

Quantum Circuit Cutting 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 Avenue, Lemont, Illinois 60439*

³*Mathematics and Computer Science Division, Argonne National
Laboratory, 9700 S. Cass Avenue, Lemont, Illinois 60439*

⁴*Computational Science Division, Argonne National
Laboratory, 9700 S. Cass Avenue, Lemont, Illinois 60439*

(Dated: 2021-01-31)

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.

* mika.perlin@gmail.com

† zsaleem@anl.gov

‡ msuchara@anl.gov

§ osborn@alcf.anl.gov

INTRODUCTION

The advent of noisy intermediate-scale quantum (NISQ) technologies [1] 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 [2–4]. Advances on the hardware front have been matched by the theoretical development of suitable hardware benchmarks [5], which have in turn enabled proof-of-principle demonstrations of a computational advantage over classical computing systems [4].

Despite tremendous progress, existing devices still lack the number and quality of qubits required for practical NISQ-era applications such as digital quantum simulation [6, 7], quantum optimization [8–10] and quantum machine learning [11, 12]. 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 [13, 14] and problem decompositions that trade quantum and classical computing resources [15, 16].

One decomposition, inspired by the fragmentation methods used for quantum molecular cluster simulations [17–19], applies fragmentation to the execution of quantum circuits [16]. 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 inter-cluster interactions. Such circuits can make appearances in the context of Hamiltonian simulation [16], 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, QAOA) [8, 9, 20, 21] and variational quantum

eigsolvers (VQE) [16, 22].

Due to the presence of fundamental shot noise (equivalently, finite sampling error), an unavoidable feature of the original fragment recombination method in Ref. [16] 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) [23], 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, 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 [24] or pyQuil [25]) we show that circuit cutting methods can outperform direct execution and sampling of a 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. [16], 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. [16]. 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 ${}_n \langle b | \psi \rangle$ 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 $|\psi\rangle = V_{23} U_{12} |000\rangle$, 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 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 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 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 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 ρ into conditional states, we can expand each $M \in \mathcal{B}$ in its eigenbasis:

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

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 $|M_s\rangle$ 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 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 Figure 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 Figure 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 $\rho_f(I) = \rho_f(M_{+1}) + \rho_f(M_{-1})$ 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 reduction (in K) of the number of floating-point operations required to recombine fragment data. Note that the recombination procedure in Ref. [16] 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 (SIC-POVM) $\{\Pi_j^{\text{SIC}} : j \in \mathbb{Z}_4\}$ (consisting of projectors Π_j^{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 ρ_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 [23] 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

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. [23], 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 maximum likelihood state tomography (MLST) [23] 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 [26–28] 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, 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

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

at its input to the bipartite state

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

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 Λ by conditioning on (i.e. fixing) a trivial state $|0_i\rangle\langle 0_i|$ 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.

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 Λ_A that approximates $\tilde{\Lambda}$, but that 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 an algorithm borrowed from MLST in Ref. [23]. 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

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

for some complete basis of operators $\{\sigma_i \otimes \sigma_o \otimes z_c\}$ on the target Hilbert space of $\tilde{\Lambda}$, where 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 [23] 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 $\{I, Z\}^{\otimes Q_o}$. 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

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

where σ_i^T denotes the transpose of σ_i . Whereas the operators σ_o and z_c may still be chosen from the set of Pauli strings, the input state σ_i^T is restricted to satisfy $\text{tr} \sigma_i^T = 1$. This restriction excludes the possibility of choosing σ_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 [29]. 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 Λ_A for $\tilde{\Lambda}$, which is the MLFT analogue of the ‘experimentally noisy’ matrix μ described in the original MLST work [23]. 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

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

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 σ_o (on the quantum outputs) when preparing the state σ_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, 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. [23], which

- (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. [23], 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 has to do with 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_{ML} from Λ_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 Λ_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 Λ_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 Λ 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 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 $\tilde{\Lambda}_s$.

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 output requires contracting $\prod_f N_c^{(f)}$ 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 $2^{O(\text{cc}(g))}$, where $\text{cc}(g)$ is the contraction complexity of g [16], 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 random unitary circuits (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 Figure 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 [24] or ‘trials’ in pyQuil [25]) 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 Figure 4, and reconstructing these fragments as prescribed by the original circuit cutting work [16], 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_{estimate} , with the actual probability distribution p_{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 ~ 1 –10% of their mean value (see the Supplementary Information [30]). Figure 5 also shows analytical estimates of infidelity for the full and direct simulation methods, derived in the Supplementary Information.

An immediate takeaway from Figure 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 Figure 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 circuit sizes. It is worth noting that the original circuit cutting work [16] proved that a reconstructed circuit output (probability distribution) can be estimated to an accuracy of ϵ 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 $\mathcal{I} \sim \sum_{f=1}^F 2^{C_o^f}$, where $C_o^f \approx Q/F$ is the number of classical outputs on fragment f and $\sum_f C_o^f = Q$. 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 [30], 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 Figure 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 random unitary circuits, 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 [31]. Having framed fragment characterization as a tomography task, it would also be interesting to adapt and apply different quantum process tomography techniques [32] 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/QSPLIT-MLFT \[33\]](#).

CODE AVAILABILITY

The codes used for numerical experiments in this work are available in the online repository [Quantum-Software-Tools/QSPLIT-MLFT \[33\]](#).

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.

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.

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. Ma. P.* **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 <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 <https://arxiv.org/abs/2010.06660> (2020).
- [21] Saleem, Z. H., Tomesh, T., Gokhale, P. & Suchara, M. Scaling up Constrained Quantum Approximate Optimization Preprint at <https://martinsuchara.com/ScalingConstrainedQA0.pdf> (2021).
- [22] Peruzzo, A. *et al.* A variational eigenvalue solver on a photonic quantum processor. *Nat. Comm.* **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 <https://arxiv.org/abs/1608.03355> (2017).
- [26] Jamiołkowski, 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] See Supplementary Information.
- [31] Ayral, T., Régent, F.-M. L., Saleem, Z., Alexeev, Y. & Suchara, M. Quantum Divide and Compute: Hardware Demonstrations and Noisy Simulations. In *2020 IEEE Computer Society Annual Symposium on VLSI (ISVLSI)*, 138–140 (2020).
- [32] Torlai, G. *et al.* Quantum process tomography with unsupervised learning and tensor networks. Preprint at <https://arxiv.org/abs/2006.02424> (2020).
- [33] QSPLIT-MLFT. <https://github.com/Quantum-Software-Tools/QSPLIT-MLFT> (2020).
- [34] Życzkowski, K. & Kus, M. Random unitary matrices. *J. Phys. A. - Math. Gen.* **27**, 4235–4245 (1994).
- [35] Emerson, J., Livine, E. & Lloyd, S. Convergence conditions for random quantum circuits. *Phys. Rev. A* **72**, 060302 (2005).

FIGURE LEGENDS

Figure 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 (pre-cut) circuit.

Figure 2

Block-diagonalizing circuit fragments. 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) 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.

Figure 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 $k\ell mn$ (i.e. a concatenation of $k, \ell, m, n \in \mathbb{Z}_2$) on the output of the fragment is given by the contraction of the diagonal blocks $\tilde{A}_k, \tilde{B}_{\ell m}, \tilde{C}_n$. The lack of classical inputs to fragment B implies that $\tilde{B} = B$.

Figure 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 [34, 35]. 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 inter-cluster entangling gate.

Figure 5


Infidelity in reconstructed circuit outputs. The infidelity $\mathcal{I} = 1 - \mathcal{F}$ as a function of sample number S (**a**, **b**, **c**) or qubit number Q (**d**, **e**, **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 $\sum_{f=1}^F 2^{C_o^f}/n$ for the direct method, where C_o^f 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 [30]. Results for each data point are averaged over 100 instances of a clustered RUC.







$$= \frac{1}{2} \sum_{\substack{M \in \mathcal{B} \\ r, s \in \lambda(M)}} r s$$




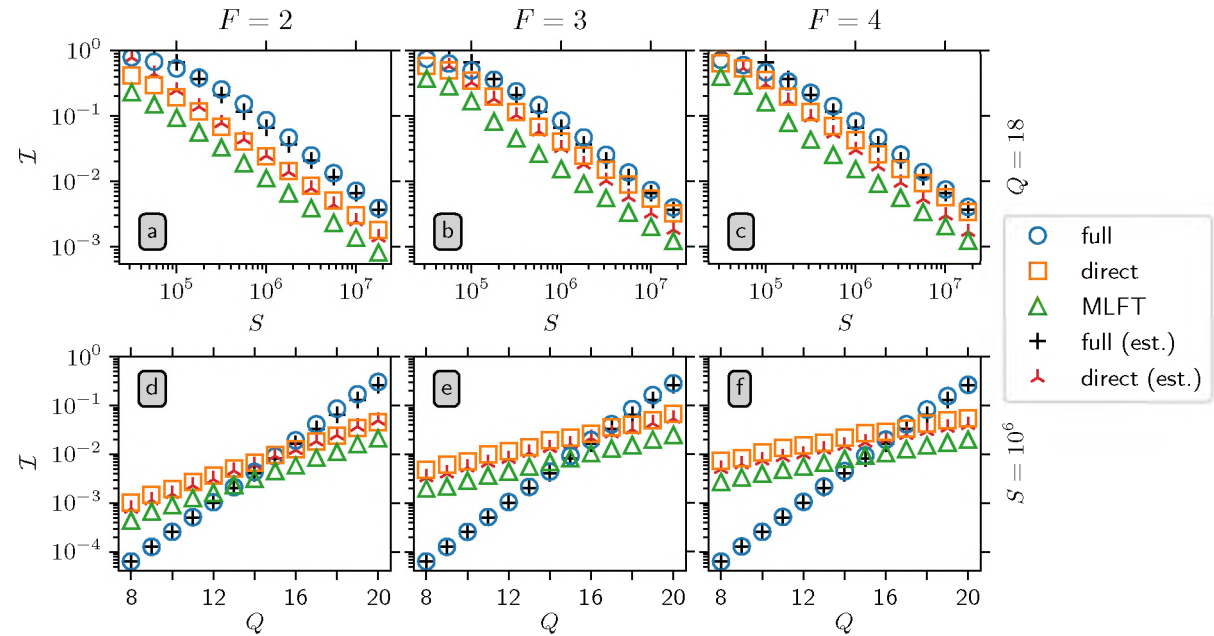
$$\approx \frac{1}{2} \sum_{\substack{M \in \mathcal{B} \\ r, s \in \lambda(M)}} r s$$















Supplementary Information for ‘Quantum Circuit Cutting 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 Avenue, Lemont, Illinois 60439*

³*Mathematics and Computer Science Division, Argonne National
Laboratory, 9700 S. Cass Avenue, Lemont, Illinois 60439*

⁴*Computational Science Division, Argonne National
Laboratory, 9700 S. Cass Avenue, Lemont, Illinois 60439*

(Dated: 2021-01-31)

SUPPLEMENTARY METHODS

In this document, we derive (i) the expected infidelity of a circuit output estimated using the ‘full’ method in the main text (i.e. full circuit execution and sampling), and (ii) an asymptotic bound on the expected infidelity of a circuit output estimated using the ‘direct’ method in the main text (i.e. with circuit cutting, but without maximum-likelihood corrections to fragment models). In practice, when comparing with numerical experiments we find that our asymptotic bound for the ‘direct’ method is overly pessimistic in its scaling with the total number of cuts K in a fragmented circuit. Nonetheless, this upper bound provides a scaling with fragment size that agrees with numerical results.

* mika.perlin@gmail.com

† zsaleem@anl.gov

‡ msuchara@anl.gov

§ osborn@alcf.anl.gov

Multinomial distribution sampling error

Let p be a classical probability distribution over a discrete (and finite) set of measurement outcomes $\{b\}$, and let p_b be the probability of outcome b . We denote an empirical estimate of p_b by \tilde{p}_b , and denote the error in this estimate by $\epsilon_b = \tilde{p}_b - p_b$. If we take n samples of p and set each \tilde{p}_b to the fraction of times that we observed outcome b , then the statistical means and covariances of the errors ϵ_b are

$$\langle\langle \epsilon_b \rangle\rangle = 0, \quad \langle\langle \epsilon_b \epsilon_c \rangle\rangle = \frac{p_b (\delta_{bc} - p_c)}{n}, \quad (1)$$

where $\langle\langle X \rangle\rangle$ denotes the expected value of X after averaging over attempts to estimate p from n samples; and $\delta_{bc} = 1$ if $b = c$ and zero otherwise.

Sampling infidelity

Let p_b be the probability of observing bitstring b at the end of a circuit, and $\tilde{p}_b = p_b + \epsilon_b$ an empirical estimate of p_b . The infidelity of the estimated probability distribution is

$$\mathcal{I} = 1 - \left[\sum_b \sqrt{p_b \tilde{p}_b} \right]^2 = 1 - \left[\sum_b p_b \sqrt{1 + \frac{\epsilon_b}{p_b}} \right]^2, \quad (2)$$

where we tentatively assume that all $p_b \neq 0$. Expanding the square root as $\sqrt{1+x} \approx 1 + x/2 - x^2/8 + O(x^3)$, up to $O(\epsilon^3)$ corrections we find that

$$\mathcal{I} \approx 1 - \left[\sum_b p_b \left(1 + \frac{1}{2} \frac{\epsilon_b}{p_b} - \frac{1}{8} \frac{\epsilon_b^2}{p_b^2} \right) \right]^2 \quad (3)$$

$$= 1 - \sum_{b,c} p_b p_c \left(1 + \frac{1}{2} \frac{\epsilon_b}{p_b} - \frac{1}{8} \frac{\epsilon_b^2}{p_b^2} \right) \left(1 + \frac{1}{2} \frac{\epsilon_c}{p_c} - \frac{1}{8} \frac{\epsilon_c^2}{p_c^2} \right) \quad (4)$$

$$\approx 1 - \sum_{b,c} p_b p_c \left[1 + \frac{1}{2} \left(\frac{\epsilon_b}{p_b} + \frac{\epsilon_c}{p_c} \right) + \frac{1}{4} \frac{\epsilon_b \epsilon_c}{p_b p_c} - \frac{1}{8} \left(\frac{\epsilon_b^2}{p_b^2} + \frac{\epsilon_c^2}{p_c^2} \right) \right] \quad (5)$$

$$= - \sum_b \epsilon_b - \frac{1}{4} \sum_{b,c} \epsilon_b \epsilon_c + \frac{1}{4} \sum_b \frac{\epsilon_b^2}{p_b}. \quad (6)$$

Infidelity for full circuit execution

Let p_b be the probability of observing bitstring b at the end of the circuit, and $\tilde{p}_b = p_b + \epsilon_b$ an empirical estimate of p_b . If we sample the probability distribution n times and set each

\tilde{p}_b to the fraction of times that we observed bitstring s , then the estimates \tilde{p}_b of p_b are normalized with $\sum_b \tilde{p}_b = 1$, so

$$\sum_b \epsilon_b = \sum_b (\tilde{p}_b - p_b) = \sum_b \tilde{p}_b - \sum_b p_b = 0. \quad (7)$$

Up to $O(\epsilon^3)$ corrections, the expected infidelity $\langle\langle \mathcal{I} \rangle\rangle$ is then

$$\langle\langle \mathcal{I} \rangle\rangle \approx \frac{1}{4} \sum_b \frac{\langle\langle \epsilon_b^2 \rangle\rangle}{p_b} = \frac{1}{4} \sum_b \frac{1 - p_b}{n} = \frac{2^Q - 1}{4n} \approx \frac{2^Q}{4n}, \quad (8)$$

where $2^Q = \sum_b 1$ is the total number of bitstrings that can be measured at the output of Q qubits.

A few comments concerning the result in Eq. (8) are in order. First, the restriction that $\mathcal{I} \in [0, 1]$ implies that Eq. (8) can only hold for $4n > 2^Q - 1$, and that if $4n$ is comparable to 2^Q then the $O(\epsilon^3)$ contributions to \mathcal{I} must become relevant. Second, if any $p_b = 0$, then the factor 2^Q should be replaced by the sample space volume $|\{b : p_b \neq 0\}|$. This second observation in particular suggests that Eq. (8) can only hold for sufficiently ‘generic’ probability distributions, as large separations of scale in the probabilities p_b should reduce 2^Q to some smaller ‘effective’ sample space volume, likely determined by the output entropy $S(p) = -\sum_b p_b \log p_b$. Finally, we point out that Eq. (8) also describes the infidelity with which n samples estimate the conditional probability distributions associated with a single fragment of a cut-up circuit. Unfortunately, the presence of quantum correlations between circuit fragments implies that the infidelity of a reconstructed circuit output is not additive in the infidelities of the fragments. Nonetheless, we show in the following section that the infidelity of a reconstructed circuit still scales inversely with the number of fragment samples, and exponentially in fragment size, i.e.

$$\langle\langle \mathcal{I} \rangle\rangle \sim O\left(\frac{1}{n} \sum_f 2^{C_o^f}\right), \quad (9)$$

where now n is the number of samples used to estimate each variant of each fragment, f indexes a single fragment, and C_o^f is the number of classical output bits on fragment f .

One cut, two fragments

We now attempt to compute the expected infidelity of a circuit output estimated from fragment data. We first work out the relatively simple case of one cut and two fragments, which we will subsequently generalize to the case of arbitrary cuts and fragments.

In order to evaluate the expected infidelity $\langle\langle\mathcal{I}\rangle\rangle$, we first need to construct an estimator \tilde{p} of p . We can take the diagonal elements of Eqs. (5) and (7) in the main text (in the computational basis) to expand the probability distribution

$$p = \frac{1}{2} \sum_{M \in \mathcal{B}} p^1(M) \otimes p^2(M) \quad p^f(M) = \sum_{r \in \lambda(M)} r p^f(M_r), \quad (10)$$

where $p^f(M_r)$ is the probability distribution at the output of fragment f conditional on measuring or initializing M_r , as appropriate. We will assume without loss of generality that the first fragment is measured and that the second fragment is initialized at the cut, in which case $p^1(M_r)$ is normalized to the probability of getting outcome M_r when measuring the appropriate qubit in the diagonal basis of M . To deal with this normalization properly, we expand

$$p^f(M_r) = a^f(M_r) q^f(M_r), \quad (11)$$

where $a^f(M_r) = \sum_b p_b^f(M_r)$ is the normalization of $p^f(M_r)$ (which is equal to 1 for initialization conditions), and $q^f(M_r) = p^f(M_r) / a^f(M_r)$ is a normalized probability distribution.

We construct estimators $\tilde{p}^f(M_r)$ of $p^f(M_r)$ as follows. For measurement conditions, the estimator $\tilde{a}^f(M_r) \approx a^f(M_r)$ is set to the fraction of times that we observe outcome M_r when measuring in the appropriate basis; for initialization conditions, $\tilde{a}^f(M_r) = a^f(M_r) = 1$. Each entry $\tilde{q}_b^f(M_r)$ of the probability distribution $\tilde{q}^f(M_r) \approx q^f(M_r)$ is set to the fraction of times that we observe bitstring b on the classical output of fragment f when conditioned on M_r . For any argument M_r of a^f , q^f , etc., we can then define the errors $\beta^f = \tilde{a}^f - a^f$ and $\gamma^f = \tilde{q}^f - q^f$ and expand

$$\tilde{p}^f = \tilde{a}^f \tilde{q}^f = p^f + \epsilon^f, \quad \epsilon^f = \beta^f q^f + a^f \gamma^f + \beta^f \gamma^f, \quad (12)$$

and define, for all symbols $X \in \{p^f, a^f, q^f, \tilde{p}^f, \tilde{a}^f, \tilde{q}^f, \epsilon^f, \beta^f, \gamma^f\}$,

$$X(M) = \sum_{s \in \lambda(M)} s X(M_s), \quad (13)$$

which allows us to construct the estimator

$$\tilde{p} = \frac{1}{2} \sum_{M \in \mathcal{B}} \tilde{p}^1(M) \otimes \tilde{p}^2(M) = p + \epsilon, \quad (14)$$

where

$$\epsilon = \frac{1}{2} \sum_{M \in \mathcal{B}} [\epsilon^1(M) \otimes p^2(M) + p^1(M) \otimes \epsilon^2(M) + \epsilon^1(M) \otimes \epsilon^2(M)]. \quad (15)$$

Strictly speaking, the definition of estimators and errors $X(I)$ are ambiguous as presented, as there is not a unique decomposition of I for the sum in Eq. (13). In practice, our implementations of circuit cutting algorithms set

$$X(I) = \frac{1}{3} \sum_{\substack{M \in \{X,Y,Z\} \\ r \in \lambda(M)}} X(M_r). \quad (16)$$

Whereas the errors $\epsilon^f(M)$ and $\epsilon^f(M')$ for $M, M' \in \{X, Y, Z\}$ are uncorrelated unless $M = M'$, the decomposition in Eq. (16) implies that $\epsilon^f(I)$ is correlated with $\epsilon^f(M)$ for all $M \in \mathcal{B}$. For simplicity, however, we will assume that the estimators $\tilde{p}^f(M)$ and $\tilde{p}^f(M')$ with $M \neq M'$ are built from independent experimental data, such that their corresponding errors $\epsilon^f(M)$ and $\epsilon^f(M')$ are uncorrelated. Crucially, this assumption does not affect the general structure of our calculations, and therefore leaves our main conclusions (namely, how reconstruction infidelity $\langle\langle \mathcal{I} \rangle\rangle$ scales with different fragment parameters) in tact.

In order to compute the expected infidelity of \tilde{p} , we now need to determine the statistical means $\langle\langle \epsilon_b \rangle\rangle$, covariances $\langle\langle \epsilon_b \epsilon_c \rangle\rangle$, and variances $\langle\langle \epsilon_b^2 \rangle\rangle$. The means $\langle\langle \epsilon_b \rangle\rangle = 0$ because all contributions to ϵ_b are either (i) proportional to a single error in the estimate of a multinomially distributed random variable, which is mean-zero, or (ii) a product of multiple independent (uncorrelated) errors, which is also mean-zero.

Covariances

To compute the covariance $\langle\langle \epsilon_b \epsilon_c \rangle\rangle$, we note that are only working to second order in error variables, whereas the contributions to $\langle\langle \epsilon_b \epsilon_c \rangle\rangle$ from the $\sim \epsilon^1 \otimes \epsilon^2$ terms of Eq. (15) are either fourth order or zero; we are therefore free to neglect these terms. Additionally throwing out terms that vanish because they are the product of uncorrelated random variables, we find that

$$\langle\langle \epsilon_b \epsilon_c \rangle\rangle \approx \frac{1}{4} \sum_{f, M} \langle\langle \epsilon_{b_f}^f(M) \epsilon_{c_f}^f(M) \rangle\rangle p_{b_f}^{\bar{f}}(M) p_{c_f}^{\bar{f}}(M), \quad (17)$$

where $f \in \{1, 2\}$ and $\bar{f} \neq f$, i.e. such that $\{f, \bar{f}\} = \{1, 2\}$; and b_f, c_f are the substrings of b, c associated with fragment f . We now expand

$$\begin{aligned} \langle\langle \epsilon_b^f(M) \epsilon_c^f(M) \rangle\rangle &\approx \sum_{r,s \in \lambda(M)} rs \langle\langle \beta^f(M_r) \beta^f(M_s) \rangle\rangle q_b^f(M_r) q_c^f(M_s) \\ &+ \sum_{r,s \in \lambda(M)} rs a^f(M_r) a^f(M_s) \langle\langle \gamma_b^f(M_r) \gamma_c^f(M_s) \rangle\rangle, \end{aligned} \quad (18)$$

where we again throw out terms that are zero or fourth order in error variables. The normalization errors $\beta^f(M_r), \beta^f(M_s)$ are zero for initialization conditions, and are always correlated for measurement conditions because they are errors in mutually exclusive measurement outcomes. The probability distribution errors $\gamma_b^f(M_r), \gamma_c^f(M_s)$, meanwhile, are independent unless $M_r = M_s$. The covariances between these errors are determined by multinomial distribution sampling errors, so

$$\langle\langle \beta^f(M_r) \beta^f(M_s) \rangle\rangle = \frac{1}{n} a^f(M_r) [\delta_{rs} - a^f(M_s)] (1 - \delta_{Q_o^f,0}), \quad (19)$$

$$\langle\langle \gamma_b^f(M_r) \gamma_c^f(M_s) \rangle\rangle = \frac{\delta_{rs}}{a^f(M_r) n} q_b^f(M_r) [\delta_{bc} - q_c^f(M_r)], \quad (20)$$

where n is the number of times that we sample fragment each variant of each fragment (i.e. each choice of initialization conditions and measurement bases on a fragment), $a^f(M_r) n$ is the expected number of times that we sample fragment f with condition M_r , Q_o^f is the number of quantum outputs (i.e. or measurement conditions) on fragment f .

If $\delta_{Q_o^f,0} = 0$, then the $\sim a^f(M_r) \delta_{rs}$ contributions from Eq. (19) cancel out with the $\sim q_b^f(M_r) q_c^f(M_r)$ contributions from Eq. (20) when substituting these results into Eq. (18). Meanwhile, if $\delta_{Q_o^f,0} = 1$ then $a^f(M_r) = 1$, so

$$n \langle\langle \epsilon_b^f(M) \epsilon_c^f(M) \rangle\rangle \approx \delta_{bc} p_b^f(I) - \left(1 - \delta_{Q_o^f,0}\right) p_b^f(M) p_c^f(M) - \delta_{Q_o^f,0} \sum_r p_b^f(M_r) p_c^f(M_r). \quad (21)$$

Altogether,

$$\sum_{b,c} \langle\langle \epsilon_b \epsilon_c \rangle\rangle \approx \frac{1}{4n} \sum_{f,M} \left[\sum_{b_f, c_f} \langle\langle \epsilon_{b_f}^f(M) \epsilon_{c_f}^f(M) \rangle\rangle \right] \left[\sum_{b_{\bar{f}}, c_{\bar{f}}} p_{b_{\bar{f}}}^{\bar{f}}(M) p_{c_{\bar{f}}}^{\bar{f}}(M) \right] \quad (22)$$

$$\approx \frac{1}{4n} \sum_{f,M} \left[a^f(I) - \left(1 - \delta_{Q_o^f,0}\right) a^f(M)^2 - 2\delta_{Q_o^f,0} a^{\bar{f}}(M)^2 \right]. \quad (23)$$

Substituting

$$a^1(I) = 1, \quad a^2(M) = 2\delta_{M,I}, \quad \delta_{Q_o^f,0} = \delta_{f,2}, \quad 1 - \delta_{Q_o^f,0} = \delta_{f,1}, \quad (24)$$

we thus find that

$$\sum_{b,c} \langle\langle \epsilon_b \epsilon_c \rangle\rangle = 0. \quad (25)$$

The infidelity $\langle\langle \mathcal{I} \rangle\rangle$ is therefore determined entirely by the variances $\langle\langle \epsilon_b^2 \rangle\rangle$.

Variances

We now consider the variances

$$\langle\langle \epsilon_b^2 \rangle\rangle \approx \frac{1}{4} \sum_{f,M} \langle\langle \epsilon_{b_f}^f(M)^2 \rangle\rangle p_{b_f}^{\bar{f}}(M)^2 \quad (26)$$

where from Eq. (21) we know that

$$n \langle\langle \epsilon_b^f(M)^2 \rangle\rangle = p_b^f(I) - p_b^f(M)^2 \left(1 - \delta_{Q_o^f,0}\right) - \delta_{Q_o^f,0} \sum_r q_b^f(M_r)^2. \quad (27)$$

In principle, we have to simplify

$$\langle\langle \mathcal{I} \rangle\rangle \approx \frac{1}{4} \sum_b \frac{\langle\langle \epsilon_b^2 \rangle\rangle}{p_b}, \quad p_b = \frac{1}{2} \sum_M p_{b_1}^1(M) p_{b_2}^2(M), \quad (28)$$

but this calculation is intractable due to the sum over $M \in \mathcal{B}$ in the denominator. We therefore settle for trying to find an upper bound on this sum, to which end we observe that

$$n \langle\langle \epsilon_b^2 \rangle\rangle \lesssim \frac{1}{4} \sum_{f,M} p_{b_f}^f(I) p_{b_f}^{\bar{f}}(M)^2 < \sum_f p_{b_f}^f(I) p_{b_f}^{\bar{f}}(I)^2. \quad (29)$$

Rather than the factors $p_{b_f}^f(I)$, we can express this bound in terms of the mean probability $\langle p_{b_f}^f \rangle$ to get bitstring b_f on fragment f , averaged over all possible conditions. This mean probability is $\langle p_{b_f}^f \rangle = p_{b_f}^f(I) / 2^{Q_i^f}$, where Q_i^f is the number of quantum inputs (initialization conditions) on fragment f , so

$$n \langle\langle \epsilon_b^2 \rangle\rangle \lesssim \sum_f 2^{Q_i^f} \langle p_{b_f}^f \rangle 4^{Q_i^{\bar{f}}} \langle p_{b_f}^{\bar{f}} \rangle^2 = 4 \langle p_{b_1}^1 \rangle \langle p_{b_2}^2 \rangle \sum_f \frac{1}{2^{Q_i^f}} \langle p_{b_f}^{\bar{f}} \rangle. \quad (30)$$

In turn, we can bound

$$\langle\langle \mathcal{I} \rangle\rangle \approx \frac{1}{4} \sum_b \frac{\langle\langle \epsilon_b^2 \rangle\rangle}{p_b} \lesssim \frac{1}{n} \sum_b \left[\frac{\langle p_{b_1}^1 \rangle \langle p_{b_2}^2 \rangle}{p_b} \right] \sum_f \frac{1}{2^{Q_i^f}} \langle p_{b_f}^{\bar{f}} \rangle. \quad (31)$$

The factor $\langle p_{b_1}^1 \rangle \langle p_{b_2}^2 \rangle / p_b$ is a measure of the quantum correlation between two fragments: it is equal to 1 if there is no correlation, and is smaller (greater) than 1 if quantum correlations cause constructive (destructive) interference for measurement outcome b on the combined circuit. In principle, this factor can be made arbitrarily large, but that requires fine tuning, and generally speaking we would expect $\langle p_{b_1}^1 \rangle \langle p_{b_2}^2 \rangle / p_b \sim O(1)$ for random circuits. We therefore expect that

$$\langle\langle \mathcal{I} \rangle\rangle \sim O \left(\frac{1}{n} \sum_{b,f} \frac{1}{2^{Q_i^f}} \langle p_{b_f}^{\bar{f}} \rangle \right) = O \left(\frac{1}{n} \sum_f 2^{C_o^f - Q_i^f} \right). \quad (32)$$

The general case

We now estimate an upper bound on expected infidelity for the general case of K cuts and F fragments. Defining the projectors

$$\mathcal{P}_{\mathcal{B}} = \bigcup_{M \in \mathcal{B}} \{M_r : r \in \lambda(M)\}, \quad (33)$$

a fragment with K_f incident cuts is nominally characterized by the conditional distributions

$$p^f(\mathbf{M}) = a^f(\mathbf{M}) q^f(\mathbf{M}), \quad (34)$$

where $\mathbf{M} \in \mathcal{P}_{\mathcal{B}}^{K_f}$, $a^f(\mathbf{M}) = \sum_b p_b^f(\mathbf{M})$ is the normalization of $p^f(\mathbf{M})$, and $q^f(\mathbf{M})$ is a normalized probability distribution. Similarly to before, the estimator $\tilde{a}^f(\mathbf{M}) \approx a^f(\mathbf{M})$ is set to the fraction of times that we observe the measurement conditions in \mathbf{M} when measuring the corresponding qubits in the appropriate bases and preparing the initialization conditions in \mathbf{M} . Each entry $\tilde{q}_b^f(\mathbf{M})$ of the probability distribution $\tilde{q}_b^f(\mathbf{M}) \approx q_b^f(\mathbf{M})$ is set to the fraction of times that we observe bitstring b on the classical output of fragment f with conditions \mathbf{M} . The estimator \tilde{p}^f and errors $\beta^f, \gamma^f, \epsilon^f$ are defined just as before, and all appropriate objects are defined for ‘conditions’ $\mathbf{M} \in \mathcal{B}^{K_f}$ by summing over the spectrum of each condition similarly to Eq. (13). Altogether, we have the estimator

$$\tilde{p} = \frac{1}{2^K} \sum_{\mathbf{M} \in \mathcal{B}^K} \bigotimes_f \tilde{p}^f(\mathbf{M}^f) = p + \epsilon, \quad (35)$$

where $\mathbf{M}^f \subset \mathbf{M}$ denotes the conditions in \mathbf{M} that are incident on fragment f , and the distribution error ϵ has components

$$\epsilon_b \approx \frac{1}{2^K} \sum_{\mathbf{M} \in \mathcal{B}^K} \sum_f \epsilon_{b_f}^f(\mathbf{M}^f) \prod_{g \neq f} p_{b_g}^f(\mathbf{M}^g), \quad (36)$$

where we neglect terms that are second order or higher in the fragment errors ϵ^f .

Covariances

By the same argument as before, the mean $\langle\langle \epsilon_b \rangle\rangle = 0$, so we turn to considering the covariance

$$\langle\langle \epsilon_b \epsilon_c \rangle\rangle \approx \frac{1}{4^K} \sum_{\mathbf{M} \in \mathcal{B}^K} \sum_f \langle\langle \epsilon_{b_f}^f(\mathbf{M}^f) \epsilon_{c_f}^f(\mathbf{M}^f) \rangle\rangle \prod_{g \neq f} p_{b_g}^g(\mathbf{M}^g) p_{c_g}^g(\mathbf{M}^g), \quad (37)$$

which leads us to expand

$$\begin{aligned} & \langle\langle \epsilon_b^f(\mathbf{M}^f) \epsilon_c^f(\mathbf{M}^f) \rangle\rangle \\ & \approx \sum_{R, S \in \lambda(\mathbf{M}^f)} \|R\| \|S\| \langle\langle \beta^f(\mathbf{M}_R^f) \beta^f(\mathbf{M}_S^f) \rangle\rangle q_b^f(\mathbf{M}_R^f) q_c^f(\mathbf{M}_S^f) \\ & \quad + \sum_{R, S \in \lambda(\mathbf{M}^f)} \|R\| \|S\| a^f(\mathbf{M}_R^f) a^f(\mathbf{M}_S^f) \langle\langle \gamma_b^f(\mathbf{M}_R^f) \gamma_c^f(\mathbf{M}_S^f) \rangle\rangle, \end{aligned} \quad (38)$$

where $R = (r_1, r_2, \dots, r_{K_f}) \in \lambda(\mathbf{M}^f)$ is a choice of eigenvalue $r \in \lambda(M)$ for each condition M in \mathbf{M}^f , \mathbf{M}_R^f is the corresponding list of (projectors onto) eigenvectors, and $\|R\| = \prod_{r \in R} r$.

The covariances between β^f and γ^f are determined by the multinomial distributions

$$\langle\langle \beta^f(\mathbf{M}_R^f) \beta^f(\mathbf{M}_S^f) \rangle\rangle = \frac{1}{n} \times a^f(\mathbf{M}_R^f) [\delta_{RS} - a^f(\mathbf{M}_S^f)] (1 - \delta_{Q_{\circ}^f, 0}), \quad (39)$$

$$\langle\langle \gamma_b^f(\mathbf{M}_R^f) \gamma_c^f(\mathbf{M}_S^f) \rangle\rangle = \frac{1}{n} \times \frac{\delta_{RS}}{a^f(\mathbf{M}_R)} \times q_b^f(\mathbf{M}_R^f) [\delta_{bc} - q_c^f(\mathbf{M}_R^f)], \quad (40)$$

which as before implies that

$$\begin{aligned} & n \langle\langle \epsilon_b^f(\mathbf{M}^f) \epsilon_c^f(\mathbf{M}^f) \rangle\rangle \\ & \approx \delta_{bc} p_b^f(\mathbf{I}^f) - (1 - \delta_{Q_{\circ}^f, 0}) p_b^f(\mathbf{M}^f) p_c^f(\mathbf{M}^f) - \delta_{Q_{\circ}^f, 0} \sum_R p_b^f(\mathbf{M}_R^f) p_c^f(\mathbf{M}_R^f), \end{aligned} \quad (41)$$

where $\mathbf{I}^f = (I, I, \dots)$ is a constant list of length K_f . Altogether,

$$\sum_{b,c} \langle\langle \epsilon_b \epsilon_b \rangle\rangle \approx \frac{1}{4^K n} \sum_{f, \mathbf{M}} \sum_{b,c} \langle\langle \epsilon_{b_f}^f(\mathbf{M}^f) \epsilon_{c_f}^f(\mathbf{M}^f) \rangle\rangle \prod_{g \neq f} p_{b_g}^g(\mathbf{M}^g) p_{c_g}^g(\mathbf{M}^g) \quad (42)$$

$$\approx \frac{1}{4^K n} \sum_{f, \mathbf{M}} \left[a^f(\mathbf{I}^f) - (1 - \delta_{Q_o^f, 0}) a^f(\mathbf{M}^f)^2 - 2^{K_f} \delta_{Q_o^f, 0} \right] \prod_{g \neq f} a^g(\mathbf{M}^g)^2. \quad (43)$$

Though unsure how to evaluate this quantity, we can use the fact that $a^f(\mathbf{M}^f) \leq a^f(\mathbf{I}^f)^2 \leq 4^{K_f}$ to bound

$$\frac{1}{4^K} \sum_{f, \mathbf{M}} a^f(\mathbf{M}^f)^2 \prod_{g \neq f} a^g(\mathbf{M}^g)^2 < F \prod_g a^g(\mathbf{I}^g)^2 < 4^K F, \quad (44)$$

and similarly

$$\frac{1}{4^K} \sum_{f, \mathbf{M}} 2^{K_f} \prod_{g \neq f} a^g(\mathbf{M}^g)^2 \leq \sum_f 2^{K_f} \prod_{g \neq f} a^g(\mathbf{I}^g)^2 \leq \sum_f 2^{K_f} \prod_{g \neq f} 4^{K_g} < 4^K F, \quad (45)$$

which implies that the contribution to $\langle\langle \mathcal{I} \rangle\rangle$ from the covariances $\langle\langle \epsilon_b \epsilon_b \rangle\rangle$ satisfies

$$-\frac{1}{4} \sum_{b,c} \langle\langle \epsilon_b \epsilon_b \rangle\rangle < \frac{4^K F}{4n}. \quad (46)$$

Variances

As in the case of one cut and two fragments, we now bound

$$n \langle\langle \epsilon_b^2 \rangle\rangle \lesssim \frac{1}{4^K} \sum_{f, \mathbf{M}} p_{b_f}^f(\mathbf{I}^f) \prod_{g \neq f} p_{b_g}^g(\mathbf{M}^g)^2 \quad (47)$$

$$< \sum_f p_{b_f}^f(\mathbf{I}^f) \prod_{g \neq f} p_{b_g}^g(\mathbf{I}^g)^2 \quad (48)$$

$$= \left[\prod_h p_{b_h}^h(\mathbf{I}^h) \right] \sum_f \prod_{g \neq f} p_{b_g}^g(\mathbf{I}^g) \quad (49)$$

$$= \left[\prod_h 2^{Q_i^h} \langle p_{b_h}^h \rangle \right] \sum_f \prod_{g \neq f} 2^{Q_i^g} \langle p_{b_g}^g \rangle \quad (50)$$

$$= 4^K \left[\prod_h \langle p_{b_h}^h \rangle \right] \sum_f 2^{-Q_i^f} \prod_{g \neq f} \langle p_{b_g}^g \rangle. \quad (51)$$

The contribution to infidelity $\langle\langle \mathcal{I} \rangle\rangle$ from the variances $\langle\langle \epsilon_b^2 \rangle\rangle$ is then

$$\frac{1}{4} \sum_b \frac{\langle\langle \epsilon_b^2 \rangle\rangle}{p_b} \lesssim \frac{4^K}{4n} \sum_b \left[\frac{\prod_h \langle p_{b_h}^h \rangle}{p_b} \right] \sum_f 2^{-Q_i^f} \prod_{g \neq f} \langle p_{b_g}^g \rangle \sim O\left(\frac{4^K}{n} \sum_f 2^{C_o^f - Q_i^f}\right), \quad (52)$$

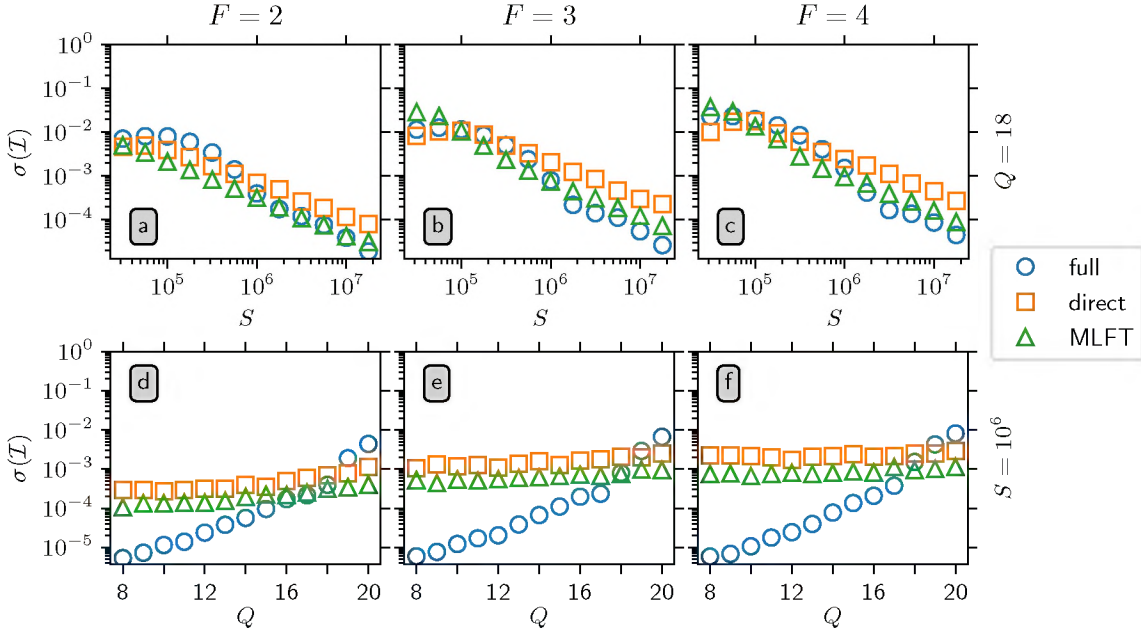
so altogether

$$\langle\langle \mathcal{I} \rangle\rangle \sim O\left(\frac{4^K}{n} \sum_f 2^{C_o^f - Q_i^f}\right). \quad (53)$$

In practice, we find this asymptotic bound to be overly pessimistic with regards to the scaling with K . There are other ways in which this bound is too *optimistic*: by assuming that fragments are weakly correlated and $\frac{1}{p_b} \prod_h \langle p_{b_h}^h \rangle \sim O(1)$, this bound does not capture the effect of errors due to noisy virtual teleportation of qubits across cuts between fragments, at their quantum inputs and outputs. Nonetheless, the bound in Eq. (53) demonstrates $\sim \frac{1}{n} \sum_f 2^{C_o^f}$ scaling with shot number n and fragment size C_o^f , which (all else equal) are not affected by these considerations.

SUPPLEMENTARY FIGURES

Figure 5 of the main text shows infidelities $\langle\langle \mathcal{I} \rangle\rangle$ of reconstructed outputs for clustered random unitary circuits (RUCs). To provide a sense of the robustness of \mathcal{I} to circuit variations, Supplementary Figure 1 shows the standard deviation $\sigma(\mathcal{I}) = \sqrt{\langle\langle (\mathcal{I} - \langle\langle \mathcal{I} \rangle\rangle)^2 \rangle\rangle}$ in the same simulations. Generally speaking, $\sigma(\mathcal{I})$ is orders of magnitude smaller than the $\langle\langle \mathcal{I} \rangle\rangle$.



Supplementary Figure 1. Standard deviation of the infidelities shown in Figure 5.