^{ab}(i) = \begin{cases} 2\hat{p}_B^\alpha(i; b) & \alpha = X, Z \\ 2\hat{p}_B^\alpha(i; 1-b) & \alpha = Y. \end{cases}$$

Thus, after relabeling $b \rightarrow 1-b$ for $\alpha = Y$ in the final formula Eq. (5), we finally obtain the final expression:

$$p(\hat{b}_0 \dots \hat{b}_{m-1}) = \sum_{\alpha=X,Y,Z} \sum_{bb' \in \{0,1\}^2} \gamma_\alpha^{bb'} p_A^\alpha(\hat{b}_0 \dots \hat{b}_{n-1}; b') p_B^\alpha(b; \hat{b}_n \dots \hat{b}_{m-1}). \tag{10}$$

where $\gamma_X^{bb'} = 2\delta_{bb'} - 1$, $\gamma_Y^{bb'} = -\gamma_X^{bb'}$ and $\gamma_Z^{bb'} = 2\delta_{bb'}$.

Fig. 2 Graphical representation of the contraction formula. **a** Two fragment case. **b** Multifragment case for the GHZ circuit shown in Fig. 1a



The graphical representation for such a contraction is shown in Fig. 2a.

Multi-fragment Case

The formula for the multi-fragment case can be inferred from that of the two-fragment case: the procedure sketched for the two-fragment case can be recast in more generic terms, as described in [12]. This is done by considering the directed acyclic graph $G = (V, E)$ corresponding to the quantum circuit at hand (see Fig. 3 for an illustration of the procedure). Its vertices V are quantum operations such as qubit initialization, measurement and gates. The cutting procedure amounts to finding a subset $E' \subset E$ of M (directed) edges in this graph whose removal leads to K disconnected directed acyclic graphs $\{G^{(i)} = (V_i, E_i)\}_{i=1 \dots K}$. In each disconnected graph, $n_i + m_i$ vertices have a dangling edge corresponding to the original n_i incoming and m_i outgoing edges connecting it to the rest of the original graph, with $\sum_i n_i = \sum_i m_i = M$. One then adds a measurement along axis α_k ($\alpha_k = X, Y, Z$) as a termination to each outgoing dangling edge ($k = 1 \dots n_i$), and a Bell-state gadget (as described in the previous section), whose ancilla line is terminated by an α'_k -measurement, to each incoming dangling edge. Translating the family of graphs $G^{\alpha_1 \dots \alpha_{n_i}, \alpha'_1 \dots \alpha'_{m_i}}$ back to quantum circuits $\mathcal{C}^{\alpha_1 \dots \alpha_{n_i}, \alpha'_1 \dots \alpha'_{m_i}}$, we can sample (using a quantum computer) the corresponding probability distributions. We denote as

$$p_i^{\alpha_1 \dots \alpha_{n_i}, \alpha'_1 \dots \alpha'_{m_i}}(b_1, \dots, b_{n_i}; s; b'_1, \dots, b'_{m_i})$$

the probability of measuring bitstring $b_1, \dots, b_{n_i}; s; b'_1, \dots, b'_{m_i}$, with $s = (\hat{b}_1 \dots \hat{b}_{p_i})$ a bitstring corresponding to the state of “final” qubits of circuit $\mathcal{C}^{(i)}$, and (b_1, \dots, b_{n_i}) (resp. b'_1, \dots, b'_{m_i}) the bitstrings corresponding to the measured value for the measurements on the incoming (resp. outgoing) edges of sub-graph $G^{(i)}$ after pre-measurement rotations along axes $\alpha_1 \dots \alpha_{n_i}, \alpha'_1 \dots \alpha'_{m_i}$.

The final probability distribution is obtained by contracting the tensor network defined by the graph $\hat{G} = (\hat{V}, \hat{E})$, with $|\hat{V}| = K + M$ and $|\hat{E}| = 2M$. Here, K “fragment” vertices

correspond to the K disconnected components $\{G^{(i)}\}$, and M “connecting” vertices to the M removed edges. The $2M$ edges connect each of the K “fragment” vertices via one of the M “connecting” vertices. To each “fragment” vertex, we associate a distribution p_i , while to each “connecting” vertex, we associate a γ tensor [as defined below Eq. (10)].

We give an example of such a tensor network for the Greenberger–Horne–Zeilinger (GHZ) circuit we considered in our previous work as well in Fig. 2 b: in this case, the underlying graph turns out to be linear. We also show, in Fig. 3, an example with a more complex circuit and the resulting, more complex tensor network. Here, $K = 3$ and $M = 3$.

The contraction of these networks yields the sought-after distribution. The classical complexity of carrying out this contraction will be discussed in “Contraction Complexity and Relation to Tensor-Network Simulation”.

Noisy Simulation

NISQ processors are characterized by a substantial level of noise. In this section, we describe what noise processes we took into account in our simulation of the IBM Johannesburg quantum processor.

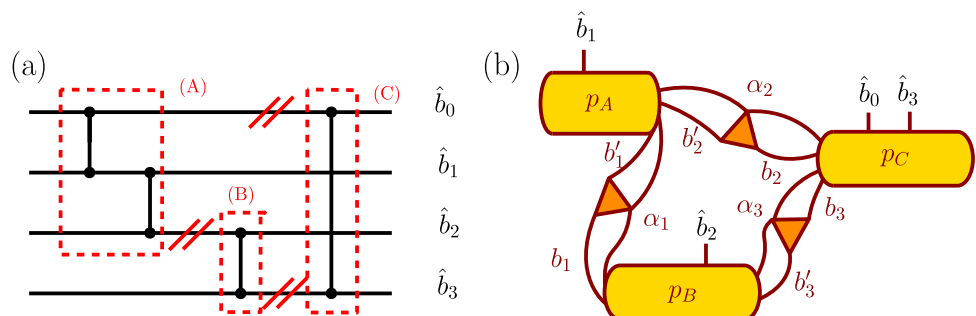
In this study, we focus on the noise processes whose quantitative levels are reported by the hardware manufacturer, IBM (see Table 1 for a summary of the numerical values used in the noisy simulations below). This pragmatic approach is justified a posteriori by the reasonable

Table 1 Johannesburg processor metrics, as retrieved from IBM Quantum Experience on March 5th, 2020

Parameter	Value
Readout error rate γ	4.1%
One-qubit-gate error rate $\epsilon_{\text{avg}}^{(1)}$	0.041%
Two-qubit-gate error rate $\epsilon_{\text{avg}}^{(2)}$	0.202%
Relaxation time T_1	65 μs
Dephasing time T_2	70 μs

All rates are averages over all the qubits/qubit pairs

Fig. 3 Graphical representation of the contraction formula for a generic case (here with three fragments). **a** Sketch of the fragmentation of a four-qubit circuit in three fragments. **b** Corresponding tensor network to contract to get final distribution



agreement of our numerical simulations with the experimental data (see Ref. [1], and “Results”). It should nevertheless be emphasized that (i) it uses rather simple noise models, that should be compared to noise models extracted from a full process tomography of the processor, and that (ii) it excludes some noise processes that are suspected to affect the final quantum state distribution in a non-negligible way, e.g., crosstalk (spatially correlated noise) and temporally correlated noise (like 1/f noise).

The most prominent source of error in today’s superconducting processors is the readout error. The duration of the dispersive readout conducted in transmon processors, of the order of a few microseconds, makes for a higher probability of error, most notably of the relaxation (or amplitude damping) type. We thus model the readout process as a two-outcome POVM corresponding to an amplitude-damping quantum channel of duration τ followed by a perfect Z-axis measurement: $\{E, I - E\}$, with $E = \begin{pmatrix} 0 & 0 \\ 0 & 1 - \gamma \end{pmatrix}$. The duration τ is adjusted so as to obtain a readout error rate $\gamma = 1 - e^{-\tau/T_1}$ that matches the readout error rate reported by IBM. With $\gamma = 4.1\%$ and $T_1 = 65 \mu\text{s}$, we find $\tau = 2.75 \mu\text{s}$, a duration that is consistent with the usual measurement durations of dispersive readout processes. Note that this noise model does not include measurement crosstalk effects [29].

Another source of error is gate noise, i.e. gate imperfections. Here, since the hardware manufacturer only reports average 1- and 2-qubit gate error rates, we picked the simplest noise process to model gate noise, namely depolarizing noise with a depolarization probability adjusted so that the average process fidelity \mathcal{F}_{avg} matches the qubit-averaged average error rates $\epsilon_{\text{avg}} = 1 - \mathcal{F}_{\text{avg}}$ reported by the hardware maker. We recall that the one-qubit depolarizing noise process is characterized by the following Kraus operators:

$$K_0^D = \sqrt{1 - p_{(1)}^D} I,$$

$$K_i^D = \sqrt{p_{(1)}^D} \sigma_i, \quad i = 1, 2, 3,$$

where σ_i denote the Pauli spin matrices. We model two-qubit depolarization processes as a tensor product of the one-qubit depolarizing noise. Let us stress that more structured, and therefore more accurate, noise models could be extracted from quantum process tomography methods, at the cost of a larger characterization overhead. Furthermore, this noise model does not include any crosstalk effects (see, e.g. [30]), despite evidence that they play some role in today’s NISQ processors.

Finally, we include the effect of decoherence on idle qubits, i.e. qubits that are not being acted upon by a quantum gate, but are nevertheless coupled to the outside environment. This decoherence can be decomposed into two main

types, namely relaxation and dephasing. Relaxation (also known as amplitude damping or, in other contexts, spontaneous emission) causes excited qubits (i.e. in state $|1\rangle$) to relax to their ground state ($|0\rangle$) with a probability that is characterized by a time T_1 : $p_{\tau_{\text{idle}}}^{\text{A.D}} = 1 - e^{-\tau_{\text{idle}}/T_1}$, namely, the longer the idling duration τ_{idle} , the higher the probability of a relaxation event. Similarly, dephasing events cause the two components $|0\rangle$ and $|1\rangle$ of a superposed state to acquire an unwanted dephasing with a certain probability. Under simplifying assumptions about the power spectral density (PSD) of the qubit-environment system, namely the assumption of a white noise PSD, this probability is given by $p_{\tau_{\text{idle}}}^{\text{P.D}} = 1 - e^{-2\tau_{\text{idle}}/T_\phi}$, with $\frac{1}{T_\phi} = \frac{1}{T_2} - \frac{1}{2T_1}$. We note that this is a quite strong simplification, as actual transmon processors are known to have a PSD that deviates from white noise, with, most notably, a sizable pink (1/f) noise component (see, e.g [31] for a review) that leads to a deviation to the exponential decay of the formula we used. Let us also stress that such a noise modeling does not take into account temporally correlated noise. As a reminder, here are the Kraus operators associated with amplitude damping and (pure) dephasing:

$$K_0^{\text{A.D}} = \begin{bmatrix} 1 & 0 \\ 0 & \sqrt{1 - p_{\tau_{\text{idle}}}^{\text{A.D}}} \end{bmatrix}, K_1^{\text{A.D}} = \begin{bmatrix} 0 & \sqrt{p_{\tau_{\text{idle}}}^{\text{A.D}}} \\ 0 & 0 \end{bmatrix},$$

$$K_0^{\text{P.D}} = \begin{bmatrix} 1 & 0 \\ 0 & \sqrt{1 - p_{\tau_{\text{idle}}}^{\text{P.D}}} \end{bmatrix}, K_1^{\text{P.D}} = \begin{bmatrix} 0 & 0 \\ 0 & \sqrt{p_{\tau_{\text{idle}}}^{\text{P.D}}} \end{bmatrix}.$$

The values we used for T_1 and T_2 are reported in Table 1.

The noisy simulations are conducted on the Atos Quantum Learning Machine (QLM), a classical supercomputing platform dedicated to writing, simulating and optimizing quantum algorithms [22].

Before simulating the circuits resulting from the fragmentation procedure described in the previous section, we use the QLM’s *Nnizer* plugin to compile the circuits, i.e. most notably to adapt them to the Johannesburg processor’s restricted qubit topology (shown in Fig. 4). Then, we perform noisy simulations using a density-matrix-based noise simulator that uses a dense representation of the density matrix ρ of the qubit register.

Results

Summary of Previous Results

In [1], we investigated the performance of the circuit-cutting procedure for a simple GHZ-type circuit shown in Fig. 1a. As a proxy for the quality of the procedure, we chose the quantity

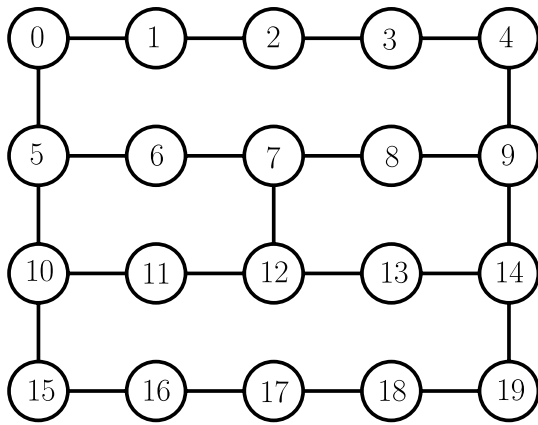


Fig. 4 Qubit connectivity map of the Johannesburg processor. Edges are shown between qubit pairs coupled via a resonator that allows application of the two-qubit CNOT gate

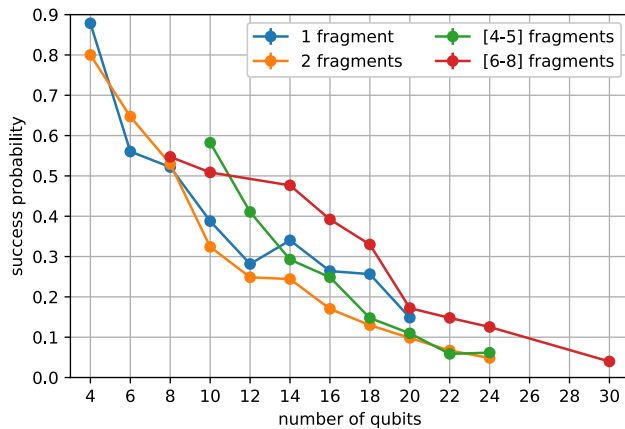


Fig. 5 Success probability as a function of circuit size (number of qubits) for various numbers of fragments using IBM’s Johannesburg processor

$$P_{\text{success}} \equiv p(|0\rangle^{\otimes m/2}|1\rangle^{\otimes m/2}) + p(|1\rangle^{\otimes m/2}|0\rangle^{\otimes m/2}), \quad (11)$$

which, given the GHZ circuit at hand, is unity in the absence of any noise.

We carried out the procedure both using an actual 20-qubit processor, IBM Johannesburg, and using the Atos Quantum Learning Machine’s noisy simulator.

The experimental success probabilities, shown in Fig. 5, display two clear trends: on the one hand, increasing the number of qubits leads to a decreasing success probability. This trend can be accounted for by the fact that increasing the number of qubits increases the number of gates of the circuit, and thus the sensitivity to gate errors and environmental decoherence. On the other hand, increasing the number of fragments in general leads to an improved success probability: the 6–8 fragment success probabilities are larger than the success probabilities obtained for lower numbers

of fragments (with some exceptions to this observation: the one-fragment success probability often exceeds that of the 2 and 4–5 fragment cases, maybe due to compiler optimizations on the hardware side for circuits with larger numbers of qubits; we also note a point at $n_{\text{qubits}} = 10$ where the 4–5 fragment success probability exceeds that of the 6–8 fragment case). This trend can be ascribed to the smaller gate count of each individual fragment, and thus a reduced sensitivity to errors. This smaller gate count not only comes from the mere cutting procedure, but also from the fact that smaller circuits better match the limited connectivity (Fig. 4) of the Johannesburg chip. Conversely, larger circuits need to be compiled to fulfill the connectivity constraints, leading to larger gate counts.

To substantiate these interpretations, we performed noisy simulations with noise models established using the constructor’s calibration data (Table 1). We show the results in Fig. 6: a 20% agreement is found between the noisy simulations and the experimental data. In particular, the drops in success probability, which can be traced back to connectivity-related insertions of SWAP gates, are reproduced. We note that the error bars coming from the finite number of shots (8192) used for each fragment are contained within the data symbols.

Analysis of the Influence of the Different Noise Types

In this section, we study and compare the differential impact of all the noise types we have previously taken into account: gate imperfections, idling and readout errors. Our goal is to understand which types of noise have a particularly severe influence on the fidelity of the fragmenting procedure and to formulate recommendations as to which noise types should be addressed first if one wants to make the most of the fragmenting procedure. Hence, we study the influence of the three noise types by simulating better readout measurements (Fig. 7), better gates (Fig. 8) and a better coherence time (Fig. 9).

Faster readout. First, we analyze the impact of readout errors by decreasing the duration τ of the measurements on all the subcircuits generated by the splitting procedure. Readout error is at present the largest source of errors in superconducting processors, with error rates as high as a few percent. It is thus reasonable to assume that large experimental efforts are going to be made to reduce this error rate. Here, we suppose the reduction in readout error rate to originate from a reduction of the readout duration (in practice by a factor 5), although it would be equivalent, in this noise model that assumes the errors to come only from an amplitude damping noise, to keep the readout duration fixed and to increase the T1 coherence time (by the same factor 5). In reality, progress is being made on both fronts (see, e.g

Fig. 6 Success probability as a function of circuit size (number of qubits) for various numbers of fragments using IBM’s Johannesburg processor (solid black lines) and Atos QLM noisy simulation (dashed blue lines). The black integers next to each black disk indicate the maximum fragment size (in number of qubits)

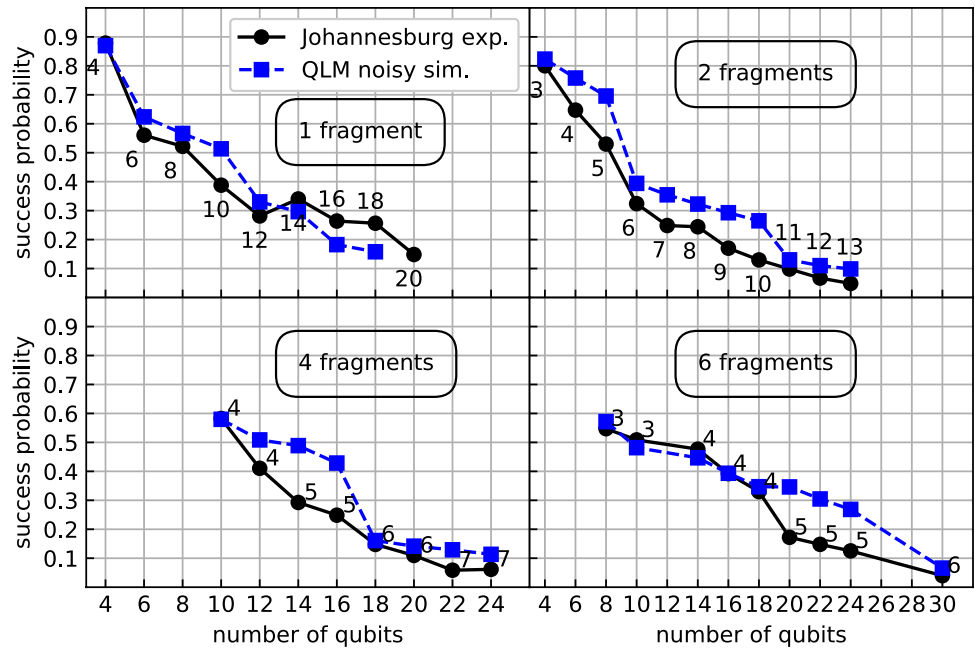
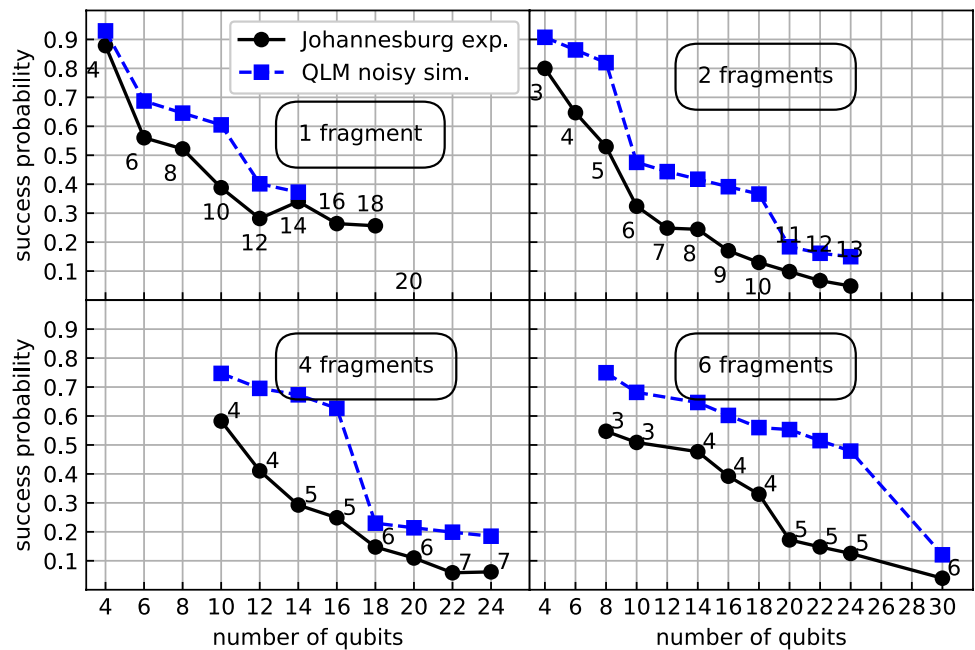


Fig. 7 Effect of better readout: Same as Fig. 5, but with a readout duration divided by 5



[32, Fig. 2.c], for the increasing T1 trend, and [33] for recent efforts towards faster measurements).

We see in Fig. 9 that better readout improves the overall success probability all the more as the fragment number is large. The difference between the solid and the dashed lines qualitatively increases with the number of readout measurements used, and consequently the number of fragments. Indeed, more fragments necessitate more measurements to characterize the quantum state of each fragment. Nevertheless, we still see drops in the evolution of the success probability with the number of qubits. It can be explained

by the topology constraints that require the use of several SWAP gates when we try to perform gates between physical qubits that are not adjacent. This calls the study of the next paragraph.

Better gates. To model the use of better gates, we choose to lower the amplitude of the depolarizing channel by dividing the depolarizing error rate by a factor of 5. The limited gate fidelity is the second major source of errors in superconducting processors. It comes from calibration errors as well as decoherence. Here, we mimic the improvement in gate quality by simply dividing the error rate by a factor of

Fig. 8 Effect of better gates: Same as Fig. 5, but with a depolarizing error per gate divided by 5

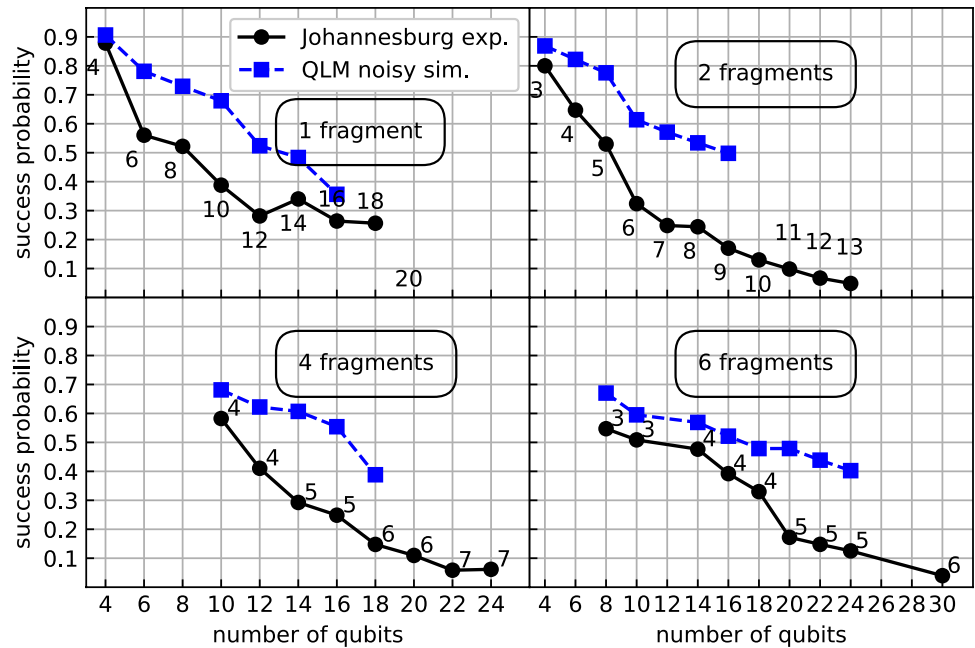
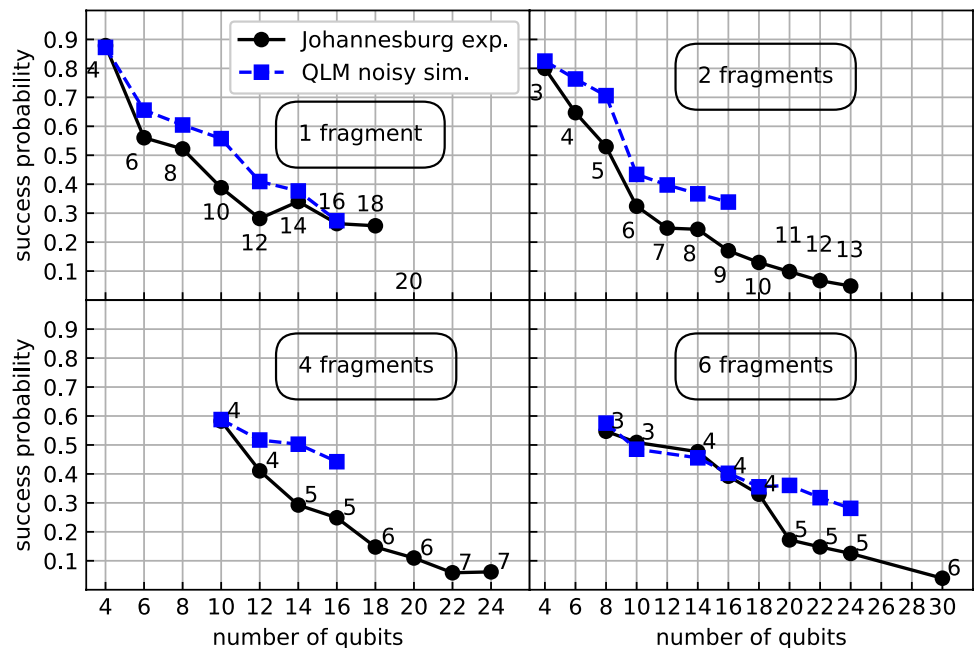


Fig. 9 Effect of better coherence: Same as Fig. 5, but with T_1 and T_2 coherence times multiplied by 5



5. Such a factor is realistic, in view of the improvements in gate qualities of superconducting processors in the recent years, and of the variability in the error rates reported by the hardware providers (the two-qubit error rates reported for IBM Johannesburg [34], Google Sycamore [35, Fig.2, Table II] and Rigetti Aspen 7 [36], are, respectively, 0.2%, 0.62% and 4.8%).

The results of this change in the noise model can be seen in Fig. 6. We notice that the slope is more regular as the number of qubits increases. Indeed, a smoothing of the “drops” in success probability is observed. These drops

were the consequence of performing a gate between qubits that are not adjacent in the connectivity map (Fig. 4) and that require using several SWAP gates. Thus, better gates help mitigate the effect of topology. The insertion of additional SWAP gates because of topology constraints becomes less detrimental to the overall success probability when the inserted gates are of good fidelity.

Better coherence. Finally, to understand the impact of coherence on the splitting procedure, we increase the relaxation time T_1 and the dephasing time T_2 by multiplying them by a factor of 5 (see “Noisy Simulation” for a

definition of the corresponding Kraus operators). Decoherence errors indeed account for another portion of the errors incurred by a quantum processor. They not only lead to a decrease in gate fidelity, but also affect idle qubits. Here, the factor of 5 we chose is compatible with the improvements of the recent years (see [32], Fig 2c for the increasing T1 trend) Doing this will delay both spontaneous emission (amplitude damping) and phase flip (dephasing) events.

As shown in Fig. 9, better coherence only has a limited impact on the fragmenting procedure: it seems to improve more the success probability of the runs with fewer fragments than the one of the runs with more fragments where the solid and dashed lines are closer one to the other. This behavior is expected. Using a larger number of fragments imply that the fragments are smaller in terms of qubits size and such small fragments are less sensitive to decoherence.

All these observations are summarized in Fig. 10, which shows the increased success probability using the new parameters compared to the success probabilities $P_{\text{success}}^{(0)}$ computed with the Johannesburg noise parameters. For each of the scenarios \mathcal{S} introduced above, we compute the increase in probability defined as:

$$\Delta P(\mathcal{S}, n_f) = \langle P_{\text{success}}(\mathcal{S}, n_f, n_q) - P_{\text{success}}^{(0)}(n_f, n_q) \rangle_{n_q}. \quad (12)$$

We see that, as discussed above, better readout is all the more helpful as the number of fragments is large, while, conversely, better coherence is more beneficial for smaller number of fragments. Achieving better gate fidelities, on the other hand, is equally beneficial with and without fragmentation since the slope of the orange line is close to 1. (We stress that because of the arbitrariness in the quantitative choice of level of improvement for the three scenarios, one cannot conclude any quantitative insight from the value of

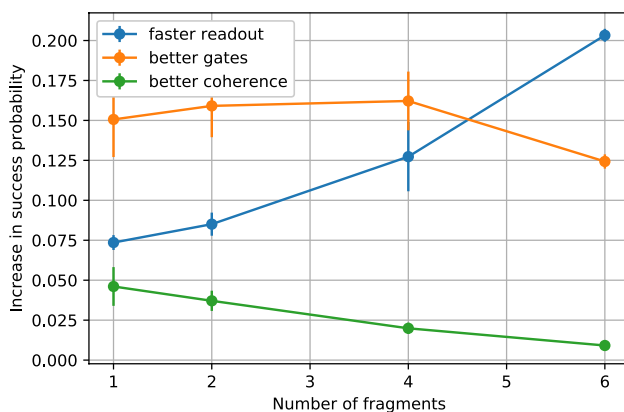


Fig. 10 Increase in success probability averaged over the number of qubits as a function of the number of fragments, $\Delta P(\mathcal{S}, n_f)$, for a faster readout (blue, parameters of Fig. 7), better gates (orange, parameters of Fig. 8), better coherence (green, parameters of Fig. 9)

the improvement; here, our conclusions are qualitative and only based on the slope with respect to the number of fragments). Consequently, to make the most of the fragmenting procedure in the case of numerous fragments, the major error source to focus on is the measurement error by designing faster readouts.

Contraction Complexity and Relation to Tensor-Network Simulation

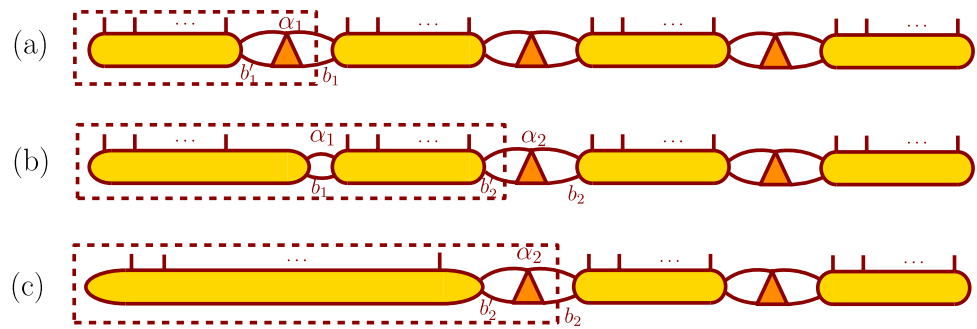
In this section, we elaborate on the complexity of the fragmentation algorithm. As presented in “Circuit Cutting”, the fragmentation method consists of a quantum and a classical step. In the quantum step, a batch of quantum circuits is executed on a Quantum Processing Unit (QPU). The number of such circuits scales as the number K of disconnected subgraphs of the original directed acyclic graph with some edges removed. The outcome of this step is a list of probability distributions p_i . In the classical step, a tensor network with nodes corresponding either to the probability distributions or to the γ tensors defined in “Circuit Cutting” needs to be contracted.

Here, we shall be interested in the contraction complexity of such a tensor network, assuming one wants to recover the probability of a single bitstring $(\hat{b}_0, \dots, \hat{b}_{m-1})$, i.e. for a fixed assignment of the external legs of the tensor network shown in Fig. 2b. A naive contraction of the tensor network at hand, namely a simultaneous summation over all internal indices $(\alpha_i, b_i, b'_i)_{i=1 \dots K-1}$, would entail a contraction complexity of 12^{K-1} , i.e. a classical computation that is exponential in the number K of fragments. In our case, however, the linear structure of the graph underlying the tensor network allows for a much more efficient sequential contraction strategy. Such a strategy, which is also widely exploited for contracting so-called Matrix Product States (see, e.g. [37, 38] for a review), consists in sequentially contracting the nodes of the network starting from one end of the linear graph. This is illustrated in Fig. 11, where we show the first three steps. The contraction complexity of the successive steps is 12, 36, 12, 36, ..., 12, 36, 12, 6. For K fragments, this yields an overall contraction complexity of $48(K-2) + 18 = 48K - 78$, i.e. a linear complexity in the fragment number K .

In the case of a general tensor network, the optimal contraction complexity can be shown to be at least of the order of $O(e^T)$, where T is the so-called treewidth of the network graph [39]. The treewidth of a graph can be defined as a combinatorial metric of closeness of the graph to a tree. There are a few ways to define the treewidth in more formal way: the minimum k for which a given graph is a partial k -tree, or the elimination width.

Tensor-network theory can also be leveraged to simulate quantum circuits classically. There are a number of tensor-network-based simulators developed for such simulations:

Fig. 11 First three contraction steps for the fragmentation of the GHZ-type circuit of Fig. 1



QFlex [40], AC-QDP [41], Quimb [42], and QTensor [43]. These simulators are typically much faster and more efficient than state vector simulators for shallow circuits [44] such as the circuits in this work. In these tensor simulators, the circuits are not directly represented by tensors, but rather use line graphs, which was proposed by Boixo et al. [45]. This approach has multiple benefits. The only disadvantage of the line graph approach is that it has limited usability to simulate sub-tensors of amplitudes, which was resolved in the work by Schutski et al. [46].

The method studied in our work, circuit cutting, has a counterpart in tensor-network-based simulation. It is called *tensor slicing*. One way to understand the slice of a tensor as an index that can be viewed as the function of many variables evaluated at some value of one variable:

$$f(x_1, x_2, \dots, x_n)|_{x_1=a} = \tilde{f}(x_2, \dots, x_n),$$

where variables can have integer values $x_i \in [0, d - 1]$. Thus, in this technique, slicing reduces the number of indices of the tensor one by one. Since all sizes of indices we use are equal to 2, removal of n vertices allows to split the expression into 2^n separate parts. This operation is also equivalent to decomposition of the full tensor expression. Each separate tensor is represented by a graph with lower connectivity than the original one. As a result, it dramatically reduces the complexity of finding the optimal elimination. Thus, it results in a lower contraction cost. It is a powerful technique that allows to simulate large circuits as does the circuit-cutting technique described in this work.

Homogeneous and Heterogeneous Quantum Computing

One exciting application of the circuit-cutting technique is to allow to execute much larger circuits. It can be done in two ways: split circuits and run sequentially on a quantum device (as we demonstrated in [1]), or run at the same time on multiple quantum devices. The latter way can lead to an exciting new era of how quantum computation is done—distributed quantum computing. It can potentially not only allow for the

execution of larger circuits, but also for a much faster execution. It is arguably a more realistic approach in the near future compared to the “true” distributed quantum computing that requires a quantum network connecting quantum devices. In our approach, indeed, we would utilize only the classical network.

Conclusions

In this work, we further investigated the Quantum Divide and Conquer approach, whose first implementation was demonstrated in a recent work of ours [1].

After giving more details as to the mathematical framework and physical models used for this implementation, we analyzed the influence of different noise sources on the success probability of a simple, GHZ-type circuit using classical noisy simulations on the Atos Quantum Learning Machine. We focused on the three main noise sources of today’s superconducting processors, namely readout errors, gate errors and decoherence (relaxation and dephasing) on idle qubits. We showed that readout errors are the most detrimental to the QDC procedure, because QDC requires additional measurements as the number of fragments increases. Conversely, the effect of idling noise is mitigated by QDC, as QDC results in smaller circuits that are less susceptible to this source of noise.

We also analyzed the computational complexity of QDC using tensor-network methods. While for a general circuit the contraction complexity increases exponentially with the number of cuts, for the GHZ-like circuit we studied, the complexity increases linearly with the number of cuts.

Finding more complex circuits in which the contraction complexity is still manageable is an interesting future direction. Circuits that have a “clustered” structure [14], that are e.g required in methods like the Dynamic Quantum Variational Ansatz [47], are promising candidates. In these methods, indeed, the ansatz has a mixer unitary that is made up of partial mixers that can have limited connectivity between each other, and can therefore form clusters.

Appendix A: Pauli-Basis Representation of Operators and Superoperators

We can decompose any Hermitian operator (including density matrices) as

$$\rho = \sum_{\alpha} \rho_{\alpha} P_{\alpha}, \quad \rho_{\alpha} = \frac{1}{d} \text{Tr}[P_{\alpha} \rho] \quad (13)$$

with $d = 2^{n_{\text{qbits}}}$ and P_{α} a generalized Pauli matrix on n_{qbits} qubits. Similarly, superoperators can be decomposed on this basis,

$$[\mathcal{R}]_{\alpha\beta} = \frac{1}{d} \text{Tr}[P_{\alpha} \cdot \mathcal{O}(P_{\beta})].$$

\mathcal{R} is called the Pauli transfer matrix (PTM) representation of \mathcal{O} . Then the coordinates of $\rho' = \mathcal{O}(\rho)$ in the Pauli basis are simply given by

$$\begin{aligned} \rho'_{\alpha} &= \frac{1}{d} \text{Tr}[P_{\alpha} \mathcal{O}(\rho)] \\ &= \frac{1}{d} \sum_{\beta} \rho_{\beta} \text{Tr}[P_{\alpha} \mathcal{O}(P_{\beta})] \\ &= \sum_{\beta} \mathcal{R}_{\alpha\beta} \rho_{\beta}. \end{aligned} \quad (14)$$

We note that

$$\begin{aligned} \text{Tr}[A^{\dagger} \cdot B] &= \sum_{\alpha\beta} A_{\alpha}^{*} B_{\beta} \text{Tr}[P_{\alpha}^{\dagger} P_{\beta}] \\ &= \sum_{\alpha\beta} A_{\alpha}^{*} B_{\beta} d \delta_{\alpha\beta} = d \sum_{\alpha} A_{\alpha}^{*} B_{\alpha}. \end{aligned} \quad (15)$$

Defining the scalar product $\langle\langle A|B \rangle\rangle \equiv \sum_{\alpha} A_{\alpha}^{*} B_{\alpha}$, we thus have

$$\text{Tr}[A^{\dagger} \cdot B] = d \langle\langle A|B \rangle\rangle = 2^{n_{\text{qbits}}} \langle\langle A|B \rangle\rangle. \quad (16)$$

Acknowledgements This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725. This research also used the resources of the Argonne Leadership Computing Facility, which is DOE Office of Science User Facility supported under Contract DE-AC02-06CH11357. Yuri Alexeev, Zain H. Saleem and Martin Suchara were supported by the DOE, Office of Science, under Contract DE-AC02-06CH11357. The compilation and noisy simulations were performed using Argonne National Laboratory's and Atos Quantum Laboratory's Quantum Learning Machines. The submitted manuscript has been created by UChicago Argonne, LLC, Operator of Argonne National Laboratory ("Argonne"). Argonne, a U.S. Department of Energy Office of Science laboratory, is operated under Contract No. DE-AC02-06CH11357. The U.S. Government retains for itself, and others acting on its behalf, a paid-up nonexclusive, irrevocable worldwide license in said article to reproduce, prepare derivative works, distribute copies to the public, and perform publicly and display publicly, by or on behalf of the Government. The Department of Energy will provide public access to these results of federally sponsored research

in accordance with the DOE Public Access Plan (<http://energy.gov/downloads/doe-public-access-plan>).

Data Availability Statement The data and materials presented in this paper are available upon request to the authors.

Compliance with Ethical Standards

Conflict of interest The authors declare that they have no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Ayril T, Le Regent FM, Saleem Z, Alexeev Y, Suchara M. Quantum divide and compute: Hardware demonstrations and noisy simulations. In: Proceedings of IEEE computer society annual symposium on VLSI, ISVLSI; 2020, pp. 138–140. <https://doi.org/10.1109/ISVLSI49217.2020.00034>
2. Preskill J. Quantum computing in the NISQ era and beyond. *Quantum*. 2018;2:79. <https://doi.org/10.22331/q-2018-08-06-79>.
3. Kelly J. A preview of Bristlecone, Google's new quantum processor, Google AI Blog, <https://ai.googleblog.com/2018/03/a-preview-of-bristlecone-googles-new.html>, Mar 2018.
4. Knight W. IBM raises the bar with a 50-qubit quantum computer, MIT Technology Review, <https://www.technologyreview.com/s/609451/ibm-raises-the-bar-with-a-50-qubit-quantum-computer/>, Nov 2017.
5. Hsu J. CES 2018: Intel's 49-qubit chip shoots for quantum supremacy, IEEE Spectrum, <https://spectrum.ieee.org/tech-talk/computing/hardware/intels-49qubit-chip-aims-for-quantum-supremacy>, Jan 2018.
6. Gambetta JM, Chow JM, Steffen M. Building logical qubits in a superconducting quantum computing system. *NPJ Quantum Inf*. 2017;3(1):2. <https://doi.org/10.1038/s41534-016-0004-0>.
7. Monroe C, Kim J. Scaling the ion trap quantum processor. *Science*. 2013;339(6124):1164–1169. <https://science.sciencemag.org/content/339/6124/1164>
8. Saffman M. Quantum computing with neutral atoms. *Nat Sci Rev*. 2018;6(1):24–5. <https://doi.org/10.1093/nsr/nwy088>.
9. Dolde F, Jakobi I, Naydenov B, Zhao N, Pezzagna S, Trautmann C, Meijer J, Neumann P, Jelezko F, Wrachtrup J. Room-temperature entanglement between single defect spins in diamond. *Nat Phys*. 2013;9:139. <https://doi.org/10.1038/nphys2545>.
10. Bernien H, Hensen B, Pfaff W, Koolstra G, Blok MS, Robledo L, Taminau TH, Markham M, Twitchen DJ, Childress L, Hanson R. Heralded entanglement between solid-state qubits separated by three metres. *Nature*. 2013;497:86. <https://doi.org/10.1038/nature12016>.

11. Bravyi S, Smith G, Smolin JA. Trading classical and quantum computational resources. *Phys Rev X*. 2016;6:021043. <https://doi.org/10.1103/PhysRevX.6.021043>.
12. Peng T, Harrow A, Ozols M, Wu X. Simulating large quantum circuits on a small quantum computer. 2019. arXiv preprint [arXiv:1904.00102](https://arxiv.org/abs/1904.00102).
13. Cross AW, Bishop LS, Sheldon S, Nation PD, Gambetta JM. Validating quantum computers using randomized model circuits. *Phys Rev A*. 2019;100(3):032328. <https://doi.org/10.1103/PhysRevA.100.032328>.
14. Perlin MA, Saleem ZH, Suchara M, Osborn JC. Quantum circuits: divide and compute with maximum likelihood tomography. 2020. arXiv preprint [arXiv:2005.12702](https://arxiv.org/abs/2005.12702).
15. Tang W, Tomesh T, Larson J, Suchara M, Martonosi M. CutQC: using small quantum computers for large quantum circuit evaluations. In: Proceedings of the ACM international conference on architectural support for programming languages and operating systems (ASPLOS); 2021.
16. Preskill J. Quantum computing and the entanglement frontier. [arXiv:1203.5813](https://arxiv.org/abs/1203.5813), Nov 2012 [Online].
17. Alexeev Y. Evaluation of the intel-QS performance on theta supercomputer. In: Argonne national laboratory—leadership computing facility, Technical report ANL/ALCF 18/2, Apr 2018.
18. Häner T, Steiger DS. 0.5 petabyte simulation of a 45-qubit quantum circuit. In: Proceedings of the international conference for high performance computing, networking, storage and analysis, ser. SC '17. New York: ACM; 2017. pp. 33:1–33:10. <https://doi.org/10.1145/3126908.3126947>
19. Boixo S, Isakov SV, Smelyanskiy VN, Babbush R, Ding N, Jiang Z, Bremner MJ, Martinis JM, Neven H. Characterizing quantum supremacy in near-term devices. *Nat Phys*. 2018;14(6):595–600. <https://doi.org/10.1038/s41567-018-0124-x>.
20. Aleksandrowicz G et al. Qiskit: An open-source framework for quantum computing. 2019.
21. Smelyanskiy M, Sawaya NPD, Aspuru-Guzik A. qHiPSTER: the quantum high performance software testing environment. 2016. [arXiv:1601.07195](https://arxiv.org/abs/1601.07195) [Online].
22. Atos quantum learning machine. <https://atos.net/wp-content/uploads/2018/07/Atos-Quantum-Learning-Machine-brochure.pdf>. Jun 2018.
23. Steiger DS, Häner T, Troyer M. ProjectQ: an open source software framework for quantum computing. *Quantum*. 2018;2:49.
24. McClean JR, Romero J, Babbush R, Aspuru-Guzik A. The theory of variational hybrid quantum-classical algorithms. *New J Phys*. 2016;18(2):023023. <https://doi.org/10.1088/1367-2630/18/2/023023>.
25. Barrett S, Hammerer K, Harrison S, Northup TE, Osborne TJ. Simulating quantum fields with cavity QED. *Phys Rev Lett*. 2013;110:090501. <https://doi.org/10.1103/PhysRevLett.110.090501>.
26. Farhi E, Goldstone J, Gutmann S. A quantum approximate optimization algorithm. [arXiv:1411.4028](https://arxiv.org/abs/1411.4028). Nov 2014.
27. Wecker D, Hastings MB, Troyer M. Progress towards practical quantum variational algorithms. *Phys Rev A*. 2015;92:042303. <https://doi.org/10.1103/PhysRevA.92.042303>.
28. Guerreschi GG, Matsuura AY. QAOA for max-cut requires hundreds of qubits for quantum speed-up. [arXiv:1812.07589](https://arxiv.org/abs/1812.07589) Dec 2018.
29. Chen Y, Farahzad M, Yoo S, Wei T-C. Detector tomography on IBM quantum computers and mitigation of an imperfect measurement. *Phys Rev A*. 2019;100(5):052315. <https://doi.org/10.1103/PhysRevA.100.052315>.
30. Sarovar M, Proctor T, Rudinger K, Young K, Nielsen E, Blume-Kohout R. Detecting crosstalk errors in quantum information processors. *Quantum* 2020;4:321. <https://quantum-journal.org/paper/q-2020-09-11-321/>
31. Paladino E, Galperin Y, Falci G, Altshuler BL. 1/f noise: implications for solid-state quantum information. *Rev Mod Phys*. 2014;86(2):361–418. <https://doi.org/10.1103/RevModPhys.86.361>.
32. Kjaergaard M, Schwartz ME, Braumüller J, Krantz P, Wang JI-J, Gustavsson S, Oliver WD. Superconducting qubits: current state of play. *Annu Rev Condens Matter Phys* 2020;11(1):031–119–050 605. <https://doi.org/10.1146/annurev-conmatphys-031119-050605>.
33. Heinsoo J, Andersen CK, Remm A, Krinner S, Walter T, Sala-thé Y, Gasparinetti S, Besse J-C, Potočnik A, Wallraff A, Eichler C. Rapid high-fidelity multiplexed readout of superconducting qubits. *Phys Rev Appl*. 2018;10(3):034040. <https://doi.org/10.1103/PhysRevApplied.10.034040>.
34. Ibm quantum experience website. <https://quantum-computing.ibm.com/>. Accessed 5 Mar 2020.
35. Arute F, Arya K, Babbush R, Bacon D, Bardin JC, Barends R, Biswas R, Boixo S, Brandao FGSL, Buell DA, Burkett B, Chen Y, Chen Z, Chiaro B, Collins R, Courtney W, Dunsworth A, Farhi E, Foxen B, Fowler A, Gidney C, Giustina M, Graff R, Guerin K, Guerin S, Habegger S, Harrigan MP, Hartmann MJ, Ho A, Hoffmann M, Huang T, Humble TS, Isakov SV, Jeffrey E, Jiang Z, Kafri D, Kechedzhi K, Kelly J, Klimov PV, Knysh S, Korotkov A, Kostitsa F, Landhuis D, Lindmark M, Lucero E, Lyakh D, Mandrà S, McClean JR, McEwen M, Megrant A, Mi X, Michielens K, Mohseni M, Mutus J, Naaman O, Neeley M, Neill C, Niu MY, Ostby E, Petukhov A, Platt JC, Quintana C, Rieffel EG, Roushan P, Rubin NC, Sank D, Satzinger KJ, Smelyanskiy V, Sung KJ, Trevithick MD, Vainsencher A, Villalonga B, White T, Yao ZJ, Yeh P, Zalcman A, Neven H, Martinis JM. Quantum supremacy using a programmable superconducting processor. *Nature*. 2019;574(7779):505–10. <https://doi.org/10.1038/s41586-019-1666-5>.
36. Rigetti computing website. <https://www.rigetti.com/what>. Accessed 23 Nov 2020.
37. Schollwöck U. The density-matrix renormalization group in the age of matrix product states. *Ann Phys*. 2011;326(1):96–192. <https://doi.org/10.1016/j.aop.2010.09.012>.
38. Orus R. A practical introduction to tensor networks: matrix product states and projected entangled pair states. *Ann Phys*. 2013;349:117–58. <https://doi.org/10.1016/j.aop.2014.06.013>.
39. Markov IL, Shi Y. Simulating quantum computation by contracting tensor networks. *SIAM J Comput*. 2008;38(3):963–81. <https://doi.org/10.1137/050644756>.
40. Villalonga B, Boixo S, Nelson B, Henze C, Rieffel E, Biswas R, Mandrà S. A flexible high-performance simulator for verifying and benchmarking quantum circuits implemented on real hardware. *NPJ Quantum Inf*. 2019;5:1–16. <https://doi.org/10.1038/s41534-019-0196-1>.
41. Huang C, Szegedy M, Zhang F, Gao X, Chen J, Shi Y. Alibaba cloud quantum development platform: applications to quantum algorithm design. arXiv preprint [arXiv:1909.02559](https://arxiv.org/abs/1909.02559) 2019.
42. Gray J. quimb: a python package for quantum information and many-body calculations. *J Open Source Softw*. 2018;3(29):819.
43. Lykov D, Ibrahim C, Galda A, Alexeev Y. Tensor network simulator QTensor. 2020. <https://github.com/danlkv/QTensor>.
44. Wu X-C, Di S, Dasgupta EM, Cappello F, Finkel H, Alexeev Y, Chong FT. Full-state quantum circuit simulation by using data compression. In: Proceedings of the high performance computing, networking, storage and analysis international conference (SC19). Denver IEEE Computer Society; 2019. <https://doi.org/10.1145/3295500.3356155>.
45. Boixo S, Isakov SV, Smelyanskiy VN, Neven H. Simulation of low-depth quantum circuits as complex undirected graphical models. 2017. arXiv preprint [arXiv:1712.05384](https://arxiv.org/abs/1712.05384).

46. Schutski R, Lykov D, Oseledets I. An adaptive algorithm for quantum circuit simulation. 2019. arXiv preprint [arXiv:1911.12242](https://arxiv.org/abs/1911.12242).
47. Saleem ZH, Tariq B, Suchara M. Approaches to constrained quantum approximate optimization. 2020. arXiv preprint [arXiv:2010.06660](https://arxiv.org/abs/2010.06660).

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.