

This is the accepted manuscript made available via CHORUS. The article has been published as:

## Real-space parallel density matrix renormalization group

E. M. Stoudenmire and Steven R. White

Phys. Rev. B **87**, 155137 — Published 22 April 2013

DOI: [10.1103/PhysRevB.87.155137](https://doi.org/10.1103/PhysRevB.87.155137)

# Real-Space Parallel Density Matrix Renormalization Group

E.M. Stoudenmire and Steven R. White

*Department of Physics and Astronomy, University of California, Irvine, CA 92697*

(Dated: April 3, 2013)

We demonstrate how to parallelize the density matrix renormalization group (DMRG) algorithm in real space through a straightforward modification of serial DMRG. This makes it possible to apply at least an order of magnitude more computational power to challenging simulations, greatly accelerating investigations of two-dimensional systems and large parameter spaces. We discuss details of the algorithm and present benchmark results including a study of valence-bond-solid order within the square-lattice  $Q_2$  model and Néel order within the triangular lattice Heisenberg model. The parallel DMRG algorithm also motivates an alternative canonical form for matrix product states.

PACS numbers: 05.10.Cc, 02.70.-c, 05.30.-d

## I. INTRODUCTION

The density matrix renormalization group (DMRG) is a method for computing ground states of one-dimensional (1d) systems<sup>1–3</sup> which has recently proven surprisingly effective for studying two-dimensional (2d) models, especially those beyond the reach of quantum Monte Carlo due to frustrated interactions or mobile fermions.<sup>4–14</sup> DMRG is also valuable in quantum chemistry where it has extended the ability of previous methods to deal with strong correlation.<sup>15–17</sup> But applying DMRG to 2d and quantum chemical systems is computationally very demanding. Even in 1d, DMRG can be costly for systems having many degrees of freedom.<sup>18–20</sup> Finding an efficient way to divide a single DMRG calculation across multiple processors would make larger calculations tractable, but a truly parallel DMRG algorithm has remained an outstanding problem.

Previous DMRG implementations have achieved a limited form of parallelism by parallelizing over different terms in the Hamiltonian.<sup>21</sup> This is effective in a quantum chemistry context where the number of terms is especially large. However, since the terms contained entirely within a single spatial region are already combined as much as possible within DMRG, the efficiency of this approach is less than ideal. A similar type of limited parallelism subdivides matrices into quantum number blocks.<sup>22,23</sup> But this approach is restricted by a limited range of quantum numbers and a large variation in sub-block size. The least powerful, but simplest form of parallelism breaks up dense matrix computations into sub-blocks and is performed automatically by standard linear algebra libraries.

Here we present a much more powerful form of parallelism, dividing a single DMRG calculation over separate regions of the system in real space. Except at the boundary of each region, the algorithm reduces to standard finite-size DMRG, making it relatively straightforward to implement in existing, highly-optimized DMRG codes. Real-space parallelism can be used independently of the other types discussed above and becomes increasingly effective as the system size increases. In practice, we ob-




FIG. 1. Timing of parallel DMRG ground state calculations for the spin 1/2 Heisenberg model on the  $24 \times 8$  square lattice (cylindrical boundary conditions<sup>14</sup>). A speedup of  $S$  indicates that the calculation using  $n$  nodes was  $S$  times faster than the same calculation using only one node. Each calculation consisted of 10 sweeps and reached a relative energy accuracy of  $10^{-5}$  by keeping  $m = 2000$  states in the final two sweeps. The speedup is so close to ideal for the two-node case because the reflection symmetry permits optimal load balancing and minimizes waiting time at the communication step.

serve close to ideal speedups as shown in Fig. 1.

Note that some extensions of DMRG are already real-space parallelizable, such as time-dependent DMRG algorithms based on factorizations of the time-evolution operator, as in a Suzuki-Trotter approximation.<sup>24,25</sup> These algorithms can be used to compute ground states through imaginary time evolution. But imaginary time evolution is much less efficient than the usual DMRG diagonalization-based method, therefore it is primarily useful when the standard DMRG ground state approach is not applicable (for example, when optimizing 2d tensor networks such as PEPS<sup>26–28</sup>).

After introducing the parallel DMRG algorithm in section II, we use it to compute the strength of the valence-bond-solid order in the pure  $Q_2$  model on the square




FIG. 2. The parallel DMRG algorithm for the simplest case of dividing a single calculation over two machines. Circles represent lattice sites (or matrix product state site-tensors) and the diamond is the matrix  $V$  on the shared bond as in Eq. (4). In step (a), the local information needed to sweep left is copied to the left machine and similarly for the right. Each machine then performs a DMRG sweep (b) in parallel over its half of the system. In step (c) the wavefunctions are merged together using the prescription preceding Eq. (5). Before repeating the algorithm, the merged state is optimized (d) using Lanczos or Davidson on the shared bond.

lattice (section III A) and the strength of the Néel order in the Heisenberg model on the triangular lattice (section III B). Both applications involve calculations which would take many weeks using serial DMRG but require only a few days using our parallel approach. Finally, in section IV we discuss how real-space parallel DMRG motivates an alternative canonical gauge for matrix product states.

## II. PARALLEL DMRG ALGORITHM

To describe the algorithm, first consider parallelizing a single DMRG calculation over just two regions. For concreteness, take the system to have six sites such that the center bond connects sites 3 and 4. As a warmup, first perform a few sweeps of the standard, serial finite-size DMRG algorithm<sup>1–3</sup> keeping only a small number of states so this non-parallel part takes relatively little time. Stop when the two exposed sites are at the center bond as in Fig. 2(a).

At this point the wavefunction within DMRG has the form

$$|\Psi\rangle = \sum_{\alpha_2 s_3 s_4 \alpha_4} \Psi^{\alpha_2 s_3 s_4 \alpha_4} |\alpha_2\rangle_L |s_3\rangle |s_4\rangle |\alpha_4\rangle_R \quad (1)$$

where  $s_3, s_4 = 1, \dots, d$  label the lattice basis on sites 3 and 4, and  $\alpha_2, \alpha_4 = 1, \dots, m$  label orthonormal many-body states approximating  $|\Psi\rangle$  within the left and right blocks, respectively.

DMRG proceeds in two steps: first, this wavefunction is optimized using a few Lanczos or Davidson steps with the Hamiltonian projected into the  $|\alpha_2\rangle_L |s_3\rangle |s_4\rangle |\alpha_4\rangle_R$

basis.<sup>3</sup> Then a renormalization group procedure is carried out based on the singular value decomposition (SVD) of the amplitudes  $\Psi$

$$\Psi^{(\alpha_2 s_3)(s_4 \alpha_4)} = \sum_{\alpha_3} A_{\alpha_3}^{(\alpha_2 s_3)} \Lambda_{\alpha_3}^{\alpha_3} B_{\alpha_3}^{(s_4 \alpha_4)}. \quad (2)$$

(with  $\Psi$  temporarily treated as an  $(md) \times (dm)$  matrix). Following the SVD, all but the largest  $m$  singular values are truncated.

In the serial DMRG algorithm, one next selects ei-




FIG. 3. Prediction fidelity  $1 - \langle \Psi'_n | \Psi_n \rangle$  at the center (shared) bond of a DMRG calculation parallelized across two nodes.  $|\Psi'_n\rangle$  is the initial merged wavefunction following sweep  $n$  as in Eq. (5).  $|\Psi_n\rangle$  is the wavefunction resulting from a fully converging Davidson calculation initialized with  $|\Psi'_n\rangle$ . The systems are, from top to bottom, the 200 site  $S = 1/2$  Heisenberg chain without warmup sweeps (dashed blue circles);  $S = 1/2$  chain with warmup sweeps (solid blue circles);  $S = 1$  Heisenberg chain without (dashed black squares) and with warmup sweeps (solid black squares). The warmup consisted of 5 serial DMRG sweeps keeping  $m = 50$  states. For the parallel sweeps,  $m$  was increased after every other sweep to a maximum of 600.

ther the transformation matrix  $B$  to grow the right block, in which case the product  $AA\Lambda$  is used to compute the wavefunction amplitudes on the bond to the left, or one chooses the transformation matrix  $A$  to grow the left block and  $\Lambda B$  to form the next wavefunction amplitudes.<sup>29,30</sup> Iterating these steps generates a sweeping procedure which improves the wavefunction at every bond sequentially.

Now assume we have two computers able to work in parallel. Following the SVD, the first machine could sweep left while the second simultaneously sweeps right. However this soon leads to a conceptual issue: each computer will be working with a different wavefunction globally. One can temporarily ignore this problem, but when the computers meet again it will be unclear how to merge their wavefunctions for the purpose of performing a DMRG step on their shared bond.

To overcome this inconsistency, rewrite Eq. (2) but insert the identity  $\Lambda V = 1$ , where  $V \stackrel{\text{def}}{=} \Lambda^{-1}$ :

$$\Psi^{\alpha_2 s_3 s_4 \alpha_4} = \sum_{\alpha_3} A_{\alpha_3}^{\alpha_2 s_3} \Lambda^{\alpha_3} V_{\alpha_3} \Lambda^{\alpha_3} B_{\alpha_3}^{s_4 \alpha_4} \quad (3)$$

$$\stackrel{\text{def}}{=} \sum_{\alpha_3} \psi_3^{\alpha_2 s_3 \alpha_3} V_{\alpha_3} \psi_4^{\alpha_3 s_4 \alpha_4}. \quad (4)$$

Again both machines can sweep in parallel, but when they return to their shared bond—the left with updated amplitudes  $\psi'_3$  and the right with amplitudes  $\psi'_4$ —there is a consistent way to merge the two wavefunctions. To define the merged wavefunction, take the DMRG basis states  $|\alpha_2\rangle_L$  for the left from the machine sweeping over that region; the states  $|\alpha_4\rangle_R$  for the right from the other machine; and for the amplitudes at the center bond choose

$$\Psi' = \psi'_3 V_3 \psi'_4, \quad (5)$$

similarly to Eq. (4), using the original  $V$ . Note that the exact ground state is a fixed point of this procedure. After merging the two wavefunctions for optimization on the shared bond, the merged wavefunction can again be split in two if more parallel sweeps are needed.

Though there is no communication between machines prior to each merge, in practice we find that Eq. (5) provides a good initial state for the Lanczos or Davidson steps on the shared bond, as shown in Fig. 3. However, because each machine otherwise updates the wavefunction independently, DMRG convergence is typically slower at the shared bond, especially near the beginning of a calculation. Although this means parallel DMRG gives slightly worse results compared to serial DMRG for the same number of sweeps, the nearly ideal speedup in calculation time more than compensates for this effect.

Though the discussion above emphasizes the wavefunction, an important part of DMRG is transforming any projected operators (such as the Hamiltonian) while sweeping. First, as each machine sweeps away from the shared bond in parallel, this transformation occurs in




FIG. 4. Sweeping pattern for one full sweep of the parallel DMRG algorithm split over four computational nodes. First, (a) the nodes are positioned in a spatially staggered pattern and sweep to the other end of their block (b) they wait for their neighboring node to arrive then communicate. Finally the nodes sweep back (c) to their starting positions and (d) communicate with their other neighbor.

the usual way. For example, the matrix  $B$  from Eq. (2) transforms the Hamiltonian into the local basis of the next pair of sites to the left. Later when the two machines merge their wavefunctions, they merge operators in an analogous way: operator terms acting in the left half of the system are approximated by their projection into the basis states from the left machine and similarly for the right.

The role of the matrix  $V$  in the merge can be understood by observing that DMRG approximately preserves the reduced density matrix over regions where it does not sweep. Assuming that the wavefunction does not change too drastically after each sweep, the matrix  $V$  normalizes  $\psi'_4$  on the right such that it approximates the reduced density matrix eigenstates on the right half of the system, and similarly for the left. Note the resemblance of Eq. (5) to the prediction step in the infinite DMRG algorithm of Ref. 31, which directly motivated the present work. As explained in Section IV, from a matrix product state point of view Eq. (3) is a transformation from a gauge having only one orthogonality center (a tensor whose indices all label orthonormal states) to a gauge with two orthogonality centers.

One concern about introducing a matrix  $V = \Lambda^{-1}$  at shared bonds is that the relative errors made in computing small singular values may be amplified when inverting them to compute the elements of  $V$ , which can be orders of magnitude larger than one. To address this issue, we have implemented an accurate SVD algorithm described in detail in the Appendix. By recursively performing an SVD on submatrices containing only the smallest singular values, one can obtain uniform relative accuracy for all singular values—not just the largest.

The two-block algorithm described above can be readily extended to  $n$  real-space blocks by first splitting the




FIG. 5. Columnar VBS order parameter  $\langle D_x \rangle$  on open cylinders of the square-lattice  $Q_2$  model Eq. (6) with strong-pinning boundary conditions.<sup>14</sup> Each order parameter value was estimated by increasing the states kept  $m$  after every other sweep and extrapolating in the truncation error. Error bars are smaller than symbol sizes except for the  $L = 10$  system.

system in two, then further splitting each sub-block until there are  $n$  total. This motivates the sweeping pattern shown in Fig. 4. Odd numbered nodes start on the left end of their block and even nodes on the right. This way, when a node reaches the end of its block the next node is ready to communicate. If a node reaches the end of its block before its neighbor arrives, it is better for the node to wait instead of immediately beginning the next half sweep. Having an updated environment far outweighs the loss in efficiency due to a node briefly remaining idle.

### III. BENCHMARK APPLICATIONS

#### A. Pure $Q_2$ Model

To demonstrate that real-space parallel DMRG can be used to accelerate very challenging two-dimensional DMRG calculations, we use it to study the  $S = 1/2$ , pure  $Q_2$  model on the square lattice. This model has been proposed as a benchmark for testing the predictive ability of DMRG for 2d systems.<sup>32</sup> Extensive quantum Monte Carlo (QMC) calculations show the model is in a phase with weak columnar valence bond solid (VBS) order.<sup>32–34</sup> Because DMRG is limited to much smaller finite-size systems than QMC, a reasonable concern is that DMRG could miss such weak order and possibly mistake the phase for a spin liquid.

The Hamiltonian of the pure  $Q_2$  model is defined as

$$H = -Q_2 \sum_{\langle ij;kl \rangle} (\mathbf{S}_i \cdot \mathbf{S}_j - 1/4)(\mathbf{S}_k \cdot \mathbf{S}_l - 1/4) \quad (6)$$

where the sum is over all pairs of bonds  $(i, j)$  and  $(k, l)$

on opposite sides of the elementary square plaquettes. To investigate the columnar VBS order following Ref. 32, we study finite-size systems of width  $L$  in the  $y$  direction and length  $2L$  in the  $x$  direction. For each system size, we measure the order parameter

$$\langle \hat{D}_x \rangle = \frac{1}{L} \sum_y \langle \mathbf{S}_r \cdot \mathbf{S}_{r+\hat{x}} \rangle - \frac{1}{2} \langle \mathbf{S}_{r-\hat{x}} \cdot \mathbf{S}_r \rangle - \frac{1}{2} \langle \mathbf{S}_{r+\hat{x}} \cdot \mathbf{S}_{r+2\hat{x}} \rangle \quad (7)$$

averaged over all  $\mathbf{r} = (L, y)$  at the central column of the system.

In the shorter  $y$  direction we take periodic boundary conditions but use open boundary conditions in the  $x$  direction for technical reasons. In contrast to Ref. 32, however, we include extra Hamiltonian terms at the edges to pin columnar VBS order. Following the strong-pinning prescription of Ref. 14, we imagine fictitious spins just beyond the edge of the system locked into an ideal columnar VBS in the  $x$  direction. Tracing over these fictitious spins induces a term

$$H_{\text{pin},\mathbf{r}} = \frac{Q_2}{4} (\mathbf{S}_r \cdot \mathbf{S}_{r+\hat{y}} - 1/4) \quad (8)$$

on each vertical bond along the edges of the real system. The combination of these terms and a 2:1 aspect ratio helps to control finite-size effects.

We carried out parallel DMRG calculations for cylinders of width  $L = 4, 6, 8$  and 10 and show the resulting  $\langle \hat{D}_x \rangle$  order parameter values in Fig. 5. Each calculation was parallelized over four real-space blocks with each block assigned to a separate 8-core Intel Harpertown 2.66 GHz node. The largest calculation—keeping up to  $m = 3000$  states for the  $L = 10$  system—took 6 days and would have taken 18–21 days without parallelization. All calculations could easily have been parallelized further with access to more nodes.

With our choice of aspect ratio and boundary conditions, we found only a weak dependence of the order parameter on system size for large enough  $L$ . By averaging  $\langle \hat{D}_x \rangle$  for the largest three systems, we estimate a value  $\langle \hat{D}_x \rangle = 0.078(3)$  for the thermodynamic limit, in good agreement<sup>35</sup> with the value 0.078 predicted by quantum Monte Carlo.<sup>32</sup> (p.10)

#### B. Triangular Heisenberg Antiferromagnet

As a second application of parallel DMRG, we measure the  $120^\circ$  Néel order of the antiferromagnetic  $S = 1/2$  Heisenberg model  $H = J \sum_{\langle ij \rangle} \mathbf{S}_i \cdot \mathbf{S}_j$  on the triangular lattice. In contrast to the  $Q_2$  model, this system is beyond the reach of QMC due to the sign problem. It is challenging even for DMRG because of the high coordination number of the lattice.

Reference 4 computed the staggered magnetization of this system using DMRG on cylinders up to width  $L_y = 9$ , but found somewhat unsatisfactory results for the largest systems. In particular, it was unclear whether




FIG. 6. Staggered magnetization  $M_z$  of the antiferromagnetic  $S = 1/2$  Heisenberg model on open cylinders of the triangular lattice with strong-pinning boundary conditions.<sup>4</sup> The aspect ratio of each cylinder is defined as  $\alpha = L_x/L_y$ . Each magnetization value was computed by gradually increasing the number of states kept and extrapolating in the truncation error. Error bars are smaller than symbol sizes except for the  $L_y = 9$  systems.

the order parameter develops a crossing point at some particular aspect ratio.

Here we revisit that calculation using parallel DMRG to study the width  $L_y = 9$  systems more accurately. Our procedure, described in more detail in Refs. 4 and 14, is to study cylinders with  $L_y$  sites and periodic boundary conditions in the  $y$  direction and  $L_x$  sites and open boundaries in the  $x$  direction. We apply a magnetic field term  $-J/2 S_i^z$  to sites  $i$  at the open edges on only one of the three sublattices in order to pin the magnetization direction and reduce the bulk entanglement. We then measure  $M_z = \langle S^z \rangle$  on the pinned sublattice averaged over the center column of the cylinder. First, we reproduce the results of Ref. 4 for cylinders of width  $L_y = 3$  and  $L_y = 6$ . These results are shown in Fig. 6 and agree with those of Ref. 4 essentially exactly.

For the width 9 systems we chose aspect ratios closer to the crossing point than Ref. 4. We kept up to  $m = 5000$  states in DMRG and extrapolated  $M_z$  in the truncation error. We also performed a second set of runs increasing the number of states according to a different pattern to test our extrapolation and found excellent agreement. The largest runs took about a week each using four nodes, and therefore would have taken 3-4 weeks without parallelization.

As shown in Fig. 6, by keeping more states and choosing aspect ratios closer to the crossing point, we obtain tighter bounds on the 2d magnetization. By averaging the two  $L_y = 9$  values in Fig. 6, we find  $\alpha_c \simeq 1.64$  and estimate  $M_z = 0.202(2)$  for the infinite 2d system.

#### IV. INVERSE CANONICAL MATRIX PRODUCT STATE GAUGE

An interesting byproduct of the parallel DMRG wavefunction transformation Eqs. (4) and (5) is that it motivates an alternative gauge for matrix product states (MPS) similar to the canonical gauge<sup>30,36</sup> but with every site an an orthogonality center rather than every bond. By orthogonality center (OC) we mean any tensor whose indices all label orthonormal states.

For example, the only OC of an MPS in the left-canonical gauge—the gauge naturally occurring at the end of a full DMRG sweep—is the first site tensor.<sup>30</sup> This MPS gauge is shown in Fig. 7(a) with the first site tensor shaded to indicate it is an OC. The arrow on each index indicates whether that index transforms as a ket (vector) or bra (covector). An outgoing arrow indicates a ket index and an incoming arrow a bra index. In traditional tensor notation this corresponds to a raised or lowered index, respectively. Because the physical site indices transform as kets by definition, they always point out of a ket MPS.

In our convention, the arrows of the virtual or link indices within an MPS do not flow according to a rigid pattern, such as left to right, but rather out of the OC (or out of each OC if there are more than one). To motivate this convention, perform a gauge transformation of the MPS 7(a). Contract the first two site tensors over their shared link index, then compute an SVD as shown in Fig. 7(b), and multiply the diagonal singular-value matrix into the second site tensor. The resulting MPS Fig. 7(c) has its OC at the second site since the first site tensor is formed from a unitary matrix. As expected, the arrows now flow out of the second site tensor.




FIG. 7. Matrix product state in (a) the left-canonical gauge where the OC is the first site. Combining (b) the first two site tensors, computing a singular-value decomposition, then multiplying the singular value matrix (shaded diamond) into the second site tensor transforms the MPS into (c) a mixed-canonical gauge with the second site as the OC.<sup>30</sup>






FIG. 8. Matrix product state in (a) the left-canonical gauge with the first site an OC. In panel (b) combine the first two site tensors and compute a singular-value decomposition but now with two copies of the singular value matrix  $\Lambda$  (shaded diamond) and its inverse  $V$  (white circle). Multiplying each singular value matrix  $\Lambda$  into the neighboring site tensor transforms the MPS into (c) a new gauge having two OCs.

Within the parallel DMRG algorithm, when computing an SVD at the shared bond between two nodes, one duplicates the singular value matrix  $\Lambda$  by also inserting the matrix  $V \stackrel{\text{def}}{=} \Lambda^{-1}$  on that bond. Multiplying each copy of  $\Lambda$  into its neighboring site tensor creates an MPS gauge with an additional OC. If we repeat the example of the previous paragraph, but now use this modified SVD scheme as shown in Fig. 8(b), the result is a gauge in which the first two sites are OCs and a matrix  $V$  appears on the first bond. Repeating this procedure at every bond results in the gauge shown in Fig. 9. This gauge resembles the canonical gauge but has diagonal matrices containing inverse Schmidt coefficients on each bond. For this reason we refer to it as the inverse canonical gauge.

A key advantage of working in this gauge is that every site tensor is simultaneously an OC. This makes operations such as computing expectation values of local operators very simple since only site tensors on which an operator acts non-trivially need to be included in the computation (all other site tensors cancel by construction since they are external to an OC and therefore represent orthonormal states). In fact, such local expectation values can be computed in parallel in this gauge. By contrast, a mixed-canonical MPS such as Fig. 7(c) must be re-gauged unless the OC is already included in the support of the operator to be measured.

Finally, we note that when conserving abelian quan-




FIG. 9. Inverse canonical matrix product state gauge. Each site tensor  $\psi$  is an OC. This gauge is similar to the canonical gauge ( $\Gamma$ - $\Lambda$  form)<sup>30,36</sup> where each bond tensor  $\Lambda$  is an OC and is a diagonal matrix containing the Schmidt decomposition weights at that bond. Here each bond tensor  $V = \Lambda^{-1}$  is a diagonal matrix containing the inverse Schmidt weights. One can directly map between the two gauges via  $V_j = \Lambda_j^{-1}$  and  $\psi_j = \Lambda_{j-1} \Gamma_j \Lambda_j$ .

tum numbers it is natural to use the same arrow convention described above to denote quantum number flux. In this convention, OCs act as quantum number sources having non-zero flux whereas in a mixed-canonical gauge like that of Fig. 7(c), for example, all other tensors have zero flux since they enact unitary basis transformations.

## V. CONCLUSION

We have presented a straightforward modification of the standard DMRG algorithm which allows it to be parallelized across real-space blocks, providing a nearly ideal speedup. The algorithm differs from serial DMRG only at block boundaries and can readily be combined with other approaches for parallelizing DMRG. This algorithm should be especially useful for DMRG studies of 2d lattice models, quantum chemical systems, and very large or otherwise difficult 1d models.

We have also presented a set of best practices for parallel DMRG simulations, such as the sweeping pattern suggested in Fig. 4, but there remains considerable freedom in implementing the algorithm. For example, in our benchmark applications we typically divided the system into real-space blocks of about 8-20 sites, but in principle the blocks could be as small as two sites, offering maximum parallelization. It is interesting to note that fully converging a parallel DMRG calculation in this limit would automatically produce an MPS in the inverse canonical gauge.

Looking ahead, we expect to see real-space parallelism become a standard tool for accelerating challenging DMRG calculations since it can be implemented in existing codes. We also hope this work encourages authors of DMRG-related papers to identify parallel aspects of their methods even more prominently.

## ACKNOWLEDGMENTS

This paper is dedicated to Ernie Compton, a gifted science teacher and inspiring role model. We thank Thomas Barthel, Bela Bauer, Bryan Clark, Stefan Depenbrock,

Adrian Feiguin, Hong-Chen Jiang, Salvatore Manmana, Ian McCulloch, Ulrich Schollwöck, Guifré Vidal, and Zhenyue Zhu for helpful discussions. This work is supported by NSF grant DMR-1161348.

## APPENDIX: ACCURATE SINGULAR VALUE DECOMPOSITION ALGORITHM

When implementing the parallel DMRG algorithm or working with inverse canonical matrix product states it is essential to compute singular value decompositions (SVD) to high accuracy. This is because of the presence of the matrix  $V$  having the inverse singular values along its diagonal. Typical vendor-provided SVD algorithms (such as SGESVD within LAPACK) may provide poor relative accuracy for the smallest singular values  $\lambda_\alpha$ , which then translates into very large errors upon computing  $\lambda_\alpha^{-1} \gg 1$ .

To overcome this problem while maintaining efficiency, we have implemented the following SVD algorithm. Though it calls itself recursively, its asymptotic cost remains  $\sim mn^2$  for an  $n \times m$  rectangular matrix  $M$  (assuming  $n < m$  without loss of generality) since each recursive call only involves a smaller submatrix.

The algorithm proceeds as follows:

1. Compute the SVD of  $M = A\Lambda B$  using a standard algorithm such as SGESVD or through the eigenvalue decomposition of  $\rho \stackrel{\text{def}}{=} M M^\dagger = A\Lambda^2 A^\dagger$  (then computing  $B$  by orthogonalizing the columns of  $A^\dagger M$ ).
2. Denote the diagonal elements of  $\Lambda$  (the singular values) as  $\{\lambda_\alpha | \alpha = 1 \dots n\}$ . For some predetermined threshold  $\epsilon > 0$ , find the smallest integer  $p$  such that  $\lambda_p/\lambda_1 < \epsilon$ . We have found  $\epsilon = 10^{-4}$  to be a good choice.
3. If no such  $p$  exists, the algorithm has converged. Return the matrices  $A$ ,  $\Lambda$ , and  $B$  from step 1.
4. If the algorithm has not converged, compute  $X = A^\dagger M B^\dagger$ , but only the last  $n - p$  rows and columns such that  $X$  is an  $(n - p) \times (n - p)$  matrix. If the SVD of step 1 could be computed exactly,  $X$  would be diagonal and contain the last  $(n - p)$  singular values. In practice,  $X$  will only be approximately diagonal due to numerical errors.
5. Recursively repeat the algorithm starting again at step 1 but with  $M$  replaced by  $X$ . Denote the resulting SVD matrices  $\tilde{A}$ ,  $\tilde{\Lambda}$ , and  $\tilde{B}$ .

6. Update  $A$ ,  $B$ , and  $\Lambda$  as follows:

$$A_{i\alpha} = \sum_{k=p}^n A_{ik} \tilde{A}_{k\alpha} \quad \alpha = p \dots n \quad (9)$$

$$B_{\alpha j} = \sum_{k=p}^n \tilde{B}_{\alpha k} B_{kj} \quad \alpha = p \dots n \quad (10)$$

$$\lambda^\alpha = \tilde{\lambda}^\alpha \quad \alpha = p \dots n \quad (11)$$

where  $\lambda^\alpha$  are the diagonal elements of  $\Lambda$ .

7. Return the updated SVD matrices  $A$ ,  $\Lambda$ , and  $B$ .

Because the SVD method used in step 1 of the algorithm is typically accurate for all but the smallest singular values (as defined by the threshold  $\epsilon$ ), by calling the method recursively on a submatrix containing only these smallest singular values the algorithm finds all the singular values accurately.



- 
- <sup>1</sup> S. R. White, Phys. Rev. Lett., **69**, 2863 (1992); S. White, Phys. Rev. B, **48**, 10345 (1993).
  - <sup>2</sup> U. Schollwöck, Rev. Mod. Phys., **77**, 259 (2005).
  - <sup>3</sup> R. M. Noack and S. R. Manmana, AIP Conference Proceedings, **789**, 93 (2005).
  - <sup>4</sup> S. R. White and A. Chernyshev, Phys. Rev. Lett., **99**, 127004 (2007).
  - <sup>5</sup> H.-C. Jiang, Z.-C. Gu, X.-L. Qi, and S. Trebst, Phys. Rev. B, **83**, 245104 (2011).
  - <sup>6</sup> S. Yan, D. A. Huse, and S. R. White, Science, **332**, 1173 (2011).
  - <sup>7</sup> P. Corboz, S. R. White, G. Vidal, and M. Troyer, Phys. Rev. B, **84**, 041108 (2011).
  - <sup>8</sup> H.-C. Jiang, H. Yao, and L. Balents, Phys. Rev. B, **86**, 024424 (2012).
  - <sup>9</sup> B. Bauer, P. Corboz, A. M. Läuchli, L. Messio, K. Penc, M. Troyer, and F. Mila, Phys. Rev. B, **85**, 125116 (2012).
  - <sup>10</sup> S. Depenbrock, I. P. McCulloch, and U. Schollwöck, Phys. Rev. Lett., **109**, 067201 (2012).
  - <sup>11</sup> L. Cincio and G. Vidal, Phys. Rev. Lett., **110**, 067208 (2013).
  - <sup>12</sup> Z. Zhu, D. A. Huse, and S. R. White, Phys. Rev. Lett., **110**, 127205 (2013).
  - <sup>13</sup> R. Ganesh, J. van den Brink, and S. Nishimoto, Phys. Rev. Lett., **110**, 127203 (2013).
  - <sup>14</sup> E.M. Stoudenmire and S. R. White, Annual Review of Condensed Matter Physics, **3**, 111 (2012).
  - <sup>15</sup> Y. Kurashige and T. Yanai, J. Chem. Phys., **135**, 094104 (2011).
  - <sup>16</sup> G. K.-L. Chan and S. Sharma, Annual Review of Physical Chemistry, **62**, 465 (2011).
  - <sup>17</sup> S. Sharma and G. K.-L. Chan, J. Chem. Phys., **136**, 124121 (2012).
  - <sup>18</sup> J. Hachmann, W. Cardoen, and G. K.-L. Chan, The Journal of Chemical Physics, **125**, 144101 (2006).
  - <sup>19</sup> M. Dolfi, B. Bauer, M. Troyer, and Z. Ristivojevic, Phys. Rev. Lett., **109**, 020604 (2012).
  - <sup>20</sup> E. M. Stoudenmire, L. O. Wagner, S. R. White, and K. Burke, Phys. Rev. Lett., **109**, 056402 (2012).
  - <sup>21</sup> G. K.-L. Chan, J. Chem. Phys., **120**, 3172 (2004).
  - <sup>22</sup> G. Hager, E. Jeckelmann, H. Fehske, and G. Wellein, Journal of Computational Physics, **194**, 795 (2004).
  - <sup>23</sup> Y. Kurashige and T. Yanai, The Journal of Chemical Physics, **130**, 234114 (2009).
  - <sup>24</sup> A. Daley, C. Kollath, U. Schollwöck, and G. Vidal, J. Stat. Mech., **04**, 005 (2004).
  - <sup>25</sup> S. R. White and A. E. Feiguin, Phys. Rev. Lett., **93**, 076401 (2004).
  - <sup>26</sup> Z. Y. Xie, H. C. Jiang, Q. N. Chen, Z. Y. Weng, and T. Xiang, Phys. Rev. Lett., **103**, 160601 (2009).
  - <sup>27</sup> I. Pizorn, L. Wang, and F. Verstraete, Phys. Rev. A, **83**, 052321 (2011).
  - <sup>28</sup> P. Czarnik, L. Cincio, and J. Dziarmaga, Phys. Rev. B, **86**, 245101 (2012).
  - <sup>29</sup> S. R. White, Phys. Rev. Lett., **77**, 3633 (1996).
  - <sup>30</sup> U. Schollwöck, Annals of Physics, **326**, 96 (2011).
  - <sup>31</sup> I. P. McCulloch, arxiv:0804.2509 (2008).
  - <sup>32</sup> A. W. Sandvik, Phys. Rev. B, **85**, 134407 (2012).
  - <sup>33</sup> A. W. Sandvik, Phys. Rev. Lett., **98**, 227202 (2007).
  - <sup>34</sup> A. W. Sandvik, Phys. Rev. Lett., **104**, 177201 (2010).
  - <sup>35</sup> Our calculations only agree up to a factor of precisely two. Up to this factor, we have successfully reproduced the results (not shown) of Ref. 32 when omitting the edge pinning terms. We therefore report twice the value stated in Ref. 32 (p. 10) as the QMC estimate of the 2d order parameter.
  - <sup>36</sup> G. Vidal, Phys. Rev. Lett., **91**, 147902 (2003).