

Quantum Circuits: Divide and Compute with Maximum Likelihood Tomography

Michael A. Perlin,^{1,2} Zain H. Saleem,³ Martin Suchara,³ and James C. Osborn^{2,4}

¹*JILA, National Institute of Standards and Technology and University of Colorado, 440 UCB, Boulder, Colorado 80309, USA*

²*Argonne Leadership Computing Facility, Argonne National Laboratory, 9700 S Cass Ave, Argonne, IL 60439*

³*Mathematics and Computer Science Division, Argonne National Laboratory, 9700 S Cass Ave, Argonne, IL 60439*

⁴*Computational Science Division, Argonne National Laboratory, 9700 S Cass Ave, Argonne, IL 60439*

(Dated: 2020-05-26)

We introduce the method of maximum likelihood fragment tomography (MLFT) as an improved circuit cutting technique for running clustered quantum circuits on quantum devices with limited quantum resources. In addition to minimizing the classical computing overhead of circuit cutting methods, MLFT finds the most likely probability distribution over measurement outcomes at the output of a quantum circuit, given the data obtained from running the circuit’s fragments. Unlike previous circuit cutting methods, MLFT guarantees that all reconstructed probability distributions are strictly non-negative and normalized. We demonstrate the benefits of MLFT with classical simulations of clustered random unitary circuits. Finally, we provide numerical evidence and theoretical arguments that circuit cutting can estimate the output of a clustered circuit with higher fidelity than full circuit execution, thereby motivating the use of circuit cutting as a standard tool for running clustered circuits on quantum hardware.

I. INTRODUCTION

The advent of Noisy Intermediate Scale Quantum (NISQ) technologies[1] makes multiqubit processors with increasing numbers of qubits available to the quantum computing community for experimentation. The rapid progress in the development and manufacturing of these devices is remarkable. Advances on the hardware front have been matched by the theoretical development of suitable hardware benchmarks[2], which have enabled proof-of-principle demonstrations of a computational advantage over classical computing systems[3]. However, existing devices are still lacking in the number and quality of qubits that are required for practical applications such as digital quantum simulation[4, 5], quantum optimization[6–8] and quantum machine learning[9, 10]. The gap between the requirements of existing quantum algorithms and the resources available on NISQ devices has motivated problem decompositions that trade-off quantum and classical computing resources[11, 12].

One such decomposition, inspired by fragmentation methods for quantum molecular cluster simulations[13–15], applies the idea of fragmentation to the execution of quantum circuits[12]. This decomposition consists of (i) “cutting” a quantum circuit into smaller sub-circuits, or *fragments*, that can be executed on smaller hardware, and then (ii) reconstructing the probability distribution over measurement outcomes in the original quantum circuit from probability distributions associated with its fragments. The severed quantum connections between circuit fragments are essentially simulated by classical post-processing of sub-circuit data, leading to a classical computing overhead that grows exponentially with the number of cuts that are made to a circuit. This approach is therefore suitable for simulating circuits that are decomposable into clusters of gates with a small number of inter-cluster interactions, which make appearances in near-term quantum computing applications such as

QAOA[6, 7] and VQE[12, 16].

Due to the presence of shot noise, an unavoidable feature of the original fragment recombination method in Ref. [12] is that the distribution over measurement outcomes obtained by characterizing and recombining circuit fragments does not generally satisfy central axioms of probability theory, namely that a probability distribution must be non-negative and normalized. A naive fix to this problem would be to simply remove all negative probabilities and normalize the reconstructed distribution in question. In the spirit of maximum likelihood state tomography (MLST)[17], however, one would like to determine the “most likely” probability distribution that is consistent with available fragment data.

In this work, we find this “most likely” probability distribution by generalizing MLST to maximum likelihood fragment tomography (MLFT), the use of which guarantees that reconstructed probability distributions are automatically non-negative and normalized. We discuss how MLFT minimizes the classical computing resources necessary to characterize circuit fragments, and introduce a tensor-network-based fragment recombination method that allows for the use of efficient numerical routines. We test our methods in numerical experiments on clustered random unitary circuits, and demonstrate that MLFT estimates probability distribution at the output of a fragmented quantum circuit with higher fidelity than the naive method of removing negative probabilities and normalizing. As an added bonus, we find that for a fixed budget of quantum computing resources, circuit cutting methods can outperform direct execution and sampling of a clustered circuit to estimate its associated probability distribution. We provide theoretical arguments to support this finding, which motivates the use of circuit cutting a standard tool for the evaluation of clustered circuits on quantum hardware, even when all hardware requirements for full circuit execution are satisfied.

In Section II we provide an overview of the circuit cutting method introduced in Ref. [12], before discussing the shortcomings of this method in Section III. We then introduce the procedure of maximum likelihood fragment tomography in Section IV, in addition to the tensor-network-based fragment recombination algorithm that it motivates. To test out our ideas and demonstrate their utility in an application-agnostic setting, we perform numerical experiments on clustered random unitary circuits in Section V, before discussing our main conclusions and future outlooks in Section VI.

II. THE GENERAL CUT-AND-STITCH PRESCRIPTION

Here we provide a basic overview of the circuit cutting procedure introduced in Ref. [12], and establish some formalism and language that we will use throughout the remainder of this work. Given an arbitrary quantum state $|\psi\rangle$ of N qubits, a straightforward resolution of the identity operator $I = \sum_{b \in \{0,1\}} |b\rangle\langle b|$ on qubit n implies that

$$|\psi\rangle = I_n |\psi\rangle \simeq \sum_{b \in \{0,1\}} |b\rangle \otimes_n \langle b | \psi \rangle, \quad (1)$$

where I_n denotes the action of I on qubit n ; the relation \simeq denotes equality up to a permutation of tensor factors (i.e. qubit order); and ${}_n \langle b | \psi \rangle$ is a sub-normalized state of $N - 1$ qubits acquired by projecting $|\psi\rangle$ onto state $|b\rangle$ of qubit n . If the structure of a quantum circuit that prepares $|\psi\rangle$ allows, a similar resolution of the identity operator I can be used to “cut” the circuit by inserting I at a location that splits the circuit into two disjoint sub-circuits. For example, if $|\psi\rangle = V_{23}U_{12}|000\rangle$ with U_{12} and V_{23} respectively the two-qubit gates U and V acting on qubits 1, 2 and 2, 3, then by inserting the identity operator I_2 between U_{12} and V_{23} we find that

$$|\psi\rangle \simeq \sum_{b \in \{0,1\}} |\psi_1(b)\rangle \otimes |\psi_2(b)\rangle, \quad (2)$$

where the factors

$$|\psi_1(b)\rangle \equiv {}_2 \langle b | U | 00 \rangle, \quad |\psi_2(b)\rangle \equiv V | b0 \rangle \quad (3)$$

are (generally sub-normalized) “conditional” states prepared by projecting onto $|b\rangle$ or preparing $|b\rangle$ as appropriate. The identity in Eq. (2) is visualized in Figure 1, albeit with the use of density operators that we discuss below.

The above splitting method relies on the capability to project qubit n onto state $|b\rangle$ while preserving phase information. Such capability is possible when running classical simulations of a circuit, but is not possible on quantum computing hardware. This limitation can be overcome by representing quantum states $|\psi\rangle$ with density operators $\rho = |\psi\rangle\langle\psi|$, whose diagonal entries in a given

measurement basis define a classical probability distribution over measurement outcomes in that basis. For ease of language, we will at times blur the distinction between a state ρ and the probability distribution defined by its diagonal entries in a fixed computational basis. In the remainder of this work, we will discuss circuit splitting and reconstruction in way that is compatible with circuit execution on quantum computing hardware. Nonetheless, our methods can be applied to the case of classical state simulation, with minor simplifying modifications to account for the added capability of performing deterministic, phase-preserving qubit projections.

The identity analogous to Eq. (1) for density operators ρ reads

$$\rho \simeq \frac{1}{2} \sum_{M \in \mathcal{B}} M \otimes \text{tr}_n (M_n \rho) \quad (4)$$

where \mathcal{B} is a basis of self-adjoint 2×2 matrices with normalization $\text{tr} [M^{(i)} M^{(j)}] = 2\delta_{ij}$ for $M^{(i)}, M^{(j)} \in \mathcal{B}$; tr_n denotes a partial trace with respect to qubit n ; and M_n denotes an operator which acts with M on qubit n and trivially (i.e. with the identity I) on all other qubits. For concreteness we will use the set of Pauli operators together with the single-qubit identity operator I , $\mathcal{B} \equiv \{X, Y, Z, I\}$, as our basis. The identity in Eq. (4) implies that the state prepared by the action of a three-qubit circuit $V_{23}U_{12}$ on the trivial state $|0\rangle\langle 0|^{\otimes 3}$ can be decomposed as

$$\rho \simeq \frac{1}{2} \sum_{M \in \mathcal{B}} \rho_1(M) \otimes \rho_2(M), \quad (5)$$

where now the factors

$$\rho_1(M) \equiv \text{tr}_2 (M_2 U |0\rangle\langle 0|^{\otimes 2} U^\dagger), \quad (6)$$

$$\rho_2(M) \equiv V (M \otimes |0\rangle\langle 0|) V^\dagger, \quad (7)$$

have no straightforward interpretation as “conditional” states, similarly to $|\psi_1(b)\rangle$ and $|\psi_2(b)\rangle$ in Eq. (2). In order to decompose ρ into conditional states, the idea in Ref. [12] was essentially to expand each $M \in \mathcal{B}$ into its eigenbasis:

$$\rho \simeq \frac{1}{2} \sum_{\substack{M \in \mathcal{B} \\ r, s \in \lambda(M)}} r s \rho_1(M_r) \otimes \rho_2(M_s), \quad (8)$$

where $\lambda(M)$ denotes the spectrum of M , i.e. $\lambda(X) = \lambda(Y) = \lambda(Z) = (+1, -1)$ and $\lambda(I) = (1, 1)$; and $M_s \equiv |M_s\rangle\langle M_s|$ is a projector onto an eigenstate of $|M_s\rangle$ of M with eigenvalue s . Note that the choice of eigenstates for the identity operator I is arbitrary as long as these two states are orthogonal, so we can re-use the eigenstates from one of the other operators.

The decomposition in Eq. (8) allows interpreting each $\rho_f(M_s)$ as a conditional state, obtained either by post-selecting onto the measurement of a qubit in state $|M_s\rangle$, or by preparing a qubit in state $|M_s\rangle$, as appropriate

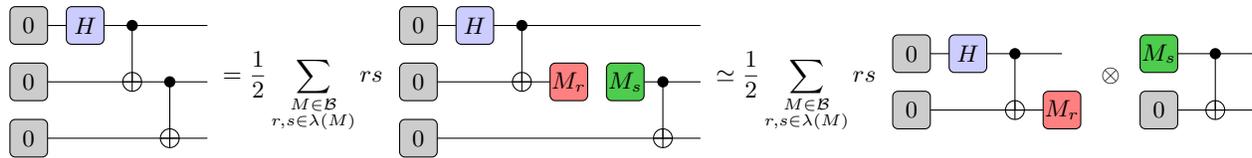


FIG. 1. A 3-qubit GHZ circuit cut into two 2-qubit fragments by inserting an identity operator. Here $\mathcal{B} \equiv \{X, Y, Z, I\}$ is the set of Pauli operators X, Y, Z and the identity I , which together form an orthogonal basis for the space of single-qubit operators; $\lambda(M)$ denotes the spectrum of M ; and $M_s \equiv |M_s\rangle\langle M_s|$ is the projector onto an eigenstate $|M_s\rangle$ of M with eigenvalue s . Green (red) boxes correspond to preparations (projections) of a qubit in the state specified within the box. After cutting a circuit, the resulting fragments can be simulated independently, and an appropriate post-processing of simulation results recovers the output of the original (pre-cut) circuit.

(see Figure 1). The decomposition in Eq. (8) thus corresponds to the following procedure for circuit cutting and reconstruction: after cutting a circuit into (say) two fragments, characterize the classical probability distributions $\rho_f(M_s)$ over measurement outcomes by running the corresponding sub-circuit and either post-selecting on measurement outcomes M_s or preparing states M_s as appropriate. Note that post-selected probability distributions are generally sub-normalized, with the normalization $\text{tr} \rho_1(M_s)$ equal to the probability of getting outcome M_s when measuring in the diagonal basis of M . After characterizing the conditional distributions $\rho_f(M_s)$ for each of $f \in \{1, 2\}$, $M \in \{X, Y, Z\}$, and $s \in \{+1, -1\}$, combine these distributions according to Eq. (8). This scenario is illustrated in Figure 1, which cuts a 3-qubit GHZ circuit preparing the state $|\psi\rangle \propto |000\rangle + |111\rangle$ into two 2-qubit fragments.

III. SHORTCOMINGS OF PREVIOUS WORK

In practice, recombining circuit fragments as prescribed by Eq. (8) is uneconomical in two ways. First, the tensor products in Eq. (8) are a computational bottleneck for fragment recombination. It is therefore faster to post-process conditional distributions by first (i) for each fragment f , combining the six independent distributions $\rho_f(M_s)$ into four distributions: $\rho_f(M) = \rho_f(M_{+1}) - \rho_f(M_{-1})$ for each $M \in \{X, Y, Z\}$ and $\rho_f(I) = \rho_f(M_{+1}) + \rho_f(M_{-1})$ for any $M \in \{X, Y, Z\}$, and then (ii) combining the fragment distributions $\rho_f(M)$ according to Eq. (5). In a circuit with K cuts, this post-processing reduces the number of tensor products that must be computed during recombination from 16^K to 4^K , which is an exponential improvement in K^a .

Second, the recombination formula in Eq. (8) nominally requires, for each fragment f incident on K_f cuts,

characterizing K_f^6 probability distributions. This characterization is overcomplete, as the K_f^6 distributions are not all linearly independent. In the case of two fragments with one cut each, for example, we can decompose

$$\rho_f(X_-) = \rho_f(Z_+) + \rho_f(Z_-) - \rho_f(X_+). \quad (9)$$

In fact, a fragment with K_f incident cuts can be completely characterized by K_f^4 distributions, which can be deduced from the fact that the space of operators on the Hilbert space of a qubit has real dimension four. The symmetric, informationally complete, positive operator-valued measure (SIC-POVM) $\{\Pi_j^{\text{SIC}} : j \in \mathbb{Z}_4\}$ consisting of projectors Π_j^{SIC} onto states represented by the four corners of a regular tetrahedron inscribed in the Bloch sphere, for example, form a mutually unbiased basis for the space of single-qubit operators. Given any single-qubit operator M , we can therefore expand $\rho_f(M) = \sum_{j \in \mathbb{Z}_4} c_{Mj}^{(f)} \rho_f(\Pi_j^{\text{SIC}})$ with real coefficients $c_{Mj}^{(f)}$.

Finally, characterizing fragments is a noisy process, due to both (i) hardware errors that are unavoidable without error correction, as in all NISQ devices, and (ii) statistical sampling (shot) noise. As a result, the “experimentally inferred” distributions $\tilde{\rho}_f(M_s)$ approximating the “true” distributions $\rho_f(M_s)$ will generally contain errors, and will fail to satisfy self-consistency conditions such as Eq. (9). As a consequence, when combining these distributions according to Eqs. (5) and (8) there is similarly no guarantee that the reconstructed probability distribution will satisfy conditions required of a probability distribution, such as non-negativity and normalization.

To fix these shortcomings, in the following section we recast the task of characterizing conditional distributions into that of performing fragment tomography, treating the fragments ρ_f , rather than distributions $\rho_f(M_s)$, as first-class objects. In addition to being “automatically” economical in terms of the classical memory footprint for the characterization of each fragment, performing fragment tomography allows us to adapt the method of maximum likelihood tomography[17] to construct a model for each fragment that is, by construction, guaranteed to satisfy all appropriate self-consistency conditions. Fragment recombination is then similarly guaranteed to yield

^a The recombination procedure in Ref. [12] involves 8^K terms, rather than 16^K , because it post-processes of “measurement” conditions, but not “preparation” conditions, essentially collapsing the sum over r in Eq. (8) while leaving the sum over s .

a probability distribution that is both non-negative and normalized. Finally, we show how the fragment models constructed via fragment tomography naturally admit a tensor-network-based method for recombination, allowing for the use of efficient numerical routines that outperform the recombination procedures discussed so far.

IV. MAXIMUM LIKELIHOOD FRAGMENT TOMOGRAPHY

Once a circuit has been cut into fragments ρ_f , rather than characterizing conditional distributions $\rho_f(M_s)$ we can perform a more systematic *maximum likelihood fragment tomography* (MLFT) procedure to characterize these fragments. The purpose of MLFT is to perform a “maximum likelihood” characterization, similar to the characterization of quantum states in Ref. [17], which guarantees that any probability distribution associated with these fragments will be (i) the “most likely” distribution consistent with available fragment data, while (ii) satisfying all necessary constraints for a valid (i.e. non-negative and normalized) probability distribution. MLFT essentially generalizes maximum likelihood state tomography (MLST)[17] to the case of channels (processes) with mixed (quantum/classical) inputs and outputs.

Any given fragment, nominally a unitary circuit on Q qubits, will generally have Q_i “quantum input” and Q_o “quantum output” qubits at the locations of cuts. We refer to these inputs and outputs as “quantum” because characterizing the fragment for circuit reconstruction will require performing full quantum tomography on the corresponding degrees of freedom. In contrast, the remaining $C_i \equiv Q - Q_i$ “classical input” qubits are always initialized in the trivial state $|0\rangle_i \equiv |0\rangle^{\otimes C_i}$, and the remaining $C_o \equiv Q - Q_o$ “classical output” qubits are always measured in a fixed computational basis. For definiteness, we can first think of a fragment as a quantum channel \mathcal{E}_Λ on the state of Q qubits. The channel-state duality[18–20] implies that this channel is uniquely defined by a 4-partite state (density operator) Λ defined by

$$\Lambda \equiv \sum_{\substack{k,\ell,m,n \\ p,q,r,s}} \Lambda_{k\ell;mn;pq;rs} |k\rangle\langle\ell| \otimes |m\rangle\langle n| \otimes |p\rangle\langle q| \otimes |r\rangle\langle s| \quad (10)$$

where the bitstrings k, ℓ (m, n ; p, q ; r, s) index states in the Hilbert space of the quantum input (classical input; quantum output; classical output) qubits of the fragment, and are implicitly summed over $\mathbb{Z}_2^{Q_i}$ ($\mathbb{Z}_2^{C_i}$; $\mathbb{Z}_2^{Q_o}$; $\mathbb{Z}_2^{C_o}$). Specifically, the channel \mathcal{E}_Λ maps a bipartite input state

$$\rho \otimes |\mathbf{0}_i\rangle\langle\mathbf{0}_i| \equiv \sum_{k,\ell} \rho_{k\ell} |k\rangle\langle\ell| \otimes |\mathbf{0}_i\rangle\langle\mathbf{0}_i| \quad (11)$$

at its input to the bipartite state

$$\mathcal{E}_\Lambda(\rho \otimes |\mathbf{0}_i\rangle\langle\mathbf{0}_i|) \equiv \sum_{k,\ell,p,q,r,s} \Lambda_{k\ell;0,0;pq;rs} \rho_{k\ell} |p\rangle\langle q| \otimes |r\rangle\langle s| \quad (12)$$

at its output. To account for the fact that classical outputs are only ever measured in a fixed computational basis, we can remove all parts of $\mathcal{E}_\Lambda(\rho)$ that are off-diagonal with respect to the measurement basis of the corresponding qubits. In total, we therefore need only characterize the channel $\mathcal{E}_{\tilde{\Lambda}}$ defined by

$$\mathcal{E}_{\tilde{\Lambda}}(\rho) \equiv \sum_{k,\ell,p,q,s} \tilde{\Lambda}_{k\ell;pq;s} \rho_{k\ell} |p\rangle\langle q| \otimes |s\rangle\langle s|, \quad (13)$$

where

$$\tilde{\Lambda}_{k\ell;pq;s} \equiv \Lambda_{k\ell;0,0;pq;ss}. \quad (14)$$

The task of performing MLFT thus reduces, in essence, to performing tomography on the tri-partite block-diagonal state

$$\begin{aligned} \tilde{\Lambda} &\equiv \sum_{k,\ell,p,q,s} \tilde{\Lambda}_{k\ell;pq;s} |k\rangle\langle\ell| \otimes |p\rangle\langle q| \otimes |s\rangle\langle s| \quad (15) \\ &= \sum_s \tilde{\Lambda}_s \otimes |s\rangle\langle s|, \quad (16) \end{aligned}$$

where Eq. (16) implicitly defines the blocks $\tilde{\Lambda}_s$. In words, the reduced state $\tilde{\Lambda}$ is acquired from the full state Λ by conditioning on (i.e. fixing) a trivial state $|\mathbf{0}_i\rangle\langle\mathbf{0}_i|$ on its classical inputs, and the block $\tilde{\Lambda}_s$ is acquired from $\tilde{\Lambda}$ by conditioning on measurement of the bitstring s on its classical outputs. The relationship between Λ , $\tilde{\Lambda}$, and $\tilde{\Lambda}_s$ is sketched out in Figure 2, which also summarizes the MLFT procedure discussed below.

Both MLST and MLFT begin by collecting experimental data to characterize the quantum state under consideration. In the case of the block-diagonal state $\tilde{\Lambda}$, one needs to essentially characterize expectation values

$$\langle \sigma_i \otimes \sigma_o \otimes z_c \rangle_{\tilde{\Lambda}} \equiv \text{tr} \left[\tilde{\Lambda} (\sigma_i \otimes \sigma_o \otimes z_c) \right] \quad (17)$$

for some complete basis of operators $\{\sigma_i \otimes \sigma_o \otimes z_c\}$ on the target Hilbert space of $\tilde{\Lambda}$, where σ_i , σ_o , and z_c are respectively operators on the quantum input, quantum output, and classical output of the fragment in question, with z_c strictly diagonal in the computational basis. MLST[17] collects data by performing informationally complete measurements of $\tilde{\Lambda}$, e.g. choosing operators $\sigma_{i,o}$ from the set of all Pauli strings $\{\mathbb{1}, X, Y, Z\}^{\otimes Q_{i,o}}$, and z_c from the set of diagonal Pauli strings $\{\mathbb{1}, Z\}^{\otimes C_o}$. In the case of fragment tomography, however, we do not have direct access to the state $\tilde{\Lambda}$, and instead have access to the channel $\mathcal{E}_{\tilde{\Lambda}}$. It is therefore not possible to directly measure the degrees of freedom in $\tilde{\Lambda}$ that are associated with inputs to the channel. Instead, MLFT characterizes

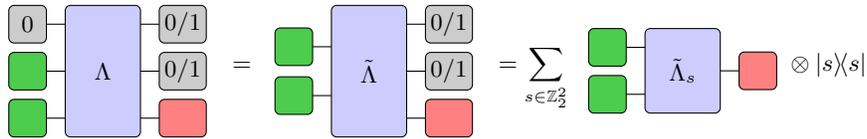


FIG. 2. Each circuit fragment can be identified with a density operator Λ on the joint Hilbert space of its input (left) and output (right) qubits. Classical inputs and outputs of a fragment (gray color) correspond to qubits that are either prepared in the trivial state $|0\rangle$ (labeled “0”) or measured in a fixed computational basis (labeled “0/1”), as appropriate. Quantum inputs (green) and outputs (red), meanwhile, correspond to qubits associated with cuts in a circuit. Due to the presence of trivial inputs, we only need to characterize a reduced state $\tilde{\Lambda}$ on the Hilbert space of the quantum inputs and all outputs. Classical outputs, meanwhile, give the reduced state $\tilde{\Lambda} = \sum_s \tilde{\Lambda}_s \otimes |s\rangle\langle s|$ a block-diagonal structure, where the block $\tilde{\Lambda}_s$ is associated with the measurement of bitstring s on the classical outputs of the fragment. Fragment tomography is performed by providing a variety of quantum inputs to $\tilde{\Lambda}$, and measuring its quantum outputs in a variety of bases. The blocks $\tilde{\Lambda}_s$ are then inferred by least-squares fitting to a linear operator on the quantum inputs, using all available data from experiments in which bitstring s was observed on the classical outputs. This procedure yields an experimental ansatz state Λ_A that approximates $\tilde{\Lambda}$, but which generally does not have the properties required of a density operator, such as a non-negative spectrum. The last step in MLFT is therefore to convert the ansatz state Λ_A into a “maximum likelihood” state Λ_{ML} by using the algorithm of Ref. [17].

the quantum input degrees of freedom in $\tilde{\Lambda}$ by preparing an informationally complete set of states, making use of the fact that

$$\begin{aligned} \langle \rho_i^T \otimes \sigma_o \otimes z_c \rangle_{\tilde{\Lambda}} &= \text{tr} [\mathcal{E}_{\tilde{\Lambda}}(\rho_i)(\sigma_o \otimes z_c)] \\ &= \langle \sigma_o \otimes z_c \rangle_{\mathcal{E}_{\tilde{\Lambda}}(\rho_i)}, \end{aligned} \quad (18)$$

where ρ_i^T denotes the transpose of ρ_i . Whereas the operators σ_o and z_c may be still be chosen from the set of Pauli strings, the input state ρ_i is restricted to satisfy $\text{tr} \rho_i = 1$. This restriction excludes the possibility of choosing states ρ_i from an *orthogonal* basis for the space of the space of Q_i -qubit operators (such as the set of Pauli strings), but any *complete* basis will suffice. For example, one can choose input states from the basis of pure states $\{|0\rangle, |1\rangle, |0\rangle + |1\rangle, |0\rangle + i|1\rangle\}^{\otimes Q_i}$. For an unbiased basis, one can take tensor products of symmetric informationally complete (SIC) states of a single qubit, although the practical advantages of using this basis generally depends on the fidelity with which one can prepare the SIC states. Overall, for a given fragment one must prepare each of 4^{Q_i} input states, and measure outputs in each of 3^{Q_o} possible bases (for each quantum output qubit, the diagonal bases of X, Y, Z), so fragment tomography requires $O(4^{Q_i} 3^{Q_o})$ experiments.

After collecting an informationally complete set of data on the state $\tilde{\Lambda}$, a straightforward least-squares fitting procedure yields an empirical ansatz Λ_A for $\tilde{\Lambda}$, which is the MLFT analogue of the “experimentally noisy” matrix μ described in the original MLST work[17]. The block diagonal structure of $\tilde{\Lambda} = \sum_s \tilde{\Lambda}_s \otimes |s\rangle\langle s|$ implies that the least-squares fitting procedure can be performed independently for each block $\tilde{\Lambda}_s$ of size $2^{Q_i+Q_o} \times 2^{Q_i+Q_o}$. Specifically, $\tilde{\Lambda}_s$ is obtained by fitting to

$$\text{tr} [\tilde{\Lambda}_s(\sigma_i \otimes \sigma_o)] = p_s \langle \sigma_i \otimes \sigma_o \rangle_s, \quad (20)$$

where p_s is the probability of observing bitstring s on the classical output of a fragment, and $\langle \sigma_i \otimes \sigma_o \rangle_s$ is the

experimental expectation value of $\sigma_i \otimes \sigma_o$ when observing bitstring s . As it is constructed from a fit to noisy experimental data, the ansatz state Λ_A will generally have negative eigenvalues, which is not allowed for density operators. The final step in both MLST and MLFT is therefore to find the closest state to Λ_A that has no negative eigenvalues. To this end, MLFT borrows the “fast algorithm for subproblem 1” in Ref. [17], which essentially

- (i) diagonalizes Λ_A ,
- (ii) eliminates the most negative eigenvalue (setting it to zero),
- (iii) adds an equal amount to all other eigenvalues to enforce $\text{tr} \Lambda_A = 1$, and
- (iv) repeats steps (ii,iii) until there are no more negative eigenvalues.

As proven in Ref. [17], this algorithm finds the closest positive semidefinite state Λ_{ML} to Λ_A with respect to the metric induced by the 2-norm: $\|A\|_2 \equiv \sqrt{\text{tr}(A^\dagger A)}$. In this sense, Λ_{ML} is the “most likely” state consistent with Λ_A . The only additional consideration for this algorithm when performing MLFT concerns making use of block diagonal structure to diagonalize Λ_A : each block of size $2^{Q_i+Q_o} \times 2^{Q_i+Q_o}$ can be diagonalized independently. The overall serial runtime of the algorithm to find Λ_{LM} from Λ_A is therefore $O(2^{3(Q_i+Q_o)} N_c)$, where N_c is the number of blocks in Λ_A , or equivalently the number of distinct bitstrings observed on the classical output of the fragment throughout tomography.

As an added bonus, the treatment of fragments and their dual states Λ as first-class objects in MLFT enables a straightforward tensor-network-based circuit reconstruction method, which in turn allows for the use of efficient numerical routines for circuit reconstruction. Rather than explicitly computing and summing over each term of the fragment recombination formula in Eq. (4),

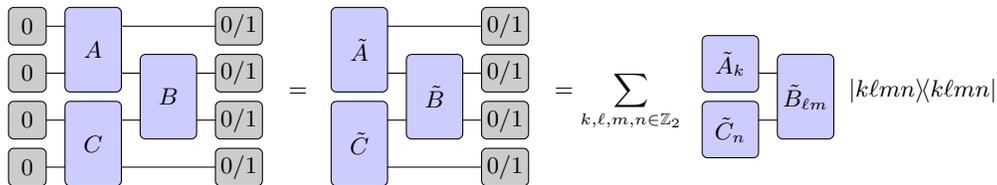


FIG. 3. Fragment recombination of three circuit fragments A, B, C as a tensor network contraction problem. The full probability distribution over measurement outcomes for a reconstructed circuit can be represented by a tensor contraction of the reduced states $\tilde{A}, \tilde{B}, \tilde{C}$, obtained by performing MLFT on the fragments. The probability to measure a given bitstring $klmn$ (i.e. a concatenation of $k, \ell, m, n \in \mathbb{Z}_2$) on the output of the fragment, meanwhile, is given by the contraction of the diagonal blocks $\tilde{A}_k, \tilde{B}_{\ell m}, \tilde{C}_n$. Note that the lack of classical inputs to fragment B implies that $\tilde{B} = B$.

the basic idea is to think of the sum over a complete basis as a contraction of two tensors. We sketch out this idea in Figure 3, making use of the relationship between fragment states Λ , their reductions $\tilde{\Lambda}$, and diagonal blocks $\tilde{\Lambda}_s$. In total, the full probability distribution over measurement outcomes for a reconstructed circuit can be acquired by a tensor network contraction of reduced states $\tilde{\Lambda}$, and the individual probabilities of measuring any given bitstring at the output of a circuit can be acquired by a similar contraction of diagonal blocks $\tilde{\Lambda}_s$.

V. EXPERIMENTAL DEMONSTRATION

In order to test the benefits of MLFT in an application-agnostic setting, we run classical simulations of random unitary circuits (RUCs). As the cost of circuit cutting scales exponentially with the number of cuts made to a circuit, we construct RUCs with a structure that makes them amenable to circuit cutting (see Figure 4). We then vary the number of qubits in a clustered RUC, as well as the number of clusters. For each choice of a number of qubits and a number of clusters, we fix a total budget of 10^6 queries to quantum hardware (known as “shots” or “trials” in e.g. Qiskit[21] or pyQuil[22]), and consider three methods to estimate the probability distribution over measurement outcomes at the end of a clustered RUC.

First, as a “standard” benchmark we consider sampling an entire circuit 10^6 times without any circuit cutting or reconstruction, which we refer to as the method of *full* circuit execution. Second, we consider cutting a circuit into fragments, with each fragment corresponding to a cluster as shown in Figure 4, and reconstructing these fragments essentially as prescribed by the original circuit cutting work[12], which we refer to as the *direct* method of circuit cutting and reconstruction^b. A fragment with Q_i quantum inputs and Q_o quantum outputs

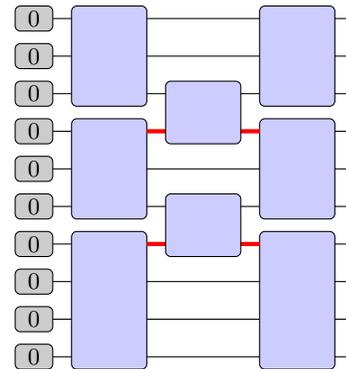


FIG. 4. Random unitary circuit (RUC) of 10 qubits split into 3 clusters. Qubits are first split among clusters as evenly as possible, and each cluster is prepared in a random state by the application of a Haar-random unitary gate. Adjacent clusters are then entangled with random 2-qubit gates, before again applying a layer of random unitaries on all clusters. A clustered RUC is cut into fragments by cutting the bottom legs (red) of every inter-cluster entangling gate.

has $4^{Q_i} \times 3^{Q_o}$ *variants* that must be simulated for circuit reconstruction, where each variant corresponds to a choice of state preparations and measurement bases on the quantum inputs and outputs of the fragment. We divide the budget of 10^6 queries evenly among all fragment variants, and simulate all variants accordingly to characterize and recombine fragments. Finally, we consider the full maximum likelihood fragment tomography and recombination procedure outlined in Section IV, which we will refer to as the MLFT method. Note that the direct and MLFT methods only differ in the classical post-processing of fragment simulation results.

To compare the efficacy of these methods, we compute the fidelity of each estimated probability distribu-

^b To minimize classical computing costs, our actual implementation of the “direct” method involves performing the fragment tomography and recombination procedures discussed in Section

IV, skipping the “maximum likelihood” corrections to fragment models. This procedure is mathematically equivalent (which is to say that it yields an identical output) to that in Ref. [12].

tion over measurement outcomes, p_{estimate} , with the actual probability distribution p_{actual} that is determined by exact classical simulations of a circuit:

$$f = \left[\sum_s \sqrt{p_{\text{actual}}(s) p_{\text{estimate}}(s)} \right]^2, \quad (21)$$

where $p_{\text{actual}}(s)$, $p_{\text{estimate}}(s)$ are respectively the probabilities of measuring the N -qubit state (bitstring) $s \in \mathbb{Z}_2^N$ according to the distributions $p_{\text{actual}}, p_{\text{estimate}}$. The fidelity defined in Eq. (21) is an analogue of the quantum state overlap $|\langle \phi | \psi \rangle|^2$ for classical probability distributions. The only caveat in our calculation of fidelities is that these fidelities only make sense when dealing with valid (non-negative and normalized) probability distributions, whereas the direct circuit cutting method generally yields an unnormalized distribution that may have negative entries. We therefore convert the distribution yielded by the direct method into a valid probability distribution by eliminating all negative entries (setting them to zero), and normalizing the distribution.

Figure 5 shows the infidelities, $1 - f$, of the probability distributions yielded by each simulation method. To ensure that results are not sensitive to the specific choice of random gates, these infidelities are averaged over 100 instances of each clustered RUC. Figure 5 also shows the standard deviation σ_f in the fidelity f over each set of 100 instances, which quantifies the robustness of each method with respect to circuit variations (for a fixed circuit structure).

The key takeaway from Figure 5 is that the MLFT method introduced in this work always outperforms the direct method from past work. MLFT infidelities are always lower than direct infidelities, and are more robust to circuit variations. This result is consistent with theoretical arguments that MLFT finds the “most likely” fragment model consistent with noisy experimental data. Though we only consider shot noise in this work, it would be interesting to see how the benefits of MLFT change with the introduction of additional noise such as measurement and gate errors. We defer a study of the effect of such errors to future work.

Both the direct and MLFT methods perform worse than full circuit execution for small circuits, but their fidelities scale more favorably with circuit size, and they eventually outperform the method of full circuit execution. This result is surprising at first glance, in light of the fact that the circuit cutting methods use strictly fewer quantum computing resources: their 10^6 query budget is spent on executing smaller circuits (namely, fragment variants). The better performance of the circuit cutting methods can be understood by the fact that they use their query budget in a targeted manner that exploits circuit structure, rather than blindly sampling the entire circuit. However, when circuits are sufficiently small for the number of queries (10^6) to explore the sample space of the entire circuit (2^Q , for a Q -qubit circuit), full circuit sampling performs better than circuit cutting

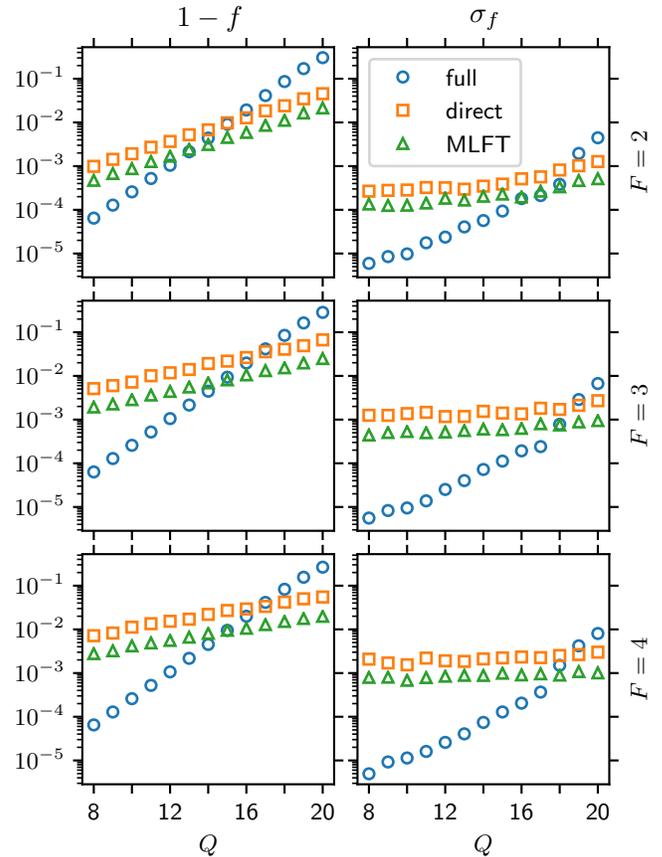


FIG. 5. Average infidelity ($1 - f$, left column) and the standard deviation of the fidelity (σ_f , right column) in simulations of clustered random unitary circuits (RUCs) with Q qubits and $F = 2, 3, 4$ clusters (top, middle, and bottom rows). The number of clusters determines the corresponding circuit structure according to Figure 4. Blue circle, orange square, and green triangle markers respectively correspond to simulations of the full circuit (“full”), and simulations via cutting and recombination before and after maximum likelihood corrections (respectively, “direct” and “MLFT”). Each point is computed by simulating 100 instances of a clustered RUC, with a budget of 10^6 quantum hardware queries for each RUC.

because it does not waste resources on characterizing the “virtual” degrees of freedom associated with the quantum inputs and outputs of fragments.

We can also understand the better scaling of the circuit cutting methods by considering the difficulty of estimating a probability distribution defined by a Q -qubit RUC by (i) sampling of the full circuit directly, versus (ii) sampling all fragment variants for circuit reconstruction. The first task requires, in principle, exploring a sample space of size 2^Q . If a circuit is cut into F fragments with K cuts, then each fragment will consist of $\sim Q/F$ qubits and have $\sim K/F$ quantum inputs and outputs, for a total of $\sim 4^{K/F} \times 3^{K/F}$ fragment variants. Altogether, circuit cutting requires exploring a sample space of size $\sim 2^{Q/F}$ for $\sim F \times 4^{K/F} \times 3^{K/F}$

fragment variants. If the number of cuts (K) is independent of the number of qubits (Q), as is the case with our clustered RUCs, then circuit cutting reduces the overall sample space volume from 2^Q to $O(2^{Q/F})$. This argument is loosely mirrored by the observation that the infidelity $1 - f$ as a function of qubit number Q in Figure 5 has a slope $\partial \log(1 - f) / \partial Q \approx 0.3/F$ (where we treat the “full” method as having only one “fragment”). The linear growth of log-infidelity must eventually end, as $\log(1 - f)$ is bounded above by 1, but away from this bound a slope $\partial \log(1 - f) / \partial Q \propto 1/F$ is consistent with a problem difficulty that scales exponentially in Q/F .

VI. CONCLUSIONS AND OUTLOOK

Circuit cutting is a promising technique for simulating clustered quantum circuits with reduced quantum resource requirements. We have introduced improved circuit cutting methods by minimizing associated classical computing costs, and using maximum likelihood fragment tomography (MLFT) to reconstruct the “most likely” probability distribution defined by a quantum circuit, given available data about its fragments. To test our ideas in an application-agnostic setting, we ran classical simulations of clustered random unitary circuits, which demonstrate the advantages of MLFT over the original circuit cutting method. Moreover, we provide both numerical evidence and theoretical arguments for the advantages of circuit cutting as a technique for running clustered circuits on quantum hardware, even when all hardware requirements are satisfied.

Our work opens several avenues for the improvement and application of circuit cutting techniques. For example, MLFT enforces that fragment models satisfy appropriate self-consistency conditions, but MLFT makes no use of the fact that each fragment corresponds to a *uni-*

tary quantum channel. Furthermore, our present work has neglected the effects of hardware errors that are important to consider in the context of NISQ devices. Due to the capability of MLFT to mitigate shot noise, we expect the advantages of MLFT over full circuit execution to be enhanced when additionally considering the effects of hardware errors. We likewise expect unitarity constraints to provide additional benefits for mitigating various sources of noise. Our work is thus complementary to on-going efforts[23] that study the benefits of circuit cutting in the presence of hardware errors.

As a final point, we note that circuit cutting in its current form estimates a probability distribution associated with a given circuit. Ideally, one would like to *sample* this probability distribution, avoiding the need to reconstruct a distribution that is ultimately defined over an exponentially large space of possible measurement outcomes. To this end, our work makes important progress in understanding the mechanics of circuit cutting, by providing a convenient and economical framework for thinking about individual circuit fragments. We hope that this framework will help in achieving the ultimate goal of sampling a quantum circuit by sampling its fragments.

ACKNOWLEDGMENTS

We acknowledge helpful discussions with Yuri Alexeev, Bradley Pearlman, Teague J. Tomesh, and Wei Tang. This material is based upon work supported by Laboratory Directed Research and Development (LDRD) funding from Argonne National Laboratory, provided by the Director, Office of Science, of the U.S. Department of Energy under contract DE-AC02-06CH11357. 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.

-
- [1] J. Preskill, Quantum Computing in the NISQ era and beyond, [Quantum](#) **2**, 79 (2018).
 - [2] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven, Characterizing quantum supremacy in near-term devices, [Nature Physics](#) **14**, 595 (2018).
 - [3] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. S. L. Brandao, D. A. Buell, *et al.*, Quantum supremacy using a programmable superconducting processor, [Nature](#) **574**, 505 (2019).
 - [4] S. Lloyd, Universal Quantum Simulators, [Science](#) **273**, 1073 (1996).
 - [5] I. M. Georgescu, S. Ashhab, and F. Nori, Quantum simulation, [Reviews of Modern Physics](#) **86**, 153 (2014).
 - [6] E. Farhi, J. Goldstone, and S. Gutmann, A Quantum Approximate Optimization Algorithm, [arXiv:1411.4028 \[quant-ph\]](#) (2014).
 - [7] S. Hadfield, Z. Wang, B. O’Gorman, E. G. Rieffel, D. Venturelli, and R. Biswas, From the Quantum Approximate Optimization Algorithm to a Quantum Alternating Operator Ansatz, [Algorithms](#) **12**, 34 (2019).
 - [8] N. Moll, P. Barkoutsos, L. S. Bishop, J. M. Chow, A. Cross, D. J. Egger, S. Filipp, A. Fuhrer, J. M. Gambetta, M. Ganzhorn, *et al.*, Quantum optimization using variational algorithms on near-term quantum devices, [Quantum Science and Technology](#) **3**, 030503 (2018).
 - [9] V. Dunjko, J. M. Taylor, and H. J. Briegel, Quantum-Enhanced Machine Learning, [Physical Review Letters](#) **117**, 130501 (2016).
 - [10] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, Quantum machine learning, [Nature](#) **549**, 195 (2017).
 - [11] S. Bravyi, G. Smith, and J. A. Smolin, Trading Classical and Quantum Computational Resources, [Physical Review X](#) **6**, 021043 (2016).

- [12] T. Peng, A. Harrow, M. Ozols, and X. Wu, Simulating large quantum circuits on a small quantum computer, [arXiv:1904.00102 \[quant-ph\]](https://arxiv.org/abs/1904.00102) (2019).
- [13] 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* **111**, 2193 (2007).
- [14] 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* **112**, 7061 (2008).
- [15] 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* **112**, 632 (2012).
- [16] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, A variational eigenvalue solver on a photonic quantum processor, *Nature Communications* **5**, 4213 (2014).
- [17] J. A. Smolin, J. M. Gambetta, and G. Smith, Efficient Method for Computing the Maximum-Likelihood Quantum State from Measurements with Additive Gaussian Noise, *Physical Review Letters* **108**, 070502 (2012).
- [18] A. Jamiolkowski, Linear transformations which preserve trace and positive semidefiniteness of operators, *Reports on Mathematical Physics* **3**, 275 (1972).
- [19] M.-D. Choi, Completely positive linear maps on complex matrices, *Linear Algebra and its Applications* **10**, 285 (1975).
- [20] M. Jiang, Channel-state duality, *Physical Review A* **87**, 10/gghbvp (2013).
- [21] H. Abraham, A. Offei, I. Y. Akhalwaya, G. Aleksandrowicz, T. Alexander, G. Alexandrowics, E. Arbel, A. Asfaw, C. Azaustre, A. Ngoueya, *et al.*, *Qiskit: An open-source framework for quantum computing* (2019).
- [22] R. S. Smith, M. J. Curtis, and W. J. Zeng, A practical quantum instruction set architecture (2016), [arXiv:1608.03355 \[quant-ph\]](https://arxiv.org/abs/1608.03355).
- [23] T. Ayrál, F.-M. L. Régent, Z. Saleem, Y. Alexeev, and M. Suchara, Quantum Divide and Compute: Hardware Demonstrations and Noisy Simulations (ISVLSI, 2020).