Analyzing Many-Body Localization with a Quantum Computer

Bela Bauer[1] and Chetan Nayak[1,2]

[1]*Station Q, Microsoft Research, Santa Barbara, California 93106, USA*
[2]*Physics Department, University of California, Santa Barbara, California 93106, USA*
(Received 16 July 2014; published 3 November 2014)

Many-body localization, the persistence against electron-electron interactions of the localization of states with nonzero excitation energy density, poses a challenge to current methods of theoretical and numerical analyses. Numerical simulations have so far been limited to a small number of sites, making it difficult to obtain reliable statements about the thermodynamic limit. In this paper, we explore the ways in which a relatively small quantum computer could be leveraged to study many-body localization. We show that, in addition to studying time evolution, a quantum computer can, in polynomial time, obtain eigenstates at arbitrary energies to sufficient accuracy that localization can be observed. The limitations of quantum measurement, which preclude the possibility of directly obtaining the entanglement entropy, make it difficult to apply some of the definitions of many-body localization used in the recent literature. We discuss alternative tests of localization that can be implemented on a quantum computer.

I. INTRODUCTION

Since the seminal contributions of Gornyi *et al.* [1] and Basko *et al.* [2,3], the question of whether Anderson localization can persist against interactions at nonzero excitation energy density has been revisited. The possible existence of such a phenomenon, dubbed many-body localization (MBL), is closely intertwined with other open fundamental questions in quantum statistical mechanics: whether the eigenstate thermalization hypothesis [4,5] holds for generic quantum systems and whether isolated quantum systems can equilibrate [6]. In studying these questions, significant theoretical and numerical evidence [7–23] has been assembled indicating that a many-body localized phase exists; for a recent review, see Ref. [24]. More recently, the existence of a many-body localized phase has been proven rigorously in a class of spin chains [25].

However, the numerical simulation of putative many-body localized systems has remained extremely challenging: While these states have low entanglement and should be well approximated by tensor-network states [21], the known methods to find such states are best suited for finding ground states, and they become inefficient when targeting states at generic energies in the spectrum, where the gaps to nearby states are exponentially small in the system size. Another approach has been to study the time evolution of such systems from easily prepared initial states, which is limited to small systems by the unbounded growth of entanglement [11]. This has limited accurate computations of the properties of many-body localization to systems of approximately 20 sites or smaller.

In this paper, we ask how a quantum computer of moderate size could be used to break through this barrier. Quantum computers are naturally suited to simulating the dynamics of quantum systems; indeed, the idea was first conceived in this context [26]. However, the devil is in the details: On a "digital" quantum computer operating within the gate model [27], the evolution of a quantum system is mapped to a series of one- and two-qubit gates chosen from a finite, but computationally universal, set. This mapping can induce enormous overhead in practice. Moreover, one does not have access to the full resulting quantum wave function but, rather, to information that can be extracted from projective measurements of individual qubits, although one has the freedom to choose the basis in which these measurements are done. Recently, it has become evident that while there are quantum algorithms that have exponential speedup over their classical counterparts, many of these algorithms have such an advantage only in an asymptotic regime that is unlikely to be reached on a quantum computer in the foreseeable future, even without taking into account additional overhead necessitated by quantum error correction. Consider two examples: Shor's celebrated algorithm for factoring [28,29] and the simulation of quantum chemistry [30]. In the first case, a quantum computer would require thousands of qubits to factor a number that cannot be factored using the most efficient classical algorithms. In the second case, a moderate number of qubits is required, but the mapping of the unitary evolution into the gate model requires the coherent execution of a very large number of gates [31–33].

In this paper, we will discuss how the power of a quantum computer can be brought to bear on understanding the phenomenon of many-body localization. We demonstrate that even a relatively small quantum computer will have significant advantages over a classical computer in analyzing Hamiltonians that may exhibit many-body localization.

In doing so, we address another important question: Are the hypothesized properties of MBL eigenstates observable in experiments? An exact eigenstate of the Hamiltonian with nonzero excitation energy density cannot be prepared efficiently—i.e., in a time that scales only polynomially with system size—since the time required depends inversely on the required energy resolution, and the energy-level spacing in the middle of the spectrum is exponentially small in the system size [34]. Even though our main interest in this paper is *not* analog *simulation* but, rather, quantum *computation* of the properties of MBL states with a general-purpose quantum computer, it is useful to momentarily regard a quantum computer as an experimental system, albeit a very idealized one. Then, our results demonstrate that the characteristic properties of exact energy eigenstates can, indeed, be observed in approximate eigenstates that can be prepared in polynomial time. Hence, these properties are likely also observable in other, less idealized, experimental setups.

In Sec. II, we first review briefly how the time evolution of a quantum system is mapped into the gate model using a Trotter-Suzuki decomposition, since this is a basic building block for all that follows. We show that typical features of many-body localized Hamiltonians, such as short-ranged interactions and on-site disorder, make this decomposition feasible for system sizes of interest. A quantum computer can thereby determine the time evolution of an easily prepared initial state such as a random product state. This may be viewed as a computation of the global quench dynamics of the system where the system is, at least initially, very far away from equilibrium. Such a computation would allow us to probe the equilibration and thermalization properties of the system, which can reflect its localization properties.

We then explore how random energy eigenstates of the system can be prepared with sufficient accuracy to observe signatures of many-body localization with a polynomial-depth quantum circuit. This is done using the quantum phase estimation algorithm. However, unlike in many applications in which one is interested in finding ground states, a generic eigenstate that results from quantum phase estimation is relevant to the study of many-body localization. This key step of preparing eigenstates greatly enhances the usefulness of a quantum computer since it opens up the study of situations other than global quenches and, in particular, gives access to the dynamical response of the system to weak perturbations. Examples include transport measurements or local quenches in systems prepared at

fixed energy. This may also allow more quantitative connections to experiments, such as states of ultracold atoms in optical traps.

However, preparing a state—either through quench dynamics or through quantum phase estimation—is only half the battle. We are now faced with extracting information from this state. In contrast to classical simulations, we cannot simply examine the wave function directly and deduce all of its properties. In particular, measuring the entanglement entropy in the state is difficult, if not impossible, within the constraints of our setup. Instead, we are limited to performing unitary operations on the state and then performing local, projective measurements, thereby obtaining a string of zeros and ones—as many classical bits of information as measured qubits. The question thus arises as to how we can characterize the eigenstates. In Sec. III, we discuss in more detail the limitations of the measurement process and propose scenarios for how the eigenstates can be characterized, where we focus mostly on measurements of transport properties at finite energy densities. The setups we consider are *local* quench dynamics, as well as probing the response of the system to a "tilt." Finally, we analyze the possible effects of errors, such as decoherence and discretization errors, on the computation of MBL states and their properties.

II. STATE PREPARATION

A. Models, representation, and time evolution

The models that we focus on here are (1) a spinless-fermion model with nearest-neighbor hopping in one dimension, nearest-neighbor interactions, and on-site disorder, and (2) an *XXZ* spin chain with a random Zeeman field in the *z* direction. The first model has the Hamiltonian

$$H_\mathrm{f} = -t\sum_{i=1}^{L-1}(c_i^\dagger c_{i+1} + c_{i+1}^\dagger c_i) + \sum_{i=1}^{L}w_i n_i + V\sum_{i=1}^{L-1}n_i n_{i+1}, \tag{1}$$

where $c_i^\dagger$ creates a spinless fermion on site i, and $n_i = c_i^\dagger c_i$. The w_i are uniformly chosen from $w_i \in [-W, W]$. The second model has

$$H_\mathrm{s} = -J_\perp \sum_{i=1}^{L-1}(S_i^x S_{i+1}^x + S_i^y S_{i+1}^y)$$
$$+ \sum_{i=1}^{L}w_i S_i^z/2 + J_z \sum_{i=1}^{L-1}S_i^z S_{i+1}^z. \tag{2}$$

For open boundary conditions, these models have the same spectrum for $J_\perp = t$ and $V = 2J_z$, while for closed (i.e., periodic or antiperiodic) boundary conditions, more care must be taken when relating them to each other through a Jordan-Wigner transformation.

We would like to compute time evolution due to these Hamiltonians on a general-purpose quantum computer operating under the circuit model. Therefore, the unitary evolution $U = \exp(-iTH)$ must be mapped to a series of gates chosen from a given set of available gates. We will call this procedure *compiling* below, in analogy to the well-known classical procedure of compiling a program in a high-level programming language into the assembly code, i.e., machine language, of the target hardware platform. While many approaches to achieve this are known (for some recent improvements, see, e.g., Refs. [35–37]), by far the most widely used is the Trotter-Suzuki decomposition [38,39]. First, the time evolution is broken up into a series of time steps $\delta t = T/N$:

$$U = e^{-i\delta t H} \ldots e^{-i\delta t H}. \qquad (3)$$

Then, we write $H = \sum_{i=1}^{m} H_i$, where the H_i are chosen such that $U_i(\delta t) = \exp(-i\delta t H_i)$ can be compiled into a series of gates exactly. Then, if we use a first-order Trotter-Suzuki decomposition, we write

$$e^{-i\delta t H} = \left(\prod_{j=1}^{m} e^{-i\delta t H_j} \right) + O(\delta t^2). \qquad (4)$$

This decomposition is only accurate to order δt^2, but a more elaborate decomposition can be found that is accurate to any given order in δt. However, the number of terms in the decomposition grows quickly with the order. Therefore, the optimal order depends on the desired accuracy ϵ (in trace-norm distance on the final state), the total time T, and the norm of the Hamiltonian operator. Reference [40] gives an upper bound on the total number $N_{\exp}$ of separate exponentials $U_i(\delta t) = \exp(-i\delta t H_i)$ that must be executed to achieve a given accuracy in trace norm distance, as well as an estimate for the optimal order of Trotter decomposition. In order to determine the number of elementary gates required to perform the evolution U, we multiply $N_{\exp}$ by the number of elementary gates needed to perform each of the $U_i(\delta t)$.

In the models that are relevant to many-body localization, both kinetic and interaction terms are generally local. Consider, for example, the Hamiltonian given in Eq. (1). It can be expressed as a sum of three noncommuting terms:

$$H_1 = -t \sum_{i=1}^{(L-2)/2} (c_{2i}^{\dagger} c_{2i+1} + c_{2i+1}^{\dagger} c_{2i})$$

$$H_2 = -t \sum_{i=1}^{L/2} (c_{2i}^{\dagger} c_{2i-1} + c_{2i-1}^{\dagger} c_{2i})$$

$$H_3 = \sum_{i=1}^{L} w_i n_i + V \sum_{i=1}^{L-1} n_i n_{i+1}, \qquad (5)$$

Here, we have taken L to be even. More generally, the number of noncommuting terms is $m = 1 + z$ on a regular lattice in d dimensions with coordination number z, where 1 accounts for the diagonal terms (interaction and on-site potential), and z accounts for the kinetic terms. Crucially, this is independent of the size of the system; however, $\|H\|$ is extensive [41], and thus the overall scaling is slightly faster than $L \cdot T$.

The individual terms in the Trotter expansion of U still need to be expressed in the available gate set. At this stage, the fermionic (1) and spin (2) Hamiltonians are equivalent, so we use terminology appropriate to the latter. We write $H_i = \sum_j h_i^j$, where h_i^j is a product of Pauli matrices, and use

$$\exp(-i\delta t H_i) = \prod_j \exp(-i\delta t h_i^j). \qquad (6)$$

These gates can be transformed into a basic gate set (see Refs. [31,42] for details). We find the following parallel gate counts, i.e., assuming that gates that operate on different qubits can be executed simultaneously: (i) For H_1 and H_2, we need 10 gates each. (ii) For H_3, we again need $1 + z$ gates, where z is the coordination number of the lattice. If the original Hamiltonian is fermionic then, for dimensions $d > 1$, there will be additional Jordan-Wigner strings, but their overhead can be greatly reduced [32,43,44].

For how much time T must we evolve the system to observe physical manifestations of many-body localization? It has been shown that within the scenario of a global quench, certain properties, such as the saturation of entanglement entropy to a volume law [11], can be observed only after time that is exponentially large in the system size. This makes it experimentally unfeasible to observe these properties and ultimately renders them unphysical. A more appealing scenario may be that of a local quench [22]. In the MBL phase, the perturbation only propagates a finite distance, and the long-time behavior is observed after a time that does not scale with system size.

As the numerical studies described in the following sections show, the Trotter error need not be kept extremely small for the purpose of detecting MBL physics; instead, other sources of error, such as limitations on the time to which quantum phase estimation can be run, are more relevant. If, however, a quench or transport scenario requires very low error bounds on the unitary evolution, the methods recently put forward in Ref. [37] may be favorable over a Trotter decomposition.

B. Quantum phase estimation

In order to make our discussion self-contained, we briefly review some essential features of the quantum phase estimation algorithm; see also Fig. 1. More detailed pedagogical discussions can be found in, for instance,
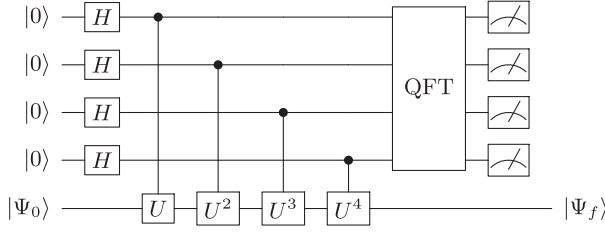
FIG. 1. Overview of the quantum phase estimation algorithm discussed in the main text for $k = 4$ ancilla qubits. The lowest line in the figure indicates the N qubits used for the physical system. Here, QFT denotes the quantum Fourier transformation. After the measurements, the readout on the k ancilla qubits contains an estimate for the energy, whereas the N qubits for the physical system contain the final state $|\Psi_f\rangle$ of Eq. (12), which is an approximation to an eigenstate and can be processed further to obtain measurements on the physical system.

Refs. [45–47]. Let us suppose that we would like to find an eigenvalue and corresponding eigenvector of a unitary operator U acting on a Hilbert space of dimension 2^N. For us, the unitary operator will be the exponential of a Hamiltonian, $U = e^{-iTH}$, that we wish to test for many-body localization. To perform quantum phase estimation, we consider a system of $N + k$ qubits, where we refer to the first N qubits as the data qubits, on which the operator U acts, and the next k qubits as ancillas. We will assume that we can perform Hadamard gates, controlled-U gates, and arbitrary controlled-phase gates. Controlled-U^k gates can be implemented by applying the controlled-U gate k times in succession. Suppose that our ancillas are all initially in the state $|0\rangle$ and the N data qubits are in an arbitrary initial state $|\psi_0\rangle$, chosen at random. Then, we perform Hadamard gates on each of the ancillas, thereby putting each in the state $(|0\rangle + |1\rangle)/\sqrt{2}$. We then act with a controlled-U gate in which the first ancilla is the control and the N data qubits are the target on which the unitary operator acts when the ancilla is in the state $|1\rangle$. We act with a controlled-U^2 gate in which the second ancilla is the control and the data qubits are the target. We continue similarly with each ancilla so that the jth ancilla is the control for a controlled $U^{2^{j-1}}$. The resulting state is

$$\sum_{\{i_n=0,1\}} U^{i_1+2i_2+\cdots+2^{k-1}i_k}|\psi_0\rangle \otimes |i_1, i_2, \ldots, i_k\rangle.$$

If T is the integer whose binary expansion is $i_1 i_2 \ldots i_k$, then this can be written in the form

$$\sum_{T=0}^{2^k-1} U^T|\psi_0\rangle \otimes |T\rangle. \tag{7}$$

Expanding the initial state of the data qubits in terms of the eigenstates of U, $|\psi_0\rangle = \sum_n c_n|n\rangle$, where $|n\rangle$ has eigenvalue e^{iE_n}, we can rewrite Eq. (7) in the form

$$\sum_n \left[c_n|n\rangle \otimes \left(\sum_{T=0}^{2^k-1} e^{iE_nT}|T\rangle \right) \right]. \tag{8}$$

We now apply the (inverse) quantum Fourier transform on the ancillas, which acts on a basis vector according to

$$|T\rangle \rightarrow \sum_{J=0}^{2^k-1} e^{-2\pi iT\frac{J}{2^k}}|J\rangle. \tag{9}$$

This results in the state

$$\sum_n \left[c_n|n\rangle \otimes \left(\sum_{J=0}^{2^k-1} g\left(E_n - \frac{2\pi J}{2^k} \right)|J\rangle \right) \right], \tag{10}$$

where $g(x) = (1 - e^{-i2^k x})/(1 - e^{-ix})$. The function $g(x)$ is peaked around $x = 0$. If we increase k so that $2^k E_n/2\pi$ approaches an integer, then it becomes more strongly peaked. If $E_n/2\pi$ is a k-bit binary number modulo integers, then $2^k E_n/2\pi$ is an integer and $g(x) = \delta_{2^k x,0}$. Then, we have

$$\sum_n [c_n|n\rangle \otimes |2^k E_n/2\pi\rangle]. \tag{11}$$

By measuring the ancilla, we obtain $2^k E_n/2\pi$ with probability $|c_n|^2$ and the data qubits are left in the state $|n\rangle$. While the eigenvalue E_n is the primary goal for applications to period finding, our main goal here is to obtain the state $|\psi_n\rangle$. Moreover, one is often interested in finding the ground state of a Hamiltonian and, therefore, needs to choose an initial state $|\psi_0\rangle$ with high overlap with the ground state, so that $|c_0|^2$ is not too small. In the present application to many-body localization, however, we are interested in generic states, so we can take a random initial state.

If $2^k E_n/2\pi$ is not an integer, then when we measure the ancillas, we obtain an approximate eigenvalue E_{approx} that is $2\pi/2^k$ times a k-bit integer. The data qubits are in the state

$$|\Psi_f\rangle = \sum_n c_n g(E_n - E_{\text{approx}})|n\rangle, \tag{12}$$

which is not an energy eigenstate but has amplitude that is sharply peaked at eigenstates that are near E_{approx}. We can make it more sharply peaked, thereby obtaining an eigenstate to within any desired accuracy, by increasing the number of ancillas.

In practice, it may be favorable to use an iterative quantum phase estimation (IQPE) algorithm, as described in Ref. [48], which effectively performs the same calculation as outlined above but requires only one ancilla qubit. This is particularly useful in the context of classical simulation of the quantum computer for validation

purposes, as the classical simulation time scales exponentially in the number of qubits, unless an approximate method is used.

It remains to be confirmed that (i) we obtain all states with sufficient probability, and (ii) we can prepare these states to sufficient accuracy to observe signatures of many-body localization with a total computation time that scales, at most, polynomially in the system size L even near the middle of the spectrum, where gaps to adjacent states are exponentially small in L.

Other approaches of obtaining eigenstates on a quantum computer seem possible: For example, one could attempt to adiabatically cool towards the ground state of $A = (H - \lambda \mathbb{1})^2$, for some $\lambda \in [-\|H\|, +\|H\|]$, or adiabatically move from an eigenstate of the diagonal part of the Hamiltonian, which can easily be prepared, to an eigenstate of the full Hamiltonian. However, these approaches have major drawbacks: In the first approach, the evolution must be performed under a nonlocal Hamiltonian. Furthermore, in both cases, the accuracy depends on whether the adiabatic evolution is performed slowly enough compared to the relevant energy scale, which is hard to control.

C. Gate counts

We now test the procedure outlined above in a numerical simulation. Our goals are (i) to confirm that we sample from the correct distribution of eigenstates, (ii) to determine the number of gates that need to be executed to obtain an eigenstate with a given energy standard deviation, and (iii) to confirm that we can obtain eigenstates to sufficient accuracy to be able to probe signatures of many-body localization. The last point will be deferred to the next section.

In our numerical simulations, we study the Hamiltonian (1) on an open chain of L sites. We perform IQPE on $U = \exp(-iHT)$. To keep eigenvalues from wrapping around the unit circle, we need to ensure that $T\|H\| < 1$ and thus set $T = [L(2 + V + W)]^{-1} < \|H\|^{-1}$. For small systems, we perform the unitary evolution exactly; for larger systems and to assess the effect of Trotter errors, we perform a first-order Trotter decomposition. We find that for the system sizes and evolution times considered here, the Trotter error is not very significant; for larger systems, a higher-order Trotter decomposition may be favorable. In this setup, the number of gates (assuming parallel execution of nonoverlapping gates) necessary to evolve the system to a time τ is given by

$$N_g = 23\tau/\delta t. \qquad (13)$$

Here, the number 23 is how many gates are needed to execute $\exp(-i\delta t H_1)\exp(-i\delta t H_2)\exp(-i\delta t H_3)$ [see Eq. (5)].

To obtain k bits of the desired eigenvalue, we need to evolve for a total time

$$T_{\text{tot}} = T \sum_{i=0}^{k} 2^i = T(2^{k+1} - 1). \qquad (14)$$

This total time effectively determines the absolute accuracy as well as computation cost of IQPE; the same accuracy can, in principle, be achieved by reducing T and, accordingly, increasing k, or vice versa. We therefore plot all results against T_{tot}.

We first confirm that we obtain states with the correct probability when starting from random initial product states. To this end, we compare the density of states $\rho(E)$ obtained using the IQPE procedure outlined above to that obtained from an exact, full diagonalization of the same Hamiltonians. We consider at least 100 disorder realizations and average over a total of 10000 states. Our results are shown in Fig. 2. Clearly, the agreement between the two approaches is excellent both in the delocalized ($W = 1$) and the localized ($W = 8$) regime.

As a simple measure of how accurately we can prepare eigenstates, we calculate the energy standard deviation

$$\Delta E = \langle \sqrt{\langle H^2 \rangle - \langle H \rangle^2} \rangle, \qquad (15)$$

where the outer average is over output states from different runs of IQPE for different disorder realizations and initial states. A naive expectation is that $\Delta E/L \cdot T_{\text{tot}} \sim 1$. To check this, we perform numerical simulations, again averaging over 10000 states for each choice of T_{tot}, W, and L. In these simulations, we perform the time evolution exactly to separate out the effect of Trotter errors. This limits the system size we can study to $L = 12$ because we exponentiate the Hamiltonian exactly. As shown in Fig. 3, we find $\Delta E/L = cT_{\text{tot}}^{-\alpha}$, where c and α are fit parameters. The best fit is obtained for $\alpha \approx 0.8$, which deviates slightly from the naive expectation of $\alpha = 1$. We observe that all curves for different L but equal W collapse, indicating that
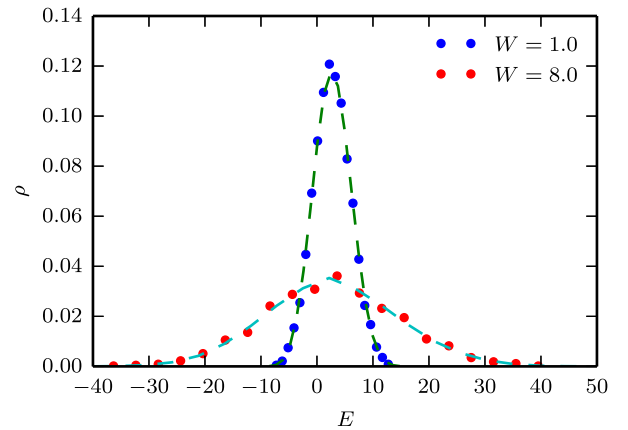

FIG. 2. Dashed lines: Density of states $\rho(E)$ obtained using full diagonalization. Points: ρ obtained using IQPE for $k = 16$ bits. All results are for $L = 12$.
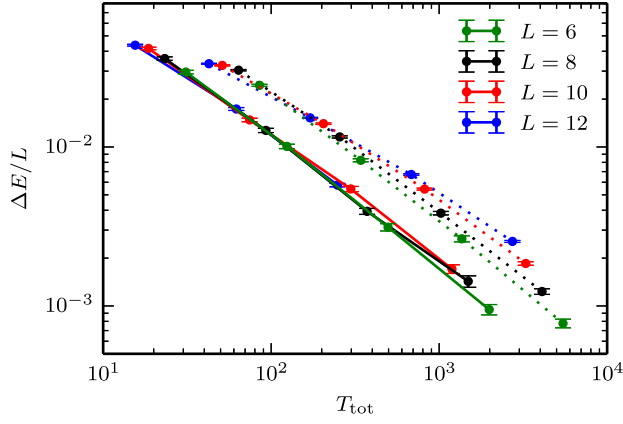
FIG. 3. Energy standard deviation density $\Delta E/L$. Dotted lines are $W = 1$; solid lines are $W = 8$. Results have been obtained using exact time evolution to exclude any Trotter errors.

c depends only on W. We observe that in the localized regime, $\Delta E/L$ of the final states is lower than in the delocalized regime.

At this point, we can also analyze Trotter errors and, using Eq. (13), obtain the gate count necessary to obtain a final state with some fixed $\Delta E/L$. For this analysis, we restrict ourselves to a first-order Trotter decomposition. In Fig. 4, we show $\Delta E/L$ against the gate count for different choices of the Trotter time step δt. We observe that the time step has to be decreased roughly as $\delta t \sim \epsilon$, where ϵ is the desired $\Delta E/L$. For example, to achieve $\Delta E/L < \epsilon = 0.01$, a time step $\delta t = 0.025$ seems necessary; to achieve $\epsilon = 0.005$, $\delta t < 0.0147$ is necessary.

We note that the total gate counts shown here are on the order of a few million and thus much more realistic than the gate counts obtained for quantum chemistry in Ref. [31]. Assuming a logical gate rate of 1 MHz, the preparation of an approximate eigenstate would require a coherence time on the order of 1 second.
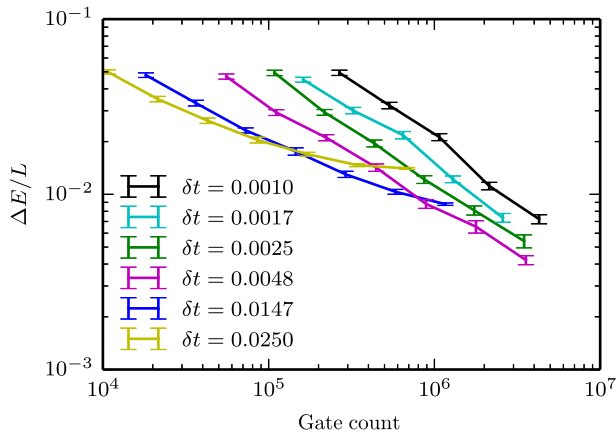

FIG. 4. Trotter errors for $L = 16$, averaging over 1000 instances. Here, we use a first-order Trotter decomposition.

D. Observation of MBL

Having established the gate counts required to reach a certain $\Delta E/L$, the question arises as to what $\Delta E/L$ must be achieved in order to observe signatures of many-body localization. Naively, one might expect that ΔE must be small compared to the mean level spacing δE. As the latter is exponentially small in the system size, this would imply an exponential scaling. In the following, however, we will argue that, in the many-body localized phase, there are states that can be prepared in polynomial (in system size) time that display key signatures of many-body localization. A similar conclusion was reached from a very different perspective in Refs. [49,50]. As a test of the localization properties of the final states obtained in our quantum algorithm, we use the entanglement entropy. The presence of an area law for the entanglement entropy has been established previously as a good indicator of many-body localization in exact eigenstates [21].

It is intuitively clear that even in a many-body localized phase, where exact eigenstates obey an area law, approximate eigenstates with $\Delta E/L > \delta E$, but small compared to other scales in the problem, may display very different entanglement properties. A superposition with random coefficients of exponentially (in the system size) many eigenstates in a given, small energy window will likely have volume-law entanglement. On the other hand, a superposition of a few eigenstates that are far away in energy can have the same $\Delta E/L$ but may still display an area law.

To explore this quantitatively in the setup described above, we apply IQPE to different initial product states. In Fig. 5, we show a comparison between (i) "Z" states that are initially polarized in the Z basis and (ii) "ZX" states where initially half of the spins are polarized in the Z and
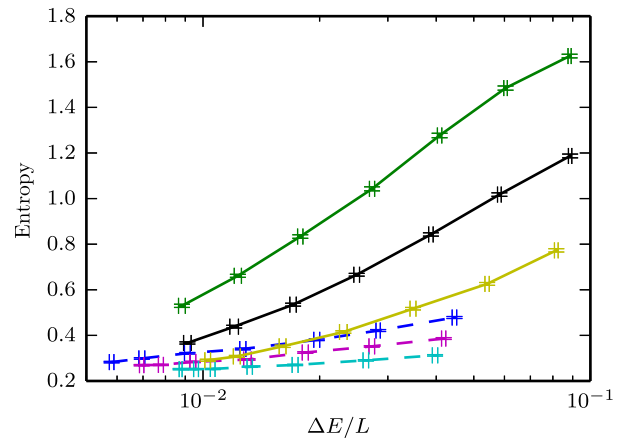

FIG. 5. The entanglement entropy at the center of the system as a function of $\Delta E/L$ for various system sizes and initial states used in the IQPE algorithm. Dashed lines show data for initial product states in the S^z basis, while solid lines show data for initial states with a mix of S^z and S^x bases, as explained in the text. System sizes are, from top to bottom, $L = 12$, 10, and 8.

the other half are in the X basis (i.e., an equal-weight superposition of $2^{L/2}$ product states in the Z basis). In Fig. 5, we observe that for the same $\Delta E/L$, the states obtained when starting from ZX initial states have drastically larger entanglement entropy.

The reason for this can be understood in the picture of local conserved constants of motion put forward for MBL states in Refs. [14,16]: The eigenstates of the effective Hamiltonian proposed there are simply product states in some fixed basis. Flipping a single spin in this basis will incur a large energy penalty; however, by flipping many spins, one can obtain another product state that is exponentially close in energy to the original state but that can be locally distinguished from the first state almost anywhere simply by measuring in the preferred basis. In the setup we consider, the local constants of motion are likely to be close to the physical σ^z operators, as the disorder is diagonal in this basis. The initial states polarized in the Z basis thus differ from exact eigenstates only by *local* fluctuations and thus have overlap with exact eigenstates that are far apart in energy. In this case, IQPE is successful at isolating one or a few eigenstates with high accuracy. The ZX states, on the other hand, can be thought of as superpositions of Z states that differ in $L/2$ spins and thus can have overlap with eigenstates that are close in energy. IQPE is thus less efficient at separating these states, and for a given energy standard deviation, the final state is a superposition of many nearby eigenstates.

Concentrating on the Z initial states which, as argued above, give a better approximation to the exact eigenstates for a fixed evolution time T_{tot}, we can ask how the entanglement entropy of the final state depends on T_{tot}. Our results are given in Fig. 6, where we show the entanglement entropy for a cut in the middle of the system averaged over 10000 disorder realizations. We find good
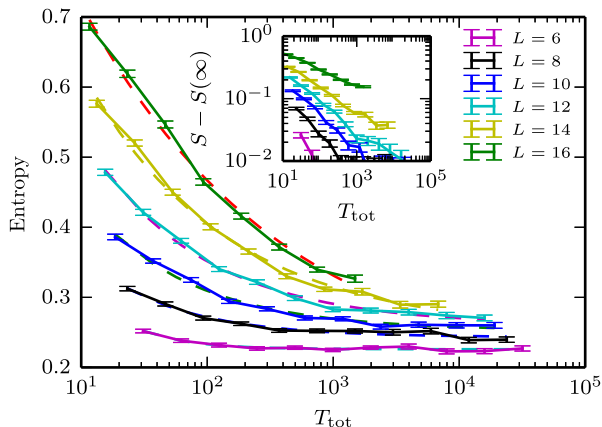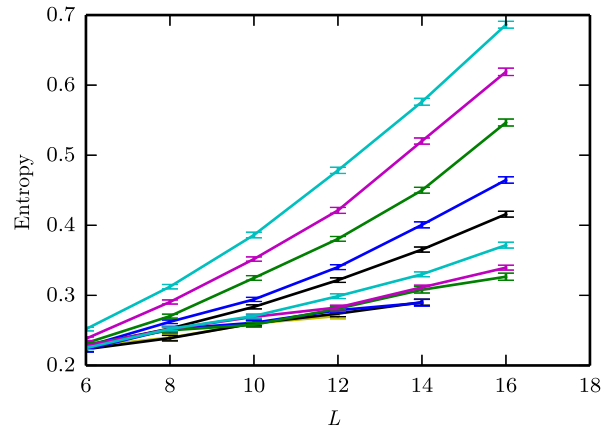

FIG. 7.   The bipartite entanglement entropy for a center cut as a function of system size L in the strong disorder limit, as obtained by IQPE to accuracy (from top to bottom) $k = 10, 11, \ldots, 20$. As the accuracy is increased, and the state begins to approach an energy eigenstate, the entanglement entropy decreases.

agreement with a power-law fit, i.e., that the entanglement entropy approaches the exact value as T^{-b} for some power b. In the inset of the figure, we show the same data after subtracting the constant term on a double-logarithmic scale. For the small system sizes accessible to our simulations, the power b appears to grow slowly as L is increased. This dependence crucially affects how the time T_{tot} necessary to reach a given error in entropy ϵ scales in the system size. Assuming that this growth is sufficiently slow once an asymptotic regime is reached, we can observe an area law by evolving for a time that is polynomial in the system size. To illustrate this, we show in Fig. 7 the entropy vs system size L for different choices of T_{tot}, and we observe that as T_{tot} is increased, an area law is observed for an increasing range of system sizes.

III. MEASUREMENTS

A. General remarks

Suppose now that we have obtained an energy eigenstate $|\psi_E\rangle$ at some random energy E by quantum phase estimation. What can we do with it to probe many-body localization? As mentioned in the Introduction, in contrast to classical simulations, we cannot simply examine the wave function to deduce its properties, and we are limited to performing unitary operations and projective measurements. Each such projective measurement yields, at most, N binary numbers, where N is the number of qubits used for representing the physical system; the expectation value is then reconstructed by averaging over many such measurements.

This gives rise to an additional complication when characterizing eigenstates obtained with the method described above: Since this method does not allow us to target a specific eigenstate, we are unlikely to encounter the

FIG. 6.   Entanglement entropy at the center of the system. Dashed lines indicate a fit to $S = S_0 + aT_{tot}^{-b}$, where a and b are fit parameters. Inset: $S(T_{tot}) - S(T_{tot} \to \infty)$, where $S(\infty)$ has been obtained from the dashed fit in the main panel, on a double-logarithmic scale.

same eigenstate more than once, and since we only measure $k < N$ bits of the energy, we cannot uniquely identify the state by its energy. We thus average *simultaneously* over eigenstates in some energy window, where the width of the energy window depends on k, and over measurement outcomes. Notably, within the constraints of this setup, there is no known way to extract the entanglement entropy of the resulting eigenstates.

Nevertheless, there are several powerful ways in which eigenstates can be characterized under these constraints. In the following, we discuss the examples of transport in weakly perturbed eigenstates (either with a weak global perturbation or a local perturbation) and how to adapt recent spin-echo proposals [23].

If, on the other hand, we consider states obtained by performing global quenches on easily prepared initial states, the preparation procedure becomes reproducible: We can prepare the same initial state over and over and apply the time evolution for the same time interval t, and thereby prepare the same final state many times and perform repeated measurements on this state. Although the signatures of many-body localization are not as clear in this setup, they are more easily obtained because of the possibility of repeated measurements on multiple copies of a state.

B. Linear response and transport in perturbed energy eigenstates

It is instructive to briefly consider how an isolated system at fixed energy can be probed in an experiment. One way is to couple the system to another, better-understood auxiliary system and see how it reacts. In a sense, this is similar to coupling it to a heat bath but in the limit in which the coupling is very small and can be turned on locally so that the auxiliary system can be used in a manner analogous to a thermometer. In a many-body localized phase, we do not expect particles (or energy) to flow into the auxiliary system.

A second possibility is to "tilt" the system. Having obtained an approximate eigenstate $|E\rangle$, we can evolve it for time T with the Hamiltonian

$$H = H_f + \sum_j V j n_j, \qquad (16)$$

where H_f is defined in Eq. (1). This would correspond, in a cold-atom experiment, to loading atoms into the trap with fixed energy and then tilting the potential in the trap, as in Ref. [51]. For small V and if the system is initially prepared close to an eigenstate, this corresponds to a weak global perturbation. We can then measure the current $i(c_{j+1}^\dagger c_j -$ H.c.) at various locations within the system. Restricting to one dimension, where no Jordan-Wigner strings need to be accounted for, this measurement is performed straightforwardly, as shown in Fig. 8. Consider the following input state with an ancilla in state $|0\rangle$:

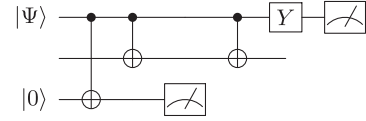

FIG. 8. Quantum circuit to measure the current operator $i(c_i^\dagger c_j - c_j^\dagger c_i)$, as described in the main text. Here, the top two qubits are the qubits corresponding to the physical sites i and j, and the bottom qubit is an ancilla qubit.

$$(a|00\rangle + b|01\rangle + c|10\rangle + d|11\rangle)|0\rangle. \qquad (17)$$

The first two qubits are the data qubits, which are the occupation numbers of sites j, $j + 1$, and the third qubit is the ancilla. The circuit in Fig. 8 implements the following operations:

$$2\,\text{CNOTs}: \ (a|00\rangle + d|11\rangle)|0\rangle + (b|01\rangle + c|10\rangle)|1\rangle, \quad (18)$$

$$\text{Measure 1}: \ b|01\rangle + c|10\rangle, \qquad (19)$$

$$\text{CNOT}: \ (b|0\rangle + c|1\rangle)|1\rangle. \qquad (20)$$

Here, in the step denoted as "Measure 1," we measure the ancilla qubit. If the measurement outcome is 0, the total measurement is 0 and the remaining steps need not be performed. If the outcome is 1, then CNOT is performed on the data qubits, followed by a measurement of the first qubit in the Y basis.

If the system is in a metallic state, then we expect the current to grow at short times until it reaches a steady-state value, at a time on the order of the mean-free time $\tau \sim t/W^2$. This value persists until the boundaries of the system reflect the current, at a time $T \sim L^2/D \sim L^2(W^2/t^3)$, where D is the diffusion constant, $D \sim t^3/W^2$. (We are working here in units in which the lattice spacing is 1.) Therefore, so long as $L \gg (t/W)^2$, there is a large interval of times over which the steady-state value can be observed. In a MBL state, on the other hand, we expect the current to grow at short times and then rebound at a time on the order of $T \sim \xi^2/D$, where ξ is the localization length. After that, it will undergo damped oscillations, before reaching a vanishing steady-state value. Therefore, apart from a short time transient, a current oscillating about and tending to zero will be observed.

Alternatively, we can study a "local quench": Starting in an energy eigenstate, we can perturb the system locally, e.g., by flipping the spins in a small region R (possibly even a single spin), and study the resulting time evolution. Unlike a global quench, which is expected to show unbounded growth of the entanglement, a local quench of a MBL state is expected to disturb the system only in a localized region. This can be traced by following the spreading of correlations in the system and observing, e.g., a zero-velocity Lieb-Robinson bound [52]. Alternatively, we can measure the energy current at distant

locations. The energy current (as per Noether's theorem) on the link between j and $j+1$ is $i(c_{j+2}^\dagger c_j - \text{H.c.}) + i(c_{j+1}^\dagger c_{j-1} - \text{H.c.})$. This vanishes exponentially with distance at long times in the MBL phase.

C. Spin echo

One interesting variant on a local quench was suggested by Serbyn et al. [23]. This is most easily described in the context of a spin model (2). In their proposal, one begins with a system in an initial product state in the σ^z basis, except for the ith spin, which is in a $\sigma_i^x = +(1/2)$ eigenstate. The system is evolved for time T. The ith spin is then reversed by applying σ^x, and the system is evolved again for time T. If the Hamiltonian were diagonal in the σ^z basis, this would return the ith spin to its initial state. If spin-flip terms do not change the physics qualitatively, i.e., if the system is adiabatically connected to one in which the Hamiltonian is diagonal in the σ^z basis, then the ith spin will not return precisely to its initial state but to a state with $\langle \sigma_i^x \rangle > 0$. Since this does not distinguish between many-body localization and single-particle localization, Serbyn et al. [23] propose that some other spin or set of spins in a region R far from i is manipulated (e.g., with a $\pi/2$ pulse). In a noninteracting localized state, this would have no additional effect because of the absence of coupling between spins in region R and the ith spin. In a many-body localized state, however, there would be an intermediate range of T values over which $\langle \sigma_i^x \rangle$ would show power-law decay before saturating to a nonzero value at large T. In a delocalized state, by contrast, $\langle \sigma_i^x \rangle$ decays rapidly to zero with T. It is a straightforward matter to stop the time evolution of a quantum state to reverse spin i and to manipulate spins in region R before continuing the time evolution.

IV. ERRORS

Thus far, we have assumed that our quantum computer has infinite coherence time and that all operations can be performed flawlessly. Obviously, this will not be the case, and errors must be taken into account in any appraisal of the prospects for applying a quantum computer. One type of error that may be relatively benign is Trotterization errors. Such errors, which are systematic, may be reinterpreted as a modification of the Hamiltonian. This effective Hamiltonian will have nonlocal terms induced by higher-order commutators of the original Hamiltonian terms. However, such commutators are exponentially suppressed in their order, and thus the effective Hamiltonian has, at worst, exponentially decaying terms. If we are interested in universal properties of MBL states, such a modification of the Hamiltonian will be unimportant.

Much more serious errors are caused by the environment. However, of these, bit-flip errors may be far more problematic than phase errors. Suppose that the

Hamiltonian for a MBL state can be written in terms of quasilocal conserved quantities τ_i^z [14,16]:

$$H = \sum_i h_i \tau_i^z + \sum_{i,j} J_{ij} \tau_i^z \tau_j^z + \cdots. \qquad (21)$$

Then, the addition of a coupling to the environment that is diagonal in this basis (and, therefore, leads only to phase errors) takes the form

$$H_{\text{sys-env}} = \sum_i \tau_i^z B(X_a). \qquad (22)$$

Here, $B(X_a)$ is the effective field, which depends on the environment degrees of freedom, X_a, and, therefore, entangles them with the quasilocal conserved quantities τ_i^z. However, such a coupling clearly has no effect on many-body localization.

Bit-flip errors, on the other hand, may have a rather drastic effect, in particular, during the final stages of the quantum phase estimation: Once the state is close to an eigenstate, a flip of a single spin will change the energy of the state by a large amount. Because of its recursive nature, this will lead to failure of the IQPE algorithm, and bit-flip errors must thus be corrected at a lower level. This may be achieved by using topological qubits [53] or by performing some error correction on the physical qubits. If we only correct bit-flip errors, then we can essentially use a classical error-correcting code. Consider, for illustrative purposes, a $[7, 4]$ Hamming code. If we assume flawless error detection and recovery, then an initial error rate of ε is lowered to an error rate of $21\varepsilon^2$. In order to perform 10^6 gates, we would, thus, need a gate fidelity of 99.99%, at a cost of encoding four sites in seven qubits, corresponding to a less than twofold increase in the required number of qubits.

V. CONCLUSION

Since classical computers are limited, for the foreseeable future, to the study of MBL systems of approximately 20 sites or less, a quantum computer need not be very large to accommodate a significantly larger system. From the preceding considerations, it appears that a system of 50 sites could be simulated with fewer than 100 physical qubits, assuming realistic bit-flip error rates. Moreover, as we have shown, a quantum computer can, in a straightforward manner, manipulate such a system in ways that would be very time-consuming with a classical computer. Two important features of many-body localization pave the way to the practical application of a quantum computer: (i) The Hamiltonian is local and can be written in terms of just three groups of noncommuting terms. This greatly reduces the number of gates required for time evolution. (ii) Since we are interested in dynamical properties and properties of highly excited states, possibly close to a phase transition where the states acquire a volume law, classical computers

are limited to very small systems on the order of 20 sites. Even deep in the localized phase, where states exhibit an area law and efficient algorithms exist to find the ground state, there are no presently known algorithms that efficiently find highly excited states. (iii) We are primarily interested in universal features and can, therefore, tolerate certain types of errors, unlike in the case of Shor's algorithm or applications to quantum chemistry. As a result, a relatively small quantum computer can, in a reasonable time, evolve an initial state to longer times t than would be possible with a classical simulation. In addition, a quantum computer can be used to find an approximation to a typical energy eigenstate of a Hamiltonian. Both applications of a quantum computer can give insights into many-body localization.

It is worth emphasizing that these results can be complementary to those obtained with a classical computer. We do not have access to a classical representation of a quantum state prepared with a quantum computer. Moreover, we cannot prepare multiple copies of the same approximate energy eigenstate. For these reasons, there is no obvious way to compute the entanglement entropy of an approximate energy eigenstate. However, we can study features of MBL systems that would be very difficult with a classical computer, such as transport, spin-echo effects, and the long-time approach to a thermal or nonthermal state.

If one regards the quantum computation as a highly idealized model for an experiment on an isolated quantum system, our results imply that the properties of eigenstates can be observed in the laboratory with resources polynomial in the system size. This is a nontrivial insight, as exact eigenstates cannot generally be prepared to very high accuracy unless exponentially large resources are used. This gives additional justification for studying the properties of a many-body localized phase in its energy eigenstates but also has implications for experiments.

Indeed, some of the proposals we discuss bear great similarity to experimental approaches, for example, in cold-atom systems. For example, the transport scenario of measuring the response of an energy eigenstate to a weak tilt is relevant to experiments such as the one reported in Ref. [51]. In this experiment, fermionic atoms are loaded into an optical lattice in a trap. A speckle pattern disorders the potential in the trap. The atoms carry spin 1/2 and interact via an on-site Hubbard-like interaction. In addition, the atoms are not in their ground state but have some energy density that is fixed when they are loaded into the trap. Both of these conditions indicate that this system is in a regime in which many-body localization could be observed. The trap is tilted, and then the momentum distribution is measured. When the disorder is weak, the momentum distribution is skewed by the tilt. When it is strong, the momentum distribution is unaffected. This is broadly consistent with many-body localization, but it is difficult to distinguish a transition from a crossover and, therefore, difficult to determine whether the putative localized phase in the

experiment is, in fact, a metallic phase with small but nonzero conductivity.

ACKNOWLEDGMENTS

[1] I. V. Gornyi, A. D. Mirlin, and D. G. Polyakov, *Interacting Electrons in Disordered Wires: Anderson Localization and Low-t Transport*, Phys. Rev. Lett. **95**, 206603 (2005).

[2] D. M. Basko, I. L. Aleiner, and B. L. Altshuler, *Metal Insulator Transition in a Weakly Interacting Many-Electron System with Localized Single-Particle States*, Ann. Phys. (Amsterdam) **321**, 1126 (2006).

[3] D. M. Basko, I. L. Aleiner, and B. L. Altshuler, *Problems of Condensed Matter Physics: Quantum Coherence Phenomena in Electron-Hole and Coupled Matter-Light Systems*, edited by A. L. Ivanov and S. G. Tikhodeev (Oxford University Press, Oxford, 2007).

[4] J. M. Deutsch, *Quantum Statistical Mechanics in a Closed System*, Phys. Rev. A **43**, 2046 (1991).

[5] M. Srednicki, *Chaos and Quantum Thermalization*, Phys. Rev. E **50**, 888 (1994).

[6] M. Rigol, V. Dunjko, and M. Olshanii, *Thermalization and Its Mechanism for Generic Isolated Quantum Systems*, Nature (London) **452**, 854 (2008).

[7] V. Oganesyan and D. A. Huse, *Localization of Interacting Fermions at High Temperature*, Phys. Rev. B **75**, 155111 (2007).

[8] G. De Chiara, S. Montangero, P. Calabrese, and R. Fazio, *Entanglement Entropy Dynamics of Heisenberg Chains*, J. Stat. Mech. (2006) P03001.

[9] M. Žnidarič, T. Prosen, and P. Prelovšek, *Many-Body Localization in the Heisenberg xxz Magnet in a Random Field*, Phys. Rev. B **77**, 064426 (2008).

[10] A. Pal and D. A. Huse, *Many-Body Localization Phase Transition*, Phys. Rev. B **82**, 174411 (2010).

[11] J. H. Bardarson, F. Pollmann, and J. E. Moore, *Unbounded Growth of Entanglement in Models of Many-Body Localization*, Phys. Rev. Lett. **109**, 017202 (2012).

[12] S. Iyer, V. Oganesyan, G. Refael, and D. A. Huse, *Many-Body Localization in a Quasiperiodic System*, Phys. Rev. B **87**, 134202 (2013).

[13] R. Vosk and E. Altman, *Many-Body Localization in One Dimension as a Dynamical Renormalization Group Fixed Point*, Phys. Rev. Lett. **110**, 067204 (2013).

[14] M. Serbyn, Z. Papić, and D. A. Abanin, *Local Conservation Laws and the Structure of the Many-Body Localized States*, Phys. Rev. Lett. **111**, 127201 (2013).

[15] M. Serbyn, Z. Papić, and D. A. Abanin, *Universal Slow Growth of Entanglement in Interacting Strongly Disordered Systems*, Phys. Rev. Lett. **110**, 260601 (2013).

[16] D. A. Huse and V. Oganesyan, *A Phenomenology of Certain Many-Body-Localized Systems*, arXiv:1305.4915.

[17] Y. Bahri, R. Vosk, E. Altman, and A. Vishwanath, *Localization and Topology Protected Quantum Coherence at the Edge of 'Hot' Matter*, arXiv:1307.4092.

[18] D. Pekker, G. Refael, E. Altman, E. Demler, and V. Oganesyan, *The Hilbert-Glass Transition: New Universality of Temperature-Tuned Many-Body Dynamical Quantum Criticality*, Phys. Rev. X **4**, 011052 (2014).

[19] N. Y. Yao, C. R. Laumann, S. Gopalakrishnan, M. Knap, M. Mueller, E. A. Demler, and M. D. Lukin, *Many-Body Localization with Dipoles*, arXiv:1311.7151 [Phys. Rev. Lett. (to be published)].

[20] A. Chandran, V. Khemani, C. R. Laumann, and S. L. Sondhi, *Many-Body Localization and Symmetry-Protected Topological Order*, Phys. Rev. B **89**, 144201 (2014).

[21] B. Bauer and C. Nayak, *Area Laws in a Many-Body Localized State and Its Implications for Topological Order*, J. Stat. Mech. 2013, P09005 (2013).

[22] J. A. Kjäll, J. H. Bardarson, and F. Pollmann, *Many-Body Localization in a Disordered Quantum Ising Chain*, Phys. Rev. Lett. **113**, 107204 (2014).

[23] M. Serbyn, M. Knap, S. Gopalakrishnan, Z. Papić, N. Y. Yao, C. R. Laumann, D. A. Abanin, M. D. Lukin, and E. A. Demler, *Interferometric Probes of Many-Body Localization*, Phys. Rev. Lett. **113**, 147204 (2014).

[24] R. Nandkishore and D. A. Huse, *Many Body Localization and Thermalization in Quantum Statistical Mechanics*, arXiv:1404.0686.

[25] J. Z. Imbrie, *On Many-Body Localization for Quantum Spin Chains*, arXiv:1403.7837.

[26] R. P. Feynman, *Simulating Physics with Computers*, Int. J. Theor. Phys. **21**, 467 (1982).

[27] D. Deutsch, *Quantum Computational Networks*, Proc. R. Soc. A **425**, 73 (1989).

[28] P. W. Shor, *Algorithms for Quantum Computation: Discrete Logarithms and Factoring*, in *Proceedings of the 35th Annual Symposium on Foundations of Computer Science, 1994* (IEEE, New York, 1994), pp. 124–134.

[29] P. W Shor, *Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer*, SIAM J. Comput. **26**, 1484 (1997).

[30] M.-H. Yung, J. D. Whitfield, S. Boixo, D. G. Tempel, and A. Aspuru-Guzik, *Introduction to Quantum Algorithms for Physics and Chemistry* (John Wiley & Sons, New Jersey, 2014).

[31] D. Wecker, B. Bauer, B. K. Clark, M. B. Hastings, and M. Troyer, *Gate Count Estimates for Performing Quantum Chemistry on Small Quantum Computers*, Phys. Rev. A **90**, 022305 (2014).

[32] M. B. Hastings, D. Wecker, B. Bauer, and M. Troyer, *Improving Quantum Algorithms for Quantum Chemistry*, arXiv:1403.1539.

[33] D. Poulin, M. B. Hastings, D. Wecker, N. Wiebe, A. C. Doherty, and M. Troyer, *The Trotter Step Size Required for Accurate Quantum Simulation of Quantum Chemistry*, arXiv:1406.4920.

[34] This may not apply if the relevant energy scale is a local scale instead of a global one.

[35] D. W. Berry and A. M. Childs, *Black-Box Hamiltonian Simulation and Unitary Implementation*, Quantum Information Computation **12**, 29 (2012).

[36] A. Daskin, A. Grama, G. Kollias, and S. Kais, *Universal Programmable Quantum Circuit Schemes to Emulate an Operator*, J. Chem. Phys., **137**, 234112 (2012).

[37] D. W Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, *Exponential Improvement in Precision for Simulating Sparse Hamiltonians*, arXiv:1312.1414.

[38] H. F. Trotter, *On the Product of Semi-groups of Operators*, Proc. Am. Math. Soc. **10**, 545 (1959).

[39] M. Suzuki, *Generalized Trotter's Formula and Systematic Approximants of Exponential Operators and Inner Derivations with Applications to Many-Body Problems*, Commun. Math. Phys. **51**, 183 (1976).

[40] D. Berry, G. Ahokas, R. Cleve, and B. Sanders, *Efficient Quantum Algorithms for Simulating Sparse Hamiltonians*, Commun. Math. Phys. **270**, 359 (2007).

[41] Note that in Ref. [31], it was pointed out that the bound of Ref. [40] can be made more tight and stated in terms of the norm of the individual noncommuting terms; however, in our setup, the norm of these is also extensive, and this tighter bound makes no difference.

[42] J. D. Whitfield, J. Biamonte, and A. Aspuru-Guzik, *Simulation of Electronic Structure Hamiltonians Using Quantum Computers*, Mol. Phys. **109**, 735 (2011).

[43] S. B. Bravyi and A. Y. Kitaev, *Fermionic Quantum Computation*, Ann. Phys. (Amsterdam) **298**, 210 (2002).

[44] N. Cody Jones, J. D. Whitfield, P. L. McMahon, M.-H. Yung, R. Van Meter, A. Aspuru-Guzik, and Y. Yamamoto, *Faster Quantum Chemistry Simulation on Fault-Tolerant Quantum Computers*, New J. Phys. **14**, 115023 (2012).

[45] A. Y. Kitaev, *Quantum Measurements and the Abelian Stabilizer Problem*, arXiv:quant-ph/9511026.

[46] A. Yu Kitaev, A. H. Shen, and M. N. Vyalyi, *Classical and Quantum Computation* (American Mathematical Society, Providence, Rhode Island, 2002), p. 47.

[47] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, England, 2010).

[48] S. Parker and M. B. Plenio, *Efficient Factorization with a Single Pure Qubit and log N Mixed Qubits*, Phys. Rev. Lett. **85**, 3049 (2000).

[49] R. Nandkishore, S. Gopalakrishnan, and D. A. Huse, *Spectral Features of a Many-Body Localized System Weakly Coupled to a Heat Bath*, Phys. Rev. B **90**, 064203 (2014).

[50] S. Johri, R. Nandkishore, and R. N. Bhatt, *Numerical Study of a Many-Body Localized System Coupled to a Bath*, arXiv:1405.5515.

[51] S. S. Kondov, W. R. McGehee, and B. DeMarco, *Interplay of Disorder and Interactions in an Optical Lattice Hubbard Model*, arXiv:1305.6072 [Phys. Rev. Lett. (to be published)].

[52] E. Hamza, R. Sims, and G. Stolz, *Dynamical Localization in Disordered Quantum Spin Systems*, Commun. Math. Phys. **315**, 215 (2012).

[53] C. Nayak, S. H. Simon, A. Stern, M. Freedman, and S. Das Sarma, *Non-abelian Anyons and Topological Quantum Computation*, Rev. Mod. Phys. **80**, 1083 (2008).

[54] B. Bauer *et al.*, *The ALPS Project Release 2.0: Open Source Software for Strongly Correlated Systems*, J. Stat. Mech. (2011) P05001.