

Classical algorithms for many-body quantum systems at finite energies

Yilun Yang , J. Ignacio Cirac , and Mari Carmen Bañuls 

Max-Planck-Institut für Quantenoptik, Hans-Kopfermann-Straße 1, D-85748 Garching, Germany
and Munich Center for Quantum Science and Technology (MCQST), Schellingstraße 4, D-80799 München, Germany



(Received 3 May 2022; revised 29 June 2022; accepted 30 June 2022; published 15 July 2022)

We investigate quantum-inspired algorithms to compute physical observables of quantum many-body systems at finite energies. They are based on the quantum algorithms proposed by S. Lu, M. C. Bañuls, and J. I. Cirac [PRX Quantum **2**, 020321 (2021)], who use the quantum simulation of the dynamics of such systems, as well as classical filtering and sampling techniques. Here, we replace the quantum simulation by standard classical methods based on matrix product states and operators. As a result, we can address significantly larger systems than those reachable by exact diagonalization or by other algorithms. We demonstrate the performance with spin chains up to 80 sites.

DOI: [10.1103/PhysRevB.106.024307](https://doi.org/10.1103/PhysRevB.106.024307)

I. INTRODUCTION

Computing the properties of quantum systems in equilibrium is of central interest in many-body physics. For a system at finite temperature, there exists a wide spectrum of techniques that are used in practice. For large systems, where exact solutions are unreachable, they typically approximate the corresponding Gibbs state using variational, sampling, or series expansion methods [1–7]. For systems at a finite energy, e.g., in the microcanonical ensemble, methods are more scarce [8–10].

A possible approach consists of simulating the dynamics and extracting the equilibrium properties from there. For instance, one can use spectral filters in order to retrieve expectation values of an observable O , by averaging them at different times. In this way, one obtains results connected to the diagonal ensemble corresponding to the initial state [11], namely,

$$\bar{O} = \sum_n |c_n|^2 \langle E_n | O | E_n \rangle, \quad (1)$$

where $|E_n\rangle$ are the energy eigenstates and c_n the coefficients of the initial state in that basis. For local Hamiltonians and initial states with finite correlation length (like product states), the values of c_n are significant if $|E_n - E| < \delta = \mathcal{O}(\sqrt{N})$, where E is the mean energy of the initial state. Under the eigenstate thermalization hypothesis (ETH) [12–14], \bar{O} converges to the equilibrium value in the thermodynamic limit, and it is not necessary to average the results at different time but just to wait for a sufficiently long time. Quantum

computers and analog quantum simulators are very well suited for that task, since they can deal with the dynamics of many-body quantum systems in a very natural way [15,16]. Classical methods to simulate the dynamics typically suffer from the linear growth of entanglement [17], which gives an exponential cost with time. Even if the (weak) ETH applies, the thermalization time can be very long [18,19], and this severely restricts the applicability of such classical algorithms, regardless of the typical slow growth of entanglement in this case.

In this work we propose and analyze a classical algorithm to compute expectation values of the form (1). In particular, for c_n that are Gaussian functions of $|E_n - E|$ with a variance δ , this can be achieved by simulating the dynamics for a time $\mathcal{O}(1/\delta)$. This allows us, for instance, to reach $\delta = \mathcal{O}(\sqrt{N})$ by just using standard time evolution techniques for tensor networks for very short times, when the entanglement is still very small and thus the techniques work well. One can also reach values of $\delta = \mathcal{O}(1)$ with modest computational resources. The algorithm is inspired by a quantum algorithm presented in [10] that allows one to compute expectation values of the form (1). This algorithm combines classical sampling (Monte Carlo) techniques with time series and Loschmidt echo-like measurements [20,21] that can be obtained by quantum simulation of the dynamics. Our main modification is to replace the latter by a classical simulation using tensor network states [22–27]. This allows us to compute (1) for times $\mathcal{O}(1/\delta)$ instead of the thermalization time, thus circumventing the problem of entanglement growth. This is done at the expense of having to sample, which just involves the repetition of the whole procedure until convergence. We apply the algorithm to one-dimensional systems, sample over a basis of product states, and use matrix product states (MPSs) and operators (MPOs) to simulate the evolution [28]. We illustrate the performance of the method for a nonintegrable Ising chain, for which we obtain convergence to the microcanonical values for systems up to 80 sites, far larger than what is possible with exact diagonalization.

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Open access publication funded by the Max Planck Society.

Apart from that, in [10] another quantum algorithm was proposed to compute physical observables in a state where an energy filter of width δ is applied. For local Hamiltonians, the computational time also scales as $\mathcal{O}(1/\delta)$. Here we also analyze a classical algorithm inspired in that method. We notice that this method typically requires much narrower δ (thus longer times) to approach thermodynamic quantities. However, the filtering achieved with a limited evolution time can be optimized if the initial state is chosen with already a reduced energy width. Here we demonstrate this possibility by applying the classical version of the first algorithm on matrix product states found by minimizing the energy variance.

The rest of the paper is organized as follows. In Sec. II we briefly review the concept of energy filters, introduce the filter ensemble, and discuss its applications to determine microcanonical and diagonal properties. We also review briefly the quantum algorithms [10] that motivate this work. In Sec. III we discuss the details of a tensor network (TN) simulation of the quantum algorithms, the different possibilities, and the associated parameter choices. Section IV presents our numerical results for Ising chains, for each of the algorithms implemented. The paper is closed with the discussion in Sec. V.

II. FILTERS AND QUANTUM ALGORITHMS

We start by recalling the definition and properties of the energy filters that are at the basis of the algorithms in [10] and this work.

A. Energy filters

The main tool used by the finite-energy algorithms discussed here is a filtering operator that suppresses energy eigenstates outside a target energy interval. In particular, given the Hamiltonian H , we define a Gaussian filter centered in energy E and of width δ as the following operator:

$$\hat{P}_\delta(E) = \exp[-(\hat{H} - E)^2/2\delta^2]. \quad (2)$$

Notice that, up to normalization, $\hat{P}_\delta(E)$ is a diagonal ensemble in the energy basis. We refer to it as the *filter ensemble*:

$$\rho_{(E,\delta)} = \frac{\hat{P}_\delta(E)}{\text{tr}[\hat{P}_\delta(E)]}. \quad (3)$$

The corresponding expectation values are precisely of the form (1), with coefficients $|c_n|^2$ distributed according to a Gaussian of width δ . For local Hamiltonians and large systems, product states have that kind of spectral decomposition, with $\delta \sim \mathcal{O}(\sqrt{N})$ [29,30], but since they contain coherences in the energy basis, the corresponding expectation values are very different. An ensemble as (3) could nevertheless be obtained from a product state, but only after evolving and averaging over a long time.

As the width δ is reduced, the filter approaches the microcanonical ensemble. Thus we can make use of the filter to access the microcanonical properties of the quantum system in the following two different ways.

(a) *Filtering a state.* Given a state $|\psi\rangle$, its local density of states (LDOS) is defined as $D_\psi(E) = \langle \psi | \delta(E - \hat{H}) | \psi \rangle$. A




FIG. 1. Approximating $\hat{P}_\delta(E_0)$ with $\hat{F}_{\delta,\alpha}(E_0)$.

broadened version can be computed with the filter as

$$D_{\delta,\psi}(E) = \frac{1}{\sqrt{2\pi}\delta} \langle \psi | \hat{P}_\delta(E) | \psi \rangle. \quad (4)$$

We can also use the filtered state to explore the microcanonical ensemble expectation value $O_{\text{micro}}(E)$ of an observable \hat{O} . In generic cases in which ETH is satisfied, and for a value E at which the LDOS does not vanish,

$$O_\delta(E, |\psi\rangle) = \frac{\langle \psi | \hat{P}_\delta(E) \hat{O} \hat{P}_\delta(E) | \psi \rangle}{\langle \psi | \hat{P}_\delta(E)^2 | \psi \rangle} \quad (5)$$

will converge to $O_{\text{micro}}(E)$ in the limit $\delta \rightarrow 0$.

(b) *Filtering the whole spectrum.* The filter ensemble itself converges to the microcanonical ensemble as the width is reduced. Hence, we can also use it without specifying a state, but directly taking its trace. In particular, the density of states (DOS) of the Hamiltonian H , defined as $D(E) = \text{tr}[\delta(E - \hat{H})]$, can be approximated through the broadening of the δ functions as

$$D_\delta(E) = \text{tr}[\hat{P}_\delta(E)]. \quad (6)$$

Moreover, the expectation values in the filter ensemble

$$O_\delta(E) = \text{tr}[\hat{O} \rho_{(E,\delta)}] = \text{tr}[\hat{O} \hat{P}_\delta(E)] / \text{tr}[\hat{P}_\delta(E)] \quad (7)$$

will converge to the microcanonical values as $\delta \rightarrow 0$.

Whereas in generic cases one can in principle approach the microcanonical expectation values with either (5) or (7), the convergence of the second with δ is much faster, since the filter ensemble is diagonal in the energy basis, while (5) contains contributions from off-diagonal matrix elements [31–33], which can converge much slower than the diagonal part [14,34].

Implementing the filter

For the purpose of numerical, but also quantum, simulations of the filter, it is convenient to substitute an approximation for the Gaussian filter (see Fig. 1). In this paper, following the quantum algorithms in [10], we focus on the cosine filter [35,36] defined as

$$\hat{F}_{\delta,\alpha}(E) = \cos[(\hat{H} - E)/\alpha]^M \approx \hat{P}_\delta(E), \quad (8)$$

where α is a parameter (with dimensions of energy) that controls the validity of the approximation, and $M = \lfloor \alpha^2/\delta^2 \rfloor_2$ with $\lfloor \dots \rfloor_2$ giving the closest smaller even integer. The approximation is valid if $\|\hat{H} - E\| \leq \alpha\pi/2$, when the spectrum of \hat{H} lies in one period of the cosine function [37]. Equation (8) can be further approximated by a truncated series of

evolution operators

$$\hat{F}_{\delta,\alpha}(E) \approx \hat{F}_{\delta,\alpha}^x(E) = \sum_{m=-R}^R c_m e^{-i(\hat{H}-E)t_m}, \quad (9)$$

where $R = \lfloor x\alpha/\delta \rfloor$, $t_m = 2m/\alpha$, x is a constant that bounds the truncation error in the operator norm as $\|\hat{P}_\delta(E) - \hat{F}_{\delta,\alpha}(E)\| \leq 2e^{-x^2/2}$, and

$$c_m = \frac{1}{2^M} \binom{M}{M/2 - m}. \quad (10)$$

With the cosine filter the problem is turned into evolving states or computing the traces of time evolution operators, which leads itself to a natural implementation in quantum simulators.

B. The quantum algorithms

In Lu *et al.*'s paper [10], two hybrid classical-quantum algorithms corresponding to the two different ways of applying the filter were introduced. We sketch them here for completeness.

The first one computes (5) for a state $|\psi\rangle$ that can be easily prepared. Suppose the quantum device can efficiently obtain the following quantities:

$$\begin{aligned} a_\psi(t) &= \langle \psi | e^{-i\hat{H}t} | \psi \rangle, \\ a_{O,\psi}(t_1, t_2) &= \langle \psi | e^{i\hat{H}t_1} \hat{O} e^{-i\hat{H}t_2} | \psi \rangle; \end{aligned} \quad (11)$$

then (5) can be determined by classical postprocessing as

$$O_\delta(E, |\psi\rangle) = \frac{\sum_{m,n=-R}^R c_m^* c_n a_{O,\psi}(t_m, t_n)}{\sum_{m,n=-R}^R c_m^* c_n a_\psi(t_n - t_m)}, \quad (12)$$

without explicitly preparing the filtered state. The required timescale is proved to be a polynomial of system size N , the inverse of the width of filter $1/\delta$, and the inverse of the error, provided the state $|\psi\rangle$ can be prepared efficiently, for a value of E in a small interval around the mean energy of $|\psi\rangle$.

In the second algorithm (quantum-assisted Monte Carlo), importance sampling is applied to compute (7). Let us rewrite that expression as

$$O_\delta(E) = \frac{\int d\mu_\phi D_{\delta,\phi}(E) O_{\delta,\phi}(E)}{\int d\mu_\phi D_{\delta,\phi}(E)}, \quad (13)$$

where $\{|\phi\rangle\}$ is an (over)complete basis, with $d\mu_\phi$ the appropriate measure to ensure the closure relation $\int d\mu_\phi |\phi\rangle \langle \phi| = \mathbb{1}$ (a simple choice is for instance the computational basis). $D_{\delta,\phi}$ is the LDOS defined in (4) and

$$O_{\delta,\phi} = \langle \phi | \hat{O} \hat{P}_\delta(E) | \phi \rangle / \langle \phi | \hat{P}_\delta(E) | \phi \rangle. \quad (14)$$

Both $D_{\delta,\phi}(E)$ and $O_{\delta,\phi}(E)$ can be obtained by measuring the quantities defined in (11), as long as we can run the first algorithm with the quantum device for the states in the basis $\{|\phi\rangle\}$. Then a Metropolis-Hastings step can be applied classically with regard to the probability distribution $D_{\delta,\phi}(E) / \int d\mu_\phi D_{\delta,\phi}(E)$, and the value of $O_\delta(E)$ can be estimated. Because $D_{\delta,\phi}(E)$ is positive, this method does not encounter a sign problem.

Given δ , this second algorithm provides access to observables in the filter ensemble at the cost of simulating time evolutions for times $\mathcal{O}(1/\delta)$, at the expense of repeating the procedure until the sampling converges. This is especially remarkable because one could obtain a similar result from the time evolution of an initial state with the same distribution of coefficients, but this would require evolving for as long as the thermalization time. As a particular application, if one chooses $\delta = o(\sqrt{N})$, the expectation values of intensive quantities will (under the ETH) already be equivalent to those in the Gibbs ensemble at the same mean energy (see Sec. III B 1), and therefore this algorithm is an inexpensive way of accessing thermal properties.

If we are interested in microcanonical expectation values, we can use either algorithm, but we need to reduce the width of the filter. Since the trace quantities converge faster in δ to the microcanonical values, a shorter evolution time is required with this second algorithm. In exchange, the procedure needs to be repeated over many states, to perform the classical Monte Carlo sampling.

Extreme values of energy

Both algorithms above rely on the evolution of easily preparable states (a requirement for the single initial state in the first algorithm or the whole sampling basis in the second). A most practical choice is that of product states. The mean energies of such states are contained in an extensive but generally restricted interval within the spectrum [38], such that values close to the edges of the spectrum may be out of reach. As indicated in [10], the accessible range of energies can be extended by considering larger sets of states. In particular, MPSs can be used to circumvent the limitation.

For the first algorithm, the initial state $|\psi\rangle$ can be found as an MPS with a small bond dimension such that its energy expectation value is close enough to E , as the MPS serves as a good representation of the ground state and low-lying excited states. This can be done, for instance, by first finding the MPS with that bond dimension and minimal energy (E_{\min}), and then changing its parameters until the desired energy $E > E_{\min}$ is reached. Another possibility is to find the MPS minimizing $(H - E)^2$ [39,40].

For the second algorithm, a different basis for Monte Carlo sampling can be chosen, where we start from any state $|\phi_0\rangle$ whose mean energy is close to E , obtained in the same way, and apply a random Pauli matrix σ^x , σ^y , or σ^z on a random site in each proposed move. This strategy gives a complete basis set, as

$$\frac{1}{2^N} \sum_{\substack{\mu_i=0 \\ 1 \leq i \leq N}}^3 \sigma_N^{\mu_N} \cdots \sigma_1^{\mu_1} |\phi_0\rangle \langle \phi_0| \sigma_1^{\mu_1} \cdots \sigma_N^{\mu_N} = \mathbb{1} \quad (15)$$

for any state $|\phi_0\rangle$. The change in mean energy is $\mathcal{O}(1)$ in each move, and thus this choice of basis ensures enough states for sampling.

III. CLASSICAL SIMULATION

The methods that we study in this paper replace the quantum simulation of the dynamics in the quantum algorithms

of [10] by classical simulations using tensor networks. The longest evolution time required in (9), which is the deciding factor for the efficiency, is determined by the width of the filter δ as $t_{\max} = t_R \approx 2x/\delta$. A quantum simulator should be able to deal efficiently with times $t = \mathcal{O}(\text{poly}(N))$, which gives access to $\delta = \Omega(\text{poly}(1/N))$ [41]. This should be sufficient for both (5) (in the case in which ETH is satisfied [34]) and (7) to converge to the microcanonical values. On a classical computer, TN techniques provide the possibility to simulate the time evolution of a local Hamiltonian [22,23,28], but the bond dimension required to do so can increase exponentially with time. Thus, starting from a product state, we can simulate times $t_{\max} \propto \log N$ with a bond dimension polynomial in system size, which would allow us to efficiently perform classical simulations of the algorithm for $\delta = \Omega(1/\log N)$. For the actual implementation of the classical simulation of the dynamics, there exist several options, some of which we discuss in this section.

A. Tensor network implementation

There are multiple different approaches to simulate time evolution with TN techniques (see [28] for a recent review). Some of the most commonly used methods are based on a Suzuki-Trotter approximation of the time evolution operator. One possibility is then to repeatedly apply the approximated short time evolution steps onto a matrix product state (MPS), to obtain a representation of the time-evolved state. Alternatively, the time evolution operator itself, $e^{-i\hat{H}t}$, can be approximated by a matrix product operator (MPO), constructed also from the iteration of Trotterized steps. Therefore, we can use various techniques for the classical simulation of the quantum methods above.

Finding the MPO representation of each term $e^{-i\hat{H}t_m}$ as an MPO allows us to estimate $\text{tr}[\hat{P}_\delta(E)]$ as a linear combination of the corresponding traces, which for MPOs can be computed very efficiently. This strategy will however fail as we approach the edge of the spectrum for large system sizes while considering small values of δ . The reason is the extremely imbalanced distribution of the DOS $D(E)$, which becomes exponentially small when E is far from the center of the spectrum. For a traceless, local, and bounded Hamiltonian, in the thermodynamic limit $D(E)$ converges weakly to a Gaussian distribution with mean energy $E = 0$ and width proportional to \sqrt{N} [29,30]:

$$\int_{-\infty}^{E_0} D(E) dE \xrightarrow{N \rightarrow \infty} \int_{-\infty}^{E_0} \frac{d^N e^{-E^2/2N\sigma_0^2}}{\sqrt{2\pi N\sigma_0}} dE, \quad (16)$$

where d is the local Hilbert space dimension and σ_0 is some constant independent of the system size. Thus, for energies $\approx \sqrt{N}$, we expect $D(E)$ to become exponentially smaller than its value at the center. Therefore given all $\text{tr}(e^{-i\hat{H}t_m})$ that are reasonably precise, the ratio in (7) could still not be properly achieved all through the whole spectrum: the applicable energy range will be proportional to the square root of the system size, and hence $\mathcal{O}(1/\sqrt{N})$ in energy density, a restriction we also observed in a previous work [9].

Fortunately, this difficulty can be overcome with the importance-sampling method described in Sec. II B. To begin

with, in this method we only need to evaluate the ratio (14) for states for which the probability factor $\langle \phi | \hat{P}_\delta(E) | \phi \rangle$ is above some threshold. In particular, when E is away from the center of the spectrum, contributions from the exponentially large maximum of the DOS will be suppressed. Of course, one needs that the chosen basis $\{|\phi\rangle\}$ has enough states around the target energy E . Since for our numerical simulations we are free to choose any basis from which we can sample efficiently and whose states can be written as MPS, we can exploit this freedom to try to ensure this condition. A product basis minimizes the cost of the contractions and is often a good choice, since, as mentioned above, it covers an extensive window of the energy spectrum. If this is not the case, we can use any of the methods mentioned at the end of Sec. II to find an MPS with small bond dimension close to the desired energy, and use it to construct a complete basis of the form (15). For the cases we consider in this paper, the computational basis is already an adequate choice, sufficient to produce accurate numerical results over the full spectrum, as we illustrate in the next section.

In the Monte Carlo simulation, the time evolution can be done either at the level of the states, i.e., directly evolving the sampled state as an MPS, or at the level of the operators, i.e., approximating the evolution unitaries for each state as MPOs, storing them in memory, and using them later to do contractions with states randomly sampled from the basis. The second option has the advantage of simulating the dynamics a single time, as the same operators can be reused when doing the sampling over different states. Hence it is faster, but it also consumes much more memory to store the required MPOs [42]. Also, the bond dimension needed to approximate an evolution operator as an MPO is significantly larger than the one used to approximate a time-evolved MPS for the same time. As a concrete example, for system size $N = 80$, we find bond dimension $D_{\text{MPS}} = 40$ to be enough for evolving MPSs, and $D_{\text{MPO}} \sim 100$ for storing the MPOs, for $(\delta, \alpha) \propto (1, \sqrt{N})$ and 5×10^4 samples. The typical timescales taken are 1 week and 1 day, respectively, with an Intel Xeon Gold 6138 processor.

B. Filter parameters and the microcanonical limit

The cosine filter depends on two parameters: the width δ and the period of the filter α . As mentioned, δ determines the maximum time we need to evolve, $t_{\max} = 2x/\delta$, while for a fixed δ , α determines the number of terms in the expansion $R = x\alpha/\delta$. Additionally, in the Monte Carlo algorithm we introduce a cutoff parameter ϵ and discard states for which the probability is found to be below this threshold.

In this section we discuss the significance of these parameters, as well as which choices ensure approaching the microcanonical limit. Note that the conclusions are valid for both the classical and the quantum version of the algorithms.

1. Filter width δ

A width δ and mean energy E_0 determine the properties of the filter ensemble $\hat{P}_\delta(E_0)$. But to estimate the corresponding energy distribution we need to take into account the density of states. Assuming a Gaussian DOS as in (16), the energy distribution of the filter ensemble will also be a Gaussian

given by

$$D_{\rho_{(E_0, \delta)}}(E) = D(E) \exp \left[-\frac{(E - E_0)^2}{2\delta^2} \right] \\ \sim \exp \left(-\frac{E_0^2}{2\gamma N \sigma_0^2} \right) \exp \left[-\frac{\gamma}{2\delta^2} (E - E_0/\gamma)^2 \right], \quad (17)$$

where $\gamma = 1 + \delta^2/N\sigma_0^2$. We omitted an energy-independent factor in (17). It can be concluded that in the thermodynamic limit, the mean energy and width of the filtered ensemble are given by

$$E_{\rho_{(E_0, \delta)}} = E_0/\gamma, \quad \Delta_{\rho_{(E_0, \delta)}} = \delta/\sqrt{\gamma}. \quad (18)$$

If we choose $\delta \propto \sqrt{N}$, the mean energy of the ensemble is shifted with respect to the parameters of the filter, as explicitly shown in Fig. 4. A filter width that scales as $\delta = o(\sqrt{N})$ is enough to ensure that $E_{\rho_{(E_0, \delta)}} \rightarrow E$ and $\Delta_{\rho_{(E_0, \delta)}} \rightarrow \delta$ as $N \rightarrow \infty$. This observation is especially relevant if we are interested in approaching the microcanonical limit: in general, assuming ETH, in the thermodynamic limit a microcanonical energy shell centered at E will yield the thermal values for intensive quantities at energy density E/N if the width Δ satisfies $\Delta/N \rightarrow 0$. This condition is already satisfied for the filter ensemble with $\delta \propto \sqrt{N}$, which means that the expectation values will converge to the thermal ones, only at shifted energies, according to the previous argument (see Fig. 4).

We can similarly estimate the energy distribution of the pure state resulting from the application of the filter onto an individual state $|\psi\rangle$. There is actually a similar argument for $\hat{P}_\delta(E)|\psi\rangle$ if $|\psi\rangle$ is a product state, as such states also have essentially Gaussian LDOSs whose widths are proportional to \sqrt{N} [43]. With a spectral decomposition $|\psi\rangle = \sum_k c_k |E_k\rangle$, where $|E_k\rangle$ are energy eigenstates, the filtered state results as

$$|\hat{P}_\delta \psi\rangle := \sqrt{\Gamma} \hat{P}_\delta |\psi\rangle = \sqrt{\Gamma} \sum_k c_k e^{-\frac{(E-E_k)^2}{2\delta^2}} |E_k\rangle, \quad (19)$$

where $\Gamma = 1/\langle \psi | \hat{P}_\delta(E)^2 | \psi \rangle$ is the normalization factor. By choosing the center of the filter at the mean energy of the state $E = \langle \psi | \hat{H} | \psi \rangle =: E_\psi$, the average energy of the filtered state does not change. Assuming the LDOS of $|\psi\rangle$ has a Gaussian form with width $\sigma_\psi \sqrt{N}$, where σ_ψ is independent of system size, the energy variance of $|\psi\rangle$ can be estimated through substituting the sum over eigenstates by an integral over energy values with Gaussian weights. We obtain

$$\Delta_{\hat{P}_\delta \psi}^2 := \langle \hat{P}_\delta \psi | H^2 | \hat{P}_\delta \psi \rangle - \langle \hat{P}_\delta \psi | H | \hat{P}_\delta \psi \rangle^2 \quad (20)$$

$$\approx \frac{\delta^2}{2 + \delta^2/N\sigma_\psi^2}. \quad (21)$$

Again, we may want to consider how this affects approaching the microcanonical limit as the width of the filter is decreased. A major difference in this respect between the filter ensemble and the filtered state is that the second contains coherent contributions from different energy eigenstates. Thus (5) includes contributions from off-diagonal matrix elements in the energy basis, which only become negligible when the width of the energy distribution decreases sufficiently fast

with N . More concretely, from canonical typicality arguments we can expect that, for nonintegrable systems, the expectation value of a local observable converges to the thermal value when the energy deviation of the state decreases as a polynomial of $1/N$ [34]. In [35] we observed a trend to convergence already with a slower decrease $\sim 1/\log(N)$. For these scalings of the filter width, according to (21), the width of the filtered state will scale in the same way.

2. Period of cosine filter α

Different from the Gaussian one, the cosine filter (8) is periodic, but it remains a good approximation of the former when the argument is bounded within one period. More concretely, operators $\hat{P}_\delta(E)$ and $\hat{F}_{\delta, \alpha}(E)$ are close to each other when $\|\hat{H} - E\| \leq \alpha\pi/2$. At the same time, because the number of terms that need to be evaluated in the sum (9) is proportional to α , it is convenient to choose the smallest possible value that ensures the previous property. For a local Hamiltonian, a value $\alpha \propto N$ is enough for the condition to hold for all values of E within the energy spectrum. If the operator acts only on a limited energy window, a smaller value of α can be chosen, as long as all relevant states are almost supported in $[E - \alpha\pi/2, E + \alpha\pi/2]$. This can be used, for instance, when the filter acts on a product state, whose energy distribution is approximately Gaussian, with support on an energy interval $\propto \sqrt{N}$. If additionally the filter is centered near the mean energy of the state and $\delta = o(\sqrt{N})$, it is enough to choose $\alpha \propto \sqrt{N}$ [10]. In practice, we find $\alpha = 3 \max(\sigma_\phi \sqrt{N}, \delta)$ to work well for all system sizes.

3. Monte Carlo cutoff threshold ϵ

For the discussions in Sec. III B 2, it should be ensured that the samples in Monte Carlo simulations not step into other energy periods of the cosine filter when $\alpha \propto \sqrt{N}$. In other words, the weights of the states whose mean energy are close to the edges of $[E - \alpha\pi/2, E + \alpha\pi/2]$ should be small enough. A cutoff threshold ϵ can be applied to the weights of samples to improve numerical stability in Monte Carlo simulations. To be more concrete, a proposed state $|\phi\rangle$ will be directly discarded (its probability assimilated to 0) if $D_{\delta, \phi} < \epsilon$. Besides making the numerics more stable, the presence of the cutoff prevents the peak of the DOS from shifting the ensemble when we target energies near the edges, because it restricts the visited energy range in more general cases, as we show next.

If we consider a product state basis, an individual state $|\phi\rangle$ will have mean energy E_ϕ and width $\sigma_\phi \sqrt{N}$. Again, assuming a Gaussian distribution, we can estimate its weight in the sum as

$$D_{\delta, \phi}(E) = \int d\mu \frac{1}{\sqrt{2\pi N \sigma_\phi}} \exp \left[-\frac{(\mu - E_\phi)^2}{2\sigma_\phi^2 N} - \frac{(\mu - E)^2}{2\delta^2} \right] \quad (22)$$

$$= \sqrt{\frac{\delta^2}{\delta^2 + N\sigma_\phi^2}} \exp[-(E - E_\phi)^2/2(\delta^2 + N\sigma_\phi^2)] \quad (23)$$

TABLE I. Possible choices of the filter parameters (δ, α) that ensure approaching the microcanonical values in the thermodynamic limit for the various methods, and corresponding maximum evolution time and number of steps.

	Method	δ	α	t_{\max}	R	Applicable energy range
(I)	Spectrum filtering, direct trace	$\mathcal{O}(\sqrt{N})$	N	$\omega(1/\sqrt{N})$	$\omega(\sqrt{N})$	$\mathcal{O}(\sqrt{N})$
(IIa)	Spectrum filtering, Monte Carlo	$\mathcal{O}(\sqrt{N})$	N	$\Omega(1/\sqrt{N})$	$\Omega(\sqrt{N})$	Full spectrum
(IIb)	Spectrum filtering, Monte Carlo	$\mathcal{O}(1)$	$\Omega(\sqrt{N})$	$\Omega(1)$	$\Omega(\sqrt{N})$	Full spectrum
(III)	State filtering	$\text{poly}(1/N)$	$\Omega(\sqrt{N})$	$\text{poly}(N)$	$\text{poly}(N)$	Full spectrum

so that $D_{\delta, \phi} < \epsilon$ holds if

$$|E - E_{\phi}| > \nu_{\phi} := \left[(\delta^2 + N\sigma_{\phi}^2) \ln \left(\frac{\delta}{\epsilon \sqrt{\delta^2 + N\sigma_{\phi}^2}} \right) \right]^{1/2}. \quad (24)$$

When $\delta = \Omega(\sqrt{N})$, it follows that $\nu_{\phi} \sim \delta$ for all states in the basis, while if $\delta = \mathcal{O}(\sqrt{N})$, $\nu_{\phi} = \mathcal{O}(\sqrt{N})$ holds as well. Hence for any $\delta = \mathcal{O}(N)$, the cutoff ϵ itself can restrict the sampling space within an energy interval of width $\mathcal{O}(N)$ that screens the peak of the DOS.

4. Choosing the parameters for microcanonical values

According to the discussions above, we can summarize in Table I some possible choices of parameters for the various algorithms, such that we obtain convergence to thermal values in the thermodynamic limit. The table shows the scaling with system size of α and δ , as well as the resulting cost (in terms of maximum evolution time and number of evolutions to run) and the energy range where the methods are applicable. The fastest method [(I) in the table] is directly computing (7) by taking the traces of the evolution operators approximated as MPOs, but, as discussed in Sec. III A, it is only applicable in an energy interval of width proportional to \sqrt{N} around the center of the spectrum. With Monte Carlo sampling (II), in contrast, it is possible to reach the whole spectrum. The filter width required to obtain convergence to thermal values in the thermodynamic limit should scale at most as $\mathcal{O}(\sqrt{N})$. A larger width corresponds to a shorter evolution time t_{\max} , and hence a smaller bond dimension required for the MPS or MPO, but it also shows slower convergence to the thermal values as the system size is increased. We thus show two possible choices (IIa) and (IIb), both of which we explore numerically, using different approaches for time evolution, in Sec. IV A. In (IIa), a cutoff threshold is applied in the Monte Carlo simulations to avoid the energy shift due to DOS when $\delta \propto \sqrt{N}$. Finally, when filtering a state (III), the time t_{\max} required to approach microcanonical values is polynomial in the system size, which means that with TN techniques we will be able to extract microcanonical values with this method only for small system sizes. For this last method, however, the achievable width can be optimized by applying the algorithm on a state with reduced energy width. We present a way to implement this improvement by using MPSs obtained after a variational minimization of the variance.

IV. RESULTS

To demonstrate and benchmark the various methods described above, we apply them to a quantum Ising chain with open boundary conditions,

$$\hat{H}_{\text{Ising}} = J \sum_{i=1}^{N-1} \sigma_i^z \sigma_{i+1}^z + \sum_{i=1}^N (g\sigma_i^x + h\sigma_i^z). \quad (25)$$

The model is integrable if either $g = 0$ or $h = 0$. Here we choose a particular set of parameters $(J, g, h) = (1, -1.05, 0.5)$ far from integrability [44]. In the thermodynamic limit, the corresponding energy density lies in the interval $E/N \in [-1.33, 1.72]$. For the observable, we focus on the average magnetization

$$\hat{m}_z = \sum_{i=1}^N \sigma_i^z / N. \quad (26)$$

A. Filter ensemble

We start by illustrating the performance of the Monte Carlo algorithm to estimate expectation values in the filter ensemble (3) at all values of energy. Since the largest time we need to simulate is $t_{\max} \propto 1/\delta$, the classical simulation can efficiently treat widths $\delta = \mathcal{O}(1/\log N)$ and larger.

For the numerical benchmarking, we choose widths $\delta \propto \sqrt{N}$ and $\delta = \text{constant}$, which, according to the discussions above, are enough to approach the microcanonical values in the thermodynamical limit [see (II) in Table I]. In the non-integrable model we consider, the values are thus expected to converge to the thermal ones. Thus, we can compare the results of the algorithm with the exact values in thermal equilibrium at the corresponding energies, which we can compute independently using standard TN techniques [6].

The calculations can be done using different options for the TN evolution (Sec. III A). As long as the results are converged in bond dimension, both approximating the evolution operators as MPOs or the individual evolved states as MPSs are valid strategies, and we show results obtained with both of them.

1. MPO version of Monte Carlo simulation

Figure 2 demonstrates the success of the method to find expectation values in the filter ensemble, for system sizes up to $N = 80$. In particular, for this plot, we chose to simulate and store the MPOs for all evolution operators before realizing the sampling over the computational basis. For the filter, we used filter parameters $(\delta, \alpha) \propto (\sqrt{N}, N)$, which, as argued in Sec. III B, in a generic case is enough for the observable to

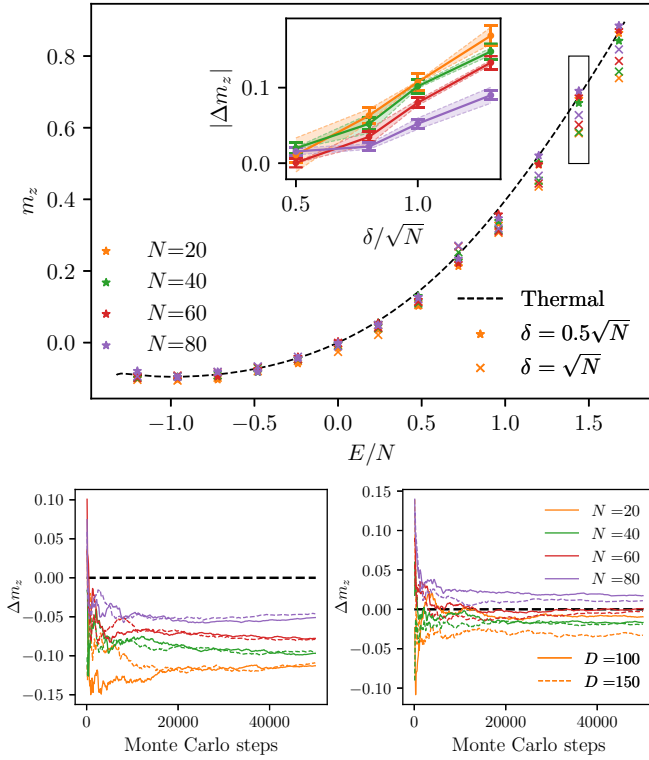


FIG. 2. Magnetization (26) in the nonintegrable Ising chain (25) with $(J, g, h) = (1, -1.05, 0.5)$, computed by the MPO version of the Monte Carlo method with parameters $(\delta \propto \sqrt{N}, \alpha = \sqrt{J^2 + g^2 + h^2 N})$ and a cutoff (see main text). The sample size is 5×10^4 for each energy and δ . Upper panel: Convergence over the whole spectral range. The black dashed line is the thermal value for $N = 80$. Data points are $(m_z)_\delta$, where different colors stand for different system sizes and different shapes for different δ . The inset plots the difference with respect to the thermal value as a function of δ/\sqrt{N} at the point $E/N = 1.44$ (indicated by a box in the main plot), with error bars indicating the standard deviation of the Monte Carlo sampling, while the shadowed region represents errors from finite MPO bond dimensions (the differences between results of $D = 100$ and $D = 150$). Lower panels: Convergence with the sample size at $E/N = 1.44$ for $\delta = \sqrt{N}$ (left) and $\delta = 0.5\sqrt{N}$ (right). $\Delta m_z = (m_z)_\delta - m_z^{\text{thermal}}$.

converge, in the thermodynamic limit, to the thermal expectation value if introducing a cutoff threshold (Sec. III B 3). We choose to sample over the computational basis and the cutoff threshold $\epsilon = 10^{-4} D_{\delta, \phi_0}(E)$, where $|\phi_0\rangle$ is the initial state in the Monte Carlo simulation, obtained by minimizing $\langle \phi | (\hat{H} - E)^2 | \phi \rangle$ for $|\phi\rangle$ in the basis set.

As shown in the upper panel of Fig. 2, the results clearly converge to the thermal value as the system size N is increased or δ is reduced. The inset shows explicitly this convergence for energy density $E/N = 1.44$, relatively close to the edge of the spectrum. Errors have two main sources, which are shown in this plot: the statistical error from the Monte Carlo sampling (error bars) and the truncation error from the finite bond dimension of the MPO (shown as shadowed region). Only for the largest system size $N = 80$ and smallest width $\delta = 0.5\sqrt{N}$ do we observe a small discrepancy, but it is compatible with our estimated errors from both sources.

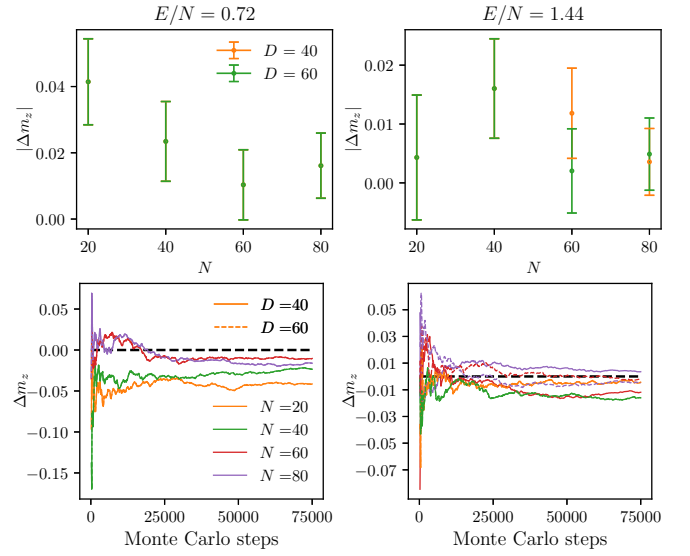


FIG. 3. Magnetization (26) obtained by the MPS version of the Monte Carlo method for $(\delta = 1, \alpha = 6\sqrt{N})$ and no cutoff, for two values of the energy density, $E/N = 0.42$ (left) and $E/N = 1.44$ (right). Upper panels: Difference between $(m_z)_\delta(E)$ and thermal value as a function of the system size for bond dimensions $D = 40, 60$. The error bars correspond to a standard deviation of Monte Carlo samplings. Lower panels: Convergence of Monte Carlo sampling with the number of steps. The results have converged for bond dimension $D = 40$ in the left plots, and hence the solid and dashed lines are on top of each other.

The lower panels of Fig. 2 show explicitly the convergence of the Monte Carlo sampling at the same energy density $E/N = 1.44$, for various system sizes and bond dimensions, and for two different values of the width. In all these cases we observe that, after 50 000 steps, the results are practically converged, even though some fluctuations can be appreciated.

2. MPS version of Monte Carlo simulation

To illustrate the performance of the algorithm when individual states, rather than operators, are evolved, we choose a narrower filter width, which should lead to values closer to the thermal ones, while still being reachable by classical simulations. Notice, however, that similar values could have been obtained with the MPO option, at a different cost in memory and time.

In particular, considering a constant value $\delta = \mathcal{O}(1)$, independent of system size, requires evolution until a constant time, which (for not too small values of δ) can be efficiently simulated using TN. Thus we choose parameters $(\delta, \alpha) \propto (1, \sqrt{N})$ [(IIB) in Table I], and explore two values of the energy density, one near the center of the spectrum ($E/N = 0.72$) and one close to the edge ($E/N = 1.44$), and both within the reach of product states from the computational basis, which we take again as our sampling basis.

Results are shown in Fig. 3. In the upper panels, we plot the difference between $(m_z)_\delta$ and the corresponding thermal value as a function of the system size, for two values of the bond dimension, with error bars indicating the statistical error. We

find that, within error bars, the distance to the thermal value decreases as the system size grows, with a relative difference smaller than 0.5% for $N = 80$ and $E/N = 1.44$.

In the lower panels of Fig. 3, we again show explicitly the convergence of the Monte Carlo sampling. Note that here the same seed for randomization was used for different bond dimensions, so the fact that the solid and dashed lines (representing $D = 40$ and $D = 60$, respectively) are on top of each other indicates the convergence with regard to bond dimension already at $D = 40$.

As illustrated above, the combination of short-time dynamics simulation and sampling provides a powerful method to compute the expectation values in the filter ensemble, as long as the basis contains vectors with substantial weight in the energy region of interest. For the computational basis that we have used in the examples, mean energies lie in the interval $E/N \in [-1, 1.5]$. According to (23), the weights of basis states $D_{\delta,\phi}$ will decay exponentially with the system size for a fixed energy density E/N outside this interval. Figure 2 shows that, for system size $N = 80$, the Monte Carlo sampling with computational basis remains valid at $E/N = 1.68$ (rightmost point) and $E/N = -1.2$ (leftmost point), much closer to the edges of the spectrum, so that we do not need to resort to the Pauli basis mentioned in Sec. II B. Using this basis may however become necessary as we keep increasing the system size, or if we consider other models or higher dimensions.

3. Exploring the center of the spectrum without sampling

The fastest alternative to evaluate (7) with TN simulations is to directly evaluate numerator and denominator from traces of the evolution MPOs, without the sampling iteration [(I) in Table I]. As discussed above, this is only feasible in the central region of the spectrum, over a width $\propto \sqrt{N}$, before the density of states becomes exponentially small. If the filter width is not much smaller than this scale, the mean energy of the filter ensemble will be effectively shifted toward the maximum of the DOS, as explicitly computed in Sec. III B 1.

Figure 4 illustrates the behavior of this alternative for $(\delta, \alpha) \propto (\sqrt{N}, N)$. The upper plot shows the results obtained for the magnetization m_z for various system sizes and filter widths, as a function of the energy density corresponding to the center of the filter. Because of the shift discussed above, the results do not converge to the thermal ones (indicated by the solid line) at that energy, but at a shifted value according to (18) (indicated by a dashed line for each δ).

We observe that, while in the central part of the spectrum better convergence is observed as δ decreases or N increases, near the edge of the spectrum the method fails to give the correct microcanonical values, especially for larger systems and smaller widths, as then it becomes sensitive to the exponentially smaller density of states. This is also visible in the lower panel of the figure, where we plot the difference between the computed values $O_\delta(E)$ and the thermal ones at the corresponding shifted energies. According to (17), the denominator of (7), $\text{tr}[\hat{P}_\delta(E)]$, should scale as $\exp(-E^2/2\gamma N\sigma_0^2)$, indicating the range of energy densities for which the method actually converges shrinks as $1/\sqrt{N}$.




FIG. 4. Magnetization (26) obtained by MPO simulation of the evolution operators and direct evaluation of the ratio (7). Upper panel: Computed magnetization as a function of the energy density E/N at the center of the filter for system sizes up to $N = 80$ and several widths. Over most of the spectral range the results converge to thermal values at the shifted energy depending on δ (18) (dashed lines) instead of the thermal value at E/N . At the edges of the spectrum this fails, due to the exponentially decaying DOS in these regions, as discussed in Sec. III B 1. Lower panel: Difference between the computed magnetization and the shifted thermal values. All results were obtained using bond dimension $D = 200$ for the MPOs.

B. Filtered pure state

As discussed in Sec. III B, we can apply the filter on a state to decrease its energy variance and, in the generic case, obtain convergence to the microcanonical properties. Also in this case, the largest time that needs to be simulated is $t_{\max} = 2x/\delta$, which can be done efficiently by TN simulations if the width is at least $\mathcal{O}(1/\log N)$. But, in contrast to the calculations for the filter ensemble, diagonal in the energy basis, a much smaller δ is required in this case to approach microcanonical values in the thermodynamic limit. More concretely, $\delta = \mathcal{O}(\text{poly}(1/N))$ should be enough, but this requires $t_{\max} = \Omega(\text{poly}(N))$, and thus a bond dimension that increases




FIG. 5. Results computed by filtering states. $E/N = 0.72$. The filtered states are MPSs of bond dimensions 1, 2, 5, and 10. We compare the results of different system sizes with $\delta \propto 1$ in the left figure and $\delta \propto 1/N^2$ in the right.

exponentially with N . Additionally, when filtering a product state, the total number of terms to probe will be $(2R + 1)^2$, where $R = x\alpha/\delta$, which grows at least as N^3 .

One way to mitigate the second problem is to apply the filter on a state with already reduced energy variance. This allows us to choose a smaller period α and correspondingly keep a more moderate value of R and to test the strategy for moderate sizes. We implement this strategy using as initial states MPSs with a given bond dimension, found by variationally minimizing $(\hat{H} - E)^2$ at the value of E we are interested in. To test this strategy, we have targeted a value $E/N = 0.72$ for system sizes $20 \leq N \leq 80$, and obtained initial MPSs with reduced widths from the minimization of $(H - E)^2$ with bond dimensions $D_0 \in \{1, 2, 5, 10\}$. For each one of these states, we compute the width σ_D , and then apply a filter with parameters $\delta = \sigma_D/2\sqrt{N}$ and $\alpha = 3\sigma_D$. We show the results in Fig. 5. To analyze the convergence toward the thermal value as the width decreases, we plot the relative error of the magnetization with respect to the thermal one as a function of $1/\delta$ (left panel) and $1/(N^2\delta)$ (right panel) for each system size. The first case shows no clear scaling laws, which indicates that $\delta = \mathcal{O}(1)$ is not enough for convergence in the thermodynamic limit. The right panel, instead, exhibits a trend to convergence for $\delta = \mathcal{O}(1/N^2)$, even though numerically it becomes difficult to reach the lower right corner for large systems.

V. SUMMARY AND DISCUSSION

We have presented a quantum-inspired classical method, based on a TN simulation of the quantum-assisted Monte Carlo algorithm proposed in [10], that allows us to compute microcanonical and diagonal values for quantum many-body systems. Our method estimates broadened spectral functions, which takes the form of the trace of an energy filter operator, or a product of the latter with an observable, via sampling over time-evolved product states. Because the longest required time is proportional to the inverse filter width, which is very short and easy to simulate with tensor networks, it allows us to find expectation values in a diagonal ensemble which would only be reached after a much longer time evolution from an initial product state. While filter widths $\mathcal{O}(\sqrt{N})$ are enough to find the diagonal ensemble values of generic product states, we can also reach energy filters of constant widths, as these

only require $\mathcal{O}(1)$ evolution times. These scalings are enough to obtain convergence to thermal equilibrium in the thermodynamic limit, in the generic case.

We have benchmarked the algorithm on the nonintegrable Ising chain, for sizes up to $N = 80$ sites (far beyond the reach of exact diagonalization), and we have checked different choices of the parameters that affect the efficiency and applicability of the method. In particular, we explicitly show diagonal expectation values for Gaussian ensembles of width $\mathcal{O}(\sqrt{N})$, obtained with low computational cost over the whole range of energies, and observe their convergence toward the thermal equilibrium values. Reducing further the width, we obtain microcanonical expectation values with high precision.

We can also classically simulate the provably efficient quantum algorithm in [10], in which the filter is applied on a fixed initial state. If we want to use this method to explore the microcanonical properties, nevertheless, the filter width needs to decrease with the system size in order to guarantee convergence in the thermodynamic limit, which results in increasing times and an exponentially growing bond dimension. We have however optimized the procedure by choosing as initial states MPSs with minimal variance. In this way, we can run the algorithm and observe convergence for reasonably sized systems.

The results shown in this paper demonstrate the potential of the algorithm. Accessing the microcanonical values would be helpful to investigate all sorts of out-of-equilibrium quantum many-body behavior, for instance many-body localization [45] and quantum scars [46], for large systems. Although we have only implemented it for a translationally invariant spin chain with short-range interactions, the method can be easily extended to systems with disorder, or long-range interactions, and also to bosonic or fermionic systems. We have recently learned that Schuckert *et al.* are using a related method to study long-range interacting models [47].

In principle, the same method can be also extended to higher-dimensional systems, by using projected entangled pair states (PEPSs) [48] or mapping the system to adopt more restricted *Ansätze* such as MPSs and other subfamilies of PEPSs, which is a promising possibility, given the absence of numerical methods for extracting microcanonical expectation values in that case, beyond exact diagonalization. The numerical challenge is then higher, due to the larger computational cost of the corresponding algorithms, and determining the preferable implementation option should be analyzed. Also further extensions are possible that consider other filter functions or combinations of filters.

ACKNOWLEDGMENTS

This work was partly supported by the Deutsche Forschungsgemeinschaft (DFG; German Research Foundation) under Germany's Excellence Strategy—EXC-2111—3908148688, by the European Union through the ERC grant QUENOCOBA, ERC-2016-ADG (Grant No. 742102), and by the German Federal Ministry of Education and Research (BMBF) through the funded project EQUAHUMO (Grant No. 13N16066) within the funding program quantum technologies—from basic research to market, in association to the Munich Quantum Valley.

- [1] A. W. Sandvik and J. Kurkijärvi, *Phys. Rev. B* **43**, 5950 (1991).
- [2] F. Verstraete, J. J. García-Ripoll, and J. I. Cirac, *Phys. Rev. Lett.* **93**, 207204 (2004).
- [3] M. Rigol, T. Bryant, and R. R. P. Singh, *Phys. Rev. Lett.* **97**, 187202 (2006).
- [4] M. B. Hastings, *Phys. Rev. B* **76**, 201102(R) (2007).
- [5] S. R. White, *Phys. Rev. Lett.* **102**, 190601 (2009).
- [6] E. M. Stoudenmire and S. R. White, *New J. Phys.* **12**, 055026 (2010).
- [7] A. Molnar, N. Schuch, F. Verstraete, and J. I. Cirac, *Phys. Rev. B* **91**, 045138 (2015).
- [8] F. Schrodi, P. Silvi, F. Tschirsich, R. Fazio, and S. Montangero, *Phys. Rev. B* **96**, 094303 (2017).
- [9] Y. Yang, S. Iblisdir, J. I. Cirac, and M. C. Bañuls, *Phys. Rev. Lett.* **124**, 100602 (2020).
- [10] S. Lu, M. C. Bañuls, and J. I. Cirac, *PRX Quantum* **2**, 020321 (2021).
- [11] M. Rigol, V. Dunjko, and M. Olshanii, *Nature (London)* **452**, 854 (2008).
- [12] J. M. Deutsch, *Phys. Rev. A* **43**, 2046 (1991).
- [13] M. Srednicki, *Phys. Rev. E* **50**, 888 (1994).
- [14] L. D'Alessio, Y. Kafri, A. Polkovnikov, and M. Rigol, *Adv. Phys.* **65**, 239 (2016).
- [15] S. Lloyd, *Science* **273**, 1073 (1996).
- [16] J. Haah, M. Hastings, R. Kothari, and G. H. Low, in *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, Vol. 2018-October (IEEE, 2018), pp. 350–360.
- [17] P. Calabrese and J. Cardy, *J. Stat. Mech.: Theory Exp.* (2005) P04010.
- [18] Z. Lan, M. van Horssen, S. Powell, and J. P. Garrahan, *Phys. Rev. Lett.* **121**, 040603 (2018).
- [19] C.-J. Lin, A. Chandran, and O. I. Motrunich, *Phys. Rev. Research* **2**, 033044 (2020).
- [20] A. Peres, *Phys. Rev. A* **30**, 1610 (1984).
- [21] R. A. Jalabert and H. M. Pastawski, *Phys. Rev. Lett.* **86**, 2490 (2001).
- [22] F. Verstraete, V. Murg, and J. I. Cirac, *Adv. Phys.* **57**, 143 (2008).
- [23] U. Schollwöck, *Ann. Phys.* **326**, 96 (2011).
- [24] T. Huckle, K. Waldherr, and T. Schulte-Herbrüggen, *Linear Algebra and its Applications* **438**, 750 (2013).
- [25] R. Orús, *Ann. Phys.* **349**, 117 (2014).
- [26] P. Silvi, F. Tschirsich, M. Gerster, J. Jünemann, D. Jaschke, M. Rizzi, and S. Montangero, *SciPost Physics Lecture Notes* **8**, 1 (2019).
- [27] K. Okunishi, T. Nishino, and H. Ueda, *J. Phys. Soc. Jpn.* **91**, 062001 (2022).
- [28] S. Paeckel, T. Köhler, A. Swoboda, S. R. Manmana, U. Schollwöck, and C. Hubig, *Ann. Phys.* **411**, 167998 (2019).
- [29] M. Hartmann, G. Mahler, and O. Hess, *J. Stat. Phys.* **119**, 1139 (2005).
- [30] J. P. Keating, N. Linden, and H. J. Wells, *Commun. Math. Phys.* **338**, 81 (2015).
- [31] W. Beugeling, R. Moessner, and M. Haque, *Phys. Rev. E* **91**, 012144 (2015).
- [32] D. J. Luitz and Y. Bar Lev, *Phys. Rev. Lett.* **117**, 170404 (2016).
- [33] R. Mondaini and M. Rigol, *Phys. Rev. E* **96**, 012157 (2017).
- [34] A. Dymarsky and H. Liu, *Phys. Rev. E* **99**, 010102(R) (2019).
- [35] M. C. Bañuls, D. A. Huse, and J. I. Cirac, *Phys. Rev. B* **101**, 144305 (2020).
- [36] Y. Ge, J. Tura, and J. I. Cirac, *J. Math. Phys.* **60**, 022202 (2019).
- [37] The approximation actually works beyond the range; see Appendix E of Ref. [10].
- [38] E. H. Lieb, *Commun. Math. Phys.* **31**, 327 (1973).
- [39] S. P. Lim and D. N. Sheng, *Phys. Rev. B* **94**, 045111 (2016).
- [40] X. Yu, D. Pekker, and B. K. Clark, *Phys. Rev. Lett.* **118**, 017201 (2017).
- [41] The asymptotic notations Ω , ω , \mathcal{O} , and o are used in this paper, which represent to be bounded from below, to dominate, to be bounded from above, and to be dominated.
- [42] The memory requirement can be reduced by storing only a constant fraction of the MPOs for all time steps, and combining several of them to compute the required matrix elements.
- [43] M. Hartmann, G. Mahler, and O. Hess, *Lett. Math. Phys.* **68**, 103 (2004).
- [44] H. Kim and D. A. Huse, *Phys. Rev. Lett.* **111**, 127205 (2013).
- [45] D. A. Abanin, E. Altman, I. Bloch, and M. Serbyn, *Rev. Mod. Phys.* **91**, 021001 (2019).
- [46] C. J. Turner, A. A. Michailidis, D. A. Abanin, M. Serbyn, and Z. Papić, *Nat. Phys.* **14**, 745 (2018).
- [47] A. Schuckert, A. Bohrdt, E. Crane, and M. Knap, [arXiv:2206.01756](https://arxiv.org/abs/2206.01756).
- [48] F. Verstraete and J. I. Cirac, [arXiv:cond-mat/0407066](https://arxiv.org/abs/cond-mat/0407066).