

Parallel time-dependent variational principle algorithm for matrix product states

Paul Secular^{1,*}, Nikita Gourianov², Michael Lubasch², Sergey Dolgov³, Stephen R. Clark^{4,5} and Dieter Jaksch^{2,6}

¹*Department of Physics, University of Bath, Claverton Down, Bath BA2 7AY, United Kingdom*

²*Clarendon Laboratory, Department of Physics, University of Oxford, Oxford OX1 3PU, United Kingdom*

³*Department of Mathematical Sciences, University of Bath, Claverton Down, Bath BA2 7AY, United Kingdom*

⁴*H.H. Wills Physics Laboratory, University of Bristol, Bristol BS8 1TL, United Kingdom*

⁵*Max Planck Institute for the Structure and Dynamics of Matter, CFEL, 22761 Hamburg, Germany*

⁶*Centre for Quantum Technologies, National University of Singapore, 3 Science Drive 2, Singapore 117543, Singapore*



(Received 10 January 2020; revised manuscript received 30 March 2020; accepted 2 April 2020; published 5 June 2020)

Combining the time-dependent variational principle (TDVP) algorithm with the parallelization scheme introduced by Stoudenmire and White for the density matrix renormalization group (DMRG), we present the first parallel matrix product state (MPS) algorithm capable of time evolving one-dimensional (1D) quantum lattice systems with long-range interactions. We benchmark the accuracy and performance of the algorithm by simulating quenches in the long-range Ising and XY models. We show that our code scales well up to 32 processes, with parallel efficiencies as high as 86%. Finally, we calculate the dynamical correlation function of a 201-site Heisenberg XXX spin chain with $1/r^2$ interactions, which is challenging to compute sequentially. These results pave the way for the application of tensor networks to increasingly complex many-body systems.

DOI: [10.1103/PhysRevB.101.235123](https://doi.org/10.1103/PhysRevB.101.235123)

I. INTRODUCTION

Although classical simulation of the quantum many-body problem is in general exponentially hard, many physically interesting, slightly entangled states can be efficiently simulated using tensor network methods [1–3]; in particular, those obeying an entanglement “area law” [4,5]. The most common approach for one-dimensional (1D) and quasi-two-dimensional (2D) systems is the matrix product state (MPS) ansatz. This forms the basis of the modern formulation [6–8] of White’s density matrix renormalization group (DMRG) algorithm [9] for computing ground states and low-lying excited states [10,11].

The introduction of the MPS-based time-evolving block decimation (TEBD) algorithm [1,2,12] (and subsequent developments) allowed the short-time dynamics of local 1D quantum lattice models to be simulated with great success. In recent years, however, there has been a renewed interest in the dynamics of models with long-range interactions, driven by experimental advances in atomic, molecular, and optical physics. Interactions that decay as $1/r^\alpha$ are now realized in experiments with polar molecules ($\alpha = 3$) [13,14] and Rydberg atoms ($\alpha = 6$) [15–17], whilst trapped ion experiments can simulate spin models with a tunable exponent ($0 \lesssim \alpha \lesssim 3$) [18–24].

In response, MPS algorithms have been developed that are able to simulate the time evolution of such models classically [25–38]. One of the most promising approaches is the time-dependent variational principle (TDVP) [31,33–35], which

has found widespread use in condensed matter physics,¹ and which is beginning to find applications in quantum chemistry [53–60].

TDVP is a serial algorithm that uses sequential sweeps for numerical stability. This means it cannot easily take advantage of multicore architectures or high-performance computing clusters. In fact the same is true of most MPS algorithms. Attempts to address this shortcoming include the use of parallel linear algebra operations [60–63], parallelization over

¹TDVP is applied to $1/r^\alpha$ models in Refs. [35,39–52].

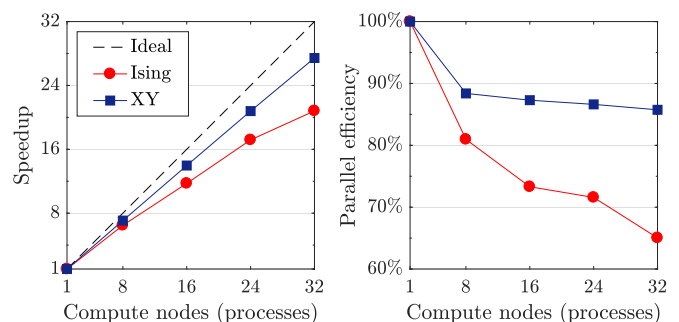


FIG. 1. Best and worst case strong scaling of the parallel time-dependent variational principle algorithm for our benchmark examples. “Ising” is the 129-site spin chain evolved under the $1/r^{2.3}$ transverse-field Ising Hamiltonian described in Sec. IV A. “XY” is the 101-site spin chain evolved under the $1/r^{0.75}$ XY Hamiltonian described in Sec. IV B. The benchmark described in Sec. IV C is excluded as it was unfeasible to run on a single process.

*paul@secular.me.uk; <http://secular.me.uk/>.

quantum number blocks [61,64], and parallelization over terms in the Hamiltonian [65]. More generally, it is desirable to parallelize over different parts of an MPS, but this network-level parallelization is nontrivial. To the best of our knowledge, the only such algorithms to date are nearest-neighbor TEBD [66–69], real-space parallel DMRG [70,71], and parallel infinite DMRG [72], although network-level parallelization has also been proposed for projected entangled pair states (PEPS) [73,74]. A parallel time evolution method for MPS capable of handling long-range interactions has remained an open problem.

Given the close relationship between TDVP and DMRG established in Refs. [34,35], it is natural to ask whether a parallel version of TDVP can be developed in a similar manner to real-space parallel DMRG. In this work, we demonstrate that it can. The parallelization of the algorithm allows state-of-the-art calculations to be sped up by a factor of 20+ (see Fig. 1). Moreover, we show how it makes larger calculations possible that may otherwise be unfeasible. For example, in Sec. IV C, we calculate the dynamical spin-spin correlation function of a 201-site long-range Heisenberg model in a matter of days, rather than weeks.

The rest of the paper is organized as follows. We start with a background section, before moving on to introduce the parallel TDVP algorithm in Sec. III. We provide increasingly complex benchmark examples for quantum spin-chains in Sec. IV, and finally conclude and suggest directions for future research in Sec. V.

II. BACKGROUND

For completeness, we start by reviewing relevant background material, covering matrix product states, the inverse canonical gauge, matrix product operators, and the time dependent variational principle. We also establish the notation used in the rest of the paper.

A. Matrix product states

Any finite-dimensional, N -partite quantum state can be expressed in a given basis $|\sigma_1\sigma_2\dots\sigma_N\rangle$ as

$$|\psi\rangle = \sum_{\sigma_1\dots\sigma_N=1}^{d_1\dots d_N} M_1^{\sigma_1} M_2^{\sigma_2} \dots M_N^{\sigma_N} |\sigma_1\sigma_2\dots\sigma_N\rangle, \quad (1)$$

where the $M_j^{\sigma_j}$ are matrices ($M_1^{\sigma_1}$ and $M_N^{\sigma_N}$ being row and column vectors, respectively). This decomposition is known as a tensor train [75] or matrix product state (MPS) [76]. To avoid index gymnastics we will often use the graphical tensor notation introduced by Penrose [77] to represent MPS as tensor networks (see, e.g., Ref. [78]). In this notation, tensors are displayed as nodes in a network, with their connected edges representing pairs of dummy indices in the usual Einstein notation. A contractible edge is referred to as a bond and the dimension of the corresponding dummy index is the bond dimension. Free indices (which may represent physical degrees of freedom) are shown as disconnected edges. In this

notation, Eq. (1) is written as

$$|\psi\rangle = \begin{array}{c} M_1 \quad \quad M_2 \quad \quad M_3 \quad \quad \dots \quad M_N \\ \circ \quad \quad \circ \quad \quad \circ \quad \quad \dots \quad \circ \\ d_1 \quad \quad d_2 \quad \quad d_3 \quad \quad \dots \quad d_N \end{array} \begin{array}{c} \chi_1 \\ \chi_2 \\ \chi_3 \\ \dots \\ \chi_{N-1} \end{array}, \quad (2)$$

where we have explicitly labeled the dimensions of all bonds χ_j , and physical edges d_j (for notational clarity, dimensions are suppressed in the rest of the paper).

Here we are interested in MPS in the context of finite 1D lattice models with open boundaries. In such a model there is a one-to-one correspondence between the lattice sites and the MPS “site tensors” M_j . In quantum chemistry, the “sites” may instead be molecular orbitals [79]. In a lattice model, the physical dimensions d_j will often be independent of j . For the spin-half chains considered in this paper, $d_j = d = 2$. The bonds between sites capture the entanglement present. In general, χ_j is not independent of j . It is thus convenient to define $\chi = \max(\chi_j)$. Exactly representing an arbitrary state as an MPS requires a $\chi = \chi_{\text{exact}}$ that is exponential in the number of sites. However, physical states are often well approximated by MPS of lower bond dimension. We denote by χ_{max} the maximum bond dimension chosen for a calculation, where typically $\chi_{\text{max}} \ll \chi_{\text{exact}}$.

Even with all χ_j fixed, an MPS representation is not unique, since

$$\dots M_j^{\sigma_j} M_{j+1}^{\sigma_{j+1}} \dots = \dots \tilde{M}_j^{\sigma_j} \tilde{M}_{j+1}^{\sigma_{j+1}} \dots,$$

where $\tilde{M}_j^{\sigma_j} = M_j^{\sigma_j} X$, and $\tilde{M}_{j+1}^{\sigma_{j+1}} = X^{-1} M_{j+1}^{\sigma_{j+1}}$ for any invertible matrix X . This “gauge freedom” can be exploited by algorithms for numerical stability. In particular, a canonical form is usually employed in which one or more tensors act as orthogonality centers [8]. The parts of the MPS to the left and right of an orthogonality center form orthonormal basis states, respectively $|\Phi_{L,k}\rangle$ and $|\Phi_{R,l}\rangle$, where k and l are the indices of the states. The orthogonality center thus defines the quantum state with respect to these bases.

Bipartitioning a system into sites $[1 : j]$ and $[j + 1 : N]$ allows us to write the Schmidt decomposition,

$$|\psi\rangle = \sum_{k=1}^{\chi_j} \lambda_k |\Phi_{L,k}^{[1:j]}\rangle \otimes |\Phi_{R,k}^{[j+1:N]}\rangle. \quad (3)$$

This can be expressed as an MPS with orthogonality center $\Lambda_j = \text{diag}(\lambda_1 \dots \lambda_{\chi_j})$ [8],

$$|\psi\rangle = \sum_{\sigma_1 \dots \sigma_N} A_1^{\sigma_1} \dots A_j^{\sigma_j} \Lambda_j B_{j+1}^{\sigma_{j+1}} \dots B_N^{\sigma_N} |\sigma_1 \dots \sigma_N\rangle, \quad (4)$$

where $A_p^{\sigma_p}$ ($1 \leq p \leq j$) and $B_q^{\sigma_q}$ ($j + 1 \leq q \leq N$) satisfy

$$\sum_{\sigma_p} (A_p^{\sigma_p})^\dagger A_p^{\sigma_p} = \sum_{\sigma_q} B_q^{\sigma_q} (B_q^{\sigma_q})^\dagger = \mathbb{1}. \quad (5)$$

The orthogonality center can alternatively be made into a site tensor. For example, setting $\Psi_j^{\sigma_j} = A_j^{\sigma_j} \Lambda_j$ gives an MPS with orthogonality center at site j ,

$$|\psi\rangle = \sum_{\sigma_1 \dots \sigma_N} A_1^{\sigma_1} \dots A_{j-1}^{\sigma_{j-1}} \Psi_j^{\sigma_j} B_{j+1}^{\sigma_{j+1}} \dots B_N^{\sigma_N} |\sigma_1 \dots \sigma_N\rangle. \quad (6)$$

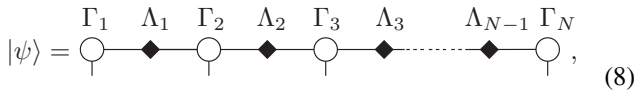
Eqs. (4) and (6) are known as the mixed canonical form [8]. This MPS gauge is useful for serial algorithms that update tensors sequentially, as the orthogonality center can be shifted by one site whilst maintaining the orthonormality of the state. For details of this approach we refer the reader to Ref. [8].

B. Inverse canonical gauge

The key to parallelizing MPS algorithms is to use a gauge with multiple orthogonality centers. This allows different site tensors to be updated simultaneously, and to be merged back into the MPS consistently. The first such gauge was introduced by Vidal for TEBD. Any N -partite state can be written in Vidal's canonical form [1,2,12],

$$|\psi\rangle = \sum_{\sigma_1 \dots \sigma_N} \Gamma_1^{\sigma_1} \Lambda_1 \Gamma_2^{\sigma_2} \Lambda_2 \dots \Lambda_{N-1} \Gamma_N^{\sigma_N} |\sigma_1 \dots \sigma_N\rangle, \quad (7)$$

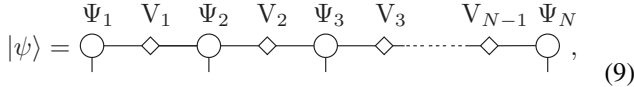
where the Λ_j are again diagonal matrices of singular values that serve as orthogonality centers. In graphical tensor notation, we write this as



$$|\psi\rangle = \Gamma_1 \Lambda_1 \Gamma_2 \Lambda_2 \Gamma_3 \Lambda_3 \dots \Lambda_{N-1} \Gamma_N, \quad (8)$$

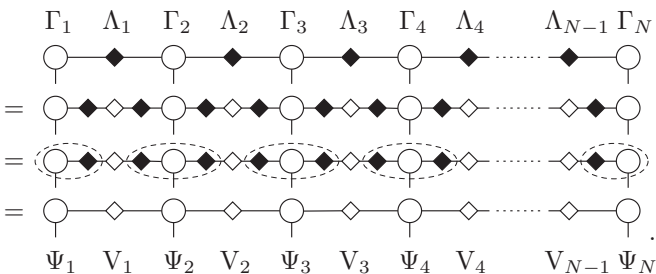
where the Γ_j are site tensors. The beauty of this canonical gauge is that it simultaneously gives the Schmidt decomposition of all bipartitions.

In this work, we use the inverse canonical gauge due to Stoudenmire and White [70], which is given by



$$|\psi\rangle = \Psi_1 V_1 \Psi_2 V_2 \Psi_3 V_3 \dots V_{N-1} \Psi_N, \quad (9)$$

where $V_j \equiv \Lambda_j^{-1}$. Although equivalent to the canonical gauge, it turns out to be a more natural choice for parallel TDVP and DMRG, as well as for TEBD, due to the site tensors, rather than the diagonal matrices, being orthogonality centers. An MPS in canonical form is transformed into inverse canonical form by inserting $V_j \Lambda_j = \mathbb{1}$ at each bond and then contracting the Λ matrices with the Γ site tensors, i.e.,



As Λ_j is diagonal, calculating V_j simply requires taking the reciprocal of the singular values.² The contractions $\Gamma_j V_j$, and $V_j \Gamma_{j+1}$ are similarly cheap. We find no issue with the inversion of Λ_j as tiny singular values corresponding to numerical noise are discarded [8,80].

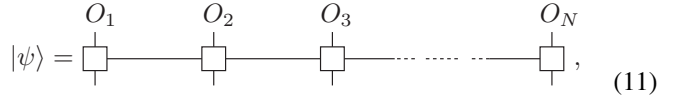
²Taking the reciprocal of a floating-point number should be “safe” as long as overflows and divisions by zero are avoided.

C. Matrix product operators

We represent N -site operators using the matrix product operator (MPO) construction,

$$O = \sum_{\substack{\sigma_1 \dots \sigma_N \\ \tau_1 \dots \tau_N}} O_1^{\sigma_1 \tau_1} O_2^{\sigma_2 \tau_2} \dots O_N^{\sigma_N \tau_N} |\sigma_1 \dots \sigma_N\rangle \langle \tau_1 \dots \tau_N|, \quad (10)$$

where the $O_j^{\sigma_j \tau_j}$ are matrices ($O_1^{\sigma_1 \tau_1}$ and $O_N^{\sigma_N \tau_N}$ being row and column vectors, respectively). In graphical tensor notation we write this as



$$|\psi\rangle = \square_{O_1} \square_{O_2} \square_{O_3} \dots \square_{O_N}, \quad (11)$$

where the O_j are site operators (represented by squares). Site operators are analogous to MPS site tensors, but have an extra physical edge. The physical edges have the same dimensions d_j as the MPS on which they act.

Local Hamiltonians are known to have a particularly compact MPO representation [81,82], which means the maximum MPO bond dimension m is independent of N . Exponentially decaying interactions can also be encoded efficiently [28,83]. As the same is not true of arbitrary long-range interactions, we follow Refs. [28,83,84] in approximating power laws by sums of exponentials, giving $m = n_H n_{\text{exps}} + 2$, where n_H is the number of long-range terms in the Hamiltonian, and n_{exps} is the number of exponentials used in the approximation. In this paper, we use the algorithm described in Ref. [28]. An alternative method [85–87] is discussed in Ref. [88].

D. Time-dependent variational principle

The McLachlan formulation of the Dirac-Frenkel-McLachlan time-dependent variational principle (TDVP) [89,90] approximates the time evolution of a state $|\psi\rangle$ under the Hamiltonian H by minimizing

$$\left\| i \frac{d}{dt} |\psi(t)\rangle - H |\psi(t)\rangle \right\|_2^2, \quad (12)$$

with $|\psi\rangle$ kept fixed while its derivative is varied (note that we set $\hbar = 1$ throughout this paper). Assuming the set of MPS with a given uniform bond dimension to be a smooth manifold (proven in Refs. [91–93]), Haegeman *et al.* used this variational principle to derive a novel algorithm for real and imaginary time evolution [31].

More recently, an improved TDVP algorithm was derived for finite MPS with open boundaries, which relies on the mixed canonical gauge [34,35]. This approach leads to an effective Schrödinger equation for states constrained to the MPS manifold,

$$i \frac{d}{dt} |\psi(t)\rangle = P_{T_{|\psi\rangle}} H |\psi(t)\rangle, \quad (13)$$

where $P_{T_{|\psi\rangle}}$ is an orthogonal projector onto the tangent space of $|\psi(t)\rangle$. The essence of this method is that the tangent space

projector can be decomposed as

$$P_{T_{|\psi\rangle}} = \sum_{j=1}^N P_L^{[1:j-1]} \otimes \mathbb{1}_j \otimes P_R^{[j+1:N]} - \sum_{j=1}^{N-1} P_L^{[1:j]} \otimes P_R^{[j+1:N]}, \quad (14)$$

where

$$P_L^{[1:j-1]} = \sum_{k=1}^{\chi_{j-1}} |\Phi_{L,k}^{[1:j-1]}\rangle \langle \Phi_{L,k}^{[1:j-1]}|, \\ P_R^{[j+1:N]} = \sum_{l=1}^{\chi_{j+1}} |\Phi_{R,l}^{[j+1:N]}\rangle \langle \Phi_{R,l}^{[j+1:N]}|, \quad (15)$$

meaning

$$|\psi(t + \delta t)\rangle = \exp(-iP_{T_{|\psi\rangle}} H \delta t) |\psi(t)\rangle \quad (16)$$

can be approximated by applying a Lie-Trotter-Suzuki decomposition [94] to the exponential. Consequently, one can sweep back and forth along the MPS (as in single-site DMRG [95–97]), time evolving one site tensor at a time. This algorithm, which we refer to as 1TDVP, is symplectic, so conserves the energy and norm of a state. However, it also restricts MPS to a fixed bond dimension.

To overcome this limitation, Haegeman *et al.* introduced a two-site variant (2TDVP) that similarly relies on the mixed canonical gauge [35]. In 2TDVP, the tangent space projector of Eq. (14) is replaced by

$$P_{T_{|\psi\rangle}^{[2]}} = \sum_{j=1}^{N-1} P_L^{[1:j-1]} \otimes \mathbb{1}_j \otimes \mathbb{1}_{j+1} \otimes P_R^{[j+2:N]} - \sum_{j=2}^{N-1} P_L^{[1:j-1]} \otimes \mathbb{1}_j \otimes P_R^{[j+1:N]}. \quad (17)$$

Because MPS of different bond dimension do not belong to the same manifold, it is no longer possible to describe the time evolution of the entire MPS by a differential equation. Instead, Haegeman *et al.* use a symmetric second-order Lie-Trotter-Suzuki decomposition with a discrete time step δt to arrive at an algorithm with the same sweeping pattern as the original two-site DMRG.³ Being able to dynamically vary the bond dimension makes 2TDVP particularly convenient. It has also been shown to give accurate results for a range of problems [37,98], so it is this variant we consider here.

III. PARALLEL TWO-SITE TDVP

In this section, we introduce the parallel two-site TDVP (p2TDVP) algorithm. As a preliminary, we describe how serial 2TDVP can be carried out in the inverse canonical gauge. This is mathematically equivalent to the usual formulation given in the literature but, crucially, allows the algorithm

³Higher-order decompositions entailing more sweeps are also possible [35].




FIG. 2. Left-to-right sweep in serial two-site TDVP using the inverse canonical gauge. Pairs of site tensors are evolved forwards in time by half a time step and single site tensors are evolved backwards in time by half a time step.

to be parallelized. We then describe how this parallelization is carried out, and finally discuss the truncation of singular values, and the algorithm’s stability.

A. Serial algorithm

A single time step in 2TDVP comprises a sweep from left to right followed by a sweep from right to left. During each of these two sweeps the MPS is evolved forwards in time by half a time step. Here we only explicitly describe the left-to-right sweep, illustrated in Fig. 2, as the right-to-left sweep is equivalent, with just the direction reversed. Both sweeps are described formally in Appendix A.

Starting from the left, the algorithm proceeds as follows: sites 1 and 2 are evolved forwards in time by half a time step; site 2 is evolved backwards in time by half a time step; sites 2 and 3 are evolved forwards in time by half a time step; site 3 is evolved backwards in time by half a time step;




FIG. 3. (a) The effective left environment β_j , and (b) the effective right environment γ_j , for an MPS site tensor Ψ_j . β_j is formed from all MPS tensors to the left of Ψ_j , along with their Hermitian conjugates. The corresponding Hamiltonian MPO tensors are sandwiched in between. γ_j is formed similarly, but from the tensors to the right of Ψ_j .




FIG. 4. Definition of the local state and its effective Hamiltonian for (a) the two-site update, and (b) the one-site update. The effective Hamiltonian is formed by connecting the left and right effective environments to the Hamiltonian MPO tensor(s) corresponding to the site(s) being updated. When carrying out a two-site update we use the L and R subscripts to refer to the left site and right site, respectively.

and so on. This pattern of local updates continues until the rightmost pair of sites is reached. The final step is to evolve this pair of sites forwards in time by half a time step, with no backwards time evolution necessary. The whole process is then reversed and the algorithm sweeps back from right to left. As the second sweep is simply the mirror image of the first, it begins by evolving the rightmost pair of sites forwards in time by another half a time step. One can therefore evolve the rightmost pair just once by a full time step instead of two half time steps (see Appendix A).

The one and two-site time evolution steps rely on “effective environments”, which are the same as in DMRG. Each site tensor Ψ_j has a left and a right effective environment, labeled β_j and γ_j , respectively. These are defined in Fig. 3. The leftmost and rightmost MPS environments (β_1 and γ_N) are trivial, corresponding to 1×1 identity matrices. It is important to note that the effective environments need not be created from scratch at every step since previous environments can be cached and updated. At the beginning of a simulation, all righthand environments $\gamma_1 \dots \gamma_N$ are created iteratively from right to left.

To time evolve two sites, Ψ_j and Ψ_{j+1} , we construct an effective local state $\Theta^{(2)}$ and an effective two-site Hamiltonian $H_{\text{eff}}^{(2)}$. These are described in Fig. 4(a). Evolving $\Theta^{(2)}$ forwards in time means calculating

$$\Theta^{(2)'} = \exp(-iH_{\text{eff}}^{(2)}\delta t/2)\Theta^{(2)}. \quad (18)$$

Using the Lanczos exponentiation [99] method⁴ means that $H_{\text{eff}}^{(2)}$ is not explicitly required, only $H_{\text{eff}}^{(2)}\Theta^{(2)}$, so a more




FIG. 5. The resplitting of $\Theta^{(2)'}$ following a two-site update. First $\Theta^{(2)'}$ is reshaped into a matrix. Then a truncated singular value decomposition (SVD) is performed, giving $A'\Lambda'B'$. A' and B' are reshaped into site tensors, and $V'\Lambda'$ is inserted between Λ' and B' . Finally, $A'\Lambda'$ is contracted to form Ψ'_L , and $\Lambda'B'$ is contracted to form Ψ'_R .

efficient tensor contraction pattern can be employed. Figure 5 explains how $\Theta^{(2)'}$ is split back into two site tensors using the singular value decomposition (SVD). Here the smallest singular values are discarded to keep the bond dimension from growing too large (see Sec. III C). After this two-site update a new lefthand environment β_{j+1} is created from the contraction of β_j , Ψ_j , V_j , Ψ_j^\dagger , V_j , and the Hamiltonian MPO tensor for site j , i.e.,

$$\beta_{j+1} = \beta_j \Psi_j V_j \Psi_j^\dagger V_j. \quad (19)$$

Time evolving one site Ψ_{j+1} requires the construction of an effective one-site Hamiltonian $H_{\text{eff}}^{(1)}$ from β_{j+1} and γ_{j+1} . This is described in Fig. 4(b). The local state $\Theta^{(1)}$ is now just the vectorization of Ψ_{j+1} . To time evolve $\Theta^{(1)}$ backwards in time means calculating

$$\Theta^{(1)'} = \exp(+iH_{\text{eff}}^{(1)}\delta t/2)\Theta^{(1)}, \quad (20)$$

but again only $H_{\text{eff}}^{(1)}\Theta^{(1)}$ is actually required by the Lanczos routine. After this one-site update, γ_{j+1} can be discarded.

A time step in 2TDVP has the same time complexity as a sweep in two-site DMRG. The most expensive operation is the contraction of the network representing $H_{\text{eff}}^{(2)}\Theta^{(2)}$, giving a bound of

$$\mathcal{O}(N\chi^3md^2 + N\chi^2m^2d^3). \quad (21)$$

This means that systems with many degrees of freedom (e.g. bosonic systems) can be very demanding. Large systems with long-range interactions are also challenging because of the size of the bond dimension required by the Hamiltonian MPO (especially if it contains multiple long-range terms). These considerations motivate the need for a parallel algorithm.

B. Parallelization

The intuition behind the parallel algorithm is the fact that local updates approximately preserve the inverse canonical

⁴An alternative is the recent algorithm from Al-Mohy and Higham based on truncated Taylor series [100].




FIG. 6. Partitioning of MPS and sweeping pattern for parallel 2TDVP with (a) two and (b) four processes. The dashed lines represent partition boundaries. At the start of a time step the sweeps proceed away from the center of the MPS, with each neighboring partition sweeping in the opposite direction (top arrows). For the second half of the time step the sweeps are reversed (bottom arrows).

gauge for small δt . We thus parallelize 2TDVP by carrying out these local updates simultaneously on separate processes. More concretely, we split the N -site MPS into p partitions, which are updated in parallel. We use a message-passing parallel programming model and assign each partition to a separate process. A partition must contain a minimum of two site tensors, so we let p be an even number between 2 and $N/2$. The full sweeps of the serial algorithm are replaced by partial sweeps carried out in parallel; each process simultaneously sweeps along the tensors in its own partition following the pattern introduced by Stoudenmire and White for parallel DMRG [70]. This is illustrated in Fig. 6 for two and four processes. Notice that the sweeping direction alternates for each neighboring partition. The two central processes always sweep away from the center of the MPS during the first half of a time step.

At the start of a simulation, the necessary initial effective environments are computed sequentially, with each being assigned to the process owning the corresponding site tensor. Processes that start by sweeping right will require righthand environments and vice-versa for those that start by sweeping left.

When sweeps reach a partition boundary, it is necessary for neighboring processes to communicate. Firstly, the processes need to exchange boundary environments, and secondly, one of the processes needs the Ψ tensor belonging to its neighbor in order to carry out the two-site update. This communication involves sending $\mathcal{O}(\chi^2(m+d))$ floating-point numbers. We let the lefthand process update the boundary sites while the righthand process waits. The lefthand process then sends the updated tensors to the righthand process and both processes update their respective effective environments. Figure 7 illustrates how the local updates proceed in parallel away from partition boundaries. We describe the algorithm formally in Appendix B.

C. Truncation of singular values

Serial 2TDVP has four sources of error: the projection, the Lie-Trotter-Suzuki decomposition, the local integration,




FIG. 7. Time evolving local sites in parallel 2TDVP. Two neighboring processes are shown sweeping away from their shared boundary (represented by a dashed line). As in Fig. 2, pairs of sites are evolved forwards in time by half a time step and single sites are evolved backwards in time by half a time step.

and the truncation of Schmidt coefficients. For a detailed discussion we refer the reader to Ref. [37]. The dependence of the different errors on the MPS bond dimension and the choice of δt is rather subtle [98] but, if these parameters are chosen with care, it is the truncation and projection errors that usually dominate due to the growth of bipartite entanglement [101]. The projection error can be prohibitively expensive to compute, especially for long-range models [37,102], but the truncation error is simply calculated from the discarded singular values, as in TEBD [12].

In both the serial and parallel algorithms, we quantify the truncation error due to the single SVD shown in Fig. 5 using the discarded weight

$$w = \sum_{j=\chi+1}^{\chi'} \lambda_j^2, \quad (22)$$

where χ' is the full rank of the matrix $\Theta^{(2)^\nu}$, λ_j are its singular values (sorted in order of descending magnitude), and χ is the truncated rank. We choose χ as follows. First, we define a truncation error tolerance w_{\max} , which is the maximum allowed discarded weight per SVD. We then find the minimum rank χ_w such that $w \leq w_{\max}$. Finally, we set $\chi = \min(\chi_w, \chi_{\max})$.

The total discarded weight w_{total} is defined as the cumulative sum of w over all SVDs, over all time steps. In the worst case, w_{total} will grow exponentially due to a linear growth of bipartite entanglement entropy [4]. However, long-range models can exhibit a logarithmic growth of entanglement entropy, even when this growth is linear in the corresponding local model [41,103–106], meaning w_{total} will grow as a power law.

D. Time step and stability

The parallelization of 2TDVP introduces two further sources of error:

(i) Information about lattice sites propagates along the MPS at a finite speed meaning that each process will always

be using at least one “out-of-date” local environment. For nonlocal 1D models, this induces an artificial locality, since instantaneous long-range interactions become effectively retarded. For p parallel processes, we expect this error to be small if the characteristic velocity v of the dynamics satisfies

$$v \ll (N/p)/\delta t, \quad (23)$$

where we have assumed that each parallel partition contains $\sim(N/p)$ sites.

(ii) Like all parallel MPS algorithms, the local updates in p2TDVP formally break the global gauge conditions, meaning the inverse canonical form only holds approximately. In serial TDVP the inverse canonical gauge is also technically broken, but the orthogonality of the state is preserved with respect to the last updated site (which remains an orthogonality center). In parallel TDVP this orthogonality may be lost, since different parts of the MPS are updated simultaneously.

These errors can be controlled by reducing δt or decreasing the number of parallel processes, meaning that it should be possible to converge a calculation, as is typically done with serial MPS algorithms. For moderate values of δt , however, the breaking of the gauge conditions at partition boundaries can cause the p2TDVP algorithm to become unstable if very small singular values are kept. To circumvent this, we define a relative SVD truncation tolerance ε . We discard singular values smaller than $\varepsilon\lambda_1$ (where λ_1 is the largest singular value), in addition to carrying out the truncation procedure described in Sec. III C.

If the error in the norm grows unacceptably large during a simulation, the MPS can also be reorthonormalized and the effective environments recomputed. As this is a serial procedure, it should be carried out infrequently to avoid affecting the algorithm’s parallel efficiency. Note that it is, however, important to ensure the initial state is orthonormal.

For the benchmark calculations described in Sec. IV, a value of $\varepsilon = 10^{-12}$ was sufficient, with no reorthonormalization necessary. However, the appropriate value of ε depends on the system and choice of time step. In Fig. 8, we describe a p2TDVP simulation carried out on a 641-site spin chain using different values of δt and ε . With $\delta t = 0.01$ and $\varepsilon = 10^{-8}$, the error in the norm is seen to blow up. However, reorthonormalization brings it back under control. In comparison, the calculation with $\delta t = 0.002$, and $\varepsilon = 10^{-12}$ remains stable, taking 2.4 times longer to run.

IV. BENCHMARKS

In this section, we describe the results of our numerical experiments. To test p2TDVP, we carried out benchmark calculations on spin-half models with one, two, and three long-range interaction terms. We utilized up to 32 processes, with one process assigned per compute node. Each compute node used up to 16 threads (e.g., for linear algebra operations). Full details of the test platform [107], software used [108–121], and simulation parameters are provided in Ref. [88].

To quantify the error introduced by the parallelization, we run our benchmark simulations on a single process and calculate the difference in the observables of interest. We also




FIG. 8. Error in the norm for a 641-site spin chain evolved using 32 processes under the long-range Ising Hamiltonian with $\alpha = 2.3$, following the quench described in Sec. IV A. The solid curve shows the result for $\delta t = 0.01$, $\varepsilon = 10^{-8}$. After about 600 time steps, the error starts to grow exponentially (dotted curve). If the MPS is reorthonormalized before this, its norm is brought back under control (righthand solid curve). At the end of the calculation $\chi = 116$ and $w_{\text{total}} = 7.0 \times 10^{-10}$. The dashed curve shows the result for $\delta t = 0.002$, $\varepsilon = 10^{-12}$. At the end of this calculation, $\chi = 90$ and $w_{\text{total}} = 8.8 \times 10^{-10}$. Both calculations used $m = 22$, and $w_{\text{max}} = 10^{-16}$.

compute the infidelity,

$$I(|\psi_1\rangle, |\psi_2\rangle) \equiv 1 - \frac{|\langle \psi_1 | \psi_2 \rangle|}{\sqrt{\langle \psi_1 | \psi_1 \rangle \langle \psi_2 | \psi_2 \rangle}}, \quad (24)$$

between the serial and parallel MPS at the end of the calculations, which bounds the error in all observables.

When truncation is the dominant source of error, $I \sim w_{\text{total}}$ [3,75]. Under this condition w_{total} can be used as a proxy for I . In general, however, we would expect w_{total} to provide a lower bound as it does not account for the projection error.

In the following, we define spin Hamiltonians in terms of Pauli X , Y , and Z operators σ_i^x , σ_i^y , and σ_i^z (where i is the index of the spin), and set the interaction strengths (and hence the energy and time scales) to unity.

A. Long-range Ising model

Our first benchmark looks at the spreading of correlations after a global quench in the ferromagnetic phase of the transverse field Ising model with short to intermediate-range interactions ($\alpha > 2$). Denoting the transverse magnetic field by B , the Hamiltonian is

$$H = - \sum_{i < j} \frac{1}{|i - j|^\alpha} \sigma_i^z \sigma_j^z - B \sum_{i=1}^N \sigma_i^x. \quad (25)$$

We track the evolution of an equal-time spin-spin correlation function to see how well p2TDVP can capture a nonlocal observable. Here we follow Liu *et al.* [106], but note that this scenario was first studied using TDVP for the antiferromagnetic interaction case in Refs. [40,41].

Using a Krylov space method [122,123], Liu *et al.* accurately simulate the quench dynamics of a 19-site lattice with periodic boundaries. Starting from the ferromagnetic product state $|\psi_0\rangle = |111 \dots 1\rangle$, correlation confinement [124] is shown to arise due to the presence of power law interactions. This confinement is stronger the longer the range of the interactions. Stronger confinement is also shown to decrease

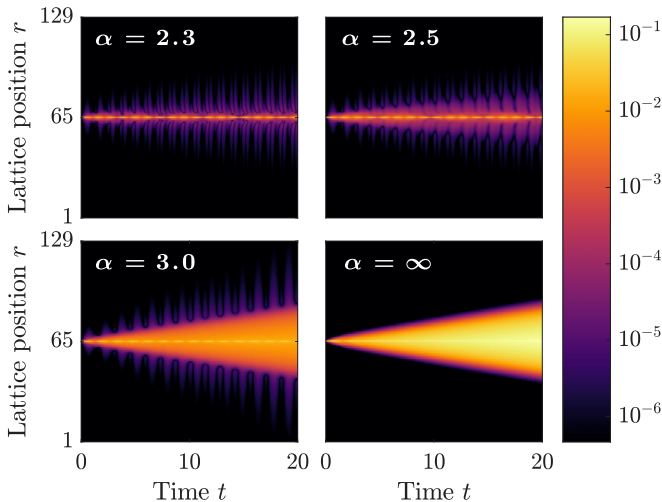


FIG. 9. Density plots of the correlation function $C_{r,65}(t)$ for a 129-site chain evolved under the long-range Ising Hamiltonian for four values of α . The data were calculated using p2TDVP with 32 processes. All values deviate from the serial calculations by less than 1% (less than 0.2% for the more accurate $\alpha = \infty$ case).

the bipartite entanglement present, meaning the maximal bond dimension required for our simulations should decrease for smaller α .

Liu *et al.* note that correlation confinement persists when the initial state is a ferromagnetic ground state of H . We observe this behavior for lattices with open boundaries when applying a quench to the ground state of Eq. (25) with $\alpha = 3.0$ and $B = 0.1$. As in Ref. [106], we quench to $B = 0.27$ for various values of α , and calculate the correlation function,

$$C_{r,k}(t) = \langle \psi(t) | \sigma_r^z \sigma_k^z | \psi(t) \rangle - \langle \psi(t) | \sigma_r^z | \psi(t) \rangle \langle \psi(t) | \sigma_k^z | \psi(t) \rangle, \quad (26)$$

where r is the lattice site index, k is the index of the central lattice site, and $|\psi(t)\rangle$ is the state at time t after the quench. As we consider chains with an odd number of spins, $k = (L + 1)/2$.

In Fig. 9, we plot $C_{r,65}(t)$ for $\alpha = 2.3, 2.5, 3.0$, and for the nearest-neighbor case ($\alpha = \infty$). These results, computed using 32 processes, give excellent agreement with the serial calculations over many orders of magnitude. We see that the correlation confinement disappears as $\alpha \rightarrow \infty$, and a linear light cone [125–128] is recovered. In fact a linear light cone already seems to be present in the $\alpha = 3.0$ case, consistent with the bounds given in Refs. [129,130]. Our numerical experiments lead us to conjecture that the light cone is sub-linear for $\alpha < 3$. The other striking difference between the local and nonlocal models is the existence of oscillatory long-range correlations in the latter, similar to Refs. [40,41]. The relatively large system size used here allows us to conjecture that these correlations decay as a power-law with an exponent approximately equal to α .

At the end of the simulations, we compute the infidelities I between the serial and parallel calculations. In Table I, we show these for 32 processes, along with the total discarded weights w_{total} from the end of the serial calculations. We find

TABLE I. Summary of results for the 129-site long-range Ising model at the end of 1000 time steps (corresponding to $t = 20$). Calculations were carried out using $\delta t = 0.02$, and $\varepsilon = 10^{-12}$, on 1 and 32 processes. χ_{max} is the maximal MPS bond dimension, m is the Hamiltonian MPO bond dimension, w_{max} is the truncation error tolerance, and w_{total} is the total discarded weight from the serial calculation. Speedups relative to the serial calculations are indicated in the last column.

α	m	χ_{max}	w_{max}	w_{total}	Infidelity (I)	Speedup
2.3	15	128	10^{-16}	3.3×10^{-11}	3.8×10^{-10}	20.8
2.5	14	192	10^{-16}	3.1×10^{-11}	3.4×10^{-10}	23.0
3.0	13	256	10^{-16}	6.1×10^{-11}	5.6×10^{-10}	25.0
∞	3	512	10^{-18}	1.7×10^{-13}	4.2×10^{-13}	23.9

that the ratio I/w_{total} grows with decreasing α , from 2.5 for the nearest-neighbor case, to 11.5 for $\alpha = 2.3$.

A strong scaling analysis for the $\alpha = 2.3$ case is shown in Fig. 1. A speedup of 20.8 was achieved using 32 processes, although by this point the parallel efficiency drops below 70%. We find greater speedups for the simulations with larger bond dimensions (summarized in Table I for 32 processes). This is to be expected as the computational complexity of the linear algebra operations asymptotically dominates over the parallel overheads.

B. Long-range XY model

We next simulate a local quench in the antiferromagnetic XY model,

$$H = \frac{1}{2} \sum_{i < j}^N \frac{1}{|i - j|^\alpha} (\sigma_i^x \sigma_j^x + \sigma_i^y \sigma_j^y), \quad (27)$$

with very long-range interactions ($\alpha < 1$). In this regime, information can spread through the system almost instantaneously [40,131–134]. This is an important test case as it is not *a priori* clear how accurately p2TDVP will capture the nonlocal [40] propagation of information from a single site.

Following Haegeman *et al.* [35,135], we calculate the ground state $|\psi_0\rangle$ of H for a 101-site spin chain and apply a $U(1)$ symmetry-breaking perturbation,

$$U = \exp(i\pi\sigma_{51}^y/4), \quad (28)$$

to the central spin. We then examine the evolution of the single-site observable [35,132]

$$C_{51,r}(t) = |\langle \psi(t) | \sigma_r^x | \psi(t) \rangle - \langle \psi_0 | \sigma_r^x | \psi_0 \rangle|, \quad (29)$$

where the perturbed state at time t is given by

$$|\psi(t)\rangle = e^{-iHt} U |\psi_0\rangle. \quad (30)$$

Using p2TDVP we reproduce the results of Ref. [35] for $\alpha = 0.75$, which is the most interesting case as it illustrates the breakdown of light cone dynamics. It is also the most numerically challenging due to the large MPO bond dimension.

As shown in Fig. 1, this calculation scales well up to 32 processes with an efficiency $\geq 86\%$. For 32 processes this corresponds to a speedup of 27.4. The scaling is near optimal because the MPS saturates our chosen χ_{max} after just one time

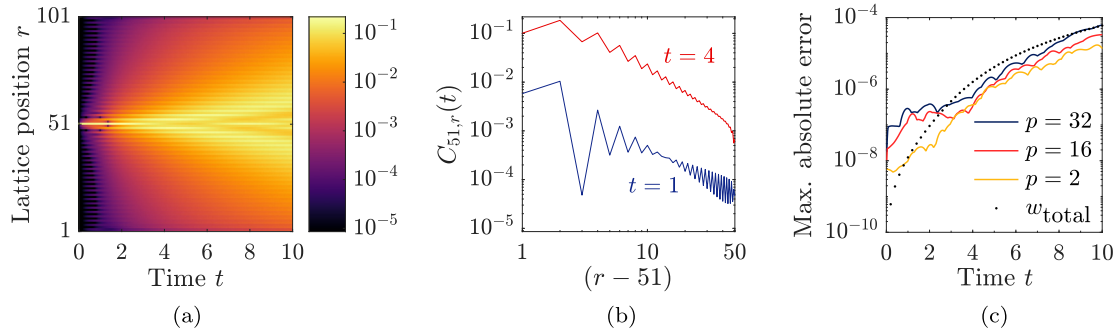


FIG. 10. Time evolution of the 101-site long-range XY spin chain with $\alpha = 0.75$. All calculations were carried out using $\delta t = 0.02$, $\varepsilon = 10^{-12}$, $\chi_{\max} = 256$, $m = 28$, and no truncation error tolerance. (a) Density plot of $C_{51,r}(t)$, calculated using 32 processes. Above the cutoff of 8.575×10^{-6} , all values deviate from the serial calculation by less than 1% (less than 0.1% after $t = 2.24$). (b) Log-log plot of $C_{51,r}(t)$ at times $t = 1$ and $t = 4$, calculated as in (a). On this scale, differences from the serial calculation are not visible. (c) Maximum absolute deviation of $C_{51,r}(t)$ from the serial calculation for p processes. w_{total} from the serial calculation is shown for comparison.

step. We have excluded the time taken to compute expectation values, but note here that we were also able to compute these in parallel, as discussed in Ref. [70], since $C_{51,r}$ depends on σ_r^x , which is a single-site observable.

The time evolution of $C_{51,r}(t)$, calculated using 32 processes, is shown in Fig. 10(a). The results deviate from the serial calculation by less than 1%, except for small values at the start of the simulation ($t < 0.8$). The spatial profile of $C_{51,r}$ is seen most clearly in Fig. 10(b). Its value oscillates, but appears to decay algebraically with r , in excellent agreement with Fig. 5 of Ref. [35].

In Fig. 10(c), we show the maximum absolute deviation in $C_{51,r}(t)$ from the serial calculation for 2, 16, and 32 processes, along with the discarded weight from the serial calculation. The deviation has a weak dependence on the number of processes p , but appears to be approximately bounded by w_{total} (except at the beginning of the simulation where the parallelization error seems to dominate). The infidelities I and total discarded weights w_{total} from the end of the calculations are shown in Table II. These also depend weakly on p , with I being less than an order of magnitude larger than w_{total} .

C. Long-range XXX model

In our final benchmark, we test p2TDVP with a U(1) symmetric MPS [110,136] by simulating the long-range isotropic Heisenberg (XXX) Hamiltonian with $\alpha = 2$,

$$H = \frac{1}{4} \sum_{i < j}^N \frac{1}{|i - j|^2} (\sigma_i^x \sigma_j^x + \sigma_i^y \sigma_j^y + \sigma_i^z \sigma_j^z). \quad (31)$$

TABLE II. Total discarded weights and infidelities for the 101-site long-range XY model at the end of 500 time steps.

Processes (p)	Total discarded weight (w_{total})	Infidelity (I)
1	6.0×10^{-5}	N/A
2	6.0×10^{-5}	5.1×10^{-4}
8	6.5×10^{-5}	5.3×10^{-4}
16	6.9×10^{-5}	5.6×10^{-4}
32	7.9×10^{-5}	6.1×10^{-4}

In the thermodynamic limit this is equivalent to the exactly solvable spin-half Haldane-Shastry model [137,138], which was argued in Ref. [30] to provide a stringent test case as it is both long-ranged and critical. Here we instead use p2TDVP to time evolve a 201-site spin chain with open boundaries in order to calculate the dynamical spin-spin correlation function,

$$C(r - k, t) = \langle \psi_0 | \sigma_r^z(t) \sigma_k^z(0) | \psi_0 \rangle, \quad (32)$$

where $|\psi_0\rangle$ is the ground state of Eq. (31), and k is the central lattice site (i.e., $k = 101$). As a σ^z perturbation does not break the U(1) symmetry of $|\psi_0\rangle$, the Z component of spin is conserved. This allows us to take advantage of symmetric block-sparse tensors [110,136], and hence use a relatively large bond dimension of $\chi_{\max} = 1024$.

In the thermodynamic limit, the dynamical spin-spin correlation function is given by [139]

$$\begin{aligned} C_{\infty}(x, t) &= \frac{(-1)^x}{4} \int_{-1}^1 d\lambda_1 \int_{-1}^1 d\lambda_2 e^{i(Qx - Et)}, \\ Q &= \pi \lambda_1 \lambda_2, \\ E &= \frac{\pi^2}{4} (\lambda_1^2 + \lambda_2^2 - 2\lambda_1^2 \lambda_2^2). \end{aligned} \quad (33)$$

In Fig. 11, we show the magnitude of $C(x, t)$ calculated using p2TDVP on 32 processes. We also show the relative difference η_{∞} from $C_{\infty}(x, t)$, where

$$\eta_{\infty} = \frac{|C(x, t) - C_{\infty}(x, t)|}{|C_{\infty}(x, t)|}. \quad (34)$$

There is quantitative difference between the calculations, but they agree well qualitatively (except towards the edges where $|C(x, t)|$ drops off exponentially in the finite system due to the open boundaries). In Fig. 12, we plot the real and imaginary parts of $C(x, t)$ for $x = 0, 2, 4, 6$. Again, we see good qualitative agreement with the analytic result, and with Fig. 4 of Ref. [30].

The simulation took 3.4 days to run, excluding the calculation of $C(x, t)$, although a different partitioning of the MPS should give a slight speedup (see Ref. [88]). The same computation would likely take weeks to run on a single




FIG. 11. (Left) Density plot of $|C(x, t)|$ in the t - x plane for the 201-site long-range XXX spin chain. At the edges, $|C(x, t)|$ drops to 1.2×10^{-3} . (Right) The relative difference η_∞ from the exact thermodynamic limit result. Towards the edges, η_∞ grows to 0.78. The calculation was carried out using p2TDVP on 32 processes with $m = 38$, $\delta t = 0.025$, $\varepsilon = 10^{-12}$, and $w_{\max} = 10^{-16}$. w_{total} at the end of the calculation was 7.4×10^{-11} .

compute node with serial 2TDVP, making a full scaling analysis impractical.

As we cannot easily calculate the error introduced by the parallel splitting for this system, we repeat the simulation on a smaller lattice of 65 sites with $\chi_{\max} = 512$. Denoting the results calculated on 1 and 32 processes by C_s and C_p , respectively, we find a maximum relative difference of $\max(\eta_p) = 3.1 \times 10^{-5}$, where

$$\eta_p = \frac{|C_p(x, t) - C_s(x, t)|}{|C_s(x, t)|}. \quad (35)$$

In contrast, $\min(\eta_\infty) = 5.3 \times 10^{-4}$ for both the serial and parallel simulations. At least for this smaller system then, the parallelization error is negligible compared to the deviation from the thermodynamic limit (see Fig. 13).




FIG. 12. Real and imaginary parts of $C(x, t)$ for $x = 0, 2, 4, 6$. The markers are the results of the p2TDVP calculation described in Fig. 11. The solid lines show the exact thermodynamic limit results for comparison.




FIG. 13. Error in $C(x, t)$ for the 65-site long-range XXX model calculation carried out with p2TDVP on 32 processes. (Left) Relative difference η_∞ from the thermodynamic limit result. (Right) Relative difference η_p from the serial 2TDVP calculation. The ratio $\eta_p/\eta_\infty < 1.3 \times 10^{-4}$ for all x and t . At the end of the calculation, w_{total} was 5.8×10^{-11} (4.5×10^{-11} for the serial calculation).

V. DISCUSSION

We have introduced a parallel version of the two-site TDVP algorithm (p2TDVP) and applied it to quenches in paradigmatic spin-half models with power law decaying interactions. To assess our algorithm's accuracy, we calculated onsite expectation values, equal-time two-point correlation functions, and a dynamical spin-spin correlation function. Remarkably, we have shown that demanding calculations can be accelerated at the cost of very little additional error. Though the parallel splitting can potentially lead to instability, we have explained how this can be worked around. Speedups are system dependent, but we have demonstrated parallel efficiencies of 65%–86% with 32 processes. This suggests that it should be possible to use our algorithm to simulate systems in a week that would otherwise take many months. The use of a dynamical load balancer may further improve this efficiency.

As a next step, p2TDVP could be applied to fermionic models. It is not yet clear how accurately p2TDVP can simulate 2D systems, but fermionic models in two dimensions can be especially challenging for all numerical methods [11,140], underlining the need for a parallel algorithm. Targeting larger 1D systems should be more straightforward. Large system sizes are important for the study of open quantum dynamics [141] and transport properties [49], and for distinguishing between many-body localized and thermal phases [142,143].

In Ref. [71], it was established that single-site DMRG can be parallelized. We therefore expect that a parallel variant of one-site TDVP (p1TDVP) could be developed using the same approach. This would enable a parallel version of the hybrid method discussed in Refs. [37,101,144], whereby a simulation starts with 2TDVP and switches to 1TDVP when χ_{\max} is saturated. 1TDVP is faster, and can give more accurate results for some observables [101]. A parallel version should scale well as the fixed bond dimension would allow for optimal load balancing. It may similarly be possible to apply the parallelization scheme presented here to other related MPS-local time-evolution methods [26,37,60].

We have focused on real time evolution, but p2TDVP can also be used for imaginary time evolution. This might prove beneficial for cases where parallel DMRG fails to converge. As this evolution is nonunitary, however, one has to pay particularly careful attention to the orthonormality of the MPS.

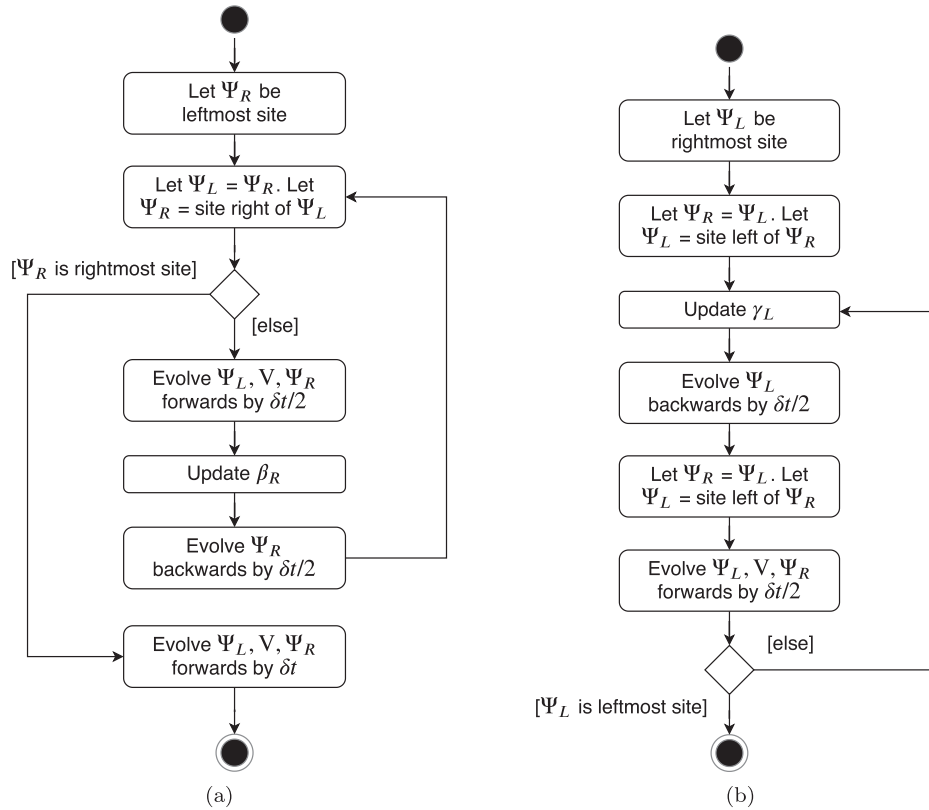


FIG. 14. UML activity diagrams illustrating the (a) left-to-right, and (b) right-to-left, sweeps constituting a single time step in the serial 2TDVP algorithm, when carried out in the inverse canonical gauge. Note that “site tensor” is abbreviated as “site”. For simplicity, we assume $N > 2$ (the trivial $N = 2$ case only requires the first sweep).

An obvious extension to this work would be the combination of parallel TDVP with established MPS techniques such as non-Abelian symmetries [81], different local bases [145–149], and infinite boundary conditions [150]. A code combining these features would be of great benefit to the community. Generalizing to other tensor network types is a further avenue to explore. TDVP can be extended to tree tensor network states [35,151], so it would be worthwhile to see if our algorithm can be modified to work with these or other networks that admit a canonical form [152,153].

Finally, another promising application is the solution of general partial differential equations (PDEs). MPS-based PDE solvers, such as the multigrid renormalization method [154], can be exponentially faster than standard PDE solvers. It is exciting to anticipate an additional speedup for such methods via the parallelization technique reported here.

ACKNOWLEDGMENTS

P.S. thanks Jonathan Coulthard, James Grant, Fangli Liu, Frank Pollmann, and Leon Schoonderwoerd for helpful discussions. This research made use of the Balena High Performance Computing Service at the University of Bath, and used Tensor Network Theory Library code written by Sarah Al-Assam, Chris Goodyer, and Jonathan Coulthard. The density plots in this paper use the “Inferno” color map by Nathaniel J. Smith, and Stefan van der Walt. P.S. is supported by ClusterVision and the University of Bath. M.L. and D.J. are grateful for funding from the UK’s Engineering and

Physical Sciences Research Council (EPSRC) for Grants No. EP/M013243/1 and No. EP/K038311/1. D.J. also acknowledges EPSRC Grant No. EP/P01058X/1. S.R.C. gratefully acknowledges support from the EPSRC under Grant No. EP/P025110/2.

S.R.C. and D.J. proposed the research direction; N.G. and M.L. conceptualized the project; N.G., M.L., and P.S. designed the algorithm; N.G. developed the serial TDVP code with supervision from M.L.; P.S. developed the parallel TDVP code, conducted the numerical experiments, and wrote the manuscript, with supervision from S.D. and S.R.C.; all authors discussed the results and edited the manuscript.

APPENDIX A: SERIAL ALGORITHM

We use unified modeling language (UML) activity diagrams to describe a single time step in the serial 2TDVP algorithm for an MPS in the inverse canonical gauge. Figures 14(a) and 14(b) describe the left-to-right sweep (illustrated schematically in Fig. 2) and right-to-left sweep, respectively. For each two-site update, the left and right site tensors are labeled Ψ_L and Ψ_R , with V being the diagonal matrix sandwiched between them.

As noted in Sec. III A, one can evolve the rightmost pair of sites at the end of the first sweep by a single full time step. This means that the second sweep does not need to carry out a forwards time evolution step for the rightmost two sites. This same approach can be used in the parallel version of the algorithm (see Fig. 15).

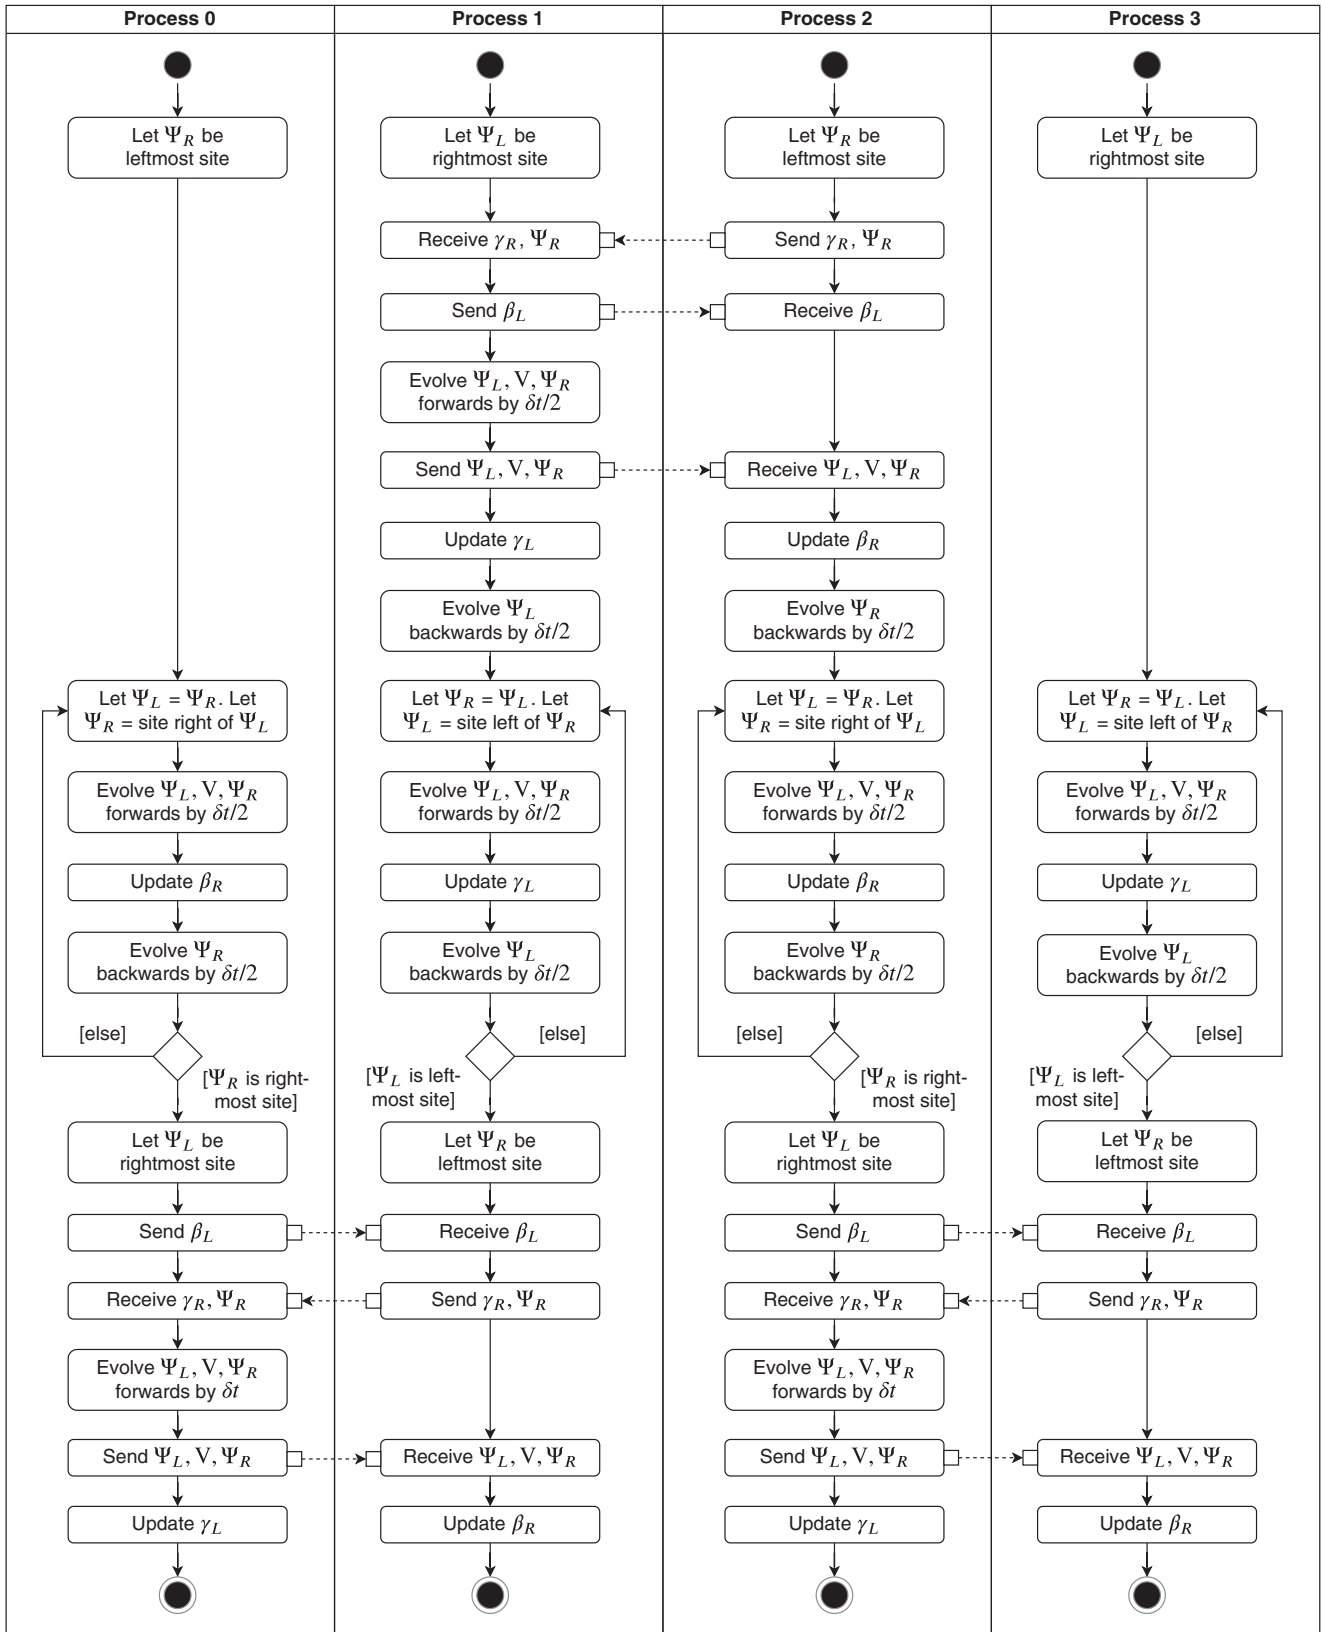


FIG. 15. UML activity diagram describing the first half of a p2TDVP time step on four processes. Dashed lines represent message-passing communication, and “site” is shorthand for “site tensor”.

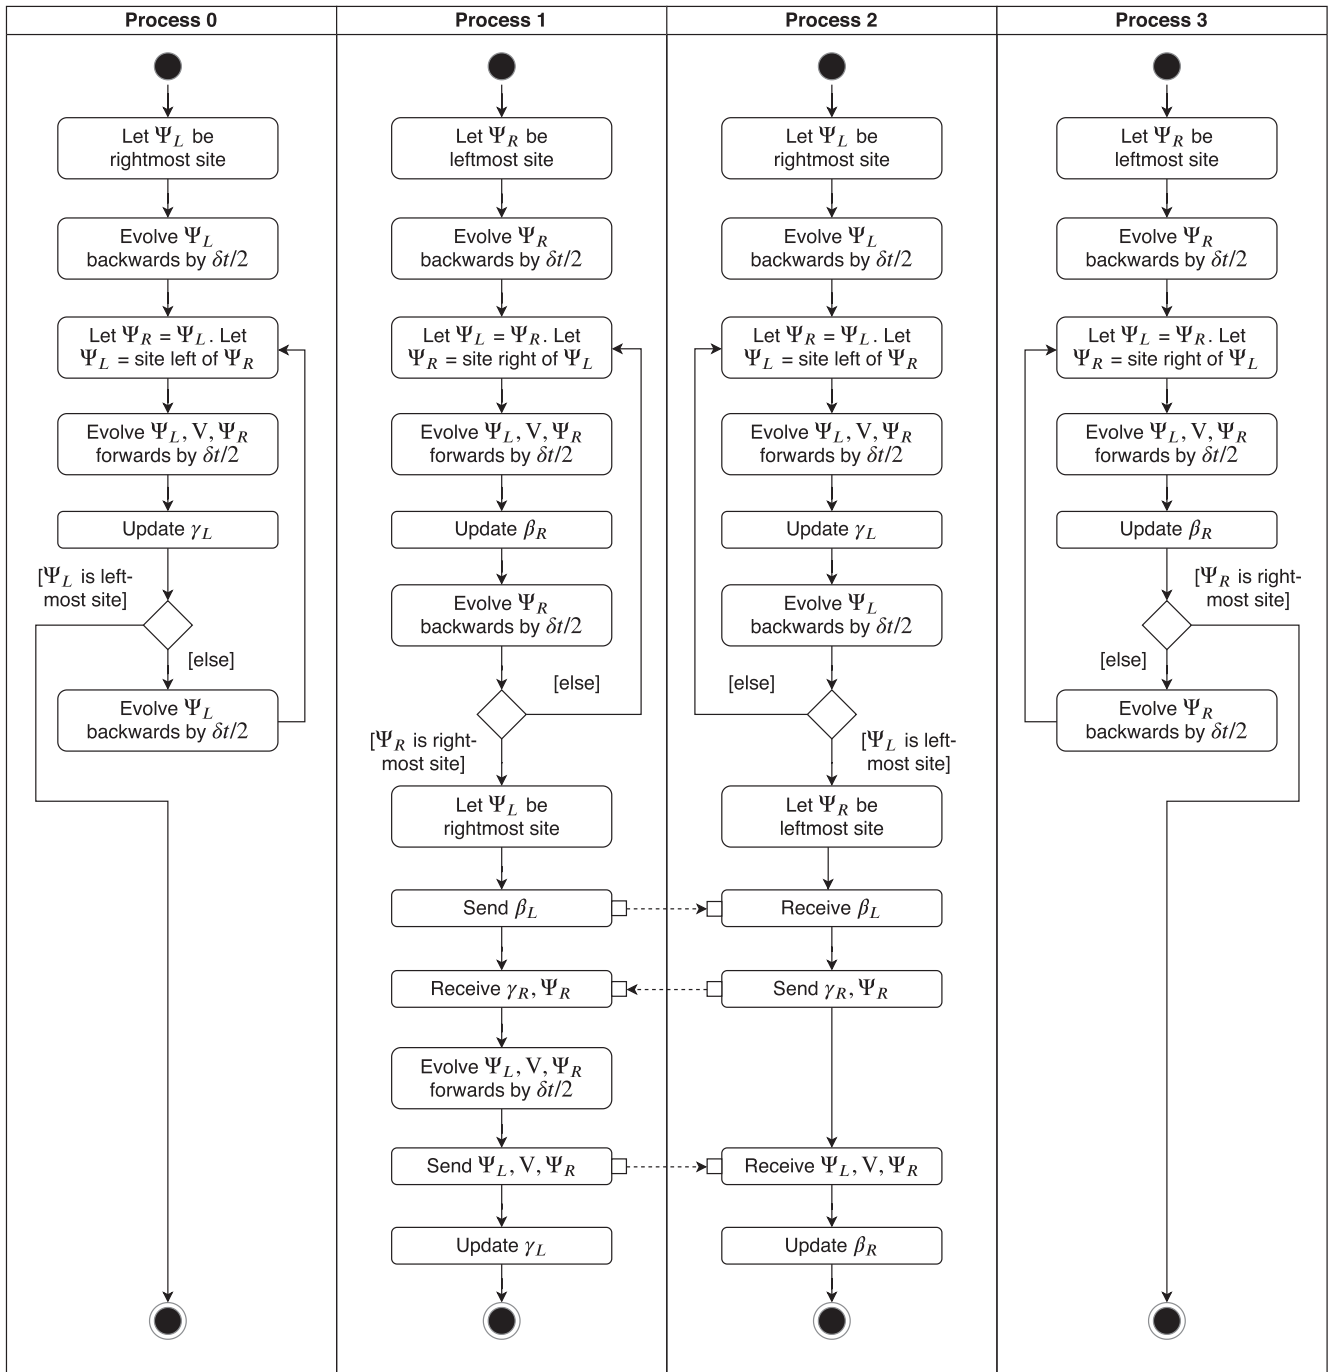


FIG. 16. UML activity diagram describing the second half of a p2TDVP time step on four processes. The sweep on each partition occurs in the opposite direction to the corresponding sweep in Fig. 15.

APPENDIX B: PARALLEL ALGORITHM

Here we use UML activity diagrams to describe a single time step in the p2TDVP algorithm. Figures 15 and

16 describe the first and second sweeps, respectively. For concreteness we show four parallel processes (illustrated schematically in Fig. 6(b)). This case contains all the logic necessary to generalize to p processes.

[1] G. Vidal, Efficient Classical Simulation of Slightly Entangled Quantum Computations, *Phys. Rev. Lett.* **91**, 147902 (2003).
 [2] G. Vidal, Efficient Simulation of One-Dimensional Quantum Many-Body Systems, *Phys. Rev. Lett.* **93**, 040502 (2004).

[3] F. Verstraete and J. I. Cirac, Matrix product states represent ground states faithfully, *Phys. Rev. B* **73**, 094423 (2006).
 [4] N. Schuch, M. M. Wolf, F. Verstraete, and J. I. Cirac, Entropy Scaling and Simulability by Matrix Product States, *Phys. Rev. Lett.* **100**, 030504 (2008).

- [5] J. Eisert, M. Cramer, and M. B. Plenio, Colloquium: Area laws for the entanglement entropy, *Rev. Mod. Phys.* **82**, 277 (2010).
- [6] S. Östlund and S. Rommer, Thermodynamic Limit of Density Matrix Renormalization, *Phys. Rev. Lett.* **75**, 3537 (1995).
- [7] J. Dukelsky, M. A. Martín-Delgado, T. Nishino, and G. Sierra, Equivalence of the variational matrix product method and the density matrix renormalization group applied to spin chains, *EPL* **43**, 457 (1998).
- [8] U. Schollwöck, The density-matrix renormalization group in the age of matrix product states, *Ann. Phys. (NY)*, **326**, 96 (2011).
- [9] S. R. White, Density Matrix Formulation for Quantum Renormalization Groups, *Phys. Rev. Lett.* **69**, 2863 (1992).
- [10] S. R. White, Density-matrix algorithms for quantum renormalization groups, *Phys. Rev. B* **48**, 10345 (1993).
- [11] E. Stoudenmire and S. R. White, Studying Two-Dimensional Systems with the Density Matrix Renormalization Group, *Annu. Rev. Condens. Matter Phys.* **3**, 111 (2012).
- [12] A. J. Daley, C. Kollath, U. Schollwöck, and G. Vidal, Time-dependent density-matrix renormalization-group using adaptive effective Hilbert spaces, *J. Stat. Mech.* (2004) P04005.
- [13] A. Micheli, G. K. Brennen, and P. Zoller, A toolbox for lattice-spin models with polar molecules, *Nat. Phys.* **2**, 341 (2006).
- [14] B. Yan, S. A. Moses, B. Gadway, J. P. Covey, K. R. A. Hazzard, A. M. Rey, D. S. Jin, and J. Ye, Observation of dipolar spin-exchange interactions with lattice-confined polar molecules, *Nature (London)* **501**, 521 (2013).
- [15] P. Schauß, M. Cheneau, M. Endres, T. Fukuhara, S. Hild, A. Omran, T. Pohl, C. Gross, S. Kuhr, and I. Bloch, Observation of spatially ordered structures in a two-dimensional Rydberg gas, *Nature (London)* **491**, 87 (2012).
- [16] H. Weimer, M. Müller, I. Lesanovsky, P. Zoller, and H. P. Büchler, A Rydberg quantum simulator, *Nat. Phys.* **6**, 382 (2010).
- [17] J. Zeiher, R. van Bijnen, P. Schauß, S. Hild, J.-y. Choi, T. Pohl, I. Bloch, and C. Gross, Many-body interferometry of a Rydberg-dressed spin lattice, *Nat. Phys.* **12**, 1095 (2016).
- [18] D. Porras and J. I. Cirac, Effective Quantum Spin Systems with Trapped Ions, *Phys. Rev. Lett.* **92**, 207901 (2004).
- [19] J. W. Britton, B. C. Sawyer, A. C. Keith, C.-C. J. Wang, J. K. Freericks, H. Uys, M. J. Biercuk, and J. J. Bollinger, Engineered two-dimensional Ising interactions in a trapped-ion quantum simulator with hundreds of spins, *Nature (London)* **484**, 489 (2012).
- [20] R. Islam, C. Senko, W. C. Campbell, S. Korenblit, J. Smith, A. Lee, E. E. Edwards, C.-C. J. Wang, J. K. Freericks, and C. Monroe, Emergence and frustration of magnetism with variable-range interactions in a quantum simulator, *Science* **340**, 583 (2013).
- [21] P. Richerme, Z.-X. Gong, A. Lee, C. Senko, J. Smith, M. Foss-Feig, S. Michalakakis, A. V. Gorshkov, and C. Monroe, Non-local propagation of correlations in quantum systems with long-range interactions, *Nature (London)* **511**, 198 (2014).
- [22] P. Jurcevic, B. P. Lanyon, P. Hauke, C. Hempel, P. Zoller, R. Blatt, and C. F. Roos, Quasiparticle engineering and entanglement propagation in a quantum many-body system, *Nature (London)* **511**, 202 (2014).
- [23] J. Smith, A. Lee, P. Richerme, B. Neyenhuis, P. W. Hess, P. Hauke, M. Heyl, D. A. Huse, and C. Monroe, Many-body localization in a quantum simulator with programmable random disorder, *Nat. Phys.* **12**, 907 (2016).
- [24] J. Zhang, G. Pagano, P. W. Hess, A. Kyrianiadis, P. Becker, H. Kaplan, A. V. Gorshkov, Z.-X. Gong, and C. Monroe, Observation of a many-body dynamical phase transition with a 53-qubit quantum simulator, *Nature (London)* **551**, 601 (2017).
- [25] A. E. Feiguin and S. R. White, Time-step targeting methods for real-time dynamics using the density matrix renormalization group, *Phys. Rev. B* **72**, 020404(R) (2005).
- [26] J. J. García-Ripoll, Time evolution of matrix product states, *New J. Phys.* **8**, 305 (2006).
- [27] J. Schachenmayer, Dynamics and long-range interactions in 1d quantum systems, Diploma thesis, Technische Universität München, 2008.
- [28] B. Pirvu, V. Murg, J. I. Cirac, and F. Verstraete, Matrix product operator representations, *New J. Phys.* **12**, 025012 (2010).
- [29] E. M. Stoudenmire and S. R. White, Minimally entangled typical thermal state algorithms, *New J. Phys.* **12**, 055026 (2010).
- [30] M. P. Zaletel, R. S. K. Mong, C. Karrasch, J. E. Moore, and F. Pollmann, Time-evolving a matrix product state with long-ranged interactions, *Phys. Rev. B* **91**, 165112 (2015).
- [31] J. Haegeman, J. I. Cirac, T. J. Osborne, I. Pižorn, H. Verschelde, and F. Verstraete, Time-Dependent Variational Principle for Quantum Lattices, *Phys. Rev. Lett.* **107**, 070601 (2011).
- [32] M. L. Wall and L. D. Carr, Out-of-equilibrium dynamics with matrix product states, *New J. Phys.* **14**, 125015 (2012).
- [33] C. Lubich, T. Rohwedder, R. Schneider, and B. Vandereycken, Dynamical approximation by hierarchical tucker and tensor-train tensors, *SIAM J. Matrix Anal. Appl.* **34**, 470 (2013).
- [34] C. Lubich, I. Oseledets, and B. Vandereycken, Time integration of tensor trains, *SIAM J. Numer. Anal.* **53**, 917 (2015).
- [35] J. Haegeman, C. Lubich, I. Oseledets, B. Vandereycken, and F. Verstraete, Unifying time evolution and optimization with matrix product states, *Phys. Rev. B* **94**, 165116 (2016).
- [36] E. Ronca, Z. Li, C. A. Jimenez-Hoyos, and G. K.-L. Chan, Time-step targeting time-dependent and dynamical density matrix renormalization group algorithms with *ab initio* hamiltonians, *J. Chem. Theory Comput.* **13**, 5560 (2017).
- [37] S. Paeckel, T. Köhler, A. Swoboda, S. R. Manmana, U. Schollwöck, and C. Hubig, Time-evolution methods for matrix-product states, *Ann. Phys. (NY)* **411**, 167998 (2019).
- [38] T. Hashizume, J. C. Halimeh, and I. P. McCulloch, Hybrid infinite time-evolving block decimation algorithm for long-range multi-dimensional quantum many-body systems, [arXiv:1910.10726](https://arxiv.org/abs/1910.10726) [cond-mat.str-el].
- [39] T. Koffel, M. Lewenstein, and L. Tagliacozzo, Entanglement Entropy for the Long-Range Ising Chain in a Transverse Field, *Phys. Rev. Lett.* **109**, 267203 (2012).
- [40] P. Hauke and L. Tagliacozzo, Spread of Correlations in Long-Range Interacting Quantum Systems, *Phys. Rev. Lett.* **111**, 207202 (2013).
- [41] A. S. Buyskikh, M. Fagotti, J. Schachenmayer, F. Essler, and A. J. Daley, Entanglement growth and correlation spreading with variable-range interactions in spin and fermionic tunneling models, *Phys. Rev. A* **93**, 053620 (2016).
- [42] J. C. Halimeh, V. Zauner-Stauber, I. P. McCulloch, I. de Vega, U. Schollwöck, and M. Kastner, Prethermalization and

- persistent order in the absence of a thermal phase transition, *Phys. Rev. B* **95**, 024302 (2017).
- [43] D. Jaschke, K. Maeda, J. D. Whalen, M. L. Wall, and L. D. Carr, Critical phenomena and Kibble–Zurek scaling in the long-range quantum Ising chain, *New J. Phys.* **19**, 033032 (2017).
- [44] J. C. Halimeh and V. Zauner-Stauber, Dynamical phase diagram of quantum spin chains with long-range interactions, *Phys. Rev. B* **96**, 134427 (2017).
- [45] V. Zauner-Stauber and J. C. Halimeh, Probing the anomalous dynamical phase in long-range quantum spin chains through Fisher-zero lines, *Phys. Rev. E* **96**, 062118 (2017).
- [46] B. Žunkovič, M. Heyl, M. Knap, and A. Silva, Dynamical Quantum Phase Transitions in Spin Chains with Long-Range Interactions: Merging Different Concepts of Nonequilibrium Criticality, *Phys. Rev. Lett.* **120**, 130601 (2018).
- [47] S. Pappalardi, A. Russomanno, B. Žunkovič, F. Iemini, A. Silva, and R. Fazio, Scrambling and entanglement spreading in long-range spin chains, *Phys. Rev. B* **98**, 134303 (2018).
- [48] A. Leroise, B. Žunkovič, J. Marino, A. Gambassi, and A. Silva, Impact of nonequilibrium fluctuations on prethermal dynamical phase transitions in long-range interacting spin chains, *Phys. Rev. B* **99**, 045128 (2019).
- [49] B. Kloss and Y. Bar Lev, Spin transport in a long-range-interacting spin chain, *Phys. Rev. A* **99**, 032114 (2019).
- [50] A. Leroise, B. Žunkovič, A. Silva, and A. Gambassi, Quasilocalized excitations induced by long-range interactions in translationally invariant quantum spin chains, *Phys. Rev. B* **99**, 121112(R) (2019).
- [51] T. Zhou, S. Xu, X. Chen, A. Guo, and B. Swingle, Operator Lévy Flight: Light Cones in Chaotic Long-Range Interacting Systems, *Phys. Rev. Lett.* **124**, 180601 (2020).
- [52] G. Piccitto, B. Žunkovič, and A. Silva, Dynamical phase diagram of a quantum Ising chain with long-range interactions, *Phys. Rev. B* **100**, 180402(R) (2019).
- [53] R. Borrelli and M. F. Gelin, Simulation of quantum dynamics of excitonic systems at finite temperature: an efficient method based on thermo field dynamics, *Sci. Rep.* **7**, 1 (2017).
- [54] R. Borrelli, Theoretical study of charge-transfer processes at finite temperature using a novel thermal Schrödinger equation, *Chem. Phys.* **515**, 236 (2018).
- [55] Y. Kurashige, Matrix product state formulation of the multi-configuration time-dependent Hartree theory, *J. Chem. Phys.* **149**, 194114 (2018).
- [56] A. Baiardi and M. Reiher, Large-scale quantum dynamics with matrix product states, *J. Chem. Theory Comput.* **15**, 3481 (2019).
- [57] M. F. Gelin, R. Borrelli, and W. Domcke, Origin of unexpectedly simple oscillatory responses in the excited-state dynamics of disordered molecular aggregates, *J. Phys. Chem. Lett.* **10**, 2806 (2019).
- [58] B. Kloss, D. R. Reichman, and R. Tempelaar, Multiset Matrix Product State Calculations Reveal Mobile Franck-Condon Excitations Under Strong Holstein-Type Coupling, *Phys. Rev. Lett.* **123**, 126601 (2019).
- [59] X. Xie, Y. Liu, Y. Yao, U. Schollwöck, C. Liu, and H. Ma, Time-dependent density matrix renormalization group quantum dynamics for realistic chemical systems, *J. Chem. Phys.* **151**, 224101 (2019).
- [60] W. Li, J. Ren, and Z. Shuai, Numerical assessment for accuracy and GPU acceleration of TD-DMRG time evolution schemes, *J. Chem. Phys.* **152**, 024127 (2020).
- [61] G. Hager, E. Jeckelmann, H. Fehske, and G. Wellein, Parallelization strategies for density matrix renormalization group algorithms on shared-memory systems, *J. Comput. Phys.* **194**, 795 (2004).
- [62] C. Nemes, G. Barcza, Z. Nagy, Ö. Legeza, and P. Szolgay, The density matrix renormalization group algorithm on kilo-processor architectures: Implementation and trade-offs, *Comput. Phys. Commun.* **185**, 1570 (2014).
- [63] A. Kantian, M. Dolfi, M. Troyer, and T. Giamarchi, Understanding repulsively mediated superconductivity of correlated electrons via massively parallel density matrix renormalization group, *Phys. Rev. B* **100**, 075138 (2019).
- [64] Y. Kurashige and T. Yanai, High-performance *ab initio* density matrix renormalization group method: Applicability to large-scale multireference problems for metal compounds, *J. Chem. Phys.* **130**, 234114 (2009).
- [65] G. K.-L. Chan, An algorithm for large scale density matrix renormalization group calculations, *J. Chem. Phys.* **120**, 3172 (2004).
- [66] A. Skaugen, Time evolution of magnetic Bloch oscillations using TEBD, Master thesis, University of Oslo, 2013, <https://www.duo.uio.no/handle/10852/38773>.
- [67] M. L. Wall, *Quantum Many-Body Physics of Ultracold Molecules in Optical Lattices*, Springer Theses (Springer, Cham, 2015).
- [68] M. Urbanek and P. Soldán, Parallel implementation of the time-evolving block decimation algorithm for the Bose–Hubbard model, *Comput. Phys. Commun.* **199**, 170 (2016).
- [69] V. Volokitin, I. Vakulchyk, E. Kozinov, A. Liniov, I. Meyerov, M. Ivanchenko, T. Lapyeva, and S. Denisov, Propagating large open quantum systems towards their asymptotic states: cluster implementation of the time-evolving block decimation scheme, *J. Phys.: Conf. Ser.* **1392**, 012061 (2019).
- [70] E. M. Stoudenmire and S. R. White, Real-space parallel density matrix renormalization group, *Phys. Rev. B* **87**, 155137 (2013).
- [71] S. Depenbrock, Tensor networks for the simulation of strongly correlated systems, Ph.D. thesis, Ludwig-Maximilians-Universität München, 2013, <https://edocub.uni-muenchen.de/15963/>.
- [72] H. Ueda, Infinite-size density matrix renormalization group with parallel Hida’s algorithm, *J. Phys. Soc. Jpn.* **87**, 074005 (2018).
- [73] M. Lubasch, J. I. Cirac, and M.-C. Bañuls, Unifying projected entangled pair state contractions, *New J. Phys.* **16**, 033014 (2014).
- [74] M. Lubasch, J. I. Cirac, and M.-C. Bañuls, Algorithms for finite projected entangled pair states, *Phys. Rev. B* **90**, 064425 (2014).
- [75] I. V. Oseledets, Tensor-train decomposition, *SIAM J. Sci. Comput.* **33**, 2295 (2011).
- [76] D. Perez-Garcia, F. Verstraete, M. M. Wolf, and J. I. Cirac, Matrix product state representations, *Quantum Inf. Comput.* **7**, 401 (2007).
- [77] R. Penrose, Applications of negative dimensional tensors, in *Combinatorial Mathematics and its Applications*, edited by D. J. A. Welsh (Academic, London, 1971).

- [78] J. C. Bridgeman and C. T. Chubb, Hand-waving and interpretive dance: An introductory course on tensor networks, *J. Phys. A: Math. Theor.* **50**, 223001 (2017).
- [79] S. Szalay, M. Pfeffer, V. Murg, G. Barcza, F. Verstraete, R. Schneider, and Ö. Legeza, Tensor product methods and entanglement optimization for *ab initio* quantum chemistry, *Int. J. Quantum Chem.* **115**, 1342 (2015).
- [80] J. Dongarra, M. Gates, A. Haidar, J. Kurzak, P. Luszczek, S. Tomov, and I. Yamazaki, The singular value decomposition: Anatomy of optimizing an algorithm for extreme scale, *SIAM Rev.* **60**, 808 (2018).
- [81] I. P. McCulloch, From density-matrix renormalization group to matrix product states, *J. Stat. Mech.* (2007) P10014.
- [82] G. M. Crosswhite and D. Bacon, Finite automata for caching in matrix product algorithms, *Phys. Rev. A* **78**, 012356 (2008).
- [83] G. M. Crosswhite, A. C. Doherty, and G. Vidal, Applying matrix product operators to model systems with long-range interactions, *Phys. Rev. B* **78**, 035116 (2008).
- [84] F. Fröwis, V. Nebendahl, and W. Dür, Tensor operators: Constructions and applications for long-range interaction systems, *Phys. Rev. A* **81**, 062337 (2010).
- [85] K. Levenberg, A method for the solution of certain non-linear problems in least squares, *Quart. Appl. Math.* **2**, 164 (1944).
- [86] D. W. Marquardt, An algorithm for least-squares estimation of nonlinear parameters, *J. Soc. Ind. Appl. Math.* **11**, 431 (1963).
- [87] The MathWorks, Inc., Least-Squares (Model Fitting) Algorithms - MATLAB & Simulink, <https://www.mathworks.com/help/optim/ug/least-squares-model-fitting-algorithms.html> (accessed 29 October 2019).
- [88] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevB.101.235123> for a comparison of different methods of approximating power laws by sums of exponentials, as well as details of our numerical experiments.
- [89] A. D. McLachlan, A variational solution of the time-dependent Schrödinger equation, *Mol. Phys.* **8**, 39 (1964).
- [90] J. Broeckhove, L. Lathouwers, E. Kesteloot, and P. Van Leuven, On the equivalence of time-dependent variational principles, *Chem. Phys. Lett.* **149**, 547 (1988).
- [91] S. Holtz, T. Rohwedder, and R. Schneider, On manifolds of tensors of fixed TT-rank, *Numer. Math.* **120**, 701 (2012).
- [92] A. Uschmajew and B. Vandereycken, The geometry of algorithms using hierarchical tensors, *Linear Algebra Its Appl.* **439**, 133 (2013).
- [93] J. Haegeman, M. Mariën, T. J. Osborne, and F. Verstraete, Geometry of matrix product states: Metric, parallel transport, and curvature, *J. Math. Phys.* **55**, 021902 (2014).
- [94] N. Hatano and M. Suzuki, Finding exponential product formulas of higher orders, in *Quantum Annealing and Other Optimization Methods*, edited by A. Das and B. K. Chakrabarti, Lecture Notes in Physics (Springer, Berlin, Heidelberg, 2005), pp. 37–68.
- [95] S. R. White, Density matrix renormalization group algorithms with a single center site, *Phys. Rev. B* **72**, 180403(R) (2005).
- [96] S. Dolgov and D. Savostyanov, Alternating minimal energy methods for linear systems in higher dimensions, *SIAM J. Sci. Comput.* **36**, A2248 (2014).
- [97] C. Hubig, I. P. McCulloch, U. Schollwöck, and F. A. Wolf, Strictly single-site DMRG algorithm with subspace expansion, *Phys. Rev. B* **91**, 155115 (2015).
- [98] D. Jaschke, M. L. Wall, and L. D. Carr, Open source matrix product states: Opening ways to simulate entangled many-body quantum systems in one dimension, *Comput. Phys. Commun.* **225**, 59 (2018).
- [99] T. J. Park and J. C. Light, Unitary quantum time evolution by iterative Lanczos reduction, *J. Chem. Phys.* **85**, 5870 (1986).
- [100] A. Al-Mohy and N. Higham, Computing the action of the matrix exponential, with an application to exponential integrators, *SIAM J. Sci. Comput.* **33**, 488 (2011).
- [101] S. Goto and I. Danshita, Performance of the time-dependent variational principle for matrix product states in the long-time evolution of a pure state, *Phys. Rev. B* **99**, 054307 (2019).
- [102] C. Hubig, J. Haegeman, and U. Schollwöck, Error estimates for extrapolations with matrix-product states, *Phys. Rev. B* **97**, 045125 (2018).
- [103] J. Schachenmayer, B. P. Lanyon, C. F. Roos, and A. J. Daley, Entanglement Growth in Quench Dynamics with Variable Range Interactions, *Phys. Rev. X* **3**, 031015 (2013).
- [104] R. Singh, R. Moessner, and D. Roy, Effect of long-range hopping and interactions on entanglement dynamics and many-body localization, *Phys. Rev. B* **95**, 094205 (2017).
- [105] A. Lerose and S. Pappalardi, Origin of the slow growth of entanglement entropy in long-range interacting systems, *Phys. Rev. Research* **2**, 012041 (2020).
- [106] F. Liu, R. Lundgren, P. Titum, G. Pagano, J. Zhang, C. Monroe, and A. V. Gorshkov, Confined Quasiparticle Dynamics in Long-Range Interacting Quantum Spin Chains, *Phys. Rev. Lett.* **122**, 150601 (2019).
- [107] The University of Bath, Balena HPC cluster, accessed 29 October 2019, <https://www.bath.ac.uk/corporate-information/balena-hpc-cluster/>.
- [108] S. Al-Assam, S. R. Clark, D. Jaksch, and TNT Development team, Tensor Network Theory Library beta version 1.2.1, 2016, <https://gitlab.physics.ox.ac.uk/tntlibrary>.
- [109] C. Goodyer, TNT Library: Tensor Manipulation and Storage, 2013, <http://www.hector.ac.uk/cse/distributedcse/reports/UniTNT/UniTNT/index.html>.
- [110] S. Al-Assam, S. R. Clark, and D. Jaksch, The tensor network theory library, *J. Stat. Mech.* (2017) 093102.
- [111] J. Coulthard, Engineering quantum states of fermionic many-body systems, D.Phil. thesis, University of Oxford, 2018, <https://ora.ox.ac.uk/objects/uuid:c2ad8834-7202-45cf-8348-4f9b07e942d1>.
- [112] The Mathworks, Inc., MATLAB version 9.3.0.713579 (R2017b), <https://www.mathworks.com/products/matlab.html>.
- [113] R. Lehoucq, K. Maschhoff, D. Sorensen, and C. Yang, ARPACK - Arnoldi Package, <https://www.caam.rice.edu/software/ARPACK/>.
- [114] opencollab, ARPACK-NG version 3.4.0, <https://github.com/opencollab/arpack-ng>.
- [115] Intel® Math Kernel Library (Intel® MKL) version 64/18.0.128, <https://software.intel.com/mkl>.
- [116] Intel® MPI Library version 64/18.0.128, <https://software.intel.com/mpi-library>.
- [117] E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen, *LAPACK Users' Guide*,

- 3rd ed. (SIAM, Philadelphia, PA, 1987), <https://www.netlib.org/lapack/lug/>.
- [118] J. J. M. Cuppen, A divide and conquer method for the symmetric tridiagonal eigenproblem, *Numer. Math.* **36**, 177 (1980).
- [119] M. Gu and S. C. Eisenstat, A stable and efficient algorithm for the rank-one modification of the symmetric eigenproblem, *SIAM J. Matrix Anal. Appl.* **15**, 1266 (1994).
- [120] J. D. Rutter, A serial implementation of Cuppen's divide and conquer algorithm for the symmetric eigenvalue problem, Tech. Rep. No. UCB/CSD-94-799, EECS Department, University of California, Berkeley, 1994, <https://www2.eecs.berkeley.edu/Pubs/TechRpts/1994/6315.html>.
- [121] The MathWorks, Inc., Numerically evaluate double integral - MATLAB integral2, accessed 29 October 2019, <https://www.mathworks.com/help/matlab/ref/integral2.html>.
- [122] A. Nauts and R. E. Wyatt, New Approach to Many-State Quantum Dynamics: The Recursive-Residue-Generation Method, *Phys. Rev. Lett.* **51**, 2238 (1983).
- [123] D. J. Luitz and Y. B. Lev, The ergodic side of the many-body localization transition, *Ann. Phys. (Berl.)* **529**, 1600350 (2017).
- [124] M. Kormos, M. Collura, G. Takács, and P. Calabrese, Real-time confinement following a quantum quench to a non-integrable model, *Nat. Phys.* **13**, 246 (2017).
- [125] E. H. Lieb and D. W. Robinson, The finite group velocity of quantum spin systems, *Commun. Math. Phys.* **28**, 251 (1972).
- [126] P. Calabrese and J. Cardy, Time Dependence of Correlation Functions Following a Quantum Quench, *Phys. Rev. Lett.* **96**, 136801 (2006).
- [127] S. Bravyi, M. B. Hastings, and F. Verstraete, Lieb-Robinson Bounds and the Generation of Correlations and Topological Quantum Order, *Phys. Rev. Lett.* **97**, 050401 (2006).
- [128] L. Cevolani, J. Despres, G. Carleo, L. Tagliacozzo, and L. Sanchez-Palencia, Universal scaling laws for correlation spreading in quantum systems with short- and long-range interactions, *Phys. Rev. B* **98**, 024302 (2018).
- [129] C.-F. Chen and A. Lucas, Finite Speed of Quantum Scrambling with Long Range Interactions, *Phys. Rev. Lett.* **123**, 250605 (2019).
- [130] T. Kuwahara and K. Saito, Strictly linear light cones in long-range interacting systems of arbitrary dimensions, [arXiv:1910.14477v3](https://arxiv.org/abs/1910.14477v3) [quant-ph].
- [131] J. Eisert, M. van den Worm, S. R. Manmana, and M. Kastner, Breakdown of Quasilocality in Long-Range Quantum Lattice Models, *Phys. Rev. Lett.* **111**, 260401 (2013).
- [132] Z.-X. Gong, M. Foss-Feig, S. Michalakis, and A. V. Gorshkov, Persistence of Locality in Systems with Power-Law Interactions, *Phys. Rev. Lett.* **113**, 030602 (2014).
- [133] D.-M. Storch, M. van den Worm, and M. Kastner, Interplay of soundcone and supersonic propagation in lattice models with power law interactions, *New J. Phys.* **17**, 063021 (2015).
- [134] D. J. Luitz and Y. Bar Lev, Emergent locality in systems with power-law interactions, *Phys. Rev. A* **99**, 010105(R) (2019).
- [135] J. Haegeman (private communication).
- [136] S. Singh, R. N. C. Pfeifer, and G. Vidal, Tensor network states and algorithms in the presence of a global U(1) symmetry, *Phys. Rev. B* **83**, 115125 (2011).
- [137] F. D. M. Haldane, Exact Jastrow-Gutzwiller Resonating-Valence-Bond Ground State of the Spin- $\frac{1}{2}$ Antiferromagnetic Heisenberg Chain with $1/r^2$ Exchange, *Phys. Rev. Lett.* **60**, 635 (1988).
- [138] B. S. Shastry, Exact Solution of an $S = 1/2$ Heisenberg Antiferromagnetic Chain with Long-Ranged Interactions, *Phys. Rev. Lett.* **60**, 639 (1988).
- [139] F. D. M. Haldane and M. R. Zirnbauer, Exact Calculation of the Ground-State Dynamical Spin Correlation Function of a $S = 1/2$ Antiferromagnetic Heisenberg Chain with Free Spinons, *Phys. Rev. Lett.* **71**, 4055 (1993).
- [140] J. P. F. LeBlanc *et al.* (Simons Collaboration on the Many-Electron Problem), Solutions of the Two-Dimensional Hubbard Model: Benchmarks and Results from a Wide Range of Numerical Algorithms, *Phys. Rev. X* **5**, 041041 (2015).
- [141] F. A. Y. N. Schröder and A. W. Chin, Simulating open quantum dynamics with time-dependent variational matrix product states: Towards microscopic correlation of environment dynamics and reduced system evolution, *Phys. Rev. B* **93**, 075105 (2016).
- [142] D. A. Abanin, E. Altman, I. Bloch, and M. Serbyn, Colloquium: Many-body localization, thermalization, and entanglement, *Rev. Mod. Phys.* **91**, 021001 (2019).
- [143] A. Safavi-Naini, M. L. Wall, O. L. Acevedo, A. M. Rey, and R. M. Nandkishore, Quantum dynamics of disordered spin chains with power-law interactions, *Phys. Rev. A* **99**, 033610 (2019).
- [144] T. Chanda, P. Sierant, and J. Zakrzewski, Time dynamics with matrix product states: Many-body localization transition of large systems revisited, *Phys. Rev. B* **101**, 035148 (2020).
- [145] C. Brockt, F. Dorfner, L. Vidmar, F. Heidrich-Meisner, and E. Jeckelmann, Matrix-product-state method with a dynamical local basis optimization for bosonic systems out of equilibrium, *Phys. Rev. B* **92**, 241106(R) (2015).
- [146] J. Motruk, M. P. Zaletel, R. S. K. Mong, and F. Pollmann, Density matrix renormalization group on a cylinder in mixed real and momentum space, *Phys. Rev. B* **93**, 155139 (2016).
- [147] L. Pastori, M. Heyl, and J. C. Budich, Disentangling sources of quantum entanglement in quench dynamics, *Phys. Rev. Research* **1**, 012007(R) (2019)..
- [148] M. M. Rams and M. Zwolak, Breaking the Entanglement Barrier: Tensor Network Simulation of Quantum Transport, *Phys. Rev. Lett.* **124**, 137701 (2020).
- [149] C. Krumnow, J. Eisert, and Ö. Legeza, Towards overcoming the entanglement barrier when simulating long-time evolution, [arXiv:1904.11999v1](https://arxiv.org/abs/1904.11999v1) [cond-mat.stat-mech].
- [150] H. N. Phien, G. Vidal, and I. P. McCulloch, Infinite boundary conditions for matrix product state calculations, *Phys. Rev. B* **86**, 245107 (2012).
- [151] D. Bauernfeind and M. Aichhorn, Time dependent variational principle for tree tensor networks, *SciPost Phys.* **8**, 024 (2020).
- [152] G. Evenbly, Gauge fixing, canonical forms, and optimal truncations in tensor networks with closed loops, *Phys. Rev. B* **98**, 085155 (2018).
- [153] M. P. Zaletel and F. Pollmann, Isometric Tensor Network States in Two Dimensions, *Phys. Rev. Lett.* **124**, 037201 (2020).
- [154] M. Lubasch, P. Moinier, and D. Jaksch, Multigrid renormalization, *J. Comput. Phys.* **372**, 587 (2018).