Matrix Product State applications for the ALPS project

Michele Dolfi[a], Bela Bauer[b], Sebastian Keller[c], Alexandr Kosenkov[a], Timothée Ewart[d,e], Adrian Kantian[d], Thierry Giamarchi[d], Matthias Troyer[a]

[a]*Theoretische Physik, ETH Zurich, 8093 Zurich, Switzerland*
[b]*Station Q, Microsoft Research, Santa Barbara, California 93106-6105, USA*
[c]*Laboratorium für Physikalische Chemie, ETH Zurich, 8093 Zurich, Switzerland*
[d]*DPMC-MaNEP, University of Geneva, 24 Quai Ernest-Ansermet, CH-1211 Geneva, Switzerland*
[e]*Blue Brain Project, Brain Mind Institute, EPFL, Switzerland*

Abstract

The density-matrix renormalization group method has become a standard computational approach to the low-energy physics as well as dynamics of low-dimensional quantum systems. In this paper, we present a new set of applications, available as part of the ALPS package, that provide an efficient and flexible implementation of these methods based on a matrix-product state (MPS) representation. Our applications implement, within the same framework, algorithms to variationally find the ground state and low-lying excited states as well as simulate the time evolution of arbitrary one-dimensional and two-dimensional models. Implementing the conservation of quantum numbers for generic Abelian symmetries, we achieve performance competitive with the best codes in the community. Example results are provided for (i) a model of itinerant fermions in one dimension and (ii) a model of quantum magnetism.

Keywords: MPS, DMRG, ground state, time evolution
PACS: 02.70.-c, 05.10.Cc, 71.27.+a

PROGRAM SUMMARY

Program Title: ALPS MPS
Journal Reference:
Catalogue identifier:
Licensing provisions: Use of 'mps_optim', 'mps_tevol', 'mps_meas' or 'mps_overlap' requires citation of this paper. Use of any ALPS program requires citation of the ALPS [1] paper.
Programming language: C++, OpenMP for parallelization.
Computer: PC, HPC cluster
Operating system: Any, tested on Linux, Mac OS X and Windows
RAM: 100 MB - 100 GB.
Number of processors used: 1 - 24.
Keywords: MPS, DMRG, TEBD.
Classification: 7.7
External routines/libraries: ALPS [1, 2], BLAS/LAPACK, HDF5.
Nature of problem: Solution of quantum many-body systems is generally a hard problem. The many-body Hilbert space grows exponentially with the system size which limits exact diagonalization results to only 20 40 spins, and the fermionic negative sign problem limits the Quantum Monte Carlo methods to a few special cases.
Solution method: The matrix product states ansatz provides a controllable truncation of the Hilbert space which makes it currently the method of choice to investigate low-dimensional systems in condensed matter physics. Our implementation allows simulation of arbitrary one-dimensional and two-dimensional models and achieve performance competitive with the best codes in the community. We implement conservation of quantum numbers for generic Abelian symmetries.
Running time: 10s – 8h per sweep.

1. Introduction

The density matrix renormalization group method (DMRG) [3] is currently the method of choice to investigate low-dimensional systems in condensed matter physics. Its applications range

arXiv:1407.0872v2 [cond-mat.str-el] 14 Oct 2014

from the study of superconductivity, quantum magnetism and exotic phases of matter to simulation of quantum circuits.

On the theoretical side, a lot of progress has been made in understanding the success of this method on the basis of matrix product states (MPS) [4], which is the class of variational states underlying the DMRG method. This has facilitated the development of new algorithms, e.g. for real-time evolution [5–7], mixed-state (finite temperature) calculations [8, 9] as well as their combination [10–12]. These have been formulated directly in the language of MPS, and are now widely being applied and further developed.

Recently, there has also been a surge of applications of DMRG to two-dimensional systems [13]. While it was realized early on that DMRG is exponentially hard for systems in more than one dimension [14], modern DMRG codes have become sufficiently efficient to combat this exponential scaling and reach system sizes where reliable conclusions about the thermodynamic limit of the 2d system become feasible [15–20].

In this paper we report the release into the ALPS project [1, 2, 21] of a set of MPS codes developed within the Swiss HP2C initiative [22], whose target was to develop new efficient and productive applications for current and future high performance computers. These applications have been designed to achieve the high performance required for the study of two-dimensional systems, and additionally to bind to the generic and well-known ALPS model and lattice libraries, which makes our MPS codes appealing for researchers in various fields. The new code has thus far been employed to study condensed matter physics problems such as ordering in an SU(3) Heisenberg model [17], the analysis of critical points for a system of supersymmetric fermions [23], the analysis of interacting mesoscopic structures [24], spin liquid phases in a Kagome-lattice spin model [20], and recent method developments such as the applications of DMRG as impurity solver within DMFT [25] and the development a multigrid DMRG scheme [26],

A number of open-source implementations of DMRG are currently available [1, 27–31]. With this release, we add a set of applications that, while implementing a generic framework amenable to many extensions, offers performance on par with the best available codes (see Table 1 for a timing example). Furthermore, it is seamlessly integrated with the widely used ALPS libraries.

Table 1: Runtime of our code compared with other state-of-the-art condensed matter DMRG codes. The benchmark is a run of 6 sweeps with a maximum bond dimension of $M = 200$ that increases linearly with every sweep for an antiferromagnetic Heisenberg chain with $L = 100$ sites. All codes are compiled with GCC 4.8.1, optimization option -O3 and linked against the sequential version of Intel MKL v11.1.1. Runtime is measured on a 2.2 GHz Intel Ivy Bridge EP E5-2660v2 processor. While we have done our best to set the parameters for all codes such as to ensure comparable executions, these timings should be understood as examples and results may depend significantly on the parameters for each code.

Code	Time
ALPS MPS (our new code)	16 sec
ALPS DMRG [1, 21]	73 sec
ITensor [27]	24 sec
OSMPS [28]	40 sec

The outline of this paper is as follows: In Section 2 we first focus on the characteristics and conventions of the new MPS code. In Section 3 we explain the main application parameters and show the results of a tutorial. For the theoretical background on MPS and DMRG we point the reader to the review [32] and references therein.

2. Algorithms and implementation

The matrix product state (MPS) is a variational ansatz wave function which for a chain of N sites reads

$$|\psi\rangle = \sum_{\boldsymbol{\sigma}} A_1^{\sigma_1} A_2^{\sigma_2} \cdots A_N^{\sigma_N} |\boldsymbol{\sigma}\rangle, \qquad (1)$$

where the exponentially growing number of parameters of the full quantum state is efficiently truncated to a polynomial number $\propto LM^2$ contained in the $M \times M$ matrices $A_i^{\sigma_i}$. This introduces an approximation that can be controlled by the matrix size M.

Matrix-product states are an example of a tensor network state. Such states are compactly represented graphically, as shown in Fig. 1. A tensor network is represented as a graph where vertices represent tensors, and edges represent indices that are contracted. Thus, in the upper panel of Fig. 1, the blue circles represent the tensors $(A_i^{\sigma_i})_{\alpha_{i-1}\alpha_i}$ of Eqn. (1); the edges connecting the tensors represent the indices that are summed over to yield matrix multiplication; and the open vertices pointing down
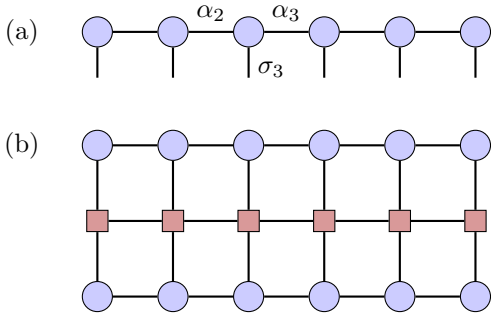
Figure 1: (a) Graphical representation of the MPS. Vertices (blue dots) represent the trivalent tensor on the i'th site, $(A_i^{\sigma_i})_{\alpha_{i-1}\alpha_i}$, where edges indicate indices. Closed edges, such as the horizontal lines indicate the indices α_{i-1}, α_i that are being contracted; vertical, open edges indicate the physical indices σ_i. (b) Contraction of the observable $\langle\psi|\hat{O}|\psi\rangle$, where $\hat{O}$ is represented as a matrix-product operator (red squares).

represent the physical indices σ_i. The lower panel of the figure shows a closed tensor network, i.e. without any open legs, that thus corresponds to a scalar. The contraction shown yields the calculation of an expectation value for an operator represented by a matrix-product operator (see Section 2.1).

It has been shown that ground states of gapped Hamiltonians in one dimension are accurately represented by a matrix-product state with a matrix size that grows only polynomially in the system size [33, 34]; if only local properties are required, one can obtain even better scaling [35]. Intuitively, this result rests on the fact that an area law holds for these systems [36], i.e. the bipartite entanglement entropy in these states is a function only of the area of the cut between two regions and not the total system size. In one dimension, this leads to constant bipartite entanglement entropy for gapped states.

The MPS ansatz is inherently a one-dimensional open chain. To simulate systems on periodic systems, ladders or even two-dimensional systems, a mapping from these lattices to a chain has to be chosen. This will lead to a Hamiltonian with longer-ranged interactions on the chain, where the range of interactions is related to the width of the ladder or 2d system. More importantly, the entanglement structure of the state is affected. The area of a cut bipartitioning the system into two halves will generally grow with the linear size L of the system. Since the bond dimension required to capture entropy S in an MPS is exponential, $M \sim e^S$, this implies an exponential cost for MPS simulations for two-

dimensional systems. The specific mapping from the 2d system to the chain therefore may have great impact on the accuracy and performance of simulations, and should therefore be chosen very carefully to minimize entanglement across each cut.

2.1. Representation of Matrix Product Operators

To introduce the notion of matrix product operators, it is useful to consider a 1d quantum lattice system with the same Hilbert space $\mathcal{H}$ on each site. Let $\{\hat{O}_k^\beta\}$ be the set of Hermitian operators that can act on the k'th site; as an example, these are identity matrix and the three Pauli matrices, $\{\mathbb{I}, \sigma^x, \sigma^y, \sigma^z\}$ for the case of $S = 1/2$ spins. We can then write any Hermitian operator $\hat{Q}$ acting on the entire lattice as

$$\hat{Q} = \sum c(\beta_1, \dots) \, \hat{O}_1^{\beta_1} \otimes \hat{O}_2^{\beta_2} \otimes \dots. \qquad (2)$$

This makes use of the fact that the operators themselves form a Hilbert space, analogous to quantum state. The above expression suggests that the coefficients $c(\alpha_1, \dots)$ can be decomposed into a matrix product, analogous to Eqn. 1. This representation is referred to as matrix-product operator (MPO). Analogous to a matrix-product state, such a matrix-product operator has an associated matrix size. Local Hamiltonians can generally be represented as an MPO of low bond dimension; even for long-range Hamiltonians, an efficient MPO description is often possible.

Consider now writing a local Hamiltonian, say for free fermions,

$$\hat{H} = -\sum_{\langle ij \rangle} \left(\hat{c}^\dagger(i) \, \hat{c}(j) + \text{h.c.} \right)$$

in this form. For simplicity, let's consider this model on three sites but with additional next-nearest neighbor hopping:

$$\hat{H} = -\hat{c}_1^\dagger \hat{c}_2 + \hat{c}_1 \hat{c}_2^\dagger - \hat{c}_2^\dagger \hat{c}_3 + \hat{c}_2 \hat{c}_3^\dagger - \hat{c}_1^\dagger \hat{c}_3 + \hat{c}_1 \hat{c}_3^\dagger \quad (3)$$

The most naive way to perform this decomposition is to assign an independent auxiliary index to each term in each of the sums:

$$H = \sum_{k,k'} A_{1k} B_{kk'} C_{k'1} \qquad (4)$$

$$A_{1k} = \left(-\hat{c}_1^\dagger, \hat{c}_1, -\hat{\mathbb{I}}, \hat{\mathbb{I}}, -\hat{c}_1^\dagger, \hat{c}_1 \right)_k \qquad (5)$$

$$B_{kk'} = \left(\hat{c}_2, \hat{c}_2^\dagger, \hat{c}_2^\dagger, \hat{c}_2, \hat{\mathbb{I}}, \hat{\mathbb{I}} \right)_k \delta_{kk'} \qquad (6)$$

$$C_{k1} = \left(\hat{\mathbb{I}}, \hat{\mathbb{I}}, \hat{c}_3, \hat{c}_3^\dagger, \hat{c}_3, \hat{c}_3^\dagger \right)_k \qquad (7)$$
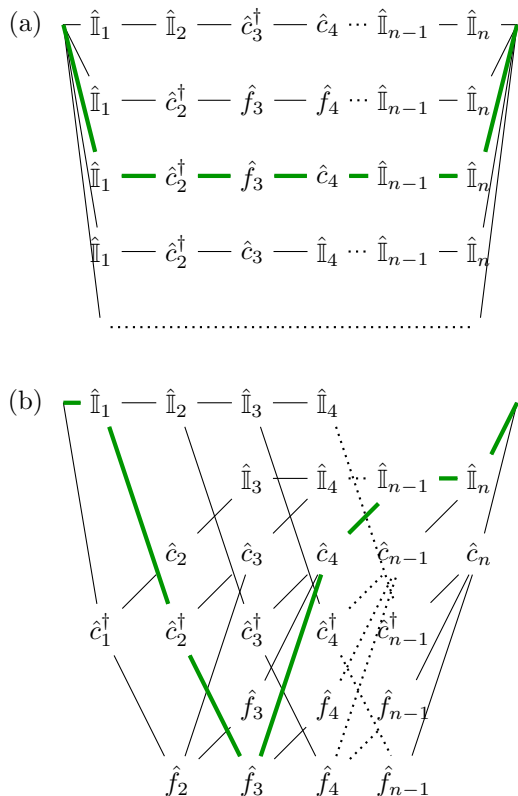
3

Figure 2: Internal structure of the MPO representing the the first part $\hat{c}^\dagger(i)\,\hat{c}(j)$ of the Hamiltonian (3) on a 3×3 square lattice. (a) trivial graph with an auxiliary bond for each term. (b) compressed version of the same graph where common branches are computed only once. For comparison, we highlight the term $\hat{c}^\dagger(2)\,\hat{c}(4)$ in both panels.

This leads to an MPO of bond dimension 6, which is very inefficient and redundant. A more efficient version is the following:

$$\tilde{A}_{1k} = \left(\hat{\mathbb{1}}, \hat{c}_1^\dagger, \hat{c}_1\right)_k \qquad (8)$$

$$\tilde{B}_{kk'} = \begin{pmatrix} 0 & \hat{c}_2 & -\hat{c}_2^\dagger \\ -\hat{c}_2 & 0 & -\hat{\mathbb{1}} \\ \hat{c}_2^\dagger & \hat{\mathbb{1}} & 0 \end{pmatrix}_{kk'} \qquad (9)$$

$$\tilde{C}_{k1} = \left(\hat{\mathbb{1}}, \hat{c}_3^\dagger, \hat{c}_3\right)_k. \qquad (10)$$

This MPO implements the same Hamiltonian, but has only bond dimension 3 by making use of the fact that the operators $\hat{c}_1^\dagger$, $\hat{c}_1$ as well as $\hat{c}_3^\dagger$, $\hat{c}_3$ occur both in the nearest- as well as next-nearest neighbor hopping term. Note that in the notation above, we have taken the elements of the matrices to be fermionic operators themselves. The antisymmetry of the central matrix in this case is a special prop-

erty of the 3-site ring and reflects the anticommutation of fermionic operators. In the actual implementation, fermionic statistics is implemented via a Jordan-Wigner transformation.

Our code implements an optimized MPO network that merges common branches, as illustrated in the above example, thus reducing the matrix size of the MPO substantially. This is sketched in the lower panel of Figure 2 (for a more extensive explanation of similar notation, see Ref. [37]). Note that in principle, the same compression techniques that are available for reducing the matrix size of MPSs can also be applied to MPOs; however, in some cases constructing the MPO in a naive way and then compressing is prohibitively time-consuming. Therefore, the code is written to perform many optimizations when the MPO is first constructed.

2.1.1. Fermionic operators

Special treatment is needed to account for the anti-commutation relations $\{\hat{c}(i), \hat{c}^\dagger(j)\} = \delta_{ij}$ characteristic of fermions. All fermionic operators have to be converted into a string of bosonic operators according to the Jordan-Wigner transformation [38]:

$$\hat{c}(i) = \left(\prod_{k<i} \hat{f}(k)\right)\hat{b}(i), \qquad (11)$$

where $\hat{f}(k)$ is the fermion sign operator, a diagonal matrix with a -1 entry for each local basis state with an odd number of fermions, and 1 for each local basis state with an even (including 0) number of fermions, and $\hat{b}(i)$ are hard-core boson operators satisfying the commutation relation $[\hat{b}(i), \hat{b}^\dagger(j)] = \delta_{ij}$ and $(\hat{b}(i))^2 = 0$. The fermion sign matrix $\hat{f}(k)$ can be written as $\hat{f}(k) = 1 - 2\hat{b}^\dagger(k)\hat{b}(k)$. To represent spinful fermions, two flavors of bosons $\hat{b}_1(k)$, $\hat{b}_2(k)$ have to be introduced and a normal-ordering of operators has to be chosen; in the ALPS libraries, the ordering is chosen as $\hat{b}_1(1)\hat{b}_2(1)\hat{b}_1(2)\hat{b}_2(2)....$ The Jordan-Wigner string must then include fermion sign operators for both flavors of bosons, i.e. two per physical site.

Noting that $\hat{f}^2 = \mathbb{1}$ and $[\hat{f}(i), \hat{b}^{(\dagger)}(j)] = 0$ for $i \neq j$, we find that for a fermionic hopping term the transformation simplifies to just a sequence of $\hat{f}(k)$ matrices between the two sites i and j:

$$\hat{c}(i)\hat{c}^\dagger(j) = \hat{b}(i)\left(\prod_{i\leq k<j} \hat{f}(k)\right)\hat{b}^\dagger(j). \qquad (12)$$

4

A similar simplification will apply to other types of fermionic bond terms; bond terms where each term is diagonal, such as density-density interactions, do not require Jordan-Wigner transformation.

2.2. Variational optimization of the MPS

Given the variational nature of the MPS ansatz, the ground state of a model described by the Hamiltonian $\hat{H}$ is obtained by an optimization process of all variational parameters looking for the energy minimum.

The most common way to optimize an MPS, as already described by S. White in '92 [3] is to optimize one or two tensors at the time, keeping all the others constant. For the case of the single site optimization the equation to solve is

$$\frac{\partial}{\partial A_i^*}\left(\langle\psi|\hat{H}|\psi\rangle - \lambda\left[\langle\psi|\,\psi\rangle - 1\right]\right) = 0, \qquad (13)$$

which maps to solving for the lowest eigenvector of the generalized eigenvalue problem $\mathcal{M}A_i - \lambda\mathcal{N}A_i = 0$, which is solved iteratively with the Jacobi-Davison algorithm provided by the IETL package in ALPS. Note that in the iterative solver the *matrix* $\mathcal{M}$ is never computed, but we efficiently contract all tensor legs on the fly in order to keep a complexity $\mathcal{O}(M^3)$.

Convergence is reached after sweeping a few times through the system. One *sweep* is here defined as optimizing twice all tensors in the chain, i.e. moving from left to right and then from right to left.

In our `mps_optim` application we implement the so called *single-site* and *two-site* optimizations. The first one optimizes one tensor at the time, whereas the latter one fuses two sites together to optimize the bond between them. The user is free to choose between them by changing the `optimization` input parameter to `singlesite` or `twosite`, respectively. It is strongly advised to always play around with both, in order to find what fits better the system to be solved.

In the case of the single site optimization we implement the procedure proposed in Ref. [39]. Here, the two shortcomings of the single-site optimization – the possibility of trapping on local minima and the inability to change the bond dimension dynamically – are remedied by enlarging the reduced density matrix after each optimization using

$$\tilde{\rho}^A = \mathrm{Tr}_B|\psi\rangle\langle\psi| + \alpha\sum_{b_l}\mathrm{Tr}_B\hat{H}_{b_l}^A|\psi\rangle\langle\psi|\hat{H}_{b_l}^A.$$
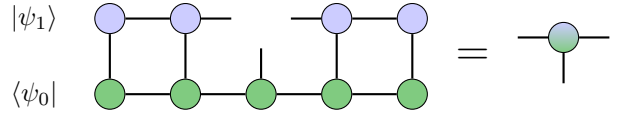

Figure 3: Construction of the environment tensor used as a local orthogonality constraint. The orthogonality constraint shown in the left network corresponds to a *mixed* MPS tensor which is used to evaluate $\langle\psi_0|\psi_1\rangle$ in the local optimization.

Assuming a partition of the system into a sub-block A and sub-block B such that the Hamiltonian is rewritten as $\hat{H} = \sum_{b_l}\hat{H}_{b_l}^A\hat{H}_{b_l}^B$, the first term in $\tilde{\rho}^A$ is the usual reduced density matrix of the subsystem A, whereas the second term introduces some reshuffling of the truncated states that protects against trapping in the optimization into local minima. The parameter α is a small number – typically 10^{-4} – which is slowly taken to zero during sweeping. In our implementation we allow three values for it `alpha_initial`, `alpha_main` and `alpha_final` (see the parameters in Table 3 for further details).

2.2.1. Excited states

In traditional DMRG approaches, the calculation of excited states is implemented through multi-state targetting, where the reduced density matrix after each iteration is truncated to describe a number of low-lying states simultaneously. While this approach has been shown to work reasonably, a more controlled approach is to sequentially find low-lying states by first finding the ground state $|\psi_0\rangle$, then finding the lowest-energy state $|\psi_1\rangle$ that is orthogonal to the ground state, $\langle\psi_0|\psi_1\rangle = 0$, etc. This approach was first described in Ref. [40].

In the MPS context, the overlap of two states can be efficiently calculated; furthermore, when optimizing a local tensor, the orthogonality constraint can be expressed as a local constraint of orthogonalizing the MPS tensor against the environment tensor shown in Fig. 3. For higher excited states, several such constraints need to be taken into account. In an iterative eigensolver, such as the Jacobi-Davidson method used here, these constraints can be taken into account by projecting newly generated Krylov vectors (or, in the case of Jacobi-Davidson, search vectors) into the orthogonal complement of the space spanned by the constraints. When sweeping through the state, the environment of Fig. 3 can be updated at very low computational cost.

2.3. Time evolution

To simulate time evolution, we follow the approach of Refs. [6, 7, 41]. The key insight to performing time evolution using MPS is that instead of time evolution, one can also variationally optimize the difference between an MPS $|\psi_1\rangle$ and an MPS with a unitary operator applied, $U|\psi_0\rangle$; that is, the optimization problem

$$\min_{|\psi_1\rangle} \left\| |\psi_1\rangle - U|\psi_0\rangle \right\| \qquad (14)$$

can be solved efficiently over normalized MPS. Here, the operator U can either act locally on a few sites, or be a sum of terms expressed as an MPO. Another key principle is that the evolution under a local Hamiltonian, $\hat{H} = \sum_n \hat{h}_n$, can be expressed as a product over local unitaries using a Trotter-Suzuki decomposition [42, 43]. Here, the time evolution operator $\exp(-i\hat{H}t)$ for a small time step Δt is decomposed into multiple products of the non-communing terms in the Hamiltonian. To first order the Trotter-Suzuki decomposition is

$$\exp(-i\hat{H}t) = \prod_{k=1}^{t/\Delta t} \prod_n e^{-i\hat{h}_n \Delta t} + \mathcal{O}(\Delta t^2), \qquad (15)$$

where $\hat{h}_n$ can themselves be sums of terms that do not overlap and hence commute. In the case of a simple nearest-neighbor Hamiltonian on a chain, $\hat{h}_1$ and $\hat{h}_2$ are the terms acting on even and odd bonds, respectively. In general better accuracy is obtained by second and fourth orders which are implemented in the `mps_evolve` application (see Appendix A.1).

Naively, the application of the unitary to an MPS increases the bond dimension to $M' = M \cdot D_h$, where $D_h > 1$ depends on the structure of the Hamiltonian; for details, see below. Therefore, after each contraction, the MPS matrices are truncated to their original size either with an singular value decomposition, or with a variational approximation (see Ref. [32]). Our application implements two types of tensor contractions for time evolution, *nearest neighbors* and *mpo*.

Nearest neighbors. Nearest neighbors time evolution algorithm is an optimized version used when the Hamiltonian contains only nearest neighbors bond term (in the case of ladders or 2d systems after having been mapped it to a Hamiltonian on a chain). The bond gates are contracted with the MPS tensors and a truncation via singular value decomposition is applied on the fly, following the TEBD method described in Ref. [41].
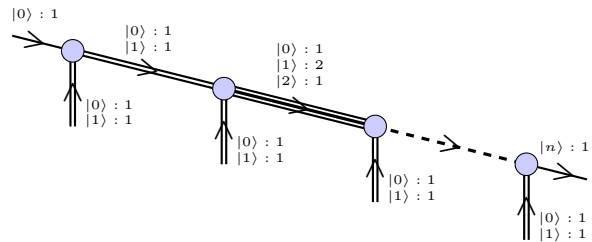

Figure 4: Labels in an MPS with conserved particles number. In the pair $(|k\rangle, s)$ we describe a charge block with label $|k\rangle$ (k particles) of size s. The MPS always starts with the trivial block $(|0\rangle, 1)$ and ends with the total charge sector that the user wants to target, n particles in this example.

MPO. MPO time evolution is a more generic algorithm where each $\exp(-i\hat{h}_n \Delta t)$ is transformed into an MPO, which allows to encode bond terms between non-neighboring sites. After multiplying the MPS wave function with the MPO, this is truncated to its original size M with a variational compression. Special care has to be taken for fermionic long-range operators, this procedure is explained in Appendix A.2.

2.4. Efficient tensor storage with abelian symmetries

Symmetries manifest themselves in many physical models, i.e. there are global operators $\hat{G}$ that commute with the system Hamiltonian, $[\hat{H}, \hat{G}] = 0$. In such a case, due to the fact that these operators can be simultaneously diagonalized, $\hat{H}$ becomes block-diagonal in a basis of eigenstates of $\hat{G}$. The blocks of the Hamiltonian are labelled by *quantum numbers*, corresponding to possible eigenvalues of $\hat{G}$. It is often computationally favorable to diagonalize the Hamiltonian separately within each quantum number sector.

In the case of DMRG, this is particularly easy when the symmetry has local generators, such that the reduced density matrices for any part of the system can also be block-diagonalized. Since the bond indices of the matrix product state essentially enumerate eigenstates of reduced density matrices, these can be associated with quantum number sectors. Each matrix of the MPS then obeys the symmetry locally, i.e. there are local conservation laws, and the matrices also become block-diagonal.

For each bond of the tensor network, a gauge choice has to be made fixing the direction of the bond. When considering, for example, particle number conservation, this is a choice of which bonds

are considered going into and which coming out of the node of the network [44]. The choice we make in our implementation is shown in Fig. 4. The MPS always starts with a trivial index on the left (e.g. only one sector with zero particles), new states are then added on every site from the local Hilbert space σ_i, and eventually the last tensor is fixed to the total quantum number sector chosen by the user.

Given our convention, locally the conservation of quantum numbers reads

$$\sigma_i \circ \alpha_i = \alpha_{i+1} \quad \text{and} \quad \alpha_i = (-\sigma_i) \circ \alpha_{i+1},$$

where α_i, σ_i are quantum numbers and $\circ$ represents the group product. The most convenient representation for the MPS tensors is in a matrix-like form where the physical index σ_i is fused together with one of the auxiliary indexes α_i or α_{i+1}. We then profit from the fact that the matrix $A_{\sigma_i \alpha_i; \alpha_{i+1}}$ is block-diagonal in the quantum number labels, and we store only the dense blocks. Linear algebra operations acting on these dense matrices are dispatched to optimized BLAS and LAPACK libraries available on the user machine[1].

The MPS codes presented in this paper are implemented to make use of Abelian symmetries. Dealing with conserved quantum numbers involves an intense calculation of indexes labels. We therefore set the symmetry at compile-time to allow for a maximum amount of compiler optimization. By default, the code is compiled for $U(1)^{\otimes n}$ symmetry, where the value n is fixed during compilation by the macro `DMRG_NUMSYMM`, as well as $\mathbb{Z}_2$ symmetry[2]. Other symmetries may easily be added.

2.5. Parallelization

We parallelize our applications for shared memory architectures using the OpenMP compiler directives. The `parallel for` loop directive is employed, for example, in the contraction of the local MPS tensor with the local MPO terms. In this routine several loops over the size of the auxiliary MPO dimension operate on independent parts. Since the MPO auxiliary dimension grows linearly with the width of a 2d system, we profit from more parallelism for large systems. The OpenMP parallelization is enabled at compile-time with the CMake option `ALPS_ENABLE_OPENMP`, and can be controlled at run-time via standard environment variables such as `OMP_NUM_THREADS`.

Many vendors provide support for multi-threaded BLAS and LAPACK libraries. For large matrices sizes our application will profit automatically from this parallelization, too. However, special care must be taken when using both the OpenMP parallelization and a parallel BLAS/LAPACK library simultaneously; this will only lead to improved scaling when using a compiler that supports such nested parallelism.

3. Codes and examples

Our set of MPS tools is composed of an application for the calculation of ground states and low energy excitations (`mps_optim`), a time evolution application (`mps_evolve`) and two utilities (`mps_meas` and `mps_overlap`) to perform measurements on the wave functions generated with the previous programs.

The interface of the applications is consistent with other ALPS applications, i.e. the same simulation can easily be run with exact diagonalization, the MPS optimization or even Monte Carlo applications. On the outer level a scheduler reads an XML input file listing all parameter sets to be computed. These parameter files can be created with the standard ALPS tools: (1) converting text files to XML with the `parameter2xml` command line tool, (2) using the convenient Python functions available in the `pyalps` package or (3) generating a workflow from VisTrails [45, 46].

Each simulation generates three output files: `simulation.out.h5` contains the simulation parameters, the iteration results and the final measurements in HDF5 format (see Appendix C for the detailed HDF5 schemes); `simulation.out.xml` contains the simulation parameters and the final results in XML format conforming with the ALPS schemas [1, 2]; `simulation.out.chkp` is the simulation checkpoint file containing the simulation status and the final MPS wave function that can be used to restore the same simulation or as input for other calculations.

Together with the applications, in the ancillary files, we ship full example simulations to illustrate how to create parameter files, run simulations and analyze the data. In sections 3.2 and 3.3 we describe some of these results.

---

[1]The reader is redirected to the ALPS documentation [1, 2] for the exact configuration options.

[2]$\mathbb{Z}_2$ symmetry is implemented only for some particular models. Refer to Appendix B for details on how to use it.

3.1. Input parameters

In Table 2, Table 3 and Table 4 we list all parameters and the default values used by the applications. They are split between common MPS parameters (valid in all codes), optimization parameters (available only in `mps_optim`) and time evolution specific parameters (for `mps_evolve`). Models, lattice and measurements are parameters of the ALPS model and lattice libraries, we point the reader to the ALPS references [1, 2] for further details.

3.1.1. Initial states

The MPS applications can be initialized with four different types of initial states.

default. The default initial state constructs an MPS which contains all allowed quantum number sectors that fulfill the total quantum numbers chosen by the user. In each MPS tensor all sectors will have a maximum size of 5×5 and they will be filled with random numbers.

thin. The thin initial state is an extension of the default option. After constructing the default MPS the routine performs a compression retaining only a total 20 truncated states per site. This initial state reduces considerably the number of sectors that the code has to tackle in the warm up sweeps. Sometimes it leads to faster convergence.

basis state. A many body basis state is easily encoded in a product state as an MPS with bond dimension $M = 1$. With this option the user can design its own initial state listing the desired local quantum numbers in the input file as illustrated in Table 2.

input file. The checkpoint file generated by a simulation can easily be reused in a new simulation by specifying it in the `initfile` parameter. This option is often used to start a time evolution from the ground state of an other model.

3.1.2. Temporary memory

An MPS simulation easily requires more than 100 GB of memory for the storage of boundary objects, i.e. partial contractions of the MPS network needed to reduce the complexity of the variational optimization to only $\mathcal{O}(M^3)$.

Our codes are capable of efficiently move this temporary memory to disk. This option is enabled by giving a non-empty value to the parameter `storagedir`. Since the application will often write > 100GB of temporary data it is advised to use fast filesystems, e.g. running on a local hard disk instead of a network mount. Once a simulation finishes the user can safely erase all temporary data.

3.2. Example 1: Itinerant fermions in the Hubbard ladder

As an example for the ground state search we show results for the pairing properties of the Hubbard model on a two-leg open ladder. The Hamiltonian of this model is

$$\hat{H} = -t \sum_{\langle ij \rangle, \sigma} \left[\hat{c}_\sigma^\dagger(i) \hat{c}_\sigma(j) + \text{h.c.} \right] \\ + U \sum_i \hat{n}_\uparrow(i) \hat{n}_\downarrow(i), \qquad (16)$$

where $\hat{c}_\uparrow^\dagger(i)$ and $\hat{c}_\downarrow^\dagger(i)$ create a fermion at site i with spin up and down, respectively, and $\hat{n}_\sigma(i) = \hat{c}_\sigma^\dagger(i) \hat{c}_\sigma(i)$ counts the number of particles with spin $\sigma \in \{\uparrow, \downarrow\}$ on site i. For a relatively strong repulsive interaction $U/t = 8$ and an average filling $\overline{n} = 0.875$ this model is expected to fall into the Luther-Emery universality class [47, 48] with a spin gap and a gapless charge excitation. To leading order both the long-range pairing correlation function $D(|i - j|) = D(l) \sim l^{-1/K_\rho}$ – defined as

$$D(i, j) = \langle \Delta(i) \, \Delta^\dagger(j) \rangle \qquad (17)$$

with $\Delta^\dagger(i) = [\hat{c}_\uparrow^\dagger(i, 1) \hat{c}_\downarrow^\dagger(i, 2) - \hat{c}_\downarrow^\dagger(i, 1) \hat{c}_\uparrow^\dagger(i, 2)]$ creating a singlet on the ith-rung – and the long-range charge density wave correlation function $C(|i-j|) = C(l) \sim l^{-K_\rho}$ – defined as

$$C(i, j) = \langle \hat{n}(i) \, \hat{n}(j) \rangle - \langle \hat{n}(i) \rangle \langle \hat{n}(j) \rangle \qquad (18)$$

– decay algebraically with an exponent that depends only on the non-universal parameter K_ρ, which we have to determine from the simulation at the chosen model parameters.

Long-range correlations have to propagate through the whole system and are therefore easily suppressed when truncating to only a few DMRG states. To treat the DMRG truncation correctly we run several instances of the simulation with increasing bond dimension M, until the observables converges or we can confidently extrapolate their value to $M \to \infty$.

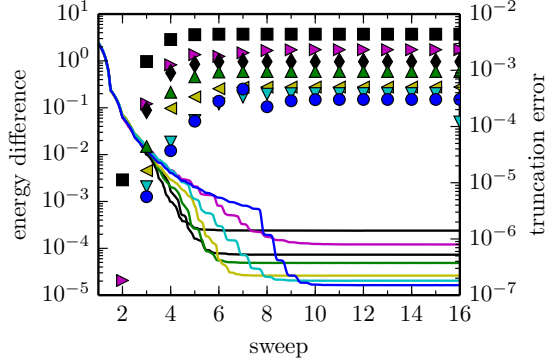In this example we simulate a 2×96 ladder with bond dimension varying from a computationally

Figure 5: Iteration history of the energy convergence (solid lines) compared to the extrapolated final energy obtained in Figure 6 and the corresponding truncation error (full points) for a 2×96 system with average filling $\overline{n} = 0.875$ and $U/t = 8$. From top to bottom the data shown correspond to a maximum bond dimension $M = 800, 1200, 1600, 2000, 2800, 3200, 3600$.
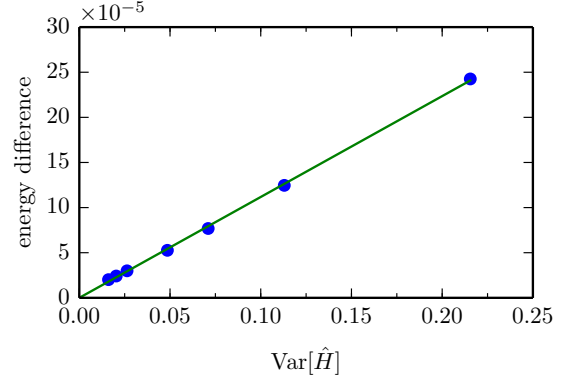

Figure 6: Extrapolation of the energy per particle as a function of the energy variance $\mathrm{Var}[\hat{H}] = \langle \hat{H}^2 \rangle - \langle \hat{H} \rangle^2$ for 2×96 system with average filling $\overline{n} = 0.875$ and $U/t = 8$. The linear extrapolation returns the ground state energy $e_0 = -0.72577\,t$ which has been subtracted from y-axis in the figure.

cheap $M = 800$ (convergence takes ca. 2.3h on 5 cores) to a computationally expensive $M = 3600$ (convergence takes 46h on 5 cores). The run script is:

```
parms = []
for M in [800, 1200, 1600, 2000, 2800, 3200,
    3600]:
    p = dict()
    p['SWEEPS'        ] = 16
    p['MAXSTATES'     ] = M
    p['init_state']  = 'thin'
    p['LATTICE'       ] = 'open ladder'
    p['L'             ] = 96
    p['MODEL_LIBRARY'] = 'mymodel.xml'
    p['MODEL'         ] = 'fermion Hubbard'
    p['t'             ] = 1
    p['U'             ] = 8
    p['CONSERVED_QUANTUMNUMBERS'] = 'Nup,Ndown'
    p['Nup_total'              ] = 84
    p['Ndown_total'            ] = 84
    p['MEASURE[EnergyVariance]'] = 1
    p['MEASURE_HALF_CORRELATIONS[pair field 1]']
        = 'field_du:fielddag_ud'
    parms.append(p)

## write the input file and run the simulation
infiles=pyalps.writeInputFiles('sim',parms)
res=pyalps.runApplication('mps_optim',infiles)
```

Results of these simulations are reported in Figures 5, 6, 7, where we illustrate the convergence history, the energy extrapolation and the decay of the pairing correlation function.

The iteration history of the simulation is loaded from the result file `simulation.out.h5` with the Python code:

```
resfiles = pyalps.getResultFiles(prefix='sim')
iters = pyalps.loadIterationMeasurements(
    resfiles, what=['Energy', 'TruncatedWeight
    '])
en_vs_iter = pyalps.collectXY(iters,
    x='iteration', y='Energy',
    foreach=['MAXSTATES'])
pyalps.plot.plot(en_vs_iter)
```

Because of the iterative optimization in DMRG, one should always check for convergence with the number of sweeps through the system. In Fig. 5, we see that both the ground state energy and the truncated weight have converged in our simulations.

Note that in the above parameters we enable the calculation of the energy variance $\mathrm{Var}[\hat{H}] = \langle \hat{H}^2 \rangle - \langle \hat{H} \rangle^2$. This is a very convenient quantity for extrapolating observables to $M \to \infty$, corresponding to $\mathrm{Var}[\hat{H}] \to 0$ for an eigenstate of $\hat{H}$. As an example the ground state energy depends to first order linearly on $\mathrm{Var}[\hat{H}]$. In Figure 6 we show the extrapolation of the energy generated with the following code:

```
## load data
resfiles = pyalps.getResultFiles(prefix='sim')
data = pyalps.loadEigenstateMeasurements(
    resfiles, ['Energy', 'EnergyVariance'])
en_vs_var = pyalps.ResultsToXY(data, '
    EnergyVariance', 'Energy')

## linear fit
xgrid = np.linspace(0, max(en_vs_var[0].y))
coeff = np.polyfit(en_vs_var[0].x, en_vs_var
    [0].y, deg=1)

## plot data and fit
pyalps.plot.plot(en_vs_var)
plt.plot(xgrid, np.polyval(coeff, xgrid))
```

The evaluation of the pairing correlation function $D(i, j)$ includes
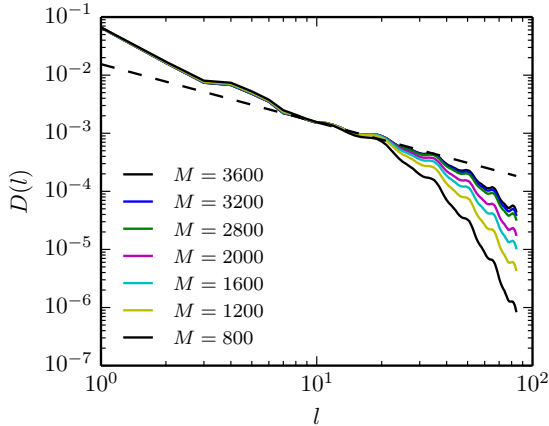
Figure 7: Rung-rung pairing correlation function $D(l)$ as a function of the distance between rungs for a 2×96 system with average filling $\bar{n} = 0.875$ and $U/t = 8$. Values at distance l are obtained by averaging 10 pairs $|i - j| = l$ around the center of the ladder. The dashed line is a reference algebraic decay with exponent 1.
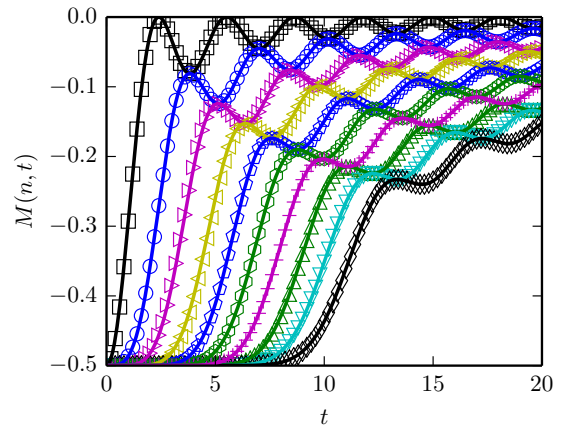

Figure 8: Magnetization $M(n, t) = \langle \psi(t) | \hat{S}_z(L/2 - n) | \psi(t) \rangle$ at time t and distance n from the middle of the chain. Points are simulation results with $n = 1 - 10$ (from top to bottom), solid lines are the corresponding exact solutions computed from Eqn. 20.

terms like `field_du : fielddag_up` $=$ $\hat{c}_\downarrow(i, 1)\hat{c}_\uparrow(i, 2)\hat{c}_\uparrow^\dagger(j, 1)\hat{c}_\downarrow^\dagger(j, 2)$ which was listed as a correlation measurement in the parameters listed above. Since `field_du` and `fielddag_ud` are defined in the file `mymodel.xml` [3] as bond operators, the application applies the four operators to all pairs of nearest neighbors sites, i.e. $i, i + 1, j, j + i$ for all $i = 1, \ldots, N - 4$ and $j = i + 2, \ldots, N$, effectively scaling as $\mathcal{O}(N^2)$.

Form the measured $D(i, j)$ we compute $D(l) = D(|i - j|)$ by averaging over a few rungs at distance $l = |i - j|$ around the center of the ladder. This technique averages out oscillations induced by the open boundaries. Results are shown in Figure 7, where one notes a very slow convergence of the pairing correlation function $D(l)$ with the retained DMRG states; for $M = 800$ pairing at large distances is underestimated. Large enough MPS matrices, like the $M = 3600$ data, are needed to correctly analyze the system.

From this example simulation we observe that $D(l)$ decays with an exponent consistent with $K_\rho > 1$, which is the condition for a dominant superconducting phase in the Luther-Emery universality class.

---

[3] The model description file mymodel.xml is available in the auxiliary files of this paper.

3.3. Example 2: Evolution of a domain wall

As a second example we run a simple time evolution simulation of the Heisenberg spin XX model, described by the Hamiltonian

$$\hat{H} = J \sum_i \left[S^x(i)S^x(i + 1) + S^y(i)S^y(i + 1) \right],$$
(19)

with $S^x(i)$, $S^y(i)$ the x, y component of the spin$-1/2$ operator $\vec{S}(i)$.

The time evolution of an initial domain wall $|\psi(t = 0)\rangle = |\downarrow\downarrow \ldots \downarrow\uparrow \ldots \uparrow\uparrow\rangle$ – a state with all up spins on the right side of the chain and the down spins on the left – under the spin XX Hamiltonian was solved exactly by a Jordan-Wigner transformation to free fermions in Ref. [49]. In particular, the local magnetization at distance n from the center after time t is described by a sum of Bessel functions as

$$M(n, t) = -\frac{1}{2} \sum_{i=1-n}^{n-1} j_i^2(t),$$
(20)

with $j_i(t)$ is the Bessel function of order i.

The setup script is very similar to the one previously used for the ground state calculations, for a simulation of 500 time steps up to a total time $\tau = 20\hbar/J$ it reads:

```
1  total_time = 20
2  nsweeps    = 500
3  p['MAXSTATES']     ] = 40
4  p['TIMESTEPS']     ] = nsteps
```

10

```
5   p['DT'                ] = nsteps/total_time
6   p['measure_each'      ] = 5
7   p['init_state'        ] = 'local_quantumnumbers'
8   p['initial_local_Sz'  ] = ','.join(   ['-0.5']*25
9                                        + ['0.5']*25)
10  p['ALWAYS_MEASURE'    ] = 'Local Magnetization'
11  p['LATTICE'           ] = 'open chain lattice'
12  p['L'                 ] = 50
13  p['MODEL'             ] = 'spin'
14  p['Jxy'               ] = 1
15  p['CONSERVED_QUANTUMNUMBERS'] = 'Sz'
16  p['Sz_total'          ] = 0
17  p['MEASURE_LOCAL[Local Magnetization]'] = 'Sz'

19  ## write the input file and run the simulation
20  infiles=pyalps.writeInputFiles('sim',parms)
21  res=pyalps.runApplication('mps_evolve',infiles
        )
```

Note the definition of the initial state `initial_local_Sz` listing the value of the S^z quantum number on each sites.

In Figure 8 the local magnetization calculated with the above simulation is compared with the exact solution from Eqn. 20 showing very good agreement. To produce a similar figure one simply has to load the results for each time step with the `loadIterationMeasurements` function and select the magnetization at distance n from the center, as shown in the following snippet.

```
1   ## simulation results
2   resfiles = pyalps.getResultFiles(prefix='sim')
3   data = pyalps.loadIterationMeasurements(
        resfiles, what=['Local Magnetization'])
4   ## select the Local Magnetization at distance
        `loc` from the center
5   numeric_mag = []
6   for d in pyalps.flatten(data):
7     L = d.props['L']
8     for loc in range(1, 11):
9       q = pyalps.DataSet()
10      q.props = deepcopy(d.props)
11      q.props['loc'] = loc
12      q.y = [ q.y[0][L/2-loc] ]
13      numeric_mag.append(q)
14  ## plot Magnetization as a function of time
15  mag_vs_time = pyalps.collectXY(numeric_mag,
16      x='Time', y='Local Magnetization',
17      foreach=['loc'])
18  pyalps.plot.plot(mag_vs_time)
```

A slightly more complex version generating a movie of the local magnetization per site is also provided in the source code of the example.

4. Acknowledgments

Appendix A. Algorithms for time evolution

Appendix A.1. Higher order Suzuki-Trotter decomposition

The time evolution unitary operator for a small time step $\exp(-i\hat{H}\Delta t)$ is conveniently transformed according to the Suzuki-Trotter decomposition. In the `mps_evolve` application the user can choose between a second or fourth order expansion with the parameter `te_order=second` or `te_order=fourth`, respectively.

The second order expands the unitary time evolution operator up to an error $\mathcal{O}(\Delta t^3)$, namely

$$\exp(-i\hat{H}\Delta t) = S\left(\frac{\Delta t}{2}\right) + \mathcal{O}(\Delta t^3), \qquad \text{(A.1)}$$

with $S(t)$ defined as

$$S(\Delta t) = \left(\prod_{n=1}^{N} e^{-i\hat{h}_n \Delta t}\right) \cdot \left(\prod_{n=N}^{1} e^{-i\hat{h}_n \Delta t}\right),$$

where the Hamiltonian operator $\hat{H}$ has been split into N non-commuting terms $\hat{h}_n$.

Fourth order expansion provides a more accurate time step $\mathcal{O}(\Delta t^5)$ at the expense of applying more operators to the MPS wave function. The decomposition reads

$$\exp(-i\hat{H}\Delta t) = \prod_{j=1}^{5} S\left(p_j \frac{\Delta t}{2}\right) + \mathcal{O}(\Delta t^5) \quad \text{(A.2)}$$

with $p_1 = p_2 = p_4 = p_5 = p = \frac{1}{4-4^{1/3}}$ and $p_3 = 1 - 4p$.

Appendix A.2. Exponential of fermionic bond terms

According to the Jordan-Wigner transformation an operator $\hat{O}(i,j) = \hat{c}_i \otimes \hat{c}_j$ acting on the bond between site i and j with the local operators $\hat{c}_i$ and $\hat{c}_j$, respectively, is decomposed into a chain of bosonic operators $\hat{b}_i$ and $\hat{b}_j$ with filling sign matrices $\hat{f}_k$ acting on sites $i < k < j$.

The exponential of this operator can be split into tensor product of exponential as

$$
\begin{aligned}
e^{\alpha \hat{O}(i,j)} &= \sum_n \frac{\alpha^n}{n!} \hat{b}_i^n \otimes \hat{f}_{i+1}^n \otimes \cdots \otimes \hat{b}_i^n \\
&= \sum_{n \in \text{even}} \frac{\alpha^n}{n!} \hat{b}_i^n \otimes \hat{\mathbb{I}}_{i+1} \otimes \cdots \otimes \hat{b}_i^n \\
&+ \sum_{n \in \text{odd}} \frac{\alpha^n}{n!} \hat{b}_i^n \otimes \hat{f}_{i+1}^n \otimes \cdots \otimes \hat{b}_i^n \\
&= \frac{1}{2} \left(B(\alpha) + F(\alpha) \right) + \frac{1}{2} \left(B(-\alpha) - F(-\alpha) \right),
\end{aligned}
\tag{A.3}
$$

with

$$
\begin{aligned}
B(\alpha) &= e^{\alpha \hat{b}_i} \otimes e^{\alpha \hat{\mathbb{I}}_{i+1}} \otimes \cdots \otimes e^{\alpha \hat{b}_j} \\
F(\alpha) &= e^{\alpha \hat{b}_i} \otimes e^{\alpha \hat{f}_{i+1}} \otimes \cdots \otimes e^{\alpha \hat{b}_j}.
\end{aligned}
$$

Appendix B. Extensions for $\mathbb{Z}_2$ symmetry

The additional $\mathbb{Z}_2$ symmetry, which is useful in the simulation of systems such as superconductors, can be enabled at compile-time using the CMake option `DMRG_BUILD_SYMMETRIES`, and setting `symmetry = "Z2"` in the parameter file. In Listing 1, we provide the skeleton of a model definition for a superconductor, where the $\mathbb{Z}_2$ symmetry is fermion parity, i.e. particle number modulo 2. The full example, along with a script that shows how to detect the topological phase transition in the Kitaev wire model [50], is included in the auxiliary files to this manuscript.

Appendix C. HDF5 schema

Simulation results are stored in the widely used Hierarchical Data Format (HDF). In this section we list the schemes used to store the simulation parameters, the measurement results and the iteration information. The ALPS Python module `pyalps` provides helper functions which load the data automatically.

Parameters

`/parameters/PARAMETERNAME`
   Value of the parameter *PARAMETERNAME*.

Helper function:
`pyalps.loadProperties(files)`

Measurements

`/spectrum/results/OBSNAME/mean/value`
   Vector of expectation values of the observable *OBSNAME*, each entry in the vector is a different eigenstate.

`/spectrum/results/OBSNAME/labels`
   Labels for the x-coordinate of the observable.

Helper function:
`pyalps.loadEigenstateMeasurements(files)`

Iterations

`/spectrum/iteration/NUM/parameters/`
   `PARAMETERNAME`

   Value of the parameter *PARAMETERNAME* specific to the iteration *NUM*. Only modified parameters appear in this group.

`/spectrum/iteration/NUM/results/OBSNAME/`
   `mean/value`

   Expectation values of the observable *OBSNAME* at the iteration *NUM*.

`/spectrum/iteration/NUM/results/OBSNAME/`
   `labels`

   Labels for the x-coordinate of the observable.

Helper function:
`pyalps.loadIterations(files)`

References

References

[1] B. Bauer, et al. (ALPS Collaboration), The ALPS project release 2.0: open source software for strongly correlated systems, J. Stat. Mech. 2011 (05) (2011) P05001. doi:10.1088/1742-5468/2011/05/P05001.
[2] http://alps.comp-phys.org.
[3] S. R. White, Density matrix formulation for quantum renormalization groups, Phys. Rev. Lett. 69 (19) (1992) 2863–2866. doi:10.1103/PhysRevLett.69.2863.
[4] S. Östlund, S. Rommer, Thermodynamic Limit of Density Matrix Renormalization, Phys. Rev. Lett. 75 (1995) 3537. doi:10.1103/PhysRevLett.75.3537.
[5] G. Vidal, Efficient simulation of one-dimensional quantum many-body systems, Phys. Rev. Lett. 93 (4) (2004) 040502. doi:10.1103/PhysRevLett.93.040502.
[6] S. R. White, A. E. Feiguin, Real-time evolution using the density matrix renormalization group, Phys. Rev. Lett. 93 (7) (2004) 076401. doi:10.1103/PhysRevLett.93.076401.
[7] A. J. Daley, C. Kollath, U. Schollwöck, G. Vidal, Time-dependent density-matrix renormalization-group using adaptive effective hilbert spaces, J. Stat. Mech.-Theory E. 2004 (04) (2004) P04005. doi:10.1088/1742-5468/2004/04/P04005.

Listing 1: Superconductor with $\mathbb{Z}_2$ symmetry.

```xml
<MODELS>
  <SITEBASIS name="spinless fermion">
    <QUANTUMNUMBER name="P" min="0" max="1" type="fermionic"/>
    <OPERATOR name="c" matrixelement="1">
      <CHANGE quantumnumber="P" change="1"/>
    </OPERATOR>
    <!-- analogously define cdag -->
  </SITEBASIS>

  <!-- define BASIS -->
  <!-- define bond operators -->

  <HAMILTONIAN name="tsc">
    <BASIS ref="spinless fermion"/>
    <BONDTERM source="i" target="j">
      -t*fermion_hop(i,j) + D*pairing(i,j)
    </BONDTERM>
  </HAMILTONIAN>

</MODELS>
```

[8] F. Verstraete, J. J. García-Ripoll, J. I. Cirac, Matrix product density operators: Simulation of finite-temperature and dissipative systems, Phys. Rev. Lett. 93 (2004) 207204. doi:10.1103/PhysRevLett.93.207204.

[9] A. E. Feiguin, S. R. White, Finite-temperature density matrix renormalization using an enlarged hilbert space, Phys. Rev. B 72 (2005) 220401. doi:10.1103/PhysRevB.72.220401.

[10] M. Zwolak, G. Vidal, Mixed-state dynamics in one-dimensional quantum lattice systems: A time-dependent superoperator renormalization algorithm, Phys. Rev. Lett. 93 (2004) 207205. doi:10.1103/PhysRevLett.93.207205.

[11] T. Barthel, U. Schollwöck, S. R. White, Spectral functions in one-dimensional quantum systems at finite temperature using the density matrix renormalization group, Phys. Rev. B 79 (2009) 245101. doi:10.1103/PhysRevB.79.245101.

[12] I. Pizorn, V. Eisler, S. Andergassen, M. Troyer, Real time evolution at finite temperatures with operator space matrix product states, ArXiv e-prints arXiv:1305.0504.

[13] E. Stoudenmire, S. R. White, Studying two-dimensional systems with the density matrix renormalization group, Ann. Rev. Cond. Mat. Phys. 3 (1) (2012) 111–128. doi:10.1146/annurev-conmatphys-020911-125018.

[14] S. Liang, H. Pang, Approximate diagonalization using the density-matrix renormalization-group method: A two-dimensional perspective, Phys. Rev. B 49 (13) (1994) 9214. doi:10.1103/PhysRevB.49.9214.

[15] S. R. White, A. L. Chernyshev, Neél Order in Square and Triangular Lattice Heisenberg Models, Phys. Rev. Lett. 99 (12) (2007) 127004. doi:10.1103/PhysRevLett.99.127004.

[16] S. Yan, D. A. Huse, S. R. White, Spin-liquid ground state of the s = 1/2 kagome heisenberg antiferromagnet, Science 332 (6034) (2011) 1173–1176. doi:10.1126/science.1201080.

[17] B. Bauer, P. Corboz, A. M. Läuchli, L. Messio, K. Penc, M. Troyer, F. Mila, Three-sublattice order in the su(3) heisenberg model on the square and triangular lattice, Phys. Rev. B 85 (2012) 125116. doi:10.1103/PhysRevB.85.125116.

[18] S. Depenbrock, I. P. McCulloch, U. Schollwöck, Nature of the spin-liquid ground state of the s = 1/2 heisenberg model on the kagome lattice, Phys. Rev. Lett. 109 (2012) 067201. doi:10.1103/PhysRevLett.109.067201.

[19] H.-C. Jiang, Z. Wang, L. Balents, Identifying topological order by entanglement entropy, Nat Phys 8 (12) (2012) 902–905. doi:10.1038/nphys2465.

[20] B. Bauer, L. Cincio, B. P. Keller, M. Dolfi, G. Vidal, S. Trebst, A. W. W. Ludwig, Chiral spin liquid and emergent anyons in a Kagome lattice Mott insulator, ArXiv e-prints arXiv:1401.3017.

[21] a. Albuquerque, F. Alet, P. Corboz, P. Dayal, a. Feiguin, S. Fuchs, L. Gamper, E. Gull, S. Gurtler, a. Honecker, The ALPS project release 1.3: Open-source software for strongly correlated systems, J. Magn. Magn. Mater. 310 (2) (2007) 1187–1193. doi:10.1016/j.jmmm.2006.10.304.

[22] http://www.hp2c.ch.

[23] B. Bauer, L. Huijse, E. Berg, M. Troyer, K. Schoutens, Supersymmetric multicritical point in a model of lattice fermions, Phys. Rev. B 87 (2013) 165145. doi:10.1103/PhysRevB.87.165145.

[24] M. Cheng, M. Becker, B. Bauer, R. M. Lutchyn, Interplay between Kondo and Majorana interactions in quantum dots, ArXiv e-prints arXiv:1308.4156.

[25] H. Shinaoka, M. Dolfi, M. Troyer, P. Werner, Hybridization expansion monte carlo simulation of multi-orbital quantum impurity problems: matrix product formalism and improved sampling, J. Stat. Mech.-Theory E. 2014 (6) (2014) P06012. doi:10.1088/1742-5468/2014/06/P06012.

[26] M. Dolfi, B. Bauer, M. Troyer, Z. Ristivojevic, Multigrid algorithms for tensor network states, Phys.

Rev. Lett. 109 (2) (2012) 020604. `doi:10.1103/PhysRevLett.109.020604`.

[27] ITensor, `http://itensor.org`.

[28] Open Source MPS, `http://sourceforge.net/projects/openmps`.

[29] G. Alvarez, The density matrix renormalization group for strongly correlated electron systems: A generic implementation, Comput. Phys. Commun. 180 (9) (2009) 1572 – 1578. `doi:10.1016/j.cpc.2009.02.016`.

[30] BLOCK, `http://www.princeton.edu/chemistry/chan/software/dmrg`.

[31] S. Wouters, W. Poelmans, P. W. Ayers, D. Van Neck, CheMPS2: A free open-source spin-adapted implementation of the density matrix renormalization group for ab initio quantum chemistry, Comput. Phys. Commun. 185 (6) (2014) 1501–1514. `doi:10.1016/j.cpc.2014.01.019`.

[32] U. Schollwöck, The density-matrix renormalization group in the age of matrix product states, Ann. Phys. (N.Y.) 326 (1) (2011) 96–192. `doi:10.1016/j.aop.2010.09.012`.

[33] M. B. Hastings, An area law for one-dimensional quantum systems, J. Stat. Mech.-Theory E. 2007 (08) (2007) P08024. `doi:10.1088/1742-5468/2007/08/P08024`.

[34] N. Schuch, M. M. Wolf, F. Verstraete, J. I. Cirac, Entropy Scaling and Simulability by Matrix Product States, Phys. Rev. Lett. 100 (3) (2008) 030504. `doi:10.1103/PhysRevLett.100.030504`.

[35] F. Verstraete, J. I. Cirac, Matrix product states represent ground states faithfully, Phys. Rev. B 73 (2006) 094423. `doi:10.1103/PhysRevB.73.094423`.

[36] J. Eisert, M. Cramer, M. B. Plenio, Colloquium: Area laws for the entanglement entropy, Rev. Mod. Phys. 82 (2010) 277–306. `doi:10.1103/RevModPhys.82.277`.

[37] G. M. Crosswhite, D. Bacon, Finite automata for caching in matrix product algorithms, Phys. Rev. A 78 (2008) 012356. `doi:10.1103/PhysRevA.78.012356`.

[38] P. Jordan, E. Wigner, Über das Paulische Äquivalenzverbot, Z. Phys. 47 (9-10) (1928) 631–651. `doi:10.1007/BF01331938`.

[39] S. R. White, Density matrix renormalization group algorithms with a single center site, Phys. Rev. B 72 (2005) 180403. `doi:10.1103/PhysRevB.72.180403`.

[40] I. P. McCulloch, From density-matrix renormalization group to matrix product states, J. Stat. Mech.-Theory E. 2007 (10) (2007) P10014. `doi:10.1088/1742-5468/2007/10/P10014`.

[41] G. Vidal, Efficient classical simulation of slightly entangled quantum computations, Phys. Rev. Lett. 91 (14) (2003) 147902. `doi:10.1103/PhysRevLett.91.147902`.

[42] H. F. Trotter, On the product of semi-groups of operators, Proc. Amer. Math. Soc. 10 (1959) 545. `doi:10.1090/S0002-9939-1959-0108732-6`.

[43] M. Suzuki, Generalized trotter's formula and systematic approximants of exponential operators and inner derivations with applications to many-body problems, Commun. Math. Phys. 51 (2) (1976) 183–190. `doi:10.1007/BF01609348`.

[44] B. Bauer, P. Corboz, R. Orús, M. Troyer, Implementing global abelian symmetries in projected entangled-pair state algorithms, Phys. Rev. B 83 (2011) 125106. `doi:10.1103/PhysRevB.83.125106`.

[45] C. T. Silva, J. Freire, S. P. Callahan, Provenance for visualizations: Reproducibility and beyond, Comput. Sci. Eng. 9 (5) (2007) 82–89. `doi:10.1109/MCSE.2007.106`.

[46] `http://vistrails.org`.

[47] R. Noack, S. White, D. Scalapino, The ground state of the two-leg hubbard ladder a density-matrix renormalization group study, Physica C 270 (3–4) (1996) 281 – 296. `doi:10.1016/S0921-4534(96)00515-1`.

[48] A. Luther, V. Emery, Backward scattering in the one-dimensional electron gas, Phys. Rev. Lett. 33 (10) (1974) 589–592, cited By (since 1996)247. `doi:10.1103/PhysRevLett.33.589`.

[49] T. Antal, Z. Rácz, A. Rákos, G. M. Schütz, Transport in the xx chain at zero temperature: Emergence of flat magnetization profiles, Phys. Rev. E 59 (1999) 4912–4918. `doi:10.1103/PhysRevE.59.4912`.

[50] A. Y. Kitaev, Unpaired majorana fermions in quantum wires, Physics-Uspekhi 44 (10S) (2001) 131. `doi:10.1070/1063-7869/44/10S/S29`.

Table 2: Common simulation parameters

Parameter	Default	Description
MAXSTATES		Maximum size of the matrices M_{σ_i}
TRUNCATION	1e-8	Smallest singular value to be kept
seed	42	Seed of the internal random number generator
CONSERVED_QUANTUMNUMBERS		Comma-separated list of the quantum numbers to be conserved
*QN*_total		Total value of the quantum number *QN*. *QN* is one of the quantum numbers defined in the parameter CONSERVED_QUANTUMNUMBERS. Example: N_total = 8 fixes the total number of particles to 8.
initfile		Path to an existing checkpoint MPS to use it as initial MPS
init_state	default	Initial state of the simulation (in case initfile is not set). Possible values: 'default', 'thin', 'local_quantumnumbers'.
initial_local_*QN*		Comma-separated list with the value of the local quantum number *QN* at every site
MEASURE_LOCAL[*NAME*]		Defines a new local measurement called *NAME*. Its value specifies the operator to be measured. Example: MEASURE_LOCAL[Local density] = "n"
MEASURE_AVERAGE[*NAME*]		Defines a new average measurement called *NAME*. Its value specifies the operator to be measured. Example: MEASURE_AVERAGE[Density] = "n"
MEASURE_CORRELATIONS[*NAME*]		Defines a new correlation measurement called *NAME*. Its value specifies the operators to be correlated separated by ":" (colon) If the operators define two or more BONDOPERATORS, e.g., Op_A and Op_B the correlation measurement compute all observables $\langle\psi\vert$ Op$_A$$(i, i+1)$ Op$_B$$(j, j+1)\vert\psi\rangle$. Example: MEASURE_CORRELATIONS[Onebody Density Matrix] = "bdag:b"
MEASURE_HALF_CORRELATIONS[*NAME*]		Same as MEASURE_CORRELATIONS but it does not exchange the order of operators. If the input is, e.g., bdag:b, the first operator will be evaluated at all locatios $i \in [0, L-2]$ but the second operator only at locations $j \in [i, L-1]$.
MEASURE_LOCAL_AT[*NAME*]		Syntax for the value: "op_1:...:op_n | $(i_{1\,1}, \dots, i_{1\,n})$, $(i_{2\,1}, \dots, i_{2\,n})$, ...". Defines a new measurement called *NAME* where the sequence of operators op_1:...:op_n is applied to all tuples of indices (of length n like the number of operators) listed after the vertical bar symbol "—".
MEASURE[EnergyVariance]	False	Measure the energy variance
MEASURE[Entropy]	False	Measure the von Neumann entropy
MEASURE[Renyi2]	False	Measure $n = 2$ Renyi entropy
ALWAYS_MEASURE		Comma-separated list of measurements to evaluate at the end of every sweep
COMPLEX	False	Use complex numbers. For time evolution simulations the default value is changed to True.
storagedir		Path to the directory to be used for temporary storage. If empty, temporary storage is disabled.

Table 3: Additional parameters for 'mps_optim'

Parameter	Default	Description
SWEEPS		Number of sweeps
NUMBER_EIGENVALUES	1	Number of eigenstate to target
optimization	twosite	Optimisation algorithm. Possible values are 'single-site', 'twosite'.
ngrowsweeps		Single site only. Number of initial sweeps where the correction factor has the value alpha_initial.
nmainsweeps		Single site only. Number of sweeps after ngrowsweeps where the correction factor has the value alpha_main.
alpha_initial	1e-2	Correction factor for single site optimization [39] in the initial part
alpha_main	1e-4	Correction factor for single site optimization [39] in the main part
alpha_final	1e-8	Correction factor for single site optimization [39] in the final part
ietl_jcd_toll	1e-8	Convergence tolerance of the Jacobi-Davidson solver
ietl_jcd_gmres	0	Convergence tolerance of the Jacobi-Davidson solver
ietl_jcd_maxiter	8	Maximum number of iterations in the Jacobi-Davidson solver

Table 4: Additional parameters for 'mps_evolve'

Parameter	Default	Description
DT		Time step
IMG_TIMESTEPS	0	Number of imaginary time sweeps to perform before the proceeding with the real time evolution
TIMESTEPS		Total number of sweeps
te_order	fourth	Order of the trotter decomposition. Possible values: 'second', 'fourth'.
te_type	nearest neighbors	Type of time evolution algorithm. Possible values: 'nearest neighbors', 'mpo'.
chkp_each	1	A checkpoint is created every chkp_each time steps.
measure_each	1	Measurements are performed every measure_each time steps.
update_each	-1	Update the Hamiltonian parameters every update_each time steps. Used while quenching Hamiltonian parameters. A negative value will never change the Hamiltonian.
P [Time]		Values assigned to the parameter P at each time step