

Qubit-efficient entanglement spectroscopy using qubit resets

Justin Yirka¹ and Yiğit Subaşı²

¹Department of Computer Science, The University of Texas at Austin, Austin, TX 78712, USA

²Computer, Computational, and Statistical Sciences Division, Los Alamos National Laboratory, Los Alamos, NM 87545, USA

One strategy to fit larger problems on NISQ devices is to exploit a tradeoff between circuit width and circuit depth. Unfortunately, this tradeoff still limits the size of tractable problems since the increased depth is often not realizable before noise dominates. Here, we develop *qubit-efficient* quantum algorithms for entanglement spectroscopy which avoid this tradeoff. In particular, we develop algorithms for computing the trace of the n -th power of the density operator of a quantum system, $\text{Tr}(\rho^n)$, (related to the Rényi entropy of order n) that use fewer qubits than any previous efficient algorithm while achieving similar performance in the presence of noise, thus enabling spectroscopy of larger quantum systems on NISQ devices. Our algorithms, which require a number of qubits independent of n , are variants of previous algorithms with width proportional to n , an asymptotic difference. The crucial ingredient in these new algorithms is the ability to measure and reinitialize subsets of qubits in the course of the computation, allowing us to reuse qubits and increase the circuit depth without suffering the usual noisy consequences. We also introduce the notion of *effective circuit depth* as a generalization of standard circuit depth suitable for circuits with qubit resets. This tool helps explain the noise-resilience of our qubit-efficient algorithms and should aid in designing future algorithms. We perform numerical simulations to compare our algorithms to the original variants and show they perform similarly when subjected to noise. Additionally, we experimentally implement one of our qubit-efficient algorithms on the Honeywell System Model H0, estimating $\text{Tr}(\rho^n)$ for larger n than possible with previous algorithms.

Justin Yirka: yirka@utexas.edu

Yiğit Subaşı: ysubasi@lanl.gov

1 Introduction

Full-scale fault-tolerant quantum computers offer eventual advantages over classical computation for a variety of tasks. While work continues toward such devices, more research is needed on how to utilize near-term devices. As we develop applications for noisy intermediate-scale quantum (NISQ) computers [14, 42], a primary limitation is the inverse relationship between the quality and the quantity of available qubits, i.e. larger devices tend to be noisier. One way to mitigate effects of noise is to design algorithms with low circuit depths, but this is often challenging. While some approaches help for specific applications [5, 19, 37, 40] or for individual circuits [39], there are few general techniques for (re)designing low-depth quantum algorithms. One technique is to trade shorter circuit depth for increased circuit width [1, 9, 50], i.e. use more qubits, but those qubits may still be unavailable or unacceptably noisy. New strategies are needed for designing low-width noise-resilient, NISQ algorithms.

One under-explored tool is qubit resetting, by which we mean the ability to reinitialize subsets of qubits in a known state, usually the $|0\rangle$ state, in the course of the computation [17]. Generally, qubits are reinitialized either to prepare the entire apparatus to run a circuit or to reuse subsets of qubits in the course of a computation; we focus on the second use. There exist methods for actively resetting qubits to their $|0\rangle$ state in time comparable to that required for a measurement [24, 36, 38, 45] — with measurement and reset generally distinct processes. The implementation depends on the particular hardware, but we are interested in resets as an algorithmic/software tool. This ability will be critical for error-correcting codes which require frequent stabilizing measurements [46], and it has recently been used to design algorithms with reduced circuit width [21, 27, 34, 44]. Work on automatically inserting resets as part of optimization and compiling should be particularly valuable for the first goal. Here, we contribute to the latter goal. We present algorithms for the application of entanglement spectroscopy which exploit qubit resets to achieve low circuit width while remaining noise-resilient.

Entanglement spectroscopy is the task of learning about the entanglement of a quantum state. The bi-

partite entanglement of a pure quantum state $|\psi\rangle$ on systems A and B can be characterized by the eigenvalues of the density operator of the reduced state $\rho_A = \text{Tr}_B(|\psi\rangle\langle\psi|)$ (equivalent to the eigenvalues of ρ_B) [3]. As noted by Li and Haldane, the entanglement spectrum (the eigenvalues of the so-called entanglement Hamiltonian H defined via $\rho_A = e^{-H}$) contains much more information than the von Neumann entropy alone [31]. For instance, it can be used to detect and characterize topological order and quantum phase transitions, as well as to determine whether a system obeys an area law and thus can be efficiently simulated classically [12, 15, 16, 20, 26, 29, 31, 43, 52]. Thus, Entanglement spectroscopy is an especially useful tool for analyzing outcomes of quantum simulation of many-body systems [28, 30, 33]. It may be similarly useful in characterizing the performance of NISQ devices. Moreover, learning just the few largest eigenvalues of ρ_A , rather than performing full tomography, is often sufficient [29]. This task is computationally hard classically due to the exponentially growing dimension of the Hilbert space, making it a clear candidate for quantum algorithms.

Known efficient quantum algorithms to approximate the n_{\max} largest eigenvalues of ρ_A generally begin by reducing the problem to computing the traces of powers of the reduced density operator, i.e. $\text{Tr}(\rho_A^n)$ for $n = 1, \dots, n_{\max}$ [29, 49].¹ These algorithms compute $\text{Tr}(\rho_A^n)$ using $O(n)$ copies of the state $|\psi\rangle$ [29, 50]. The standard algorithm is an extension of the Swap Test [11, 25] by [29] which we call the Entanglement Spectroscopy Hadamard Test (HT) (Fig. 2). It uses n copies of the state and has depth linear in n and the size of the state (number of qubits). The recent Entanglement Spectroscopy Two-Copy Test (TCT) (Fig. 3) by [50] uses $2n$ copies of the state and achieves constant depth. Although the latter algorithm achieves a depth suitable for NISQ devices, the linear width in n of both algorithms will likely restrict their application to small n in the NISQ era.

Having reduced entanglement spectroscopy to computing $\text{Tr}(\rho_A^n)$, we may state the task we study in this paper formally. **Problem:** Given as input a parameter n and black-box access to a circuit preparing a pure state $|\psi\rangle$ on subsystems A, B , estimate $\text{Tr}(\rho_A^n)$.

In this work, we introduce new *qubit-efficient* variants of the HT and TCT algorithms that require a number of qubits sufficient to prepare three or fewer copies of $|\psi\rangle$, independent of n . This is an asymptotically lower width than any previous efficient algorithm for com-

¹An exception to this is qPCA [35], where the spectrum of a state ρ is obtained using a different approach requiring phase estimation. This approach is not NISQ-friendly, so we do not discuss it here.

	Width	Depth	Effective depth	Ref.
HT	$kn + 1$	$T_{\text{sp}} + O(kn)$	$O(kn)$	[29]
$4k$ qe-HT	$4k + 1$	$\Theta(n \times (T_{\text{sp}} + k))$	$\Theta(n \times (T_{\text{sp}} + k))$	*
$3k$ qe-HT	$3k + 1$	$\Theta(n \times (T_{\text{sp}} + k))$	$\Theta(n \times (T_{\text{sp}} + k))$	*
TCT	$4kn$	$T_{\text{sp}} + O(1)$	$T_{\text{sp}} + O(1)$	[50]
$6k$ qe-TCT	$6k$	$\Theta(n \times (T_{\text{sp}} + 1))$	$T_{\text{sp}} + O(1)$	*
$4k$ qe-TCT	$4k$	$\Theta(n \times (T_{\text{sp}} + 1))$	$T_{\text{sp}} + O(1)$	*

Table 1: A summary of the algorithms in this paper for computing $\text{Tr}(\rho_A^n)$ of a state $|\psi\rangle_{AB}$ on $2k$ qubits. Algorithms marked with ‘*’ in the final column are new. See Sections 2 and 3 for more details.

puting $\text{Tr}(\rho_A^n)$ or the Rényi entropy of order n .² We achieve this by using qubit resets and preparing additional copies of the state in previously used registers, allowing us to perform computations on many copies of the state while using few qubits.

The depths of our qubit-efficient algorithms are linear in n and the size of the state, but, crucially, our new algorithms do not suffer as much in the presence of noise as their increased depth suggests. Intuitively, one hopes that periodically resetting qubits prevents errors from accumulating, but because the resets only affect a subset of the qubits at a time, errors might still carry over. By carefully choreographing the resets in our new algorithms, we try to prevent this from happening as much as possible.

We test our algorithms numerically and find that our qubit-efficient algorithms perform nearly identically well in the presence of noise as their higher-width analogs. We also implement one of our qubit-efficient TCT algorithms experimentally on the Honeywell System Model H0 [41], estimating $\text{Tr}(\rho^n)$ for larger n than possible on the device using previous algorithms.

Motivated by our results, we propose a generalization of circuit depth, which we call *effective circuit depth*, for predicting the performance of quantum algorithms that use qubit resets on noisy devices. This new attribute helps explain why our qubit-efficient algorithms perform comparably to their original counterparts; for example, while the depths of our qubit-efficient TCT variants are asymptotically greater than that of the original TCT, their effective depths match up to a constant factor. Effective circuit depth is a better descriptor of quantum circuits with qubit resets than standard circuit depth and should aid in analyzing and designing future qubit-efficient algorithms.

We begin by reviewing previous algorithms in Sec-

²The algorithm in [51] also uses a number of qubits independent of n , and in contrast to the polynomial-time algorithms described in this work, it can be used to compute $\text{Tr}(\rho_A^\alpha)$ for non-integer α . However, its time complexity scales exponentially in the system size.

tion 2. In Section 3, we introduce our qubit-efficient variants. These algorithms are summarized in Table 1. In Section 4, we present numerical simulations comparing the performance of the new and original algorithms in the presence of noise. In Section 5, we report the results of an experimental implementation of the qubit-efficient TCT on the Honeywell System Model H0. We introduce effective circuit depth in Section 6, followed by discussion in Section 7. Details of the numerical simulations are in the Appendix.

2 Previous Work

Given an entangled pure state $|\psi\rangle$ defined on subsystems A and B , discarding the qubits associated with subsystem B produces a mixed state ρ_A . If subsystem A of interest has k qubits, a subsystem B of equal size is sufficient to create any mixed state on A (the converse of purification). In what follows, we will assume that registers A and B each have k qubits — so $|\psi\rangle$ is $2k$ qubits and ρ_A is k qubits. It is straightforward to generalize the algorithms discussed in this paper to cases where the registers are different sizes, in part because A subsystems only ever interact with other A subsystems and B subsystems with other B subsystems.

Standard methods for entanglement spectroscopy begin with the observation from [29, 49] that traces of powers of the reduced density operator, i.e. $\text{Tr}(\rho_A^n)$ for $n = 1, \dots, n_{\max}$, can be used to approximately reconstruct the largest n_{\max} eigenvalues of ρ_A via the Newton-Girard method. This is especially useful given we are often interested in a small number $n_{\max} \ll 2^k$ of the largest eigenvalues. Alternatively, $\text{Tr}(\rho_A^n)$ can be used to exactly compute the Rényi entropy of order n (see, e.g. [50]).

Such traces can be expressed as the expectation values of the unitary cyclic permutation operators P_A^{cyc} over n copies of the state $\rho = |\psi\rangle\langle\psi|$,³ where the subscript indicates that it acts only on the A subsystems of the copies of ρ , i.e.

$$\text{Tr}(\rho_A^n) = \langle P_A^{\text{cyc}} \rangle_{\rho_A^{\otimes n}} = \langle P_A^{\text{cyc}} \rangle_{\rho^{\otimes n}}. \quad (1)$$

Requiring n copies is optimal for computing an n -th degree polynomial of ρ since unitary evolution is linear [10]. Using the permutation operator is a generalization of the premise for the well-known Swap Test [11, 25], where the SWAP gate is a cyclic permutation over two qubits. Together, the previous two paragraphs reduce the problem of entanglement spectroscopy to estimating the expectation value of a unitary operator.

³See, e.g. Theorem 4.13 and Definition 3.15 of [7] for a proof.

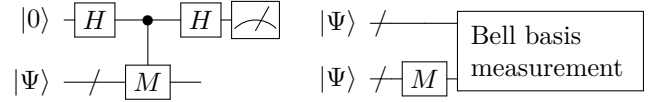


Figure 1: (a) The Hadamard Test is on the left. It applies controlled- M and requires one copy of $|\Psi\rangle$ and one ancilla qubit. The expectation values of the Pauli Z and $-Y$ operators of the ancilla qubit are the real and imaginary parts of $\langle\Psi|M|\Psi\rangle$, respectively. (b) The Two-Copy Test is on the right. It requires two copies of $|\Psi\rangle$ and applies M to one of them. Performing an overlap measurement using the Bell basis algorithm gives $|\langle\Psi|M|\Psi\rangle|^2$.

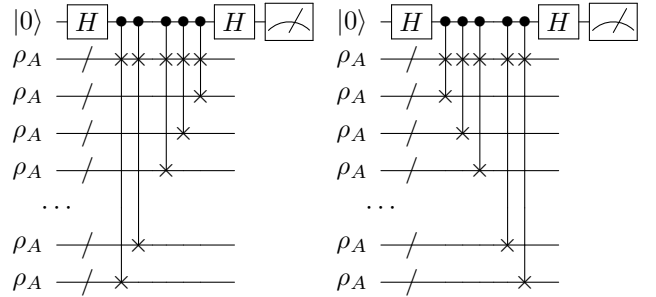


Figure 2: The HT algorithm, which is the Hadamard Test with a cyclic permutation operator, computes $\text{Tr}(\rho_A^n)$. As the two circuits show, P_A^{cyc} can be implemented as either a left shift or right shift, respectively. Each CSWAP is implicitly implemented by k sequential CSWAPs. The circuit depth is $T_{\text{sp}} + \Theta(kn)$ and the width is $2kn + 1$.

Note that there are several definitions and implementations of the cyclic permutation operator that are equivalent for the purposes of entanglement spectroscopy. First, a cyclic permutation may be a left shift or a right shift, shifting the contents of the first register back to the n -th register or forwards to the second register, respectively. These different definitions are illustrated in Fig. 2. Second, either type of shift can be implemented using $n - 1$ transpositions (swaps), but there are many possible decompositions. For example, using cycle notation, (4123) is equivalent to $(34)(23)(12)$ and to $(12)(13)(14)$. Eq. 1 holds for all of these variants. Our choices of when to use left shift and right shift and our choice of decomposition are arbitrary.

2.1 Entanglement Spectroscopy Hadamard Test (HT)

The standard algorithm for estimating the expectation value of an arbitrary unitary operator M on a state $|\Psi\rangle$ is the Hadamard Test, illustrated and described in Fig. 1(a). To be clear, the real part of $\langle\Psi|M|\Psi\rangle$ is calculated by $p_0 - p_1$, where p_i denotes the probability

that the ancilla qubit is in state $|i\rangle$. The estimates of p_0, p_1 converge with accuracy $O(1/\sqrt{S})$ where S is the number of times the circuit is run (this scaling can be improved using quantum amplitude estimation [29], but this is impractical in the NISQ era).

For entanglement spectroscopy, we substitute the cyclic permutation operator P_A^{cyc} for M and $|\psi\rangle^{\otimes n}$ for $|\Psi\rangle$. This is the algorithm of Johri, Steiger, and Troyer [29], which we refer to as the Entanglement Spectroscopy Hadamard Test (HT).⁴ It is illustrated in Fig. 2. HT is a generalization of the Swap Test [11, 25], where the Swap Test is the Hadamard Test with $M = \text{SWAP}$ acting on two states ρ, σ to compute the overlap $\text{Tr}(\text{SWAP}\rho \otimes \sigma) = \text{Tr}(\rho\sigma)$, which equals the purity $\text{Tr}(\rho^2)$ when $\rho = \sigma$.

A swap of two k -qubit registers can be implemented with k swaps of individual qubits, so the total number of controlled-SWAP (CSWAP) gates in HT is $k(n-1)$. A downside of HT is that all of the swap operations are controlled on the same ancilla, so the CSWAP gates must be applied sequentially. Given that CSWAP has constant depth, the circuit depth of HT scales linearly with kn . Specifically, it is $T_{\text{sp}} + T_{\text{cs}}kn + 2T_H = T_{\text{sp}} + \Theta(kn)$, where T_{sp} is the time it takes to prepare a single copy of the state, T_{cs} is the time to implement a CSWAP, and T_H is the time to implement a Hadamard gate. Note that we treat T_{cs} and T_H as constants which depend on the hardware and decompositions used, independent of the input; we also assume all-to-all connectivity. The circuit width of HT is $2kn + 1 = \Theta(kn)$, including the qubits for the B subsystems — recall $|\psi\rangle$ is $2k$ qubits and each subsystem is k qubits.

We just stated the depth of HT as $T_{\text{sp}} + \Theta(kn)$, but previous work has usually stated the depth as being linear in kn , dropping the time for state preparation. In the model considered in previous work, algorithms accept many copies of $|\psi\rangle$ as input at the beginning, as in Fig. 1. Thus, state preparation and the algorithm are considered independently. Our new algorithms will require a setting in which state preparation is intertwined with the rest of the algorithm. This is also a reasonable setting since the states used as input in the previous setting are prepared by some physical procedure which we may represent by a subcircuit as in Fig. 4. Rather than many copies of $|\psi\rangle$, we assume that a description of the state preparation circuit is given as input. (In fact, our algorithms work in the more restricted black box model, in that we only consider the output of the state preparation circuit, never examining the circuit itself.) To fairly compare the new and previous algorithms, we assume the same setting, so we include state

⁴In [50], this was referred to as the JST algorithm after the initials of the authors of [29]. Here, we opt to use a more descriptive name which is easier to extend to our new algorithms.

preparation in the depths of all algorithms.

Finally, we note that in the setting where state preparation is included as part of the algorithm, a simple modification can improve noise resilience. Instead of preparing all copies at the start of the algorithm, state preparation should be delayed until needed. For example, in the second HT circuit in Fig. 2, the preparation of the third copy could be delayed by T_{cs} compared to the preparation of the first copy.

2.2 Entanglement Spectroscopy Two-Copy Test (TCT)

Recently, Cincio, Subaşı, Sornborger, and Coles [13] rediscovered an algorithm of Garcia-Escartin and Chamorro-Posada [23] for computing the overlap between two states using Bell basis measurements of corresponding pairs of qubits from each state followed by efficient classical post-processing. For intuition, the Bell basis is an eigenbasis of the SWAP operator, allowing a Bell basis measurement to reproduce the result of the Swap Test. Ref. [23] related this algorithm to the Hong-Ou-Mandel effect in the context of quantum optics and referred to it as a destructive Swap Test, while [13] emphasized that it can be implemented with constant depth in a quantum computer. We refer to this algorithm as the Bell basis algorithm.

For completeness, a Bell basis measurement on a pair of qubits involves applying a controlled-not (CNOT) gate and then a Hadamard on one of the qubits followed by measuring in the standard basis; see Fig. 3(a) for an example. Importantly, because each CNOT acts on a different pair of qubits, the measurement can be performed with a single layer of CNOTs and a single layer of Hadamards, which is constant depth. To then compute the overlap between two states ρ and σ each of size m , the classical post-processing step is to compute the linear function

$$\sum_{r,s \in \{0,1\}^m} (-1)^{r_1 s_1 + \dots + r_m s_m} p_{r,s}, \quad (2)$$

where $p_{r,s}$ is the experimentally measured frequency that the first m qubits, corresponding to ρ , are measured in state $|r\rangle$ and that the second m qubits, corresponding to σ , are measured in state $|s\rangle$. This classical step can be performed in time linear in the number of trials.

Building on the Bell basis algorithm, Subaşı, Cincio, and Coles [50] introduced an algorithm for estimating $|\langle \Psi | M | \Psi \rangle|^2$ for a unitary operator M , which is called the Two-Copy Test. Classical post-processing can then yield the magnitude of the expected value, $|\langle \Psi | M | \Psi \rangle| = |\text{Tr}(M |\Psi\rangle\langle\Psi)|$. This algorithm relies on

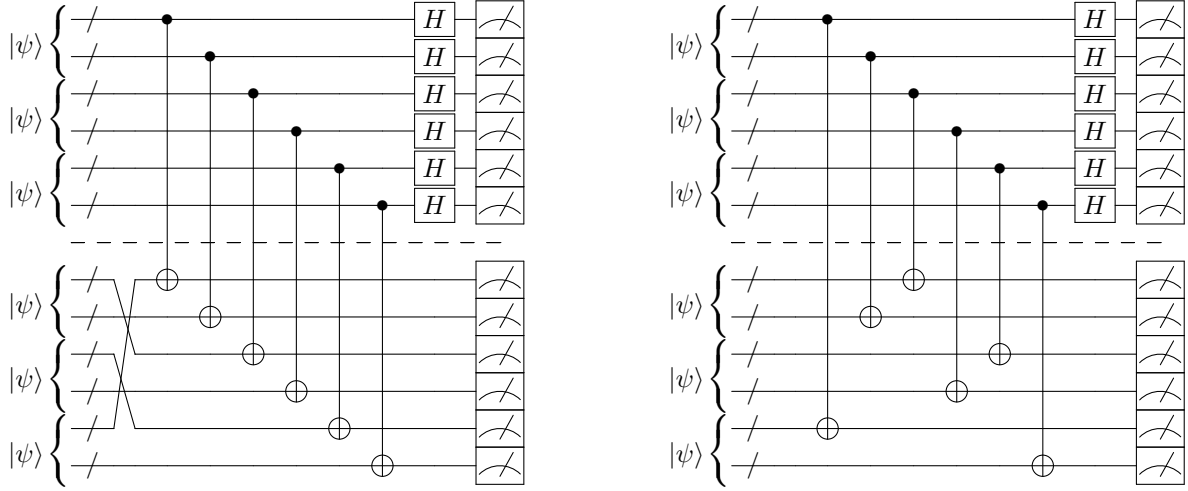


Figure 3: (a) On the left is the Two-Copy Test with a P_A^{cyc} gate for computing $\text{Tr}(\rho_A^n)^2$ for $n = 3$. Gates on registers are implicitly applied to all qubits in parallel: for example, H is $H^{\otimes k}$. The top and bottom wires for each state contain subsystems A and B , respectively. The classical post-processing step is as in Eq. (2). (b) On the right is the TCT algorithm, developed from the circuit on the left by indirectly implementing the cyclic permutation by reindexing the CNOT gates. The classical post-processing is as in Eq. (3). The circuit depth is $T_{\text{sp}} + O(1)$ and the width is $4kn$.

the observation that for pure states, the squared expectation value $|\langle \Psi | M | \Psi \rangle|^2$ is equivalent to the overlap between states $|\Psi\rangle$ and $M|\Psi\rangle$. As depicted in Fig. 1(b), the Two-Copy Test accepts two copies of the state $|\Psi\rangle$, applies M to one, and performs an overlap measurement using the Bell basis algorithm. This requires enough qubits for two copies of the state. Because the Bell basis measurement is constant-depth, the depth of the overall algorithm only depends on M and T_{sp} . Unlike the Hadamard Test, the Two-Copy Test cannot be used to obtain the real and imaginary parts of $\langle \Psi | M | \Psi \rangle$. Also, while the Hadamard Test works both for pure states and for mixed states, the Two-Copy Test can only be used to compute expectation values for pure states. Estimating $|\langle \Psi | M | \Psi \rangle|$ using the Two-Copy Test converges with accuracy $O(1/\sqrt{S})$, where S is the number of times the circuit is run.

A crucial difference between the Hadamard Test and the Two-Copy Test is that the latter uses the unitary M instead of controlled- M . Recalling that HT has linear depth because the controlled gates have to be applied sequentially, eliminating the control not only reduces the gate count, it also allows for the possibility of parallelization. Given the right operator M , this can lead to applications of the Two-Copy Test with very shallow circuits.

The Two-Copy Test can be applied to entanglement spectroscopy by observing that Eq. (1) can be recovered from $|\langle \Psi | M | \Psi \rangle|^2 = \text{Tr}(\rho_A^n)^2$ with the choices $M = P_A^{\text{cyc}}$ and $|\Psi\rangle = |\psi\rangle^{\otimes n}$ combined with the fact $\text{Tr}(\rho_A^n)$ is real and nonnegative. This circuit is depicted

in Fig. 3(a). Unlike HT, this requires access to the full state $|\psi\rangle$. Next, since an (uncontrolled) permutation operator is equivalent to a relabeling of the registers on which it acts, the cyclic permutation can be implemented without any gates by carefully changing the registers which the CNOTs in the Bell basis measurement act on and reindexing the classical post-processing formula. We refer to this algorithm as the Entanglement Spectroscopy Two-Copy Test (TCT) [50]. This circuit is depicted in Fig. 3(b).

To be clear, let $|\psi_i\rangle$ denote the i -th state in the first copy of $|\Psi\rangle$ and $|\psi'_i\rangle$ denote the i -th state in the second copy of $|\Psi\rangle$, where the operator P_A^{cyc} is applied to the second copy. Then, the B subsystem of $|\psi_i\rangle$ is paired with the B subsystem of $|\psi'_i\rangle$, and the A subsystem of $|\psi_i\rangle$ is paired with the A subsystem of $|\psi'_{i-1}\rangle$ when the permutation is a right shift ($|\psi'_{i+1}\rangle$ when a left shift). The edge case of $|\psi_1\rangle$ ($|\psi_n\rangle$ when a left shift) is handled by performing indexing modulo n . The post-processing calculation, derived from Eq. (2), is more complicated but still efficient. Assuming the permutation is a right shift, the formula for $\text{Tr}(\rho_A^n)$ is the square root of

$$\sum_{\substack{\vec{A}_j, \vec{B}_j, \vec{A}'_j, \vec{B}'_j \in \{0,1\}^k \\ j \in [n]}} (-1)^{\sum_{\ell=1}^n \vec{A}_\ell \cdot \vec{A}'_{\ell-1} + \vec{B}_\ell \cdot \vec{B}'_\ell} p_{\vec{A}_1, \vec{B}_1, \dots}, \quad (3)$$

where $p_{\vec{A}_1, \vec{B}_1, \dots}$ is the experimentally measured frequency that for all $i \in [n]$, the qubits initially containing the A and B subsystems of $|\psi_i\rangle$ are measured in the states $|\vec{A}_i\rangle$ and $|\vec{B}_i\rangle$, respectively, and that the qubits initially containing $|\psi'_i\rangle$ are measured in the states $|\vec{A}'_i\rangle$

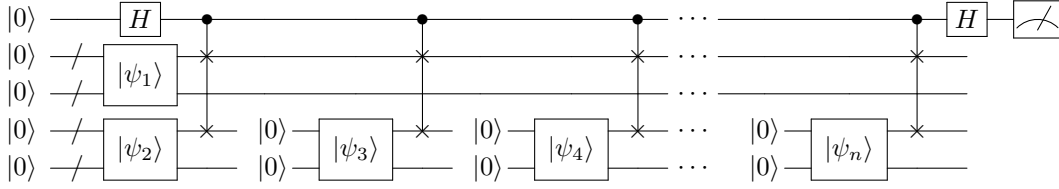


Figure 4: The $4k$ qubit-efficient HT. A break in a wire followed by a new $|0\rangle$ indicates a reset. Each $|\psi_i\rangle$ indicates the preparation of another copy of $|\psi\rangle$; the subscripts are only for guidance. The circuit depth and effective depth are $\Theta(n \times (T_{\text{sp}} + k))$ and the circuit width is $4k + 1$.

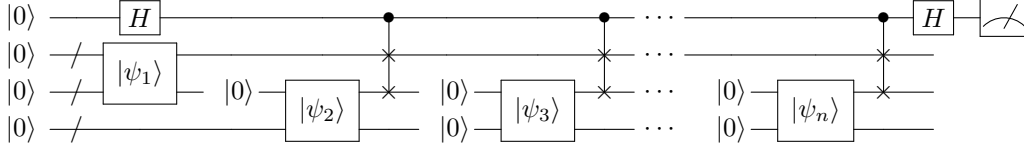


Figure 5: The $3k$ qubit-efficient HT. The circuit depth and effective depth are $\Theta(n \times (T_{\text{sp}} + k))$, asymptotically the same as the $4k$ qe-HT, and the circuit width is $3k + 1$.

and $|\vec{B}_i'\rangle$, respectively. The circuit depth of TCT is $T_{\text{sp}} + T_{\text{cn}} + T_H = T_{\text{sp}} + O(1)$, independent of k and n , where T_{cn} is the time to implement a CNOT gate. This is asymptotically better than HT, but comes with the tradeoff that the circuit width is $4kn$, almost twice the width of HT.

3 Qubit-Efficient Algorithms

In this work, we give variations of the HT and TCT algorithms which achieve asymptotically lower circuit width — proportional to k but independent of n — without significantly increasing the susceptibility to noise. We refer to these as *qubit-efficient* HT and TCT algorithms. For both the HT and TCT, we give two variants where one achieves lower width than the other; we do this in part because the higher-width variants are easier to understand. The high-level idea we rely on is to prepare only as many copies of the state $|\psi\rangle$ at a time as necessary. The structures of both HT and TCT are such that every time a new copy is needed to interact with existing copies, one of the existing copies is finished, with no gates left to act on it. So, the latter copy's qubits may be measured, reset, and used to prepare the new copy of the state (the measurement is optional depending on the particular algorithm: for example, some error mitigation methods might require measurement results). This allows us to run the HT and TCT algorithms with a circuit width independent of n . These algorithms rely on the ability to reset qubits in the course of a quantum computation.

3.1 Qubit-Efficient HT

Observe that every register in the HT circuit (Fig. 2) except for the ancilla qubit interacts with just two other registers and the ancilla. The state of the ancilla qubit, and so the output of the algorithm, is not affected by discarding other registers, so they can be reset and recycled once the last gate on them has been applied. At any time, we just need enough qubits to prepare two copies of the state and the ancilla qubit. So, by resetting qubits when we are done with their contents, we can implement HT using a constant number of registers. Note that measuring the qubits before resetting them is unnecessary unless one wants to perform postselection [33, 50].

Our first algorithm implementing this qubit-efficient strategy is given in Fig. 4. Recalling that $|\psi\rangle$ is a state on $2k$ qubits, the circuit width is $4k + 1$, independent of n . We refer to this algorithm as the $4k$ qe-HT.

The action of the algorithm can be verified by computing the reduced density matrix of the ancilla qubit after the m -th controlled-SWAP operation:

$$\rho_{\text{anc}} = \frac{1}{2}I + \frac{1}{2}\text{Tr}(\rho_A^m)X. \quad (4)$$

Thus, after n controlled-SWAP operations, a measurement in the X basis yields $\text{Tr}(\rho_A^n)$, as desired.

Our second algorithm comes from the observation that in the $4k$ qe-HT, the third register stays idle after the first state preparation. So, instead of preparing two copies simultaneously, we modify the algorithm to prepare one copy, reset the qubits associated with subsystem B , and reuse them to prepare successive copies. This saves k qubits. This algorithm is given in Fig. 5.

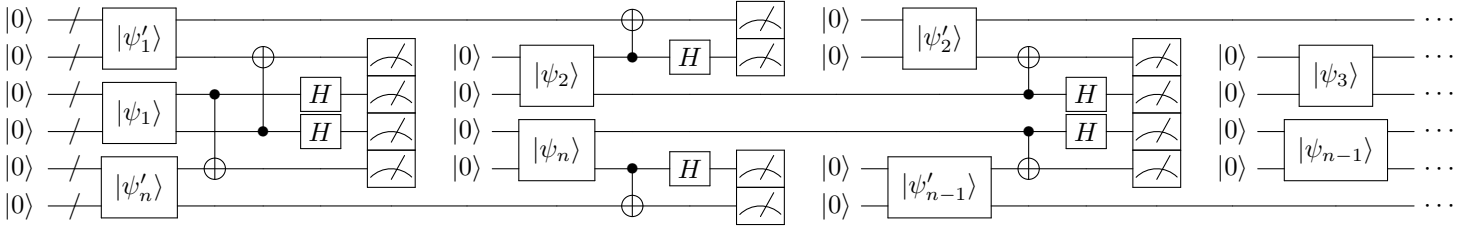


Figure 6: The $6k$ qubit-efficient TCT. The circuit depth is $\Theta(n \times (T_{\text{sp}} + 1))$, the effective depth is $2(T_{\text{sp}} + O(1)) = \Theta(T_{\text{sp}} + 1)$, and the width is $6k$.

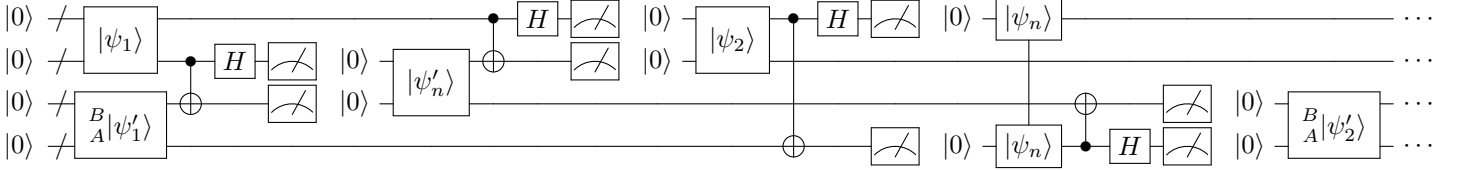


Figure 7: The $4k$ qubit-efficient TCT. Some of the copies of $|\psi\rangle$ are prepared “upside-down”, preparing the B subsystem on the upper wire, and some copies are prepared using non-adjacent wires. The circuit depth is $\Theta(n \times (T_{\text{sp}} + 1))$, about twice that of $6k$ qe-TCT, the effective depth is $3(T_{\text{sp}} + O(1)) = \Theta(T_{\text{sp}} + 1)$, and the width is $4k$.

Here, the circuit width is $3k + 1$ qubits. We refer to this as the $3k$ qe-HT.

Our two qubit-efficient versions differ only slightly. The second version requires k fewer qubits than the first one. This savings come at the cost that the second wire will have to wait longer before gates are applied, exposing it to more thermal noise. The length of the extra wait depends on how long state preparation takes, but compared to the depth of the $n - 1$ other state preparations, the effect should be negligible. After the first two state preparations, the circuits are effectively the same.

Next, we compare the two qubit-efficient versions to the original HT algorithm (Fig. 2). First, all of the circuits have the same number of gates and measurements, so we expect gate and readout errors to affect them similarly. If the fidelity of qubit reinitialization, i.e. qubit reset, is significantly worse than the fidelity of initialization in the beginning of computation, the qubit-efficient algorithms will have a disadvantage. The depth of the original algorithm is $T_{\text{sp}} + \Theta(kn)$ while the depths of the two qubit-efficient algorithms are $\Theta(n \times (T_{\text{sp}} + k))$. Thus, when T_{sp} is small, the original and new algorithms have similar depth. Fortunately, even short-depth circuits have the potential to prepare many interesting states; indeed, the recent quantum supremacy experiment by [6] used 53-qubit circuits with just forty layers of gates. These observations suggest that our new algorithms may perform similarly to the original

algorithm, given small T_{sp} , even as they achieve asymptotically lower circuit width.

3.2 Qubit-Efficient TCT

In the TCT (Fig. 3), each copy of the state interacts with two other copies of the state, one via its A subsystem and one via its B subsystem. After these interactions and, in the case of the A subsystem, a Hadamard gate, the registers containing that copy can be measured, reset, and reused. Therefore, we just need enough qubits to maintain three copies of the state. However, we must be careful, since while the HT did not require any particular ordering of the n copies of $|\psi\rangle$, the TCT does. Fortunately, the TCT is structured such that simply following a greedy strategy of preparing whichever copy is needed to interact with the current longest-lived copy is sufficient.

Our first qubit-efficient variant is given in Fig. 6. The circuit width is $6k$ qubits, so we refer to this algorithm as the $6k$ qe-TCT. Recall that we refer to the first n copies of the state by $|\psi_i\rangle$ and to the second n copies (which are acted on by the permutation operator) by $|\psi'_i\rangle$.

To further reduce the number of qubits, we observe that it is unnecessary to simultaneously prepare both copies needed by the current one. For example, after preparing $|\psi_1\rangle$, it is sufficient to first prepare $|\psi'_1\rangle$, interact the B subsystems of those copies, and then pre-

pare $|\psi'_n\rangle$ and interact the A subsystems. The register containing the B subsystem of $|\psi'_1\rangle$ can be measured, reset, and reused to prepare $|\psi'_n\rangle$. In this way, four such registers is sufficient. Our second variant is given in Fig. 7. The circuit width is $4k$, and we refer to it as the $4k$ qe-TCT.

Our two qubit-efficient variants are similar. The second version uses $2k$ fewer qubits. In both versions, half of the wires are measured quickly, after just $T_{\text{sp}} + O(1)$ timesteps. While the remaining wires in the first algorithm are used for $2T_{\text{sp}} + O(1)$ timesteps between initialization and measurement, the wires in the second algorithm must be maintained for $3T_{\text{sp}} + O(1)$ time. So, the second algorithm may suffer from thermal noise more than the first.

Next, we compare the two qubit-efficient versions to the original TCT algorithm. First, all the circuits have the same number of gates and measurements, so we expect gate and readout errors to affect them similarly. If the fidelity of qubit reinitialization, i.e. qubit reset, is significantly worse than the fidelity of initialization in the beginning of computation, the qubit-efficient algorithms will have a disadvantage. The original TCT has depth $T_{\text{sp}} + O(1)$, while our qubit-efficient variants have depth $\Theta(n \times (T_{\text{sp}} + 1))$ (note the constant term is only asymptotic, like $O(1)$, rather than a literal 1). Based on this observation, the qubit-efficient versions might appear like they should perform significantly worse in the presence of noise.

Given the results of [50] demonstrating that the TCT is more noise-resilient than the HT, we expect that each of the variants of the TCT should outperform their HT analogs, e.g. we expect the $6k$ qe-TCT to outperform the $4k$ qe-HT. However, it is unclear a priori whether the qubit-efficient variants of the TCT will still outperform the original HT.

4 Numerical simulations

In this section, we test the performance of our qubit-efficient algorithms for entanglement spectroscopy and compare them to the original versions. The most significant observation from our results is that our qubit-efficient algorithms perform similarly to the originals in the presence of noise.

We use IBM's Qiskit [2] and QASM simulator to numerically simulate noisy quantum circuits. Our simulations include thermal relaxation and decoherence error, readout error, and gate noise in the form of depolarizing and Pauli errors. See the Appendix for details and parameters. These simulations and selection of noise parameters are independent of the experimental results in the next section. The postselection methods introduced by [33, 50] for improving the accuracy of the HT and

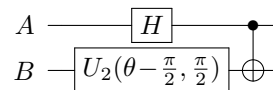


Figure 8: State preparation circuit. The $U_2(\phi, \lambda)$ gate has the matrix form $\frac{1}{\sqrt{2}}((1, -e^{i\lambda}), (e^{i\phi}, e^{i(\phi+\lambda)}))$.

TCT apply in the same way to our qubit-efficient variants, but because we expect them to affect the old and new algorithms similarly, we do not implement postselection here.

The number of qubits we can simulate is limited by memory, and the circuit widths of the original algorithms scale with n and k while the widths of our new algorithms scale only with k . So, in order to simulate the prior algorithms for many values of n , we restrict our simulations to $k = 1$, which corresponds to two-qubit states $|\psi\rangle$ and single-qubit density matrices ρ_A . In this case, knowing $\text{Tr}(\rho_A^n)$ for $n = 2$ is sufficient to reconstruct the entire entanglement spectrum. Although the values for $n > 2$ are redundant, we compute them in order to assess the performance of the algorithms.

For each n , we generate twenty quantum states with varying levels of entanglement ranging from product states to maximally entangled using the circuit in Fig. 8. We choose the twenty angles θ therein such that the associated $\text{Tr}(\rho_A^n)$ are evenly spaced from the minimum to the maximum possible values, from 2^{1-n} (fully mixed) to 1 (pure state).

After simulating the algorithms for all twenty states corresponding to a particular n , we first plot the values of the ideal $\text{Tr}(\rho_A^n)$ versus the value estimated by each quantum algorithm. Fig. 9 shows these plots for simulations of all algorithms, including original algorithms and our qubit-efficient algorithms, for $n = 2$ to $n = 6$. Note that an ideal set of results would lie on a straight line from $(2^{1-n}, 2^{1-n})$ to $(1, 1)$ with slope equal to one. Our results deviate from this line due both to simulated hardware noise and to statistical noise due to finite sampling. Intuitively, random hardware noise leads pure states to appear more mixed, leading the results to concentrate about a flatter line, and statistical noise causes the results to deviate about that line. Observe that as long as the data concentrate about some line, it is easier to confidently identify a state as more or less mixed based on the algorithm's estimate for $\text{Tr}(\rho_A^n)$ when the slope of the line is closer to one, i.e. when the line is steeper. This is in contrast to an error in the vertical intercept of a line, which can be corrected by learning the error and shifting future results. Therefore, we characterize the performance of the algorithms by their slopes. For each value of n that we tested, we compute the slope of each line in plots like in Fig. 9

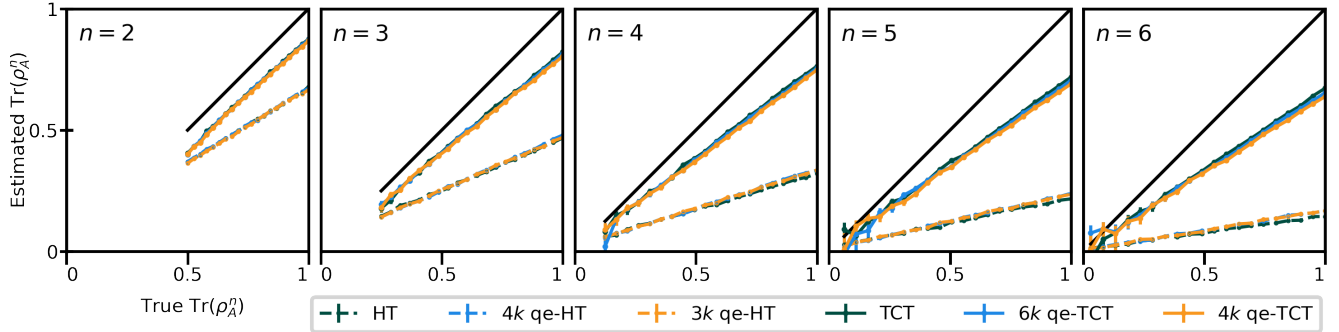


Figure 9: (Color) Numerical simulations of prior algorithms and our variants estimating twenty values of $\text{Tr}(\rho_A^n)$ under simulated hardware noise for several values of n . Noise parameters are in the Appendix. True values are on the horizontal axis, estimated values on the vertical. The topmost line shows ideal results. All points are marked. Error bars, some of which are smaller than the lines, are based on expected statistical noise only. The upper group of lines are the variants of the TCT. The lower group of dashed lines are the variants of the HT. We compare the slopes of these lines in Fig. 10.

using a linear regression. We plot the values of n versus those computed slopes.

Results for all algorithms, including original algorithms and our qubit-efficient versions, for $n = 2$ to $n = 6$ are given in Fig. 10. Note that a noiseless implementation would have slope equal to one for all n . Decreasing values indicate an algorithm’s performance degrading for larger values of n . We were limited to $n = 6$ by the Two-Copy Test, for which the circuit width scales as $4kn$; simulating 28 qubits was impractical due to time-constraints, and 32 qubits would be impractical due to memory constraints. In contrast, the number of qubits required for our qubit-efficient algorithms is independent of n , so we are able to simulate these algorithms for much larger values of n .

Results for the qubit-efficient algorithms for $n = 2$ to $n = 20$ are given in Fig. 11. The noise strength is reduced compared to the previous simulations (see the Appendix for details).

The most significant observation from these results is that our qubit-efficient algorithms perform similarly to the original variants. In Fig. 10, the qubit-efficient variants of the HT perform very similarly to the original algorithm. In fact, the original HT performs slightly worse than the qubit-efficient variants, likely due to Qiskit’s “as soon as possible” gate scheduling which prepares copies of $|\psi\rangle$ for the original HT earlier than is optimal. As we stated previously, we expected the qubit-efficient variants to perform similarly to the original when T_{sp} is small. In the case of the TCT, the qubit-efficient variants suffer almost no degradation compared to the original algorithm. For both the HT and TCT, the wider, lower-depth qubit-efficient variants perform better than the corresponding lower-width algorithms. We note that, as explored further in [50], the TCT and its variants are more susceptible to statistical noise than the HT. The TCT is most affected by statistical noise

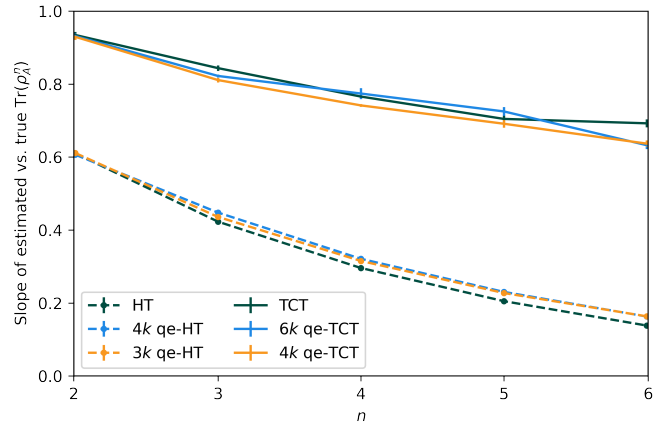


Figure 10: (Color) Comparison of the performance of all HT and TCT variants in the presence of simulated noise for $n = 2$ to $n = 6$. Values of n are on the horizontal axis, slopes computed by a linear regression on the corresponding lines in Fig. 9 are on the vertical axis. Ideal results would produce a horizontal line at height 1. Error bars, some of which are smaller than the lines, are based on the quality of the linear fit for the corresponding data in Fig. 9. The upper group of lines are the TCT variants. The lower group of dashed lines, with marked points, are the HT variants.

when estimating small values of $\text{Tr}(\rho_A^n)$, which is the case for highly entangled states $|\psi\rangle$ and exasperated by large powers n ; this is visible in Fig. 9.

In Fig. 11, simulating larger n , the two qubit-efficient variants of HT continue to perform almost identically, as expected. For the TCT variants, the $6k$ qe-TCT slightly outperforms the $4k$ qe-TCT. Notably, both the qubit-efficient variants of the TCT still appear to produce meaningful results when $n = 20$ (as good as the qe-HT when $n = 8$).

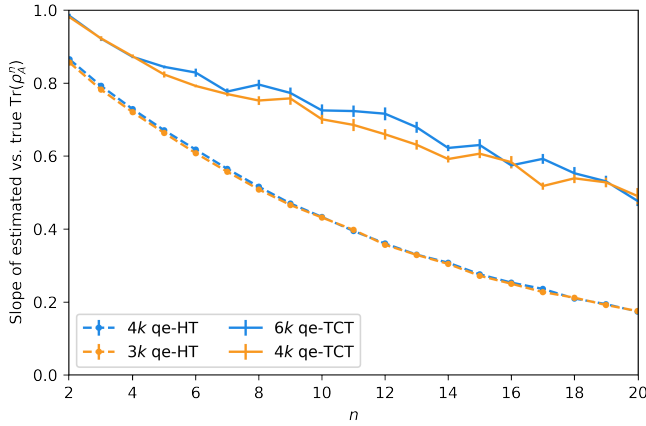


Figure 11: (Color) Comparison of the performance of our qubit-efficient variants of the HT and TCT for $n = 2$ to $n = 20$. The noise strength is reduced here compared to the previous simulations (see the Appendix for details).

5 Experimental demonstration on Honeywell System Model H0

In this section, we report the results of testing one of our qubit-efficient algorithms on the Honeywell System Model H0 [41]. We were able to estimate $\text{Tr}(\rho_A^n)$ for larger n than would have been possible on the device using any previous algorithm, and our results correctly distinguish more and less entangled states.

This quantum computer is a trapped-ion quantum charge-coupled device architecture; for details, see [41]. At the time of access (September 2020), the device supported six qubits and supported mid-circuit measurements and qubit resets. In order to test the widest variety of parameters possible given our limited time on the device, we chose to test one qubit-efficient algorithm, choosing the one which performed best in simulations: the $6k$ qe-TCT (Fig. 6).

As in our numerical simulations, we set $k = 1$, corresponding to two-qubit states $|\psi\rangle$ and one-qubit ρ_A . We prepare three states with varying levels of bipartite entanglement using the state preparation circuit of Fig. 8, setting the angle θ therein to $\theta = 1.33, 1.05, 0.87$. Because the TCT is more sensitive to statistical noise when estimating smaller values of $\text{Tr}(\rho_A^n)$, corresponding to more mixed ρ_A , and because we had only a limited number of runs available, we chose these states to be closer to pure than to fully mixed.

For each of the three states, we run the $6k$ qe-TCT for $n = 2, \dots, 7$ for 1,000 runs. Note that given six qubits, the original TCT would not fit on the device even for $n = 2$. Each circuit was sent via the HQS API, specified using operations U_2 , CNOT, Measure, and Reset (U_2 defined in Fig. 8). From there, each

circuit was compiled to the device’s native gate set, including standard optimization according to Honeywell’s software stack, and submitted to the device. Circuits were sent in batches, with calibration performed within and between each batch.

Results are shown in Fig. 12. Rather than comparing several algorithms, here we test the performance of the $6k$ qe-TCT on several different inputs. For each of the three states, we plot the values of n versus the estimates for $\text{Tr}(\rho^n)$.

After receiving the results from our tests, we found that two data points, for $\theta = 1.33, n = 3$ and for $\theta = 1.05, n = 4$, were outliers compared to the rest of the data. Honeywell offered to rerun these tests. Both the initial and second points are shown in Fig. 12.

Because of noise, the results from our tests are insufficient to recover the true, analytical values of $\text{Tr}(\rho_A^n)$. However, results for each of the three states are clearly distinguishable from each other and are correctly ordered according to their degree of entanglement. The data is remarkably smooth across varying n , with simulations predicting more varied outcomes and with these tests using only 1,000 runs versus the simulations in Fig. 9 using 100,000. Although we only tested the algorithm on states closer to pure than fully mixed (recall the minimum value of $\text{Tr}(\rho_A^n)$ is 2^{1-n}), the results appear promising for more entangled states. They also suggest that tests with larger values of n should produce results along a similar trend.

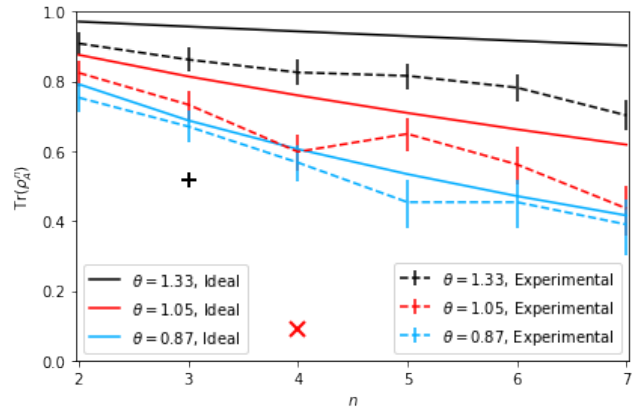


Figure 12: (Color) Results from our experimental test of the $6k$ qe-TCT on the Honeywell System Model H0. Values of n are on the horizontal axis, $\text{Tr}(\rho_A^n)$ on the vertical. Three different states ρ_A corresponding to different values of θ were used. For each state, there is a solid line of true values and a dashed line of experimental results. Tests were repeated for two data points because the results were outliers; the initial points, one for $\theta = 1.33$ and one $\theta = 1.05$, are respectively marked by ‘+’ and ‘x’. Error bars are based on expected statistical noise.

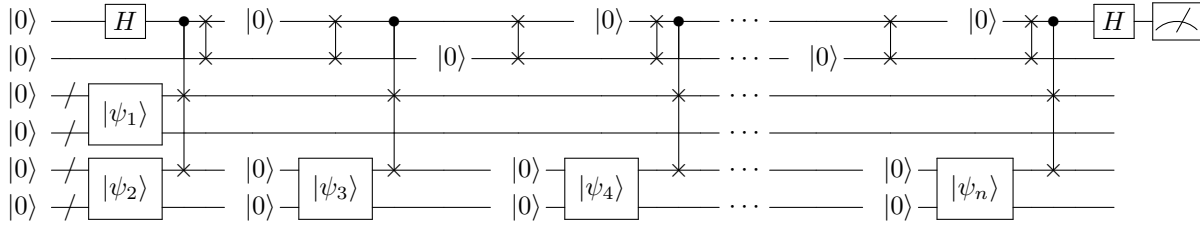


Figure 13: An alternative implementation of the qe-HT that uses an extra qubit to disguise how long the ancilla qubit is required to remain coherent.

6 Effective circuit depth

In this section, we introduce a generalization of circuit depth which is more useful for circuits using qubit resets, which we call effective depth. The depth of a circuit is defined as the number of timesteps assuming that gates can be applied in parallel, or equivalently as the maximum length of a path from the input to the output. Circuit depth is often used to quantitatively judge how susceptible a quantum computation will be to thermal decoherence and relaxation noise. Intuitively, the higher a circuit’s depth, the more time during which the circuit may be affected by noise. This is especially relevant in the NISQ era, as coherence times remain a primary limiting factor on tractable problem sizes. However, depth is only a heuristic for judging noise-resilience. Circuits may be affected by various sources of noise besides thermal noise, and comparing the depths of two circuits does not perfectly predict the relative performance of the circuits even when the noise model is restricted to thermal noise. For example, circuits which produce highly entangled states will be significantly more affected by decoherence than circuits which remain in computational basis states (entirely classical information) even when those circuits have the same depth. Nevertheless, considered alongside other factors, circuit depth is a convenient, often-used tool for assessing quantum algorithms.

In the setting of circuits that use qubit resets, circuit depth is no longer useful for assessing noise resilience. Consider, for example, that the depth of the original TCT is $T_{\text{sp}} + O(1)$ while the depths of the qubit-efficient versions are $\Theta(n \times (T_{\text{sp}} + 1))$, an asymptotic increase. But, as shown numerically in Section 4, the algorithms perform similarly in the presence of noise. Circuit depth judges circuits with resets too harshly. Anticipating increased use of qubit resets, we would like a measure which incorporates their presence.

Defining such an attribute is subtle. A naive idea for a depth-like predictor of noise-resilience for circuits with qubit resets might be the largest amount of time between resets of any particular qubit. However, consider the alternative implementation of the qe-HT shown in

Fig. 13. This implementation utilizes two ancilla qubits where our previous implementations used one, making frequent swaps between the two ancilla qubits. It is designed to obfuscate the long time for which the ancilla qubit in the qe-HT must be kept coherent. This naive measure would rate this circuit as $\Theta(k)$, and with further changes this could be made $O(1)$. Clearly though, information stored in the ancilla is just as exposed to thermal noise in this new circuit as in the other qubit-efficient circuits.

Instead, our definition is inspired by the idea of information flow and locality. At a high level, quantum information is only transferred between qubits when multi-qubit gates are applied.⁵ In particular, the corruption of quantum information due to noise on one register cannot propagate to another register except through future multi-qubit interactions. These ideas are considered further in [47, 48]. By focusing on information flow, the shortcomings of traditional circuit depth for circuits using resets can be eliminated.

We define the *effective depth* of a circuit to be the maximum length of a path along which there is information flow. Equivalently, it is the maximum number of timesteps for which some quantum information is propagated. Such directed paths can be constructed by beginning from any qubit (re)initialization, following the qubit, optionally crossing from one qubit to another when there is a two-qubit gate between them, and terminating when there is a reset or when the last operation is reached;⁶ the longest path which can be formed in this way gives the effective depth.

To justify our definition, consider the following. First, effective depth reduces to the standard definition of depth for circuits which do not use resets. Second, observe that the graph of paths stemming from any set of qubit initializations can effectively be viewed as a subcircuit. Then, an equivalent definition of effective depth is the maximum (standard) depth of such a subcircuit.

⁵We consider a level of abstraction that ignores potential crosstalk.

⁶A similar but distinct construction is used in [9] with the different motivation of parallelizing circuits.

From this perspective, effective depth is a natural extension of depth, rating a circuit with resets according to the depth of its largest complete subcircuit. Third, although effective depth is not a perfect tool, it is a heuristic which provides a worst-case assessment just as standard circuit depth does. The length of the longest path may be unusually long compared to the rest of the paths, some multi-qubit gates may transfer information asymmetrically, or the distribution of inputs may mean certain paths are more significant than others: these factors should be considered alongside effective depth just as additional factors are needed alongside traditional depth. Finally, effective depth sensibly rates all of the circuits in this article, as we discuss next.

Since effective depth reduces to standard depth for circuits which do not use resets, the effective depth of HT is $T_{\text{sp}} + \Theta(nk)$ and the effective depth of TCT is $T_{\text{sp}} + O(1)$. For the qubit-efficient variants of HT, the ancilla qubit leads to effective depths the same as their depths, $\Theta(n \times (T_{\text{sp}} + k))$, which asymptotically matches that of HT when T_{sp} is small. Now, for the qubit-efficient TCT variants, standard circuit depth is insufficient to explain our numerical results. The effective depth of $6k$ qe-TCT is $2(T_{\text{sp}} + O(1))$ and the effective depth of $4k$ qe-TCT is $3(T_{\text{sp}} + O(1))$. These values asymptotically match the effective depth of the original TCT, $T_{\text{sp}} + O(1)$, which helps explain why our qe-TCTs perform just as well as the original TCT: their effective depths are the same.⁷ Finally, effective depth assigns the same value to the contrived qe-HT circuit (Fig. 13) as standard depth, thus avoiding the potential pitfall of assessing this circuit too gently. Our numerical results in Section 4 are consistent with all these observations.

7 Discussion

In this work, we introduced new qubit-efficient algorithms for performing entanglement spectroscopy via computing $\text{Tr}(\rho_A^n)$ that use qubit resets to achieve asymptotically lower width than previous algorithms. Our numerical results show that the performance of our algorithms is only slightly degraded by noise even as they save a significant number of qubits. First, the qubit-efficient HT requires as few as $3k + 1$ qubits and achieves similar performance to the original HT algorithm; we expect this to hold given small state preparation time T_{sp} . Second, and in particular, the qubit-efficient TCT requires as few as $4k$ qubits while achiev-

⁷Intuitively, the $4k$ qe-TCT should experience three times more thermal noise (two times for the $6k$ qe-TCT algorithm) than the original TCT because its effective depth is three (two) times greater. We have tested this intuition using numerical simulations with only thermal noise, multiplying the gate times for the original TCT by three, and found it correct.

ing similar performance to the original TCT algorithm, and we expect this to hold in general. Our algorithms demonstrate the usefulness of the as yet understudied tool of qubit resets.

Just as the HT algorithm of [29] may be better than the TCT for the case of $n = 2$ (i.e. the Swap Test), the original TCT algorithm of [50] may remain preferable to our new variants for small powers n . Our approach is preferable for values of k and n where at least $4kn$ qubits are unavailable or when a smaller circuit width is desired.

To demonstrate the practicality of our qubit-efficient algorithms, we experimentally implemented the $6k$ qe-TCT for $k = 1$ and $n = 2, \dots, 7$ on the Honeywell System Model H0, which supported (at the time of implementation) six qubits. As a comparison, the original TCT would require 8 qubits for $n = 2$ and so could not be run for any n , while the original HT would require 5 qubits for $n = 2$ and could not be run for any larger n . Although the results of the experiment are too noisy to immediately recover the spectrum, they successfully differentiate and rank states with different amounts of entanglement, which could be useful for quantum simulation applications in the near future.

Traditional circuit depth is insufficient for assessing our algorithms or future algorithms using qubit resets. In contrast, effective depth justifies the performance of our new algorithms; for example, the qubit-efficient variants of the TCT have the same asymptotic effective depth as the original TCT. Our definition is a simple and useful heuristic for predicting noise-resilience, just as traditional circuit depth is for circuits without resets. Notably, when there are no qubit resets present, effective depth reduces to standard circuit depth. Effective depth will be a useful tool in the future design and analysis of qubit-efficient algorithms and it should be preferred over circuit depth for describing circuits with qubit resets.

Quantum error mitigation techniques developed for use in NISQ devices [18] may improve the performance of our algorithms. Notably, postselection strategies for the HT [33] and the TCT [50] can also be used with the corresponding qubit-efficient algorithms (these methods fit into the broader framework of symmetry verification [8]). These postselection strategies generally shift estimates of $\text{Tr}(\rho_A^n)$ upward by a constant independent of the input ρ_A (see Fig. 6 of [50]). This is indeed useful for obtaining more accurate estimates. However, not having a particular application in mind, in this work we decided to compare the performance of various algorithms by their ability to distinguish varying degrees of entanglement (see Sec. 4). Postselection does not seem to improve this ability due to the uniform improvement in estimates. Extrapolation [32, 53] techniques may

be particularly helpful in this regard. Throughout our tests, we observed that noise shifts the algorithms’ estimates of $\text{Tr}(\rho_A^n)$ proportionally to the ideal value such that results plotted as in Fig. 9 consistently remain linear; this effect can be leveraged by testing the algorithm on some known states and comparing the ideal and experimental outputs to extrapolate the effect of noise and correct for the error. We leave further improvement and error mitigation, which will depend on a range of factors including the particular hardware and inputs, for future work.

Similarly, further work on analytically modeling the effect of various potential errors on our algorithms would help to improve their performance. The analysis of [22] on the robustness of the Swap Test would be a good starting point for analyzing the HT and the TCT. The TCT is a special case of convolutional circuits recently studied by [4]. In addition to being convolutional, all of the variants of TCT have constant depth and constant effective depth. These features of TCT provide a framework for understanding its noise resilience.

Developing qubit-efficient algorithms will be critical in the NISQ era. Similar devices with fewer qubits tend to be less noisy than those with more qubits, so it is advantageous to be able to run an algorithm on the smallest quantum device possible. Given a particular device, carefully choreographing operations, qubits resets, and the resulting flow of information will help increase the size of the largest problems that can be solved. Additionally, because these algorithms use fewer qubits, they will benefit from requiring fewer swaps to implement gates between arbitrary qubits on architectures with limited connectivity. The performance could be further improved by designing special purpose devices optimized to run these algorithms. Ongoing work on compiling and optimizing quantum algorithms may enable automatically using qubit resets to reduce circuit width, as well as optimizing reset placement based on qubit connectivity and noise.

As shown in this work, entanglement spectroscopy is one application for which qubit-efficient algorithms are possible. Efficient characterization of the entanglement in quantum states will be useful in many areas. In particular, it is well-suited to the promising NISQ application of quantum simulation. In this context, our qubit-efficient algorithms might be paired with quantum simulation methods which utilize qubit resets in order to reduce the necessary number of qubits, such as recent work on simulating correlated spin systems [21]. Our algorithms may also prove helpful in characterizing the performance of NISQ devices themselves.

Additional algorithms, known and future, may be implemented with fewer qubits using qubit resets. Promising candidates include algorithms which are already

low-depth and which have a structure such that registers generally do not require interaction with many other registers.

Acknowledgments

We thank the Honeywell Quantum Solutions team and Andrew Potter for inspiring discussions. We thank anonymous referees for helpful suggestions of additional details and improvements to the presentation. Part of this work was completed while JY was a participant in the 2019 Quantum Computing Summer School at Los Alamos National Laboratory (LANL), sponsored by the LANL Information Science & Technology Institute. JY acknowledges support from a Vannevar Bush Faculty Fellowship from the Department of Defense, Opportunity No. N0001415RFO11. YS acknowledges funding from LANL ASC Beyond Moore’s Law project and the LDRD program at LANL. Los Alamos National Laboratory is managed by Triad National Security, LLC, for the National Nuclear Security Administration of the US Department of Energy under Contract No. 89233218CNA000001.

A Numerical simulation details

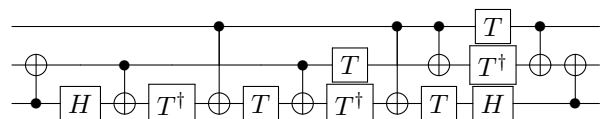
Code available. The code we developed for our numerical simulations is available at the GitHub repository [lanl/qubit-efficient-entanglement-spectroscopy](https://github.com/lanl/qubit-efficient-entanglement-spectroscopy)

We use IBM’s Qiskit [2] to perform our numerical tests. Qiskit is an open-source Python SDK for working with quantum circuits. We implement our circuits using Qiskit and simulate them in the presence of noise using the QASM simulator from the Qiskit Aer module. All simulations were performed locally using Python version 3.6.9 and Qiskit version 0.12.1.

All circuits are implemented using a native gate set of I, U_1, U_2 , and CNOT and with operations Measurement and Reset, where

$$U_1(\lambda) := \begin{pmatrix} 1 & 0 \\ 0 & e^{i\lambda} \end{pmatrix}, \quad U_2(\phi, \lambda) := \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -e^{i\lambda} \\ e^{i\phi} & e^{i(\phi+\lambda)} \end{pmatrix}$$

are gates provided by Qiskit. Note that $H = U_2(0, \pi)$ and $T = U_1(\pi/4)$. We decompose the CSWAP gate as



where the top qubit is the control.

By default, Qiskit applies gates “as soon as possible”, minimizing circuit depth by shifting gates to the left. In order to correctly apply thermal noise, we insert identity gates to fill any gaps when a register must wait for operations to finish on other registers, taking into account the duration of each operation. Thermal noise is applied on a gate-by-gate basis, but no gate noise, i.e. Pauli and depolarizing errors, is applied to the identity gates. Other than the changes mentioned in this and the previous paragraph, all circuits are implemented as they appear in the figures.

The duration of each single-qubit gate is set to one timestep, the duration of a CNOT gate five timesteps, the duration of a measurement three timesteps, and the duration of a qubit reset two timesteps (we always performed a measurement before performing a reset).

In all plots, each value of $\text{Tr}(\rho_A^n)$ is estimated using 100,000 runs. For the plots which include the original HT and TCT, the probability of readout error is 2%, which means that for each single-qubit measurement, there is a 2% probability that the measurement result would be recorded incorrectly. Thermal relaxation and decoherence errors are applied using parameters $T_1 = T_2 = 2000$ and $T_{\text{pop}} = 10^{-7}$. For an operation which takes time t , let $p_{\text{rel}} = 1 - \exp(-t/T_1)$. Then, this means that for each qubit acted on, the probability that the qubit relaxes to $|1\rangle$ is $p_{\text{rel}}T_{\text{pop}}$ and the probability that the qubit relaxes to $|0\rangle$ is $p_{\text{rel}}(1 - T_{\text{pop}})$ (Qiskit can also apply a Z operator to simulate decoherence, but for $T_1 = T_2$, the probability of this is zero). A Pauli error channel is applied to all gates except identity such that for one-qubit gates, the probabilities of an X, Y , or Z operator being applied are each 0.001. A depolarizing error channel $E(\rho) = (1 - \lambda)\rho + \lambda \text{Tr}(\rho) \frac{I}{2^m}$ is applied to all m -qubit gates except identity such that for single-qubit gates, $\lambda = 0.001$. For the CNOT gate, the Pauli and depolarizing error parameters are multiplied by five.

For the plots which only include qubit-efficient algorithms, all of the noise parameters are set the same as above except for the Pauli and depolarization error parameters, which are reduced by a factor of ten. The gate noise is reduced from the previous simulations in order to produce meaningful results for n as high as twenty; we chose to reduce the gate noise because reducing the readout or thermal errors by a similar factor was not as effective.

All plots include error bars, although some of the bars may be too small to see. The error bars in the plots of $\text{Tr}(\rho_A^n)$ versus experimental $\text{Tr}(\rho_A^n)$ are based on the expected statistical noise due to finite sampling and its effect on the post-processing formulas. For the algorithms based on HT, we use Hoeffding’s inequality and a 68% confidence level to calculate an additive error of at most $\pm 2\sqrt{-\ln(0.16)/(2S)}$, where S is the number of

trials performed. For the algorithms based on TCT, we calculate a confidence interval $[c_{\text{low}}, c_{\text{high}}]$ for the raw output, $|\langle \Psi | M | \Psi \rangle|^2$, (before taking the square root) in the same way and set the final confidence interval to $[\sqrt{c_{\text{low}}}, \sqrt{c_{\text{high}}}]$. Note that unlike for HT, the confidence intervals for TCT are affected by $\text{Tr}(\rho_A^n)$, enlarging for smaller values. However, when $\text{Tr}(\rho_A^n)$ is treated as a constant, the size of the error bars scales as $O(1/\sqrt{S})$ in both cases.

The error bars in the plots of n versus computed slopes are influenced by both statistical and simulated hardware noise. The error bar for each point (a value of n versus a slope) is calculated by applying a t -test with a 68% confidence level to the linear regression which produced that slope. Intuitively, more linear underlying data produces smaller error bars.

References

- [1] N. Abdessaïed, R. Wille, M. Soeken, and R. Drechsler. Reducing the depth of quantum circuits using additional circuit lines. In G.W. Dueck and D.M. Miller, editors, *Reversible Computation. RC 2013.*, page 221–233, Berlin, Heidelberg, 2013. Springer-Verlag. DOI: [10.1007/978-3-642-38986-3_18](https://doi.org/10.1007/978-3-642-38986-3_18).
- [2] Héctor Abraham et al. Qiskit: An open-source framework for quantum computing, 2019. URL <https://github.com/Qiskit>.
- [3] Luigi Amico, Rosario Fazio, Andreas Osterloh, and Vlatko Vedral. Entanglement in many-body systems. *Rev. Mod. Phys.*, 80(2):517, 2008. DOI: [10.1103/RevModPhys.80.517](https://doi.org/10.1103/RevModPhys.80.517).
- [4] Galit Anikeeva, Isaac H Kim, and Patrick Hayden. Recycling qubits in near-term quantum computers. *Physical Review A*, 103(4):042613, 2021. DOI: [10.1103/PhysRevA.103.042613](https://doi.org/10.1103/PhysRevA.103.042613).
- [5] Javier Argüello-Luengo, Alejandro González-Tudela, Tao Shi, Peter Zoller, and J Ignacio Cirac. Analogue quantum chemistry simulation. *Nature*, 574(7777):215–218, Oct 2019. DOI: [10.1038/s41586-019-1614-4](https://doi.org/10.1038/s41586-019-1614-4).
- [6] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando G.S.L. Brandao, David A. Buell, et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, 2019. DOI: [10.1038/s41586-019-1666-5](https://doi.org/10.1038/s41586-019-1666-5).
- [7] Costin Bădescu, Ryan O’Donnell, and John Wright. Quantum state certification. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 503–514, 2019. DOI: [10.1145/3313276.3316344](https://doi.org/10.1145/3313276.3316344).

- [8] X. Bonet-Monroig, R. Sagastizabal, M. Singh, and T.E. O'Brien. Low-cost error mitigation by symmetry verification. *Physical Review A*, 98(6):062339, 2018. DOI: [10.1103/PhysRevA.98.062339](https://doi.org/10.1103/PhysRevA.98.062339).
- [9] Anne Broadbent and Elham Kashefi. Parallelizing quantum circuits. *Theoretical Computer Science*, 410(26):2489–2510, 2009. DOI: [10.1016/j.tcs.2008.12.046](https://doi.org/10.1016/j.tcs.2008.12.046).
- [10] Todd A Brun. Measuring polynomial functions of states. arXiv:quant-ph/0401067v2, 2004.
- [11] Harry Buhrman, Richard Cleve, John Watrous, and Ronald de Wolf. Quantum fingerprinting. *Phys. Rev. Lett.*, 87:167902, Sep 2001. DOI: [10.1103/PhysRevLett.87.167902](https://doi.org/10.1103/PhysRevLett.87.167902).
- [12] Pasquale Calabrese and Alexandre Lefevre. Entanglement spectrum in one-dimensional systems. *Phys. Rev. A*, 78:032329, Sep 2008. DOI: [10.1103/PhysRevA.78.032329](https://doi.org/10.1103/PhysRevA.78.032329).
- [13] Lukasz Cincio, Yiğit Subaşı, Andrew T Sornborger, and Patrick J Coles. Learning the quantum algorithm for state overlap. *New Journal of Physics*, 20(11):113022, 2018. DOI: [10.1088/1367-2630/aae94a](https://doi.org/10.1088/1367-2630/aae94a).
- [14] Antonio D Córcoles, Abhinav Kandala, Ali Javadi-Abhari, Douglas T McClure, Andrew W Cross, Kristan Temme, Paul D Nation, Matthias Steffen, and JM Gambetta. Challenges and opportunities of near-term quantum computing systems. *Proceedings of the IEEE*, 2019. DOI: [10.1109/JPROC.2019.2954005](https://doi.org/10.1109/JPROC.2019.2954005).
- [15] M. Dalmonte, B. Vermersch, and P. Zoller. Quantum simulation and spectroscopy of entanglement Hamiltonians. *Nature Physics*, 14(8):827–831, May 2018. DOI: [10.1038/s41567-018-0151-7](https://doi.org/10.1038/s41567-018-0151-7).
- [16] G. De Chiara, L. Lepori, M. Lewenstein, and A. Sanpera. Entanglement spectrum, critical exponents, and order parameters in quantum spin chains. *Phys. Rev. Lett.*, 109:237208, Dec 2012. DOI: [10.1103/PhysRevLett.109.237208](https://doi.org/10.1103/PhysRevLett.109.237208).
- [17] D.J. Egger, M. Werninghaus, M. Ganzhorn, G. Salis, A. Fuhrer, P. Mueller, and S. Filipp. Pulsed reset protocol for fixed-frequency superconducting qubits. *Phys. Rev. Applied*, 10:044030, Oct 2018. DOI: [10.1103/PhysRevApplied.10.044030](https://doi.org/10.1103/PhysRevApplied.10.044030).
- [18] Suguru Endo, Zhenyu Cai, Simon C Benjamin, and Xiao Yuan. Hybrid quantum-classical algorithms and quantum error mitigation. *Journal of the Physical Society of Japan*, 90(3):032001, 2021. DOI: [10.7566/JPSJ.90.032001](https://doi.org/10.7566/JPSJ.90.032001).
- [19] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A quantum approximate optimization algorithm. arXiv:1411.4028v1 [quant-ph], 2014.
- [20] Lukasz Fidkowski. Entanglement spectrum of topological insulators and superconductors. *Phys. Rev. Lett.*, 104:130502, Apr 2010. DOI: [10.1103/PhysRevLett.104.130502](https://doi.org/10.1103/PhysRevLett.104.130502).
- [21] Michael Foss-Feig, David Hayes, Joan M Dreiling, Caroline Figgatt, John P Gaebler, Steven A Moses, Juan M Pino, and Andrew C Potter. Holographic quantum algorithms for simulating correlated spin systems. *Phys. Rev. Research*, 3(3):033002, 2021. DOI: [10.1103/PhysRevResearch.3.033002](https://doi.org/10.1103/PhysRevResearch.3.033002).
- [22] Steph Foulds, Viv Kendon, and Tim Spiller. The controlled SWAP test for determining quantum entanglement. *Quantum Science and Technology*, 6(3):035002, 2021. DOI: [10.1088/2058-9565/abe458](https://doi.org/10.1088/2058-9565/abe458).
- [23] Juan Carlos Garcia-Escartin and Pedro Chamorro-Posada. SWAP test and Hong-Ou-Mandel effect are equivalent. *Phys. Rev. A*, 87:052330, May 2013. DOI: [10.1103/PhysRevA.87.052330](https://doi.org/10.1103/PhysRevA.87.052330).
- [24] K. Geerlings, Z. Leghtas, I.M. Pop, S. Shankar, L. Frunzio, R.J. Schoelkopf, M. Mirrahimi, and M.H. Devoret. Demonstrating a driven reset protocol for a superconducting qubit. *Phys. Rev. Lett.*, 110:120501, Mar 2013. DOI: [10.1103/PhysRevLett.110.120501](https://doi.org/10.1103/PhysRevLett.110.120501).
- [25] Daniel Gottesman and Isaac Chuang. Quantum digital signatures. arXiv:quant-ph/0105032v2, 2001.
- [26] Maria Hermanns. Entanglement in topological systems. arXiv:1702.01525v1 [cond-mat.str-el], 2017.
- [27] William Huggins, Piyush Patil, Bradley Mitchell, K Birgitta Whaley, and E Miles Stoudenmire. Towards quantum machine learning with tensor networks. *Quantum Science and Technology*, 4(2):024001, Jan 2019. DOI: [10.1088/2058-9565/aaea94](https://doi.org/10.1088/2058-9565/aaea94).
- [28] Rajibul Islam, Ruichao Ma, Philipp M. Preiss, M. Eric Tai, Alexander Lukin, Matthew Rispoli, and Markus Greiner. Measuring entanglement entropy in a quantum many-body system. *Nature*, 528(7580):77–83, Dec 2015. DOI: [10.1038/nature15750](https://doi.org/10.1038/nature15750).
- [29] Sonika Johri, Damian S. Steiger, and Matthias Troyer. Entanglement spectroscopy on a quantum computer. *Phys. Rev. B*, 96:195136, Nov 2017. DOI: [10.1103/PhysRevB.96.195136](https://doi.org/10.1103/PhysRevB.96.195136).
- [30] Christian Kokail, Rick van Bijnen, Andreas Elben, Benoît Vermersch, and Peter Zoller. Entanglement hamiltonian tomography in quantum simulation. *Nature Physics*, pages 936–942, 2021. DOI: [10.1038/s41567-021-01260-w](https://doi.org/10.1038/s41567-021-01260-w).
- [31] Hui Li and F.D.M. Haldane. Entanglement spectrum as a generalization of entanglement entropy: Identification of topological order in non-Abelian fractional quantum Hall effect states. *Phys. Rev. Lett.*, 101:010504, Jul 2008. DOI: [10.1103/PhysRevLett.101.010504](https://doi.org/10.1103/PhysRevLett.101.010504).

- [32] Ying Li and Simon C Benjamin. Efficient variational quantum simulator incorporating active error minimization. *Physical Review X*, 7(2):021050, 2017. DOI: [10.1103/PhysRevX.7.021050](https://doi.org/10.1103/PhysRevX.7.021050).
- [33] N.M. Linke, S. Johri, C. Figgatt, K.A. Landsman, A.Y. Matsuura, and C. Monroe. Measuring the Rényi entropy of a two-site Fermi-Hubbard model on a trapped ion quantum computer. *Phys. Rev. A*, 98:052334, Nov 2018. DOI: [10.1103/PhysRevA.98.052334](https://doi.org/10.1103/PhysRevA.98.052334).
- [34] Jin-Guo Liu, Yi-Hong Zhang, Yuan Wan, and Lei Wang. Variational quantum eigensolver with fewer qubits. *Phys. Rev. Research*, 1:023025, Sep 2019. DOI: [10.1103/PhysRevResearch.1.023025](https://doi.org/10.1103/PhysRevResearch.1.023025).
- [35] Seth Lloyd, Masoud Mohseni, and Patrick Rebentrost. Quantum principal component analysis. *Nature Physics*, 10(9):631–633, 2014. DOI: [10.1038/nphys3029](https://doi.org/10.1038/nphys3029).
- [36] P. Magnard, P. Kurpiers, B. Royer, T. Walter, J.-C. Besse, S. Gasparinetti, M. Pechal, J. Heinsoo, S. Storz, A. Blais, and A. Wallraff. Fast and unconditional all-microwave reset of a superconducting qubit. *Phys. Rev. Lett.*, 121:060502, Aug 2018. DOI: [10.1103/PhysRevLett.121.060502](https://doi.org/10.1103/PhysRevLett.121.060502).
- [37] Jarrod R McClean, Jonathan Romero, Ryan Babush, and Alán Aspuru-Guzik. The theory of variational hybrid quantum-classical algorithms. *New Journal of Physics*, 18(2):023023, feb 2016. DOI: [10.1088/1367-2630/18/2/023023](https://doi.org/10.1088/1367-2630/18/2/023023).
- [38] D.T. McClure, Hanhee Paik, L.S. Bishop, M. Steffen, Jerry M. Chow, and Jay M. Gambetta. Rapid driven reset of a qubit readout resonator. *Phys. Rev. Applied*, 5:011001, Jan 2016. DOI: [10.1103/PhysRevApplied.5.011001](https://doi.org/10.1103/PhysRevApplied.5.011001).
- [39] Tianyi Peng, Aram W Harrow, Maris Ozols, and Xiaodi Wu. Simulating large quantum circuits on a small quantum computer. *Phys. Rev. Lett.*, 125(15):150504, 2020. DOI: [10.1103/PhysRevLett.125.150504](https://doi.org/10.1103/PhysRevLett.125.150504).
- [40] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J Love, Alán Aspuru-Guzik, and Jeremy L O’Brien. A variational eigenvalue solver on a photonic quantum processor. *Nature Communications*, 5(1):4213, Jul 2014. DOI: [10.1038/ncomms5213](https://doi.org/10.1038/ncomms5213).
- [41] J.M. Pino, J.M. Dreiling, C. Figgatt, J.P. Gaebler, S.A. Moses, M.S. Allman, C.H. Baldwin, M. Foss-Feig, D. Hayes, K. Mayer, C. Ryan-Anderson, and B. Neyenhuis. Demonstration of the trapped-ion quantum CCD computer architecture. *Nature*, 592(7853):209–213, 2021. DOI: [10.1038/s41586-021-03318-4](https://doi.org/10.1038/s41586-021-03318-4).
- [42] John Preskill. Quantum Computing in the NISQ era and beyond. *Quantum*, 2:79, August 2018. DOI: [10.22331/q-2018-08-06-79](https://doi.org/10.22331/q-2018-08-06-79).
- [43] Emil Prodan, Taylor L. Hughes, and B. Andrei Bernevig. Entanglement spectrum of a disordered topological Chern insulator. *Phys. Rev. Lett.*, 105:115501, Sep 2010. DOI: [10.1103/PhysRevLett.105.115501](https://doi.org/10.1103/PhysRevLett.105.115501).
- [44] Arthur G Rattew, Yue Sun, Pierre Minssen, and Marco Pistoia. Quantum simulation of Galton machines using mid-circuit measurement and reuse. arXiv:2009.06601v1 [quant-ph], 2020.
- [45] M.D. Reed, B.R. Johnson, A.A. Houck, L. DiCarlo, J.M. Chow, D.I. Schuster, L. Frunzio, and R.J. Schoelkopf. Fast reset and suppressing spontaneous emission of a superconducting qubit. *Applied Physics Letters*, 96(20):203110, 2010. DOI: [10.1063/1.3435463](https://doi.org/10.1063/1.3435463).
- [46] Diego Ristè, Luke CG Govia, Brian Donovan, Spencer D Fallek, William D Kalfus, Markus Brink, Nicholas T Bronn, and Thomas A Ohki. Real-time processing of stabilizer measurements in a bit-flip code. *npj Quantum Information*, 6(1):1–6, 2020. DOI: [10.1038/s41534-020-00304-y](https://doi.org/10.1038/s41534-020-00304-y).
- [47] Benjamin Schumacher and Michael D. Westmoreland. Locality and information transfer in quantum operations. *Quantum Information Processing*, 4(1):13–34, Feb 2005. DOI: [10.1007/s11128-004-3193-y](https://doi.org/10.1007/s11128-004-3193-y).
- [48] Benjamin Schumacher and Michael D. Westmoreland. Isolation and information flow in quantum dynamics. *Foundations of Physics*, 42(7):926–931, 2012. DOI: [10.1007/s10701-012-9651-y](https://doi.org/10.1007/s10701-012-9651-y).
- [49] H. Francis Song, Stephan Rachel, Christian Flindt, Israel Klich, Nicolas Lafflorencie, and Karyn Le Hur. Bipartite fluctuations as a probe of many-body entanglement. *Phys. Rev. B*, 85:035409, Jan 2012. DOI: [10.1103/PhysRevB.85.035409](https://doi.org/10.1103/PhysRevB.85.035409).
- [50] Yiğit Subaşı, Lukasz Cincio, and Patrick J Coles. Entanglement spectroscopy with a depth-two quantum circuit. *Journal of Physics A: Mathematical and Theoretical*, 52(4):044001, 2019. DOI: [10.1088/1751-8121/aaf54d](https://doi.org/10.1088/1751-8121/aaf54d).
- [51] Sathyawageeswar Subramanian and Min-Hsiu Hsieh. Quantum algorithm for estimating Rényi entropies of quantum states. arXiv:1908.05251v1 [quant-ph], 2019.
- [52] Brian Swingle and T. Senthil. Geometric proof of the equality between entanglement and edge spectra. *Phys. Rev. B*, 86:045117, Jul 2012. DOI: [10.1103/PhysRevB.86.045117](https://doi.org/10.1103/PhysRevB.86.045117).
- [53] Kristan Temme, Sergey Bravyi, and Jay M Gambetta. Error mitigation for short-depth quantum circuits. *Physical review letters*, 119(18):180509, 2017. DOI: [10.1103/PhysRevLett.119.180509](https://doi.org/10.1103/PhysRevLett.119.180509).