

## ARTICLE OPEN



# Quantum circuit cutting with maximum-likelihood tomography

Michael A. Perlin<sup>1,2</sup>✉, Zain H. Saleem<sup>3</sup>, Martin Suchara<sup>3</sup> and James C. Osborn<sup>2,4</sup>

We introduce maximum-likelihood fragment tomography (MLFT) as an improved circuit cutting technique for running clustered quantum circuits on quantum devices with a limited number of qubits. In addition to minimizing the classical computing overhead of circuit cutting methods, MLFT finds the most likely probability distribution for the output of a quantum circuit, given the measurement data obtained from the circuit's fragments. We demonstrate the benefits of MLFT for accurately estimating the output of a fragmented quantum circuit with numerical experiments on random unitary circuits. Finally, we show 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.

*npj Quantum Information* (2021)7:64; <https://doi.org/10.1038/s41534-021-00390-6>

## INTRODUCTION

The advent of noisy intermediate-scale quantum (NISQ) technologies<sup>1</sup> makes quantum 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, with state-of-the-art superconducting quantum processors reaching ~50 qubits with percent-level gate and readout errors<sup>2–4</sup>. Advances on the hardware front have been matched by the theoretical development of suitable hardware benchmarks<sup>5</sup>, which have in turn enabled proof-of-principle demonstrations of a computational advantage over classical computing systems<sup>4</sup>.

Despite tremendous progress, existing devices still lack the number and quality of qubits required for practical NISQ-era applications such as digital quantum simulation<sup>6,7</sup>, quantum optimization<sup>8–10</sup>, and quantum machine learning<sup>11,12</sup>. Without error correction, these applications are severely limited by the accumulation of errors that will only compound as devices scale up to more qubits and deeper circuits. Bridging the gap between the requirements of NISQ-era quantum algorithms and the capabilities of NISQ devices will require error mitigation techniques<sup>13,14</sup> and problem decompositions that trade quantum and classical computing resources<sup>15,16</sup>.


One decomposition, inspired by the fragmentation methods used for quantum molecular cluster simulations<sup>17–19</sup>, applies fragmentation to the execution of quantum circuits<sup>16</sup>. This decomposition consists of first “cutting” a quantum circuit into smaller subcircuits, or “fragments”, that can be executed on processors with fewer qubits, and then reconstructing the probability distribution over measurement outcomes for the original quantum circuit from probability distributions associated with its fragments. The severed quantum connections between circuit fragments are simulated by classical post-processing of fragment data, which leads 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 intercluster interactions. Such circuits can make appearances in the context of Hamiltonian simulation<sup>16</sup>, as well as near-term applications based on a variational ansatz that allows for some freedom in choosing circuit structure, such as the quantum approximate optimization algorithm (also the quantum alternating operator ansatz)<sup>8,9,20,21</sup> and variational quantum eigensolvers<sup>16,22</sup>.

Due to the presence of fundamental shot noise (equivalently, finite sampling error), an unavoidable feature of the original fragment recombination method in ref. <sup>16</sup> 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)<sup>23</sup>, 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 and introducing maximum-likelihood fragment tomography (MLFT), the use of which guarantees that reconstructed probability distributions are non-negative and normalized. We discuss how MLFT minimizes the classical computing resources necessary to characterize circuit fragments, and provide a tensor-network-based method for fragment recombination. We test our methods in numerical experiments with random unitary circuits (RUCs), and demonstrate that MLFT estimates the probability distribution at the output of a fragmented quantum circuit with higher fidelity than the naive method of removing negative probabilities and normalizing. These benefits come at no cost to the computational complexity of circuit cutting, as they are achieved by post-processing fragment data in a manner that has a smaller computational cost than that of recombining fragment data to reconstruct a circuit output. As an added bonus, for a fixed number of queries to quantum hardware (known as “shots” or “trials” in e.g., Qiskit<sup>24</sup> or pyQuil<sup>25</sup>) we show that circuit cutting methods can outperform direct execution and sampling of a

<sup>1</sup>JILA, National Institute of Standards and Technology and University of Colorado, Boulder, CO, USA. <sup>2</sup>Argonne Leadership Computing Facility, Argonne National Laboratory, Lemont, IL, USA. <sup>3</sup>Mathematics and Computer Science Division, Argonne National Laboratory, Lemont, IL, USA. <sup>4</sup>Computational Science Division, Argonne National Laboratory, Lemont, IL, USA. ✉email: mika.perlin@gmail.com



**Fig. 1 Circuit cutting example.** A 3-qubit GHZ circuit can be 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 labeled by the state  $M_s$  ( $M_r$ ) correspond to preparations (projections) of a qubit in the corresponding state. 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 (precut) circuit.

clustered circuit in order to estimate its associated probability distribution. We provide theoretical arguments to support this finding, which motivates the use of circuit cutting as a standard tool for evaluating clustered circuits on quantum hardware, even when all hardware requirements for full circuit execution are satisfied.

In the remainder of this section, we provide a basic overview and discussion of the circuit cutting procedure first introduced in ref. <sup>16</sup>, and establish terminology that we will use throughout the rest of this work. We note that our overview of circuit cutting will use the language of quantum states and channels, rather than the language of tensor networks that was used in ref. <sup>16</sup>. These two formalisms are mathematically equivalent, but the former will allow for a more seamless integration with the material in the following section.

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  $\langle b| \psi\rangle$  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 subcircuits. For example, if  $|\psi\rangle = V_{23}U_{12}|000\rangle$ , where  $U_{12}$  and  $V_{23}$  are the two-qubit gates  $U$  and  $V$  acting on qubits 1, 2 and 2, 3, then by inserting the identity operator  $I_2$  (on qubit 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 Fig. 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  $\rho$  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 just as well to 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  $\rho$  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  $2 \times 2$  matrices with normalization  $\text{tr}[M^{(i)}M^{(j)}] = 2\delta_{ij}$  for  $M^{(i)}, M^{(j)} \in \mathcal{B}$ ;  $\text{tr}_n$  denotes a partial trace with respect to qubit  $n$ ; and  $M_n$  with  $n \in \mathbb{Z}_n$  denotes an operator that acts with  $M$  on qubit  $n$  and trivially (i.e., with the identity  $I$ ) on all other qubits. To be concrete, we will use the set of Pauli operators together with the single-qubit identity operator,  $\mathcal{B} \equiv \{X, Y, Z, I\}$ , 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

$$\rho_1(M) \equiv \text{tr}_2(M_2 U |0\rangle\langle 0|^{\otimes 2} U^\dagger), \quad \rho_2(M) \equiv V(M \otimes |0\rangle\langle 0|) V^\dagger, \quad (6)$$

have no straightforward interpretation as “conditional” states, as with  $|\psi_1(b)\rangle$  and  $|\psi_2(b)\rangle$  in Eq. (2). In order to decompose  $\rho$  into conditional states, we can expand each  $M \in \mathcal{B}$  in its eigenbasis

$$\rho \simeq \frac{1}{2} \sum_{\substack{M \in \mathcal{B} \\ r,s \in \lambda(M)}} rs \rho_1(M_r) \otimes \rho_2(M_s), \quad (7)$$

where  $\lambda(M)$  denotes the spectrum of  $M$ , i.e.,  $\lambda(X) = \lambda(Y) = \lambda(Z) = (+1, -1)$  and  $\lambda(I) = (1, 1)$ ; and  $M_s \equiv |M_s\rangle\langle M_s|$  with  $s \in \lambda(M)$  is a projector onto an eigenstate 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 reuse the eigenstates from one of the other operators.

The decomposition in Eq. (7) 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 (see Fig. 1). This decomposition 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, and the normalization  $\text{tr}\rho_1(M_s)$  is 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. (7). This scenario is illustrated in Fig. 1, which cuts a 3-qubit GHZ circuit preparing the state  $|\psi\rangle \propto |000\rangle + |111\rangle$  into two 2-qubit fragments.

In practice, recombining circuit fragments as prescribed by Eq. (7) is inefficient in two ways. First, the tensor products in Eq. (7) 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  $\rho_f(l) = \rho_f(M_{+1}) + \rho_f(M_{-1})$  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 reduction (in  $K$ ) of the number of floating-point operations required to recombine fragment data. Note that the recombination procedure in ref. <sup>16</sup> involves  $8^K$  tensor products, rather than  $16^K$ , because it consolidates “measurement” conditions, but not “preparation” conditions, which is equivalent to collapsing the sum over  $r$  in Eq. (7) but leaving the sum over  $s$ .

Second, the recombination formula in Eq. (7) nominally requires, for each fragment  $f$  incident on  $K_f$  cuts, characterizing  $K_f^6$  probability distributions. This characterization is overcomplete, because the  $K_f^6$  distributions are not all linearly independent. In the case of a fragment with a single incident cut, for example, we can use the fact that  $X_+ + X_- = Z_+ + Z_- = I$  to decompose

$$\rho_f(X_-) = \rho_f(Z_+) + \rho_f(Z_-) - \rho_f(X_+). \quad (8)$$

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  $\{\Pi_j^{\text{SIC}} : j \in \mathbb{Z}_4\}$  (consisting of projectors  $\Pi_j^{\text{SIC}}$  onto the states represented by the four corners of a regular tetrahedron inscribed in a 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. (8). When combining these distributions according to Eqs. (5) and (7) 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 address these shortcomings, in the following section we recast the task of characterizing conditional distributions into the task of performing fragment tomography, treating the fragments  $\rho_f$ , rather than distributions  $\rho_f(M_s)$ , as first-class objects. In addition to being automatically efficient in terms of the classical memory footprint of characterizing each fragment, performing fragment tomography allows us to adapt the method of maximum-likelihood tomography<sup>23</sup> 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 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.

## RESULTS

### Maximum-likelihood fragment tomography (MLFT)

Once a circuit has been cut into fragments  $\rho_f$ , rather than characterizing conditional distributions  $\rho_f(M_s)$  we can perform a more systematic 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. <sup>23</sup>, 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 is a type of quantum process tomography, which generalizes MLST<sup>23</sup> 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  $|\mathbf{0}\rangle_i \equiv |0\rangle_i^{\otimes C_i}$ , 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  $\mathcal{E}_\Lambda$  on the state of  $Q$  qubits. The channel-state duality<sup>26–28</sup> implies that this channel is uniquely determined by a 4-partite state (density operator) of the form

$$\Lambda \equiv \sum_{\substack{k,\ell,m,n \\ p,q,r,s}} \Lambda_{k\ell;mn;pq;rs} |k\rangle \langle \ell| \otimes |m\rangle \langle n| \otimes |p\rangle \langle q| \otimes |r\rangle \langle s|, \quad (9)$$

where the bitstrings  $k, \ell$  ( $m, n; p, q; r, s$ ) 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  $\mathcal{E}_\Lambda$  maps a bipartite input state

$$\rho \otimes |\mathbf{0}\rangle_i \langle \mathbf{0}\rangle_i \equiv \sum_{k,\ell} \rho_{k\ell} |k\rangle \langle \ell| \otimes |\mathbf{0}\rangle_i \langle \mathbf{0}\rangle_i \quad (10)$$

at its input to the bipartite state

$$\mathcal{E}_\Lambda(\rho \otimes |\mathbf{0}\rangle_i \langle \mathbf{0}\rangle_i) \equiv \sum_{k,\ell,p,q,r,s} \Lambda_{k\ell;0,0;pq;rs} \rho_{k\ell} |p\rangle \langle q| \otimes |r\rangle \langle s| \quad (11)$$

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

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

where


$$\tilde{\Lambda}_{k\ell;pq;s} \equiv \Lambda_{k\ell;0,0;pq;ss}. \quad (13)$$

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

$$\tilde{\Lambda} \equiv \sum_{k,\ell,p,q,s} \tilde{\Lambda}_{k\ell;pq;s} |k\rangle \langle \ell| \otimes |p\rangle \langle q| \otimes |s\rangle \langle s| = \sum_s \tilde{\Lambda}_s \otimes |s\rangle \langle s|, \quad (14)$$

where Eq. (14) implicitly defines the blocks  $\tilde{\Lambda}_s$ . In words, the reduced state  $\tilde{\Lambda}$  is acquired from the full state  $\Lambda$  by conditioning on (i.e., fixing) a trivial state  $|\mathbf{0}\rangle_i \langle \mathbf{0}\rangle_i$  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  $\Lambda$ ,  $\tilde{\Lambda}$ , and  $\tilde{\Lambda}_s$  is sketched out in Fig. 2.

In a nutshell, MLFT is performed by providing a variety of quantum inputs to  $\mathcal{E}_{\tilde{\Lambda}}$ , and measuring its quantum outputs in a variety of bases. The blocks  $\tilde{\Lambda}_s$  are inferred by least-squares fitting to a linear operator that maps quantum inputs to quantum outputs, using all available data from experiments in which bitstring  $s$  was observed on the classical outputs of a fragment. This procedure yields an experimental ansatz state  $\Lambda_A$  that approximates  $\tilde{\Lambda}$ , but that generally does not have the properties



**Fig. 2 Block-diagonalizing circuit fragments.** Each circuit fragment can be identified with a density operator  $\Lambda$  on the joint Hilbert space of its input (left) and output (right) qubits. Classical inputs and outputs of a fragment (gray) 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”). Quantum inputs (left, green) and outputs (right, red) 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 give this reduced state a block-diagonal structure:  $\tilde{\Lambda} = \sum_s \tilde{\Lambda}_s \otimes |s\rangle\langle s|$ , where the block  $\tilde{\Lambda}_s$  is associated with the measurement of bitstring  $s$  on the classical outputs of the fragment.

required of a density operator, such as a non-negative spectrum. The last step in MLFT is therefore to convert the ansatz state  $\Lambda_A$  into a “maximum-likelihood” state  $\Lambda_{ML}$  by using an algorithm borrowed from MLST in ref. <sup>23</sup>. We describe MLFT in more detail below.

MLFT (and MLST) begins by collecting measurement data to characterize the quantum state under consideration. In the case of the block-diagonal state  $\tilde{\Lambda}$ , one needs to characterize the expectation values

$$\langle \sigma_i \otimes \sigma_o \otimes z_c \rangle_{\tilde{\Lambda}} \equiv \text{tr}[\tilde{\Lambda}(\sigma_i \otimes \sigma_o \otimes z_c)] \quad (15)$$

for some complete basis of operators  $\{\sigma_i \otimes \sigma_o \otimes z_c\}$  on the target Hilbert space of  $\tilde{\Lambda}$ , where  $\sigma_i$ ,  $\sigma_o$ , 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<sup>23</sup> collects data by performing informationally complete measurements of  $\tilde{\Lambda}$ , for example by choosing operators  $\sigma_{i,o}$  from the set of all Pauli strings  $\{I, X, Y, Z\}^{\otimes Q_{i,o}}$ , and choosing  $z_c$  from the set of diagonal Pauli strings  $\{I, Z\}^{\otimes C_o}$ . 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 the quantum input degrees of freedom in  $\tilde{\Lambda}$  by preparing an informationally complete set of states, making use of the fact that

$$\text{tr}[\tilde{\Lambda}(\sigma_i \otimes \sigma_o \otimes z_c)] = \text{tr}[\mathcal{E}_{\tilde{\Lambda}}(\sigma_i^T)(\sigma_o \otimes z_c)] = \langle \sigma_o \otimes z_c \rangle_{\mathcal{E}_{\tilde{\Lambda}}(\sigma_i^T)}, \quad (16)$$

where  $\sigma_i^T$  denotes the transpose of  $\sigma_i$ . Whereas the operators  $\sigma_o$  and  $z_c$  may still be chosen from the set of Pauli strings, the input state  $\sigma_i^T$  is restricted to satisfy  $\text{tr}\sigma_i^T = 1$ . This restriction excludes the possibility of choosing  $\sigma_i^T$  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, or even consider bases of multi-qubit SIC states. The practical advantages of using these bases, however, generally depend on the fidelity with which one can prepare SIC states. Similar considerations apply for the choice of measurement basis for quantum outputs<sup>29</sup>. Overall, in order to characterize a fragment with  $Q_i$  quantum inputs and  $Q_o$  quantum outputs 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  $\Lambda_A$  for  $\tilde{\Lambda}$ , which is the MLFT analogue of the “experimentally noisy” matrix  $\mu$  described in the original MLST work<sup>23</sup>. 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

$$\text{tr}[\tilde{\Lambda}_s(\sigma_i \otimes \sigma_o)] = p_s \langle \sigma_i \otimes \sigma_o \rangle_{z_c=s}, \quad (17)$$


where  $p_s$  is the probability of observing bitstring  $s$  on the classical output of a fragment, and  $\langle \sigma_i \otimes \sigma_o \rangle_{z_c=s}$  is the expectation value of  $\sigma_o$  (on the quantum outputs) when preparing the state  $\sigma_i^T$  (on the quantum inputs) and observing bitstring  $s$  (on the classical outputs) of the fragment. Because the ansatz state  $\Lambda_A \approx \tilde{\Lambda}$  is constructed from a fit to noisy measurement data,  $\Lambda_A$  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  $\Lambda_A$  that has no negative eigenvalues. To this end, MLFT borrows the “fast algorithm for subproblem 1” in ref. <sup>23</sup>, which

- (i) diagonalizes  $\Lambda_A$ ,
- (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. <sup>23</sup>, this algorithm finds the closest positive semidefinite state  $\Lambda_{ML}$  to  $\Lambda_A$  with respect to the metric induced by the 2-norm  $\|A\|_2 \equiv \sqrt{\text{tr}(A^\dagger A)}$ . In this sense,  $\Lambda_{ML}$  is the “most likely” state consistent with  $\Lambda_A$ . The only additional consideration for this algorithm when performing MLFT has to do with making use of block-diagonal structure to diagonalize  $\Lambda_A$ : 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  $\Lambda_{ML}$  from  $\Lambda_A$  is therefore  $O(2^{3(Q_i+Q_o)} N_c)$ , where  $N_c \leq 2^{C_o}$  is the number of blocks in  $\Lambda_A$ , or equivalently the number of distinct bitstrings observed on the classical output of the fragment throughout tomography. As we will see, the maximum-likelihood corrections to  $\Lambda_A$  are responsible for the benefits of MLFT in estimating a circuit’s output. Moreover, the cost of computing these corrections is smaller than the unavoidable cost of fragment recombination, so the benefits of MLFT are free as far as the computational complexity of circuit cutting is concerned.

The treatment of fragments and their dual states  $\Lambda$  as first-class objects in MLFT enables a straightforward tensor-network-based circuit reconstruction method. Rather than explicitly computing and summing over each term of the fragment recombination formula in Eq. (4), the basic idea is to think of the entire sum as a contraction of two tensors. We sketch out this idea in Fig. 3, making use of the relationship between fragment states  $\Lambda$ , 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  $\tilde{\Lambda}_s$ .

If a circuit has  $K$  cuts and  $F$  fragments, and  $N_c^{(f)}$  distinct bitstrings were observed on the classical output of fragment  $f \in \{1, 2, \dots, F\}$  throughout fragment tomography, then reconstructing the circuit’s



**Fig. 3** **Fragment recombination as a tensor network contraction problem.** The full probability distribution over measurement outcomes for a circuit reconstructed from fragments  $A, B, C$  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  $k, l, m, n \in \mathbb{Z}_2$ ) on the output of the fragment is given by the contraction of the diagonal blocks  $\tilde{A}_k, \tilde{B}_{lm}, \tilde{C}_n$ . The lack of classical inputs to fragment  $B$  implies that  $\tilde{B} = B$ .



**Fig. 4** **Random unitary circuit (RUC) of ten qubits split into three 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<sup>33,34</sup>. Adjacent clusters are then entangled with random two-qubit gates, before again applying a layer of random unitaries on all clusters. A clustered RUC is cut into fragments (labeled  $A, B, C$ ) by cutting the bottom legs (shown in red) of every intercluster entangling gate.

output requires contracting  $\prod_i N_c^{(f)}$  tensor networks, each of which nominally involves summing over  $4^K$  terms. Whereas the  $4^K$  cost to contract a single tensor network  $g$  can be reduced to  $2^{\text{cc}(g)}$ , where  $\text{cc}(g)$  is the contraction complexity of  $g$ <sup>16</sup>, the overall multiplicative cost in  $N_c^{(f)}$  is unavoidable. In comparison, performing maximum-likelihood corrections to fragment models comes at a cost that is additive in  $N_c^{(f)}$ . For this reason, fragment recombination is generally the computational bottleneck of circuit cutting, and maximum-likelihood corrections add no significant overhead.

### Numerical experiments

In order to test the benefits of MLFT in an application-agnostic setting, we run classical simulations of RUCs. Because 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 Fig. 4). We then vary the number of qubits and clusters in our RUCs, as well as the total number of samples (known as “shots” in Qiskit<sup>24</sup> or “trials” in pyQuil<sup>25</sup>) in a simulation, where the result of each sample is a single bitstring representing one measurement outcome. In this way, we compare 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  $S$  times without any circuit cutting, 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 Fig. 4, and reconstructing these fragments as prescribed by the original circuit cutting work<sup>16</sup>, namely without maximum-likelihood corrections. We refer to this second method as the “direct” method of circuit cutting and reconstruction.

A fragment with  $Q_i$  quantum inputs and  $Q_o$  quantum outputs 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 therefore divide the budget of  $S$  samples evenly among all fragment variants. Finally, we consider the full MLFT and recombination procedure, which we refer to as the “MLFT” method. The direct and MLFT methods only differ in the classical post-processing of fragment simulation results. Specifically, the differences between the final outputs of the direct and MLFT methods are entirely due to the application (or non-application) of maximum-likelihood corrections to fragment models.

To compare the efficacy of the full, direct, and MLFT methods, we compute the fidelity of reconstructed probability distributions over measurement outcomes,  $p_{\text{estimate}}$ , with the actual probability distribution  $p_{\text{actual}}$  that is determined by exact classical simulations of a circuit

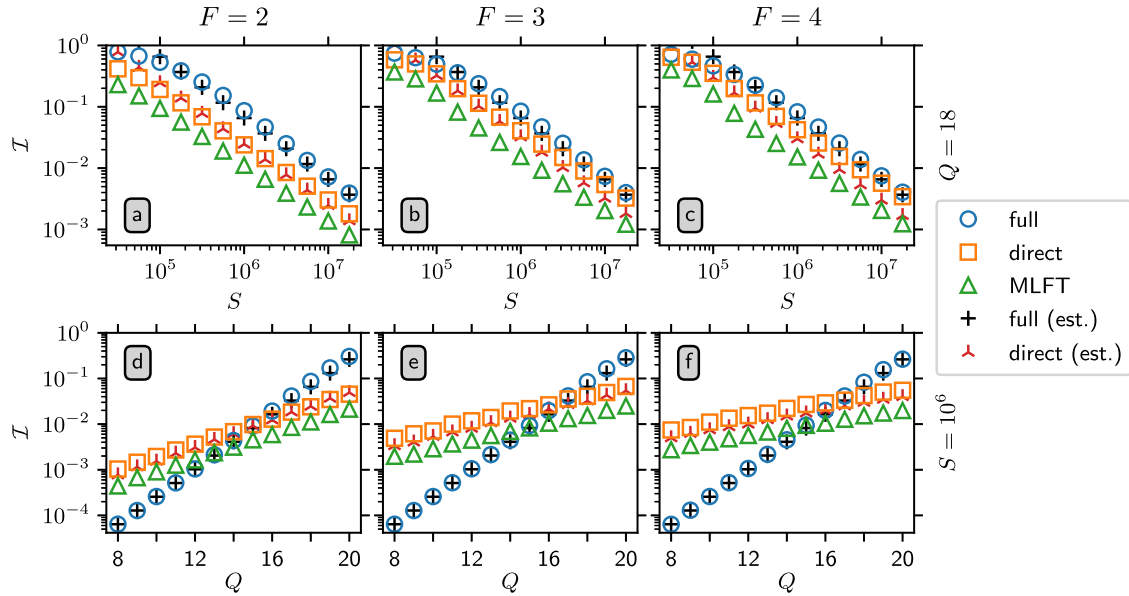
$$\mathcal{F} = \left[ \sum_s \sqrt{p_{\text{actual}}(s)p_{\text{estimate}}(s)} \right]^2, \quad (18)$$

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  $\mathcal{F}$  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 they are only well-defined 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  $\mathcal{I} = 1 - \mathcal{F}$  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, although in practice we find that these infidelities vary by only  $\sim 1$ –10% of their mean value (see Supplementary Information). Figure 5 also shows analytical estimates of infidelity for the full and direct simulation methods, derived in Supplementary Information.

An immediate takeaway from Fig. 5 is that the MLFT method introduced in this work always outperforms the direct method: MLFT infidelities are always lower than direct infidelities. This result is consistent with theoretical arguments that MLFT finds the “most likely” fragment model consistent with noisy measurement data. Although 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.

Figure 5 also shows that the infidelity  $\mathcal{I}$  for all simulation methods scales more or less identically with the sample number  $S$ , namely  $\mathcal{I} \sim 1/S$  for large  $S$ . Though some of the numerical data in Fig. 5 may better be fit by  $\mathcal{I} \propto 1/S^{(1+\eta)}$  for some  $\eta \neq 0$ , the deviation from  $\eta = 0$  are minor, and may be an artifact of small



**Fig. 5** Infidelity in reconstructed circuit outputs. The infidelity  $\mathcal{I} = 1 - \mathcal{F}$  as a function of sample number  $S$  (a–c) or qubit number  $Q$  (d–f) for clustered random unitary circuits (RUCs) with  $F = 2$  (a, d), 3 (b, e), or 4 (c, f) fragments. Open markers correspond to simulations of the full circuit (“full”), or simulations via circuit cutting before (“direct”) and after (“MLFT”) maximum-likelihood corrections to fragment models. The last two markers in the legend correspond to analytical estimates of infidelity:  $2^Q/S$  for the full method, and  $\sum_{f=1}^F 2^{C_f^c}/n$  for the direct method, where  $C_f^c$  is the number of classical outputs on fragment  $f$  and  $n = S/V$  is the number of samples devoted to each of  $V$  total fragment variants. Whereas the estimates for the full method are quantitatively accurate, the estimates for the direct method are provided only to highlight approximate scaling relationships (Supplementary Information). Results for each data point are averaged over 100 instances of a clustered RUC.

circuit sizes. It is worth noting that the original circuit cutting work<sup>16</sup> proved that a reconstructed circuit output (probability distribution) can be estimated to an accuracy of  $\epsilon$  with  $S = O(1/\epsilon^2)$  samples, which by dimensional analysis suggests that  $\mathcal{I} \sim \epsilon^2 \sim 1/S$  in all cases.

Though scaling with sample number does not strongly distinguish these methods, it is clear that the direct and MLFT methods scale much more favorably with circuit size: the full method has an infidelity  $\mathcal{I} \sim 2^Q$  for  $Q$  qubits, whereas cutting a circuit into  $F$  fragments results in  $\mathcal{I} \sim \sum_{f=1}^F 2^{C_f^c}$ , where  $C_f^c \approx Q/F$  is the number of classical outputs on fragment  $f$  and  $\sum_f C_f^c = Q$ . The more favorable scaling for circuit cutting methods is surprising at first glance, as these methods require strictly fewer quantum computing resources: their sample 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 sample budget in a targeted manner that exploits circuit structure, rather than blindly sampling the entire circuit. However, when circuits are sufficiently small for the fixed number of samples to explore the sample space of the entire circuit, full circuit sampling performs better than circuit cutting because it does not waste resources on characterizing numerous variants of nearly identical fragments.

Deferring a detailed derivation of expected infidelities to the Supplementary Information, we can make the above intuition more quantitative by considering the difficulty of estimating a probability distribution defined by a  $Q$ -qubit RUC by (i) sampling 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  $2^Q$  with  $S$  samples, so one might reasonably expect (as is indeed the case) that  $\mathcal{I} \sim 2^Q/S$ . If a circuit is cut into  $F$  fragments, meanwhile, then each fragment will have  $\sim Q/F$  qubits, and if the number of qubits is independent of the number of fragments, then the overall sample space volume is reduced from  $2^Q$  to  $2^{O(Q/F)}$ . Indeed, this argument agrees with the estimate of infidelity for the direct method of circuit cutting in Fig. 5, where

we show that  $\mathcal{I} \sim F \times 2^{Q/F}/n$  with  $n = S/V$  the number of samples devoted to each of  $V$  total fragment variants.

## DISCUSSION

Circuit cutting is a promising technique for reducing the qubit requirements of running clustered quantum circuits. We have introduced improved circuit cutting methods by minimizing associated classical computing costs (with an exponential improvement over previous methods), and by using MLFT to reconstruct the “most likely” probability distribution defined by a quantum circuit, given the measurement data obtained from its fragments. To test our ideas in an application-agnostic setting, we ran classical simulations of RUCs, which demonstrate the advantages of MLFT compared to the original circuit cutting method. Moreover, we also show that circuit cutting has advantages as a standard technique for running clustered circuits on quantum hardware, even when full circuit execution is possible.

Our work opens several avenues for the improvement and application of circuit cutting techniques. For example, MLFT guarantees that fragment models satisfy appropriate self-consistency conditions, but MLFT makes no use of the fact that each fragment corresponds to a unitary quantum channel. Furthermore, our present work neglects the effects of hardware errors that are important to consider in the context of NISQ devices. Because MLFT has the capability 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 sources of noise. Our work thus complements ongoing efforts that study the benefits of circuit cutting in the presence of hardware errors, which have generally found that circuit cutting helps mitigate the effects of noise<sup>30</sup>. Having framed fragment characterization as a tomography task, it would also be interesting to adapt and apply different quantum process tomography techniques<sup>31</sup> to the task of circuit cutting, and compare their performance and cost to that of MLFT.

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 (defined over an exponentially large space of possible measurement outcomes) without having to reconstruct it in full. To this end, our work makes important progress in understanding the mechanics of circuit cutting, by providing a convenient and efficient 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.

## DATA AVAILABILITY

The circuits and simulation data that we used for numerical experiments in this work are available in the online repository [Quantum-Software-Tools/QSPPLIT-MLFT](https://github.com/Quantum-Software-Tools/QSPPLIT-MLFT)<sup>32</sup>.

## CODE AVAILABILITY

The codes used for numerical experiments in this work are available in the online repository [Quantum-Software-Tools/QSPPLIT-MLFT](https://github.com/Quantum-Software-Tools/QSPPLIT-MLFT)<sup>32</sup>.

Received: 5 August 2020; Accepted: 1 February 2021;  
Published online: 23 April 2021

## REFERENCES

1. Preskill, J. Quantum computing in the NISQ era and beyond. *Quantum* **2**, 79 (2018).
2. Wendin, G. Quantum information processing with superconducting circuits: a review. *Rep. Prog. Phys.* **80**, 106001 (2017).
3. Kjaergaard, M. et al. Superconducting qubits: current state of play. *Annu. Rev. Condens. Matter Phys.* **11**, 369–395 (2020).
4. Arute, F. et al. Quantum supremacy using a programmable superconducting processor. *Nature* **574**, 505–510 (2019).
5. Boixo, S. et al. Characterizing quantum supremacy in near-term devices. *Nat. Phys.* **14**, 595–600 (2018).
6. Lloyd, S. Universal quantum simulators. *Science* **273**, 1073–1078 (1996).
7. Georgescu, I. M., Ashhab, S. & Nori, F. Quantum simulation. *Rev. Mod. Phys.* **86**, 153–185 (2014).
8. Farhi, E., Goldstone, J. & Gutmann, S. A quantum approximate optimization algorithm. Preprint at [http://arxiv.org/abs/1411.4028](https://arxiv.org/abs/1411.4028) (2014).
9. Hadfield, S. et al. From the quantum approximate optimization algorithm to a quantum alternating operator ansatz. *Algorithms* **12**, 34 (2019).
10. Moll, N. et al. Quantum optimization using variational algorithms on near-term quantum devices. *Quantum Sci. Technol.* **3**, 030503 (2018).
11. Dunjko, V., Taylor, J. M. & Briegel, H. J. Quantum-enhanced machine learning. *Phys. Rev. Lett.* **117**, 130501 (2016).
12. Biamonte, J. et al. Quantum machine learning. *Nature* **549**, 195–202 (2017).
13. Endo, S., Benjamin, S. C. & Li, Y. Practical quantum error mitigation for near-future applications. *Phys. Rev. X* **8**, 031027 (2018).
14. Kandala, A. et al. Error mitigation extends the computational reach of a noisy quantum processor. *Nature* **567**, 491–495 (2019).
15. Bravyi, S., Smith, G. & Smolin, J. A. Trading classical and quantum computational resources. *Phys. Rev. X* **6**, 021043 (2016).
16. Peng, T., Harrow, A. W., Ozols, M. & Wu, X. Simulating large quantum circuits on a small quantum computer. *Phys. Rev. Lett.* **125**, 150504 (2020).
17. Li, W., Li, S. & Jiang, Y. Generalized energy-based fragmentation approach for computing the ground-state energies and properties of large molecules. *J. Phys. Chem. A* **111**, 2193–2199 (2007).
18. Li, H., Li, W., Li, S. & Ma, J. Fragmentation-based QM/MM simulations: length dependence of chain dynamics and hydrogen bonding of polyethylene oxide and polyethylene in aqueous solutions. *J. Phys. Chem. B* **112**, 7061–7070 (2008).
19. Gordon, M. S., Fedorov, D. G., Pruitt, S. R. & Slipchenko, L. V. Fragmentation methods: a route to accurate calculations on large systems. *Chem. Rev.* **112**, 632–672 (2012).
20. Saleem, Z. H., Tariq, B. & Suchara, M. Approaches to constrained quantum approximate optimization. Preprint at [http://arxiv.org/abs/2010.06660](https://arxiv.org/abs/2010.06660) (2020).

21. Saleem, Z. H., Tomesh, T., Perlin, M. A., Gokhale, P. & Suchara, M. Divide and conquer for constrained quantum approximate optimization. In preparation (2021).
22. Peruzzo, A. et al. A variational eigenvalue solver on a photonic quantum processor. *Nat. Commun.* **5**, 4213 (2014).
23. Smolin, J. A., Gambetta, J. M. & Smith, G. Efficient method for computing the maximum-likelihood quantum state from measurements with additive gaussian noise. *Phys. Rev. Lett.* **108**, 070502 (2012).
24. Abraham, H. et al. Qiskit: an open-source framework for quantum computing. (2019).
25. Smith, R. S., Curtis, M. J. & Zeng, W. J. A practical quantum instruction set architecture. Preprint at [http://arxiv.org/abs/1608.03355](https://arxiv.org/abs/1608.03355) (2017).
26. Jamiolkowski, A. Linear transformations which preserve trace and positive semidefiniteness of operators. *Rep. Math. Phys.* **3**, 275–278 (1972).
27. Choi, M.-D. Completely positive linear maps on complex matrices. *Linear Algebra Appl.* **10**, 285–290 (1975).
28. Jiang, M., Luo, S. & Fu, S. Channel-state duality. *Phys. Rev. A* **87**, 022310 (2013).
29. Adamson, R. B. A. & Steinberg, A. M. Improving quantum state estimation with mutually unbiased bases. *Phys. Rev. Lett.* **105**, 030406 (2010).
30. Ayril, T., Régent, F.-M. L., Saleem, Z., Alexeev, Y. & Suchara, M. Quantum divide and compute: exploring the effect of different noise sources. *SN Comp. Sci.* **2**, 132 (2021).
31. Torlai, G. et al. Quantum process tomography with unsupervised learning and tensor networks. Preprint at [http://arxiv.org/abs/2006.02424](https://arxiv.org/abs/2006.02424) (2020).
32. QSPPLIT-MLFT. <https://github.com/Quantum-Software-Tools/QSPPLIT-MLFT> (2020).
33. Zyczkowski, K. & Kus, M. Random unitary matrices. *J. Phys. A. - Math. Gen.* **27**, 4235–4245 (1994).
34. Emerson, J., Livine, E. & Lloyd, S. Convergence conditions for random quantum circuits. *Phys. Rev. A* **72**, 060302 (2005).

## ACKNOWLEDGEMENTS

We acknowledge helpful discussions with Yuri Alexeev, Bradley Pearlman, Teague J. Tomesh, Wei Tang, Thomas Ayril, and Francois-Marie Le Régent. This material is based on 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; the Argonne Leadership Computing Facility, which is DOE Office of Science User Facility supported under Contract DE-AC02-06CH11357; and the National Science Foundation under Award No. 2037984.

## AUTHOR CONTRIBUTIONS

M.S. motivated the study of circuit cutting, and brought attention to the problem of unphysical (negative and unnormalized) reconstructed probability distributions. J.C.O. had the idea to build maximum-likelihood fragment models. Z.H.S. thought to adapt existing maximum-likelihood state tomography techniques to this task. M.A.P. worked out the theory, wrote the codes, and drafted the manuscript, with aid and consulting from Z.H.S. M.S. and J.C.O. supervised the project. All authors discussed the numerical experiments, interpreted results, and provided critical feedback and contributions to the final manuscript.

## COMPETING INTERESTS

The authors declare no competing interests.

## ADDITIONAL INFORMATION

**Supplementary information** The online version contains supplementary material available at <https://doi.org/10.1038/s41534-021-00390-6>.

**Correspondence** and requests for materials should be addressed to M.A.P.

**Reprints and permission information** is available at <http://www.nature.com/reprints>

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



**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 license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the

article's Creative Commons license 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 license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2021