Sliced Basis Density Matrix Renormalization Group for Electronic Structure

E. Miles Stoudenmire and Steven R. White

Phys. Rev. Lett. **119**, 046401 — Published 24 July 2017

DOI: 10.1103/PhysRevLett.119.046401

Sliced Basis Density Matrix Renormalization Group for Electronic Structure

E. Miles Stoudenmire and Steven R. White
Department of Physics and Astronomy, University of California, Irvine, CA 92697-4575 USA

(Dated: June 9, 2017)

We introduce a hybrid approach to applying the density matrix renormalization group (DMRG) to continuous systems, combining a grid approximation along one direction with a finite Gaussian basis set for the remaining two directions. This approach is especially useful for chain-like molecules, where the grid is used in the long direction. For hydrogen chain systems, the computational time scales approximately linearly with the number of atoms, as we show with near-exact minimal basis set calculations with up to 1000 atoms. The linear scaling comes from both the localization of the basis and a compression method for the long-ranged two-electron interaction. For shorter hydrogen chains, we show results with up to triple ζ bases.

Thousands of publications are written every year dealing with electronic structure calculations of solids or molecules. Improved electronic structure algorithms can have an enormous impact. One of the more recent approaches for molecules is the density matrix renormalization group (DMRG) which originated as a method for model condensed matter systems. In its quantum chemistry form (QCDMRG) [1], it is able to treat some of the most strongly correlated molecules where most other methods fail [2]. This has led to important applications, such as elucidating the excited state structure of the oxygen evolving photosynthesis complex [3]. However, it is difficult to scale DMRG up to larger systems and more complete bases: about 100-200 active basis functions is the current practical limit. In this paper we dramatically extend the capability of QCDMRG for a particular class of systems: long chains. Surprisingly, this is accomplished by replacing the usual Gaussian basis in the long direction by a fine grid. Despite increasing the number of active orbitals by an order of magnitude, we obtain *linear scaling* of the calculation time with system length, which easily makes up for the increased basis size. The results we present here from our initial implementation include chains of up to 1000 hydrogen atoms—well beyond QCDMRG and other strongly correlated methods.

The Hilbert space used in QCDMRG is the same as that of the Hubbard model, when equating a Hubbard site with a single basis function. But the rapid growth of computation time with the number of basis functions in QCDMRG does not occur for the one-dimensional Hubbard model, where the calculation time is approximately linear (when keeping a fixed number of DMRG states). The main reason for the poor scaling of QCDMRG is the complexity of the Hamiltonian in the basis, particularly the two-electron terms. The electron-electron Coulomb interaction terms are defined by two-electron integrals

$$V_{ijkl} = \int_{\mathbf{r}_1} \int_{\mathbf{r}_2} \frac{\phi_i(\mathbf{r}_1)\phi_l(\mathbf{r}_1)\phi_j(\mathbf{r}_2)\phi_k(\mathbf{r}_2)}{|\mathbf{r}_1 - \mathbf{r}_2|} \tag{1}$$

where the $\phi_i(\mathbf{r})$ are orthonormal basis functions. If the basis functions are delocalized, as when using molecular orbitals from a Hartree Fock calculation, the number of significant V_{ijkl} terms scales as N_b^4, where N_b is the number of basis functions. This leads to a computation time for QCDMRG scaling as $N_b^4 m^2 + N_b^3 m^3$, where m is the number of many-body states kept [4].

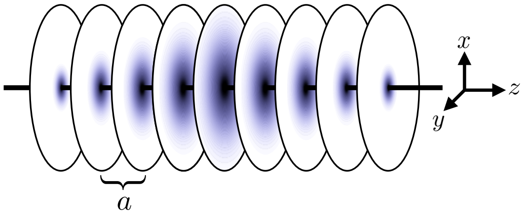

FIG. 1. The sliced basis set approach can be viewed as finely slicing the continuum into parallel two-dimensional planes, each spanned by a small set of transverse functions.

The nonlocality of the orthogonal basis functions also increases the number of states m needed for a given accuracy. DMRG is a low-entanglement approximation, and the entanglement of ground states is governed by the *area law* [5, 6]. The area law holds for ground states described in terms of local, "real space" degrees of freedom. In a delocalized basis, a volume law of entanglement holds instead (except for non-interacting systems, a special point where the entanglement is zero in the eigenstate basis). To capture volume-law states, m must grow exponentially with the system size, even in one dimension. For this reason, some effort should be made to localize the basis before applying standard QCDMRG, except on very small molecules. The localization is always imperfect—the basis functions have oscillating tails which tend to be slowly decaying.

Hypothetically, one could get rid of both the N_b^4 scaling and the increase in entanglement from extended basis functions by going to a real-space grid defined by finite differences. In such a grid the interactions are defined as $V_{ij}\hat{n}_i\hat{n}_j$, where $\hat{n}_i$ is the density operator on site i. For model one-dimensional continuum systems, this is currently the most powerful approach, and we have used it to simulate systems of 100 pseudo-hydrogen atoms with about 20 grid points per atom [7]. A key part of using a one-dimensional grid is compressing the interactions by

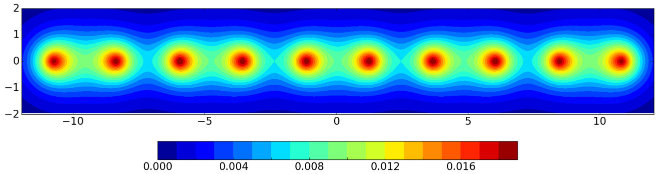

FIG. 2. Electronic density in the $y-z$ plane of a linear chain of 10 hydrogen atoms, equally spaced at a near neighbor distance $R = 2.4$ a.u., calculated in a sliced cc-pVDZ basis (with $N_o = 4$). Dimerization induced by the open boundaries is visible, representing the strong tendency to form H_2 molecules.

approximating long-range interactions as a sum of exponentials [8, 9]. With this compression, the calculation time grows only linearly with the number of atoms. The problem with such a grid approach for three dimensions is that the number of grid points would be very high, for example of order 10^6 for a system of modest size.

Here we introduce a hybrid approach, which we call sliced basis DMRG (SBDMRG). Along one particular "z" direction over which the molecule extends furthest we use a grid. At each grid point, the remaining transverse dimensions, x and y, are captured by a small number of basis functions derived from standard Gaussian basis sets, making what we call a "slice"—see Fig. 1. The total number of DMRG "sites" is therefore $N_b = N_z N_o$, where N_z is the number of grid points, and N_o is the number of transverse functions ("orbitals") per grid point. The DMRG path progresses through all orbitals on a slice, then moves to the next. This approach has several major advantages. First, all interaction terms V_{ijkl} where i and l are not on the same slice are zero, and similarly for j and k. Thus the number of terms scales as N_z^2. Second, the remaining interactions can be compressed very efficiently, making the dominant part of the calculation time linear in N_z. Third, since there is no spatial extent of the basis functions in the z direction, there is no extra entanglement due to nonlocality, potentially reducing the number of states m needed for a given accuracy. The cost of applying DMRG in this setup scales as $(N_z m^3 D + N_z m^3 N_o^3 + N_z m^2 N_o^4)$ where m is the number of states kept in DMRG, and D is a parameter controlling the compression of the long-range Coulomb interaction, which scales very weakly with N_z [10].

We demonstrate our method by simulating linear chains of hydrogen atoms. Although these are three-dimensional systems, their linear nature makes them especially well suited for both SBDMRG and QCDMRG. They also exhibit strong correlation, and can be quite challenging for electronic structure methods. The electronic density in a plane through the nuclei for a typical calculation is presented in Fig. 2.

To define the sliced basis approach, consider the electronic structure Hamiltonian for fixed nuclei in atomic units:

$$\hat{H}_{\rm el} = \int_{\mathbf{r}} \hat{\psi}^\dagger_\sigma(\mathbf{r}) \left[-\frac{1}{2}\nabla^2 + v(\mathbf{r}) \right] \hat{\psi}_\sigma(\mathbf{r}) + \frac{1}{2} \int_{\mathbf{r},\mathbf{r}'} \frac{1}{|\mathbf{r}-\mathbf{r}'|} \hat{\psi}^\dagger_\sigma(\mathbf{r}) \hat{\psi}^\dagger_{\sigma'}(\mathbf{r}') \hat{\psi}_{\sigma'}(\mathbf{r}') \hat{\psi}_\sigma(\mathbf{r}) \; . \tag{2}$$

Summation over spin labels σ is implied and $v(\mathbf{r})$ is the potential generated by the nuclei.

Along the z direction, make a grid approximation $z_n = n \cdot a$ with n an integer and a a small grid spacing. On each slice n, introduce a finite, orthonormal basis of functions $\{\phi_j(x,y)\}$ where $j = 1, 2, \ldots, N_o$. For simplicity here, assume the same N_o and functions $\{\phi_j(x,y)\}$ on every slice n. Later one can perform a change of basis to adapt the basis for each slice, possibly reducing the number of functions. In terms of discrete operators $\hat{c}^\dagger_{nj\sigma}$ and $\hat{c}_{nj\sigma}$ which create and destroy electrons in a slice orbital, the Hamiltonian takes the form

$$\hat{H} = \frac{1}{2} \sum_{nn'} \sum_{ij} t^{nn'}_{ij} \hat{c}^\dagger_{ni\sigma} \hat{c}_{n'j\sigma} \tag{3}$$

$$+ \frac{1}{2} \sum_{nn'} \sum_{ijkl} V^{nn'}_{ijkl} \hat{c}^\dagger_{ni\sigma} \hat{c}^\dagger_{n'j\sigma'} \hat{c}_{n'k\sigma'} \hat{c}_{nl\sigma} \; . \tag{4}$$

Defining $\boldsymbol{\rho} = (x, y)$ for convenience, the interaction integrals are

$$V^{nn'}_{ijkl} = \int_{\boldsymbol{\rho},\boldsymbol{\rho}'} \frac{\phi_i(\boldsymbol{\rho}) \phi_j(\boldsymbol{\rho}') \, \phi_k(\boldsymbol{\rho}') \phi_l(\boldsymbol{\rho})}{\sqrt{|\boldsymbol{\rho}-\boldsymbol{\rho}'|^2 + (z_n - z_{n'})^2}} \; . \tag{5}$$

The i, j, k, l indices only run over the small number of functions N_o on each slice. Thus, the Hamiltonian is defined by just $N_z^2 N_o^4$ interaction integrals. The single-particle couplings are defined as

$$t^{nn'}_{ij} = \delta_{nn'} \int_{\boldsymbol{\rho}} \phi_i(\boldsymbol{\rho}) \left[-\frac{1}{2}\nabla^2_{\boldsymbol{\rho}} + v(\boldsymbol{\rho}, z_n) \right] \phi_j(\boldsymbol{\rho}) \tag{6}$$

$$- \delta_{ij} \frac{1}{2a^2} \Delta_{nn'} \; . \tag{7}$$

Our discrete Hamiltonian treats the z-direction kinetic energy terms Eq. (7) on a different footing than the "integral" terms. For the z-direction kinetic energy, we treat the basis functions as smooth functions of z, and think of the slices as sampling them. Thus we use standard finite difference approximations, defined via $\Delta_{nn'}$. In the following we use an fourth-order approximation with grid error scaling as a^4. For the "integral" terms, we treat the basis functions as completely localized and nonoverlapping between slices The finite-difference grid approximation makes our results not strictly variational at finite a. But we find finite-a errors for hydrogen chains of only about 0.1 mH per atom for $a = 0.1$, and in the limit of $a \to 0$, the results are variational.

In what follows, we construct the transverse basis functions on a slice $\{\phi_j(x,y)\}$ out of standard atom-centered

Gaussian basis sets and assume the atoms are identical. In going from the spherical symmetry used in standard Gaussians to slices, we switch to cylindrical symmetry. Thus, an S-function becomes a σ function, P-functions become π functions, etc. Whereas there are $2\ell+1$ functions in a spherical set with angular momentum ℓ, there are only two cylindrical functions for any $\ell > 0$. For example, a set of D functions, with coefficient ζ, becomes the two slice basis functions

$$(x^2-y^2)\exp[-\zeta(x^2+y^2)] \tag{8}$$

$$2xy\exp[-\zeta(x^2+y^2)]\ . \tag{9}$$

We leave out functions like P_z, which looks like a σ function on a slice, or any other function looking like a function of smaller ℓ. (In principle, P_z could be kept as an additional S function.) The slice basis functions are only orthogonal between different slices. This means the functions within each slice must be orthogonalized.

In the parent 3D Gaussian bases, usually some of the functions (particularly S-type) are contracted, meaning out of N_g original Gaussians, one uses a smaller number N_o of linear combination of functions for each atom: $\phi^j = \sum_{m=1}^{N_g} c_m^j \exp[-\zeta_m(\vec{r}-\vec{r}_A)^2]$ where $j = 1 \ldots N_o$, and $N_o < N_g$. In this case, to define the transverse basis on a slice, we follow an approach that is useful very generally: we form a local orbital density matrix for each slice. Let i and i' run over an orthonormal uncontracted basis for the slice at z_n, defined by functions $\xi_i(x,y)$. Let $\phi^k(x,y,z)$ be a particular 3D contracted basis function attached to one of the atoms, and let

$$\eta_i^k = \int_{x,y} \phi^k(x,y,z_n)\xi_i(x,y) \tag{10}$$

$$\rho_{ii'} = \sum_k \eta_i^k \eta_{i'}^k. \tag{11}$$

The leading eigenvectors of ρ form optimal local functions for representing the contracted 3D basis. More generally, ρ could come from the interacting ground state, as a block of the single particle reduced density matrix $\langle c_i^\dagger c_{i'}\rangle$, and we would call the eigenvectors of ρ "slice natural orbitals" (SNOs). A subset with only N_j of these SNOs would be an ideal reduced local basis. Our procedure for contractions is conceptually similar, but with equal weighting for all 3D contracted basis functions. In this case, for example, the sharp Gaussians used to represent the nuclear cusps only appear significantly in the slices close to nuclei. In our hydrogen chain calculations, if the basis has N_S contracted S functions per atom, we keep N_S contracted functions per slice.

We perform DMRG with the Hamiltonian represented as a sum of matrix product operators (MPOs), one of which represents the long-ranged two-electron interactions. For this MPO we use a compression technique giving an MPO with matrix dimension D which is nearly independent of system length, resulting in a linear scaling computation time. (The other MPOs, say for $v(\mathbf{r})$, are naturally of constant dimension.) Consider the simplest case of a single basis function per slice such that the interactions Eq. (4) simplify to

$$\sum_{n\le n'} V_{nn'}\hat{n}_n\hat{n}_{n'}\ . \tag{12}$$

Here we focus on the compression of the upper triangle of the matrix $V_{nn'}$, giving just an outline here and more details in the Supplemental Material [11]. For a typical sliced basis, V is not translationally invariant; otherwise, an MPO could be constructed based on fitting $V(n-n')$ to a sum of exponentials [8]. Instead we present a more flexible method using a sequence of singular value decompositions (SVDs) that is a simplification of the very general approaches described in Refs. 12 and 13.

For a particular diagonal index k, let $V^{(k)}$ be the rectangular block of V with the lower left corner at V_{kk}, and extending to the upper right corner of V. An SVD gives

$$V^{(k)} = U^{(k)}S^{(k)}W^{(k)} \tag{13}$$

where $S^{(k)}$ is the diagonal matrix of singular values. The smoothness of $V(n-n')$ away from the diagonal makes this SVD have a small number D of significant singular values, allowing us to approximate $S^{(k)}$ as a $D\times D$ matrix, reducing the number of columns of $U^{(k)}$ and rows of $W^{(k)}$ accordingly.

This factorized representation at index k can be related to a similar representation at $k+1$. Define $P(U^{(k)})$ to be the direct sum of $U^{(k)}$ and a 1×1 identity matrix, that is add an extra column and row of zeros to the bottom and right of $U^{(k)}$ and set the new diagonal element to 1.

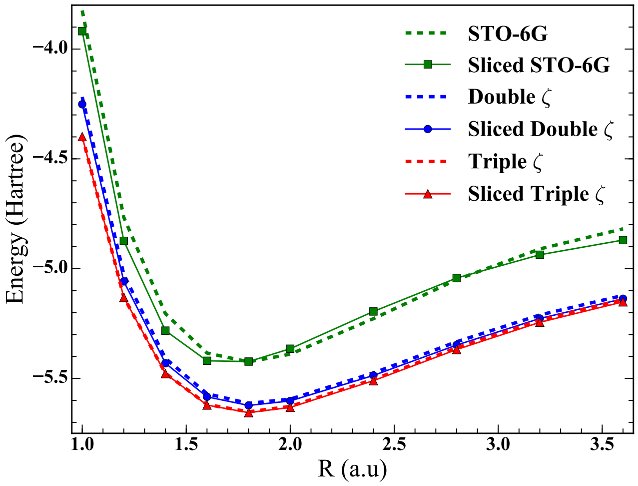

FIG. 3. Energy of linear chains of 10 hydrogen atoms, equally spaced by a distance R. Dashed lines show results using QCDMRG in standard basis sets. Solid lines with symbols are SBDMRG results in a sliced version of each basis set.

Then a matrix $X^{(k+1)}$ can be computed such that

$$U^{(k+1)} = P(U^{(k)})X^{(k+1)} \; . \tag{14}$$

The matrix $X^{(k+1)}$ has dimension $(D+1) \times D$. We see that we can recover all the $U^{(k)}$ from the $X^{(k)}$ and $U^{(1)}$. Similarly, all of the $W^{(k)}$ can be generated in terms of a reverse recursion involving $D \times (D+1)$ matrices $Y^{(k)}$. So we can reconstruct every $V^{(k)}$, and thus the entire $N_z \times N_z$ matrix V out of the $O(N_z D^2)$ parameters in $X^{(k)}$, $S^{(k)}$, and $Y^{(k)}$. In the Supplemental Material, we show how these parameters lead to an MPO representation of the interactions with MPO matrix dimension $D+2$.

Fig. 3 shows results for chains of 10 equally-spaced hydrogen atoms as a function of separation R, for several different basis sets with $a = 0.1$ and for comparison, standard QCDMRG results for parent 3D basis sets [14]. The STO-6G basis is a minimal basis, contracting 6 Gaussians to one function per atom; the sliced version also has one function per slice. One can see that the completeness of the standard and sliced bases are similar; which basis gives a lower energy varies with R. The double ζ basis (cc-pVDZ) has five functions per atom [15], and the sliced version has four per slice (no P_z). Here the energies are even closer, but the sliced version is consistently slightly lower. The triple ζ basis (cc-pVTZ) has 14 functions per atom, or 140 functions total, making this a somewhat challenging QCDMRG calculation. The sliced version has 9 functions per slice, with up to 561 slices. To get the SBDMRG total energy errors to within 1 mH took from 4-10 days (depending on R), with bond dimensions $m \sim 300 - 1000$, running on a 2013 quad core Mac mini with 16Gb. For triple ζ the sliced and non-sliced energies are also very close, but with the sliced version slightly lower. All DMRG calculations were performed using the ITensor library [16].

In Fig. 4, we present results for very long chains, demonstrating the linear scaling of SBDMRG. These calculations were at the stretched distance $R = 3.6$, using a sliced STO-6G basis with one basis function per slice, and grid spacing $a = 0.2$. The inset shows the calculation time per sweep on a single core of a 2013 3.5GHz Mac Pro, for a sweep keeping $m = 100$ states. The calculation time not only grows very close to linearly in the number of atoms, it is also quite modest. The largest system, with 1000 atoms, had over 18,000 sliced basis functions, and an $m = 100$ sweep took little more than an hour. We slowly ramped up the number of states kept, with 30 smaller-m, faster sweeps then three $m = 100$ sweeps. Subsequent sweeps up to $m = 400$ showed that at $m = 100$, the energy per atom was in error by only 0.06 mH (DMRG error only, excluding the finite basis and finite a errors). The main part of the figure shows the energy per site, in comparison with QCDMRG STO-6G. The energy results show the modest difference in completeness of STO-6G and sliced STO-6G, and also demonstrate that the sliced DMRG is converged to high accuracy.

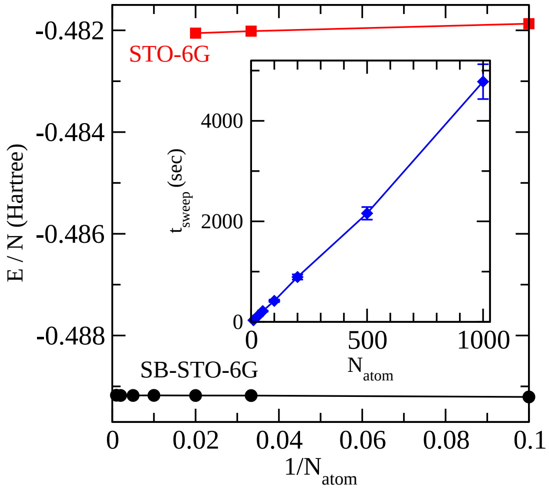

FIG. 4. Energy per atom of N hydrogen atom chains, equally spaced by a distance $R = 3.6$ a.u. using standard versus sliced STO-6G basis sets. Inset: average time per DMRG sweep with $m = 100$, demonstrating linear scaling up to $N = 1000$ atoms.

The sliced basis set approach we have introduced here is very well suited to DMRG calculations. Coupled with a compression method for the interactions, the approach gives linear scaling of computation time with system length, allowing very long systems to be treated. This formulation brings DMRG for electronic structure closer to DMRG for models, and new approaches introduced for models, such as working directly with an infinite chain [17, 18], or parallel DMRG which works best for systems with many sites [19], can probably be adapted to SBDMRG with little difficulty. We also anticipate that extending SBDMRG to more complicated molecules will be reasonably straightforward.

We acknowledge support from the Simons Foundation through the Many-Electron Collaboration, and from the U.S. Department of Energy, Office of Science, Basic Energy Sciences under award #DE-SC008696.

[1] Steven R. White and Richard L. Martin, "Ab initio quantum chemistry using the density matrix renormalization group," The Journal of Chemical Physics **110**, 4127–4130 (1999).

[2] Garnet Kin-Lic Chan and Sandeep Sharma, "The density matrix renormalization group in quantum chemistry," Annual Review of Physical Chemistry **62**, 465–481 (2011).

[3] Yuki Kurashige, Garnet Kin-Lic Chan, and Takeshi Yanai, "Entangled quantum electronic wavefunctions of the Mn_4CaO_5 cluster in photosystem II," Nat Chem **5**, 660–666 (2013).

[4] In long molecules, with minimal basis sets, truncation of the interactions can improve the scaling of QCDMRG to $O(N_b^2)$.

[5] G. Evenbly and G. Vidal, "Tensor network states and geometry," Journal of Statistical Physics **145**, 891–918 (2011).

[6] M B Hastings, "An area law for one-dimensional quantum systems," J. Stat. Mech. **2007**, P08024 (2007).

[7] E. M. Stoudenmire, Lucas O. Wagner, Steven R. White, and Kieron Burke, "One-dimensional continuum electronic structure with the density-matrix renormalization group and its implications for density-functional theory," Phys. Rev. Lett. **109**, 056402 (2012).

[8] Gregory M. Crosswhite, A. C. Doherty, and Guifré Vidal, "Applying matrix product operators to model systems with long-range interactions," Phys. Rev. B **78**, 035116 (2008).

[9] B. Pirvu, V. Murg, J. I. Cirac, and F. Verstraete, "Matrix product operator representations," New J. Phys. **12**, 025012 (2010).

[10] For a fixed transverse basis set size, increasing N_z by an order of magnitude barely increases the bond dimension D of the compressed long-range interaction MPO. As an example, for a compression of hydrogen chains of N atoms in a sliced double-ζ basis, to obtain the same compression error level requires $D = 50$ for $N = 2$; $D = 51$ for $N = 10$; and $D = 53$ for $N = 100$. Thus the scaling of D with N_z is possibly logarithmic or even weaker.

[11] See Supplemental Material for more information about computing integrals for a sliced basis and compressing the interaction terms into an MPO. The Supplemental Material includes the additional Refs. [20, 21].

[12] Michael P. Zaletel, Roger S. K. Mong, Frank Pollmann, and Edward H. Rezayi, "Infinite density matrix renormalization group for multicomponent quantum hall systems," Phys. Rev. B **91**, 045115 (2015).

[13] Garnet Kin-Lic Chan, Anna Keselman, Naoki Nakatani, Zhendong Li, and Steven R. White, "Matrix product operators, matrix product states, and ab initio density matrix renormalization group algorithms," The Journal of Chemical Physics **145**, 014102 (2016).

[14] Boxiao Zheng, Ushnish Ray, and Garnet Chan, Private communication.

[15] Thom H. Dunning Jr., "Gaussian basis sets for use in correlated molecular calculations. I. the atoms boron through neon and hydrogen," The Journal of Chemical Physics **90**, 1007–1023 (1989).

[16] ITensor Library (version 2.0.11) http://itensor.org.

[17] I. P. McCulloch, "Infinite size density matrix renormalization group, revisited," arxiv:0804.2509 (2008).

[18] R. Orús and G. Vidal, "Infinite time-evolving block decimation algorithm beyond unitary evolution," Phys. Rev. B **78**, 155117 (2008).

[19] E. M. Stoudenmire and Steven R. White, "Real-space parallel density matrix renormalization group," Phys. Rev. B **87**, 155137 (2013).

[20] Ian P. McCulloch, "From density-matrix renormalization group to matrix product states," J. Stat. Mech. , P10014 (2007).

[21] Gregory M. Crosswhite and Dave Bacon, "Finite automata for caching in matrix product algorithms," Phys. Rev. A **78**, 012356 (2008).