

Programmable Simulations of Molecules and Materials with Reconfigurable Quantum Processors

Nishad Maskara,^{1,2} Stefan Ostermann,¹ James Shee,^{3,4} Marcin Kalinowski,¹ Abigail McClain Gomez,¹ Rodrigo Araiza Bravo,¹ Derek S. Wang,⁵ Anna I. Krylov,⁶ Norman Y. Yao,¹ Martin Head-Gordon,^{2,3} Mikhail D. Lukin,¹ and Susanne F. Yelin¹

¹*Department of Physics, Harvard University, Cambridge, MA 02138, USA*

²*Lawrence Berkeley National Lab, Berkeley, CA 94720, USA*

³*College of Chemistry, University of California Berkeley, Berkeley, CA 94720, USA*

⁴*Department of Chemistry, Rice University, Houston, TX 77005, USA*

⁵*Harvard John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138, USA*

⁶*Department of Chemistry, University of Southern California, Los Angeles, CA 90089, USA*

(Dated: December 6, 2023)

Simulations of quantum chemistry and quantum materials are believed to be among the most important potential applications of quantum information processors, but realizing practical quantum advantage for such problems is challenging. Here, we introduce a simulation framework for strongly correlated quantum systems that can be represented by model spin Hamiltonians. Our approach leverages reconfigurable qubit architectures to programmably simulate real-time dynamics and introduces an algorithm for extracting chemically relevant spectral properties via classical co-processing of quantum measurement results. We develop a digital-analog simulation toolbox for efficient Hamiltonian time evolution utilizing digital Floquet engineering and hardware-optimized multi-qubit operations to accurately realize complex spin-spin interactions, and as an example present an implementation proposal based on Rydberg atom arrays. Then, we show how detailed spectral information can be extracted from these dynamics through snapshot measurements and single-ancilla control, enabling the evaluation of excitation energies and finite-temperature susceptibilities from a single-dataset. To illustrate the approach, we show how this method can be used to compute key properties of a polynuclear transition-metal catalyst and 2D magnetic materials.

A major thrust of quantum chemistry and material science involves the quantitative prediction of electronic structure properties of molecules and materials. While powerful computational techniques have been developed over the past decades, especially for weakly correlated systems [1–4], the development of tools for understanding and predicting the properties of materials that feature strongly correlated electrons remains a challenge [5–7]. Quantum computing is a promising route to efficiently capturing such quantum correlations [8–10], and algorithms for Hamiltonian simulation and energy estimation [11] with good asymptotic scaling have been developed. However, existing methods for simulating large-scale electronic structure problems are prohibitively expensive to run on near-term quantum hardware [12], highlighting the need for more efficient approaches.

One approach to capturing the complexity of strongly correlated systems utilizes model Hamiltonians [13], such as the generalized Ising, Heisenberg, and Hubbard models, which describe the interactions between the active degrees of freedom at low temperatures. Like other coarse-graining or effective Hamiltonian approaches [14], model parameters can be computed from an *ab initio* electronic structure problem using a number of classical techniques [15–20]. Furthermore, model Hamiltonians exhibit features such as low-degree connectivity that simplify implementation, making them particularly promising candidates for quantum simulation [21–26]. Though approximate, simplified model Hamiltonians have proved valuable in analyzing strongly correlated problems [27–29] for small system sizes, where accurate but costly classical methods can be applied. However, as the system size increases, classical numerical methods struggle to reliably solve strongly correlated model systems, as the relevant low-

energy states often exhibit a large degree of entanglement. In this Article, we focus on the programmable quantum simulations of *spin* models. These correspond to a class of Hamiltonians that describe compounds where unpaired electrons become localized at low-temperatures and can therefore be represented as effective local spins with $S \geq 1/2$. These include many polynuclear transition metal compounds and materials containing d- and f-block elements, which play a central role in chemical catalysis and magnetism [20, 28–34].

Recent advances in quantum simulation [35, 36] have enabled the study of paradigmatic model Hamiltonians with local connectivity. In particular, experiments have probed non-equilibrium quantum dynamics [37–39], exotic forms of emergent magnetism [40–43], and long-range entangled topological matter [44, 45], in regimes that push the limits of state-of-the-art classical simulations [46]. The model Hamiltonians describing realistic molecules and materials however often contain more complex features, including anisotropy, non-locality, and higher-order interactions [29], demanding a higher degree of programmability [18]. While universal quantum computers can in principle simulate such systems, standard implementations based on local two-qubit gates require large circuit depths [23] to realize complex interactions and long-range connectivity. Thus, for optimal performance in devices with limited coherence, it is essential to utilize hardware-efficient capabilities to simulate such systems.

In what follows we introduce a framework to simulate model spin Hamiltonians (Fig. 1) on reconfigurable quantum devices. The approach combines two elements. First, we describe a hybrid digital-analog simulation toolbox for realizing complex spin interactions, which combines the programmability of digital simulation with the efficiency of hardware-




FIG. 1. **Model Hamiltonian approach to quantum simulation of strongly-correlated matter.** (a) The procedure starts with a description of the target molecule or material structure, whose electronic structure problem is reduced using classical computational chemistry techniques to a simpler effective Hamiltonian that captures the relevant low-energy behaviors. (b) Here we study problems which are modeled by spin Hamiltonians with potentially non-local connectivity and generic on-site spin $S \geq 1/2$, where each spin is comprised of localized, unpaired electrons in the original molecule. The key simplification comes from capturing charge fluctuations perturbatively, which is a good approximation in certain contexts. (c) Programmable quantum simulation is then used to calculate properties of the model Hamiltonian. We develop a simulation framework, based on encoding spins into clusters of qubits, that can be readily implemented on existing hardware. The toolbox enables efficient generation of complex spin interactions by leveraging dynamical reconfigurability and hardware-optimized multi-qubit gates (Figs. 2 and 3). The quantum simulator performs time-evolution under the spin Hamiltonian for various simulation times t , and each qubit is projectively measured to produce an set of snapshots. Subsequent classical processing extracts properties like the low-lying excitation spectrum and magnetic susceptibilities, all from the same dataset (see Figs. 4-6).

optimized multi-qubit analog operations. Then, we introduce an algorithm, dubbed “many-body spectroscopy” that leverages time-dynamics and snapshot measurements to extract detailed spectral information of the model Hamiltonian in a resource-efficient way [47]. We describe in detail how these methods can be implemented using Rydberg atom arrays [48–52], and discuss its applicability to other emerging platforms which can support multi-qubit control and dynamic, programmable connectivity [53], such as the ion QCCD architecture [54, 55]. Finally, we illustrate potential applications of the framework on model Hamiltonians describing a prototypical biochemical catalyst and 2D materials.

Engineering spin Hamiltonians

The general Hamiltonian we consider is

$$H = \sum_{i,\alpha} B_i^\alpha \hat{S}_i^\alpha + \sum_{ij,\alpha\beta} J_{ij}^{\alpha\beta} \hat{S}_i^\alpha \hat{S}_j^\beta + \sum_{ijk,\alpha\beta\gamma} K_{ijk}^{\alpha\beta\gamma} \hat{S}_i^\alpha \hat{S}_j^\beta \hat{S}_k^\gamma + \text{higher order}, \quad (1)$$

where $\hat{S}_i^\alpha$, $\alpha = x, y, z$ are spin- S operators ($S \geq 1/2$) acting on the i -th spin, and the interaction coefficients ($J_{ij}^{\alpha\beta}$, $K_{ijk}^{\alpha\beta\gamma}$, etc.) are potentially long-range. Hamiltonians of this form can capture the effects of many processes arising in physical com-

pounds, including super-exchange, spin-orbit coupling, ring-exchange, and more [29]. In our approach, spin- S variables are encoded into the collective spin of $2S$ qubits, such that the i -th spin in (1) is rewritten as

$$\hat{S}_i^\alpha = \sum_{a=1}^{2S_i} \hat{s}_{i,a}^\alpha, \quad (2)$$

where $\hat{s}_{i,a}^\alpha$ are the spin-1/2 operators of the a -th qubit in the i -th spin. Valid spin- S states live in the symmetric subspace with maximum total spin per site $\langle (\hat{S}_i)^2 \rangle = S_i(S_i + 1)$. While several alternate approaches to encoding spins with hardware-native qudits have been proposed recently [58–61], the cluster approach introduced here uses the same controls as qubit-based computations ensuring compatibility with existing setups [57], and naturally supports simulation of models with mixed on-site spin.

The core of our protocol involves applying a K step sequential evolution under simpler interaction Hamiltonians $H_i = \sum_{g \in G_i} h_{i,g}$, $i = 1, \dots, K$, acting on disconnected groups G_i of a few qubits each. The combined sequence realizes an effective Floquet Hamiltonian H_F which approximates (1). To controllably generate effective Hamiltonians, we use the average Hamiltonian approach. In the limit of small step-sizes $\tau \ll 1/K$, the evolution is well-approximated by $H_F^{(0)} = \frac{1}{K} \sum_{i=1}^K H_i$ to leading order, and contributions from higher-order terms are bounded [62].




FIG. 2. **Hardware-efficient implementation with neutral-atom tweezer arrays.** (a) Our protocol for programmable simulation of generic spin Hamiltonian is based on applying sequences of interactions between non-overlapping few-qubit groups. Here, we illustrate an implementation of a complex spin model using the dynamical projection approach. Spin- S_i variables are encoded in the collective spin of a cluster of $2S_i$ qubits. Then, interactions between spins are generated by evolving pairs of qubits from each cluster under an interaction Hamiltonian H_I . Second-order interactions J_{ij} act on two qubits, while higher-order interactions act on multiple qubits, such as the fourth-order coefficient L_{ijk} of $\hat{S}_i^\alpha \hat{S}_i^\beta \hat{S}_j \hat{S}_k$. Then, interactions are dynamically projected into the symmetric encoding space within each cluster by evolving under Hamiltonian H_P . The target large spin Hamiltonian H is simulated by alternately evolving under H_I and H_P . This protocol can be realized in any reconfigurable quantum processor. Here, we present an implementation for Rydberg atom arrays, which has two long-lived qubit states $|0\rangle, |1\rangle$, and an excited Rydberg state $|r\rangle$ with strong interactions. The interaction connectivity is dynamically changed by moving optical tweezers [56], and interactions are generated using all-to-all Rydberg blockade interactions within each cluster and simultaneous global driving of the qubit $\Omega_q(t)$ and Rydberg $\Omega_r(t)$ transitions. (b) Fast and efficient multi-qubit spin operations U_S, U_P are identified using optimal control to optimize pulse sequences. Gate times are measured in units of the Rydberg driving frequency ΩT , where the two-qubit CZ gate from Ref. [57] takes $\Omega T/2\pi \approx 1.2$. The alternating ansatz (solid lines) decomposes the target operations into symmetric diagonal gates and single qubit rotations, which can be individually optimized (see Methods). Simultaneous (dual) driving of both transitions (dotted lines) enables even faster realization of approximate U_S and U_P gates with ideal error rates below $\sim 10^{-3}$. (c) We compare against decomposition of U_S and U_P into a two-qubit gate set composed of CPhase gates and single-qubit rotations. Such a decomposition rapidly becomes very costly as the cluster size grows, in contrast to the optimized hardware-native operation.

In general, the performance of this approach will be limited by simulation errors, characterized by the difference between H_F and the target Hamiltonian H , and gate errors, determined by the hardware overhead required to implement individual evolutions $e^{-i\hbar i, g\tau}$. To mitigate the leading sources of error, we next develop Hamiltonian engineering protocols that leverage multi-qubit spin operations to realize (1) with short periodic sequences.

Dynamical projection with digital Floquet engineering

Our Hamiltonian engineering approach is based on the cluster encoding (2). The key idea is to first generate interactions in the full Hilbert space using a spin-1/2 (qubit) version H_I of (1), and then dynamically project back onto the encoding space. We select H_I such that projection recovers the target Hamiltonian, by mapping each n -body large spin interaction in (1) onto an equivalent n -body qubit interaction acting on representatives from the n -clusters (see Methods and Fig. 2a).

While this generates the target interactions between spin clusters, it also moves encoded states out of the symmetric subspace. Therefore, to prevent evolution under H_I from destroying the encoding, we alternately apply evolution under

$$H_P = -\lambda \sum_i P[(\hat{S}_i)], \quad (3)$$

composed of projectors $P[(\hat{S}_i)]$ onto the symmetric spin states, by applying multi-qubit gates within spin clusters. This Hamiltonian creates an energy gap λ between symmetric and non-symmetric states. For $\lambda \gg J_{ij}$, the static Hamiltonian $H_I + H_P$ is effectively projected into the ground-space of H_P at low-energies, realizing H . However, accurate Trotter simulation in this regime, by alternating H_I and H_P , requires two separations of time-scale between the step-size and the target interactions $\tau \ll \lambda^{-1} \ll J_{ij}^{-1}$, leading to a large number of gates.

Instead, we develop an approach which enables projection of H_I onto the ground-state of H_P with significantly fewer

gates, using ideas inspired by dynamical decoupling [63]. This is achieved by using large-angle rotations $e^{-i\tau H_P}$, where $\tau\lambda \sim 1$, to generate a time-dependent phase on the parts of H_I that couple encoded states to non-symmetric states. These phases cancel out on average, leaving the symmetric part which commutes with H_P . More precisely, the combined evolution is described by

$$U_F = \prod_{p=1}^{N_p} e^{-i\theta_p H_P} \prod_{j=1}^D e^{-i\tau H_{I,j}} \\ = \prod_{p=1}^{N_p} \prod_{j=1}^D e^{-i\tau H_{I,j}(k)} = e^{-iK\tau H_F}, \quad (4)$$

where H_I has been decomposed into D non-overlapping groups $H_{I,j}$, $K = N_p D$ is the full length of the sequence, and θ_p, τ parameterize the evolution times. In the second line, we transform into an interaction picture, such that intermediate terms are evolved by H_P with a cumulative phase $\Theta_k = \lambda^{-1} \sum_{k' < k} \theta_{k'}$; for the rotating frame to be periodic, we require $\Theta_K \pmod{2\pi} = 0$.

Then the transformed interactions can be written as (see Methods):

$$H_{I,j}(k) = e^{-i\Theta_k H_P} H_{I,j} e^{i\Theta_k H_P} \\ = \underbrace{H_{I,j}^{(0)}}_{\text{symmetric}} + \underbrace{\left(\sum_{n=1}^{n_{\max}} e^{-in\Theta_k} H_{I,j}^{(n)} + h.c. \right)}_{\text{symmetry violating}}. \quad (5)$$

Here, $H_I^{(n)}$ captures unwanted terms that change the total spin on n sites, and the symmetric part is the target Hamiltonian $H_I^{(0)} = H_T$. To cancel out symmetry-violating terms at leading order in τ , we choose an appropriate sequence of cumulative phases $\Theta_k = \lambda \sum_{k' < k} \theta_{k'}$, such that $\sum_k e^{-i\Theta_k} = 0$. When $n_{\max} = 2$, i.e., up to two-spin interactions, the sequence $\Theta = (0, 2\pi/3, 4\pi/3)$ satisfies this condition. This sequence, being p -independent, produces a Floquet Hamiltonian which commutes with H_P in a dressed frame [64], and therefore preserves the encoding up to exponentially long times in τ^{-1} .

This approach provides significant benefits in simulating complex spin-models, where the number of overlapping terms in H grows rapidly with parameters such as the spin-size S , number of interactions per spin d , and interaction weight $n_{\max}$. In this regime, a Trotter decomposition of H into non-overlapping terms $h_{i,g}$ would require a sequence of length $K \geq d(2S)^{n_{\max}-1}$, where d measures the number of interactions each spin S_i is involved in. Instead, the projection approach produces decompositions of H into sequences of length $K = (n_{\max} + 1) \lceil \frac{d}{2S} \rceil$ (see Methods), leading to performance improvements of several orders of magnitude for models with large spins and higher-order interactions.

Similar tools can also be used to realize a large class of spin circuits, by generating different evolutions $H_F^{(i)}$ during each cycle. This effectively implements discrete time-dependent evolution

$$U_{\text{circ}} = \prod_{i=1}^T e^{-i\tau_i H_F^{(i)}}. \quad (6)$$

Variational optimization can be further used to engineer higher-order terms in H_F , enabling generation of more complex spin gates at no additional cost (see Extended Data Fig. 9). Although the classical variational optimization procedure is limited to small circuits, Hamiltonian learning protocols could be used to perform larger-scale optimization on a quantum device directly [65, 66]. Such circuits can be used for operations besides Hamiltonian simulation, including state-preparation [21].

Hardware-efficient implementation

The digital Hamiltonian engineering sequence can, in principle, be realized on any universal quantum processor, but it is especially well suited for reconfigurable processors with native multi-qubit interactions [53, 55, 56, 69]. Neutral atom arrays are a particularly promising candidate for realizing these techniques, for which we develop a detailed implementation proposal. In this platform, two long-lived atomic states encode the qubit degree of freedom $\{|0\rangle, |1\rangle\}$, which can be individually manipulated with fidelity above 99.99% [70, 71]. Strong interactions between qubits are realized by coupling to a Rydberg state $|r\rangle$ [72, 73], which enables parallel multi-qubit operations [73], with state-of-the-art two-qubit gate fidelities exceeding 99.5% [57]. Further, qubits can be transported with high fidelity by moving optical tweezers [56], to realize arbitrary groupings G_i . By placing atoms sufficiently close together, atoms within a group can undergo strong all-to-all interactions, while interactions between groups can be made negligible by placing them far apart.

The key ingredient required for efficiently implementing H_I and H_P are hardware-efficient multi-qubit spin operations. We show how these can be realized by using pulse engineering to transform the native Rydberg-blockade interaction [72, 74] into the desired form. We illustrate this on two families of representative spin operations

$$U_S(\theta) = e^{-i\theta(\hat{S}^2/2S)}, \quad U_P(\theta) = e^{-i\theta P[\hat{S}^2]} \quad (7)$$

where $\hat{S}^2$ is the total-spin operator for a cluster of $2S$ atoms, and $P[\hat{S}^2]$ are the projectors appearing in (3).

One approach to engineering these operations is based on an ansatz which naturally extends Ref. [57], where U_S and U_P are found by optimizing an alternating sequence of diagonal phase gates and single-qubit rotations. As in prior




FIG. 3. Efficiency of Hamiltonian simulation framework. Estimates of the quantum simulation’s coherence time T_{sc} , in the target Hamiltonian’s units $||H_T||_{local}$ for various models. We consider Hamiltonian simulation implemented using the dual driving gates from Fig. 2b, and assume an error proportional to the gate time, such that $\Omega T/2\pi = 1$ incurs an error of 0.1% [57, 67, 68]. In all cases, we compare against an implementation using two-qubit CPhase gates with fidelity 99.9% and perfect single-qubit rotations (see Methods for detailed descriptions of the estimation procedure). (a) The first two two models are (i) the spin-1/2 Kagome Heisenberg model, and (ii) two interacting spin-5/2’s with Heisenberg and Dzyaloshinskii–Moriya (DM) terms, both of which are composed of only two-qubit interactions. In (i), a significant speedup is achieved by utilizing three-qubit multi-qubit gates $e^{-i\tau S^2}$, which more efficiently generates Heisenberg interactions and reduces the period of the Floquet cycle from $K = 4$ to $K = 2$. In (ii) improvement is achieved using dynamical projection, which reduces K from $2S$ to 2 but at the cost of additional multi-qubit gates. (b) Two complex spin-models which include spin-interactions up to (iii) bi-quadratic interactions $J_1(\mathbf{S}_i \cdot \mathbf{S}_j) + J_2(\mathbf{S}_i \cdot \mathbf{S}_j)^2$ and (iv) bi-quartic interactions $(\mathbf{S}_i \cdot \mathbf{S}_j)^4$. These correspond to four-body and eight-body qubit interactions respectively. In (iii), the dramatic speedup originates from using dynamical projection to significantly reduce the Floquet period, as well as the hardware-efficiency of a native four-qubit gate. The individual contribution to the speedup from both sources is also analyzed in Methods. In (iv), the speedup arises fully from the hardware-efficiency of native eight-qubit operations.

works [57, 75, 76], the pulse profiles generating symmetric diagonal operations can be obtained with numerical optimization via GRAPE [77]. For this alternating ansatz, we find a roughly linear scaling of total gate time T_{gate} with size of the cluster (see Fig. 2b). Similar gates can also be implemented in ion-trap architectures, where coupling to collective motional modes can be used to implement diagonal phase gates [78, 79].

However, Rydberg atom arrays offer additional control, which allows us to go beyond the alternating ansatz. Specifically, we consider simultaneously driving the qubit transi-

tion $\Omega_q(t)$ in addition to the usual Rydberg transition $\Omega_r(t)$. We find that this dual driving enables significantly faster realizations of U_S and U_P . After optimizing with GRAPE to identify approximate gates with ideal fidelities above 99.9%, we find total gate times below $\Omega T_{gate}/2\pi = 6.0$ up to cluster sizes $n=8$ and nearly constant scaling with n (see Fig. 2b). These gates generate complex interactions in a very hardware-efficient way, making them ideal for accelerating spin-Hamiltonian simulations. Finally, we develop optimized decompositions of target spin operations into two-qubit gates (see SM), and find they are still orders of magnitude more costly than the hardware-efficient implementation (Fig. 2c).

In Fig. 3, we illustrate the performance of this method on four representative examples that lie within the family of Hamiltonians (1). To quantify the performance of the simulation, we present heuristic estimates of the accessible coherent simulation time, measured in units of the target Hamiltonian’s local energy scale. We leverage access to multi-qubit spin operations of the form $e^{-i\sum_n \theta_n S^n}$, and estimate gate errors based on the physical evolution time necessary to realize the target operation. The step-size τ is chosen to maximize coherent simulation time, balancing simulation and gate errors (see Methods). In the representative examples, we find the combination of dynamical projection and optimized multi-qubit gates outperforms a similarly constructed implementation based on Trotterized interactions and two-qubit gate decomposition. Our approach significantly extends the available simulation time (up to $\sim$ two orders of magnitude), and enables much more efficient generation of complex spin Hamiltonians.

Spectral information from dynamics

Having described a toolbox for implementing time-evolution and state-preparation, we next present a general-purpose algorithm for calculating chemically relevant information. The approach, dubbed “many-body spectroscopy”, leverages dynamical snapshot measurements and classical co-processing [47] to compute a wide variety of spectral quantities including low-lying states and finite-temperature properties. The procedure, combining insights from statistical phase estimation [80–83] and shadow tomography [84], is noise-resilient and sample-efficient, making it especially promising for near-term experiments.

Specifically, the quantity we extract is an operator-resolved density of states

$$D^A(\omega) = \sum_n \langle n|A|n\rangle \delta(\omega - \epsilon_n), \quad (8)$$

where ϵ_n and $|n\rangle$ are the energies and eigenstates of the evolution Hamiltonian H , and A denotes an arbitrary operator. Spectral functions like equation (8) can be used to access detailed information about the properties of H : the location of peaks provides information about energies [80, 82, 83], and




FIG. 4. **Many-body spectroscopy of model Hamiltonians.** (a) Schematic quantum circuit diagram for the algorithm. The first step is to apply a state-preparation circuit S to prepare a reference state $|S\rangle = S|0\rangle$, followed by an ancilla-controlled perturbation R preparing a superposition of $|S\rangle$ and the probe state $|R\rangle = R|S\rangle$. This superposition is time-evolved by the target Hamiltonian, and then each qubit is projectively measured to produce a snapshot. By repeating the procedure for various evolution times t , different perturbations R , and potentially different measurement bases, this setup provides access to detailed information about the spectrum of H . (b) Consider a spin Hamiltonian H_2 with two anti-ferromagnetically coupled spin-3/2's. We simulate 20,000 snapshot measurements and classical post processing to calculate the density-of-states $D^I(\omega)$ (black line), and total-spin resolved versions $D^{P_s}(\omega)$ (colored lines). Vertical dashed lines correspond to exact energies and colored regions represent 95% confidence intervals. Peaks are broadened due to finite (coherent) simulation time $JT_{\text{sim}} = 0.26$, which sets the spectral resolution. Hardware-efficient simulation schemes, which extend the simulation time (e.g. Fig. 3), are favorable because they improve spectral resolution. Many-body spectroscopy further improves the effective spectral resolution, by leveraging multiple observables to distinguish overlapping peaks. Here, we see that spin-resolution significantly sparsifies the signal, enabling accurate peak detection and energy estimation, while the bare spectrum $D^I(\omega)$ is too broad to resolve all states. (c) The magnetic susceptibility χ , can also be computed from snapshot measurements using S^z -resolved density-of-states (here $JT_{\text{sim}} = 1.04$). For these calculations, it is important to prevent exponential amplification of shot-noise. We therefore use a simple empirical truncation procedure which introduces a small amount of bias (see Methods), but enables rapid convergence with number of snapshots to the ideal value (black line).

properties of eigenstates can be computed by choosing A appropriately [81]. For example, we can compute total-spin of an eigenstate using $A = (\sum_i \hat{S}_i)^2$, or the local spin polarization with $A = \hat{S}_i$.

The output of the quantum computation are projective measurements (snapshots), produced by the circuit depicted in Fig. 4a. First, we initialize a system of N qubits and apply a state-preparation procedure to prepare a reference state $|S\rangle = S|0^N\rangle$. Next, we use a single ancilla qubit to apply a controlled perturbation, preparing a superposition of $|S\rangle$ and a probe state $|R\rangle = R|S\rangle$, which is then evolved under H for time t . Finally each qubit is projectively measured, producing a sequence of $N + 1$ bits — a snapshot. By measuring the ancilla in the X or Y basis, this circuit effectively performs an interferometry experiment between the reference and probe states. The resulting snapshot measurements enable parallel estimation of two-time correlation functions of the form $C_{O,R}(t) = \langle S|e^{iHt}Oe^{-iHt}R|S\rangle$ for all 2^N operators O that are diagonal in the measurement basis.

To estimate the spectral function $D^A(\omega)$, we use a hybrid quantum-classical computation based on the following expression,

$$D^A(\omega) = \underbrace{\mathbb{E}_{R \sim \mathcal{R}}}_{\text{circuit average}} \underbrace{\int dt e^{i\omega t} \sum_s \langle S|R^\dagger A O_s e^{-iHt}|S\rangle}_{\text{classical processing}} \times \underbrace{\langle S|e^{iHt}O_s e^{-iHt}R|S\rangle}_{\text{quantum evolution (Fig. 4)}}. \quad (9)$$

For this to be formally equivalent to (8), the distribution over perturbations $\mathcal{R}$ and ensemble of observables $\{O_s\}$ must couple uniformly to all eigenstates, to ensure unbiased estimation (see Methods). For example, given a polarized reference state $|0\rangle^{\otimes N}$, X -basis measurements are sufficient.

To realize the distributional average and time-integral, the quantum circuit has to be executed (Fig. 4a) for randomly sampled perturbations R and evolution times t . In order to efficiently evaluate the classical part of (9), we require efficient classical representations of $|R\rangle$ and $e^{-iHt}|S\rangle$. A good choice is to select $|S\rangle$ to be a known eigenstate of H , such that the time-evolution is trivial, and preferentially sample $|R\rangle$ to maximize the overlap with relevant target states. In contrast, the quantum part of (9) includes time-evolution of $|R\rangle$, which has overlap with unknown eigenstates, and often includes large amounts of entanglement. Therefore, this is estimated from snapshot measurements produced by quantum simulation (see Methods).

As an example, consider two interacting spin-3/2 particles, described by $H_2 = JS_1 \cdot S_2$. We prepare a polarized reference eigenstate $|S\rangle = |0\rangle^{\otimes 6}$, and sample perturbations R from an

ensemble of random single-spin rotations. This corresponds to a trivial state-preparation circuit and a simple controlled-perturbation composed of two-qubit gates in Fig. 4. Next, we measure the system in the X -basis, which provides access to the full spectrum for this choice of $|S\rangle$ (Methods). Finally, during classical processing, the bare density of states is obtained by choosing $A = \mathbb{1}$, and evaluating (9). The result contains peaks at frequencies ω associated with eigenstates of H_2 (see Fig. 4b). Further, we can isolate individual contributions of total spin sectors by instead choosing $A = P_S$, the projectors onto $S = 0, 1, 2$, and 3. This not only allows us to identify the total spin of the eigenstates, but also increases the effective spectral resolution in the presence of noise as it sparsifies the signal. Finite-temperature response functions [85], like the z -component of the zero-field magnetic susceptibility $\chi(T) = \frac{1}{Z} \text{Tr}[\frac{1}{T}(S_z)^2 e^{-H/T}]$, can also be computed from the same dataset, by integrating the S_z -projected density-of-states $D^{S_z}(\omega)$ (see Methods). To illustrate this, the magnetic susceptibility is extracted from the same dataset and shown in Fig. 4c. The algorithm is especially promising for near-term devices, having favorable resource requirements quantified by the number of snapshots (sample complexity) and maximum evolution time (coherence) required for accurate spectral computation (see Methods for further discussion).

Application to transition metal clusters and magnetic solids

As an illustration of a relevant computation in chemical catalysis, we consider the $\text{Mn}_4\text{O}_5\text{Ca}$ core of the oxygen-evolving complex (OEC), a transition metal catalyst central to photosynthesis which is still not fully understood [86, 87]. Classical chemistry calculations have been used to fit model Heisenberg Hamiltonians, containing three spin-3/2 sites and one spin-2 site [27, 28, 88] (Fig. 5a). While this spin representation cannot directly capture chemical reactions, it can capture the ground and low-lying spin-states, i. e., the spin-ladder, which are important in catalysis because reaction pathways depend critically on the spin multiplicity [89]. We simulate our framework applied to the $\text{S}_2\text{H-1b}$ structural model from Ref. [28], by first computing the bare density-of-states $D^{\mathbb{1}}(A)$ (Fig. 5b). Then, we identify a low-lying cluster of eigenstates, and compute spin-projected densities $D^{P_s}(A)$ to resolve the spin ladder (Fig. 5c). The spin ladder can also be measured experimentally, providing a way to evaluate candidate models of reaction intermediates. To highlight this, we simulate the Heisenberg model for an alternate pathway ($\text{S}_2\text{H-2b}$) [28], and observe the modification reverses the ordering of the spin ladder, indicating $\text{S}_2\text{H-1b}$ is more consistent with measurements (see Fig. 5d).

The framework can also be applied to study low-energy properties of extended systems, including strongly correlated materials. We illustrate this on the ferromagnetic, square lattice Heisenberg model (Fig. 6). For such large systems, we envision utilizing an approximate ground-state preparation method for $|S\rangle$, so that low-energy properties can be accessed




FIG. 5. Application to the oxygen-evolving complex (OEC). Our programmable quantum simulation framework can be used to compute detailed model spin Hamiltonian properties. (a) Here, we illustrate the procedure on the OEC, an organometallic catalyst with strong spin correlations. In particular, we simulate model spin Hamiltonians for two structures $\text{S}_2\text{H-1b}$ and $\text{S}_2\text{H-2b}$, which have three spin-3/2 and one spin-2 Mn (purple) active sites (reproduced from Ref. [28] with permission from the Royal Society of Chemistry). Model Heisenberg coefficients for both hypothetical structures have been computed from broken-symmetry DFT [28]. (b) A density of states $D^{\mathbb{1}}(\omega)$ calculation is simulated for the $\text{S}_2\text{H-1b}$ model spin Hamiltonian. Here, we use a polarized reference state $|S\rangle = |0\rangle^{\otimes 13}$ a probe states $|R\rangle$ generated by random single-site rotations, and a evolution times t . We select 50,000 circuits with independently chosen $|R\rangle, t$ pairs, and draw 10 snapshots from each circuit. (c) Focusing on the lowest-lying states, we see three distinct peaks in $D^{\mathbb{1}}(\omega)$. However, by evaluating spin-resolved quantities $D^{P_s}(\omega)$ on the same set of measurements, we identify three additional peaks, whose energies and total-spin match exact diagonalization results (vertical dotted lines). (d) This information is known as the spin-ladder, and can be computed using many-body spectroscopy for both the 1b and 2b states. Importantly, the spin-ladder can also be measured experimentally, and therefore can be used to help determine which structure appears in nature. In this example, experimental measurements indicate a spin-5/2 ground state and spin-7/2 first excited state. However, the ordering of low-energy states is flipped in the 2b configuration, indicating the $\text{S}_2\text{H-1b}$ hypothesis is more likely [28]. We note that quantities beyond total-spin can also be readily evaluated in low-lying eigenstates by inserting different operators A (see Methods).

in a noise-resilient manner via local controlled-perturbations R . Then, local Green's functions — two-point operators at different positions and times — can be measured to access properties of low-lying quasi-particle excitations, such as the dispersion relation of single-particle excitations (see Methods, 2D Heisenberg, for details).




FIG. 6. **Application to 2D magnetic materials.** (a) Properties of extended materials can also be investigated using our quantum simulation framework. As an example, we study the square-lattice ferromagnetic ($J > 0$) Heisenberg model $H_{2D} = -J \sum_{\langle ij \rangle} \hat{\mathbf{s}}_i \cdot \hat{\mathbf{s}}_j$. By preparing the polarized ground state $|S\rangle = |0\rangle^{\otimes N}$, applying a single-site perturbation $R = X_0$ on the central site, and measuring the system in the X -basis after time evolution, we can estimate the single-particle Green's function $G(\mathbf{r}, t) = \langle S | X_{\mathbf{r}}(t) X_0 | S \rangle$. Therefore we select $O = X_{\mathbf{r}}$ for various positions $\mathbf{r}$ as the observables in (9), all of which are diagonal in the measurement basis. We visualize the real-part of $G(\mathbf{r}, t)$, where the plotted intensity and color denotes the magnitude and sign at $Jt = 0, 0.5$, and 1.0 . (b) The structure of excited states is extracted by classical post-processing of these measurements. Even though the spectrum is continuous, additional structure can be identified by computing the momentum-resolved density of states $D^{P_{\mathbf{k}}}(\omega)$, where $P_{\mathbf{k}}$ is a projector onto plane-wave states (see Methods). Restricting to $k_y = 0$ and evolving to maximum time $JT_{\max} = 8.0$, we see $D^{P_{\mathbf{k}}}(\omega)$ forms a band-like structure, from which a peak ω can be estimated for each $\mathbf{k}$ (black line). (c) This peak extraction allows us to directly estimate the single-particle dispersion $\omega(\mathbf{k})$ across the 2D Brillouin zone.

Outlook

These considerations indicate that reconfigurable quantum processors enable a powerful, hardware-efficient framework for quantum simulation of problems from chemistry and materials science, illustrating potential directions for the search for useful quantum advantage. Specifically, in addition to the OEC, other organometallic catalysts could be studied with this approach, including iron-sulfur clusters [23, 90], for which bi-quadratic terms appear in the model Hamiltonian to capture higher-order perturbative charge fluctuation effects [91]. Another promising direction involves quantum simulation of low-energy properties of 2D and 3D frustrated spin systems, including model Hamiltonians for Kitaev materials [92–94] and molecular magnets [95, 96]. The ability to realize non-local interactions further opens the door to simulation of spin-Hamiltonians defined on non-Euclidean interaction geome-

tries [97, 98].

The efficiency of the Hamiltonian engineering approach originates from co-designing Floquet engineering and hardware-specific multi-qubit gates. Extending this approach to larger classes of strongly-correlated model Hamiltonians is an outstanding and exciting frontier. In particular, it would be especially interesting to further develop the toolbox to incorporate charge transport and electron-phonon interactions. This could enable simulation of more complex model Hamiltonians, such as the t-J, Hubbard, and Hubbard-Holstein models, and expand the class of accessible chemistry problems [9, 99]. Application to other settings, including lattice gauge theories [100, 101] and quantum optimization problems [102] are also of interest. Incorporation of error mitigation and correction into the Hamiltonian simulation should be considered; specifically, the present method can potentially be generalized to control logical, encoded degrees of freedom in a hardware-efficient way [103].

Finally, characterization and development of the model Hamiltonian approach itself is an interesting and challenging problem. Key challenges include development of efficient schemes to compute parameters for higher-order interactions [91], estimation of corrections arising from coupling to states outside the model space [104], and validation of the model Hamiltonian approximation [19]. Feedback between the classical and quantum parts of the computation is an important part of these developments [25, 105, 106]. For these reasons, the large-scale simulation of model Hamiltonians on quantum processors will be invaluable for testing approximations by comparing simulation outputs with experimental measurements. Hence, the approach proposed in this work facilitates exciting directions in computational chemistry and quantum simulation, aiming towards constructing a novel *ab initio* simulation pipeline that utilizes hybrid quantum-classical resources.

-
- [1] A. Szabo and N. S. Ostlund, *Modern quantum chemistry: introduction to advanced electronic structure theory* (Courier Corporation, 2012).
 - [2] R. J. Bartlett and M. Musiał, Coupled-cluster theory in quantum chemistry, *Rev. Mod. Phys.* **79**, 291 (2007).
 - [3] N. Mardirossian and M. Head-Gordon, Thirty years of density functional theory in computational chemistry: an overview and extensive assessment of 200 density functionals, *Mol. Phys.* **115**, 2315 (2017).
 - [4] K. Burke, Perspective on density functional theory, *J. Chem. Phys.* **136** (2012).
 - [5] Y. Aoto, A. P. de Lima Batista, A. Kohn, and A. G. de Oliveira-Filho, How to arrive at accurate benchmark values for transition metal compounds: Computation or experiment?, *J. Chem. Theory Comput.* **13**, 5291 (2017).
 - [6] D. Hait, N. M. Tubman, D. S. Levine, K. B. Whaley, and M. Head-Gordon, What levels of coupled cluster theory are appropriate for transition metal systems? a study using near-exact quantum chemical values for 3d transition metal binary

- compounds, *J. Chem. Theory Comput.* **15**, 5370 (2019).
- [7] C. Ahn, A. Cavalleri, A. Georges, S. Ismail-Beigi, A. J. Millis, and J.-M. Triscone, Designing and controlling the properties of transition metal oxide quantum materials, *Nat. Mater.* **20**, 1462 (2021).
 - [8] Y. Alexeev, D. Bacon, K. R. Brown, R. Calderbank, L. D. Carr, F. T. Chong, B. DeMarco, D. Englund, E. Farhi, B. Fefferman, A. V. Gorshkov, A. Houck, J. Kim, S. Kimmel, M. Lange, S. Lloyd, M. D. Lukin, D. Maslov, P. Maunz, C. Monroe, J. Preskill, M. Roetteler, M. J. Savage, and J. Thompson, Quantum Computer Systems for Scientific Discovery, *PRX Quantum* **2**, 017001 (2021).
 - [9] B. Bauer, S. Bravyi, M. Motta, and G. K.-L. Chan, Quantum algorithms for quantum chemistry and quantum materials science, *Chem. Rev.* **120**, 12685 (2020).
 - [10] S. McArdle, S. Endo, A. Aspuru-Guzik, S. C. Benjamin, and X. Yuan, Quantum computational chemistry, *Rev. Mod. Phys.* **92**, 015003 (2020).
 - [11] L. Lin, Lecture notes on quantum algorithms for scientific computation, [arXiv:2201.08309 \[quant-ph\]](https://arxiv.org/abs/2201.08309) (2022).
 - [12] M. E. Beverland, P. Murali, M. Troyer, K. M. Svore, T. Hoefler, V. Kliuchnikov, G. H. Low, M. Soeken, A. Sundaram, and A. Vashchillo, Assessing requirements to scale to practical quantum advantage, [arXiv:2211.07629 \[quant-ph\]](https://arxiv.org/abs/2211.07629) (2022).
 - [13] A. Auerbach, *Interacting electrons and quantum magnetism* (Springer Science & Business Media, 1998).
 - [14] N. P. Bauman, G. H. Low, and K. Kowalski, Quantum simulations of excited states with active-space downfolded Hamiltonians, *J. Chem. Phys.* **151**, 234114 (2019).
 - [15] H. Bolvin, From ab Initio Calculations to Model Hamiltonians: The Effective Hamiltonian Technique as an Efficient Tool to Describe Mixed-Valence Molecules, *J. Phys. Chem. A* **107**, 5071 (2003).
 - [16] N. J. Mayhall and M. Head-Gordon, Computational quantum chemistry for single Heisenberg spin couplings made simple: Just one spin flip required, *J. Chem. Phys.* **141**, 134111 (2014).
 - [17] N. J. Mayhall and M. Head-Gordon, Computational Quantum Chemistry for Multiple-Site Heisenberg Spin Couplings Made Simple: Still Only One Spin-Flip Required, *J. Phys. Chem. Lett.* **6**, 1982 (2015).
 - [18] P. Pokhilko and A. I. Krylov, Effective Hamiltonians derived from equation-of-motion coupled-cluster wave functions: Theory and application to the Hubbard and Heisenberg Hamiltonians, *J. Chem. Phys.* **152**, 094108 (2020).
 - [19] S. Kotaru, S. Kähler, M. Alessio, and A. I. Krylov, Magnetic exchange interactions in binuclear and tetranuclear iron (iii) complexes described by spin-flip dft and heisenberg effective hamiltonians, *J. Comput. Chem.* **44**, 367 (2023).
 - [20] D.-T. Chen, P. Helms, A. R. Hale, M. Lee, C. Li, J. Gray, G. Christou, V. S. Zapf, G. K.-L. Chan, and H.-P. Cheng, Using Hyperoptimized Tensor Networks and First-Principles Electronic Structure to Simulate the Experimental Properties of the Giant {Mn84} Torus, *J. Phys. Chem. Lett.* **13**, 2365 (2022).
 - [21] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets, *Nature* **549**, 242 (2017).
 - [22] A. Chiesa, F. Tacchino, M. Grossi, P. Santini, I. Tavernelli, D. Gerace, and S. Carretta, Quantum hardware simulating four-dimensional inelastic neutron scattering, *Nat. Phys.* **15**, 455 (2019).
 - [23] R. N. Tazhigulov, S.-N. Sun, R. Haghshenas, H. Zhai, A. T. Tan, N. C. Rubin, R. Babbush, A. J. Minnich, and G. K.-L. Chan, Simulating Models of Challenging Correlated Molecules and Materials on the Sycamore Quantum Processor, *PRX Quantum* **3**, 040318 (2022).
 - [24] D. Wecker, M. B. Hastings, N. Wiebe, B. K. Clark, C. Nayak, and M. Troyer, Solving strongly correlated electron models on a quantum computer, *Physical Review A* **92**, 062318 (2015).
 - [25] B. Bauer, D. Wecker, A. J. Millis, M. B. Hastings, and M. Troyer, Hybrid Quantum-Classical Approach to Correlated Materials, *Physical Review X* **6**, 031045 (2016).
 - [26] H. Ma, M. Govoni, and G. Galli, Quantum simulations of materials on near-term quantum computers, *npj Computational Materials* **6**, 85 (2020).
 - [27] V. Krewald, F. Neese, and D. A. Pantazis, On the Magnetic and Spectroscopic Properties of High-Valent Mn₃CaO₄ Cubanes as Structural Units of Natural and Artificial Water-Oxidizing Catalysts, *J. Am. Chem. Soc.* **135**, 5726 (2013).
 - [28] V. Krewald, M. Retegan, N. Cox, J. Messinger, W. Lubitz, S. DeBeer, F. Neese, and D. A. Pantazis, Metal oxidation states in biological water splitting, *Chem. Sci.* **6**, 1676 (2015).
 - [29] J. P. Malrieu, R. Caballol, C. J. Calzado, C. de Graaf, and N. Guihéry, Magnetic Interactions in Molecules and Highly Correlated Materials: Physical Content, Analytical Derivation, and Rigorous Extraction of Magnetic Hamiltonians, *Chem. Rev.* **114**, 429 (2014).
 - [30] J. Shee, M. Loipersberger, D. Hait, J. Lee, and M. Head-Gordon, Revealing the nature of electron correlation in transition metal complexes with symmetry breaking and chemical intuition, *J. Chem. Phys.* **154**, 194109 (2021).
 - [31] D. A. Pantazis, M. Orto, T. Petrenko, S. Zein, W. Lubitz, J. Messinger, and F. Neese, Structure of the oxygen-evolving complex of photosystem II: information on the S2 state through quantum chemical calculation of its magnetic properties, *Phys. Chem. Chem. Phys.* **11**, 6788 (2009).
 - [32] C. J. Calzado, J. M. Clemente-Juan, E. Coronado, A. Gaitarino, and N. Suaud, Role of the Electron Transfer and Magnetic Exchange Interactions in the Magnetic Properties of Mixed-Valence Polyoxovanadate Complexes, *Inorg. Chem.* **47**, 5889 (2008).
 - [33] X. Li, H. Yu, F. Lou, J. Feng, M.-H. Whangbo, and H. Xiang, Spin Hamiltonians in Magnets: Theories and Computations, *Molecules* **26**, 803 (2021).
 - [34] P. Pokhilko, D. S. Bezrukov, and A. I. Krylov, Is Solid Copper Oxalate a Spin Chain or a Mixture of Entangled Spin Pairs?, *J. Phys. Chem. C* **125**, 7502 (2021).
 - [35] C. Monroe, W. Campbell, L.-M. Duan, Z.-X. Gong, A. Gorshkov, P. Hess, R. Islam, K. Kim, N. Linke, G. Pagano, P. Richerme, C. Senko, and N. Yao, Programmable quantum simulations of spin systems with trapped ions, *Rev. Mod. Phys.* **93**, 025001 (2021).
 - [36] A. J. Daley, I. Bloch, C. Kokail, S. Flannigan, N. Pearson, M. Troyer, and P. Zoller, Practical quantum advantage in quantum simulation, *Nature* **607**, 667 (2022).
 - [37] H. Bernien, S. Schwartz, A. Keesling, H. Levine, A. Omran, H. Pichler, S. Choi, A. S. Zibrov, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin, Probing many-body dynamics on a 51-atom quantum simulator, *Nature* **551**, 579 (2017).
 - [38] A. Keesling, A. Omran, H. Levine, H. Bernien, H. Pichler, S. Choi, R. Samajdar, S. Schwartz, P. Silvi, S. Sachdev, P. Zoller, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin, Quantum Kibble-Zurek mechanism and critical dynamics on a programmable Rydberg simulator, *Nature* **568**, 207 (2019).
 - [39] 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 (2021).
- [40] H. Labuhn, D. Barredo, S. Ravets, S. de Léséleuc, T. Macrì, T. Lahaye, and A. Browaeys, Tunable two-dimensional arrays of single Rydberg atoms for realizing quantum Ising models, *Nature* **534**, 667 (2016).
- [41] S. Ebadi, T. T. Wang, H. Levine, A. Keesling, G. Semeghini, A. Omran, D. Bluvstein, R. Samajdar, H. Pichler, W. W. Ho, S. Choi, S. Sachdev, M. Greiner, V. Vuletić, and M. D. Lukin, Quantum phases of matter on a 256-atom programmable quantum simulator, *Nature* **595**, 227 (2021).
- [42] P. Scholl, M. Schuler, H. J. Williams, A. A. Eberharter, D. Barredo, K.-N. Schymik, V. Lienhard, L.-P. Henry, T. C. Lang, T. Lahaye, A. M. Läuchli, and A. Browaeys, Quantum simulation of 2D antiferromagnets with hundreds of Rydberg atoms, *Nature* **595**, 233 (2021).
- [43] C. Chen, G. Bornet, M. Bintz, G. Emperauger, L. Leclerc, V. S. Liu, P. Scholl, D. Barredo, J. Hauschild, S. Chatterjee, M. Schuler, A. M. Läuchli, M. P. Zaletel, T. Lahaye, N. Y. Yao, and A. Browaeys, Continuous symmetry breaking in a two-dimensional Rydberg array, *Nature* **616**, 691 (2023).
- [44] G. Semeghini, H. Levine, A. Keesling, S. Ebadi, T. T. Wang, D. Bluvstein, R. Verresen, H. Pichler, M. Kalinowski, R. Samajdar, A. Omran, S. Sachdev, A. Vishwanath, M. Greiner, V. Vuletić, and M. D. Lukin, Probing topological spin liquids on a programmable quantum simulator, *Science* **374**, 1242 (2021).
- [45] K. J. Satzinger, Y.-J. Liu, A. Smith, C. Knapp, M. Newman, C. Jones, Z. Chen, C. Quintana, X. Mi, A. Dunsworth, C. Gidney, I. Aleiner, F. Arute, K. Arya, J. Atalaya, R. Babush, J. C. Bardin, R. Barends, J. Basso, A. Bengtsson, A. Bilmes, M. Broughton, B. B. Buckley, D. A. Buell, B. Burkett, N. Bushnell, B. Chiaro, R. Collins, W. Courtney, S. Demura, A. R. Derk, D. Eppens, C. Erickson, L. Faoro, E. Farhi, A. G. Fowler, B. Foxen, M. Giustina, A. Greene, J. A. Gross, M. P. Harrigan, S. D. Harrington, J. Hilton, S. Hong, T. Huang, W. J. Huggins, L. B. Ioffe, S. V. Isakov, E. Jeffrey, Z. Jiang, D. Kafri, K. Kechedzhi, T. Khattar, S. Kim, P. V. Klimov, A. N. Korotkov, F. Kostritsa, D. Landhuis, P. Laptev, A. Locharla, E. Lucero, O. Martin, J. R. McClean, M. McEwen, K. C. Miao, M. Mohseni, S. Montazeri, W. Mruczkiewicz, J. Mutus, O. Naaman, M. Neeley, C. Neill, M. Y. Niu, T. E. O'Brien, A. Opremcak, B. Pató, A. Petukhov, N. C. Rubin, D. Sank, V. Shvarts, D. Strain, M. Szalay, B. Villalonga, T. C. White, Z. Yao, P. Yeh, J. Yoo, A. Zalcman, H. Neven, S. Boixo, A. Megrant, Y. Chen, J. Kelly, V. Smelyanskiy, A. Kitaev, M. Knap, F. Pollmann, and P. Roushan, Realizing topologically ordered states on a quantum processor, *Science* **374**, 1237 (2021).
- [46] Y. Kim, A. Eddins, S. Anand, K. X. Wei, E. Van Den Berg, S. Rosenblatt, H. Nayfeh, Y. Wu, M. Zaletel, K. Temme, and A. Kandala, Evidence for the utility of quantum computing before fault tolerance, *Nature* **618**, 500 (2023).
- [47] H. H. S. Chan, R. Meister, M. L. Goh, and B. Koczor, Algorithmic shadow spectroscopy, *arXiv:2212.11036 [quant-ph]* (2023).
- [48] D. Barredo, S. De Léséleuc, V. Lienhard, T. Lahaye, and A. Browaeys, An atom-by-atom assembler of defect-free arbitrary two-dimensional atomic arrays, *Science* **354**, 1021 (2016).
- [49] A. Cooper, J. P. Covey, I. S. Madjarov, S. G. Porsev, M. S. Safronova, and M. Endres, Alkaline-Earth Atoms in Optical Tweezers, *Phys. Rev. X* **8**, 041055 (2018).
- [50] S. Ma, A. P. Burgers, G. Liu, J. Wilson, B. Zhang, and J. D. Thompson, Universal Gate Operations on Nuclear Spin Qubits in an Optical Tweezer Array of ^{171}Yb Atoms, *Phys. Rev. X* **12**, 021028 (2022).
- [51] K. Singh, S. Anand, A. Pocklington, J. T. Kemp, and H. Bernien, Dual-Element, Two-Dimensional Atom Array with Continuous-Mode Operation, *Phys. Rev. X* **12**, 011040 (2022).
- [52] A. Jenkins, J. W. Lis, A. Senoo, W. F. McGrew, and A. M. Kaufman, Ytterbium Nuclear-Spin Qubits in an Optical Tweezer Array, *Physical Review X* **12**, 021027 (2022).
- [53] O. Katz, L. Feng, A. Risinger, C. Monroe, and M. Cetina, Demonstration of three- and four-body interactions between trapped-ion spins, *Nat. Phys.* **19**, 1452 (2023).
- [54] H. Häffner, C. F. Roos, and R. Blatt, Quantum computing with trapped ions, *Phys. Rep.* **469**, 155 (2008).
- [55] S. A. Moses, C. H. Baldwin, M. S. Allman, R. Ancona, L. Ascarunz, C. Barnes, J. Bartolotta, B. Bjork, P. Blanchard, M. Bohn, J. G. Bohnet, N. C. Brown, N. Q. Burdick, W. C. Burton, S. L. Campbell, J. P. C. I. au2, C. Carron, J. Chambers, J. W. Chan, Y. H. Chen, A. Chernoguzov, E. Chertkov, J. Colina, J. P. Curtis, R. Daniel, M. DeCross, D. Deen, C. Delaney, J. M. Dreiling, C. T. Ertsgaard, J. Esposito, B. Estey, M. Fabrikant, C. Figgatt, C. Foltz, M. Foss-Feig, D. Francois, J. P. Gaebler, T. M. Gatterman, C. N. Gilbreth, J. Giles, E. Glynn, A. Hall, A. M. Hankin, A. Hansen, D. Hayes, B. Higashi, I. M. Hoffman, B. Horning, J. J. Hout, R. Jacobs, J. Johansen, L. Jones, J. Karcz, T. Klein, P. Lauria, P. Lee, D. Liefer, C. Lytle, S. T. Lu, D. Lucchetti, A. Malm, M. Matheny, B. Mathewson, K. Mayer, D. B. Miller, M. Mills, B. Neyenhuis, L. Nugent, S. Olson, J. Parks, G. N. Price, Z. Price, M. Pugh, A. Ransford, A. P. Reed, C. Roman, M. Rowe, C. Ryan-Anderson, S. Sanders, J. Sedlacek, P. Shevchuk, P. Siegfried, T. Skripka, B. Spaun, R. T. Sprengle, R. P. Stutz, M. Swallows, R. I. Tobey, A. Tran, T. Tran, E. Vogt, C. Volin, J. Walker, A. M. Zolot, and J. M. Pino, A race track trapped-ion quantum processor, *arXiv:2305.03828 [quant-ph]* (2023).
- [56] D. Bluvstein, H. Levine, G. Semeghini, T. T. Wang, S. Ebadi, M. Kalinowski, A. Keesling, N. Maskara, H. Pichler, M. Greiner, V. Vuletić, and M. D. Lukin, A quantum processor based on coherent transport of entangled atom arrays, *Nature* **604**, 451 (2022).
- [57] S. J. Evered, D. Bluvstein, M. Kalinowski, S. Ebadi, T. Manovitz, H. Zhou, S. H. Li, A. A. Geim, T. T. Wang, N. Maskara, H. Levine, G. Semeghini, M. Greiner, V. Vuletić, and M. D. Lukin, High-fidelity parallel entangling gates on a neutral-atom quantum computer, *Nature* **622**, 268 (2023).
- [58] A. Kruckenhauser, R. v. Bijnen, T. V. Zache, M. D. Liberto, and P. Zoller, High-dimensional $\text{SO}(4)$ -symmetric Rydberg manifolds for quantum simulation, *Quantum Sci. Technol.* **8**, 015020 (2022).
- [59] Y. Chi, J. Huang, Z. Zhang, J. Mao, Z. Zhou, X. Chen, C. Zhai, J. Bao, T. Dai, H. Yuan, M. Zhang, D. Dai, B. Tang, Y. Yang, Z. Li, Y. Ding, L. K. Oxenløwe, M. G. Thompson, J. L. O'Brien, Y. Li, Q. Gong, and J. Wang, A programmable qudit-based quantum processor, *Nat. Commun.* **13**, 1166 (2022).
- [60] P. Hrmo, B. Wilhelm, L. Gerster, M. W. van Mourik, M. Huber, R. Blatt, P. Schindler, T. Monz, and M. Ringbauer, Native qudit entanglement in a trapped ion quantum processor, *Nat. Commun.* **14**, 2242 (2023).
- [61] D. González-Cuadra, T. V. Zache, J. Carrasco, B. Kraus, and P. Zoller, Hardware Efficient Quantum Simulation of Non-Abelian Gauge Theories with Qudits on Rydberg Platforms, *Phys. Rev. Lett.* **129**, 160501 (2022).
- [62] D. A. Abanin, W. De Roeck, W. W. Ho, and F. Huveneers, Ef-

- fective Hamiltonians, prethermalization, and slow energy absorption in periodically driven many-body systems, *Phys. Rev. B* **95**, 014112 (2017).
- [63] J. Choi, H. Zhou, H. S. Knowles, R. Landig, S. Choi, and M. D. Lukin, Robust Dynamic Hamiltonian Engineering of Many-Body Spin Systems, *Phys. Rev. X* **10**, 031002 (2020).
- [64] D. V. Else, B. Bauer, and C. Nayak, Prethermal Phases of Matter Protected by Time-Translation Symmetry, *Phys. Rev. X* **7**, 011026 (2017).
- [65] J. Carrasco, A. Elben, C. Kokail, B. Kraus, and P. Zoller, Theoretical and Experimental Perspectives of Quantum Verification, *PRX Quantum* **2**, 010102 (2021).
- [66] M. Benedetti, M. Fiorentini, and M. Lubasch, Hardware-efficient variational quantum algorithms for time evolution, *Physical Review Research* **3**, 033083 (2021), arXiv:2009.12361 [quant-ph].
- [67] P. Scholl, A. L. Shaw, R. B.-S. Tsai, R. Finkelstein, J. Choi, and M. Endres, Erasure conversion in a high-fidelity Rydberg quantum simulator, *Nature* **622**, 273 (2023).
- [68] S. Ma, G. Liu, P. Peng, B. Zhang, S. Jandura, J. Claes, A. P. Burgers, G. Pupillo, S. Puri, and J. D. Thompson, High-fidelity gates and mid-circuit erasure conversion in an atomic qubit, *Nature* **622**, 279 (2023).
- [69] J. W. Lis, A. Senoo, W. F. McGrew, F. Rönchen, A. Jenkins, and A. M. Kaufman, Mid-circuit operations using the omg-architecture in neutral atom arrays (2023), arXiv:2305.19266 [cond-mat, physics:physics, physics:quant-ph].
- [70] C. Sheng, X. He, P. Xu, R. Guo, K. Wang, Z. Xiong, M. Liu, J. Wang, and M. Zhan, High-fidelity single-qubit gates on neutral atoms in a two-dimensional magic-intensity optical dipole trap array, *Phys. Rev. Lett.* **121**, 240501 (2018).
- [71] H. Levine, D. Bluvstein, A. Keesling, T. T. Wang, S. Ebadi, G. Semeghini, A. Omran, M. Greiner, V. Vuletić, and M. D. Lukin, Dispersive optical systems for scalable raman driving of hyperfine qubits, *Phys. Rev. A* **105**, 032618 (2022).
- [72] 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 (2000).
- [73] H. Levine, A. Keesling, G. Semeghini, A. Omran, T. T. Wang, S. Ebadi, H. Bernien, M. Greiner, V. Vuletić, H. Pichler, and M. D. Lukin, Parallel implementation of high-fidelity multi-qubit gates with neutral atoms, *Phys. Rev. Lett.* **123**, 170503 (2019).
- [74] E. Urban, T. A. Johnson, T. Henage, L. Isenhower, D. D. Yavuz, T. G. Walker, and M. Saffman, Observation of Rydberg blockade between two atoms, *Nat. Phys.* **5**, 110 (2009).
- [75] S. Jandura and G. Pupillo, Time-Optimal Two- and Three-Qubit Gates for Rydberg Atoms, *Quantum* **6**, 712 (2022).
- [76] A. Pagano, S. Weber, D. Jaschke, T. Pfau, F. Meinert, S. Montangero, and H. P. Büchler, Error budgeting for a controlled-phase gate with strontium-88 rydberg atoms, *Phys. Rev. Res.* **4**, 033019 (2022).
- [77] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, and S. J. Glaser, Optimal control of coupled spin dynamics: design of NMR pulse sequences by gradient ascent algorithms, *J. Magn. Reson.* **172**, 296 (2005).
- [78] A. Sørensen and K. Mølmer, Quantum computation with ions in thermal motion, *Phys. Rev. Lett.* **82**, 1971 (1999).
- [79] O. Katz, M. Cetina, and C. Monroe, N-body interactions between trapped ion qubits via spin-dependent squeezing, *Phys. Rev. Lett.* **129**, 063603 (2022).
- [80] L. Lin and Y. Tong, Heisenberg-limited ground state energy estimation for early fault-tolerant quantum computers, *PRX Quantum* **3**, 010318 (2022), arXiv:2102.11340 [quant-ph].
- [81] S. Lu, M. C. Bañuls, and J. I. Cirac, Algorithms for quantum simulation at finite energies, *PRX Quantum* **2**, 020321 (2021).
- [82] R. D. Somma, Quantum eigenvalue estimation via time series analysis, *New J. Phys.* **21**, 123025 (2019).
- [83] T. E. O’Brien, B. Tarasinski, and B. M. Terhal, Quantum phase estimation of multiple eigenvalues for small-scale (noisy) experiments, *New J. Phys.* **21**, 023022 (2019).
- [84] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, *Nat. Phys.* **16**, 1050 (2020).
- [85] M. Alessio and A. I. Krylov, Equation-of-Motion Coupled-Cluster Protocol for Calculating Magnetic Properties: Theory and Applications to Single-Molecule Magnets, *J. Chem. Theory Comput.* **17**, 4225 (2021).
- [86] M. Askerka, G. W. Brudvig, and V. S. Batista, The o₂-evolving complex of photosystem ii: Recent insights from quantum mechanics/molecular mechanics (qm/mm), extended x-ray absorption fine structure (exafs), and femtosecond x-ray crystallography data, *Acc. Chem. Res.* **50**, 41 (2017).
- [87] S. Paul, F. Neese, and D. A. Pantazis, Structural models of the biological oxygen-evolving complex: achievements, insights, and challenges for biomimicry, *Green Chem.* **19**, 2309 (2017).
- [88] V. Krewald, M. Retegan, F. Neese, W. Lubitz, D. A. Pantazis, and N. Cox, Spin State as a Marker for the Structural Evolution of Nature’s Water-Splitting Catalyst, *Inorg. Chem.* **55**, 488 (2016).
- [89] S. Shaik, H. Hirao, and D. Kumar, Reactivity of high-valent iron-oxo species in enzymes and synthetic reagents: a tale of many states, *Acc. Chem. Res.* **40**, 532 (2007).
- [90] Z. Li, S. Guo, Q. Sun, and G. K.-L. Chan, Electronic landscape of the P-cluster of nitrogenase as revealed through many-electron quantum wavefunction simulations, *Nat. Chem.* **11**, 1026 (2019).
- [91] S. Sharma, K. Sivalingam, F. Neese, and G. K.-L. Chan, Low-energy spectrum of iron-sulfur clusters directly from many-particle quantum mechanics, *Nat. Chem.* **6**, 927 (2014).
- [92] H.-K. Jin, W. M. H. Natori, F. Pollmann, and J. Knolle, Unveiling the S=3/2 Kitaev honeycomb spin liquids, *Nat. Commun.* **13**, 3813 (2022).
- [93] C. Xu, J. Feng, M. Kawamura, Y. Yamaji, Y. Nahas, S. Prokhorenko, Y. Qi, H. Xiang, and L. Bellaiche, Possible Kitaev Quantum Spin Liquid State in 2D Materials with S = 3 / 2, *Phys. Rev. Lett.* **124**, 087205 (2020).
- [94] H. Takagi, T. Takayama, G. Jackeli, G. Khaliullin, and S. E. Nagler, Concept and realization of Kitaev quantum spin liquids, *Nat. Rev. Phys.* **1**, 264 (2019).
- [95] L. Bogani and W. Wernsdorfer, Molecular spintronics using single-molecule magnets, *Nature Materials* **7**, 179 (2008).
- [96] D. N. Woodruff, R. E. P. Winpenny, and R. A. Lay, Lanthanide Single-Molecule Magnets, *Chem. Rev.* **113**, 5110 (2013).
- [97] A. Anshu, N. P. Breuckmann, and C. Nirkhe, NLTS Hamiltonians from good quantum codes, in *Proceedings of the 55th Annual ACM Symposium on Theory of Computing* (2023) pp. 1090–1096.
- [98] T. Rakovszky and V. Khemani, The physics of (good) ldpc codes i. gauging and dualities, arXiv:2310.16032 [quant-ph] (2023).
- [99] M. Motta and J. E. Rice, Emerging quantum computing algorithms for quantum chemistry, *Wiley Interdiscip. Rev. Comput. Mol. Sci.* **12**, e1580 (2022).
- [100] A. Celi, B. Vermersch, O. Viyuela, H. Pichler, M. D. Lukin, and P. Zoller, Emerging Two-Dimensional Gauge Theories in Rydberg Configurable Arrays, *Phys. Rev. X* **10**, 021057 (2020).

- [101] T. V. Zache, D. González-Cuadra, and P. Zoller, Fermion-qudit quantum processors for simulating lattice gauge theories with matter, [Quantum](#) **7**, 1140 (2023).
- [102] S. Ebadi, A. Keesling, M. Cain, T. T. Wang, H. Levine, D. Bluvstein, G. Semeghini, A. Omran, J.-G. Liu, R. Samajdar, X.-Z. Luo, B. Nash, X. Gao, B. Barak, E. Farhi, S. Sachdev, N. Gemelke, L. Zhou, S. Choi, H. Pichler, S.-T. Wang, M. Greiner, V. Vuletić, and M. D. Lukin, Quantum optimization of maximum independent set using Rydberg atom arrays, [Science](#) **376**, 1209 (2022).
- [103] D. Bluvstein, S. J. Evered, A. A. Geim, S. H. Li, H. Zhou, T. Manovitz, S. Ebadi, M. Cain, M. Kalinowski, D. Hangleiter, J. P. Bonilla Ataides, N. Maskara, I. Cong, X. Gao, P. S. Rodriguez, T. Karolyshyn, G. Semeghini, M. J. Gullans, M. Greiner, V. Vuletic, and M. D. Lukin, Logical quantum processor based on reconfigurable atom arrays, Submitted to Nature.
- [104] J. Günther, A. Baiardi, M. Reiher, and M. Christandl, More quantum chemistry with fewer qubits, [arXiv:2308.16873 \[quant-ph\]](#) (2023).
- [105] C. Vorwerk, N. Sheng, M. Govoni, B. Huang, and G. Galli, Quantum embedding theories to simulate condensed systems on quantum computers, [Nature Computational Science](#) **2**, 424 (2022).
- [106] W. J. Huggins, B. A. O’Gorman, N. C. Rubin, D. R. Reichman, R. Babbush, and J. Lee, Unbiasing fermionic quantum Monte Carlo with a quantum computer, [Nature](#) **603**, 416 (2022).
- [107] M. Kalinowski, N. Maskara, and M. D. Lukin, Non-Abelian Floquet Spin Liquids in a Digital Rydberg Simulator, [Phys. Rev. X](#) **13**, 031008 (2023).
- [108] M. S. J. Tepaske, D. J. Luitz, and D. Hahn, [Optimal compression of constrained quantum time evolution](#) (2023), [arXiv:2311.06347 \[cond-mat, physics:quant-ph\]](#).
- [109] Qiskit contributors, [Qiskit: An open-source framework for quantum computing](#) (2023).
- [110] M. Hermele, M. P. A. Fisher, and L. Balents, Pyrochlore photons: The $U(1)$ spin liquid in a $S = 1/2$ three-dimensional frustrated magnet, [Physical Review B](#) **69**, 064404 (2004).
- [111] K. Wan, M. Berta, and E. T. Campbell, Randomized Quantum Algorithm for Statistical Phase Estimation, [Physical Review Letters](#) **129**, 030503 (2022).
- [112] I. Zintchenko and N. Wiebe, Randomized gap and amplitude estimation, [Physical Review A](#) **93**, 062306 (2016), [arXiv:1603.03672 \[quant-ph\]](#).
- [113] R. Haghshenas, J. Gray, A. C. Potter, and G. K.-L. Chan, The Variational Power of Quantum Circuit Tensor Networks, [Physical Review X](#) **12**, 011047 (2022), [arXiv:2107.01307 \[cond-mat, physics:quant-ph\]](#).
- [114] M. Foss-Feig, D. Hayes, J. M. Dreiling, C. Figgatt, J. P. Gaebler, S. A. Moses, J. M. Pino, and A. C. Potter, Holographic quantum algorithms for simulating correlated spin systems, [Physical Review Research](#) **3**, 033002 (2021).
- [115] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, The theory of variational hybrid quantum-classical algorithms, [New Journal of Physics](#) **18**, 023023 (2016).
- [116] A. Schuckert, A. Bohrdt, E. Crane, and M. Knap, [Probing finite-temperature observables in quantum simulators with short-time dynamics](#) (2022), [arXiv:2206.01756 \[cond-mat, physics:quant-ph\]](#).
- [117] M. Knap, A. Kantian, T. Giamarchi, I. Bloch, M. D. Lukin, and E. Demler, Probing real-space and time resolved correlation functions with many-body Ramsey interferometry, [Physical Review Letters](#) **111**, 147205 (2013), [arXiv:1307.0006 \[cond-mat, physics:quant-ph\]](#).
- [118] I. Affleck, T. Kennedy, E. H. Lieb, and H. Tasaki, Rigorous results on valence-bond ground states in antiferromagnets, [Physical Review Letters](#) **59**, 799 (1987).
- [119] F. T. Arecchi, E. Courtens, R. Gilmore, and H. Thomas, Atomic Coherent States in Quantum Optics, [Physical Review A](#) **6**, 2211 (1972).
- [120] L. H. Pedersen, N. M. Møller, and K. Mølmer, Fidelity of quantum operations, [Phys. Lett. A](#) **367**, 47 (2007).
- [121] A. Elben, S. T. Flammia, H.-Y. Huang, R. Kueng, J. Preskill, B. Vermersch, and P. Zoller, The randomized measurement toolbox, [Nature Reviews Physics](#) **5**, 9 (2022).

Acknowledgments. We thank A. Aldossary, T. Betley, D. Bluvstein, M. Cain, L. Cunha, J. Feldmeier, C. Kokail, J. Lee, N. Leitao, S. Gant, J. Haber, S. Hollerith, T. Manovitz, J. Neaton, S. Sachdev, A. Schuckert, K. Seethram, and P. Zoller for insightful discussions. This work was supported by the US Department of Energy [DE-SC0021013 and DOE Quantum Systems Accelerator Center (contract no. 7568717)], the Defense Advanced Research Projects Agency (grant no. W911NF2010021), the National Science Foundation (via CUA PFC, Q-IDEAS HDR, PHY-2207972), the Department of Defense Multidisciplinary University Research Initiative (ARO MURI, grant no. W911NF2010082), Wellcome Leap under the Q for Bio program, the DARPA IMPAQT programme (grant number HR0011-23-3-0030) and the Harvard-MIT Center for Ultracold Atoms. N.M. acknowledges support from the Department of Energy Computational Science Graduate Fellowship under Award Number DE-SC0021110. A.M.G and R.A.B acknowledges support from the NSF through the Graduate Research Fellowships Program, and A.M.G also acknowledges support through the Theodore H. Ashford Fellowships in the Sciences.

APPENDIX

Hamiltonian Engineering

The Hamiltonian engineering toolbox introduced here is based on the average Hamiltonian approach. This approach uses the fact that in the high-frequency limit, the effective Floquet cycle period $K\tau$ is much smaller than the inverse local energy scales of $H(k)$, and the Floquet Hamiltonian H_F can be well-approximated by expanding in a small parameter $\frac{K\tau}{||H(k)||_{\text{local}}}$ [62]. The leading contribution is the *average Hamiltonian*

$$H_F^{(0)} = \frac{1}{K} \sum_{k=0}^{K-1} H(k). \quad (10)$$

The second order term also takes a simple form,

$$H_F^{(1)} = \frac{\tau}{2K} \sum_{k < k'} [H(k), H(k')]. \quad (11)$$

The results presented in the main text involve engineering the average Hamiltonian $H_F^{(0)}$ to reproduce the target (1). In this setting, the second order term generates simulation errors, and is order $O(\tau)$. It can be cancelled by selecting time-reversal symmetric sequences of length $2K$, where the second half of the pulse is defined by $\Theta_k = \Theta_{K-1-k}$. This reduces simulation errors to order $O(\tau^2)$, but might potentially alter the pre-thermal properties of the Floquet Hamiltonian, which is an interesting problem for further research.

Going beyond average Hamiltonian engineering is also possible by optimizing the Floquet sequence [66, 107, 108]. We demonstrate engineering of higher-order terms for a small system composed of two interacting spin-3/2's, to controllably engineer up to bi-cubic terms $(\hat{S}_i \cdot \hat{S}_j)^3$ using only two- and three-qubit operations in Extended Data Fig. 7 and SM.

High-spin Hamiltonian engineering with dynamical Floquet projection

To compute the form of the interaction Hamiltonian in the rotated frame $H_I(k) = e^{i\Theta_k H_P} H_I e^{-i\Theta_k H_P}$ as used in (5) let's consider a single spin-1/2 particle belonging to a spin- S_i cluster, and define projectors $P_i = P[\hat{S}_i^2]$ onto the symmetric space and $Q_i = 1 - P_i$ onto its complement. The spin-1/2 term can be split into four parts

$$\hat{s}_{i,1} = \underbrace{P_i \hat{s}_{i,1} P_i + Q_i \hat{s}_{i,1} Q_i}_{\text{symmetric}} + \underbrace{P_i \hat{s}_{i,1} Q_i + Q_i \hat{s}_{i,1} P_i}_{\text{non-symmetric}} \quad (12)$$

where the first two terms preserve the on-site total spin, and the second two change the on-site total spin. Since H_P acts as $1 - P_i$ on the i -th spin, we label terms by how they change the expectation value of $(1 - P_i)$

$$\hat{s}_{i,1}^{(0)} = P_i \hat{s}_{i,1} P_i + Q_i \hat{s}_{i,1} Q_i \quad (13)$$

$$\hat{s}_{i,1}^{(+1)} = Q_i \hat{s}_{i,1} P_i \quad (14)$$

$$\hat{s}_{i,1}^{(-1)} = P_i \hat{s}_{i,1} Q_i \quad (15)$$

This rule can be extended to higher-weight operators. For example, two-spin interactions between clusters $h_{ij} = \hat{s}_{i,1} \cdot \hat{s}_{j,1}$ decompose into five parts $h_{ij}^{(n)}$, $n = -2, -1, 0, 1, 2$,

$$h_{ij}(k) = \underbrace{\hat{s}_{i,1}^{(0)} \cdot \hat{s}_{j,1}^{(0)} + \hat{s}_{i,1}^{(+1)} \cdot \hat{s}_{j,1}^{(-1)} + \hat{s}_{i,1}^{(-1)} \cdot \hat{s}_{j,1}^{(+1)}}_{h_{ij}^{(0)}} \quad (16)$$

$$+ \underbrace{\left(\hat{s}_{i,1}^{(+1)} \cdot \hat{s}_{j,1}^{(0)} + \hat{s}_{i,1}^{(0)} \cdot \hat{s}_{j,1}^{(+1)} \right)}_{h_{ij}^{(+1)}} + h.c. \quad (17)$$

$$+ \underbrace{\left(\hat{s}_{i,1}^{(+1)} \cdot \hat{s}_{j,1}^{(+1)} \right)}_{h_{ij}^{(+2)}} + h.c., \quad (18)$$

corresponding to the different ways to change the expectation value of $2 - P_i - P_j$. n -spin interactions will have parts running from $h^{(-n)}$ to $h^{(+n)}$. Therefore, in the rotating frame, the Hamiltonian terms transform as

$$h(k) = \sum_{n=-n_{\max}}^{n_{\max}} h^{(n)} e^{i\Theta_k n}. \quad (19)$$

We further choose a sequence of Θ_k 's such that only the $h^{(0)}$ contribution is non-zero on average. The simplest sequence which satisfies these conditions are a family of cyclic pulses of order P .

$$\Theta_k = \frac{2\pi i}{P} k \quad k = 0, \dots, P-1 \quad (20)$$

which satisfy the cancellation condition as long as $n_{\max} < P$. In the two-body case, of the terms which contribute to $h_{ij}^{(0)}$, only $\hat{s}_{i,1}^{(0)} \cdot \hat{s}_{j,1}^{(0)}$ acts non-trivially in the symmetric subspace, so we focus on this term. An explicit form can be computed by decomposing $\hat{s}_{i,1}$ into its permutation-symmetric and orthogonal components,

$$\hat{s}_{i,1} = \underbrace{\frac{1}{2S_i} \sum_{a=1}^{2S_i} \hat{s}_{i,a}}_{\text{symmetric}} + \underbrace{\left(\frac{2S_i - 1}{2S_i} \hat{s}_{i,1} - \frac{1}{2S_i} \sum_{a'=2}^{2S_i} \hat{s}_{i,a'} \right)}_{\text{non-symmetric}}. \quad (21)$$

Therefore, the symmetric part $\hat{s}_{i,1}^{(0)} = \frac{1}{2S_i} \hat{S}_i$ is proportional to the collective spin.

With this understanding, we can construct a spin-1/2 interaction Hamiltonian H_I that recovers (1) under projection, by replacing each n -site high-spin interaction with an analogous spin-1/2 one. For example, for $n = 2$ the replacement proceeds as

$$\begin{aligned} J_{ij}^{\alpha\beta} \hat{S}_i^\alpha \hat{S}_j^\beta &\rightarrow \overline{J}_{ij}^{\alpha\beta} \hat{s}_{i,a}^\alpha \hat{s}_{j,b}^\beta \\ \overline{J}_{ij}^{\alpha\beta} &= 4S_i S_j J_{ij}^{\alpha\beta}, \end{aligned} \quad (22)$$

where intra-cluster indexes a, b encode which representative from spin's i and j are used to generate the interaction. The interaction strength is further boosted to $\overline{J}_{ij}$, to account for the $\frac{1}{2S}$ factor in (21). A straightforward calculation shows that for higher-weight interactions (e.g. $n_{\max}$), the interaction should also be boosted (e.g. $\overline{K}_{ijk} = 8S_i S_j S_k K_{ijk}$) to recover the target large-spin operator under projection.

In general, it may not be feasible to uniquely assign each spin- S_i interaction in H to qubits in H_I especially when a spin- S_i is involved in more than $2S$ interactions. In this case, H_I can be implemented by splitting it into a sequence of non-overlapping groups $H_{I,1}, \dots, H_{I,D}$ that approximate H_I on average. Each sequence can handle up to $D(2S)$ interactions

per spin, so if d is the interaction degree then we require a sequence of length $D = \lceil \frac{d}{2S} \rceil$. While manual decompositions sufficed for the models studied here, automated methods to determine efficient decompositions will have to be developed for more complex systems.

Multi-qubit gates with Rydberg blockade

The Rydberg Hamiltonian governing a cluster of N atoms is

$$H_{\text{cluster}} = \frac{\Omega_q(t)}{2} \sum_i |1\rangle_i \langle 0| + \frac{\Omega_r(t)}{2} \sum_i |r\rangle_i \langle 1| + h.c. + \sum_{i<j} V_{ij} |r\rangle_i \langle r| \otimes |r\rangle_j \langle r| \quad (23)$$

where $\Omega_q(t), \Omega_r(t)$ are complex valued driving fields. In the blockade approximation, which is valid when $V_{ij} \gg \Omega_r$, there is at most one atom in state $|r\rangle$. Therefore, at leading order in Ω_r/V_{ij} , H_{cluster} is approximated by projecting into the manifold of blockade consistent states. This produces an interacting model with an emergent permutation symmetry (see SM). This symmetry allows us to write H_{cluster} in a low-dimensional representation of the Hilbert space scaling as $O(NN_S)$ for a representation including the N_S largest total-spin sectors (see SM).

Figure 2b shows optimization results for an alternating ansatz with separate $\Omega_r(t)$ and $\Omega_q(t)$ applications (solid lines) versus a dual driving scheme with simultaneous field control (dashed lines). For the alternating ansatz the optimization process begins with finding short sequences of symmetric diagonal gates $D(\phi)$ and global single-qubit rotations $Q(\theta)$ that combined realize U_S and U_P (see SM). While $Q(\theta)$ uses global $\Omega_q(t)$ control, $D(\phi)$ involves multi-qubit interactions, and the pulse sequences to realize $D(\phi)$ can be optimized through gradient ascent pulse engineering (GRAPE) [57, 75, 77]. Numerical optimization is feasible due to the manageable size of the low-dimensional basis. We find that the maximum gate time $T^*(n)$ for generic phases (see SM). Gate times in Fig. 2 are computed by multiplying $T^*(n)$ with the shortest sequence length determined in the first step. The operations $Q(\theta)$ and $D(\phi)$ can also be promoted to controlled operations (see SM), as required for the controlled perturbation in Fig. 4a.

Dual driving gate profiles are directly optimized using GRAPE, with an added smoothness regularization to ensure driving profiles can be implemented with available classical controls (see SM and Extended data Fig. 7). Interestingly the optimized fidelity is roughly system size independent, up to error rates around 10^{-3} . We select this fidelity threshold, and plot the resulting gate times in Fig. 2b (dotted lines). For more stringent thresholds, such as 10^{-6} , we recover an approximately linear dependence in n , producing gate times comparable to the alternating decomposition.

For comparison, we estimate gate counts for a decomposition of U_S and U_P into single and two qubit gates (see Fig. 2c). Using the Qiskit transpiler [109], we find two-qubit decompositions into single qubit rotations and CPhase gates for U_P . For U_S , the diagonal gates $D(\phi)$ found numerically in the alternating ansatz only require two-qubit interactions, and can be decomposed into $\binom{n}{2}$ two-qubit CPhase gates, outperforming the Qiskit result for $n \geq 3$. Finally, we show in SM that symmetric operations like U_S and U_P can be implemented with $\text{poly}(n)$ two-qubit gates using ancilla qubits, by constructing an efficient MPO representation of arbitrary n -spin operations.

The fidelity of the multi-qubit gates is subject to errors like spontaneous emission and dephasing due to a finite T_2^* , as seen in current Rydberg gates [57]. These errors can be mitigated by improving the excitation schemes. Specifically, single-photon schemes as used in as in Refs. [67, 68], avoiding intermediate-state scattering, may enhance performance for bigger clusters. This is because the rate of decay from the Rydberg state does not depend on cluster size, since the number of Rydberg excitations is never larger than one.

Estimating simulation time

The simulation time is defined as the maximum evolution time, before which the typical error-per-qubit is below some target threshold ϵ . We account for both coherent Hamiltonian simulation errors and incoherent gate errors. For a symmetrized sequence, we estimate the scaling of both contributions to be

$$\epsilon_{\text{sim}} = (c_2 \tau^2)^2 T^2, \epsilon_{\text{gate}} = \frac{gT}{\tau}, \quad (24)$$

with target evolution time T and step-size τ . Here, the coefficient c_2 depends on the detail of the Hamiltonian simulation protocol, and can be estimated from numerics or the third order term in the Magnus expansion (see SM). The coefficient g measures the gate error probability per cycle, and is determined by assuming each multi-qubit gate has a fixed probability of failure $T_{\text{gate}} g_0$ which scales linearly in the time of the gate. For simplicity, we utilize the estimated T_{gate} for large-angle unitary U_P with $\theta = \pi$. Both c_2 and g are estimated to grow extensively in system size. Thus, we work instead with the intensive version of these quantities, $\tilde{c}_2 = c_2/L$ and $\tilde{g} = g/L$.

Optimizing the step-size τ to minimize error (see SM), the maximum evolution time scales as

$$T_{\text{opt}} = \frac{2^{2/3}}{5^{5/6}} \frac{\epsilon^{5/6}/L}{(\tilde{c}_2 \tilde{g}^2)^{1/3}} \quad (25)$$

for a target error rate ϵ .

In Fig. 3, we use this formula, along with numerical estimates for c_2 in models (i) and (ii), and heuristic estimates

for models (iii) and (iv). We further estimate g using the simultaneous driving gate times of Fig. 2b, and select an error per Rabi cycle of $g_0 = 10^{-3}$. The target error we select is $\varepsilon^{5/6}/L = 0.1$, which grows approximately extensively with system size.

We illustrate the benefits of our approach on four example Hamiltonians (Fig. 3). First we consider the spin-1/2 Heisenberg model on a Kagome lattice (i). The standard two-qubit decomposition involves applying a sequence of six-steps, each applying a $e^{-i\theta \mathbf{s}_i \cdot \mathbf{s}_j}$ gate along an edge (i, j) (solid lines in Fig. 3). However, using native three-qubit Heisenberg gates, we group the interactions into upwards and downwards facing triangles (dotted lines Fig. 3), reducing the sequence to only two-steps, improving c_2 . The three-qubit gate also generates interactions more efficiently, further improving g . Note that four-qubit version of this scheme could simulate the spin-1/2 Heisenberg model on a pyrochlore lattice using $K = 2$ [110].

The second model (ii) are two interacting high-spins with $S_i = 5/2$, which interact via a mix of Heisenberg and Dzyaloshinskii-Moriya (DM) interactions,

$$J_{ij}^{\alpha\beta} = J\delta_{\alpha\beta} + D \sum_{\gamma} \epsilon_{\alpha\beta\gamma}.$$

The conventional decomposition splits the 25 pairwise interactions into five groups of five, which are applied sequentially. In contrast, there is a dynamical projection scheme for this model with $K = 2$ (see SM). This improvement in Floquet engineering efficiency compensates the overhead from introducing a five-qubit gate, leading to slightly higher accessible simulation times. Scaling with S_i is discussed in SM.

The final two models (iii and iv) are composed of spin-2 particles on a square lattice. In model (iii), nearest-neighbor spins interact via a Heisenberg and bi-quadratic interaction, $\sum_{k=1}^2 J_k (\mathbf{S}_i \cdot \mathbf{S}_j)^k$. The conventional Trotter approach requires realizing many four-qubit interactions per spin-2 producing very long Floquet periods. In contrast, dynamical projection can be implemented by applying a single four-qubit gate per edge, resulting in smaller K (see SM). The multi-qubit implementation is also significantly more efficient, as the two-qubit decomposition of a four-qubit interaction comes with large overhead, as illustrated in Fig. 2c. To highlight the separate contribution from Floquet projection and multi-qubit gates, we compute the effective simulation time for four cases in the table below.

$T_{\text{sc}} \times \ H\ _{\text{local}}$	Two-qubit gates	Four-qubit gates
Trotterization	0.15	1.0
Dynamical Projection	0.25	2.5

TABLE I. The quantum simulation's achievable coherence time T_{sc} in the target Hamiltonian's units $\|H\|_{\text{local}}$ (see also Fig. 3) for four different simulation schemes applied to model (iii). This comparison illustrates the advantages of the techniques we propose here. Hamiltonian engineering based on dynamical projection in combination with efficient, hardware-optimized interactions outperforms alternative approaches based on bare Trotter evolution of two- or multi-qubit gates, as well as dynamical projection just using two-qubit gates.

Finally, we consider a model including up to bi-quartic terms $k \leq 4$. Here, we further consider using an eight-qubit multi-qubit gate to directly realize the interaction between two large-spins (iv). This interaction naturally preserves the symmetry, so dynamical projection is not needed in this case. As such, the speedup comes entirely from the efficiency of the optimized eight-qubit operation, compared against the large overhead associated with a two-qubit decomposition of an eight-qubit interaction.

Many-body Spectroscopy

To complete our simulation framework, we also develop tools for resource-efficient readout of Hamiltonian properties. First, we illustrate how to compute two-time correlation functions of the form

$$C_{O,R}(t) = \langle S|O(t)R(0)|S\rangle, \quad (26)$$

using the circuit in Fig. 4a. The real and imaginary parts of $C_{O,R}(t)$ are independently accessed by measuring the ancilla in the X and Y basis, respectively [80, 81, 83, 111]. More concretely, consider the state of the system right before measurement, including both the ancilla qubit and the system,

$$|\psi_f\rangle = \frac{1}{\sqrt{2}} (|0\rangle \otimes U(t)|S\rangle + |1\rangle \otimes U(t)R|S\rangle), \quad (27)$$

where U is the time-evolution operator. Measuring $X \otimes O$ or $Y \otimes O$ results in

$$\langle X \otimes O \rangle_{\psi_f} = \frac{1}{2} (\langle S| R^\dagger O(t) |S\rangle + \langle S| O(t) R |S\rangle), \quad (28)$$

$$\langle Y \otimes O \rangle_{\psi_f} = \frac{i}{2} (\langle S| R^\dagger O(t) |S\rangle - \langle S| O(t) R |S\rangle), \quad (29)$$

which together gives the full complex-valued $C_{O,R}(t)$ by taking a linear combination of the two,

$$C_{O,R} = \langle (X + iY) \otimes O \rangle_{\psi_f}. \quad (30)$$

For observables O diagonal in the measurement basis, $C_{O,R}$ can be efficiently estimated in parallel from snapshots. During the i -th run, let $\mu^{(i)} = \{x, y\}$ be the randomly sampled ancilla measurement basis, and $a^{(i)} = \{0, 1\}$, and $|b^{(i)}\rangle$ be the ancilla and system measurement outcomes respectively. Then, the estimator can be written as

$$\overline{C_{O,R}(t)} = \frac{1}{M} \sum_{i=1}^M 2\sigma(\mu^{(i)}, a^{(i)}) \langle b^{(i)} | O | b^{(i)} \rangle \quad (31)$$

where σ is a function taking on the values

$$\begin{aligned}\sigma(x, 0) &= +1, & \sigma(x, 1) &= -1, \\ \sigma(y, 0) &= +i, & \sigma(y, 1) &= -i.\end{aligned}\quad (32)$$

and $|b^{(i)}\rangle$ is the measured projected state.

These measurements can be used to compute the operator-resolved density of states (8), which can be rewritten as follows

$$\begin{aligned}D^A(\omega) &= \int dt e^{i\omega t} \sum_n \text{Tr}[A|n\rangle e^{-i\epsilon_n t} \langle n|] \\ &= \int dt e^{i\omega t} \text{Tr}[AU(t)],\end{aligned}\quad (33)$$

where we have replaced $\delta(\omega - \epsilon_n) \rightarrow \int dt e^{i(\omega - \epsilon_n)t}$. In practice, we will sample evolution times t from a probability distribution $p(t)$, such that the integral is normalized to one when $\omega = \epsilon_n$, i. e., $\int dt e^{i\omega t} p(t) = 1$. To arrive at (9), we can replace the trace with an average over probe states [112]

$$\text{Tr}[AU(t)] = \tilde{\mathbb{E}}_{R \sim \mathcal{R}} \langle R|AU(t)|R \rangle, \quad (34)$$

where $\tilde{\mathbb{E}}_{R \sim \mathcal{R}} = \text{Tr}[\mathbb{1}] \mathbb{E}_{R \sim \mathcal{R}}$ is a normalized expectation value, and $\text{Tr}[\mathbb{1}]$ is the dimensionality of the Hilbert space. This is valid, as long as the ensemble forms a 2-design,

$$\mathbb{E}_{R \sim \mathcal{R}} |R\rangle \langle R| = \frac{\mathbb{1}}{\text{Tr}[\mathbb{1}]}, \quad (35)$$

Observables such as (34) can in principle be computed via a modified Hadamard test by applying controlled-time evolution (see Ref. [81]). Since time-evolution is generally the most costly step, we avoid the overhead associated with controlled-evolution and instead utilize a reference state $|S\rangle$ with simple time-evolution. In particular, we select an ensemble of observables O_s such that

$$\frac{1}{\mathcal{N}(O_s)} \sum_s O_s U(t) |S\rangle \langle S| U^\dagger(t) O_s = \mathbb{1}. \quad (36)$$

The normalization factor $\mathcal{N}(O_s)$ depends on the choice of ensemble. Then, we can insert this resolution of the identity into (34) to get (9). In Fig. 4 and Fig. 5, we consider the polarized reference state $|S\rangle = |0\rangle^{\otimes N}$ which is an exact eigenstate of the Heisenberg Hamiltonian, and the ensemble of Pauli- X operators $O_s = X_s = \bigotimes_{i=1}^N (X_i)^{s_i}$ where s is an N -bit string. This satisfies the condition, and has $\mathcal{N}(X_s) = 1$ since $X_s|0\rangle^{\otimes N}$ is an orthonormal basis. For generic reference states $|S\rangle$ prepared by applying S to $|0\rangle^{\otimes N}$, the ensemble $O_s = S X_s S^\dagger$ satisfies the condition, and can be measured by applying the inverse preparation circuit $S^\dagger$ before measuring in the X -basis. Lastly, an ensemble which is independent of the reference state, is the set of Pauli strings

$P_s = \bigotimes_{i=1}^N \sigma_i^{s_i}$, where s is a base-four string, and σ^{s_i} denotes the four Pauli operators I, X, Y, Z ; this can be accessed with randomized measurements (see SM), and has a normalization factor $\mathcal{N}(P_s) = 2^N$.

Efficient estimation of $D_R^A(t)$ from snapshots

One of the key advantages of our approach, is the ability to compute many estimators $D_R^A(t)$ from one dataset, including complex operators such as projectors onto spin-sectors $A = P_S$. This enhances the sample efficiency and reduces simulation time requirements, which are crucial limited resources in quantum computation.

To estimate density-of-states in a sample efficient way, we introduce a classical co-processing algorithm which utilizes knowledge of the prepared reference state. In particular, we discuss the estimation procedure for

$$\begin{aligned}D_R^A(t) &= \sum_s \underbrace{\langle R|A O_s U(t)|S\rangle}_{\text{classical}} \underbrace{\langle S|U^\dagger(t) O_s U(t)|R\rangle}_{\text{quantum}} \\ &= e^{-iE_S t} \sum_s \langle R|A O_s |S\rangle \langle S|O_s(t)|R\rangle\end{aligned}\quad (37)$$

which becomes (9) after averaging over times t and perturbations R . In the second line, we assumed $|S\rangle$ is a known eigenstate, so $U(t)|S\rangle = e^{-iE_S t}|S\rangle$, and the classical part becomes equivalent to a zero-time correlation function. One challenge is that the sum in (37) involves exponentially many observables. However, by using the estimator (31) this reduces to a sum over M snapshots. In particular, for the polarized reference state, and Pauli- X measurements, a simple calculation (see SM) shows that

$$\overline{D_R^A(t)} = \frac{1}{M} \sum_{i=1}^M e^{-iE_S t} 2^{N/2} \sigma(\mu^{(i)}, a^{(i)}) \langle R|A|b^{(i)}\rangle \quad (38)$$

is an unbiased estimator. The variance of this estimator is independent of system size for appropriately chosen ensembles $\mathcal{R}$, and scales as $O(1/\sqrt{M})$. While $C_{O_s, AR}^*(0)$ could instead be estimated by measuring the unevolved state, this would produce an estimator that converges very slowly, leading to a sample complexity that is exponential in N . However, classical methods can compute $C_{O_s, AR}^*(0)$ with no error, making the procedure much more sample-efficient. In the SM, we discuss how such calculations can be efficiently performed for any pair of reference and probe states with an efficient matrix product state (MPS) description, and any observable A with a matrix product operator (MPO) description. Further, we show that projectors P_S onto spin-sectors can be written as MPOS with poly(N) bond-dimension, making evaluation of $D^{P_S}(\omega)$ classically efficient and scalable. We further discuss how entangled measurements between two systems could be used to estimate these correlators efficiently, even when $|S\rangle$ does not have a known classical description.

Thermal expectation values

The operator-resolved density of states can be used to compute thermal expectation values via [81]

$$\langle A \rangle_\beta = \frac{\int e^{-\beta\omega} D^A(\omega) d\omega}{\int e^{-\beta\omega} D^1(\omega) d\omega}. \quad (39)$$

For example, to compute the magnetic susceptibility, we simply select the operator $A = \beta(S^z)^2$, where $\beta = 1/T$ is the inverse temperature. Interestingly, this method of estimating thermal expectation values is insensitive to uniform spectral broadening of each peak, due to a cancellation between the numerator and denominator (see SM). However, it is highly sensitive to noise at low ω , which is exponentially amplified by $e^{-\beta\omega}$. To address this, we estimate the SNR for each $D^A(\omega)$ independently, and zero-out all points with SNR below three times the average SNR. This potentially introduces some bias by eliminating peaks with low-signal, but ensures the effects of shot-noise are well controlled.

Noise Modelling

To quantify the effect of noise on the engineered time-dynamics, we simulate a microscopic error model by applying a local depolarizing channel with an error probability p at each gate. This results in a decay of the obtained signals for the correlator $D_R^A(t)$. The rate of the exponential decay grows roughly linearly with the weight of the measured operators (see Extended Data Fig. 8). This scaling with operator weight can be captured by instead applying a single depolarizing channel at the end of the time-evolution, with a per-site error probability of γt with an effective noise rate γ . This effective γ also scales roughly linear as a function of the single-qubit error rate per gate p (see Extended Data Fig. 8).

Scaling the approach

Quantum simulations are constrained by the required number of samples and the simulation time needed to reach a certain target accuracy. These factors are crucial for determining the size of Hamiltonians which can be accessed for particular quantum hardware.

Focusing on a single gapped eigenstate we determine the number of snapshots C_M needed to distinguish a spectral peak from noise (Fig. 10). The signal arises from the overlap of the probe states with the target eigenstate. The noise is given by the variance of the estimator (38), and decays as $\epsilon \sim M^{-1/2}$. For certain ensembles of probe states, the variance can be made system size independent (see SM). However, a random probe state will have exponentially vanishing overlap with any specific eigenstate. One approach to mitigate this, is to initialize probe states with higher overlap. In Fig. 9b we show

that for a spin-1 AFM chain a simple bond-dimension two MPS can outperform product states by orders-of-magnitude in ground-state estimation. While bond-dimension two states can be efficiently prepared with simple circuits of two qubit gates, more general ansatz can also be efficiently realized using the simulation techniques described here [113, 114]. Optimized ansatz could be further combined with importance sampling [115], to improve the sample-efficiency of computing finite-temperature or excited state properties (see SM).

The simulation time $T_{\max}$ will depend on the required spectral resolution, which does not scale with system size for a gapped eigenstate. However, the rate of spectral broadening depends sensitively on the weight of measured observables (Fig. 8). When the reference state is high in energy, such as the polarized state for an AFM chain, the relevant observables typically have extensive weight, requiring $T_{\max} \sim N$ to maintain constant spectral resolution. In contrast, preparing a low-energy reference state, such as the ground-state $|S\rangle = |GS\rangle$, allows coupling to other low-energy states using low-weight operators. This results in a noise-resilient and system-size independent procedure (Fig. 9). We further note that ground-state preparation can be approximate, which would result in additional spectral broadening in the computation of $D^A(\omega)$. While the spectral resolution requirements should also grow as the gap shrinks, we have illustrated that operator-resolution can mitigate this in certain settings (e.g. Fig. 5 and Fig. 6). As such, understanding the general capabilities of this approach is an interesting direction for continued research.

OEC Hamiltonians

The two candidates for the closed S2 state of the oxygen evolving complex (OEC) are parameterized with Heisenberg models $H = -\sum_{ij} J_{ij} \hat{S}_i \cdot \hat{S}_j$ [27, 88]. The coupling constants for the two cases used to generate the data presented in Fig. 5 are summarized in Table II. The corresponding spin sizes of the different sites are $S_1 = 3/2$, $S_2 = 3/2$, $S_3 = 3/2$, $S_4 = 2$.

	J_{12}	J_{13}	J_{14}	J_{23}	J_{24}	J_{34}
S₂H-1b	30.5	12.9	4.5	36.5	1.3	-7.3
S₂H-2b	32.6	11.7	4.0	37.3	1.5	-2.6

TABLE II. Exchange coupling constants in units of wavenumbers $(\text{cm})^{-1}$ for the two different candidates for the S_2 stage of the OEC used in Fig. 5.

Additional information about eigenstates can be calculated by choosing the operator A in the operator resolved density of states appropriately, and multiplying by a narrow “band-pass” filter in Fourier space to isolate a small set of frequencies [81, 116]. For example, we investigate the total spin of the cubane sub-unit, i. e., the three magnetic sites supported on opposite vertices of the cube, using $A = (\hat{S}_1 + \hat{S}_2 + \hat{S}_3)^2$, and compute

$$A(\omega) = \frac{D^A(\omega)}{D^I(\omega)} \quad (40)$$

evaluated at energies ω_n of the eigenstates, which can be extract from peaks in the spectral functions (Extended data fig. 11). We further use spin-projection to improve the estimate in the presence of broadening. For example, if the peak at ω_n occurs in spin-sector P_S , we insert the projector P_S in the numerator and denominator, $A(\omega_n) = \frac{D^{AP_S}(\omega_n)}{D^{PS}(\omega_n)}$.

2D Heisenberg calculations

The square lattice Heisenberg calculation was performed on a large ($L \times L$) system, with Hamiltonian


$$H_{2D} = -J \sum_{\mathbf{r}} \left(\hat{\mathbf{S}}_{\mathbf{r}} \cdot \hat{\mathbf{S}}_{\mathbf{r}+\hat{x}} + \hat{\mathbf{S}}_{\mathbf{r}} \cdot \hat{\mathbf{S}}_{\mathbf{r}+\hat{y}} \right) \quad (41)$$

We measure the Green's function $G(\mathbf{r}, t) = \langle S | X_{\mathbf{r}}(t) X_0 | S \rangle$ from a polarized reference $|S\rangle^{\otimes L^2}$ [117]. Since S^z is conserved under the dynamics $G(\mathbf{r}, t)$ is classically simulated by restricting it to the space containing the $|S\rangle$ and single spin-flip states $X_{\mathbf{r}}|S\rangle$, which has dimension $L^2 + 1$. We evolve under equally spaced times up to $JT_{\max} = 8$, and select $L = 23$ which is large enough such that $G(r, t)$ vanishes far from the boundaries. Therefore, by letting $G(r, t) = 0$ outside the simulated region, this provides a good approximation for the $L \rightarrow \infty$ limit. As such, we define projectors onto (unnormalized) plane-wave states $P_{\mathbf{k}} = |\mathbf{k}\rangle\langle\mathbf{k}|$, where $\langle\mathbf{k}|X_{\mathbf{r}}|S\rangle = e^{-i\mathbf{k}\cdot\mathbf{r}}$. Then, $D^{P_{\mathbf{k}}}(\omega)$ can be written as


$$D^{P_{\mathbf{k}}}(\omega) = \int dt e^{i(\omega - \omega_0)t} \sum_{\mathbf{r}} \underbrace{\langle S | X_0 P_{\mathbf{k}} X_{\mathbf{r}} | S \rangle}_{e^{i\mathbf{k}\cdot(\mathbf{r}-0)}} \underbrace{\langle S | X_{\mathbf{r}}(t) X_0 | S \rangle}_{G(r, t)} \quad (42)$$

which reduces to the Fourier transform $G(k, \omega)$ when the energy of the polarized state ω_0 is set to zero. Plotting this for a continuous set of k and ω produces the spectral weight depicted in Fig. 6. This further shows that finite-size systems are sufficient to simulate extended systems at finite evolution times.


By computing the peak value of ω for each k in the vicinity of the single-particle excitations, we estimate the dispersion relation $\omega(k)$ associated with a single spin-flip excitation. Computations of the many-particle Green's function could be performed similarly by applying multi-site perturbations, to extract finite-temperature properties and characterize the interactions between the quasi-particles.




Extended Data – FIG. 7. **Gate optimization procedure and variational Hamiltonian Engineering.** (a) To find smooth gates we perform the GRAPE optimization procedure in two steps. In the first step, we penalize rapid changes in the pulse profile by introducing an extra term in the cost-function. In this case, the resulting relationship between ideal gate-error vs. time (yellow circles) saturates around 10^{-6} . For the second step we initialize the search with the smooth gates found in the previous step, which are re-optimized by removing the smooth penalty. This significantly reduces the ideal gate error (orange). The data are shown for a $N = 4$ qubit cluster. Since the first step already confined the problem into a subspace of the search space with smooth gates, the resultant pulses also remain smooth after the second step. On the right we show an example smoothened pulse profile for the hyperfine angle θ and the Rydberg phase ϕ (see Methods) with an ideal gate error rate $\mathcal{E} = 10^{-3}$. (b) Higher order interaction terms can be controllably engineered via a simple modification of the $K = 6$ Floquet projection sequence. For example, reducing the second time-step and increasing the fourth time-step by the same amount, δ preserves the target Hamiltonian at leading order. (c) By tuning τ and δ , for a system of two interacting spin-3/2's, a very large family of coefficients J_1, J_2, J_3 in the general Hamiltonian $H_T = J_1(\mathbf{S}_1 \cdot \mathbf{S}_2) + J_2(\mathbf{S}_1 \cdot \mathbf{S}_2)^2 + J_3(\mathbf{S}_1 \cdot \mathbf{S}_2)^3$ can be engineered. In particular, the roughly horizontal gray lines correspond to constant δ grid lines, and roughly vertical gray lines correspond to values with constant τ . We see that changing δ primarily modifies J_2 and J_3 , while changing τ primarily modifies J_1 , consistent with our analysis that δ picks out certain higher-order terms. We further note that the especially interesting AKLT family, with $J_2/J_1 = 116/243$ and $J_3/J_1 = 16/243$ lies among the family of efficiently engineerable interactions.




Extended Data – FIG. 8. **Gate-level noise simulations.** (a)-(c) Time evolution of the correlations $D_R^A(t)$ for a polarized reference state $|S\rangle = |0\rangle$ and A chosen to be projectors onto different S^z sectors. This involves measuring operators of fixed weight. We see that the simulation of gate-level noise modelled as a depolarizing channel with a gate error probability $p = 0.001$ (orange dash dotted curve) matches well with the dynamics obtained by adding an additional phenomenological noise $\propto \exp(-\gamma t)$ with rate $\gamma = (0.128, 0.251, 0.36)$ for $N_{\text{spin flip}} = (1, 4, 8)$ (solid yellow lines) to the loss-less time evolution $D_{R,\text{Floquet}}^A$ (green dashed). (d) Relationship between the optimal phenomenological noise rate γ and the gate error probability p for different spin flip sectors. We see a roughly linear-relationship of the decay rate γ with the gate error rate p and operator weight S_z . All data are obtained using time-steps $J\tau = 0.05$.



Extended Data – FIG. 9. **Noise susceptibility and eigenstate overlaps.** (a) Density-of-states for a spin-1 AFM chain, computed from a polarized reference state $|S\rangle = |0\rangle^{\otimes N}$ (top), and ground state reference $|S\rangle = |GS\rangle$ (bottom). The spectrum is separated into sectors distinguished by their operator weight from $|S\rangle$. For the polarized state, these correspond to sectors with well-defined S^z . For the ground-state, each sector is orthogonalized with respect to lower-weight sectors. Each sector is phenomenologically broadened by $e^{-\gamma t n_{\text{flips}}}$ to simulate the operator-weight dependence of decoherence. When computing low-energy properties, the ground-state reference is more robust to noise, since the low-energy eigenstates can be reached with lower-weight operators. (b) The amplitude of the corresponding spectral peak is determined by the eigenstate overlap. We analyze the ground-state overlap for an AFM spin-1 chain performing DMRG for low bond dimensions ($D = 1, 2$) and find that it is significantly larger for bond-dimension $D = 2$ matrix product states (red diamonds) compared to bond-dimension $D = 1$, i.e., mean-field states (green circles). Interestingly, the ground state overlap decays slower with the chain length for $D = 2$ indicating that the fidelity density is large. This feature makes low bond dimension states a promising direction for efficient state preparation within our scheme since they can efficiently be decomposed into short circuits.



Extended Data – FIG. 10. **Shot noise scaling with system size and convergence with number of snapshots.** (a) Numerically computed standard deviations of the estimator (9) for the density of states $D^A(\omega)$ of a spin-1 AFM chain, for different chain lengths and observables. Here, we consider a polarized reference state $|S\rangle$ and random single-qubit rotations R for the controlled-perturbation. For the bare density of states the standard deviation slowly scales with the chain length (blue circles). For the projectors into the zero (gray diamonds) and three (orange circles) spin-flip sectors the standard deviation is consistent with being independent of system size. (b) Convergence of estimator with number of samples for two different observables. Shaded regions correspond to 2σ error bars around the mean, and decrease with the number of measurements as $1/\sqrt{M}$. Dark lines are running averages for a specific sampled dataset. The sample complexity is defined as the number of samples needed such that the error bars around the mean do not include zero. After this point, spectral peaks can be reliably distinguished from noise.



Extended Data – FIG. 11. **Additional observables calculated for OEC.** By choosing A to be the local cubane spin S_{123}^2 , and calculating $D^A(\omega)$ at each of the peaks in the density of states, we can infer the local structure of spin correlations within each eigenstate. (a) The cubane spin for each eigenstate for the open cubane configuration ($S_2\text{H-1a}$). This model has coefficients: $J_{12} = -15.8$, $J_{13} = 1.9$, J_{14} , $J_{23} = 23.1$, $J_{24} = 1.9$, $J_{34} = -13.9$ in cm^{-1} , and spin-sizes $S_1 = 2$, $S_2 = S_3 = S_4 = 3/2$. The triangles, circles and crosses are obtained using three different methods. The crosses indicate the exact value obtained via exact diagonalization (ED), the green circles indicate the values obtained via the spin resolved (projected) density of states (DOS), and the orange triangles are obtained by the bare DOS. The projected DOS, which is able to resolve individual eigenstates, matches well with the ED result. In contrast the bare DOS, which has overlapping peaks (see Fig. 5), does not capture the expectation values accurately. (b) Same as in (a) but for the closed cubane configuration $S_2\text{H-1b}$. This model has $S_1 = S_2 = S_3 = 3/2$ and $S_4 = 2$. We see that for the lowest-lying cluster, all spins in the cubane subunit are approximately polarized with a total spin value $S = 9/2 = 4.5$. In contrast, the second cluster of eigenstates seems to differ by a single spin flip as indicated by the value $S = 7/2 = 3.5$.

Supplemental Material

HAMILTONIAN ENGINEERING

Variational spin-gate optimization.

An exciting route to efficiently generating more complex interactions is to go beyond average Hamiltonian engineering. In particular, we introduce a p -dependent evolution time τ_p , and consider tuning the parameters τ_p and θ_p in the pulse sequence (4) to control certain higher-order terms, with no additional gate overhead. As a concrete example (see also Extended data Fig. 7), we consider two spin-3/2 particles, evolved by $H_I = \frac{9J}{4} \hat{\mathbf{S}}_{1,1} \cdot \hat{\mathbf{S}}_{2,1}$ and a $K = 6$ Floquet sequence with parameters

$$\Theta = (0, 2\pi/3, 4\pi/3, 4\pi/3, 2\pi/3, 0) \quad (S1)$$

$$\boldsymbol{\tau} = (\tau, \tau - \delta, \tau, \tau, \tau + \delta, \tau). \quad (S2)$$

We find that this sequence naturally engineers an SU(2) symmetric Hamiltonian with Heisenberg interactions, but also bi-quadratic and bi-cubic terms

$$H_{2,extended} = J_1(\hat{\mathbf{S}}_1 \cdot \hat{\mathbf{S}}_2) + J_2(\hat{\mathbf{S}}_1 \cdot \hat{\mathbf{S}}_2)^2 + J_3(\hat{\mathbf{S}}_1 \cdot \hat{\mathbf{S}}_2)^3. \quad (S3)$$

To characterize the values of J_1, J_2 , and J_3 which can be accurately realized, we simulate the action of U_F on the state ρ_{sym} which is maximally mixed in the encoded subspace, for various pairs of values τ, δ . Then, as long as the average population leaving the subspace is below 10^{-4} , we project U_F into the encoded subspace, compute the effective Hamiltonian $H_F = -i \log(U_F)$ by exact diagonalization, and fit the coefficients J_1, J_2, J_3 by solving a set of linear equations. The resulting family of engineerable coefficients forms a 2D surface in a 3D space, and is visualized in Extended data Fig. 7b. Interestingly, the AKLT family of Hamiltonians [S118]

$$H_{\text{AKLT},3/2} = \sum_i \sum_{k=1}^3 J_k^{\text{AKLT}} \left(\vec{S}_i \cdot \vec{S}_{i+1} \right)^k, \quad (S4)$$

lies within the space of engineerable Hamiltonians. Here, the coefficients J_k^{AKLT} are chosen such that $H_{\text{AKLT},3/2}$ is equivalent to the projection Hamiltonian H_P on six qubits. In other words the groundstate of $H_{\text{AKLT},3/2}$ is the total spin-3 subspace, and is separated by a constant gap to the other lower spin subspaces.

The procedure outlined above could also be applied to anisotropic Hamiltonians H_I , to try and engineer the full-space of two-spin operations. Extensions beyond pairwise interactions could also be considered, for example to efficiently realize three-spin interactions as in Ref. [S107]. We note that, higher-order interactions in our approach emerge similarly to those in realistic materials. In our case they arise from virtual total-spin fluctuations, whereas in a material they arise from virtual charge-fluctuations.

Procedures for Multi-Qubit Gate Optimization

The Hamiltonian engineering approach described in this work stands out for its hardware efficiency, owing to the strategic use of efficient multi-qubit gates. To perform optimization efficiently, we first introduce a low-dimensional basis for each cluster, leveraging the permutation symmetry of the blockade interaction. Then, we illustrate two complementary approaches to determining pulse profiles which realize the target gate. First, we show how to use the GRAPE algorithm to find smooth off-diagonal gates where the pulses $\Omega_q(t)$ and $\Omega_r(t)$ are applied simultaneously. Second we describe the two-stage optimization used to identify alternating sequences of diagonal multi-qubit interaction gates and single-qubit gates to generate the target off-diagonal operations.

Low-dimensional construction of blockade cluster basis

To efficiently perform numerical optimization of the pulse sequences $\Omega_q(t)$ and $\Omega_r(t)$ (see below), we use a low-dimensional representation of the Hilbert space, leveraging the permutation-symmetry of the time-dependent Rydberg Hamiltonian (23). In the blockade approximation, states with more than one Rydberg excitation within the cluster are disallowed. When the Rydberg Hamiltonian is projected into this manifold, it becomes an all-to-all interacting ‘‘PXP’’ model, where a spin-flip is only allowed

if all other atoms are in the qubit ($|0\rangle, |1\rangle$) states. Define $P_i = |1\rangle_i\langle 1| + |0\rangle_i\langle 0|$ to be the projector onto the qubit state at site i , and $P = \prod_{i=1}^n P_i$ to be the global projector. Then, the constrained Hamiltonian can be written as

$$H_{\text{proj}} = \frac{\Omega_q(t)}{2} \Sigma_q^+ + \frac{\Omega_r(t)}{2} \Sigma_r^+ P + h.c. \quad (\text{S5})$$

where we have defined collective raising and lowering operators $\Sigma_q^+ = \sum_i |0\rangle_i\langle 1|$ and $\Sigma_r^+ = \sum_i |r\rangle_i\langle 0|$. This is a good approximation for the dynamics at leading order in Ω_r/V , and has an emergent permutation symmetry. In particular, notice that H_{proj} is invariant under exchange of any two sites. In general, the Hilbert space of an N -atom cluster has 3^N states, and the blockade projection further reduces this to an $O(2^N)$ scaling. However, to efficiently perform numerical optimization for larger clusters, we construct a basis whose dimension grows quadratically in N by leveraging the permutation symmetry of H_{proj} .

We start by defining a class of symmetric states which are fully specified their occupation numbers $|n_1, n_0, n_r\rangle$. Each state is constructed by taking a symmetric superposition of all configurations with the corresponding occupations. For example, the $N = 2$ triplet state is given by

$$|1, 1, 0\rangle = (|01\rangle + |10\rangle)/\sqrt{2}.$$

More generally, the states spanned by $|n_1, n_0, 0\rangle$ form the degenerate ground space of $\hat{S}^2$, and have total spin $N/2$ (where $N = n_1 + n_0$). Excited states of S^2 live outside the symmetric space, and have total spin $< N/2$; thus, to form the lowest-lying excited states, i.e., states with total spin $N/2 - 1$, we need to expand our basis.

In general, there are $N - 1$ independent irreducible representations (irrep) with $S = N/2 - 1$ [S119]. One irrep can be formed by constructing states that are symmetric under permutation of spins $i = 1, \dots, N - 1$, but orthogonal to the fully symmetric state. A simple way to construct a basis spanning this $S = N/2 - 1$ irrep and the symmetric $S = N/2$ irrep is to take a tensor product of symmetric spaces on $N - 1$ atoms and 1 atom. To do this, we split the atoms into two groups, A and B , and label states by six numbers $|n_1^A, n_0^A, n_r^A, n_1^B, n_0^B, n_r^B\rangle$, associated with the occupation of each group. Therefore, states must satisfy $n_1^A + n_0^A + n_r^A = N - 1$ and $n_1^B + n_0^B + n_r^B = 1$. The blockade constraint is imposed by further requiring $n_r^A + n_r^B \leq 1$. Then, matrix elements of H_{proj} can be efficiently constructed in this manifold. More generally, we can construct a basis which spans irreps of spin $S = N/2 - k, S = N/2 - k + 1$, to $S = N/2$, by splitting the system into $N - k$ atoms and k atoms. Simulations are still efficient, since the dimensionality scales as $\Theta(N^2)$ when $k = \lfloor N/2 \rfloor$.

Simultaneous drive – Smooth gates with GrAPE

The primary tool we utilize to identify fast, hardware-optimized gate operations is gradient ascent pulse engineering (GrAPE) [S77], which has recently been utilized to develop high-fidelity diagonal gates based on Rydberg driving [S57, S75]. Here, we extend the optimization results to include off-diagonal gates based on simultaneous Rydberg $\Omega_r(t)$ and hyperfine $\Omega_q(t)$ driving. The core of the algorithm involves approximating time-dependent evolution under $H(t)$ with discrete, Trotterized time evolution. The time-dependent profile we consider has two parameters, a phase on the Rydberg drive $\Omega_r(t) = \Omega_r e^{-i\phi(t)}$, and a hyperfine rotation angle $\theta(t)$ defined by $\Omega_q(t) = \partial_t \theta(t)$. The action of the continuous pulse is approximated by discretizing evolution of n_T steps with step-size τ each, resulting in a unitary of the form

$$U_F(\theta, \phi) = \prod_{k=1}^{n_T} e^{i\theta(k\tau)(\Sigma_q^+ + \Sigma_q^-)} e^{-i\tau(e^{i\phi(k\tau)}\Sigma_r^+ + h.c.)} e^{-i\theta(k\tau)(\Sigma_q^+ + \Sigma_q^-)}. \quad (\text{S6})$$

Optimization is performed by fixing $T = n_T \tau$, and tuning discrete set of variational parameters $\theta(k\tau), \phi(k\tau)$ by computing the distance between U_F and the target unitary U_{target} . Since the initial and final states should be fully supported in the qubit manifold, we evaluate the distance between U_{target} and U_F within this manifold. Then, we can measure the gate fidelity using the expression

$$\mathcal{F}(\{\theta, \phi\}) = \frac{N_q + |\sum_q \langle q | U_{\text{target}}^\dagger U_F(\theta, \phi) | q \rangle|^2}{N_q(N_q + 1)} \quad (\text{S7})$$

where $|q\rangle$ is an orthonormal basis for the qubit states and N_q is the dimensionality of the associated subspace [S120]. We note that in our implementation we use the reduced symmetric basis described above, with $k = \lfloor N/2 \rfloor$, so the gate fidelity differs slightly from a calculation in the full qubit space. In particular, it weights the symmetric subspace more heavily than lower-spin subspaces. Nevertheless, the basis is complete, so if $\mathcal{F} = 1$ in the reduced subspace, it also is unity in the full qubit subspace.

When we optimize dual-driving gates, we further add a regularization to the cost function to bias the solution towards smooth gates. Unlike Refs. [S75], we initialize the optimization with random profiles to avoid local minima, so the smoothening

regularization is crucial for obtaining continuous pulse profiles which can realistically be implemented experimentally. The combined cost function is therefore

$$C_\lambda(\{\theta, \phi\}) = (1 - \mathcal{F}(\{\theta, \phi\})) + \lambda \sum_k \left[\left(\frac{\phi(k\tau) - \phi((k-1)\tau)}{2} \right)^2 + \left(\frac{\theta(k\tau) - \theta((k-1)\tau)}{2} \right)^2 \right]. \quad (\text{S8})$$

Here, we penalize the derivative of ϕ and θ , which correspond to the detuning of the Rydberg drive and amplitude of the hyperfine drive.

Next, we discuss a systematic procedure for calculating high-fidelity smooth gates (see Extended Data Fig. 7). To bias the optimizer towards finding smooth gates, we first minimize the cost function C_λ with $\lambda = 10^{-3}$, and evaluate each profile by calculating C_0 . For sufficiently long evolution times T , we are able to reduce C_0 to be about 10^{-3} , and see the gate is reasonably smooth. However, the regularization introduces a small amount of bias. As such, we subsequently re-optimize θ, ϕ using C_0 as the cost function, and see the ideal fidelity improves dramatically but the final gates remain smooth (see Extended Data Fig. 7). Intuitively, the first step where rapid changes in the pulse profile are penalized confines the subsequent second search into a region of search space where gates are smooth.

Alternating Ansatz

A simpler but less time efficient method to realize the required multi-qubit gates involves applying $\Omega_r(t)$ and $\Omega_q(t)$ separately. This alternating approach is a natural extension of the capabilities demonstrated in Ref. [S57], and further generalizes more readily to other reconfigurable architectures [S53, S55].

We parameterize a generic symmetric diagonal operation, and a global qubit rotation, as

$$D(\phi) = \exp \left(-i \sum_{j=0}^N \phi_j \hat{\mathcal{Z}}_j \right) \quad (\text{S9})$$

$$Q(\theta) = \exp \left(-i \sum_{\alpha=\{x,y,z\}} \theta_\alpha \hat{S}^\alpha \right) \quad (\text{S10})$$

where $\mathcal{Z}_j = \sum_{i_1, \dots, i_j} Z_{i_1} Z_{i_2} \dots Z_{i_j}$ is a linear combination of all weight- j Pauli- Z strings supported on the N qubits. In particular, a diagonal gate is characterized by time-evolution under the Rydberg part of the Hamiltonian $\Omega_r(t)$, such that all qubit states return to qubit states at the end of the evolution, with a path-dependent phase. Different qubit states acquire different phases, producing a phase gate. The intuition is that $D(\phi)$ parameterizes all possible diagonal gates which can be generated by the all-to-all Rydberg blockade interaction. After removing the global phase ϕ_0 , and the part which can be generated by a single-qubit Hamiltonian ϕ_1 , there are $N - 1$ remaining parameters which characterize the interaction. The global qubit rotations $Q(\theta)$ generate single-qubit off-diagonal rotations since $\hat{S}^\alpha = \sum_{a=1}^{2^S} \hat{s}_a^\alpha$, but generate no interactions between qubits. As such, alternately applying diagonal and off-diagonal operations enables generation of a large class of off-diagonal interaction operations. To find sequences which realize U_S and U_P , we again use GRAPE-style optimization. In particular, we parameterize the unitary as

$$U_{\text{alternating}} = \prod_{k=1}^{n_{\text{seq}}} Q(\theta_k) D(\phi_k) \quad (\text{S11})$$

and variationally optimize θ_k, ϕ_k to try and match $U_{\text{alternating}}$ to the target unitary. We identify the shortest sequences which are capable of implementing the target unitary, and present the corresponding sequence lengths (n_{seq}) in Table I. We further note

n_{seq}	$N = 2$	$N = 3$	$N = 4$	$N = 5$	$N = 6$	$N = 7$
U_S	2	3	6	6	6	6
U_P	2	3	7	7	10	10

TABLE I. Optimized sequence lengths n_{seq} for target error below $\sim 10^{-3}$.

that for U_S , the sequences only require using the two-qubit phase ϕ_2 , likely because the evolution Hamiltonian $\hat{\mathbf{S}}^2$ is composed

of two-qubit operators. We use this fact below to construct shorter two-qubit decompositions for U_S compared to U_P . Similarly, we find for all sequences (U_S and U_P), the odd-qubit phases $\phi_3, \phi_5, \dots$ can be set to zero.

In neutral-atom setups, the global qubit rotation $Q(\theta)$ can be realized with very high fidelity [S73]. As such, the gate errors are dominated by the interaction step $D(\phi)$, which involves coupling to the Rydberg state. To minimize gate errors, we therefore use GrAPE to identify Rydberg driving profiles $\Omega_r(t)$ that can implement a generic symmetric diagonal gate $D(\phi)$. For a given system size N , by randomly sampling target phases ϕ , and using GrAPE to optimize the profile independently for each setting, we identify a threshold evolution time $T^*(n)$ above which the GrAPE algorithm always reliably identifies a high-fidelity gate (Fig. S1). These times, combined with n_{seq} are used to compute the values in Fig. 2b.

Finally, we discuss how gates $D(\phi)$ can be conveniently promoted to controlled- $D(\phi)$ gates. A generic control gate can be parameterized by

$$CD(\phi) = \exp \left(-i \left(\frac{\mathbb{1} + Z_a}{2} \right) \sum_{j=0}^n \phi_j Z_j \right) \quad (\text{S12})$$

where $\frac{1+Z_a}{2} = |0\rangle\langle 0|$ is a projector onto the $|0\rangle$ state. We see that CU_{diag} can be decomposed into two symmetric diagonal gates.

$$U_1 = \exp \left(\frac{-i}{2} \sum_{j=1}^{n+1} \phi_{j-1} \tilde{Z}_j \right) \quad (\text{S13})$$

$$U_2 = \exp \left(\frac{-i}{2} \sum_{j=1}^n (\phi_j - \phi_{j-1}) Z_j \right) \quad (\text{S14})$$

The first gate applies a symmetric diagonal operation on the ancilla and target qubits, where $\tilde{Z}_j$ is the linear combination of weight- j Pauli- Z strings supported on the $n+1$ ancilla and target qubits. This realizes the interactions terms with Z_a on the ancilla site. However, it also creates additional unwanted terms within the target qubits. Therefore, the second term simultaneously produces interactions terms with $\mathbb{1}$ on the ancilla site, and removes the additional terms from the first step, such that $CU_{\text{diag}} = U_1 U_2$. This diagonal gate can be combined with the optimized alternating ansatz to construct controlled-spin operations CU_S and CU_P . Such operations can be used for controlled probe-state preparation in Fig. 2a.

Lastly, we note that the MPO construction for the symmetric spin projectors P_s , presented in the final section, could be turned into an ancilla-mediated scalable construction for generating spin-gates using two-qubit gates. In particular, the constructed MPO corresponds to a sequence of controlled-gates from the system qubits onto an ancilla register which effectively measures the total-spin information. Then, spin-operations can be systematically generated by applying the MPO circuit between the system and ancilla, applying a diagonal gate on the ancilla, and then undoing the MPO circuit [S11].

Heuristic Simulation Time Estimates

In this section, we provide more details on the heuristics used to estimate coherent simulation times in Fig. 3. The estimates involve breaking down the error contribution into two parts: simulation error and gate error. Simulation errors arise from imperfect Hamiltonian engineering, due to mismatch between the engineered Hamiltonian H_F and the target Hamiltonian H_T . Gate errors arise from experimental imperfections which accumulate during implementation of few-qubit interactions. Our goal is to study qualitatively how our Hamiltonian engineering techniques, including both novel Hamiltonian simulation schemes and hardware-optimized multi-qubit operations, can improve the performance of a quantum simulation.


To characterize simulation errors, we look at the decay of the many-body fidelity between the state evolved under H_F and H_T . For symmetrized pulse sequences the decay should exhibit a characteristic t^2 and τ^4 dependence.

$$\int d\psi [1 - |\langle \psi | e^{iH_T t} e^{-iH_F t} | \psi \rangle|^2] = \int d\psi [t^2 (\langle \psi | (H_T - H_F)^2 | \psi \rangle - \langle \psi | (H_T - H_F) | \psi \rangle^2)] \quad (\text{S15})$$

$$= \int d\psi [t^2 \langle \psi | (\Delta^2 H_F^{(2)}) | \psi \rangle] \quad (\text{S16})$$

$$\approx (c_2 \tau^2 t)^2 \quad (\text{S17})$$

where we have only kept leading contributions in t and τ , averaged over valid initial states $|\psi\rangle$, and introduced a coefficient c_2 which captures the magnitude of the leading order error term in $H_F - H_T$. This coefficient depends on the details of the



Extended Data – FIG. S1. We determine the maximum evolution time $T^*(n)$ required to implement arbitrary symmetric diagonal gates $D(\phi)$ using the blockade interaction for clusters of $N = 2, \dots, 7$ qubits. To do this, for each evolution time ΩT we randomly select a target rotation angles ϕ , and use GrAPE to find phase profiles implementing $D(\phi)$. Above the critical times T^* (vertical dotted lines), the infidelity achieved was always below 10^{-6} . Identified values for $T^*(N)$ are also presented in the inset.

Hamiltonian simulation protocol, and can be understood as arising from the third order term in the Magnus expansion

$$H_F^{(2)} = \frac{\tau^2}{K} \sum_{k_1=1}^K \sum_{k_2=1}^{k_1} \sum_{k_3=1}^{k_2} \left(1 - \frac{\delta_{k_1, k_2} + \delta_{k_2, k_3}}{2} \right) ([H(k_1), [H(k_2), H(k_3)]] + [[H(k_1), H(k_2)], H(k_3)]). \quad (\text{S18})$$

In what follows, we characterize simulation errors via c_2 , and focus on the effect of the cycle length K on the value of c_2 . At a first glance, we observe that in (S18) there are $\binom{K}{3}$ terms with $k_1 \neq k_2 \neq k_3$, and $\binom{K}{2}$ terms with $k_1 = k_2$ or $k_2 = k_3$, and $\binom{K}{1}$ terms where $k_1 = k_2 = k_3$. In general, the nested commutators contain a lot of additional structure which will affect c_2 , including large cancellations between individual terms. For example, the terms where $k_1 = k_2 = k_3$ automatically vanish. Further, for symmetrized pulse sequences, $H(k) = H(K - k + 1)$, and many pairs of triple commutators also vanish. Finally, the fact that each $H(k)$ can be decomposed as $\sum_n H^{(n)} e^{-i\Theta_k n}$, paired with choose pulse sequences such that $\sum_k e^{-i\Theta_k n}$, implies many of the phases will interfere destructively. A more detailed analysis of the Floquet engineering would take these structures into account explicitly, but is beyond the scope of this work. Nevertheless, the structure of the third-order term suggests that c_2 scales as most as $O(K^2)$, but also may have a large $O(K)$ contribution. As such, we will utilize numerical simulations to determine a heuristic relationship between K and c_2 .

To characterize the contribution from incoherent gate errors, we assume the error is accumulated linearly in time. If we let g be the typical gate error per cycle, then the total incoherent error accumulated during simulation to time T is approximately gT/τ . We neglect coherent gate errors, which will grow as T^2 , for simplicity. Our estimates are based on using hardware-optimized operations for the Rydberg platform, so we estimate g based on the dual-driving gate time from Fig. 2b, and assume an error rate $g_0 = 10^{-3}$ per physical Rabi time $\Omega T = 2\pi$.

When combined with the Floquet error, sourced by c_2 , the total error of a simulation can be written as


$$\varepsilon = (c_2 \tau^2)^2 T^2 + \frac{gT}{\tau} \quad (\text{S19})$$

There is a tradeoff between these two error sources, which can be tuned by adjusted the step-size τ . Minimizing the total error ε for a fixed evolution time T , we see the optimal step-size is

$$\tau_{\text{opt}} = \left(\frac{g}{4Tc_2^2} \right)^{1/5} \quad (\text{S20})$$

which results in a minimum error of

$$\varepsilon_{\text{opt}} = 5 \left(\frac{c_2^2 g^4 T^6}{2^4} \right)^{1/5} \quad (\text{S21})$$



Extended Data – FIG. S2. Visualization of procedure used to fit c_2 from numerical simulations, for two interacting spins (a) and the Kagome Heisenberg model (b). In (a), we compare a Trotterization approach to one using dynamical projection and multi-qubit gates. For an ensemble of randomly chosen initial states in the symmetric subspace, we time evolve to a target t and compute the average fidelity $F(t) = |\langle \psi | U_{\text{idgal}}^\dagger U_F(t) | \psi \rangle|^2$. We further vary the stepsize τ , and observe the characteristic τ^4 scaling associated with symmetrization, which eliminates the τ^2 term. The overall coefficient c_2 is also extracted from the same fits.

or equivalently, in a maximum evolution time given a particular error tolerance

$$T_{\max} = \frac{2^{2/3}}{5^{5/6}} \frac{\epsilon^{5/6}}{(cg^2)^{1/3}} \quad (\text{S22})$$

Notice that both c_2 and g are extensive quantities, growing linearly in the number of qubits N . As such, we will typically work with intensive quantities, $\tilde{c}_2 = c_2/N$ and $\tilde{g} = g/N$, and absorb the N dependence into the target error tolerance

$$T_{\max} = \frac{2^{2/3}}{5^{5/6}} \frac{\epsilon^{5/6}/L}{(\tilde{c}_2 \tilde{g}^2)^{1/3}}. \quad (\text{S23})$$

This corresponds to selecting an error which grows (almost) extensively as well, i.e. a (nearly) constant error density.

Two interacting spin- S 's

We next describe how c_2 and g are estimated heuristically for the four models studied in Fig. 3. We start by simulating the two spin- S model with generic two-body interactions described by $H_2 = \sum_{\alpha\beta} J^{\alpha\beta} \hat{S}_1^\alpha \hat{S}_2^\beta$. In particular, we consider a mixture of Heisenberg and DM interactions, $J^{\alpha\beta} = J\delta^{\alpha\beta} + D\sum_\gamma \epsilon^{\alpha\beta\gamma}$. The c_2 coefficient is fitted numerically from noise-free dynamics, for both a two-qubit Trotter and multi-qubit dynamical projection approach. In particular, we simulate ideal and simulated Hamiltonian dynamics from random product states $|\psi\rangle$ in the symmetric subspace, and compute the error at short times t . Then, we extract the coefficient $c_2\tau^2$ by averaging over a small time window, and fit c_2 by varying τ (see Fig. S2).

In the Trotter approach, the Hamiltonian simulation protocol involves applying a sequence of $2S$ Hamiltonians, each of which act pairwise between the two clusters.

$$(H_{2,\text{Trotter}})_i = \sum_{\alpha\beta} \sum_{a=1}^{2S} \bar{J}_{ij}^{\alpha\beta} \hat{s}_{i,a}^\alpha \hat{s}_{j,a+i \bmod(2S)}^\beta \quad (\text{S24})$$

This decomposition ensures every pair of constituent qubits between the two clusters interacts exactly once during the $2S$ step-sequence. Note that we say this sequence has $K = 2S$, even though the symmetrized sequence we simulate actually has period $2K$.

In the dynamical projection approach, we select $H_{2,I} = (H_{2,\text{Trotter}})_0$ to be the first pairing of the cluster of qubits. It turns out, for this problem, a $K = 2$ sequence is sufficient to symmetrize $H_{2,I}$. The key observation, is to notice that $H_{2,I}^{(\pm 1)} = 0$, due to the specific structure of the Hamiltonian. Decomposing into all terms we get,

$$\begin{aligned}
 H_{2,I}(k) = & \sum_{a=1}^{2S} \bar{J}^{\alpha\beta} (\hat{s}_{1,a}^{\alpha(0)} \hat{s}_{2,a}^{\beta(0)} \\
 & + \hat{s}_{1,a}^{\alpha(+1)} \hat{s}_{2,a}^{\beta(0)} e^{i\Theta_k} + \hat{s}_{1,a}^{\alpha(0)} \hat{s}_{2,a}^{\beta(+1)} e^{i\Theta_k} \\
 & + \hat{s}_{1,a}^{\alpha(-1)} \hat{s}_{2,a}^{\beta(0)} e^{-i\Theta_k} + \hat{s}_{1,a}^{\alpha(0)} \hat{s}_{2,a}^{\beta(-1)} e^{-i\Theta_k} \\
 & + \hat{s}_{1,a}^{\alpha(1)} \hat{s}_{2,a}^{\beta(1)} e^{2i\Theta_k} + \hat{s}_{1,a}^{\alpha(-1)} \hat{s}_{2,a}^{\beta(-1)} e^{-2i\Theta_k}).
 \end{aligned} \tag{S25}$$

However, notice that $\hat{s}_{j,a}^{\beta(0)} = \frac{1}{2S} S_j$, and $\sum_a \hat{s}_{1,a}^{\alpha(\pm n)} = 0$ for any $n \neq 0$ due to symmetry. As such, the terms which change H_P by ± 1 vanish, leaving only the ± 2 terms. This means the cancellation condition can be satisfied using the two-step sequence $\Theta_k = (0, \pi/2)$.

The fitted values of c_2 from numerical simulations are presented in Table. II. In particular, we present an intensive version $\tilde{c}_2$ which is normalized by the total number of qubits in the system. Furthermore, note that the value of c_2 depends on the characteristic energy scale of the Hamiltonian. To handle this systematically, we scale the simulation time by a simple measure of the local Hamiltonian strength $|H|_{\text{local}}$. This is computed by simply adding the number of Pauli operators which intersect with any given qubit i weighted by the magnitude of their coefficients, and dividing by the weight of each operator. For example, for H_2 , there are $2S \times 3$ two-body Pauli operators with coefficient J , and $2S \times 6$ two-body operators with coefficient D . Therefore, $|H|_{\text{local}} = 6S(J + 2D)/2 = 9SJ$ for $D = J$.

$\tilde{c}_2$	$S = 1$	$S = 3/2$	$S = 2$	$S = 5/2$
Trotter	0.135	0.215	0.285	0.397
Dynamical Projection	0.12	0.17	0.17	0.18

TABLE II. Numerically estimated values for the normalized, intensive coefficients $\tilde{c}_2$, for two interacting large spin S 's. The Trotter error is consistent with roughly linear growth, arising the growth of the sequence length $K = 2S$. In contrast, dynamical projection appears consistent with a constant $\tilde{c}_2$, matching the constant sequence length $K = 2$.

The computation of g involves counting the total number of gate per step. In the Trotter approach, there are $2S$ two-qubit interactions. Because each interaction involves a symmetric and anti-symmetric part, it needs to be further decomposed, e.g. into a sequence of symmetric gates and single-qubit rotations. In particular, a three step sequence of symmetric interactions $h_1 = J(s_1^x s_2^x + s_1^y s_2^y + s_1^z s_2^z)$, $h_2 = D(s_1^x s_2^y + s_1^y s_2^x + s_1^z s_2^z)$, and $h_3 = D(s_1^x s_2^z + s_1^y s_2^x + s_1^z s_2^y)$ can realize the target Hamiltonian on average. Importantly, h_2 and h_3 can be brought into the same form as h_1 by applying a single-qubit rotation on the second site around the $s^x + s^y + s^z$ axis. Note that, while this introduces additional simulation error at higher-order, variational gate optimization as in the alternating ansatz could be used to identify higher-fidelity local sequences. As such, each cycle in the Trotter approach utilizes $2S \times 3$ two-qubit symmetric gates. In the dynamical projection approach, there are still $2S \times 3$ two-qubit gates per cycle, but also two additional $2S$ -qubit gates to apply H_P . However, in practice some pairs of gates applied during separate steps can be merged, such as when the interaction frame $\Theta_k = \Theta_{k+1}$ is the same. In symmetrized sequences, this happens during the 1st and K -th gates in the $2K$ cycle. As such, we multiply g by $\frac{K-1}{K}$ to account for this reduction. The intensive quantity $\tilde{g}$ is presented in Table. III.

Gate error $\tilde{g}$	$S = 1$	$S = 3/2$	$S = 2$	$S = 5/2$
Trotter	3.1	4.1	4.7	5.0
Dynamical Projection	4.1	4.2	4.0	3.9

TABLE III. Numerically estimated values for the gate error per cycle ($\times 10^{-3}$), per qubit, $\tilde{g}$. Note that despite the overhead, the projection approach becomes more efficient than Trotter due to the ability to combine gates during the 1st and K -th cycle. Further, the projection overhead decreases with $2S$ since this quantity is intensive.

Since K is constant for dynamical projection, and grows with $2S$ for Trotterization, the dynamical projection approach becomes more favorable as the spin-size grows. Combining the estimated values for c_2 and g produces the simulation times in Table IV.

Sim. Times ($\times H _{\text{local}}$)	$S = 1$	$S = 3/2$	$S = 2$	$S = 5/2$
Trotter	5.6	3.9	3.3	2.8
Dynamical Projection	4.8	4.2	4.3	4.3

TABLE IV. Numerically estimated simulation times, for an (approximately) per-qubit error tolerance of $\epsilon = 0.1$, in units of $|H|_{\text{local}}$.

Kagome Heisenberg

Next, we study the Kagome Heisenberg model. Since this model has no large-spins, there is no need for the projection approach. However we can still utilize multi-qubit gates to improve the simulation. To illustrate this, we consider two cases, one with two-qubit gates and another with three-qubit gates. In the two-qubit case, we consider a $K = 6$ protocol, where two-qubit $\hat{S}_i \cdot \hat{S}_j$ gates are applied during each step. In the three-qubit case, we develop a $K = 2$ protocol, applying three-qubit $\hat{S}^2$ evolution along upwards-facing triangles and downwards-facing triangles in alternating fashion.

We follow the same fitting procedure for c_2 with $N = 6$ and $N = 12$ size systems with periodic boundary conditions, and again estimate g using the gate-times from Fig. 2b. The $N = 12$ values are used for the estimates in Fig. 3.

N	$\tilde{c}_2$	$\tilde{g}$	T_{max}
$N = 6$, two-qubit gates	2.6	1.7	3.1
$N = 6$, three-qubit gates	0.25	1.7	6.9
$N = 12$, two-qubit gates	3.5	1.7	2.8
$N = 12$, three-qubit gates	0.36	1.7	6.1

TABLE V. Parameters used to estimate effective simulation time for the Kagome Heisenberg

We notice that the $\tilde{c}_2$ coefficients are significantly higher for the $K = 6$ protocol compared to the similar $K = 5$ spin-5/2 value. This discrepancy could be partly due to the fact that the local Hamiltonian norm is overestimated in the two spin-5/2 case, which has more frustrated (non-commuting) interaction terms. In particular, under a time rescale $t \rightarrow \lambda t$, the c_2 coefficient varies by λ^3 , making its specific value quite sensitive to the choice of units. Nevertheless, the overall trend is clear, shortening the sequence appears to accumulate simulation errors more slowly.

Spin-2 models

Finally, we use this approach to present heuristic estimates of the simulation time for a representative complex spin-model. In particular, we consider a system of spin-2's on a square lattice, with pairwise higher-order terms $(\mathbf{S}_i \cdot \mathbf{S}_j)^k$, making the interactions quite complicated.

$$H_{\text{sq},n} = \sum_{k=1}^n \sum_{\langle ij \rangle} \frac{J_n}{(2S)^{n-1}} (\mathbf{S}_i \cdot \mathbf{S}_j)^k \quad (\text{S26})$$

One of the key limitations towards realizing this model with the conventional two-qubit Trotter approach stems from the fact that each term $(\mathbf{S}_i \cdot \mathbf{S}_j)^k$ decomposes into a sum of $2k$ -qubit interactions. Further, decomposing higher-order interactions into two-qubit gates comes with large overhead, as depicted in Fig. 2c. The second is the rapid growth of the number of terms with the order of the interaction. For example, since each $S = 2$ site has four qubits, there are $\binom{4}{2}^2 = 36$ groups of four-qubit interactions coupling adjacent sites i and j in $(\mathbf{S}_i \cdot \mathbf{S}_j)^2$. Both of these problems are mitigated in the multi-qubit approach. In particular, the hardware-level optimization introduces native multi-qubit gates to directly realize higher-qubit operations. Further, dynamical projection realizes the multitude of terms using simpler operations.

We first consider the H_2 model with Heisenberg $k = 1$ and bi-quadratic $k = 2$ interactions. The Trotter version involves cycling through the 144 four-qubit interactions, between each spin and its four neighbors. Since each interaction acts on two-qubits per cluster, this can be parallelized into a cycle of period $K = 72$. In contrast, the dynamical projection approach alternates between vertical interactions, horizontal interactions, and an intra-cluster interaction (H_P). Since $n_{\text{max}} = 2$, this needs to be cycled three times, leading to $K = 6$. In order to get a heuristic estimate for c_2 , we extrapolate a linear-scaling, and simply approximate $\tilde{c}_2 \sim K$, which is within an order-of-magnitude of the estimates from simulation above.

Decomposition of the four-qubit spin interaction into two-qubit operations requires roughly 100 gates, resulting in a gate-error per qubit per step of $100g_0/4 = 25g_0$. In contrast, the large-angle four-qubit simultaneous driving gate takes time $3.74g_0$. Since there are two interaction gates for every intra-cluster gate, this results in a gate-error per qubit per step of $3.74 * 1.5g_0/4 = 1.395g_0$. Note for simplicity, we neglect the additional $\frac{K-1}{K}$ factor arising from combining gates.

Putting this together, our estimated simulation time is $T_{\max} = 0.17$ for the conventional, two-qubit Trotterization approach, and $T_{\max} = 2.7$ for the multi-qubit dynamical projection approach. We can also isolate the relative contribution from multi-qubit gates and dynamical projection. The multi-qubit Trotter approach, (using $c_2 = 77$ and $g = 1.395g_0$) is estimated to be $T_{\max} = 1.2$. The two-qubit, dynamical projection approach, (using $c_2 = 77$ and $g = 1.395g_0$) is estimated to have $T_{\max} = 0.4$.

Finally, we consider a model with up to bi-quartic $k = 4$ interactions. Here, the $k = 4$ terms act on all eight qubits between two clusters, so no dynamical projection is needed. Instead, all of the advantage comes from the efficiency of the multi-qubit approach. The trotterization period is $K = 4$, and the two-qubit decomposition estimates suggests over 10^4 gates are needed to realize the $k = 4$ interaction, so we plug in $g \sim 10^4 g_0/8 \sim 10^3 g_0$. In contrast, the eight-qubit gate times is $\Omega T/2\pi = 5.0$, producing $g = 5.0g_0/8$. Again using $c_2 = K$ heuristically, this results in estimated evolution times of 0.03 and 5.3 respectively.

MANY-BODY SPECTROSCOPY

The algorithmic framework we develop utilizes classical co-processing to extract detailed spectral information from time-evolution and snapshot measurements in a resource-efficient way. In the following sections, we discuss the following key aspects. First, we describe how measurements from quantum simulation (Fig. 4a) enable parallel measurement of exponentially many two-time correlation functions of the form $C_{O,R}(t)$. Next, we explain how from M snapshots, and for simple reference states $|S\rangle$, this information is in principle sufficient to efficiently reconstruct *any linear functional* of the time-evolved probe states $U(t)|R\rangle$. We show how this enables efficient calculation of the density-of-states, and how the convergence of the estimator can be system size independent for appropriately chosen ensembles. Then, we discuss the estimation of thermal expectation values, and finally present a poly(N) MPO construction for computing spin-projected DOS, enabling efficient classical evaluation of $D^{Ps}(\omega)$.

Estimator (31) for $C_{O,R}(t)$ from snapshots

The core quantities which enable access to spectral properties of the Hamiltonian are correlation functions of the form $C_{O,R}(t) = \langle S|O(t)R(0)|S\rangle$. For a given evolution time t and perturbation R , these can be efficiently obtained using the circuit in Fig. 4a. Further assume the system qubits are measured in a fixed basis, such as the X -basis, and O is diagonal in this basis. Then, $C_{O,R}$ can be efficiently estimated in parallel from the resulting snapshot measurements. During the i -th run of the experiment, let $\mu^{(i)} = \{x, y\}$ be the randomly selected ancilla measurement basis each with probability $1/2$, and $a^{(i)} = \{0, 1\}$, and $b^{(i)}$ be the ancilla and system measurement outcomes respectively. Then, as presented in (31), the estimator can be written as

$$\overline{C_{O,R}(t)} = \frac{1}{M} \sum_{i=1}^M 2\sigma(\mu^{(i)}, a^{(i)}) \langle b^{(i)} | O | b^{(i)} \rangle. \quad (\text{S27})$$

To study the performance of this estimator, first we show that it is unbiased, in the sense that it's expectation value matches the target correlation function. Then, we compute it's variance, which determines how quickly the estimator converges to it's expectation value.

To show the estimator is unbiased, we need to compute the probability that during any given run, the output of the simulation is μ, a, b . The conditional measurement probabilities $P(a, b|\mu)$ can be straightforwardly computed using the form of $|\psi_f\rangle$.

$$P(0, b|x) = \frac{1}{2} \langle \mathbb{1} \otimes \Pi_b \rangle_{\psi_f} + \frac{1}{2} \langle X \otimes \Pi_b \rangle_{\psi_f} \quad (\text{S28})$$

$$P(1, b|x) = \frac{1}{2} \langle \mathbb{1} \otimes \Pi_b \rangle_{\psi_f} - \frac{1}{2} \langle X \otimes \Pi_b \rangle_{\psi_f} \quad (\text{S29})$$

$$P(0, b|y) = \frac{1}{2} \langle \mathbb{1} \otimes \Pi_b \rangle_{\psi_f} + \frac{1}{2} \langle Y \otimes \Pi_b \rangle_{\psi_f} \quad (\text{S30})$$

$$P(1, b|y) = \frac{1}{2} \langle \mathbb{1} \otimes \Pi_b \rangle_{\psi_f} + \frac{1}{2} \langle Y \otimes \Pi_b \rangle_{\psi_f} \quad (\text{S31})$$

where $\Pi_b = |b\rangle\langle b|$ is the projector onto the measured state. For observables O which are diagonal in the X -basis, it can be uniquely decomposed in the operator basis spanned by the projectors Π_b ,

$$O = \sum_b \Pi_b \langle b | O | b \rangle. \quad (\text{S32})$$

Therefore, we see that the expectation value of the estimator is equal to

$$\mathbb{E}[2\sigma(\mu, a)\langle b|O|b\rangle] = \sum_a \sum_b (P(a, b|x)P(x)2\sigma(x, a) + P(a, b|y)P(y)2\sigma(y, a)) \langle b|O|b\rangle \quad (\text{S33})$$

$$= \left(\frac{1}{2}\langle \mathbb{1} \otimes \Pi_b \rangle_{\psi_f} + \frac{1}{2}\langle X \otimes O \rangle_{\psi_f} \right) - \left(\frac{1}{2}\langle \mathbb{1} \otimes \Pi_b \rangle_{\psi_f} - \frac{1}{2}\langle X \otimes O \rangle_{\psi_f} \right) \quad (\text{S34})$$

$$+ i \left(\frac{1}{2}\langle \mathbb{1} \otimes \Pi_b \rangle_{\psi_f} + \frac{1}{2}\langle Y \otimes O \rangle_{\psi_f} \right) - i \left(\frac{1}{2}\langle \mathbb{1} \otimes \Pi_b \rangle_{\psi_f} - \frac{1}{2}\langle Y \otimes O \rangle_{\psi_f} \right) \quad (\text{S35})$$

$$= \langle (X + iY) \otimes O \rangle_{\psi_f} = C_{O,R} \quad (\text{S36})$$

as desired. Note that the factor of two in the coefficient $2\sigma(\mu, a)$ accounts for the fact that each basis is only sampled half the time, cancelling the contribution from $P(\mu)$.

The variance of the estimator can be computed by performing similar calculations. Formally, the second moment of the estimator is given

$$\mathbb{E}[|2\sigma(\mu, a)\langle b|O|b\rangle|^2] = \mathbb{E}[4|\langle b|O|b\rangle|^2] \quad (\text{S37})$$

$$= 4 \sum_b \langle \Pi_b \rangle_{\psi_f} |\langle b|O|b\rangle|^2 \quad (\text{S38})$$

$$= 2 \sum_b (\langle R(t)|b\rangle \langle b|R(t)\rangle + \langle S(t)|b\rangle \langle b|S(t)\rangle) |\langle b|O|b\rangle|^2 \quad (\text{S39})$$

Here, we used the fact that $|\sigma(\mu, a)|^2 = 1$ for all choices of μ and a ; thus, the properties of the ancilla drop out of the expression. Finally, subtracting the expected value gives an expression for the variance.

Beyond diagonal operators

When the operator O is not necessarily diagonal in the measurement basis, estimating $C_{O,R}(t)$ will require varying the measurement basis of the system qubits as well. State-of-the-art shadow tomography methods could be utilized for this purpose [S84, S121]. We discuss here the simplest case, random Pauli measurements, where the joint ancilla and system measurement basis is $\mu\nu = \mu\nu_1 \dots \nu_n$. As before, μ determines whether the ancilla is measured in X or Y , while $\nu_1, \dots, \nu_n$ determines whether each system qubit is measured in X , Y , or Z . An operator O can be decomposed into a linear combination of Pauli operators, $O = \sum_P O(P)P$, where $P = P_1 \dots P_n$ is a product of Pauli operators, and $O(P) = \text{tr}(OP)$. Furthermore, define $|b^{(i)}\rangle$ as the system state measured during the i -th run, which depends on the choice of basis setting $\nu^{(i)}$ and measured bitstring. Then, the estimator $\overline{C_{O,R}(t)}$ can be constructed as

$$\overline{C_{O,R}(t)} \simeq \frac{1}{M} \sum_{i=1}^M 2\sigma(\mu^{(i)}, a^{(i)}) \sum_P 3^{|P|} O(P) \langle b^{(i)}|P|b^{(i)}\rangle \quad (\text{S40})$$

The additional factor of $3^{|P|}$ comes from inverting the probability that the measurement basis is aligned with the Pauli operator P , where $|P|$ is the number of sites non-identity sites [S84]. Indeed, if the measurement basis is not aligned, then $\langle b^{(i)}|P|b^{(i)}\rangle = 0$ instead of ± 1 , and so no information about the expectation value of P is learned.

Interpretation as shadow tomography of $|R(t)\rangle$ from single-basis measurements

Snapshot measurements enable construction of a classical representation, i.e., a classical shadow [S84], of $|R(t)\rangle = U(t)R|S\rangle$, from which linear functionals (such as wavefunction amplitudes $\langle \phi|R(t)\rangle$) can be efficiently estimated. The amplitude used for the DOS calculation $\langle R|R(t)\rangle$ is just one such example.

To illustrate this formalism, let us still again consider the case where $|S\rangle = |0\rangle^{\otimes N}$, and snapshot measurements of the system are performed in the X -basis, where all Pauli strings are enumerated as X_s . Then, $\{X_s|0\rangle^{\otimes N}, s \in \{0, \dots, 2^N - 1\}\}$ forms a complete, orthonormal basis for the Hilbert space and we can again utilize the resolution of the identity in terms of X_s acting on a classical state,

$$\mathbb{1} = \sum_s X_s |0^N\rangle \langle 0^N| X_s, \quad (\text{S41})$$

to rewrite $|R(t)\rangle$ as a linear combination correlation functions $C_{X_s,R}(t)$, which can be measured simultaneously for each s ,

$$\begin{aligned} |R(t)\rangle &= \sum_s X_s |0^N\rangle \langle 0^N | X_s |R(t)\rangle \\ &= \sum_s X_s |0^N\rangle C_{X_s,R}(t) e^{-iE_0 t}, \end{aligned} \quad (\text{S42})$$

where we additionally use the fact that $|0\rangle^{\otimes N}$ is an eigenstate of the Hamiltonian with energy E_0 and does not evolve in time. We can use the estimator of $C_{X_s,R}(t)$ from Eq. (31) to form an estimator for linear functionals of $|R(t)\rangle$, which makes Eq. (S42) analogous to a classical shadow. The classical shadow is an object in Ref. [S84], composed of snapshot measurements, which replicates the properties of a density matrix ρ as the number of measurements $M \rightarrow \infty$ tends to infinity. The version here, tends to the wavefunction $|R(t)\rangle$, and is hence estimators for wavefunction amplitudes $\langle \phi | R(t) \rangle$ can be written as

$$\overline{\langle \phi | R(t) \rangle} = \frac{e^{-iE_0 t}}{M} \sum_{i=1}^M 2\sigma(\mu^{(i)}, a^{(i)}) 2^{N/2} \langle \phi | b^{(i)} \rangle \quad (\text{S43})$$

where $|b^{(i)}\rangle$ is the i -th measured X -basis bitstring. To be able to efficiently evaluate this expression, it is important that $|\phi\rangle$ has an efficient classical representation, so that the coefficient $\langle \phi | b^{(i)} \rangle$ are computable classically. Note that, while any wavefunction amplitude can be efficiently computed.

Unlike the original versions of shadow tomography [S84], here we utilize partial knowledge of the measured state (namely the form of $|S\rangle$) in order to extract all of the unknown information (functionals of $|R(t)\rangle$) from a single-basis measurement. Further, our procedure is sensitive to global phase information in $|R(t)\rangle$, since it works by performing interferometry relative to a reference state $|S\rangle$. This phase information is crucial for estimation of spectral information via the Fourier-transformed state $|R(\omega)\rangle$. In Ref. [S47], snapshot measurements were combined with time-series data, to extract a different Hamiltonian spectrum that has peaks corresponding to energy differences $\epsilon_i - \epsilon_j$ between eigenstates. Our method, utilizing the ancilla and reference state, is capable of detecting individual energies ϵ_i and filtering their corresponding eigenstates [S81].

The operator-resolved DOS estimator

The key quantity required to evaluate the DOS estimator (9) is the calculation of D_R^A defined in (37). One way of arriving at the expression for the estimator, is by noticing that $D_R^A = \langle R | A | R(t) \rangle$, which is equivalent to selecting $|\phi\rangle = A | R \rangle$ in (S42). This estimator is also equivalent to computing $C_{O,R}$ for a non-hermitian operator O . Focusing for simplicity on the polarized reference state $|S\rangle = |0\rangle^{\otimes N} = |0^N\rangle$, the observable is

$$O := O_R^A = \sum_s \langle R | A X_s | 0^N \rangle \hat{X}_s, \quad (\text{S44})$$

and can be written as a linear combination of Pauli- X operators. We can verify that $C_{O_R^A,R}$ reproduces D_R^A up to the global time-dependent phase shift,

$$C_{O_R^A,R}(t) = \sum_s \langle R | A X_s | 0^N \rangle \langle 0^N | \hat{X}_s(t) | R \rangle = e^{iE_0 t} D_R^A(t), \quad (\text{S45})$$

where we assumed that $|0^N\rangle$ is an eigenstate of H with energy E_S .

Despite the exponentially large sum over $\{s\}$, the density-of-states estimator can often be evaluated efficiently from snapshot data, as long as the states $|R\rangle$ and $|S\rangle$ are chosen correctly. For the polarized reference state the diagonal matrix elements take a simple form,

$$\begin{aligned} \langle b | O_R^A | b \rangle &= \sum_s \langle R | A X_s | 0^N \rangle \langle b | \hat{X}_s | b \rangle \\ &= \langle R | A \sum_s (-1)^{b \cdot s} X_s | 0^N \rangle \\ &= 2^{N/2} \langle R | A | b \rangle, \end{aligned} \quad (\text{S46})$$

which is straightforward to evaluate for many choices of the probe state $|R\rangle$ and operators A . A large class of states for which this can be efficiently evaluated classically, is when $|R\rangle, |S\rangle$ are matrix product states (MPS) with $\text{poly}(N)$ bond-dimension, and A is an matrix product operator (MPO) with $\text{poly}(N)$ bond-dimension. For example, the projectors onto different total spin-sectors P_s we leverage in Fig. 4 and Fig. 5 can be written efficiently using the MPO formalism, as we discuss further below.

Convergence of the DOS estimator

To compute how quickly the DOS ($A = \mathbb{1}$) calculation will converge for polarized reference $|S\rangle = |0\rangle^{\otimes N}$ and $O_s = X_s$, we bound the variance of the estimator, and discuss the dependence on the ensemble of probe states. For simplicity, let us focus on $t = 0$, where $C_{O_R, R}(0) = 1$ by construction. The second moment is equal to

$$\mathbb{E}[\overline{C_{O_R, R}}]^2 = 2^{N+1} \sum_b (|\langle R|b\rangle|^2 + |\langle S|b\rangle|^2) |\langle R|b\rangle|^2 = 2 + 2^{N+1} \sum_b |\langle R|b\rangle|^4, \quad (\text{S47})$$

where we used $|\langle S|b\rangle|^2 = 2^{-N}$. We evaluate this expression in a few relevant limits. First, we consider states $|R\rangle$ which come from applying random single-qubit rotations of the form $\prod_i e^{-i\eta_i X_i}$, where η_i encodes the random rotation angles. For this choice, it is straightforward to show that $\langle R|b\rangle\langle b|R\rangle = 2^{-N}$ for every X -basis eigenstate $|b\rangle$. Intuitively, each qubit of $|R\rangle$ lives in the YZ -plane, and hence is orthogonal to the X -axis. Then, the second moment evaluates to

$$\mathbb{E}[\overline{C_{O_R, R}}]^2 = 2 + 2^{N+1} \sum_b |\langle R|b\rangle|^4 \rightarrow 2 + 2^{N+1} \times 2^N \times 2^{-2N} = 4, \quad (\text{S48})$$

and the variance is $\text{Var}(\overline{C_{O_R, R}}) = \mathbb{E}[\overline{C_{O_R, R}}^2] - |\mathbb{E}[\overline{C_{O_R, R}}]|^2 = 3$. More generally, states $|R\rangle$ generated by evolving under a Hamiltonian diagonal in the X basis will satisfy these results.

Next, we consider another limiting case, whereby $|R\rangle$ is an N -qubit Haar random state. The Haar-averaged second moment can be written as,

$$\mathbb{E}_{R \in \text{Haar}} \mathbb{E}[\overline{C_{O_R, R}}]^2 = 2 + 2^{N+1} \sum_b \int dR |\langle R|(|b\rangle\langle b|)|R\rangle|^2, \quad (\text{S49})$$

which is expressed via a random-state average of the $|b\rangle\langle b|$ operator. We use the formula for Haar-random averages (Eq. (1) in Ref. [S120]) with $M = |b\rangle\langle b|$,

$$\int dR |\langle R|M|R\rangle|^2 = \frac{1}{2^N(2^N + 1)} [\text{Tr}(MM^\dagger) + |\text{Tr}(M)|^2], \quad (\text{S50})$$

which results in the variance,

$$\text{Var}(\overline{C_{O_R, R}})_{\text{Haar}} = 2 + 2^{N+1} \times 2^N \times \frac{1 + 1}{2^N(2^N + 1)} - 1 = 1 + \frac{4}{1 + 2^{-N}}, \quad (\text{S51})$$

which is of order 1 for large systems. Moreover, the above formula has a product structure, which allows us to also analyze the case where the perturbation R is a product of locally Haar-random operations. Concretely, we can divide N qubits into n groups of $k = N/n$ qubits and assume that $|R\rangle$ is a 2-design within each group (e.g. is a Clifford circuit). Then, we get

$$\text{Var}(\overline{C_{O_R, R}})_{\text{local Haar}} = 2 + 2 \prod_{i=1}^n 2^k \sum_{b_i} |\langle R_i|b_i\rangle\langle b_i|R_i\rangle|^2 - 1 = 1 + 2 \left(\frac{2}{1 + 2^{-k}} \right)^n, \quad (\text{S52})$$

which reproduces the previous result for $n = 1$. Now, if we consider a fixed subsystem size k that is independent of N , the variance will grow exponentially with N . However, if the subsystem size increases with N (fixed number of partitions α), such that $k = \alpha N$, then the expression will saturate to a constant 2^α . Thus, the variance is larger than for $k = N$, but still independent of the system size. Since variance is directly related to sample complexity, it is important to avoid exponentially growing variances in practical settings. We note that maintaining low variance while considering optimized state-preparation strategies, such as MPS state preparation, is an interesting problem we leave for future work.

More general reference states

Spectroscopy from more general reference states $|S\rangle$ can also be performed, by considering different measurement strategies. One approach is to apply the inverse state-preparation unitary $S^\dagger$ and then measure in the X -basis. This enables parallel access to the 2^N operators $\{S X_s S^\dagger\}$, which by construction form an orthonormal basis when applied to the reference state $|S\rangle$. In particular, we can use this ensemble to construct an alternate resolution of the identity

$$\mathbb{1} = \sum_s S X_s S^\dagger |S\rangle\langle S| S X_s S^\dagger = S \left(\sum_s X_s |0^N\rangle\langle 0^N| X_s \right) S^\dagger = S S^\dagger, \quad (\text{S53})$$

which can be used as before to compute $D_R^A(t)$,

$$\begin{aligned} D_R^A(t) &= \sum_s \langle R | A S X_s S^\dagger | S \rangle \langle S | S X_s S^\dagger | R(t) \rangle \\ &= \sum_s \langle R | A S X_s | 0^N \rangle C_{S X_s S^\dagger, R}(t) e^{-i E_0 t}. \end{aligned} \quad (\text{S54})$$

As with (S46), this sum can be efficiently evaluated classically as long as the state-preparation unitary S has an efficient MPO description.

Randomized measurements also provide an alternate route for extracting the density-of-states, where the measurement basis can be independent of the specific reference state $|S\rangle$. Here, we consider random Pauli measurements. The set of all Pauli operators can be used to construct a resolution of the identity from the over-complete basis $\{P_r | S\rangle\}$

$$\frac{1}{2^N} \sum_r P_r | S\rangle \langle S | P_r = \mathbb{1}. \quad (\text{S55})$$

For this choice, the expression for $D_R^A(t)$ becomes

$$\langle S | A U(t) R | S \rangle = \frac{1}{2^N} \sum_r \langle S | A P_r U(t) | S \rangle \langle S | U^\dagger(t) P_r U(t) R | S \rangle \quad (\text{S56})$$

$$= e^{-i E_S t} \sum_r \langle S | A P_r | S \rangle C_{P_r, R}(t) \quad (\text{S57})$$

where E_S is the energy of the eigenstate $|S\rangle$. This quantity can be efficiently evaluated classically as well. However, note the variance of the estimator for this quantity is generically large, due to the exponential “shadow norm” for high-weight operators [S84], see also (S40). Improving sample complexity for the randomized-measurement version of many-body spectroscopy is an exciting direction for continued research.

Beyond MPS reference states with doubled Hilbert spaces

Here, we have proposed a hybrid classical-quantum approach for evaluating $D^A(\omega)$, which leverages an efficient classical representation for $|R\rangle$ and $|S\rangle$, that is suitable for near-term applications. In some cases, this requirement can be replaced by a fully quantum measurement, while preserving the sample-efficiency of the procedure, by coherently simulating two systems in parallel and performing entangled measurements between them [S39].

We illustrate such a scheme for estimation of the bare-DOS $D^{\mathbb{1}}(\omega)$. Then, we observe that the sum over Pauli operators in (S56) can be re-written as a two-copy measurement of the form

$$D_R^{\mathbb{1}}(t) = \langle S | U(t) R | S \rangle = \frac{1}{2^N} \sum_r \langle R | P_r | S \rangle \langle S | P_r U(t) R | S \rangle \quad (\text{S58})$$

$$= \frac{1}{2^N} \sum_r \langle R | \langle S | (P_r \otimes P_r) (\mathbb{1} \otimes U(t)) | S \rangle | R \rangle \quad (\text{S59})$$

This resulting correlation function has the same form as the two-time correlators measured by Fig. 4a. Further, all the observables commute $P_r \otimes P_r$, indicating that they can be written as diagonal operators in the correct measurement basis. This basis corresponds to performing pairwise Bell measurements between the two systems. To see this, notice that a single pairwise Bell measurement simultaneously measures $\mathbb{1}, XX, YY, ZZ$. Therefore, all 4^N Pauli-pairs $P_r \otimes P_r$ can be measured in parallel by factorizing into individual pairs between the two systems. These entangled measurements can be readily implemented in reconfigurable architectures, as demonstrated in Ref. [S56].

Thermal expectation values

Once we have $D^A(\omega)$, a simple way to compute thermal expectation values is by reweighting the density of states, and computing the following ratio

$$\langle A \rangle_\beta = \frac{\int e^{-\beta \omega} D^A(\omega) d\omega}{\int e^{-\beta \omega} D^{\mathbb{1}}(\omega) d\omega} \quad (\text{S60})$$

as discussed also in Ref. [S81].

We further note that this ratio is theoretically robust to uniform spectral broadening, e.g. arising from a finite maximum evolution time. To illustrate this, we can broaden the delta function in $D^A(\omega)$ to be

$$D_T^A(\omega) = \sum_n \langle n|A|n \rangle \delta_T(\omega - \epsilon_n) \quad (\text{S61})$$

where δ_T is a broadened delta function with width $\sim 1/T$. When we add the Boltzmann weight $e^{-\beta\omega}$ and integrate over frequencies, the numerator can be rewritten as

$$\int d\omega D_T^A(\omega) e^{-\beta\omega} = \int d\omega \sum_n \langle n|A|n \rangle \delta_T(\omega - \epsilon_n) e^{-\beta\omega} \quad (\text{S62})$$

$$= \sum_n \langle n|A|n \rangle \int d\omega \delta_T(\omega - \epsilon_n) e^{-\beta\omega} \quad (\text{S63})$$

$$= \sum_n \langle n|A|n \rangle e^{-\beta\epsilon_n} \underbrace{\int d\omega \delta_T(\omega) e^{-\beta\omega}}_{C_T} \quad (\text{S64})$$

$$= C_T \sum_n \langle n|A|n \rangle e^{-\beta\epsilon_n}. \quad (\text{S65})$$

Therefore, we see this is equal to the the numerator when $T \rightarrow \infty$, rescaled by a T -dependent constant C_T . The denominator is similarly rescaled by the same constant,

$$\int d\omega D_T^{\mathbb{1}}(\omega) e^{-\beta\omega} = \int d\omega \sum_n \delta_T(\omega - \epsilon_n) e^{-\beta\omega} \quad (\text{S66})$$

$$= \sum_n e^{-\beta\epsilon_n} \int d\omega \delta_T(\omega) e^{-\beta\omega} \quad (\text{S67})$$

$$= C_T \sum_n e^{-\beta\epsilon_n}. \quad (\text{S68})$$

As such, the constant C_T cancels out in the ratio, producing

$$\langle A \rangle_{\beta, T} = \frac{\sum_n \langle n|A|n \rangle e^{-\beta\epsilon_n}}{\sum_n e^{-\beta\epsilon_n}} = \langle A \rangle_{\beta, T \rightarrow \infty} \quad (\text{S69})$$

which is equal to the ideal value.

Magnetic Susceptibility

The magnetic susceptibility for the Heisenberg models studied in Fig. 4 and Fig. 5 can be computed straightforwardly by selecting $A = (S^z)^2$. To understand why this is the case, consider coupling the system to an external magnetic field $\mathbf{B}$ [S85]. Define global spin operators $\hat{\mathbf{S}} = \sum_i \hat{\mathbf{S}}_i$. Then, the Hamiltonian in the presence of a field becomes

$$H \rightarrow H + \mathbf{B} \cdot \hat{\mathbf{S}} \quad (\text{S70})$$

The density matrix at finite temperature is further given by


$$\rho_\beta = \frac{1}{Z} e^{-\beta H}, Z = \text{Tr}[e^{-\beta H}]. \quad (\text{S71})$$

The magnetization is a vector quantity measured by

$$M^\mu(\mathbf{B}, \beta) = \frac{1}{Z} \text{Tr}[S^\mu e^{-\beta H}] \quad (\text{S72})$$

and the susceptibility is a response matrix

$$\chi_\nu^\mu = \frac{\partial}{\partial B^\nu} M^\mu(\mathbf{B}, \beta) \quad (\text{S73})$$



Extended Data – FIG. S3. **Importance sampling to improve SNR.** We consider the signal-to-noise ratio associated with calculating ground-state ($T = 0$) properties of the S2H-1b model of the OEC. We consider perturbed states $|R\rangle = R|0^N\rangle$ generated by applying random single-qubit rotations (i.e. random product states), and sample from an ensemble $\mathcal{R}'$ with probabilities proportional to $e^{-\beta_{ensemble} E(R)}$. We find that there is an optimal temperature $\beta_{ensemble}$ which approximately doubles the signal-to-noise ratio compared to a fully random ensemble ($\beta_{ensemble} = 0$). Further increasing $\beta_{ensemble}$ actually causes the signal-to-noise ratio to become worse, since the ground state cannot be expressed as a product state. Initial states supporting some entanglement (e.g. MPS) should further improve the SNR.

which depends on both the temperature β and external field $\mathbf{B}$. Using the fact that $\mathbf{S}$ commutes with H in the Heisenberg models, the derivative can be simplified to

$$\chi_{\beta}^{\alpha}(\mathbf{B}, \beta) = \beta \text{Tr}[\hat{S}^{\alpha} \hat{S}^{\beta} \rho_{\beta}]. \quad (\text{S74})$$

In the main text, we focus on χ at zero-field $\mathbf{B} = \mathbf{0}$, where the rotational symmetry implies the matrix χ_{β}^{α} is proportional to δ_{β}^{α} . Thus, we focus on one component, and plot

$$\chi(T) = \chi_z^z(\mathbf{B} = \mathbf{0}, \beta) = \beta \text{Tr}[(\hat{S}^z)^2 \rho_{\beta}]. \quad (\text{S75})$$

Sampling schemes for low temperatures

As discussed above, it is important for sample efficiency that the ensemble of perturbations $\mathcal{R}$ prepare probe states that have significant overlap with target eigenstates. Classical optimization can be used to prepare variational states with larger overlap on, e.g. the ground state.

However, to accurately compute finite-temperature properties, while preferentially targetting the low-energy sector, it is crucial that we accurately compute the density of states $D^A(\omega)$, without introducing bias by restricting to a smaller subspace. Importance sampling is a method to simulate uniform sampling over all of Hilbert space, while still preferentially sampling low-energy states. The key intuition is that we can estimate the low-energy part of $D^A(\omega)$ with lower-noise, at the cost of making the high-energy part of $D^A(\omega)$ more noisy. Then, since the contribution to low-temperature properties is dominated by the low-energy part of $D^A(\omega)$, this produces an overall less noisy estimate.

Specifically, we consider an importance sampling scheme, where the probability of sampling a state $|R\rangle$ from a uniform distribution $\mathcal{R}$ is weighted by the Boltzmann factor $P_{\mathcal{R}'}(R) \propto e^{-\beta E(R)}$ where $E(R) = \langle R|H|R\rangle$ is a classically computed energy. Then, we can replace the average over $\mathcal{R}$ with a weighted average over $\mathcal{R}'$.

$$\mathbb{E}_{R \sim \mathcal{R}} 2^N \rightarrow \mathbb{E}_{R \sim \mathcal{R}'} \frac{2^N}{P_{\mathcal{R}'}(R)} \quad (\text{S76})$$

Substitutions of this form can reduce the total sample complexity, as we demonstrate for the random product state case in Extended data Fig. S3.

Efficient evaluation of spin-sector projectors

Our efficient estimator (S46) of the operator-resolved density of states involves evaluating expectation values of the form $\langle R|A|b\rangle$ where b is the measured (X -basis) product state, $|R\rangle$ is the initially prepared probe state, and A is the operator to measure. As interesting choices for A , we utilize the projector onto a total S^z sector, and the projector onto total-spin sectors P_s .

To evaluate these overlaps efficiently, we utilize the formalism of 1D tensor networks. In particular, write the projector onto an S^z sector with m spin-flips as a matrix-product operator

$$P(N = N, S^z = \frac{N}{2} - m) = \langle m|_a \prod_{i=1}^N (\pi_i^0 \mathbb{1}_a + \pi_i^1 \sigma_a^+) |0\rangle_a \quad (\text{S77})$$

where in the term $\pi_i^1 \sigma_a^+$, the first index acts on the physical site i , and the second term acts on the ancillary mode, which starts at $|0\rangle$ on the right and at the end is projected onto $\langle m|$. Intuitively, the ancilla mode counts the number of qubits in the state $|1\rangle$, and the ket $\langle m|$ only picks out strings with exactly m down spins. Thus, this forms the projector onto all states with exactly m down spin, or $S^z = \frac{N}{2} - m$. This MPO has bond dimension $\Theta(N)$. Therefore, overlaps of the form $\langle R|P_{S^z}|b\rangle$, where $|R\rangle$ and $|b\rangle$ are both finite-bond dimension MPS, can be efficiently evaluated in $O(N^2)$ time via tensor contraction.

Next, we generalize the projector construction to sectors with well-defined total-spin and z -component, S and S^z respectively. The key insight underlying the construction is to utilize three ancilla modes, which track the total S^z of the bra and ket, as well as their Hamming distance.

$$P(N = N, S = \frac{N}{2} - s, S^z = \frac{N}{2} - m) = \langle N, s, m|_a \times \prod_{i=1}^N (|0\rangle_i \langle 0|_i (\mathbb{1} \mathbb{1} \mathbb{1})_a + |0\rangle_i \langle 1|_i (\mathbb{1} \sigma^+ \sigma^+)_a + |1\rangle_i \langle 0|_i (\sigma^+ \mathbb{1} \sigma^+)_a + |1\rangle_i \langle 1|_i (\sigma^+ \sigma^+ \mathbb{1})_a) \times |000\rangle_a \quad (\text{S78})$$

The bond-dimension of this MPO scales as $O(N^3)$ in general. To understand this expression, notice that we can always write the matrix elements of the projector as $P = \sum_{s,s'} P_{s,s'} |s\rangle \langle s'|$ in the z -basis, where the coefficient $P_{s,s'} = \langle s|P|s'\rangle$. When the MPO for P is contracted with the two bistrings $\langle s| \cdot |s'\rangle$, the expression reduces to

$$P_{s,s'} = \langle N, s, m|w(s)w(s')w(s-s')\rangle_a \quad (\text{S79})$$

where $w(s)$ measures the number of down-spins (weight) of each configuration s . As in the S^z case, for this expression to be non-zero the number of down-spins in both s and s' must equal m , which imposes the condition $w(s) = w(s') = m$. However, in general, the projector also contains off-diagonal terms, where the Hamming distance $w(s-s')$ between the bra and ket state can be non-zero.

A simple example is the projector onto the triplet state,

$$P(N = 2, S = 1, S^z = 0) = \frac{1}{2} (|01\rangle \langle 01| + |01\rangle \langle 10| + |10\rangle \langle 01| + |10\rangle \langle 10|) \quad (\text{S80})$$

For this example, we see the termination vector should take the form

$$|2, 0, 1\rangle = \frac{1}{2} (|110\rangle + |112\rangle) \quad (\text{S81})$$

Note that this object is a virtual state, not a physical state, and hence does not need to be normalized. To simplify notation, we will in general write each termination vector as

$$|N, s, m\rangle_a = \sum_{d=0}^m P_{s,m,d}^N |(m)(m)(2d)\rangle_a \quad (\text{S82})$$

where $P_{s,m,d}^N$ encodes all of the relevant weights. We note that, in this expression, we use the fact that the Hamming distance between two configurations with the same S^z is always an even number $2d$, and is can be no greater than twice the number of spin-flips m . As such, we can interpret d as the number of exchanges needed to connect the two-strings s and s' .

We present an efficient ($\text{poly}(N)$ time) recursive construction for the coefficients $P_{s,m,d}^N$

To illustrate the idea, let us start with the simplest projector, onto $S = \frac{N}{2}$ and $S^z = \frac{N}{2}$, or $s = 0, m = 0$. The only state which contributes is the all-up state, so $P_{0,0,0}^N = 1$. The projectors onto all other symmetric subspaces, with $S = \frac{N}{2}, S^z = \frac{N}{2} - m$ can be computed using a ladder-like construction. For one spin-flip ($m = 1$), the projector can be written as

$$P(N = N, S = \frac{N}{2}, S^z = \frac{N}{2} - 1) = \frac{1}{\mathcal{N}} \left(\sum_i \sigma_i^+ \right) P(S = \frac{N}{2}, S^z = \frac{N}{2} - 1) \left(\sum_i \sigma_i^- \right). \quad (\text{S83})$$

By plugging in representative configurations with a single down-spin, we can compute the relevant coefficients to be

$$P_{0,1,0}^N = 1/\mathcal{N} \quad P_{0,1,1}^N = 1/\mathcal{N} \quad (\text{S84})$$

To compute the normalization factor $\mathcal{N}$, we use the fact that $P^2 = P$. In particular, we can again plug in representative configurations to match

$$\langle s | P | s' \rangle = \langle s | P^2 | s' \rangle \quad (\text{S85})$$

$$= \sum_{s''} \langle s | P | s'' \rangle \langle s'' | P | s' \rangle \quad (\text{S86})$$

Taking $|s\rangle = |s'\rangle = |10^{N-1}\rangle$ as the representative states, we could explicitly sum over the N relevant intermediate states. However, to simplify the calculation, we note the intermediate states naturally split into two groups. The first is distance-0 from s and s' , i.e. $|s''\rangle = |s\rangle$, which contributes $(P_{0,1,0}^N)^2$. The second group is distance-2 from s and s' , such as $|s''\rangle = |010^{N-2}\rangle$. This contributes $(P_{0,1,1}^N)^2$, and there are $N - 1$ such contributions. Therefore, we get the equality

$$P_{0,1,0}^N = (P_{0,1,0}^N)^2 + (N - 1)(P_{0,1,1}^N)^2 \quad (\text{S87})$$

$$1/\mathcal{N} = N/\mathcal{N}^2 \quad (\text{S88})$$

$$\mathcal{N} = N. \quad (\text{S89})$$

We could similarly select $|s\rangle$ and $|s'\rangle$ so the left-hand side is the term $P_{0,1,1}^N$ corresponding to Hamming distance 2. There, a similar calculation over intermediate paths leads to the constraint

$$P_{0,1,1}^N = 2(P_{0,1,0}^N P_{0,1,1}^N) + (N - 2)(P_{0,1,1}^N)^2 \quad (\text{S90})$$

$$\mathcal{N} = N \quad (\text{S91})$$

resulting in the same normalization factor. Indeed, the condition $P = P^2$ implies the computed value for $\mathcal{N}$ must be the same for all non-zero coefficients appearing in the left-hand side. The key intuition underlying these calculations, is that the valid paths can be efficiently partitioned into groups, and the coefficients can be efficiently computed via combinatorics.

The same intuition can be generalized to a recursive algorithm for computing $P_{s,m,d}$ from $P_{s,m-1,d'}$, which allows us to compute the full ladder of states with $m = 0, \dots, N$ from the root cases. First, we compute the unnormalized coefficients using the formula

$$P'_{s,m,d} = d^2 P_{s,m-1,d-1} + (m-d)(2d+1)P_{s,m-1,d} + (m-d)(m-d-1)P_{s,m-1,d+1} \quad (\text{S92})$$

Where we assume $P_{s,m,d} = 0$ when $d > m$. To understand the coefficient associated with each term, we consider a pair of representative states $|s\rangle = |1^{m-d}1^d0^{N-m-d}\rangle$ and $|s'\rangle = |1^{m-d}0^d1^{N-m-d}\rangle$, which have m down-spins and Hamming distance $2d$. For convenience, we label the four groups of size $m-d, d, d$, and $N-m-d$ as A, B, C , and D respectively. The formula comes from the following case-by-case breakdown:

1. The term $|s\rangle\langle s'|$ can come from a pair with smaller Hamming distance $2(d-1)$ by flipping one spin from B in s and from C in s' . There are d^2 such choices.
2. The term $|s\rangle\langle s'|$ can come from a pair with the same Hamming distance $2d$, by flipping one spin from A in s and one spin from C in s' , vice versa, or by flipping the same spin from A for s and s' . There are $2(m-d)d + (m-d) = (m-d)(2d+1)$ such choices.
3. The term $|s\rangle\langle s'|$ can come from a pair with larger Hamming distance $2(d+1)$ by flipping two different spins from A in s and s' . There are $(m-d)(m-d-1)$ of these pairs.

Then, we need to compute the normalization factor. In order to do this, consider the same pair of representative states, and enumerate over all intermediate states s'' . In particular, we can efficiently label intermediate states by counting the number of down-spins in each of the four groups, m_A, m_B, m_C, m_D , where the total number of down-spins satisfies $m = m_A + m_B + m_C + m_D$. For each intermediate configuration, the distance from s to s'' and s' to s'' is given by

$$d(s, s'') = (m - d - m_A) + (d - m_B) + m_C + m_D \quad (\text{S93})$$

$$d(s', s'') = (m - d - m_A) + m_B + (d - m_C) + m_D \quad (\text{S94})$$

The number of intermediate states s'' associated with each label is

$$Z(m_A, m_B, m_C, m_D) = \binom{m-d}{m_A} \binom{d}{m_B} \binom{d}{m_C} \binom{N-m-d}{m_D}. \quad (\text{S95})$$

Taken together, by enumerating over valid assignments for the number of down-spins in each group, we can compute the normalization factor

$$\mathcal{N} = \frac{\sum_{m_A+m_B+m_C+m_D=m} Z(m_A, m_B, m_C, m_D) P'_{s,m,d(s,s'')} P'_{s',m,d(s',s'')}}{P'_{s,m,d(s,s')}} \quad (\text{S96})$$

and the normalized projector coefficients $P_{s,m,d} = P'_{s,m,d} / \mathcal{N}$.

Next, we discuss projectors onto states where the total-spin is smaller than $N/2$. To illustrate this, we start by computing $P(S = \frac{N}{2} - 1, S^z = \frac{N}{2} - 1)$. To do this, we see that there are only two total-spin sectors which contribute to $S^z = \frac{N}{2} - 1$, $s = 0$ and $s = 1$. Therefore, we know that

$$P(N = N, S = \frac{N}{2} - 1, S^z = \frac{N}{2} - 1) = P(N = N, S^z = \frac{N}{2} - 1) - P(N = N, S = \frac{N}{2}, S^z = \frac{N}{2} - 1) \quad (\text{S97})$$

The termination vector associated with $S^z = \frac{N}{2} - m$ sector is $|(m)(m)(0)\rangle$. Therefore, the coefficients of the projector can be written as $P_{1,1,d}^N = \delta_{d,0} - P_{0,1,d}^N$. As an example, consider the $N = 2$ singlet state, for which this formula implies $P_{1,1,0}^2 = 1/2$ and $P_{1,1,2}^2 = -1/2$. This leads to the familiar expression

$$P(N = 1, S = 0, S^z = 0) = \frac{1}{2} (|01\rangle\langle 01| - |10\rangle\langle 01| - |01\rangle\langle 10| + |10\rangle\langle 10|). \quad (\text{S98})$$

More generally, to construct each of the terms in the root projector with $s = m$, we orthogonalize w.r.t all higher-spin sectors.

$$P_{m,m,d} = \delta_{d,0} - \sum_{0 \leq s < m} P_{s,m,d} \quad (\text{S99})$$

Here, each term with $s < m$ can be computed from root terms where $s = m$ using the recursive ladder approach described above.