

Preparing quantum many-body scar states on quantum computers

Erik J. Gustafson^{1,2}, Andy C. Y. Li^{1,2}, Abid Khan^{1,3,4,5}, Joonho Kim^{1,6}, Doğa Murat Kürkçüoğlu^{1,2}, M. Sohaib Alam^{1,4,5}, Peter P. Orth^{1,7,8,9}, Armin Rahmani¹⁰, and Thomas Iadecola^{1,7,8}

¹Superconducting Quantum Materials and Systems Center (SQMS), Fermi National Accelerator Laboratory, Batavia, IL 60510, USA

²Fermi National Accelerator Laboratory, Batavia, IL, 60510, USA

³Department of Physics, University of Illinois Urbana-Champaign, Urbana, IL, United States 61801

⁴USRA Research Institute for Advanced Computer Science (RIACS), Mountain View, CA, 94043, USA

⁵Quantum Artificial Intelligence Laboratory (QuAIL), NASA Ames Research Center, Moffett Field, CA, 94035, USA

⁶Rigetti Computing, Berkeley, CA, 94710, USA

⁷Department of Physics and Astronomy, Iowa State University, Ames, IA 50011, USA

⁸Ames National Laboratory, Ames, IA 50011, USA

⁹Department of Physics, Saarland University, 66123 Saarbrücken, Germany

¹⁰Department of Physics and Astronomy and Advanced Materials Science and Engineering Center, Western Washington University, Bellingham, WA 98225, USA

Quantum many-body scar states are highly excited eigenstates of many-body systems that exhibit atypical entanglement and correlation properties relative to typical eigenstates at the same energy density. Scar states also give rise to long-lived coherent dynamics when the system is prepared in a special initial state having finite overlap with them. Many models with exact scar states have been constructed, but the fate of scarred eigenstates and dynamics when these models are perturbed is difficult to study with classical computational techniques. In this work, we propose state preparation protocols that enable the use of quantum computers to study this question. We present protocols both for individual scar states in a particular model, as well as superpositions of them that give rise to coherent dynamics. For superpositions of scar states, we present both a system-size-linear depth unitary and a finite-depth nonunitary state preparation protocol, the latter of which uses measurement and postselection to reduce the circuit depth. For individual scarred eigenstates, we formulate an exact state preparation approach based on matrix product states that yields quasipolynomial-depth circuits, as well as a variational approach with a polynomial-depth ansatz circuit. We also provide proof of principle state-preparation demonstrations on superconducting quantum hardware.

1 Introduction

Recent decades have seen remarkable advances in our understanding of how quantum statistical mechanics can emerge from isolated strongly-interacting quantum mechanical systems. One of the most fundamental of these advances is the so-called eigenstate thermalization hypothesis (ETH) [1, 2, 3, 4], which posits that individual quantum mechanical eigenstates at finite energy density become locally equivalent in the thermodynamic limit to equilibrium Gibbs ensembles at a corresponding temperature. Such eigenstates also govern the approach to this local equilibrium

Thomas Iadecola: iadecola@iastate.edu

under unitary dynamics [5]. In parallel with these developments, there have been enormous strides towards realizing quantum technologies based on coherent quantum systems that are approximately isolated on experimentally relevant time scales. This gives rise to the possibility of testing the ETH and the related phenomenon of quantum information scrambling experimentally using analog quantum simulators [6, 7, 8, 9, 10] and digital quantum computers [11].

Along with this progress in understanding the ETH and quantum thermalization has come the realization that there are quantum systems that do not thermalize under certain conditions. Two notable examples include integrable [12, 13] and many-body localized systems [14, 15, 16], where an extensive number of conserved quantities preclude the possibility of reaching a conventional locally thermalized state. An alternative means of avoiding thermalization is provided by quantum many-body scars (QMBS), a phenomenon whereby nonintegrable quantum systems exhibit a set of rare finite-energy-density eigenstates that do not satisfy the ETH [17, 18, 19]. Such eigenstates have been found in a variety of contexts, including the Affleck-Kennedy-Lieb-Tasaki (AKLT) model [20, 21], ensembles of Rydberg atoms [22, 23, 24, 25], and various other interacting spin [26, 27, 28, 29, 30, 31, 32, 33, 34], bosonic [35, 36], and fermionic [37, 38, 39, 40, 41, 42, 43] models. These eigenstates can also give rise to coherent periodic dynamics from certain initial states, which has allowed QMBS to be observed in quantum simulation experiments [22, 25, 36, 44].

Substantial effort has been devoted to formulating mathematical criteria for the emergence of QMBS [28, 29, 31, 34, 35, 45], and all examples so far require some degree of fine tuning. Understanding the fate of scarred eigenstates and the associated coherent dynamics under perturbations is thus an important research direction, but relatively little progress has been made so far. Ref. [46] found a general lower bound on the thermalization timescale for a scarred state in the presence of a perturbation, $t_* = O(\epsilon^{-1/(d+1)})$, where ϵ is the perturbation strength and d is the system's spatial dimension. However, this bound relies only on the underlying Hamiltonian's spatial locality, and numerical studies of specific models in, e.g., Refs. [46, 34] found lifetimes that substantially exceed this bound.

Studying the lifetime of QMBS under perturbations is challenging because, although scarred eigenstates often have modest entanglement and can be represented efficiently in terms of matrix product states (MPSs) [21, 30, 47], perturbations couple them to states nearby in energy which typically have extensive volume-law entanglement. Analytical perturbation theory is impractical here owing both to the complexity of these highly excited eigenstates and to the exponentially large density of states at finite energy density. Numerical methods to directly evaluate the real time evolution of a scarred state under perturbations also encounter challenges. Exact methods are limited to small system sizes, while approximate tensor network methods [48, 49] are generally limited to early times [50].

A natural question is whether one could exploit quantum computers to investigate the behavior of scarred states under perturbations. It has long been known that quantum computers can evaluate real-time dynamics efficiently [51]; indeed, simulating quantum dynamics is one of the leading candidates for near-term practical quantum advantage [52, 53]. However, even if we assume access to a noiseless quantum computer that can perform accurate time evolution, we are still left with the challenge of state preparation: how can we prepare scarred states on a digital quantum computer, and what are the resources required to do so? This is the subject of the present work.

There are two general state preparation tasks that one faces in this context, depending on whether one wants to investigate the lifetime of scarred *dynamics* or scarred *eigenstates*. In the first case, the aim is to time-evolve a *superposition* of scarred states that exhibits periodic dynamics in the unperturbed limit and extract the lifetime of the observed oscillations. In many cases of interest, a product state is sufficient for these purposes and the state preparation is therefore trivial [22, 23, 26, 25, 36, 44, 34, 35, 54]. However, in other cases, the simplest superposition of scar states is area-law entangled and has a nontrivial MPS representation with a finite correlation length [55, 27, 56], and here some thought must be put into the most efficient method to prepare such states.

In the second case, the aim is to prepare an *individual* scarred eigenstate and evolve it under the perturbed Hamiltonian. In many cases of interest, scarred eigenstates have entanglement scaling logarithmically with system size [57, 21, 24, 26, 27, 55, 37, 38], similar to critical states in one dimension that are described by conformal field theories. The entanglement content of these states is modest compared to typical volume-law states at the same energy density, suggesting the

possibility of an efficient state preparation circuit. In fact, in many of the simplest examples of QMBS the scarred eigenstates are equivalent to Dicke states [58, 26, 37, 38, 34, 35], for which a polynomial-depth state preparation circuit is known [59]. However, in other examples, the states of interest can be viewed as projected or finite-momentum Dicke states [27, 55, 32], about which much less is known. Thus in general it is an open question what are the minimal quantum resources needed to prepare scarred eigenstates.

In this work we address both of the above cases for a specific model with QMBS. In Sec. 2 we define the model and state preparation tasks in more detail. In Sec. 3 we consider the problem of preparing a particular class of superpositions of scarred eigenstates that can be realized as a one-parameter family of MPSs with bond dimension $\chi = 2$. We consider two related approaches. First, we explicitly construct a linear-depth circuit that prepares the desired state with perfect fidelity. Second, we discuss a probabilistic method that prepares the desired states in constant depth using measurements and postselection. The latter method has a postselection success probability that decays exponentially with system size, albeit with a base that can be tuned by adjusting the circuit depth. This allows for a flexible tradeoff between circuit depth and success probability that is advantageous for implementation on near-term quantum hardware. In Sec. 4, we discuss two strategies for the preparation of individual scarred eigenstates. In the first strategy, we identify MPS representations for the scarred eigenstates and convert these to quantum circuits that prepare the states with perfect fidelity in quasipolynomial depth. In the second, we propose a polynomial-depth variational ansatz that we show numerically captures the scarred eigenstates with at least 99% fidelity at numerically accessible system sizes. In both Secs. 3 and 4, we provide proof-of-concept demonstrations of the respective state preparation tasks on Rigetti and IBM quantum processing units (QPUs), often with some additional simplifications in order to obtain better results on hardware. Finally, we provide a conclusion and outlook in Sec. 5. In the appendices, we discuss an alternative state preparation protocol based on adiabatic evolution (Appendix A), details about state preparation circuit derivations (Appendix B), and additional QPU results on IBM hardware (Appendix C).

2 Model and State Preparation Tasks

In this work, we select a particular reference model with QMBS to exemplify our state preparation techniques for scarred eigenstates and their superpositions. We consider the spin-1/2 model defined in Ref. [27], whose Hamiltonian reads

$$H_0 = \lambda \sum_{i=2}^{N-1} (X_i - Z_{i-1} X_i Z_{i+1}) + \Delta \sum_{i=1}^N Z_i + J \sum_{i=1}^{N-1} Z_i Z_{i+1}. \quad (1)$$

This Hamiltonian acts on a chain of N qubits with open boundary conditions, and each qubit is equipped with Pauli operators X_i, Y_i , and Z_i . Note that H_0 commutes with the operators Z_1 and Z_N —the Z -basis projections of the edge qubits are therefore conserved quantities. H_0 also conserves the number of Ising domain walls, measured by the operator $n_{\text{DW}} = \sum_{i=1}^{N-1} (1 - Z_i Z_{i+1})/2$. The λ term in Eq. (1) can be viewed as a kinetic term for domain walls, while the Ising interaction J serves as a chemical potential for the domain walls. The Δ term induces nonlocal interactions between domain walls and makes the model nonintegrable. It is interesting to note that this model is dual to a $\mathbb{Z}_2$ lattice gauge theory coupled to fermionic matter [60] and can be realized in Rydberg atom quantum simulators in the antiblockade regime [61].

The model (1) has two towers of QMBS states related by the global spin-flip operation $G = \prod_{i=1}^N X_i$. The first tower is given by

$$|\mathcal{S}_k\rangle = \frac{1}{k! \sqrt{\mathcal{N}(N, k)}} (Q^\dagger)^k |\Omega\rangle, \quad (2)$$

where $\mathcal{N}(N, k) = \binom{N-k-1}{k}$, $|\Omega\rangle = |0 \dots 0\rangle$, and $k = 0, \dots, N/2 - 1$ (we take N even for simplicity). The raising operator for this tower of states is given by

$$Q^\dagger = \sum_{i=2}^{N-1} (-1)^i P_{i-1} \sigma_i^+ P_{i+1}, \quad (3)$$

where $P_j = |0\rangle_j \langle 0|_j = (1 + Z_j)/2$ and $\sigma_j^\pm = (X_j \mp iY_j)/2$. The states $|\mathcal{S}_k\rangle$ are eigenstates of H_0 with energies $E_k = \Delta N + J(N-1) - (2\Delta + 4J)k$. The second tower of states

$$|\mathcal{S}'_k\rangle = G|\mathcal{S}_k\rangle, \quad k = 0, \dots, \frac{N}{2} - 1 \quad (4)$$

has energies $E'_k = -\Delta N + J(N-1) + (2\Delta - 4J)k$. Both towers of states therefore have an extensive energy bandwidth, so that typical states in each tower correspond to highly excited states of the model H_0 . Nevertheless, states in either tower for which k/N is finite as $N \rightarrow \infty$ have bipartite entanglement entropy scaling as $\ln(N)$, in contrast to the volume-law entanglement entropy of typical eigenstates at the same energy density [27]. In the remainder of the paper we will restrict our attention to the tower $\{|\mathcal{S}_k\rangle\}$, but all results we obtain for this tower hold equally well for the tower $\{|\mathcal{S}'_k\rangle\}$.

Dynamical consequences of these scar states can be observed by evolving the system from a suitable family of initial states with support on the towers of interest. In Ref. [27], it was shown that the following family of states parameterized by $\xi \in \mathbb{C}$ is exclusively supported on the tower $\{|\mathcal{S}_k\rangle\}$:

$$|\xi\rangle = \frac{1}{\sqrt{Z(|\xi|^2)}} \mathcal{P}_{\text{fib}} \prod_{i=2}^{N-1} [1 + (-1)^i \xi \sigma_i^+] |\Omega\rangle, \quad (5)$$

where the projection operator

$$\mathcal{P}_{\text{fib}} = \prod_{i=1}^{N-1} (1 - P'_i P'_{i+1}), \quad P'_i = 1 - P_i \quad (6)$$

excludes any computational basis state in Eq. (5) containing the local configuration $|1\rangle_i |1\rangle_{i+1}$, and where the normalization factor

$$Z(|\xi|^2) = \sum_{k=0}^{N/2-1} |\xi|^{2k} \mathcal{N}(N, k). \quad (7)$$

The state $|\xi\rangle$ can be decomposed onto the tower $\{|\mathcal{S}_k\rangle\}$ as follows:

$$|\xi\rangle = \sum_{k=0}^{N/2-1} \xi^k \sqrt{\frac{\mathcal{N}(N, k)}{Z(|\xi|^2)}} |\mathcal{S}_k\rangle. \quad (8)$$

As such, for any value of ξ , evolving the initial state $|\xi\rangle$ under H_0 yields perfect periodic revivals of $|\xi\rangle$ with period $T = \pi/(\Delta + 2J)$ set by the energy spacing between consecutive states in the tower. Unlike typical states in the tower $\{|\mathcal{S}_k\rangle\}$, the state $|\xi\rangle$ is area-law entangled. In fact, it can be written as an MPS with bond dimension $\chi = 2$ [27], so its bipartite entanglement entropy is upper bounded by $\ln 2$.

The state $|\xi\rangle$ has connections to several interesting problems in condensed matter and atomic physics. For instance, it is unitarily equivalent to a state found in Ref. [62] to be a good approximation to the ground state of a system of Rydberg atoms in the so-called Rydberg blockade regime [63, 64], where the computational basis states $|0\rangle_i$ and $|1\rangle_i$ correspond to atom i being in its ground state or a highly excited Rydberg state, respectively. Due to strong interactions, such systems are subjected to an energetic penalty for having two excited atoms next to one another. This constraint, sometimes called the ‘‘Fibonacci constraint’’ because the number of states of m qubits that satisfy it is given by F_{m+2} (where F_ℓ is the ℓ -th Fibonacci number), is implemented by the projector $\mathcal{P}_{\text{fib}}$ in Eq. (5). Intriguingly, this constraint also emerges in theoretical descriptions of the $\nu = 1/3$ Laughlin fractional quantum Hall (FQH) state in a particular quasi-1D limit [65, 66, 67, 68]. In fact, the ground state of the system in this limit is unitarily equivalent to $|\xi\rangle$ for a particular choice of the parameter ξ [65]. Thus, our state preparation results for $|\xi\rangle$ will also be applicable to the seemingly disparate settings of Rydberg-atom quantum simulators and FQH liquids.

Studying the stability of scarred eigenstates and dynamics on quantum computers requires algorithms for high-fidelity preparation of the states $|\mathcal{S}_k\rangle$ and $|\xi\rangle$, respectively. This paper provides a survey of approaches to both state preparation tasks and a snapshot of their feasibility with current quantum hardware. We first address the preparation of the superposition state $|\xi\rangle$ in Sec. 3, before moving onto the preparation of the scar states $|\mathcal{S}_k\rangle$ in Sec. 4.

The alternating signs $(-1)^i$ appearing in Eqs. (3) and (5) can be removed by the simple unitary circuit $\prod_{i \text{ odd}} Z_i$. Therefore, it is useful to define the states

$$|\tilde{\xi}\rangle = \left(\prod_{i \text{ odd}} Z_i \right) |\xi\rangle \quad (9a)$$

and

$$|\tilde{\mathcal{S}}_k\rangle = \left(\prod_{i \text{ odd}} Z_i \right) |\mathcal{S}_k\rangle, \quad (9b)$$

which are now equal-amplitude and equal-sign superpositions of computational basis states. We will at times find it more convenient to work with these “tilde” states than with the original states.

Before proceeding, we highlight that preparing scarred eigenstates and superpositions in our reference model (1) is particularly challenging due to the Fibonacci constraint. In many scarred models of interest, the scarred eigenstates can be realized as Dicke states [58, 26, 37, 38, 34, 35], for which a polynomial-depth circuit construction is known [59]. Likewise, low-entanglement superpositions of scar states in these models can be realized as product states, for which state preparation is trivial. Our motivation in selecting this more complicated model is that kinetic constraints and the more elaborate entanglement structures they enable are present in a variety of other models, including the AKLT model [20, 21, 56, 55]. Thus, grappling with such constraints and assessing their impact on state preparation complexity is necessary in order to study the full variety of QMBS on quantum computers.

3 Preparing the State $|\xi\rangle$

There are several possible approaches to preparing the superposition state $|\xi\rangle$ from Eq. (5). In Ref. [27] it was shown that $|\xi\rangle$ is the unique ground state of a local parent Hamiltonian with finite correlation length (see also Ref. [62]). Thus one approach is to prepare $|\xi\rangle$ adiabatically using an appropriate parameter sweep from a zero-correlation-length paramagnetic Hamiltonian. We investigate this approach in App. A, where we find evidence that the gap of the interpolating Hamiltonian does not close with increasing N and the state $|\xi\rangle$ is a suitable candidate for adiabatic state preparation. In practice, adiabatic state preparation on a digital quantum computer suffers from Trotter error, even for a fully gapped interpolation Hamiltonian and assuming perfect implementation of the Trotter circuit. For this reason a finite depth circuit always incurs finite error. This motivates the consideration of alternative state preparation strategies.

In Sec. 3.1, we demonstrate that perfect state preparation can be achieved with a circuit of depth $O(N)$. Sec. 3.2 shows that a stochastic strategy [69, 70, 71, 72, 73, 74] using measurements and post-selection can reduce the circuit depth to a constant at the price of an exponential post-selection overhead. Finally, in Sec. 3.3, we benchmark both state preparation strategies on Rigetti QPUs. We show additional QPU results from IBM hardware in Appendix C.

3.1 Linear-Depth Unitary Circuit

To facilitate our discussion, we rewrite the superposition state $|\xi\rangle$ with open-boundary conditions on N qubits [Eq. (5)] as

$$|\xi\rangle = |0\rangle \otimes |\xi; N-2\rangle \otimes |0\rangle. \quad (10)$$

In this section, we show that the state $|\xi; m\rangle$ as well as its counterpart $|\tilde{\xi}; m\rangle$ with alternating phases removed per Eqs. (9) can be prepared in linear depth $O(m)$ by a unitary circuit $\mathcal{U}_\xi(m)$




Figure 1: Linear-depth circuit showing the preparation gate $\mathcal{U}_\xi(m)$ to prepare $|\xi; m\rangle$ from the zero state $|\Omega\rangle$. circuit consists of $(m - 1)$ control-rotation gates and one rotation gate. The rotation angles θ_j are given in Eq. (12).

consisting of $(m - 1)$ controlled Y -rotation gates and one Y -rotation gate. We note that a similar linear depth circuit was obtained in the FQH context in Ref. [67], but that the nonunitary approach explored in Sec. 3.2 has not been discussed in this context.

To understand this preparation circuit, we recall that $|\xi\rangle$ or $|\xi; m\rangle$ is the superposition of all computational basis states excluding any local configuration $|1\rangle_{j-1}|1\rangle_j$ with the weights of each state determined by the parameter ξ . Starting with all qubits in the $|0\rangle$ state, the absence of $|1\rangle_{j-1}|1\rangle_j$ pairs is guaranteed under a sequence of controlled Y -rotations $C_0 R_{Y,j-1,j}(\theta_j)$ by an angle θ_j targeted on the j -th qubit that is triggered only by $|0\rangle$ of the $(j - 1)$ -th qubit. For a set of properly chosen rotation angles θ_j , $|\xi; m\rangle$ can be prepared by this preparation circuit:

$$|\xi; m\rangle = \mathcal{U}_\xi(m) |\Omega\rangle, \quad \mathcal{U}_\xi(m) = \left[\prod_{j=2}^m C_0 R_{Y,j-1,j}(\theta_j) \right] R_{Y,1}(\theta_1). \quad (11)$$

The corresponding circuit diagram of the preparation gate $\mathcal{U}_\xi(m)$ is shown in Fig. 1. The angles are determined as a function of ξ according to the following recursion relation:

$$\theta_j = 2 \arg \left[1 + i \frac{(-1)^{j+1} \xi}{\phi_{j+1}} \right], \quad \phi_j = \sqrt{1 + \left| \frac{(-1)^{j+1} \xi}{\phi_{j+1}} \right|^2} \quad \text{and} \quad \phi_{m+1} = 1. \quad (12)$$

The preparation circuit $\mathcal{U}_{\tilde{\xi}}(m)$ for the state $|\tilde{\xi}; m\rangle$ is obtained by simply removing the alternating phase factors $(-1)^{j+1}$ from the above definition. The proof of the preparation circuit and the above recursion relation for the rotation angles is provided in Appendix B.

3.2 Probabilistic Constant-Depth Circuit with Postselection

While the linear-depth circuit explored in the previous section is sufficient for state preparation, it may be challenging to implement on near-term devices for large N owing to the linear circuit depth. Here, we show that it is possible to prepare the equal-amplitude superposition state $|\tilde{\xi}; N-2\rangle$ stochastically in constant depth using measurements and postselection. The idea is to prepare k m -site blocks in the state $|\tilde{\xi}; m\rangle$, which can be achieved in depth $O(m)$ using the circuit $\mathcal{U}_{\tilde{\xi}}(m)$ [Eq. (11) with alternating phases removed in the recursive formula Eq. (12)]. The resulting state $|\tilde{\xi}; m\rangle^{\otimes k}$ obeys the Fibonacci constraint enforced by the projector $\mathcal{P}_{\text{fib}}$ [Eq. (6)] within each m -site block, but adjacent blocks need not obey the constraint. To “stitch” the adjacent blocks together into a state that globally satisfies the constraint, adjacent m -site blocks are each coupled to an ancilla qubit using an appropriate unitary operation. Measuring the ancilla qubit and postselecting onto an appropriate measurement outcome prepares the desired state $|\tilde{\xi}; km\rangle$, which can then be trivially converted to $|\xi; km\rangle$ using Eq. (9).

To illustrate, let us set $\xi = 1$ for simplicity (this is not strictly necessary but does simplify the analytical expressions for the states and success probabilities). As an example, consider stitching

together two copies of the two-qubit state $|1; 2\rangle$. The two-qubit state is given by

$$|\tilde{1}; 2\rangle = \frac{1}{\sqrt{3}}(|00\rangle + |10\rangle + |01\rangle), \quad (13)$$

and the state we aim to prepare is

$$|\tilde{1}; 4\rangle = \frac{1}{\sqrt{8}}(|0000\rangle + |0001\rangle + |0010\rangle + |0100\rangle + |1000\rangle + |1010\rangle + |0101\rangle + |1001\rangle). \quad (14)$$

We first initialize the system in $|\Omega\rangle = |0000\rangle$ and apply the circuit $\mathfrak{U}_1(2)$ to the two consecutive two-qubit blocks to prepare the state

$$|\tilde{1}; 2\rangle \otimes |\tilde{1}; 2\rangle = \frac{1}{3}(|00\rangle + |10\rangle + |01\rangle) \otimes (|00\rangle + |10\rangle + |01\rangle). \quad (15)$$

There is exactly one configuration in the above superposition that would be projected away by $\mathcal{P}_{\text{fib}}$, namely $|0110\rangle$. Therefore if we apply a Toffoli (CCNOT) gate controlled on qubits 2 and 3 and targeted on an ancilla qubit initialized in the state $|0\rangle_a$ to the above state, we obtain

$$\frac{\sqrt{8}}{3}|\tilde{1}; 4\rangle|0\rangle_a + \frac{1}{3}|0110\rangle|1\rangle_a. \quad (16)$$

Measuring the ancilla qubit in the computational basis, we obtain the desired state $|\tilde{1}; 4\rangle$ whenever the measurement outcome is 0, which occurs with probability 8/9.

The same procedure can be generalized to stitch together k copies of $|\tilde{1}; 2\rangle$ into the state $|\tilde{1}; 2k\rangle$ using $k-1$ ancilla qubits initialized in the $|0\rangle$ state. Each ancilla qubit is coupled to the array of $2k$ primary qubits using a Toffoli gate controlled by the neighboring qubits from consecutive two-site blocks. In this way, the state of each ancilla qubit records whether the Fibonacci constraint is violated for the pair of primary qubits to which it is coupled. The probability that the Fibonacci constraint is satisfied (such that the ancilla register is in the state $|0\dots 0\rangle_a$) is given in terms of Fibonacci numbers F_ℓ by F_{2k+2}/F_4^k , which decays exponentially with k .

The above success probability can be improved by using the same strategy to stitch together larger blocks. For example, suppose we wish to prepare the state $|\tilde{1}; 2m\rangle$ by stitching together two m -site blocks prepared in the state $|\tilde{1}; m\rangle$ using the $O(m)$ -depth circuit $\mathfrak{U}_1(m)$. The stitching can again be achieved by coupling the two neighboring qubits from the two blocks to an ancilla qubit using a Toffoli gate. An illustration of this procedure for general ξ is shown in Fig. 2. The success probability can be obtained by noting that there are F_{2m+2} states of the full $2m$ -qubit system that satisfy the Fibonacci constraint (for which a measurement of the ancilla qubit would yield 0), while the initial tensor product state contains F_{m+2}^2 configurations; the success probability is then F_{2m+2}/F_{m+2}^2 .

Applying the same logic to k blocks of m sites yields a postselection success probability

$$p_{\text{success}}(m, k) = F_{km+2}/F_{m+2}^k, \quad (17)$$

which is plotted for various m against $N = km$ in Fig. 3. Although this expression still decays exponentially with k for fixed m , it grows with m at fixed k . Thus, the exponential sampling overhead can be mitigated by increasing m at the price of increasing the depth of the state preparation circuit. Fig. 3 shows the success probability for a lattice of length N using blocks of length m . On present-day NISQ devices, gate error and qubit decoherence rates are sufficiently high that the reduction of circuit depth at the expense of postselection may be an acceptable tradeoff. We analyze this tradeoff further in Sec. 3.3 when we implement this state preparation protocol on quantum hardware.

It is natural to ask whether a deterministic (i.e., postselection-free) finite-depth state preparation protocol can be formulated in which unwanted ancilla measurement outcomes are corrected by local unitary operations instead of being discarded. Such a scheme can leverage the fact that the local form of the state in the vicinity of an ancilla measurement outcome of 1 is fixed. For example, consider stitching together two states of the form $|\tilde{1}; m\rangle$ with a single ancilla measurement, and suppose the measurement outcome was 1. Then the primary qubit register is in the state

$$|\tilde{1}; m-2\rangle \otimes |0110\rangle \otimes |\tilde{1}; m-2\rangle, \quad (18)$$




Figure 2: Example circuit for stitching two copies of $|\tilde{\xi}; m\rangle$ together into $|\tilde{\xi}; 2m\rangle$ using a single Toffoli gate targeted onto an ancilla, which is then measured and postselected onto the $|0\rangle$ state. The circuit $\mathcal{U}_{\tilde{\xi}}(m)$ is simply Eq. (11) with the alternating signs removed from the angles in Eq. (12).




Figure 3: Success probability of the finite-depth postselection-based protocol for $N = km$ qubits and initial block size m . The inset shows the same data with a logarithmic scale on the y -axis, indicating that the success probability decays exponentially with N with a base set by m ; increasing m increases the success probability by bringing the base closer to 1.

since the state away from the central two sites (which are projected onto $|11\rangle$ by the measurement) obeys the Fibonacci constraint. One can imagine trying to correct the four site block $|0110\rangle$ using a unitary circuit controlled by the states of the two adjacent qubits from the $(m-2)$ -site blocks to its left and right. However, the target states in the four cases (labeled by the four possible states of the adjacent qubits) have different normalization factors due to the Fibonacci constraint; therefore, the state of the qubit chain after feedback will not be an equal-amplitude superposition state. This situation should be contrasted with a recently proposed finite-depth deterministic scheme to prepare the AKLT ground state [74], which is able to correct undesired measurement outcomes with a local circuit. However, the correction operation designed in that work makes use of the fact that the AKLT state is a symmetry-protected topological state [75, 76], which is a property that is not shared by the state $|\xi\rangle$. We leave the question of whether an undesired measurement outcome in our probabilistic protocol can be corrected by a finite-depth circuit for future work.

3.3 QPU Results

We now implement the $|\xi\rangle$ state preparation protocols on Rigetti's Aspen QPUs (shown in Fig. 4) and on IBM hardware. We discuss the execution results of the linear-depth circuit in Sec. 3.1 and its probabilistic post-selection variant in Sec. 3.2. We set $\xi = 1$ for numerical evaluations, in which case the Pauli rotation angles in Fig. 1 can be summarized as

$$\theta_{m-i+1} = 2 \tan^{-1}(\sqrt{F_{i+1}/F_i}) \quad (19)$$

with Fibonacci coefficients $F_i = (1, 1, 2, 3, 5, 8, \dots)_i$. The `quilc` compiler generates an optimized program for the instruction set architecture of QPU chips, compiling all logical gates into the




Figure 4: Aspen M-2 and M-3 device layout graphs, with each node representing a superconducting qubit and each edge indicating connectivity via two-qubit gates. Circled nodes denote qubits used in the QPU experiments in Sec. 3.3 and 4.3.

following group of native gates:

$$\text{RZ}(\theta), \text{RX}(\pi/2), \text{RX}(\pi), \text{CPHASE}(\theta), \text{CZ}, \text{XY}(\theta).$$

Results obtained from the NISQ hardware are typically prone to multiple sources of error, leading to deviations from ideal unitary calculations. Here we quantify these errors by estimating two figures of merit: first, the Bhattacharyya distance

$$D(p, q) = -\ln \sum_x \sqrt{p(x)q(x)} \quad (20)$$

that captures the divergence between the exact and the observed distributions, $p(x)$ and $q(x)$, of the measured bitstrings x . Second, we use the expectation value of the parent Hamiltonian of the state $|\xi\rangle$, namely [27]

$$H_\xi = \sum_{i=2}^{N-1} P_{i-1} [\xi^{-1} P'_i + \xi P_i - (-1)^i X_i] P_{i+1}, \quad (21)$$

that vanishes $\langle H_\xi \rangle = 0$ in the case of ideal noiseless preparation of the state $|\xi\rangle$ ¹. We collect 10 observed samples of these quantities, along with the success probability of the post-selection protocol, for a range of different system sizes N and block sizes m . Each sample measurement is made with 10^4 shots for the energy and the success probability, and 10^5 shots for the Bhattacharyya distance. The individual samples are depicted as round dots in Fig. 5 and 6, and their average values are connected with solid lines. See the captions of Fig. 5 and 6 for the specification of used device nodes.

When the qubits are measured immediately after state preparation, without performing further unitary operations for, e.g., time evolution, the two-step process of applying Toffoli gates and post-selecting on ancilla bits can be reduced to the classical post-selection of N -qubit output bitstrings. Such replacement reduces the room for error, since Toffoli (CCNOT) is a non-native gate that must be compiled into multiple imperfect two-qubit gates (e.g., CPHASE, CZ, XY) before running on QPUs². It also “unstitches” the whole circuit back into $\frac{N-2}{m}$ decoupled blocks of $\mathcal{U}_\xi(m)$, which introduces certain computational advantages. For example, the asynchronous execution of fragmented circuits allows us to simulate systems larger than the hardware size, or to reduce the overall error by avoiding the usage of low-fidelity qubits, at an exponential overhead of classical post-processing [82, 83]. With this in mind, we repeatedly use the same m -qubit sublattice that exhibit the best average readout and gate fidelities, and then classically combine the measurement outcomes with the Fibonacci constraint to simulate $N = \alpha m + 2$ qubit observables ($\alpha \in \mathbb{N}$).

Fig. 5 displays the results of the $N = 14$ state preparation experiment with different block sizes m , to which the circuit $\mathcal{U}_\xi(m)$ is applied. We first discuss the raw values without error

¹We note that $|\xi\rangle$ is the unique zero-energy eigenvector of H_ξ in the Fibonacci Hilbert space, and that the spectrum of H_ξ is positive semi-definite.

²An alternative to decomposing a Toffoli gate into native two-qubit gates is to tune up a two-qutrit gate pulse that maps $|11\rangle$ to $|02\rangle$. The XY_{02} gate can serve as a building block for many doubly-controlled gates [80]. See [81] for the implementation and benchmarking of three-qubit gates, e.g., Toffoli, with three-level control pulses.




Figure 5: The QPU experiment results of the $|\xi = 1\rangle$ state preparation circuit on Aspen-M devices, which uses the probabilistic post-selection protocol that combines multiple m -qubit blocks $\mathcal{U}_\xi(m)$ to build the $N = 14$ state. The indices of the participating Aspen qubits are: $\{105, 104\}$ (M-3) for $m = 2$, $\{106, 107, 100, 101\}$ (M-3) for $m = 4$, $\{16, 17, 10, 11, 26, 27\}$ (M-2) for $m = 6$. Four distinct sets of measured values are presented: EXACT, RAW, REM, and REM+SYMM+RC, where REM, SYMM, and RC are acronyms for readout error mitigation, readout symmetrization [77], and randomized compilation [78], respectively. The RAW samples are derived directly from the post-selection protocol using unprocessed bitstrings. On the other hand, the REM samples utilize corrected bitstrings, which undergo iterative Bayesian unfolding [79] prior to the post-selection. The REM+SYMM+RC samples are averages of REM outcomes over 30 logically equivalent circuits compiled with Pauli twirling and readout symmetrization. For comparison, the noiseless expected values are also presented in the EXACT curve. Results in the middle and right panels are obtained with 10^4 shots per sample and in the left panel with 10^5 shots per sample. Specifically, an individual RC circuit contributes approximately $\frac{10^4}{30}$ or $\frac{10^5}{30}$ shots to the collected REM+SYMM+RC samples. We gather 10 sample data points for each quantity, varying the value of m . Error bars denote error of the mean over the different samples. (a) Bhattacharyya distance, which is a measure of the difference between the ideal and measured bitstring probabilities versus block size m . (b) Energy versus m , which ideally vanishes. (c) Observed postselection success probability versus m .




Figure 6: The $|\xi = 1\rangle$ state preparation experiment on Aspen M-3 that combines $\mathcal{U}_\xi(m = 2)$ to construct the $4 \leq N \leq 14$ states. The sample data points are obtained using the $\{105, 104\}$ device qubits. We refer to the caption of Figure 5 for a detailed explanation of the data labels (EXACT, RAW, REM, REM+SYMM+RC) and the number of repeated measurements. We collect 10 sample points for each quantity, with the value of N varying, where the error bars denote the error of the mean across different samples. (a) Bhattacharyya distance, a measure of the difference between the ideal and measured bitstring probabilities versus system size N . (b) Energy versus N , which ideally vanishes. (c) Observed post-selection success probability versus N .




Figure 7: State tomography of the $|\xi = 1\rangle$ states for $N = 4$ and 14, constructed from multiple copies of $\mathcal{U}_\xi(m = 2)$ with the classical post-processing. It is computed using $\{105, 104\}$ qubits on Aspen M-3. The (i, j) block of the left / right density subplots shows $||\rho_{i,j}||$ / $\arg(\rho_{i,j})$, respectively, where $\rho_{i,j} \equiv \langle i|\xi = 1\rangle\langle\xi = 1|j\rangle$ and the integers i, j are in the binary representation. The nearly empty density matrix plot in the right panel ($N = 14$) illustrates the sparsity of the state $|\xi\rangle$. The inset plots show the ideal density matrices for comparison.

mitigation, shown as green dots. As m increases, both the energy expectation value $\langle H_\xi \rangle$ and the Bhattacharyya distance deviate farther from their ideal values of 0. At the same time, the success probability grows with increasing m . These manifest the trade-off between the error and the post-selection success probability: the use of a greater number of smaller circuit fragments facilitates higher-fidelity execution of the circuit, while demanding an increased sample complexity. Note that our experimental protocol is a special case of quantum divide-and-conquer [82, 83], where each block does not depend on another block’s measurement result.

Fig. 6 summarizes the output of the $4 \leq N \leq 14$ state preparation circuit that merges $\frac{N-2}{2}$ copies of $\mathcal{U}_\xi(m = 2)$. As the system size N increases, the Bhattacharyya distance stays close to 0. While the energy deviation shows an apparent $O(N)$ scaling, its values are considerably less than those with larger m building blocks. We also illustrate in Fig. 7 the state tomography of $|\xi = 1\rangle\langle\xi = 1|$ in the computational basis. In particular, its right panel highlights the sparsity of $|\xi\rangle$ states at $N = 14$.

Finally, we attempt to improve the precision of experimental results by adopting error mitigation techniques, which produce the blue and orange dots in Fig. 5 and 6. The estimated expectations usually deviate from true underlying values due to readout errors. One straightforward approach to mitigate the impact of readout errors is to calibrate the confusion matrix,

$$M_{ij} \equiv \Pr(\text{observed bits } i | \text{prepared bits } j) \quad (22)$$

and estimate the correct measurement bits that would be obtained in the absence of readout errors by applying the inverse of the confusion matrix. A more refined version of this approach is to replace the direct matrix inversion with the constrained optimization, called the iterative Bayesian unfolding [79]. In our QPU experiments, we apply this technique by calibrating the confusion matrices for each qubit group associated with the octagonal sublattice in the Aspen-M device layout.

The samples obtained from “corrected” bitstrings are colored in blue. They amplify the error measured in the Bhattacharyya distance and underestimate $\langle H_\xi \rangle$ when compared to the raw samples.

We note that, for the $\mathcal{U}_\xi(m = 2)$ block calculation, most blue samples have negative energy that violates the positive semi-definiteness of the Hamiltonian Eq. (21). The presence of the spurious samples is apparently an artifact of the Bayesian unfolding. It likely occurs due to a phase error since the sample Bhattacharyya distances are close to 0. When using larger circuit blocks $\mathcal{U}_\xi(m > 2)$, however, the Bayesian unfolding, combined with randomized compilation and readout symmetrization, helps to lower the estimated energy $\langle H_\xi \rangle$ while maintaining the Bhattacharyya distance near the same level.

Readout symmetrization is a technique that mitigates the asymmetric effects of noise [77]. It compiles multiple symmetrized programs for each original measurement. These compiled programs resemble the original program, except for a quantum bit-flip performed just before measuring the

qubit and a subsequent classical bit-flip of the measurement result. The outcomes of these programs are then combined to generate a set of ‘symmetrized’ readout results. Randomized compiling also adopts a similar approach by compiling multiple logically identical circuits, in which gates are randomized through virtual conjugation with Pauli operators [78]. By aggregating the outcomes of these equivalent unitary programs, this technique enhances circuit execution performance by suppressing coherent errors, which grow quadratically with circuit depth, and instead converting them into stochastic errors that accumulate linearly. The measurement results obtained from both circuit compilation techniques, using 30 logically equivalent circuits, are presented as individual orange dots.

Moreover, we implemented the full state preparation protocol on an IBM QPU (“ibmq_montreal”) without using fragmented circuits to reduce the circuit size. The IBM-QPU results show that the protocol’s performance increases with decreasing block size m and system size N , consistent with the abovementioned Rigetti-QPU results. More details of the IBM-QPU experiments are provided in Appendix C. We have also tried to simulate the system dynamics with the prepared state being the initial state, but the coherence properties of the QPU limit this effort.

4 Preparing the States $|\mathcal{S}_k\rangle$

Since generic states $|\mathcal{S}_k\rangle$ in this tower have entanglement scaling logarithmically with system size (see Sec. 2), they cannot be prepared with quantum circuits of constant depth. In this section, we demonstrate two polynomial-depth algorithms for preparing these states. The first, described in Sec. 4.1, relies on building an MPS representation of the states $|\mathcal{S}_k\rangle$ and then converting this representation to a quasi-polynomial depth quantum circuit. The second, described in Sec. 4.2, is a variational strategy that uses a polynomial-depth variational ansatz circuit to represent the states. Finally, we provide a proof-of-principle realization of these states on quantum hardware in Sec. 4.3, where we also write down a simplified linear-depth circuit for preparing the highest-weight state in the tower, $|\mathcal{S}_{N/2-1}\rangle$.

4.1 Quasi-Polynomial Depth Circuit from Matrix Product State Representation

To facilitate the discussion below, we rewrite the tower of scarred eigenstates for N sites as follows [see Eqs. (2), (3), and (9)]:

$$|\tilde{\mathcal{S}}_k\rangle = |0\rangle \otimes |\mathcal{D}_k^{N-2}\rangle \otimes |0\rangle, \quad (23)$$

where $|\mathcal{D}_k^{N-2}\rangle$ is an equal-amplitude superposition of all bitstrings of length $N-2$ with Hamming weight k (i.e., containing k 1s) that obey the Fibonacci constraint. The state $|\mathcal{D}_k^m\rangle \propto \mathcal{P}_{\text{fib}} |d_k^m\rangle$, where the Dicke state $|d_k^m\rangle$ is the equal-amplitude superposition of *all* length- m bitstrings with Hamming weight k . While Dicke states can be prepared in depth $O(mk)$ using the recursive strategy proposed in Ref. [59], we have not found an analogous strategy to prepare the projected Dicke states $|\mathcal{D}_k^m\rangle$ (see Sec. 5 for further discussion on this point). We therefore resort to an alternative approach that prepares the desired states in depth $O(mk \log^2 k)$, which is identical to the Dicke state preparation circuit up to an asymptotically minor multiplicative factor of $\log^2 k$. For generic scarred eigenstates $|\mathcal{S}_k\rangle$, which are written in terms of $|\mathcal{D}_k^m\rangle$ with $m = N-2$ and $k = O(N)$, this translates to a circuit with depth $O(N^2 \log^2 N)$. However, for states in the tails of the tower (i.e., ones for which k is finite), the resulting circuits are of linear depth.

Standard algorithms exist for converting an MPS with m sites and bond dimension χ into m unitaries of width $\lceil \log_2 2\chi \rceil$ [84, 85]. We first describe how to represent $|\mathcal{D}_k^m\rangle$ as an MPS, and then convert this MPS to a unitary circuit for state preparation.

One can express a matrix product state on m sites with open boundary conditions as

$$|\psi\rangle = \sum_{\vec{s}} \langle L | M^{s_1} M^{s_2} \dots M^{s_m} | R \rangle |\vec{s}\rangle, \quad (24)$$

where $|\vec{s}\rangle = \bigotimes_{i=1}^m |s_i\rangle$ are computational basis states labeled by bitstrings, M^{s_i} are square matrices of size χ , and $\langle L |$ and $|R\rangle$ are respectively χ -dimensional row and column vectors implementing




Figure 8: A graphical representation of the DFA transition matrix (31) for the state $|\mathcal{D}_4^*\rangle$.

open boundary conditions. One can view an MPS as a representation of a deterministic finite automaton (DFA), where the matrices M^{s_i} correspond to a transition matrix M for the DFA states [86]. Since, for any i , M^{s_i} is a square matrix of size χ , this DFA will have χ states, and $M_{j,k}^{s_i}$ will be nonzero if the character s_i takes the state j to the state k . The bra $\langle L|$ and ket $|R\rangle$ denote the initial and final states of the DFA, respectively. It is important to note that the states in the DFA live in the auxiliary bond space of the MPS and are not quantum states themselves.

DFAs are used to represent regular expressions, or string expressions using a finite set of characters. To see how we can use this language to represent quantum states, consider $|\mathcal{D}_2^4\rangle$:

$$|\mathcal{D}_2^4\rangle = \frac{1}{\sqrt{3}} (|0101\rangle + |1001\rangle + |1010\rangle) \quad (25)$$

The set of computational basis states in the superposition forms a language $\{|0101\rangle, |1001\rangle, |1010\rangle\}$ from which a regular expression can be constructed, where the characters in the regular expression are spin states $\{|0\rangle, |1\rangle\}$. For the states $|\mathcal{D}_k^*\rangle = \sum_{n=0}^{\infty} |\mathcal{D}_k^n\rangle$, the corresponding regular expression is

$$|0\rangle^* [|1\rangle|0\rangle|0\rangle^*]^{k-1} |1\rangle|0\rangle^*, \quad (26)$$

where $*$ is a Kleene star ($a^* = I + a + aa + aaa + \dots$ where I is identity (or empty string)) and is only applied to $|0\rangle$, and $[\cdot]^m$ means to repeat the term inside the bracket m times.

An example of a DFA representing the regular expression for $|\mathcal{D}_k^*\rangle$ (with $k = 4$) is shown in Fig. 8. Note that if we limit the number of transitions through the DFA to m , the regular expression for $|\mathcal{D}_k^*\rangle$ reduces to that for $|\mathcal{D}_k^m\rangle$. For general k , the DFA is defined on $2k$ states $\{S_0, A_j, B_j, F_k\}_{j=1\dots k-1}$. The DFA starts at S_0 and ends at F_k . The first computational state can either be $|0\rangle$, which transitions the DFA back to S_0 , or $|1\rangle$, which transitions the DFA to A_1 . We denote this by

$$\begin{aligned} S_0(|0\rangle) &\rightarrow S_0 \\ S_0(|1\rangle) &\rightarrow A_1 \end{aligned} \quad (27)$$

For the intermediate states, we have

$$\begin{aligned} A_j(|0\rangle) &\rightarrow B_j \\ B_j(|0\rangle) &\rightarrow B_j \\ B_j(|1\rangle) &\rightarrow A_{j+1} \end{aligned} \quad (28)$$

Finally, the DFA terminates at the final state F_k , where

$$\begin{aligned} B_{k-1}(|1\rangle) &\rightarrow F_k \\ F_k(|0\rangle) &\rightarrow F_k \end{aligned} \quad (29)$$

All these transition rules can be summarized by a transition matrix $\mathcal{M}$, where $\mathcal{M}_{i,j}$ gives the character needed to take the DFA from state i to state j :

$$\begin{aligned}\mathcal{M}_{S_0, S_0} &= |0\rangle, & \mathcal{M}_{S_0, A_1} &= |1\rangle, \\ \mathcal{M}_{A_j, B_j} &= |0\rangle, & \mathcal{M}_{B_j, A_{j+1}} &= |1\rangle, & \mathcal{M}_{B_j, B_j} &= |0\rangle, \\ \mathcal{M}_{B_{k-1}, F_k} &= |1\rangle, & \mathcal{M}_{F_k, F_k} &= |0\rangle.\end{aligned}\tag{30}$$

The full transition matrix is then

$$\mathcal{M} = \begin{pmatrix} S_0 & A_1 & B_1 & \cdots & A_{k-1} & B_{k-1} & F_k \\ |0\rangle & |1\rangle & 0 & \cdots & 0 & 0 & 0 \\ 0 & 0 & |0\rangle & \cdots & 0 & 0 & 0 \\ 0 & 0 & |0\rangle & \cdots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 0 & |0\rangle & 0 \\ 0 & 0 & 0 & \cdots & 0 & |0\rangle & |1\rangle \\ 0 & 0 & 0 & \cdots & 0 & 0 & |0\rangle \end{pmatrix} \begin{matrix} S_0 \\ A_1 \\ B_1 \\ \vdots \\ A_{k-1} \\ B_{k-1} \\ F_k \end{matrix}\tag{31}$$

To obtain the matrices from Eq. (24), we separate this transition matrix as $\mathcal{M} = M^0 |0\rangle + M^1 |1\rangle$, where

$$M^0 = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & \cdots & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & \cdots & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & \cdots & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & \cdots & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & \cdots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & 0 & \cdots & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & \cdots & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & \cdots & 0 & 0 & 1 \end{pmatrix}, \quad M^1 = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & \cdots & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \cdots & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & \cdots & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \cdots & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \cdots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & 0 & \cdots & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \cdots & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & \cdots & 0 & 0 & 0 \end{pmatrix}\tag{32}$$

Alternatively, M^0 is obtained by applying $\langle 0|$ to every entry in $\mathcal{M}$, and M^1 is obtained by applying $\langle 1|$ to every entry in $\mathcal{M}$. These are $2k \times 2k$ matrices, and so the overall bond dimension of the MPS will be $2k$. Also, since we start at state S_0 and end at state F_k , the boundary vectors become

$$\begin{aligned}\langle L| &= (1 \quad 0 \quad \cdots \quad 0) \\ |R\rangle &= (0 \quad \cdots \quad 0 \quad 1)^T\end{aligned}\tag{33}$$

One can now read Eq. (24) from left to right as a traversal through the DFA in Fig. 8, starting at DFA state S_0 , making m transitions according to the spin configuration $\vec{s}$, and ending at DFA state F_k .

The resulting MPS can be converted into a quantum circuit by following the sequential preparation procedure of Ref. [84]. Figure 9 illustrates how this is done through a series of QR decompositions and tensor contractions. Specifically, each node in the MPS $M_{jk}^{(s_i)}$ is a rank-3 tensor, except those at the boundaries which are rank-2 tensors. To perform a QR decomposition, we group together a physical index of M with a virtual bond index, denoted as $M_{(s_i j),k}$. A QR decomposition will decompose this $d\chi_r \times \chi_l$ matrix as $M_{(s_i j),k} = U_{(s_i j),l} R_{l,k}$ where $U_{(s_i j),l}$ is an $d\chi_l \times d\chi_l$ unitary matrix, and $R_{l,k}$ is a $d\chi_l \times \chi_r$ upper triangular matrix, where χ_l and χ_r are the left and right bond dimension of $M_{jk}^{(s_i)}$.

Figure 9 depicts a 4-site MPS with $\chi = 4$ and $d = 2$. The top-most qubit has $\chi_l = 1$ and $\chi_r = 4$, so we decompose $M_1 = U_1 R_1$, where U_1 is a 2×2 unitary, and R_1 is a 2×4 upper triangular matrix. R_1 is then contracted with the rank-3 tensor M_2 to form the rank-3 tensor




Figure 9: (top) Sequential generation of an MPS (a) into a quantum circuit (i) of 4 sites. The MPS is first ordered as shown in (b). A QR decomposition is performed on the top site giving a unitary U_1 and an upper triangular matrix R_1 resulting in (c). R_1 is then contracted with M_2 resulting in (d). The series of QR decompositions and contractions are performed as shown (e)-(h) until we reach the final MPS circuit (i). Circuit diagram for the $|\mathcal{D}_k^m\rangle$ state with $m = 6$ and $k = 2$ obtained from the MPS representation. (bottom) Once the unitaries are formed, we use Gray codes to decompose the unitaries into quantum gates.

$(M_{12})_{(ij)k}$. We then perform another QR decomposition $M_{12} = U_{12}R_{12}$, where both U_{12} and R_{12} are 4×4 , and then contract R_{12} with M_3 to form the rank-4 tensor M_{123} . The 8×4 matrix M_{123} allows us to perform a thin QR factorization, where we can decompose $M_{123} = V_{123}R_{123}$, where V_{123} is an 8×4 isometry ($V_{123}^\dagger V_{123} = I$) and R_{123} is a 4×4 matrix. To write V_{123} as a circuit, we embed this matrix into an 8×8 unitary U_{123} such that when we apply the state $|0\rangle$ to the extra degree of freedom in U_{123} , we get the isometry V_{123} . This process of QR factorization followed by contraction is performed until all MPS sites have been written as a unitary.

Once this procedure is done, we are left with m unitaries, each of which can be decomposed into $O(k)$ two-level unitaries. Using Gray codes, a two-level unitary, which here spans over $O(\log k)$ qubits, can be decomposed into $O(\log k)$ Toffoli gates [87]. These $O(\log k)$ -qubit Toffoli gates require $O(\log k)$ CNOT gates [88]. The depth of the resulting circuit is then $O(mk \log^2 k)$. Fig. 10 depicts an example preparation circuit for the state $|\mathcal{D}_2^6\rangle$. The circuit makes use of the following single-qubit unitaries:




Figure 10: Circuit diagram for the $|\mathcal{D}_k^m\rangle$ state with $m = 6$ and $k = 2$ obtained from the MPS representation. The unitaries U_1 through U_5 are listed in Eq (34)

$$\begin{aligned}
 U_1 &= \begin{pmatrix} -\sqrt{\frac{2}{5}} & -\sqrt{\frac{3}{5}} \\ -\sqrt{\frac{3}{5}} & \sqrt{\frac{2}{5}} \end{pmatrix} \\
 U_2 &= \begin{pmatrix} \sqrt{\frac{1}{2}} & \sqrt{\frac{1}{2}} \\ \sqrt{\frac{1}{2}} & -\sqrt{\frac{1}{2}} \end{pmatrix} \\
 U_3 &= \begin{pmatrix} -\sqrt{\frac{1}{4}} & \sqrt{\frac{3}{4}} \\ \sqrt{\frac{3}{4}} & \sqrt{\frac{1}{4}} \end{pmatrix} \\
 U_4 &= \begin{pmatrix} \sqrt{\frac{2}{3}} & \sqrt{\frac{1}{3}} \\ \sqrt{\frac{1}{3}} & -\sqrt{\frac{2}{3}} \end{pmatrix} \\
 U_5 &= \begin{pmatrix} -\sqrt{\frac{1}{3}} & \sqrt{\frac{2}{3}} \\ \sqrt{\frac{2}{3}} & \sqrt{\frac{1}{3}} \end{pmatrix}.
 \end{aligned} \tag{34}$$

4.2 Polynomial Depth Variational Ansatz

A strategy similar to the recursive linear circuit used in the preparation of the $|\xi\rangle$ state suggests a variational circuit architecture for creating the $|\mathcal{S}_k\rangle$ state should exist. Up to unimportant relative phase factors, the $|\mathcal{S}_k\rangle$ state is an equal-weight superposition of all bitstrings of N sites, such that there are k 1s, the first and last sites are 0, and there are no two neighboring 1s. All such basis states can be generated by a four-qubit building-block unitary operator that transforms the state $|0100\rangle$ to a linear superposition of states $|0100\rangle$ and $|0010\rangle$ [89], and acts trivially on all other 14 basis states. One can construct the unitary by applying a two-qubit gate on the two middle sites controlled by the first and the last qubit such that the gates act nontrivially only when the first and the last qubits are 0, in which case the two middle qubits transform as

$$\begin{aligned}
 |0\rangle_{j+1}|1\rangle_{j+2} &\rightarrow \cos(\theta_j)|0\rangle_{j+1}|1\rangle_{j+2} + \sin(\theta_j)|1\rangle_{j+1}|0\rangle_{j+2} \\
 |1\rangle_{j+1}|0\rangle_{j+2} &\rightarrow \cos(\theta_j)|1\rangle_{j+1}|0\rangle_{j+2} - \sin(\theta_j)|0\rangle_{j+1}|1\rangle_{j+2}.
 \end{aligned}$$

We have chosen a building block with only real elements for easier optimization.

Mathematically, we can write the unitary $U_j(\theta_j)$ acting on qubits j to $j + 3$ as

$$U_j(\theta) = \exp \left[i \frac{\theta}{2} P_j (X_{j+1} Y_{j+2} - Y_{j+1} X_{j+2}) P_{j+3} \right], \tag{35}$$

where the operators $P_j = |0\rangle_j \langle 0|_j$ implement the control on the first and last qubits and $(X_{j+1} Y_{j+2} - Y_{j+1} X_{j+2})$ acts as a Pauli- Y on the 01 and 10 states of the two middle qubits, while annihilating 00 and 11. The building block U_j is illustrated in Fig. 11(a).

N	k	d	n_a	$1 - \langle \psi \mathcal{S}_k \rangle ^2$
13	5	21	20	2.9×10^{-13}
14	2	55	45	4.9×10^{-5}
14	3	120	41	1.2×10^{-3}
14	4	126	35	2.4×10^{-3}
14	5	56	27	2.2×10^{-13}
14	6	7	17	1.6×10^{-15}
15	2	66	52	3.1×10^{-5}
15	3	165	48	1.6×10^{-3}
15	4	210	42	3.4×10^{-3}
15	5	126	34	2.1×10^{-3}
15	6	28	24	5.3×10^{-14}
16	2	78	60	1.7×10^{-5}
16	3	220	56	2.6×10^{-3}
16	4	330	50	4.4×10^{-3}
16	5	252	42	6.9×10^{-3}
16	6	84	32	6.2×10^{-13}
16	7	8	20	2.4×10^{-15}

Table 1: The dimension of the constrained Hilbert space d , the number of variational angles n_a and the minimum error $1 - |\langle \psi | \mathcal{S}_k \rangle|^2$ achieved with 2000 local constrained optimization runs starting from random initial angles, for several N and k .

We now construct an ansatz circuit. Since the gates are designed to conserve both the total number of 1s, i.e., k , and the Fibonacci constraint of no neighboring 1s, we start with the following initial state that has k 1s:

$$|\psi_0\rangle = |01\rangle^{\otimes k} |0\rangle^{\otimes N-2k}. \quad (36)$$

In Fig. 11(b), we show an example of the circuit acting on the above initial state for $k = 2$ and $N = 8$. Our ansatz consists of multiple staircases of the U_j gates. In the first layer, we apply $U_1 \dots U_{2k-3} U_{2k-1}$ to move the last 1 from qubit $2k$ to qubit $2k+1$ and generate all states with the last 1 at or before qubit $2k+1$. The next layer, $U_2 \dots U_{2k-1} U_{2k}$, generates all states with the last 1 at or before $2k+2$. These staircases continue until we reach the staircase starting with U_{N-3} , acting on the last four qubits. All these staircases have unitaries acting on every other site. These layers generate all d states. However, we have found that the number of variational parameters is insufficient to obtain an equal-weight superposition. Thus the ansatz also contains a complete staircase at the end with a unitary $U_1 U_2 \dots U_{N-3}$ to equalize the probabilities. We then optimize the variational angles in these four-qubit unitaries to generate an equal-weight superposition of all d basis states by minimizing $1 - |\langle \psi | \mathcal{S}_k \rangle|^2$, where $|\psi\rangle$ is the state after the application of the variational unitaries. A final layer of Z gates corrects the signs of the amplitudes.

Two important characteristics of this circuit are the dimension d of the constrained Hilbert space and the number n_a of variational parameters. While d scales exponentially in the asymptotic limit, n_a is quadratic in system size. The dimension of the constrained Hilbert space is $d = \mathcal{N}(N, k) = \binom{N-k-1}{k}$. We need our circuit to generate all the states subject to the Fibonacci constraint and the conserved number of 1s. Thus, the basis states can be viewed as k blocks of 01 as partitions between $N - 2k - 1$ zeros with one zero tagged at the end of the chain. The number of variational angles n_a is easy to calculate. For even N , the first and the second layer have k gates each, the third and fourth layer $k+1$ gates, and so forth until we reach the staircase with $N/2 - 2$ gates containing every site from 2 to $N - 4$. We then have a staircase with $N/2 - 1$ gates and then a full staircase with $N - 3$ gates. Thus $n_a = 2[k + (k+1) + \dots + N/2 - 2] + (N/2 - 1) + (N - 3)$, which simplifies to

$$n_a = N^2/4 - k(k-1) - 2, \quad N \text{ even}. \quad (37)$$

Similarly for odd N , we have $n_a = 2[k + (k+1) + \dots + (N+1)/2 - 2] + (N - 3)$, which leads to

$$n_a = (N^2 - 1)/4 - k(k-1) - 2, \quad N \text{ odd}. \quad (38)$$




Figure 11: (a) The four-qubit gate with one variational angle serves as a building block for constructing the variational ansatz. These gates conserve the number of ones and preserve the Fibonacci constraint while allowing transformations between all basis states in the constrained Hilbert space. (b) An example of the general architecture of the variational circuit for $N = 8$ and $k = 2$.




Figure 12: Variational circuit to prepare $|\mathcal{D}_{k=2}^{m=5}\rangle$. (a) The circuit consists of four single-controlled XY and four double-controlled XY gates, each parameterized by a single parameter. This circuit is pre-optimized by classical computers with machine-precision infidelity. (b) The controlled XY gates are decomposed into CNOTs and single-qubit gates to be implemented on “ibmq.kolkata”.

Importantly, the number of variational gates is quadratic in both N and k .

The numerical optimization results are listed for several N and k in Table 1. The system exhibits many local minima (local minimization routines with different initial angles tend to converge to different optimal angles), making global optimization challenging. However, we have found that constrained local optimization with random initial angles between 0 and 2π finds solutions with reasonably small error $1 - |\langle \psi | \mathcal{S}_k \rangle|^2 < 0.01$. For many cases with smaller n_a and d , and less complex optimization, the minimum error is zero to within machine precision, suggesting the $|\mathcal{S}_k\rangle$ state may be exactly reachable with the ansatz architecture. In practice, we can only find very good approximations for large d and n_a .

It would be interesting to determine whether the variational ansatz circuit proposed here can represent the $|\mathcal{S}_k\rangle$ states exactly. This would conclusively demonstrate that the scarred eigenstates of the model (1) can be prepared in polynomial depth (as opposed to quasipolynomial depth as shown in Sec. 4.1), which is known to be possible for, e.g., Dicke states [59]. Attempts to find an exact solution for the rotation angles in the circuit architecture proposed here yield n_a coupled nonlinear equations that are intractable in all but the simplest cases. An interesting question for future work would be to investigate whether a modified circuit architecture involving gates of




Figure 13: Probability distribution of $|\mathcal{D}_{k=2}^{m=5}\rangle$ prepared on “ibmq_kolkata”. The state is prepared on qubits [14, 13, 12, 15, 18] and measured with 196608 shots. The errors are reduced by dynamical decoupling and readout error mitigation.

the form in Eq. (35) can be used to obtain an analytically tractable system of equations for the rotation angles.

4.3 QPU Results

4.3.1 Pre-Optimized Variational Circuit

In this subsection, we present a proof-of-principle experiment preparing $|\mathcal{S}_k\rangle$ using the variational circuit proposed in Section 4.2 for one of the simplest cases, $N = 7$ and $k = 2$. Similar to $|\xi\rangle$, we simulate only the middle five spins, as the two boundary spins are always fixed to $|0\rangle$ by the open-boundary conditions. In other words, we prepare $|\mathcal{D}_k^{m=N-2}\rangle$ defined in Eq. (23).

The five-qubit experiment involves the variational circuit shown in Fig. 12(a). The variational parameters $\{\theta_j\}_{j=1,2,\dots,8}$ are pre-optimized on classical computers to obtain an infidelity of 0 to machine precision. The double-controlled XY gates are equivalent to U_j defined by Eq. (35). The U_j gates at the edge turn into single-controlled XY gates since the two boundary spins are always in $|0\rangle$. The single-controlled and double-controlled XY gates are decomposed into CNOT and single-qubit gates as shown in Fig. 12(b). With this decomposition, each single-controlled XY gate is implemented by six CNOTs. The double-controlled XY gate consists of 17 CNOTs, and it is further compiled down by Qiskit [90] to 12 CNOTs. Hence, the five-qubit variational circuit consists of 72 CNOT gates.

The pre-optimized circuit is experimentally simulated using an IBM QPU (“ibmq_kolkata”). In order to reduce errors, the circuit is further optimized by inserting dynamical decoupling pulse sequences $\sqrt{X}\sqrt{X}\sqrt{X}\sqrt{X}$ to idling qubits, and readout error mitigation is performed on the measurement outcomes. The probability distribution of the five-qubit state is measured with 196608 shots and is shown in Fig. 13.

The experimental result (blue bars) qualitatively matches the exact distribution (orange bars) with a few incorrectly high probabilities at 00010, 00100, 11010, and 11100. These incorrect bitstrings are one bit flip away from the correct bitstrings, and these errors are expected due to the high CNOT gate count of the variational circuit given the average CNOT infidelity of around 1%. The Bhattacharyya distance between the experimental and exact distributions is 0.3809, which is comparable to the state-preparation result of $|\xi\rangle$ shown in Appendix C, albeit more qubits (up to 20 with ancilla qubits) are considered in that case.

Extending the experiment to larger N on currently available QPUs is difficult due to the increased CNOT gate count exceeding the capacity of these QPUs. In order to perform the extended experiments on present-day hardware with high fidelity, it is necessary to carry out further simplifications like those discussed in the next subsection.

4.3.2 Exact Circuit for $k = k_{\max}$

It is noticeable that the doubly-controlled gates are essential pieces in building up the state preparation circuits described in Secs. 3.1 and 4.2. As we saw in the previous section, compiling them into the set of native gates often produces quantum programs relatively expensive in the two-qubit gate complexity. An important simplification occurs for the specific case of $k = k_{\max} \equiv N/2 - 1$ for which an alternative circuit, shown in Fig. 14, can be used that involves only $(N - 3)$ two-qubit gates.

We can deterministically realize the maximal k condition by setting the spin state of the $(i+1)$ -th qubit to be opposite to that of the i -th qubit for all $i \in \{2, 4, \dots, N-2\}$. This can be implemented by initializing the target qubit in $|0\rangle$ and then acting with the controlled bit-flip operator,

$$C_0\text{NOT}_{i,i+1} = X_i \cdot \text{CNOT}_{i,i+1} \cdot X_i, \quad (39)$$

triggered only when the control qubit is in the $|0\rangle$ state. Next, the superposition state can be prepared with a variant of the unitary block Eq. (11) where we modify the controlled rotations to be triggered only if the control qubit is in the $|1\rangle$ state. Running it on the even-indexed qubits between $2 \leq i \leq N-2$, we obtain a superposition of the following bitstring states,

$$|0(10)^a(01)^b0\rangle \text{ satisfying } a + b = k_{\max}. \quad (40)$$

Upon acting with Eq. (39) on all the adjacent $(i, i+1)$ pairs with $i \in \{2, 4, \dots, N-2\}$, the superposition state becomes

$$\sum_{a+b=k_{\max}} c_{a,b} |0(10)^a(01)^b0\rangle \quad (41)$$

where $c_{a,b} \in \mathbb{R}$ and $\sum_{a,b} c_{a,b}^2 = 1$. Lastly, we observe that the following Pauli rotation parameters,

$$\theta_i = 2 \tan^{-1}(\sqrt{N/2 - i}), \quad (42)$$

leads to the equal-amplitude superposition state $|0\rangle \otimes |\mathcal{D}_{k_{\max}}^{N-2}\rangle \otimes |0\rangle$, since all coefficients become the same:

$$c_{a,b} = \prod_{i=1}^a \sqrt{\frac{\frac{N}{2} - i}{\frac{N}{2} - i + 1}} \cdot \sqrt{\frac{1}{\frac{N}{2} - a}} = \sqrt{\frac{1}{N/2}}. \quad (43)$$

Figure 14 illustrates the $|\tilde{\mathcal{S}}_{k_{\max}}\rangle$ state preparation circuit.

We execute the state preparation circuit on Aspen-M-2 and Aspen-M-3. Ideally, if we perform Pauli-Z measurements after running the circuit, the resulting N -bitstrings must include $k_{\max}$ non-consecutive 1's and begin and end with 0 by construction. However, due to limitations of current NISQ hardware, the circuit simulation of the $|\mathcal{S}_{k_{\max}}\rangle$ state often induces an error in the total spin measurement $k = k_{\max}$ and may violate the boundary condition or the Fibonacci constraint. Therefore, to increase the fidelity of the simulation, we post-select the measurement bitstrings along with optional application of the readout error mitigation and randomized compiling techniques.

Assuring the total spin selection rule allows us to introduce the compact encoding scheme that uses only $k_{\max}$ qubits, i.e., indexed as $\{2, 4, \dots, N-2\}$, to encode the N -qubit states satisfying $k = k_{\max}$ without consecutive 1's. It is equivalent to decoupling the non-entangled boundary spins and replacing the error-prone quantum operator Eq. (39) with deterministic post-processing. Its use of fewer qubits also brings the advantage of circumventing non-linear qubit connectivity and considerable error reduction. To make sure the Fibonacci constraint is satisfied, we apply the following projection operator

$$\mathcal{P} = \prod_{i=2,4,\dots,N-4} (1 - P_i P'_{i+2}) \quad (44)$$

that removes $|01\rangle$ configurations in the compressed $k_{\max}$ -qubit encoding, or equivalently, $|0110\rangle$ configurations in the full N -qubit representation.




Figure 14: The $|\tilde{S}_k\rangle$ state preparation circuit with $k = k_{\max} \equiv \frac{N}{2} - 1$. The rotation angles are $\theta_i = 2 \tan^{-1}(\sqrt{N/2 - i})$. The scar state $|\mathcal{S}_{k_{\max}}\rangle$ is obtained by applying Z gates at the end of the circuit following Eqs. (9).




Figure 15: The execution results of the $|\mathcal{S}_{k_{\max}}\rangle$ state preparation circuit on Aspen-M devices, which uses the $k_{\max} = \frac{N}{2} - 1$ qubit encoding for $6 \leq N \leq 14$. The indices of the participating Aspen nodes are: $\{147, 140\}$ (M-3) for $N = 6$, $\{147, 140, 141\}$ (M-3) for $N = 8$, $\{16, 17, 10, 11\}$ (M-2) for $N = 10$, $\{11, 10, 17, 16, 1\}$ (M-2) for $N = 12$, $\{26, 11, 10, 17, 16, 1\}$ (M-2) for $N = 14$. We refer to the caption of Figure 5 for a detailed explanation of the data labels (EXACT, RAW, REM, REM+SYMM+RC) and the number of repeated measurements. We gather 5 sample points for each quantity, with the value of N varying, where the error bars denote the error of the mean across different samples. (a) Bhattacharyya distance, which is a measure of the difference between the ideal and measured bitstring probabilities versus system size N . (b) Energy versus N , which ideally takes the value of $-(N - 3)$. (c) Energy error versus N , computed from $100 \frac{|\langle H_I \rangle + N - 3|}{N - 3}$. (d) Observed postselection success probability versus N .

Before	After	Before	After	Before	After	Before	After
Z_1 and Z_N	1	$X_{2i}Z_{2i+1}$	0	$X_{2i}I_{2i+1}$	0	$X_{2i}X_{2i+1}$	X_{2i}
$Z_{2i}I_{2i+1}$	Z_{2i}	$Y_{2i}Z_{2i+1}$	0	$Y_{2i}I_{2i+1}$	0	$X_{2i}Y_{2i+1}$	$-Y_{2i}$
$I_{2i}Z_{2i+1}$	$-Z_{2i}$	$Z_{2i}X_{2i+1}$	0	$I_{2i}X_{2i+1}$	0	$Y_{2i}X_{2i+1}$	Y_{2i}
$Z_{2i}Z_{2i+1}$	-1	$Z_{2i}Y_{2i+1}$	0	$I_{2i}Y_{2i+1}$	0	$Y_{2i}Y_{2i+1}$	X_{2i}

Table 2: Projection rule for N -qubit Pauli strings to Pauli operators on the $k_{\max}$ -qubit Hilbert space with $1 \leq i \leq k_{\max}$.




Figure 16: The $|\mathcal{S}_{k_{\max}}\rangle$ state tomography in the compressed $k_{\max}$ -qubit encoding at $N = 6$ and 14 . It uses the $\{147, 140\}$ ($N = 6$) and $\{26, 11, 10, 17, 16, 1\}$ ($N = 14$) device qubits on Aspen M-3 and M-2, respectively. The (i, j) block of the left and right density subplots shows $\|\rho_{i,j}\|$ and $\arg(\rho_{i,j})$, respectively, where $\rho_{i,j} \equiv \langle i|\mathcal{S}_{k_{\max}}\rangle\langle\mathcal{S}_{k_{\max}}|j\rangle$ and the integers i, j are in the binary representation. The nearly empty appearance of the density plot in the right panel ($N = 14$) manifests the sparsity of the scar state $|\mathcal{S}_{k_{\max}}\rangle$ which only occupies $N/2$ out of $2^{k_{\max}}$ possible bitstrings. The inset plots show the ideal density matrices for comparison.

Having run the state preparation program, we measure the errors through the Bhattacharyya distance defined in Eq. (20) between the ideal and empirical bitstring distributions and the expectation value of the spin-1/2 Hamiltonian from Eq. (1). With the $k_{\max}$ -qubit encoding, the Hamiltonian is mapped to

$$H_0 = \lambda H_\lambda + \Delta H_\Delta + J H_J \quad (45)$$

where

$$\begin{aligned} H_\lambda &= 0, \quad H_\Delta = 2, \quad \text{and} \\ H_J &= Z_2 - Z_{N-2} - \frac{N}{2} + 1 - \sum_{i=1}^{N/2-2} Z_{2i} Z_{2i+2}. \end{aligned} \quad (46)$$

It is obtained by applying the encoding rules of Table 2, where vanishing substitutions on the second and third columns impose the total spin selection rule $k = k_{\max}$. This reduces the task of computing $\langle H_0 \rangle$ to measuring $\langle H_J \rangle$.

Fig. 15 collects the results of the $|\mathcal{S}_{k_{\max}}\rangle$ state preparation experiment, showing the Bhattacharyya distance, the energy $\langle H_J \rangle$ and its error, and the success probability of the Fibonacci projection, see Eq. (44). Individual sample points are drawn as round dots, and their averages are connected as the solid line. Each dot aggregates 10^4 shots for the energy and success probability, and 10^5 shots for the Bhattacharyya distance. We record the qubit indices used for the QPU experiment in the caption of Fig. 15.

We consider the raw samples without the readout error mitigation and randomized compiling techniques, denoted as green dots. We note a rapid drop of the success probability as N increases. Unlike the probabilistic $|\xi\rangle$ state preparation protocol in Section 3.2, the noiseless probability should be 1, and thus the decrease from 1 highlights the detrimental effects of the noise. While the measured Bhattacharyya distance after the Fibonacci projection remains close to 0, the estimated energy $\langle H_J \rangle$ exhibits a linearly growing error rate as compared to the ideal value $-(N-3)$, reaching around 10% at $N = 14$. Also note that the noise can create spurious samples whose $\langle H_J \rangle$ is less than the theoretical minimum of the H_J in Eq. (46).

Next, we try various error mitigation techniques to improve the estimate accuracy. To reduce the measurement error, the iterative Bayesian unfolding [79] uses the confusion matrix of 8-bit strings that reflects Aspen-M's octagonal layout. The samples obtained from the modified bitstrings are marked as blue dots and feature interesting changes. They have a higher success probability—i.e., they remain more in the Fibonacci subspace—and feature less error in estimating $\langle H_J \rangle$ after the Fibonacci projection. However, their Bhattacharyya distance also amplifies significantly with increased sample variance. Application of randomized compilation [78] and readout symmetrization [77] can help to reduce the Bhattacharyya distance while maintaining other improvements, as seen from the orange entries in Fig. 15.

Figure 16 displays the results of the $|\mathcal{S}_{k_{\max}}\rangle$ state tomography after applying the aforementioned error mitigation methods and the Fibonacci post-projection. It illuminates the sparsity of the $|\mathcal{S}_{k_{\max}}\rangle$ state that populates only $N/2$ bitstrings out of $2^{k_{\max}}$ different choices.

4.4 Alternative State Preparation Protocols

Here we consider two possible alternative strategies to prepare the scarred eigenstates $|\mathcal{S}_k\rangle$. First, it is natural to ask whether the recursive method of Ref. [59] to prepare Dicke states $|d_k^m\rangle$ in depth $O(mk)$ can be adapted to exactly prepare the $|\mathcal{S}_k\rangle$ states, or equivalently the projected Dicke states $|\mathcal{D}_k^m\rangle$ [see Eq. (23)], in polynomial depth. The construction of a state preparation circuit in Ref. [59] hinges on a recursion relation between Dicke states with different magnetizations. The projected Dicke states also obey a recursion relation:

$$\begin{aligned} |\mathcal{D}_k^m\rangle = & \sqrt{\frac{k}{m-k+1}} |\mathcal{D}_{k-1}^{m-2}\rangle \otimes |01\rangle \\ & + \sqrt{\frac{m-2k+1}{m-k+1}} |\mathcal{D}_k^{m-1}\rangle \otimes |0\rangle. \end{aligned} \quad (47)$$

However, unlike the recursion relation for Dicke states, this one relates $|\mathcal{D}_k^m\rangle$ to both $|\mathcal{D}_k^{m-1}\rangle$ and $|\mathcal{D}_{k-1}^{m-2}\rangle$, which are defined on qubit registers of different sizes. This complicates the possibility of a straightforward modification of the protocol of Ref. [59], and we leave this question for future work. Furthermore, we reiterate that the $O(mk)$ depth scaling for the Dicke states is only marginally smaller in an asymptotic sense than the $O(mk \log^2 k)$ scaling obtained within our MPS-to-circuit approach.

Second, it is interesting to consider nonunitary methods to prepare the scarred eigenstates $|\mathcal{S}_k\rangle$. In particular, these eigenstates can be prepared from the state $|\xi\rangle$ by projectively measuring the magnetization operator $M_z = \sum_{i=1}^N Z_i$. For any $k = 0, \dots, N/2-1$, $|\mathcal{S}_k\rangle$ is an eigenstate of M_z with eigenvalue $N-2k$. Since each of these eigenvalues is unique, we see from Eq. (8) that the projection of $|\xi\rangle$ into an M_z eigenspace must yield one of the $|\mathcal{S}_k\rangle$. The probability of a given measurement outcome k is peaked around a value that depends on $|\xi|^2$; this most probable outcome can be tuned from $k=0$ as $|\xi|^2 \rightarrow 0$ to $k=N/2-1$ as $|\xi|^2 \rightarrow \infty$. Thus, given the ability to measure M_z without fully collapsing the system into a z -basis eigenstate, one can in principle prepare the state $|\mathcal{S}_k\rangle$ in constant depth. One way to achieve such a measurement is by a dispersive coupling $\chi_d(\sum_i Z_i)a^\dagger a$ between the qubits and a single cavity mode, where the measurement outcome would be recorded as a shift of the cavity frequency by an integer multiple of χ_d . A projective measurement of M_z can also be achieved using an ancilla register containing $\lceil \log_2 N \rceil$ qubits, one for each digit of the binary representation of the magnetization eigenvalue, using methods proposed in Refs. [91] and [92]. This state preparation method also suffers from postselection overhead beyond what is needed to prepare $|\xi\rangle$: the probability of the most probable measurement outcome decays as a power law in N for large N . However it may have the advantage of comparatively modest circuit depth relative to the unitary state preparation protocols considered here.

Finally, we mention an alternative to the above strategy based on imaginary time evolution. Since projecting $|\xi\rangle$ into an eigenspace of M_z with eigenvalue $m_k = N-2k$ yields $|\mathcal{S}_k\rangle$, one can imagine implementing the projection by applying $e^{-\tau(M_z - m_k)^2}$ to $|\xi\rangle$ in the limit $\tau \rightarrow \infty$. This imaginary time evolution under the Hamiltonian $H_k = (M_z - m_k)^2$ can be implemented using a number of quantum algorithms [93, 94, 95, 96]. This approach could be well adapted to, e.g., trapped ion devices, which can implement the nonlocal interactions contained in H_k using their native all-to-all connectivity.

5 Conclusion and Outlook

In this work we have explored various approaches to preparing quantum many-body scar states and their superpositions on digital quantum computers with a focus on the spin-1/2 chain model of Ref. [27], where an emergent Fibonacci constraint (two “1”s may not appear next to each other) ensures that superpositions of scarred eigenstates must be entangled. Linear-depth circuits

to prepare a one-parameter family of area-law-entangled superpositions of scarred eigenstates, as well as a nonunitary preparation scheme that uses measurement and postselection to prepare the same superposition in constant depth are provided. We also derived an MPS representation of the tower of scarred eigenstates and used this to generate a quasipolynomial-depth circuit to prepare individual scar states; additionally, we proposed a variational scheme based on a polynomial-depth ansatz that captures the scar states with fidelity at least 99% in all numerically accessible cases. Proof-of-concept demonstrations of scar-state preparation were executed on quantum hardware.

While the state preparation protocols explored in this work were formulated with a specific model in mind, they can readily be adapted to other towers of scar states. For instance, entangled superposition states analogous to $|\xi\rangle$ are known for both the AKLT model [56] and the bond-bimagnon tower in the spin-1 XY model [26, 55]. In both cases, the initial state is a finite-bond-dimension MPS which can be viewed as the projection of a simple wavefunction on χ -dimensional qudits. One expects that nonunitary methods along the lines of Sec. 3 could be used to prepare such states (see also the method of Ref. [74] to prepare the AKLT ground state, which is in a similar spirit). Likewise, both towers—and many other examples of QMBS [30]—admit exact (or approximate [97]) MPS representations that can be converted to quasipolynomial-depth quantum circuits [84, 85]. It is also worth noting that simpler state preparation protocols can be used in some cases—for example, the bimagnon tower in the spin-1 XY model [26] can be superposed into a product state that is trivial to prepare, and the scarred eigenstates themselves are effective spin-1/2 Dicke states that can be generated using a slight variation on the approach of Ref. [59].

Our results can feed into future studies of the stability of scarred eigenstates and their dynamics on quantum computers. There are two experiments one would like to perform—one in which the state $|\xi\rangle$ is evolved under a perturbation of the Hamiltonian in Eq. (1), and another in which the state $|\mathcal{S}_k\rangle$ is evolved. In the former case, the goal is to extract the lifetime of the oscillatory scarred dynamics from the time series of a local operator’s expectation value³, while in the latter it is to extract the lifetime of the scarred eigenstate from similar data. One relevant class of perturbations to consider is an external magnetic field in the x -direction, $\epsilon \sum_i X_i$, which breaks the conservation of the Ising domain wall number n_{DW} by disrupting the balance of X and ZXZ in the first line of Eq. (1) [27]. The evolution operator over a small time step δt can be written in a Trotter product form as (setting $J = 0$ for simplicity)

$$U(\delta t) \approx e^{-i\Delta\delta t \sum_i Z_i} e^{i\lambda\delta t \sum_{i \text{ odd}} Z_{i-1} X_i Z_{i+1}} e^{-i(\lambda+\epsilon)\delta t \sum_i X_i} e^{i\lambda\delta t \sum_{i \text{ even}} Z_{i-1} X_i Z_{i+1}}. \quad (48)$$

The nontrivial ZXZ rotations above can be compiled using standard methods into four CNOT gates, two H gates, and an R_Z gate. An informative figure of merit to track while evolving the state $|\xi\rangle$ is the expectation value of, e.g., $Z_{i-1} X_i Z_{i+1}$, which exhibits coherent oscillations under evolution by H_0 that should decay in the presence of a nonzero ϵ . For $|\mathcal{S}_k\rangle$, it is informative to track the expectation value of M_z , which is not conserved under H_0 for generic initial states but is conserved for the initial state $|\mathcal{S}_k\rangle$; thus tracking $\langle M_z(t) \rangle$ provides information about the fidelity decay of $|\mathcal{S}_k\rangle$. These calculations are in principle straightforward to carry out on present-day quantum computers, but the QPU results in this work demonstrate that a significant fraction of the QPU’s coherence budget will be expended on state preparation. Thus, any near-term demonstration of this calculation must employ substantial problem-tailored optimizations and error mitigation strategies, including symmetry protection [98, 99], which should be explored in future work.

Acknowledgments

This material is based upon work supported by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Superconducting Quantum Materials and Systems Center (SQMS) under the contract No. DE-AC02-07CH11359 (E.G., A.C.Y.L., T.I., P.P.O., J.K., D.M.K.). M.S.A. acknowledges support from USRA NASA Academic Mission Services under contract No. NNA16BD14C with NASA, with this work funded under the NASA-DOE interagency agreement SAA2-403602 governing NASA’s work as part of the SQMS center, as was

³A version of this calculation in a case where the initial state is a product state was performed in Ref. [54].




Figure 17: Energy difference between instantaneous ground and first excited states, $E_1 - E_0$, versus interpolation parameter s/T . The total number of spins are $N = 4$ (blue), $N = 6$ (red), $N = 8$ (green), $N = 10$ (black), $N = 12$ (magenta), $N = 14$ (cyan), and $N = 16$ (brown). The minimal energy gap occurs at $s = T$ and approaches the exact value $1 - 1/\sqrt{2} = 0.292 \dots$ (dashed horizontal line) for large N . This is also shown in the inset, which contains the final energy gap between the first excited and the ground state at $s = T$ versus $1/N$. The red line shows a fit to a quadratic function with y -intercept 0.2920, which is consistent with the analytical value of the gap.

A.K.'s work as part of the Feynman Quantum Academy internship program at USRA. Ames National Laboratory is operated for the U.S. Department of Energy by Iowa State University under Contract No. DE-AC02-07CH11358 (P.P.O., T.I.). Variational calculations and manuscript writing by A.R. were supported by NSF Award No. DMR-1945395. T.I. and A.R. acknowledge the hospitality of the Aspen Center for Physics, which is supported by NSF Grant No. PHY-1607611. This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725. We acknowledge helpful discussions with Andrew Arrasmith, Emanuele Dalla Torre, Bram Evert, Pouyan Ghaemi, Lesik Motrunich, Sanjay Moudgalya, Zlatko Papić, and Matt Reagor.

A Adiabatic Preparation of $|\xi\rangle$

In this Appendix, we investigate the possibility of preparing the state $|\xi = 1\rangle$ [Eq. (5)] by performing a linear interpolation between the simple paramagnetic Hamiltonian

$$H_X = \sum_{i=2}^{N-1} \frac{I - (-1)^i X_i}{2}, \quad (\text{A1})$$

and the target Hamiltonian H_P ,

$$H_P = \sum_{i=2}^{N-1} \left(P'_{i-1} P'_i + P'_i P'_{i+1} + P_{i-1} \frac{I - (-1)^i X_i}{2} P_{i+1} \right). \quad (\text{A2})$$

Note that H_P is simply $H_{\xi=1}/2$ [see Eq. (21)], plus terms involving the projectors $P'_i = |1\rangle_i \langle 1|_i$ that ensure the ground state is in the Fibonacci Hilbert space. Both Hamiltonians are positive semidefinite (i.e., their ground states have zero energy) and conserve the Z -basis projection of the first and last qubit, which can be fixed to be in the 0 state. The ground state of H_X is then the product state $|\psi(0)\rangle = |0\rangle \otimes |+-+\dots+-\rangle \otimes |0\rangle$, where $|\pm\rangle$ satisfy $(1 \mp X_i)|\pm\rangle = 0$, while the ground state of H_P is $|\xi = 1\rangle$. To adiabatically prepare the desired state, we evolve $|\psi(0)\rangle$ under the time-dependent interpolating Hamiltonian

$$H(s) = \left(1 - \frac{s}{T}\right) H_X + \frac{s}{T} H_P \quad (\text{A3})$$




Figure 18: Total adiabatic evolution time T_* needed to prepare state $|\xi = 1\rangle$ with 99% fidelity as a function of system size N . Results are obtained using exact diagonalization with a total number of steps $n_s = 10^3$ and $\delta s = T_*/n_s$, where T_* is the minimal T value to reach at least 99% fidelity.

such that $H(0) = H_X$ and $H(T) = H_P$. Provided $H(s)$ has an energy gap for all s , then if T is sufficiently large, the time-evolved state $|\psi(T)\rangle \approx |\xi = 1\rangle$ to a good approximation. The adiabatic time evolution was carried out with exact diagonalization of the Hamiltonian (A3). After $|\psi(0)\rangle$ is prepared, the operator $U(s, s + \delta s)$ is applied $n_s = T/(\delta s)$ times with $s = n\delta s$, $0 \leq n \leq n_s - 1$. Here, n_s is the number of iterations, T is the total evolution time and δs is the time separation. The time evolution operator $U(s, s + \delta s) = e^{-i(\delta s)H(s)}$ and the overlap $|\langle\psi(s)|\xi = 1\rangle|$ are calculated at each iteration.

In Fig. 17, we plot the difference between the instantaneous energies of the ground and first excited states $E_1 - E_0$ of $H(s)$ against the interpolation parameter, s/T , for several values of N . The minimum energy gap occurs at $s = T$ and appears to saturate to a constant as N increases. In fact, an exact calculation of the energy gap of H_P gives $1 - 1/\sqrt{2} = 0.292\dots$ [62], which appears to be in reasonable agreement with the limiting value of the numerically observed gap as obtained from a quadratic extrapolation in $1/N$ (see Fig. 17 inset). This indicates that high-fidelity adiabatic state preparation should be possible in a finite total time T . In Fig. 18, we plot the total time T_* required to achieve state preparation with 99% fidelity as a function of N . T_* is calculated by solving for the root of the function $f(T) = |\langle\psi(T)|\xi = 1\rangle| - 0.99$ using Newton's method. We find that T_* appears to approach a constant value $T_* \approx 25$. This is consistent with the behavior of the gap observed in Fig. 17.

These results lead to the conclusion that adiabatic state preparation should be possible for the state $|\xi = 1\rangle$. However, it is important to compare the quantum resources required to implement adiabatic evolution on a QPU against those of the methods discussed in Sec. 3. In particular, if the evolution operator is approximated using a Trotter product formula, the number of Trotter steps required to achieve Trotter error ϵ after evolution by a time t scales as t^2/ϵ . Inserting $t = T_* \approx 25$ and demanding a Trotter error $\epsilon = 0.01$ yields an estimated 62 500 Trotter steps. Thus the needed circuit depth is constant, but prohibitively large for execution on present-day QPUs. This fact is a well-known shortcoming of the adiabatic algorithm [100] which partially motivated the development of alternative approaches like QAOA [101].

B Derivation of Stochastic Preparation Circuit Recursion Relations

In this appendix, we prove the linear-depth circuit (11) preparing $|\xi\rangle$ and the associated recursion relation (12) for the preparation circuit rotation angles.

An alternative definition of $|\xi\rangle$ is given in Ref. [27] as

$$|\xi\rangle = \frac{1}{Z} \prod_{j=2}^{N-1} [1 + (-1)^j \xi P_{j-1} \sigma_j^+ P_{j+1}] |\Omega\rangle, \quad (\text{A4})$$

where Z is the normalization factor in Eq. (7). Without loss of generality, we choose the convention where $j = 2$ is the rightmost term in the above product, i.e., $\prod_{j=2}^{N-1} V_j |\Omega\rangle = V_{N-1} V_{N-2} \cdots V_2 |\Omega\rangle$.

In this convention, the $(j+1)$ -th qubit is always in its $|0\rangle$ state when applying the j -th term, and hence the projector P_{j+1} can be omitted in the above expression. This means that

$$|\xi\rangle = \frac{1}{Z} \prod_{j=2}^{N-1} [1 + (-1)^j \xi P_{j-1} \sigma_j^+] |\Omega\rangle. \quad (\text{A5})$$

With the definition of $|\xi; m = N-2\rangle$ [Eq. (10)], we can get rid of the two qubits in the $|0\rangle$ state at both ends to write

$$|\xi; m\rangle = \frac{1}{Z} \prod_{j=1}^m [1 + (-1)^{j+1} \xi P_{j-1} \sigma_j^+] |\Omega\rangle, \quad (\text{A6})$$

where we fix $P_0 \equiv 1$ to be consistent with open boundary conditions. Note that in this equation, the j -th qubit of $|\xi; m\rangle$ corresponds to the $(j+1)$ -th qubit of $|\xi\rangle$.

To facilitate the derivation, we recall the projector $P'_j = |1\rangle_j \langle 1|_j = (1 - Z_j)/2$ and introduce $x_j = (-1)^{j+1} \xi$ to rewrite the equation as

$$|\xi; m\rangle = \frac{1}{Z} \prod_{j=1}^m [P'_{j-1} (1) + P_{j-1} (1 + x_j \sigma_j^+)] |\Omega\rangle, \quad (\text{A7})$$

where $P'_0 = 0$ for open boundary conditions. Note that to prepare the state $|\tilde{\xi}; m\rangle$, one can simply pick a different definition of x_j to be $x_j = \xi$.




Figure 19: The QPU experiment results to prepare $|\xi = 1\rangle$ on “ibmq_montreal” device with $m = 2, 4, 6, 12$ and $N = 14$. The $|\xi = 1\rangle$ states are prepared on qubits [26, 25, 22, 19, 16, 14, 13, 12, 10, 7, 4, 1, 2, 3] with ancilla qubits [0, 6, 15, 11, 20, 24]. We show the raw results (green diamond lines) without error mitigation to compare with the exact results (dotted line). Each data point is obtained with eight samples and 24576 shots per sample, and the error bars denote the standard deviation over eight samples.

The term in square brackets in the above equation is in the form of a controlled-rotation gate except that $(1 + x_j \sigma_j^+)$ is not a unitary operator. To make it unitary, we first introduce a set of real constants ϕ_j whose values will be determined later in this subsection. We can divide each term in the product by ϕ_j to get

$$|\xi; m\rangle = \frac{1}{Z'} \prod_{j=1}^m [P'_{j-1} \left(\frac{1}{\phi_j}\right) + P_{j-1} \frac{1}{\phi_j} (1 + x_j \sigma_j^+)] |\Omega\rangle,$$

where $Z' = Z / \prod_{j=1}^m \phi_j$. We recognize that the local state $|1\rangle$ is only generated by the raising operator σ_j^+ . Hence, the factor $1/\phi_j$ associated with P'_{j-1} can be moved to be associated with σ_{j-1}^+ in the previous product term. We can then write

$$|\xi; m\rangle = \frac{1}{Z'} \prod_{j=1}^m \left[P'_{j-1} + P_{j-1} \frac{1}{\phi_j} \left(1 + \frac{x_j}{\phi_{j+1}} \sigma_j^+ \right) \right] |\Omega\rangle,$$




Figure 20: The QPU experiment results to prepare $|\xi = 1\rangle$ on “ibmq_montreal” device with $m = 2$ and $6 \leq N \leq 16$. The experimental setup is the same as described in Fig. 19.

where $\phi_{m+1} = 1$. Since σ_j^+ only acts on $|0\rangle$ in this equation, we have $\sigma_j^+|0\rangle = -i\sigma_j^y|0\rangle$ where $\sigma_j^y = -i|0\rangle_j\langle 1|_j + i|1\rangle_j\langle 0|_j$. This allows us to write

$$|\xi; m\rangle = \frac{1}{Z'} \prod_{j=1}^m \left[P'_{j-1} + P_{j-1} \frac{1}{\phi_j} \left(1 - i \frac{x_j}{\phi_{j+1}} \sigma_j^y \right) \right] |\Omega\rangle. \quad (\text{A8})$$

If we pick

$$\phi_j = \sqrt{1 + \left| \frac{x_j}{\phi_{j+1}} \right|^2} \quad \text{and} \quad \phi_{m+1} = 1, \quad (\text{A9})$$

the operator associated with P_{j-1} is unitary such that

$$\frac{1}{\phi_j} \left(1 - i \frac{x_j}{\phi_{j+1}} \sigma_j^y \right) = \cos \frac{\theta_j}{2} - i \sin \frac{\theta_j}{2} \sigma_j^y = R_{Y,j}(\theta_j),$$

where $R_{Y,j}$ is the Y -rotation gate on the j -th qubit and $\theta_j = 2 \arg(1 + ix_j/\phi_{j+1})$ is the rotation angle. Rewriting Eq. (A8) in terms of Y -rotation gates, we have

$$|\xi; m\rangle = \frac{1}{Z'} \prod_{j=1}^m [P'_{j-1} + P_{j-1} R_{Y,j}(\theta_j)] |\Omega\rangle. \quad (\text{A10})$$

Since all terms are unitary, $\frac{1}{Z'}$ is just a global phase factor not relevant for the state preparation and hence can be dropped. Note also that $P'_0 = 0$ and $P_0 = 1$ and hence the $j = 1$ term is just a Y -rotation gate. Further rewriting the above equation using the controlled- Y rotation gates, we arrive at Eq. (11) with θ_j given by Eq. (12).

C QPU results for the complete constant-depth circuit with postselection on a IBM processor

In this appendix, we present the QPU results of implementing the complete constant-depth circuit (11) with postselection shown in Fig. 2 to prepare the superposition state $|\xi\rangle$ on an IBM quantum processor without introducing the fragmentation reduction used in Section 3.3.

We compile the constant-depth and postselection circuits into native gates respecting the connectivity of the “ibmq_montreal” device. In particular, each controlled Y -rotations (Toffoli gate) is decomposed into two (10) CNOT gates and single-qubit gates. The compiled circuit consists of $2 \frac{(N-2)(m-1)}{m} + 10 \left(\frac{N-2}{m} - 1 \right)$ CNOT gates and has a circuit depth of $2(m-1) + 10$ (counting only CNOT gates). Similar to the experiment described in Section 3.3, we measure the Bhattacharyya distance, the expectation value of the parent Hamiltonian H_ξ and the postselection success probability on “ibmq_montreal” for various block size m and system size N . The QPU results are shown in Figs. 19 and 20, and each data point is measured with eight samples and 24576 shots per sample.

The performance of the state preparation protocol improves with decreasing block size m indicated by the decreasing deviation between the measure values and the ideal value (0) of the Bhattacharyya distance and energy shown in Fig. 19(a) and (b). This comes with the cost of decreasing the postselection success probability shown in Fig. 19(c). The improvement is expected since the depth of the circuit decreases with decreasing m even though the number of CNOT increases. The protocol's performance deteriorates with increasing systems size N as shown in Fig. 20 since the circuit depth and the number of CNOT increase with N . The trends in varying m and N discussed in this appendix are consistent with the results shown in Section 3.3 using the fragmented circuits.

References

- [1] J. M. Deutsch. "Quantum statistical mechanics in a closed system". *Phys. Rev. A* **43**, 2046–2049 (1991).
- [2] Mark Srednicki. "Chaos and quantum thermalization". *Phys. Rev. E* **50**, 888–901 (1994).
- [3] Luca D'Alessio, Yariv Kafri, Anatoli Polkovnikov, and Marcos Rigol. "From quantum chaos and eigenstate thermalization to statistical mechanics and thermodynamics". *Adv. Phys.* **65**, 239–362 (2016).
- [4] Joshua M Deutsch. "Eigenstate thermalization hypothesis". *Rep. Prog. Phys.* **81**, 082001 (2018).
- [5] M. Rigol, V. Dunjko, and M. Olshanii. "Thermalization and its mechanism for generic isolated quantum systems". *Nature* **452**, 854 (2008).
- [6] Adam M. Kaufman, M. Eric Tai, Alexander Lukin, Matthew Rispoli, Robert Schittko, Philipp M. Preiss, and Markus Greiner. "Quantum thermalization through entanglement in an isolated many-body system". *Science* **353**, 794–800 (2016).
- [7] Christian Gross and Immanuel Bloch. "Quantum simulations with ultracold atoms in optical lattices". *Science* **357**, 995–1001 (2017).
- [8] C. Monroe, W. C. Campbell, L.-M. Duan, Z.-X. Gong, A. V. Gorshkov, P. W. Hess, R. Islam, K. Kim, N. M. Linke, G. Pagano, P. Richerme, C. Senko, and N. Y. Yao. "Programmable quantum simulations of spin systems with trapped ions". *Rev. Mod. Phys.* **93**, 025001 (2021).
- [9] Qingling Zhu, Zheng-Hang Sun, Ming Gong, Fusheng Chen, Yu-Ran Zhang, Yulin Wu, Yangsen Ye, Chen Zha, Shaowei Li, Shaojun Guo, Haoran Qian, He-Liang Huang, Jiale Yu, Hui Deng, Hao Rong, Jin Lin, Yu Xu, Lihua Sun, Cheng Guo, Na Li, Futian Liang, Cheng-Zhi Peng, Heng Fan, Xiaobo Zhu, and Jian-Wei Pan. "Observation of thermalization and information scrambling in a superconducting quantum processor". *Phys. Rev. Lett.* **128**, 160502 (2022).
- [10] J.-H. Wang, T.-Q. Cai, X.-Y. Han, Y.-W. Ma, Z.-L. Wang, Z.-H. Bao, Y. Li, H.-Y. Wang, H.-Y. Zhang, L.-Y. Sun, Y.-K. Wu, Y.-P. Song, and L.-M. Duan. "Information scrambling dynamics in a fully controllable quantum simulator". *Phys. Rev. Research* **4**, 043141 (2022).
- [11] Xiao Mi, Pedram Roushan, Chris Quintana, Salvatore Mandrà, Jeffrey Marshall, Charles Neill, Frank Arute, Kunal Arya, Juan Atalaya, Ryan Babbush, Joseph C. Bardin, Rami Barends, Joao Basso, Andreas Bengtsson, Sergio Boixo, Alexandre Bourassa, Michael Broughton, Bob B. Buckley, David A. Buell, Brian Burkett, Nicholas Bushnell, Zijun Chen, Benjamin Chiaro, Roberto Collins, William Courtney, Sean Demura, Alan R. Derk, Andrew Dunsworth, Daniel Eppens, Catherine Erickson, Edward Farhi, Austin G. Fowler, Brooks Foxen, Craig Gidney, Marissa Giustina, Jonathan A. Gross, Matthew P. Harrigan, Sean D. Harrington, Jeremy Hilton, Alan Ho, Sabrina Hong, Trent Huang, William J. Huggins, L. B. Ioffe, Sergei V. Isakov, Evan Jeffrey, Zhang Jiang, Cody Jones, Dvir Kafri, Julian Kelly, Seon Kim, Alexei Kitaev, Paul V. Klimov, Alexander N. Korotkov, Fedor Kostritsa, David Landhuis, Pavel Laptev, Erik Lucero, Orion Martin, Jarrod R. McClean, Trevor McCourt, Matt McEwen, Anthony Megrant, Kevin C. Miao, Masoud Mohseni, Shirin Montazeri, Wojciech Mruczkiewicz, Josh Mutus, Ofer Naaman, Matthew Neeley, Michael Newman, Murphy Yuezhen Niu, Thomas E. O'Brien, Alex Opremcak, Eric Ostby, Balint Pato, Andre Petukhov, Nicholas Redd, Nicholas C. Rubin, Daniel Sank, Kevin J. Satzinger, Vladimir Shvarts, Doug Strain, Marco Szalay, Matthew D. Trevithick, Benjamin Villalonga, Theodore

- White, Z. Jamie Yao, Ping Yeh, Adam Zalcman, Hartmut Neven, Igor Aleiner, Kostyantyn Kechedzhi, Vadim Smelyanskiy, and Yu Chen. “Information scrambling in quantum circuits”. *Science* **374**, 1479–1483 (2021).
- [12] Anatoli Polkovnikov, Krishnendu Sengupta, Alessandro Silva, and Mukund Vengalattore. “Colloquium: Nonequilibrium dynamics of closed interacting quantum systems”. *Rev. Mod. Phys.* **83**, 863–883 (2011).
 - [13] Lev Vidmar and Marcos Rigol. “Generalized Gibbs ensemble in integrable lattice models”. *Journal of Statistical Mechanics: Theory and Experiment* **2016**, 064007 (2016).
 - [14] Rahul Nandkishore and David A. Huse. “Many-body localization and thermalization in quantum statistical mechanics”. *Annu. Rev. Condens. Matter Phys.* **6**, 15–38 (2015).
 - [15] Ehud Altman and Ronen Vosk. “Universal dynamics and renormalization in many-body-localized systems”. *Annu. Rev. Condens. Matter Phys.* **6**, 383–409 (2015).
 - [16] Dmitry A. Abanin, Ehud Altman, Immanuel Bloch, and Maksym Serbyn. “Colloquium: Many-body localization, thermalization, and entanglement”. *Rev. Mod. Phys.* **91**, 021001 (2019).
 - [17] Maksym Serbyn, Dmitry A Abanin, and Zlatko Papić. “Quantum many-body scars and weak breaking of ergodicity”. *Nature Physics* **17**, 675–685 (2021).
 - [18] Sanjay Moudgalya, B. Andrei Bernevig, and Nicolas Regnault. “Quantum many-body scars and Hilbert space fragmentation: a review of exact results”. *Reports on Progress in Physics* **85**, 086501 (2022). [arXiv:2109.00548](https://arxiv.org/abs/2109.00548).
 - [19] Anushya Chandran, Thomas Iadecola, Vedika Khemani, and Roderich Moessner. “Quantum many-body scars: A quasiparticle perspective”. *Annual Review of Condensed Matter Physics* **14**, 443–469 (2023).
 - [20] Sanjay Moudgalya, Stephan Rachel, B. Andrei Bernevig, and Nicolas Regnault. “Exact excited states of nonintegrable models”. *Phys. Rev. B* **98**, 235155 (2018).
 - [21] Sanjay Moudgalya, Nicolas Regnault, and B. Andrei Bernevig. “Entanglement of exact excited states of Affleck-Kennedy-Lieb-Tasaki models: Exact results, many-body scars, and violation of the strong eigenstate thermalization hypothesis”. *Phys. Rev. B* **98**, 235156 (2018).
 - [22] Hannes Bernien, Sylvain Schwartz, Alexander Keesling, Harry Levine, Ahmed Omran, Hannes Pichler, Soonwon Choi, Alexander S Zibrov, Manuel Endres, Markus Greiner, et al. “Probing many-body dynamics on a 51-atom quantum simulator”. *Nature* **551**, 579 (2017).
 - [23] Christopher J Turner, Alexios A Michailidis, Dmitry A Abanin, Maksym Serbyn, and Zlatko Papić. “Weak ergodicity breaking from quantum many-body scars”. *Nature Physics* **14**, 745–749 (2018).
 - [24] C. J. Turner, A. A. Michailidis, D. A. Abanin, M. Serbyn, and Z. Papić. “Quantum scarred eigenstates in a Rydberg atom chain: Entanglement, breakdown of thermalization, and stability to perturbations”. *Phys. Rev. B* **98**, 155134 (2018).
 - [25] D. Bluvstein, A. Omran, H. Levine, A. Keesling, G. Semeghini, S. Ebadi, T. T. Wang, A. A. Michailidis, N. Maskara, W. W. Ho, S. Choi, M. Serbyn, M. Greiner, V. Vuletić, and M. D. Lukin. “Controlling quantum many-body dynamics in driven Rydberg atom arrays”. *Science* **371**, 1355–1359 (2021).
 - [26] Michael Schecter and Thomas Iadecola. “Weak Ergodicity Breaking and Quantum Many-Body Scars in Spin-1 XY Magnets”. *Phys. Rev. Lett.* **123**, 147201 (2019).
 - [27] Thomas Iadecola and Michael Schecter. “Quantum many-body scar states with emergent kinetic constraints and finite-entanglement revivals”. *Phys. Rev. B* **101**, 024306 (2020).
 - [28] Nicholas O’Dea, Fiona Burnell, Anushya Chandran, and Vedika Khemani. “From tunnels to towers: Quantum scars from Lie algebras and q -deformed Lie algebras”. *Phys. Rev. Research* **2**, 043305 (2020).
 - [29] K. Pakrouski, P. N. Pallegar, F. K. Popov, and I. R. Klebanov. “Many-Body Scars as a Group Invariant Sector of Hilbert Space”. *Phys. Rev. Lett.* **125**, 230602 (2020).
 - [30] Sanjay Moudgalya, Edward O’Brien, B. Andrei Bernevig, Paul Fendley, and Nicolas Regnault. “Large classes of quantum scarred Hamiltonians from matrix product states”. *Phys. Rev. B* **102**, 085120 (2020).
 - [31] Jie Ren, Chenguang Liang, and Chen Fang. “Quasisymmetry Groups and Many-Body Scar Dynamics”. *Phys. Rev. Lett.* **126**, 120604 (2021).

- [32] Long-Hin Tang, Nicholas O’Dea, and Anushya Chandran. “Multimagnon quantum many-body scars from tensor operators”. *Phys. Rev. Res.* **4**, 043006 (2022).
- [33] Jie Ren, Chenguang Liang, and Chen Fang. “Deformed symmetry structures and quantum many-body scar subspaces”. *Phys. Rev. Research* **4**, 013155 (2022).
- [34] Christopher M. Langlett, Zhi-Cheng Yang, Julia Wildeboer, Alexey V. Gorshkov, Thomas Iadecola, and Shenglong Xu. “Rainbow scars: From area to volume law”. *Phys. Rev. B* **105**, L060301 (2022).
- [35] Julia Wildeboer, Christopher M. Langlett, Zhi-Cheng Yang, Alexey V. Gorshkov, Thomas Iadecola, and Shenglong Xu. “Quantum many-body scars from Einstein-Podolsky-Rosen states in bilayer systems”. *Phys. Rev. B* **106**, 205142 (2022).
- [36] Guo-Xian Su, Hui Sun, Ana Hudomal, Jean-Yves Desaulles, Zhao-Yu Zhou, Bing Yang, Jad C. Halimeh, Zhen-Sheng Yuan, Zlatko Papić, and Jian-Wei Pan. “Observation of many-body scarring in a Bose-Hubbard quantum simulator”. *Phys. Rev. Res.* **5**, 023010 (2023).
- [37] Daniel K. Mark and Olexei I. Motrunich. “ η -pairing states as true scars in an extended hubbard model”. *Phys. Rev. B* **102**, 075132 (2020).
- [38] Sanjay Moudgalya, Nicolas Regnault, and B. Andrei Bernevig. “ η -pairing in Hubbard models: From spectrum generating algebras to quantum many-body scars”. *Phys. Rev. B* **102**, 085140 (2020).
- [39] K. Pakrouski, P. N. Pallegar, F. K. Popov, and I. R. Klebanov. “Group theoretic approach to many-body scar states in fermionic lattice models”. *Phys. Rev. Research* **3**, 043156 (2021).
- [40] Jean-Yves Desaulles, Debasish Banerjee, Ana Hudomal, Zlatko Papić, Arnab Sen, and Jad C. Halimeh. “Weak ergodicity breaking in the Schwinger model”. *Phys. Rev. B* **107**, L201105 (2023).
- [41] Jean-Yves Desaulles, Ana Hudomal, Debasish Banerjee, Arnab Sen, Zlatko Papić, and Jad C. Halimeh. “Prominent quantum many-body scars in a truncated Schwinger model”. *Phys. Rev. B* **107**, 205112 (2023).
- [42] Maarten Van Damme, Torsten V. Zache, Debasish Banerjee, Philipp Hauke, and Jad C. Halimeh. “Dynamical quantum phase transitions in spin- $SU(1)$ quantum link models”. *Phys. Rev. B* **106**, 245110 (2022).
- [43] Jesse Osborne, Bing Yang, Ian P. McCulloch, Philipp Hauke, and Jad C. Halimeh. “Spin- S $U(1)$ Quantum Link Models with Dynamical Matter on a Quantum Simulator” (2023). [arXiv:2305.06368](https://arxiv.org/abs/2305.06368).
- [44] Pengfei Zhang, Hang Dong, Yu Gao, Liangtian Zhao, Jie Hao, Jean-Yves Desaulles, Qiujiang Guo, Jiachen Chen, Jinfeng Deng, Bobo Liu, Wenhui Ren, Yunyan Yao, Xu Zhang, Shibo Xu, Ke Wang, Feitong Jin, Xuhao Zhu, Bing Zhang, Hekang Li, Chao Song, Zhen Wang, Fangli Liu, Zlatko Papić, Lei Ying, H. Wang, and Ying-Cheng Lai. “Many-body Hilbert space scarring on a superconducting processor”. *Nature Physics* **19**, 120–125 (2023).
- [45] Sanjay Moudgalya and Olexei I. Motrunich. “Exhaustive Characterization of Quantum Many-Body Scars using Commutant Algebras” (2022). [arXiv:2209.03377](https://arxiv.org/abs/2209.03377).
- [46] Cheng-Ju Lin, Anushya Chandran, and Olexei I. Motrunich. “Slow thermalization of exact quantum many-body scar states under perturbations”. *Phys. Rev. Research* **2**, 033044 (2020).
- [47] Shun-Yao Zhang, Dong Yuan, Thomas Iadecola, Shenglong Xu, and Dong-Ling Deng. “Extracting Quantum Many-Body Scarred Eigenstates with Matrix Product States”. *Phys. Rev. Lett.* **131**, 020402 (2023).
- [48] Ulrich Schollwöck. “The density-matrix renormalization group in the age of matrix product states”. *Ann. Phys. (N. Y.)* **326**, 96–192 (2011).
- [49] Román Orús. “A practical introduction to tensor networks: Matrix product states and projected entangled pair states”. *Annals of Physics* **349**, 117–158 (2014).
- [50] David J. Luitz and Yevgeny Bar Lev. “The ergodic side of the many-body localization transition”. *Annalen der Physik* **529**, 1600350 (2017).
- [51] Seth Lloyd. “Universal Quantum Simulators”. *Science* **273**, 1073–1078 (1996).
- [52] Andrew M. Childs, Dmitri Maslov, Yunseong Nam, Neil J. Ross, and Yuan Su. “Toward the first quantum simulation with quantum speedup”. *Proceedings of the National Academy of Sciences* **115**, 9456–9461 (2018).

- [53] Andrew J Daley, Immanuel Bloch, Christian Kokail, Stuart Flannigan, Natalie Pearson, Matthias Troyer, and Peter Zoller. “Practical quantum advantage in quantum simulation”. *Nature* **607**, 667–676 (2022).
- [54] I-Chi Chen, Benjamin Burdick, Yongxin Yao, Peter P. Orth, and Thomas Iadecola. “Error-mitigated simulation of quantum many-body scars on quantum computers with pulse-level control”. *Phys. Rev. Res.* **4**, 043027 (2022).
- [55] Sambuddha Chattopadhyay, Hannes Pichler, Mikhail D. Lukin, and Wen Wei Ho. “Quantum many-body scars from virtual entangled pairs”. *Phys. Rev. B* **101**, 174308 (2020).
- [56] Daniel K. Mark, Cheng-Ju Lin, and Olexei I. Motrunich. “Unified structure for exact towers of scar states in the Affleck-Kennedy-Lieb-Tasaki and other models”. *Phys. Rev. B* **101**, 195131 (2020).
- [57] Oskar Vafek, Nicolas Regnault, and B. Andrei Bernevig. “Entanglement of Exact Excited Eigenstates of the Hubbard Model in Arbitrary Dimension”. *SciPost Phys.* **3**, 043 (2017).
- [58] Soonwon Choi, Christopher J. Turner, Hannes Pichler, Wen Wei Ho, Alexios A. Michailidis, Zlatko Papić, Maksym Serbyn, Mikhail D. Lukin, and Dmitry A. Abanin. “Emergent SU(2) Dynamics and Perfect Quantum Many-Body Scars”. *Phys. Rev. Lett.* **122**, 220603 (2019).
- [59] Andreas Bärttschi and Stephan Eidenbenz. “Deterministic Preparation of Dicke States”. In Leszek Antoni Gasieniec, Jesper Jansson, and Christos Levcopoulos, editors, *Fundamentals of Computation Theory*. Pages 126–139. Cham (2019). Springer International Publishing.
- [60] Umberto Borla, Ruben Verresen, Fabian Grusdt, and Sergej Moroz. “Confined Phases of One-Dimensional Spinless Fermions Coupled to Z_2 Gauge Theory”. *Phys. Rev. Lett.* **124**, 120503 (2020).
- [61] Maike Ostmann, Matteo Marcuzzi, Juan P. Garrahan, and Igor Lesanovsky. “Localization in spin chains with facilitation constraints and disordered interactions”. *Phys. Rev. A* **99**, 060101 (2019).
- [62] Igor Lesanovsky. “Liquid Ground State, Gap, and Excited States of a Strongly Correlated Spin Chain”. *Phys. Rev. Lett.* **108**, 105301 (2012).
- [63] D. Jaksch, J. I. Cirac, P. Zoller, S. L. Rolston, R. Côté, and M. D. Lukin. “Fast Quantum Gates for Neutral Atoms”. *Phys. Rev. Lett.* **85**, 2208–2211 (2000).
- [64] M. D. Lukin, M. Fleischhauer, R. Cote, L. M. Duan, D. Jaksch, J. I. Cirac, and P. Zoller. “Dipole Blockade and Quantum Information Processing in Mesoscopic Atomic Ensembles”. *Phys. Rev. Lett.* **87**, 037901 (2001).
- [65] Masaaki Nakamura, Zheng-Yuan Wang, and Emil J. Bergholtz. “Exactly Solvable Fermion Chain Describing a $\nu = 1/3$ Fractional Quantum Hall State”. *Phys. Rev. Lett.* **109**, 016401 (2012).
- [66] Sanjay Moudgalya, B. Andrei Bernevig, and Nicolas Regnault. “Quantum many-body scars in a Landau level on a thin torus”. *Phys. Rev. B* **102**, 195150 (2020).
- [67] Armin Rahmani, Kevin J. Sung, Harald Putterman, Pedram Roushan, Pouyan Ghaemi, and Zhang Jiang. “Creating and Manipulating a Laughlin-Type $\nu = 1/3$ Fractional Quantum Hall State on a Quantum Computer with Linear Depth Circuits”. *PRX Quantum* **1**, 020309 (2020).
- [68] Ammar Kirmani, Kieran Bull, Chang-Yu Hou, Vedika Saravanan, Samah Mohamed Saeed, Zlatko Papić, Armin Rahmani, and Pouyan Ghaemi. “Probing Geometric Excitations of Fractional Quantum Hall States on Quantum Computers”. *Phys. Rev. Lett.* **129**, 056801 (2022).
- [69] Jay Hubisz, Bharath Sambasivam, and Judah Unmuth-Yockey. “Quantum algorithms for open lattice field theory”. *Phys. Rev. A* **104**, 052420 (2021).
- [70] 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”. *Physical Review Research* **3**, 033002 (2021).
- [71] Nathanan Tantivasadakarn, Ryan Thorngren, Ashvin Vishwanath, and Ruben Verresen. “Long-range entanglement from measuring symmetry-protected topological phases” (2022). [arXiv:2112.01519](https://arxiv.org/abs/2112.01519).
- [72] Tsung-Cheng Lu, Leonardo A. Lessa, Isaac H. Kim, and Timothy H. Hsieh. “Measurement as a Shortcut to Long-Range Entangled Quantum Matter”. *PRX Quantum* **3**, 040337 (2022).
- [73] Aaron J. Friedman, Chao Yin, Yifan Hong, and Andrew Lucas. “Locality and error correction in quantum dynamics with measurement” (2022) [arXiv:2205.14002](https://arxiv.org/abs/2205.14002).

- [74] Kevin C. Smith, Eleanor Crane, Nathan Wiebe, and S. M. Girvin. “Deterministic constant-depth preparation of the AKLT state on a quantum processor using fusion measurements” (2022)[arXiv:2210.17548](#).
- [75] Frank Pollmann, Ari M. Turner, Erez Berg, and Masaki Oshikawa. “Entanglement spectrum of a topological phase in one dimension”. *Phys. Rev. B* **81**, 064439 (2010).
- [76] Frank Pollmann, Erez Berg, Ari M. Turner, and Masaki Oshikawa. “Symmetry protection of topological phases in one-dimensional quantum spin systems”. *Phys. Rev. B* **85**, 075125 (2012).
- [77] Alistair W. R. Smith, Kiran E. Khosla, Chris N. Self, and M. S. Kim. “Qubit readout error mitigation with bit-flip averaging”. *Sci. Adv.* **7**, abi8009 (2021). [arXiv:2106.05800](#).
- [78] Joel J. Wallman and Joseph Emerson. “Noise tailoring for scalable quantum computation via randomized compiling”. *Phys. Rev. A* **94**, 052325 (2016).
- [79] Benjamin Nachman, Miroslav Urbanek, Wibe A. de Jong, and Christian W. Bauer. “Unfolding quantum computer readout noise”. *npj Quantum Information* **6** (2020).
- [80] Deanna M. Abrams, Nicolas Didier, Blake R. Johnson, Marcus P. da Silva, and Colm A. Ryan. “Implementation of the XY interaction family with calibration of a single pulse”. *Nature Electronics* **3**, 744 (2020).
- [81] Alexander D Hill, Mark J Hodson, Nicolas Didier, and Matthew J Reagor. “Realization of arbitrary doubly-controlled quantum phase gates” (2021). [arXiv:2108.01652](#).
- [82] Tianyi Peng, Aram W. Harrow, Maris Ozols, and Xiaodi Wu. “Simulating large quantum circuits on a small quantum computer”. *Physical Review Letters* **125** (2020).
- [83] Daniel T. Chen, Zain H. Saleem, and Michael A. Perlin. “Quantum Divide and Conquer for Classical Shadows” (2022). [arXiv:2212.00761](#).
- [84] 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**, 024001 (2019).
- [85] Shi-Ju Ran. “Encoding of matrix product states into quantum circuits of one- and two-qubit gates”. *Phys. Rev. A* **101**, 032310 (2020).
- [86] Gregory M. Crosswhite and Dave Bacon. “Finite automata for caching in matrix product algorithms”. *Phys. Rev. A* **78**, 012356 (2008).
- [87] Michael A. Nielsen and Isaac L. Chuang. “Quantum computation and quantum information: 10th anniversary edition”. *Cambridge University Press*. (2010).
- [88] Vivek V. Shende and Igor L. Markov. “On the CNOT-cost of TOFFOLI gates” (2008). [arXiv:0803.2316](#).
- [89] Zhi-Cheng Yang, Fangli Liu, Alexey V. Gorshkov, and Thomas Iadecola. “Hilbert-Space Fragmentation from Strict Confinement”. *Phys. Rev. Lett.* **124**, 207602 (2020).
- [90] Qiskit contributors. “Qiskit: An open-source framework for quantum computing” (2023).
- [91] Ludmila Botelho, Adam Glos, Akash Kundu, Jarosław Adam Mischczak, Özlem Salehi, and Zoltán Zimborás. “Error mitigation for variational quantum algorithms through mid-circuit measurements”. *Phys. Rev. A* **105**, 022441 (2022).
- [92] Emanuele G. Dalla Torre and Matthew J. Reagor. “Simulating the Interplay of Particle Conservation and Long-Range Coherence”. *Phys. Rev. Lett.* **130**, 060403 (2023). [arXiv:2206.08386](#).
- [93] Sam McArdle, Tyson Jones, Suguru Endo, Ying Li, Simon C Benjamin, and Xiao Yuan. “Variational ansatz-based quantum simulation of imaginary time evolution”. *npj Quantum Inf.* **5**, 75 (2019).
- [94] Mario Motta, Chong Sun, Adrian TK Tan, Matthew J O’Rourke, Erika Ye, Austin J Minnich, Fernando GSL Brandão, and Garnet Kin-Lic Chan. “Determining eigenstates and thermal states on a quantum computer using quantum imaginary time evolution”. *Nat. Phys.* **16**, 205–210 (2020).
- [95] Niladri Gomes, Feng Zhang, Noah F Berthusen, Cai-Zhuang Wang, Kai-Ming Ho, Peter P Orth, and Yong-Xin Yao. “Efficient step-merged quantum imaginary time evolution algorithm for quantum chemistry”. *J. Chem. Theory Comput.* **16**, 6256–6266 (2020).
- [96] Niladri Gomes, Anirban Mukherjee, Feng Zhang, Thomas Iadecola, Cai-Zhuang Wang, Kai-Ming Ho, Peter P Orth, and Yong-Xin Yao. “Adaptive Variational Quantum Imaginary

- Time Evolution Approach for Ground State Preparation”. [Adv. Quantum Technol.](#) **4**, 2100114 (2021).
- [97] Shun-Yao Zhang, Dong Yuan, Thomas Iadecola, Shenglong Xu, and Dong-Ling Deng. “Extracting quantum many-body scarred eigenstates with matrix product states”. [Phys. Rev. Lett.](#) **131**, 020402 (2023).
- [98] Jad C. Halimeh, Luca Barbiero, Philipp Hauke, Fabian Grusdt, and Annabelle Bohrdt. “Robust quantum many-body scars in lattice gauge theories”. [Quantum](#) **7**, 1004 (2023).
- [99] Minh C. Tran, Yuan Su, Daniel Carney, and Jacob M. Taylor. “Faster Digital Quantum Simulation by Symmetry Protection”. [PRX Quantum](#) **2**, 010323 (2021).
- [100] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, and Michael Sipser. “Quantum Computation by Adiabatic Evolution” (2000). [arXiv:quant-ph/0001106](#).
- [101] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. “A quantum approximate optimization algorithm” (2014)[arXiv:1411.4028](#).