



CHORUS

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

Lanczos-boosted numerical linked-cluster expansion for quantum lattice models

Krishnakumar Bhattaram and Ehsan Khatami

Phys. Rev. E **100**, 013305 — Published 11 July 2019

DOI: [10.1103/PhysRevE.100.013305](https://doi.org/10.1103/PhysRevE.100.013305)

Lanczos Boosted Numerical Linked-Cluster Expansion for Quantum Lattice Models

Krishnakumar Bhattaram^{1,2} and Ehsan Khatami²

¹*Lynbrook High School, 1280 Johnson Ave, San Jose, CA 95129, USA*

²*Department of Physics and Astronomy, San Jose State University, San Jose, CA 95192, USA*

Numerical linked-cluster expansions allow one to calculate finite-temperature properties of quantum lattice models directly in the thermodynamic limit through exact solutions of small clusters. However, full diagonalization is often the limiting factor for these calculations. Here, we show that a partial diagonalization of the largest clusters in the expansion using the Lanczos algorithm can be as useful as full diagonalization for the method while mitigating some of the time and memory issues. As test cases, we consider the frustrated Heisenberg model on the checkerboard lattice and the Fermi-Hubbard model on the square lattice. We find that our approach can surpass state of the art in performance in a parallel environment.

I. INTRODUCTION

Over the past decade, the numerical linked-cluster expansion (NLCE) [1–3] has become a versatile and powerful technique for solving quantum lattice models and studying finite-temperature properties of strongly-correlated systems in the thermodynamic limit. So far, it has seen applications to a variety of problems including magnetic models [2, 4–6], itinerant electron models on several different geometries [3, 7–9], quantum quenches in the thermodynamic limit [10, 11], entanglement [12–14], driven-dissipative open systems [15], disordered systems [11, 16], and even the hard problem of nonequilibrium dynamics [17, 18] or dynamical properties for systems in equilibrium in the thermodynamic limit [19].

At the heart of the method lies the exact solution of model Hamiltonians on relatively small clusters, which can have unusual topologies, using exact diagonalization. This step is the most computationally expensive and therefore the limiting factor of the method. For models such as the Fermi-Hubbard model, the growth in the size of the Hilbert space with cluster size puts the calculations quickly up against an exponential wall. In other cases, the factorially growing number of clusters that need to be diagonalized is the limiting factor.

In cases where only the ground state is of interest, and in the absence of long-range entanglement, which is unfavorable to the convergence of NLCEs, other techniques such as the Lanczos algorithm or the density matrix renormalization group have been employed [6, 14].

The main idea behind our study is to employ the Lanczos technique to obtain not only the ground state, but also a part of the low-energy spectrum of the model in order to access nonzero, but low, temperatures for large clusters in the NLCE. The goal is to combine the partial-diagonalization for solving clusters in the last order of the series with the traditional full diagonalization of smaller cluster at lower orders to push the convergence beyond temperatures that have been accessible before. The larger the clusters at higher orders, the lower the temperatures they contribute to in the NLCE [20], and the smaller the percentage of the lowest-lying

eigenenergies that is needed.

We benchmark our results from this technique against those obtained from full diagonalization in the NLCE for the antiferromagnetic Heisenberg model on the checkerboard lattice with frustration and the Fermi-Hubbard model on the square lattice. We find that in a massively parallel environment, the Lanczos-based method even with the very expensive full re-orthogonalization step can offer a significant speedup over the traditional diagonalization method.

The exposition is as follows. In Sec. II, we introduce the models. In Sec. III, we provide the basics of the NLCE and discuss how it can benefit from the Lanczos algorithm. We discuss the results for the two models, scaling and some details of the parallelization scheme in Sec. IV

II. MODELS

A. Checkerboard Lattice Heisenberg Model

As the first case study, we consider the spin-1/2 quantum Heisenberg model on the checkerboard lattice whose Hamiltonian is written as

$$\hat{H} = \sum_{ij} J_{ij} \hat{\mathbf{S}}_i \cdot \hat{\mathbf{S}}_j, \quad (1)$$

where $\hat{\mathbf{S}}_i$ is the spin-1/2 vector at site i , and $J_{ij} = J$ for nearest neighbor i and j and $J_{ij} = J'$ for next nearest neighbor i and j on every other plaquette in a checkerboard pattern as shown in Fig. 1. $J' = 0$ represents the square lattice Heisenberg model. Here, we consider the test case $J = 0.5$ and $J' = 1.0$. So, J' sets the unit of the energy when this model is discussed in the paper.

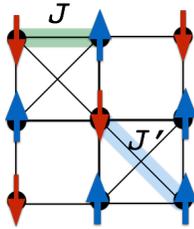


FIG. 1. A 3×3 section of the checkerboard lattice Heisenberg model. J is the strength of the exchange interaction on nearest neighbor bonds and J' is the interaction on next nearest neighbor bonds in every other 2×2 plaquette.

B. Square Lattice Fermi-Hubbard Model

As a second case study, we consider the Fermi-Hubbard model on the square lattice, expressed as

$$H = -t \sum_{\langle ij \rangle \sigma} c_{i\sigma}^\dagger c_{j\sigma} + U \sum_i n_{i\uparrow} n_{i\downarrow}, \quad (2)$$

where $c_{i\sigma}$ ($c_{i\sigma}^\dagger$) annihilates (creates) a fermion with spin σ on site i , $n_{i\sigma} = c_{i\sigma}^\dagger c_{i\sigma}$ is the number operator, t is the hopping integral and the unit of energy for this model, U is the onsite Coulomb interaction, which we take to be $8t$, and $\langle \dots \rangle$ denotes nearest neighbors. We choose the chemical potential to be $4t$, corresponding to an average density of one fermion per site (the half-filled limit).

III. NUMERICAL LINKED-CLUSTER EXPANSION

In the numerical linked-cluster expansion a given extensive property $P(\mathcal{L})$ of the lattice model is expressed as a sum over the contributions to that property from every cluster that can be embedded in the lattice \mathcal{L} . While a version of the method can be employed for finite system sizes, its main advantage has been in its applications to the models in the thermodynamic limit, $\mathcal{L} \rightarrow \infty$. In that limit, the series expansion is given as:

$$P(\mathcal{L})/\mathcal{L} = \sum_c L(c) W_P(c) \quad (3)$$

where c are topologically-distinct clusters that can be embedded in \mathcal{L} , $L(c)$ is the number of ways per site cluster c can be embedded in the lattice, and $W_P(c)$ is the contribution to property P computed via the inclusion-exclusion principle:

$$W_P(c) = P(c) - \sum_{s \subset c} W_P(s), \quad (4)$$

where s is a cluster that can be embedded in c (a sub-cluster of c), and $P(c)$ is the property calculated for cluster c using exact diagonalization (ED). $L(c)$ can

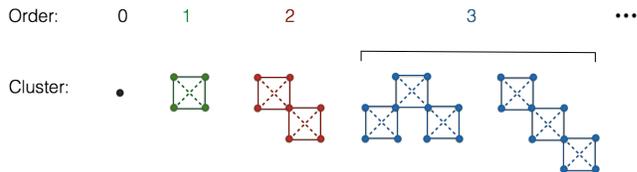


FIG. 2. Square expansion for the checkerboard lattice Heisenberg model, where clusters at higher orders are generated by adding corner-sharing 2×2 blocks with next nearest neighbor bonds to smaller clusters in lower orders. Shown are the topologically distinct clusters up to the third order of the expansion.

be thought of as the number of point group symmetry operations for the underlying lattice that give the cluster a distinct orientation. More information on how $L(c)$ are calculated and other details of the algorithm can be found in Ref. [20].

The calculation of properties $P(c)$ using ED is the most computationally expensive part of the NLCEs. For example, in a site expansion, for which order l of the series includes all clusters up to l sites, calculations for the square or honeycomb lattice Fermi-Hubbard models have so far been limited to $l = 9$ [7, 8, 21, 22]. Even though there are only 112 topologically-distinct clusters to diagonalize for the square lattice in the 9th order in that case, the calculations are limited by memory and time requirements for diagonalization of the largest matrices for clusters with a size larger than 9 sites. For quantum magnetic models in which the size of the Hilbert space ($n_{Hilbert}$) grows significantly slower with the system size than for Fermi-Hubbard models, the calculations have been limited to about 15 orders on the square lattice [20, 23], not by memory requirements, but by the large number of clusters that need to be solved in higher orders.

In those cases, one can put the limitation back on the ED and push the convergence to lower temperatures by employing expansion schemes that access much larger clusters more quickly as the order increases. That is accomplished through increasing the size of the building block used to generate clusters. For example, an expansion with squares as building blocks was used to study the checkerboard and square lattice Heisenberg models with clusters up to 19 sites in only 6 orders [5]. Figure 2 shows the first five topologically-distinct clusters in the square expansion. In previous studies, triangular building blocks were used to study magnetic models on the Kagome lattice [6, 24], and rectangular clusters were considered in the study of entanglement at a two-dimensional critical point [12].

Here, we aim to address, to the extent possible, the limitation on time and memory for full diagonalization, which remains the main issue in the NLCE. We combine the site and square expansion NLCE with an efficient partial diagonalization offered by the Lanczos algorithm to access low temperatures much faster than

previously possible. We base our approach on a key observation that only small clusters in low orders of the expansion contribute to properties at the highest temperatures; the inclusion of larger clusters in the series at higher orders does not change the results at high energies/temperatures. Therefore, a partial knowledge of the eigenvalues and corresponding eigenvectors at the low-energy end of the spectrum may be sufficient to deduce new information from larger clusters

A. Efficiency of the Lanczos Algorithm for the NLCE

The Lanczos algorithm [25] has had profound applications in solving discrete quantum systems, notably when the ground state is of interest. The following section will give an overview of the Lanczos method in terms of its applicability to the NLCE. For further information concerning the mechanics and derivation of the Lanczos algorithm, see Ref. [25]

The algorithm in its most basic form is an iterative Krylov subspace method in which an orthogonal basis is built and approximations of eigenvectors are found by projecting our matrix onto the subspace. It is well known for its superior space and time efficiency in solving for the ground state of very large systems. However, the method is known to have numerical instabilities in finite-precision mathematics [26], which manifest themselves in a loss of orthogonality of the Lanczos vectors which define the subspace. This leads to spurious degeneracies of the largest-magnitude eigenvalues, which means that while the ground-state behavior may be fully captured, finite-temperature properties are very difficult to ascertain.

The solutions to maintaining orthogonality are varied [27], however, the most basic in capturing the whole set of eigenvalues to high precision is the application at every iteration of a Gram-Schmidt orthogonalization with all previous vectors, a full reorthogonalization scheme. Figure 3 shows the exact eigenvalues of the Heisenberg Hamiltonian for a 16-site cluster in the spin sector with four spin-ups and twelve spin-downs compared with the values from the Lanczos method with and without full reorthogonalization and two different numbers of iterations (n_{iter}). It is expected that with $n_{iter} = n_{Hilbert}$ the Lanczos algorithm can find most (but not all) eigenvalues with high degree of precision. The severe degeneracy of the lowest eigenvalues in the basic Lanczos method and its resolution after reorthogonalization can be seen in the figure.

The reorthogonalization process requires the storage of about n_{iter} Lanczos vectors, which has similar memory requirements to storing the Hamiltonian matrix when $n_{iter} \sim n_{Hilbert}$, and also greatly increases the time complexity of the process due to a $n_{iter}^2 \times n_{Hilbert}$ scaling. The question thus raised is whether or not the reorthogonalized Lanczos scheme is in fact advantageous

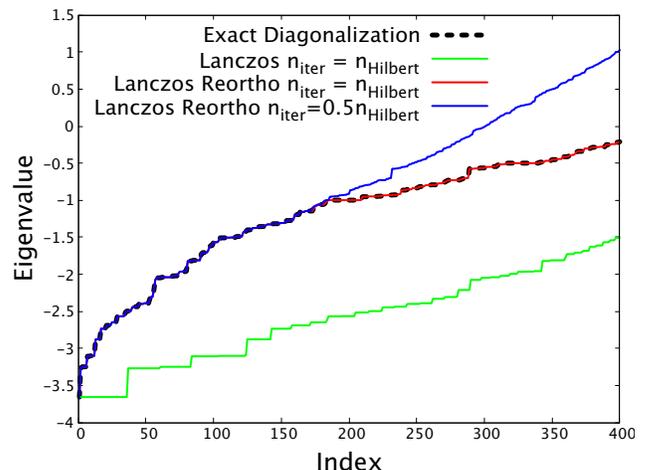


FIG. 3. Sorted first 400 eigenvalues of the Heisenberg Hamiltonian for a 16-site cluster in the four spin-up sector with $n_{Hilbert} = 1820$. Results from the Lanczos algorithm deviate from the exact ones very early on in the spectrum due to the existence of degeneracies and the loss of orthogonality between vectors in the Krylov space. A full re-orthogonalization step in the algorithm restores orthogonality of the Lanczos vector at every iteration and captures a larger percentage of the exact spectrum as the number of iterations increases.

over full diagonalization. Figure 3 includes a plot for the case where the number of iterations is half the size of the Hilbert space. It can be seen that a large fraction of low-energy states have been captured with high accuracy in that case.

A key feature of the NCLE is that with the first few orders, the series is already convergent at high temperatures where correlations in the system are short ranged. Larger clusters in higher orders provide information for properties in the thermodynamic limit only at lower temperatures where correlations grow larger. On the other hand, the Lanczos algorithm can provide very accurate information about the low-temperature behavior of the largest clusters in the series even when the number of iterations is less than the size of the Hilbert space as low-lying eigenvalues can be converged, leading to much smaller memory and time requirements in comparison with full ED. In other words, as system sizes increase by increasing the order, and the region of convergence of the series is pushed to lower temperatures, one can also expect that fewer iterations are necessary. Furthermore, the Lanczos algorithm is readily parallelizable both in the Hamiltonian operation and in the reorthogonalization step in order to further reduce the diagonalization time. We provide a discussion about scaling in Sec. IV A.

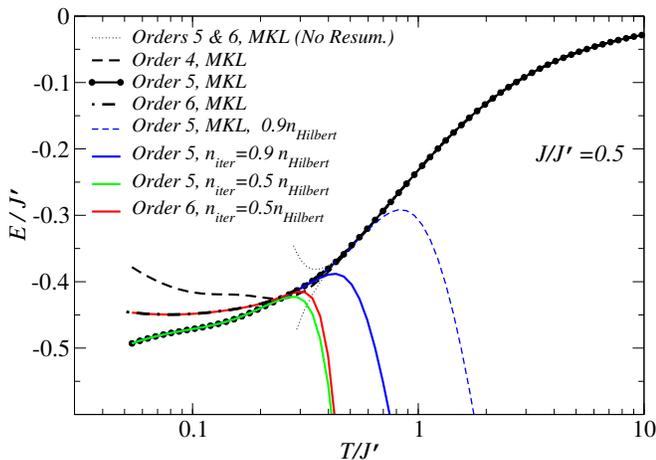


FIG. 4. Energy per site of the Heisenberg model on a checkerboard lattice with $J = 0.5$ and $J' = 1.0$ as a function of temperature. The results within the convergence region are valid in the thermodynamic limit. Thin dotted lines are bare sums in the 5th and 6th orders. The latter is taken from Ref. [5]. Thick black lines are results in the 4th, 5th, and 6th orders from Euler resummations after regular sum of the first three terms in the series. These results are obtained through full diagonalization using a Lapack routine. Color lines show the results from our Lanczos algorithm for the 5th or 6th orders with different number of iterations. As the latter increases, a larger percentage of the low-lying eigenvalues for every cluster in that order approach their exact values and the corresponding curves match the exact curve up to higher temperatures. Blue dashed line represents results from ED for the case where 90% of the eigenvalues are used in the thermal averages.

IV. RESULTS

To benchmark the performance of the Lanczos-boosted NCLE, we start with the energy per site of our Heisenberg model with $J = 0.5$ and $J' = 1.0$ calculated up to the sixth order of the square expansion NLCE (see Fig. 4). We use Intel’s math kernel library (MKL) for full ED as well as our Lanczos method with full reorthogonalization. Full ED results for the sixth order are taken from Ref. [5] where ScaLapack, a distributed-memory version of the linear algebra package (Lapack), was used for diagonalization to meet the enormous memory and time demands. We have used the Lanczos diagonalization for clusters only in the last order of the expansion while other smaller clusters were solved using MKL. We have also used MKL to diagonalize matrices smaller than 2,000 in linear size in the last order to avoid issues in Lanczos that may arise with small matrices. We have used the Euler method for numerical resummation of the series [2, 20]. Results for the fifth and sixth orders before the resummations are also shown as thin dotted lines. Results from another resummation method agree with the outcome of the Euler resummation in its region of convergence.

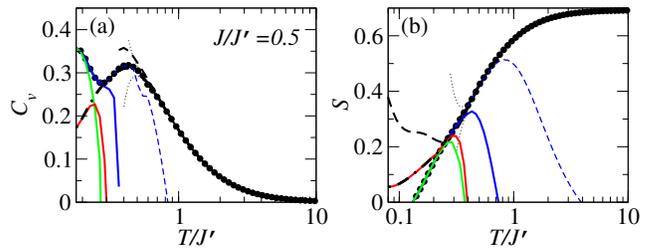


FIG. 5. (a) The heat capacity, and (b) entropy of the Heisenberg model on a checkerboard lattice with $J = 0.5$ and $J' = 1.0$ as a function of temperature. Lines are the same as Fig. 4. In case of the heat capacity, partial diagonalization using the Lanczos algorithm barely reaches the minimum convergence temperature of the series with 6 orders. However, a n_{iter} of 50% of the size of the Hilbert space is enough to reach the convergence temperature for the entropy with 5 or 6 orders.

We find that Lanczos with $n_{iter} = 0.5n_{Hilbert}$ is sufficient to reach the lowest convergence temperature of the series with five or six orders. They are shown as green and red curves in Fig. 4. One can see that they agree with results from MKL up to temperatures around $0.3J'$, slightly above the temperature at which the convergence of the series with six orders is lost. The exact energies in the thermodynamic limit beyond this point are already accessible to lower orders. We also show that by increasing n_{iter} one can systematically reach higher temperatures. With $n_{iter} = 0.9n_{Hilbert}$ the agreement with exact results extends to $T \sim 0.4J'$.

As noted above, n_{iter} does not represent the number of exact eigenvalues obtained in the Lanczos algorithm. We demonstrate that by using the smallest 90% of exact eigenvalues from MKL in the Boltzmann average of the energy in the fifth order, which results in the blue dashed curve in Fig. 4, surpassing the performance of the Lanczos algorithm with $n_{iter} = 0.9n_{Hilbert}$.

As the order increases and the convergent region is extended to lower temperatures, the number of iterations can be increasingly smaller fractions of $n_{Hilbert}$, leading to the superiority in both space and time complexity of the Lanczos-assisted diagonalization over MKL. This is due to the ability of the Lanczos method to capture the lowest eigenvalues with the least number of iterations. The calculation of properties in the sixth order with clusters up to 19 sites was only possible previously through the use of Scalapack routines as done in Ref. [5]. Here, we have obtained results for the sixth order with $n_{iter} = 0.5n_{Hilbert}$ in significantly less time. As one can see in Fig. 4, those results already capture the average energy before the series loses convergence around $T = 0.3J'$.

The efficiency of our approach in using the Lanczos algorithm for treating clusters in the last order of the NLCE depends on the property of interest, and most likely, on the model under investigation. In Fig. 5(a), we show that using half as many iterations in Lanczos as

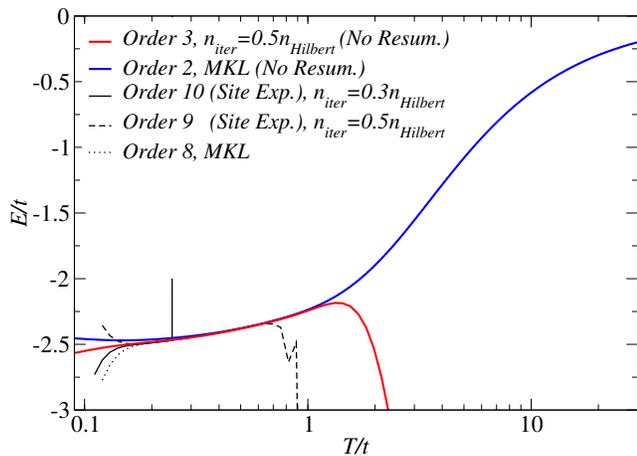


FIG. 6. The average energy per site of the half-filled Fermi-Hubbard model with $U = 8t$ as a function of temperature. Thick color lines are the bare results of the square expansion with 2 and 3 squares. The two clusters in the third order, shown in Fig. 2, have 10 sites each and are diagonalized using the Lanczos algorithm with $n_{iter} = 50\%$ of the size of the Hilbert space. Thin black lines are Euler resummations after regular sum of the first six terms in the site expansion for the same model up to the 10th order. Both orders 9 and 10 are obtained using the Lanczos partial diagonalization algorithm with n_{iter} equal to 50% and 30% of $n_{Hilbert}$, respectively. Results from order 3 of the square expansion and the 10th order of the site expansion agree down to $T \sim 0.15t$.

the size of the Hilbert space is insufficient to capture the heat capacity of the same system studied in Fig. 4 around the lowest convergence temperature of the series with up to six orders. The same is not true for the entropy, shown in Fig. 5(b), where we observe a behavior similar to that of the energy in terms of temperature access with the Lanczos algorithm. We note that properties such as heat capacity, which are higher order derivatives of the partition function than the energy, usually have a worse convergence in comparison to the energy in the NLCE or high-temperature series expansions since nontrivial terms enter the series at higher orders. This is true even with full diagonalization. For example, the base sums for the energy in Fig. 4 converge to $T \sim 0.4J'$, whereas for C_v in Fig. 5(a), they converge to a slightly higher $T \sim 0.6J'$.

To demonstrate that our approach is not limited to quantum magnetic models in which the local Hilbert space size grows as 2^N , where N is the system size, we apply the method to the half-filled Fermi-Hubbard model to obtain the average energy as a function of temperature. As mentioned above, previous NLCE calculations for the model have been limited to 9 orders in the site expansion. Here, we are able to obtain results for the 3rd order of the square expansion, containing two 10-site clusters as shown in Fig. 2. As can be seen in Fig. 6, this is enough to push the convergence temperature below that achieved with 9 orders in the site expansion ($T \sim 0.18t$) even after resummations are used for the latter.

Moreover, using n_{iter} equal to 50% of the size of the Hilbert space (again, for cases where $n_{Hilbert} > 2,000$) in order 3, allows us to include contributions from 10-site clusters to a NLCE for the Hubbard model for the first time. We discuss the performance of our method in Sec. IV A.

We also apply our Lanczos partial diagonalization algorithm to orders 9 and 10 of the site expansion for the same model. Results are shown as thin dashed and solid lines in Fig. 6 and demonstrate that n_{iter} of only 50% and 30%, respectively for orders 9 and 10, are needed for obtaining new information in those orders. Results from order 3 of the square expansion and order 10 of the site expansion agree down to $T \sim 0.15t$. Note that result from order 8, which take only few tens of seconds to compute using 51 cores, already contain the exact information at $T > 0.7t$.

Results away from half filling can be obtained as well by properly adjusting the chemical potential [7]. In that case, the convergence temperature is generally higher than that at half filling, and hence, higher percentage of excited energy levels may be needed.

A. Parallelization and Scaling

We use message passing interface (MPI) parallelization for full ED in our NLCE implementation. We assign a collection of clusters in a given order to a different core to be diagonalized. This works insofar as we do not reach the random access memory (RAM) capacity (128 gigabytes) of a node with 28 cores in the high-performance computer cluster we have used for our calculations. In this picture, the run time is expected to be proportional to the number of clusters assigned to each core for diagonalization, or inversely proportional to the number of cores requested, in the limit of large number of clusters. In Fig. 7, we show the run time as a function of inverse number of cores for MKL in the fifth order of the square expansion for the Heisenberg model. Since there are only 11 clusters in that order, run time generally decreases with increasing the number of cores, until we reach 11 cores. Beyond that, there cannot be any improvement in the run time of MKL.

We also use MPI to parallelize the inner loops of our Lanczos algorithm. Noting that the largest matrices we diagonalize for the Heisenberg model in the fifth order have a linear size of 12870, and therefore, every core can store at least up to $n_{Hilbert}$ double-precision Lanczos vectors in the RAM, we distribute the computational tasks for both the operation of the Hamiltonian on a Lanczos vector and the reorthogonalization step over all cores available. For this reason, ideally the run time would be inversely proportional to the number of cores. However, as shown in Fig. 7, we find that the overhead due to communication rapidly increases beyond one node and the run time eventually saturates. Nevertheless, for $n_{iter} = n_{Hilbert}$ our Lanczos-based

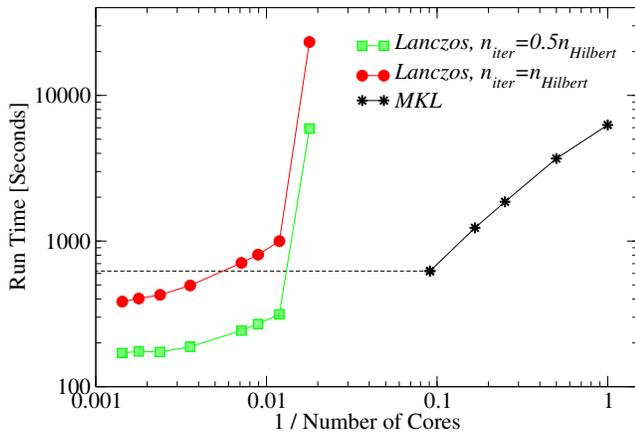


FIG. 7. Scaling of the Lanczos partial diagonalization algorithm with full re-orthogonalization for up to the 5th order. For comparison, we also show the same scaling for ED for which we use the Intel’s MKL to diagonalize every cluster on a core. The performance of MKL improves rapidly by increasing the number of cores until we have as many cores as the number of clusters in the last order of the series. No speedup can be achieved beyond that. The Lanczos algorithm with $n_{iter} = n_{Hilbert}$ quickly surpasses MKL in performance and continues to offer some speedup by increasing the number of cores. Choosing $n_{iter} = 0.5n_{Hilbert}$ reduces the computational time by roughly a factor of three.

method is already faster than our MKL scheme when using 2 nodes and remains faster than MKL as the number of nodes increases. The reorthogonalization step scales like $n_{iter}^2 \times n_{Hilbert}$ and is the most expensive part of our routine. Therefore, decreasing n_{iter} can significantly speed up the calculations. We find that choosing $n_{iter} = 0.5n_{Hilbert}$, cuts the run time by about a factor of three in the case of the fifth order as shown in Fig. 7.

For the Hubbard model, using $n_{iter} = 0.5n_{Hilbert}$ in Lanczos for the third order cuts the computational time by a factor of eight when using two nodes in comparison to MKL. So, as expected, the gain is even higher when the local size of the Hilbert space is larger.

In the site expansion for the Hubbard model, there are 300 clusters to be diagonalized in the tenth order. The largest matrices in the tenth order will take 16 times more RAM (~ 16 gigabytes) to be stored, and about 64 times longer to be diagonalized, than the largest matrices in the ninth order. Therefore, we know that the 43 nodes required to simultaneously diagonalize all the clusters in the tenth order will need about 28 hours.

Using 10 nodes, we were able to complete the partial Lanczos diagonalization with $n_{iter} = 0.3n_{Hilbert}$ for all the clusters in that order sequentially in about 78 hours. So, while in this case, our method slightly reduces the overall computational resources needed, it does not lead to a faster access to the results in the tenth order, if 43 nodes are available. This demonstrates that the efficiency of our scheme depends on how quickly one reaches to cluster sizes that are expensive to solve, and how many they are, as the expansion order increases.

In summary, we have shown that the Lanczos algorithm can be an efficient and flexible technique to partially diagonalize Hamiltonian matrices corresponding to the largest clusters in the NLCE, allowing one to access temperatures not previously accessible using Lapack full diagonalization routines. We benchmark our approach using the square expansion for the checkerboard lattice Heisenberg model with frustration and the half-filled square lattice Fermi-Hubbard model and show that by computing less than 50% of the eigenvalues, through a parallelized Lanczos algorithm in significantly less time than with MKL, we can reach temperatures relevant to the convergence regions of the energy and entropy of the models in the thermodynamic limit.

More efficient implementations of the Lanczos algorithm, such as restart Lanczos or smart partial reorthogonalization, can be employed in the future to further reduce the computational cost. Our Lanczos approach can be generalized to produce eigenvectors in addition to eigenvalues, at an additional cost, for the calculation of other properties of interest. It can also be employed for NLCE studies of other quantum lattice models. Finally, highly precise results from quantum Monte Carlo methods in limited temperature regions can also be used in the last order in a similar fashion to speedup the calculations, at least in cases where the calculations are not hindered by the fermion sign problem.

ACKNOWLEDGMENTS

We acknowledge Marcos Rigol for his early idea of employing the Lanczos algorithm for full diagonalization of clusters in the NLCE. EK acknowledges support from the National Science Foundation (NSF) under Grant No. DMR-1609560. The computations were performed on the Spartan high-performance computing facility at San José State University supported by the NSF under Grant No. OAC-1626645.

[1] M. Rigol, T. Bryant, and R. R. P. Singh, Phys. Rev. Lett. **97**, 187202 (2006).

[2] M. Rigol, T. Bryant, and R. R. P. Singh, Phys. Rev. E **75**, 061118 (2007).

- [3] M. Rigol, T. Bryant, and R. R. P. Singh, Phys. Rev. E **75**, 061119 (2007).
- [4] M. Rigol and R. R. P. Singh, Phys. Rev. Lett. **98**, 207204 (2007).
- [5] E. Khatami and M. Rigol, Phys. Rev. B **83**, 134431 (2011).
- [6] E. Khatami, R. R. P. Singh, and M. Rigol, Phys. Rev. B **84**, 224411 (2011).
- [7] E. Khatami and M. Rigol, Phys. Rev. A **84**, 053611 (2011).
- [8] B. Tang, T. Paiva, E. Khatami, and M. Rigol, Phys. Rev. Lett. **109**, 205301 (2012).
- [9] E. Khatami, Phys. Rev. B **94**, 125114 (2016).
- [10] M. Rigol, Phys. Rev. Lett. **112**, 170601 (2014).
- [11] B. Tang, D. Iyer, and M. Rigol, Phys. Rev. B **91**, 161109(R) (2015).
- [12] A. B. Kallin, K. Hyatt, R. R. P. Singh, and R. G. Melko, Phys. Rev. Lett. **110**, 135702 (2013).
- [13] A. B. Kallin, E. M. Stoudenmire, P. Fendley, R. R. P. Singh, and R. G. Melko, Journal of Statistical Mechanics: Theory and Experiment **2014**, P06009 (2014).
- [14] E. M. Stoudenmire, P. Gustainis, R. Johal, S. Wessel, and R. G. Melko, Phys. Rev. B **90**, 235106 (2014).
- [15] A. Biella, J. Jin, O. Viyuela, C. Ciuti, R. Fazio, and D. Rossini, Phys. Rev. B **97**, 035103 (2018).
- [16] T. Devakul and R. R. P. Singh, Phys. Rev. Lett. **115**, 187201 (2015).
- [17] K. Mallayya and M. Rigol, Phys. Rev. Lett. **120**, 070603 (2018).
- [18] I. G. White, B. Sundar, and K. R. A. Hazzard, (2017), arXiv:1710.07696.
- [19] M. A. Nichols, L. W. Cheuk, M. Okan, T. R. Hartke, E. Mendez, T. Senthil, E. Khatami, H. Zhang, and M. W. Zwierlein, Science (2018), 10.1126/science.aat4387, NoStop
- [20] B. Tang, E. Khatami, and M. Rigol, Computer Physics Communications **184**, 557 (2013).
- [21] E. Khatami and M. Rigol, Phys. Rev. A **86**, 023633 (2012).
- [22] B. Tang, T. Paiva, E. Khatami, and M. Rigol, Phys. Rev. B **88**, 125127 (2013).
- [23] B. Tang, D. Iyer, and M. Rigol, Phys. Rev. B **91**, 174413 (2015).
- [24] E. Khatami, J. S. Helton, and M. Rigol, Phys. Rev. B **85**, 064401 (2012).
- [25] C. Lanczos, J. Res. Nat. BUT. Standards **45255**, 282 (1950).
- [26] C. Paige, I. Inst. Math. Appl. **10**, 373 (1972).
- [27] H. D. Simon, Linear Algebra and its Applications **61**, 101 (1984).