Exact Numerical Treatment of Finite Quantum Systems using Leading-Edge Supercomputers

G. Hager[1], E. Jeckelmann[2], H. Fehske[3], and G. Wellein[1]

[1] Regionales Rechenzentrum Erlangen (RRZE), Martensstraße 1, D-91058 Erlangen, Germany
[2] Johannes-Gutenberg-Universität Mainz, Institut für Physik, KOMET 337, Staudingerweg 7/9, D-55099 Mainz, Germany
[3] Ernst-Moritz-Arndt-Universität Greifswald, Institut für Physik, Domstr. 10a, D-17489 Greifswald, Germany

Summary. Using exact diagonalization and density matrix renormalization group techniques a finite-size scaling study in the context of the Peierls-insulator Mott-insulator transition is presented. Program implementation on modern supercomputers and performance aspects are discussed.

1 Introduction

In the last few years solid-state physics has benefited a lot from scientific computing. The use of numerical techniques is likely to keep on growing quickly in almost all fields of physics as well. Because of the high complexity of many physical problems this will become possible only by the use of modern high-performance supercomputers that are powerful enough to simulate complex systems. Not only do the computational results provide us with important clues on the behaviour of specific materials but they are also widely used as a touchstone to test theoretical approaches, especially in the very difficult regime of strong correlations where the different energy scales in the problem are not well separated. In highly correlated systems, the interaction between the constituents is so strong that they can no longer be considered separately. In other words, the whole is greater than its parts. As a result, the collective behaviour of the microscopic particles, e.g., the electrons in a solid, may scale up to a macroscopic ensemble, exhibiting new and fascinating properties such as high-temperature superconductivity or colossal magnetoresistance [1].

Quasi-onedimensional strongly coupled electron-phonon systems like MX-chain compounds are further examples of electronic systems that are very different from traditional ones [2]. They are particularly rewarding to study for a number of reasons. First they exhibit a remarkably wide range of strengths of competing forces, which gives rise to a rich variety of symmetry-broken ground states. Second these systems share fundamental features with higher-dimensional novel materials, such as high-T_c cuprates or charge-ordered nickelates, i. e. they are complex enough to investigate the interplay of charge, spin, and lattice degrees of freedom which is

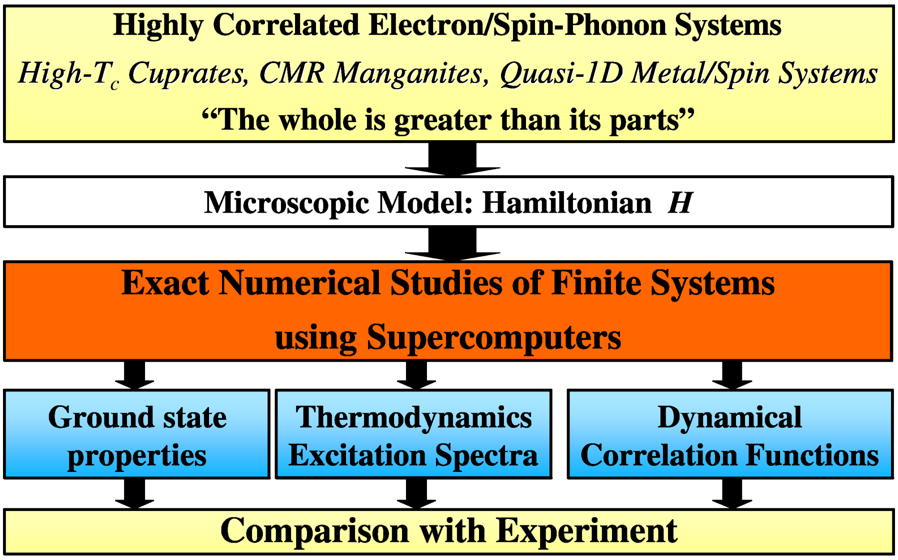

Fig. 1. *Numerical approach to strongly correlated systems.*

important for strongly correlated electronic systems in two and three dimensions as well. Nevertheless they are simple enough to allow for a nearly microscopic modeling.

In this context, a frequently used starting point has been the one dimensional Holstein-Hubbard Hamiltonian

$$H_{\mathrm{HHM}} = -t \sum_{\langle i,j \rangle, \sigma} c_{i\sigma}^{\dagger} c_{j\sigma} + U \sum_{i} n_{i\uparrow} n_{i\downarrow} + g\omega_0 \sum_{i,\sigma} (b_i^{\dagger} + b_i) n_{i\sigma} + \omega_0 \sum_{i} b_i^{\dagger} b_i \,. \quad (1)$$

Here the first two terms constitute the conventional Hubbard Hamiltonian with hopping amplitude t and on-site Coulomb repulsion strength U; $c_{i\sigma}^{\dagger}$ creates a spin-σ electron at Wannier site i and $n_{i\sigma} = c_{i\sigma}^{\dagger} c_{i\sigma}$. Since dynamical phonon effects are known to be particularly important in quasi-1D materials, the third and fourth terms take into account the electron-phonon coupling and the elastic energy of a harmonic lattice, respectively, where $g = \sqrt{\varepsilon_p/\omega_0}$ is a dimensionless electron-phonon coupling constant (ε_p gives the polaron binding energy), and ω_0 denotes the frequency of an optical phonon mode. $b_i^{(\dagger)}$ are the usual phonon annihilation (creation) operators. The Hubbard model [3], originally designed to describe the ferromagnetism of transition metals, was subsequently investigated in the context of metal–insulator (Mott) transition, heavy fermions and high-temperature superconductivity as the probably most simple model to account for strong correlation effects. In the single-electron case, the resulting Holstein model [4] has been studied extensively as a paradigmatic model for polaron formation [5]. At half-filling the electron-phonon coupling may lead to a Peierls instability related to the appearance of charge-density-wave (CDW) order (in competition with the spin-density-wave (SDW) instability triggered by U) [6].

As indicated above, in an attempt to close the gap between a microscopic model and its actual ground state, spectral and thermodynamical properties, theorists have turned to the use of large-scale computers (cf. Fig. 1). Nowadays finite-cluster exact diagonalizations (ED), density matrix renormalization group (DMRG) calculations and quantum Monte Carlo simulations have become very powerful and important tools for solving many-body problems with high accuracy. In what follows, we discuss the basic principles, the advantages and weaknesses of the former two techniques by analyzing the 1D half-filled Holstein-Hubbard model exemplarily.

2 Numerical Methods: ED and DMRG

2.1 Exact Diagonalization

In principle, ED is presently the only numerical method which allows an approximation-free study of the Holstein-Hubbard Hamiltonian in the whole parameter range. A full orthonormal basis set $\{|\Phi_i\rangle\}$ is used to build up the matrix representation of the Holstein-Hubbard Hamiltonian:

$$H_{i,j} = \langle\Phi_i|H_{\mathrm{HHM}}|\Phi_j\rangle\,. \tag{2}$$

Due to the locality of interactions, the matrices $H_{i,j}$ are extremely sparse in a real space basis and standard algorithms such as Lanczos or Davidson can be used to compute exact ground states or low lying excited states while good estimations for excitation spectra are provided by a kernel polynomial method in combination with a maximum entropy method (see [7]). The CPU-time and memory requirements of the three methods are determined by a sparse matrix-vector multiplication (MVM) involving the sparse matrix representation of H_{HHM}. However, the number of degrees of freedom to be considered for a full basis set grows exponentially with increasing cluster size. Even for rather small clusters, e.g. with $N = 8$ to 16 lattice sites, the memory requirements can already exceed the resources of present-day supercomputers. The dimension of the matrix to be diagonalized can be somewhat reduced by exploiting conservation laws. For instance, the conservation of electron number ($\sum_i n_{i\uparrow} + n_{i\downarrow}$) and the z component of the total spin ($\frac{1}{2}\sum_i n_{i\uparrow} - n_{i\downarrow}$) is usually easy to use.

At that point a peculiarity of the Holstein-Hubbard model, where the full basis set can be constructed as a direct product of electronic and phononic basis sets

$$\{|\Phi_{u,v}\rangle = |u\rangle_{\mathrm{el}} \otimes |v\rangle_{\mathrm{ph}};\ u = 1,\ldots,D_{\mathrm{el}};\ v = 1,\ldots,D_{\mathrm{ph}}\}, \tag{3}$$

becomes visible: Because of the phonon degrees of freedom, the total matrix dimension ($D_{\mathrm{tot}} = D_{\mathrm{el}} \times D_{\mathrm{ph}}$) is infinite in principle. However, a well controlled truncation procedure for the phononic basis states

$$|v\rangle_{\mathrm{ph}} = \prod_{i=1}^{N} \frac{1}{\sqrt{m_{i,v}!}} \left(b_i^{\dagger}\right)^{m_{i,v}} |0\rangle_{\mathrm{ph}},\ m_{i,v} \in [0,\infty] \tag{4}$$

can be applied by introducing an upper limit for the number of phonons M contained in each basis state, i. e. only states with $\sum_i m_{i,v} \leq M$ are incorporated in the basis set, which thus gets a dimension of $D_{\mathrm{ph}}^{M} = (M+N)!/M!N!$. Truncating the phonons in such a way is equivalent to setting an upper limit $M\omega_0$ for the elastic energy of the harmonic lattice (fourth term in Eq. (1)). Of course, only the lower part of the spectrum of H_{HHM} with $E \ll M\omega_0$ is described well in this approximation and convergence with respect to M has to be checked carefully for each parameter set $\{t, U, g, \omega_0\}$.

Although a lot of work has been done to reduce the number of basis states in the ED algorithms, even for clusters with only $N = 8$ sites matrix dimensions beyond $D_{\mathrm{tot}}^M = D_{\mathrm{el}} \times D_{\mathrm{ph}}^M \sim 10^{10}$ may be required to achieve sufficient convergence with respect to the truncation of the phonon basis states. Thus the use of powerful supercomputers is indispensable for ED studies of the Holstein-Hubbard model. Our parallel ED package is optimized with respect to memory consumption and does not even store the non-zero elements of the sparse matrix but recomputes them in each MVM step. The parallelization of the MVM step is done by exploiting the natural parallelism of the direct product formulation of the basis set (Eq. 3) (for a detailed discussion cf. [8, 9]).

2.2 Density Matrix Renormalization Group

The DMRG algorithm [10, 11, 12] tries to overcome the limitations of the ED approach by implementing a variational scheme that truncates the Hilbert space used to represent H in an "optimal" way. It is the selection of the basis states that lays the groundwork on which DMRG is built.

DMRG splits the physical system (usually in real space, although a momentum space approach is possible) into two pieces, the so-called *system block* and the *environment block*. Both together form the *superblock* (see Fig. 2).

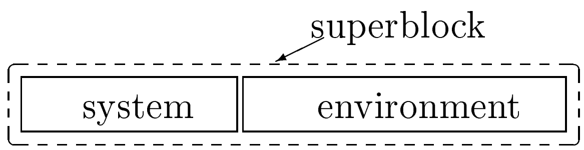

Fig. 2. *Division of the complete physical system into "system block" and "environment block". Both blocks together form the "superblock" whose Hamiltonian matrix is diagonalized.*

The central entity in the algorithm is the *reduced density matrix*

$$\rho_{ii'} = \sum_j \psi_{ij}^* \psi_{i'j} \quad , \tag{5}$$

where i and j label the states of the system and environment blocks, respectively, so that a superblock state $\left|\psi\right\rangle$ can be composed:

$$\left|\psi\right\rangle = \sum_{ij} \psi_{ij} \left|i\right\rangle \left|j\right\rangle \quad . \tag{6}$$

It can now be shown [12] that the eigenstates of ρ with the largest eigenvalues are those that have the most significant impact on observables, i.e. in order to get a good guess at an optimal basis set for the superblock Hamiltonian one has to

- diagonalize the reduced density matrix for a system block of size l and extract the m eigenvectors with largest eigenvalue,
- construct all relevant operators (system block and environment Hamiltonians, observables) for a system block of size $l+1$ in the reduced density matrix eigenbasis,
- form a superblock Hamiltonian from the system and environment block (size $l-1$) Hamiltonians plus two single sites (see Fig. 3) and determine its ground state by diagonalization. This is usually the most time-consuming step, performed by a Davidson procedure with sparse MVM as its core.

These steps must be repeated several times, shifting the interface between system block and environment block back and forth until some convergence criterion is fulfilled. This might be e. g. stationarity of the ground-state energy or a sufficiently small *discarded weight*, which is the sum of all density matrix eigenvalues that were not considered when forming the basis. The procedure can be generalized to two

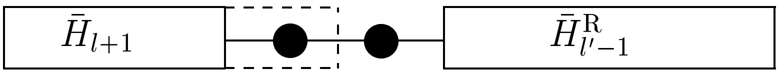

Fig. 3. *One step of the finite system DMRG algorithm (left-to-right phase).* $\bar{H}_{l+1}$ *and* $\bar{H}^{\mathrm{R}}_{l'-1}$ *are system block and environment block Hamiltonians in the reduced density matrix eigenbasis.*

dimensions, although it is not quite clear as to how the best "path" for the sweeps through the grid should be chosen [12].

The accuracy of observables like the ground-state energy depends on the number m of density matrix states kept. The discarded weight gives some hint for choosing the right m for a particular problem. Usually one starts with m rather small and increases m every time the ground-state energy has converged. Sensible values for m depend on the physical model under consideration. In the one-dimensional case, where DMRG usually performs best, $m = 500$ to 1000 is often sufficient to get decent data, even for models with electron-phonon interaction like the HHM (1).

It must be stressed that many complications show up in implementing the algorithm for a real-world problem. Fermionic and bosonic commutation rules, reflection and other symmetries, boundary conditions, degeneracies etc. all require special attention [13, 11]. Here we wish to concentrate on the performance aspects alone.

Due to the fact that the Hamiltonian is constructed using operators from the system and environment blocks, respectively, the dominant operation in the Davidson diagonalization procedure, i.e. the sparse MVM, is not memory-bound as in the ED case but DGEMM-based and thus cache-bound [14]. This makes DMRG very well suited for contemporary RISC architectures with large caches, because the dense matrix-matrix multiplication can be optimized in a way that yields near-peak speed on virtually all architectures [15]. Table 1 shows some important features of the RISC systems on which calculations were performed. For a standard benchmark case (2-dimensional isotropic half-filled 4x4 Hubbard model with periodic boundary conditions at $U = 4$ and $m = 2000$), Fig. 4 displays absolute performance numbers in MFlop/s for all systems used. Obviously, a significant fraction of peak performance is achievable.

Table 1. *One-CPU peak performance in GFlop/s, max. SMP node size and organization of shared memory accesses for all systems used for DMRG calculations. Proprietary, vendor-supplied BLAS and LAPACK implementations were used in all cases. In the case of SGI Origin we ran the calculations on 28 CPU system; the SGI Origin ccNUMA architecture however can be scaled up to 512 processors.*

System	**Peak Perf. [GFlop/s]**	**Max. #CPUs per SMP node**	**Memory access**
IBM p690/Power4 (1.3 GHz)	5.2	32	ccNUMA
HP rx5670/Itanium2 (1 GHz)	4.0	4	UMA
Intel Xeon DP (2.4 GHz)	4.8	2	UMA
SunFire 3800 (900 MHz)	1.8	24	ccNUMA
SGI Origin 3400 (500 MHz)	1.0	28 (512)	ccNUMA

Several parallelization approaches are conceivable for the DMRG algorithm. In the majority of cases, OpenMP parallelization of the MVM step in the Davidson procedure is most suitable, and that is what was used here to obtain the non-dynamical DMRG results.

DMRG can be used to obtain information about the system that requires more than the ground-state properties. This *dynamical* DMRG (DDMRG) approach can calculate important dynamical observables like optical conductivity [16]. Unfortunately, the superblock diagonalization is not any more the dominant step in DDMRG but other parts like the calculation of correction vectors require more computing time. In the present state of our software it does not make sense to use DDMRG in a parallel fashion, i e. all dynamical results have been obtained with 1-CPU runs.

3 Detecting Peierls versus Mott insulating phases in the 1D Holstein Hubbard model

Over the past years the ground-state phase diagram of the Holstein-Hubbard model (1) has been intensely studied using numerical methods. However, despite the apparent simplicity of the model, the quantum lattice and interaction effects are still not completely understood. At half-filling, electron-phonon and electron-electron interactions tend to localize the charge carriers by establishing CDW and SDW ground states, respectively. As a result, Peierls or Mott insulating phases are energetically favoured over the metallic state (Fig. 5). At $U = 0$, the ground state is a Peierls insulator (PI) above a critical electron-phonon coupling $g_c(\omega_0)$, with a vanishing threshold in the adiabatic limit $\omega_0 \to 0$. The PI state has equal spin and charge excitation gaps, and site-parity eigenvalue $P = +1$ [17]. It is mainly characterized by alternating doubly-occupied and empty sites and exhibits long-range order because the CDW phase has broken discrete symmetry. By contrast, at $g = 0$

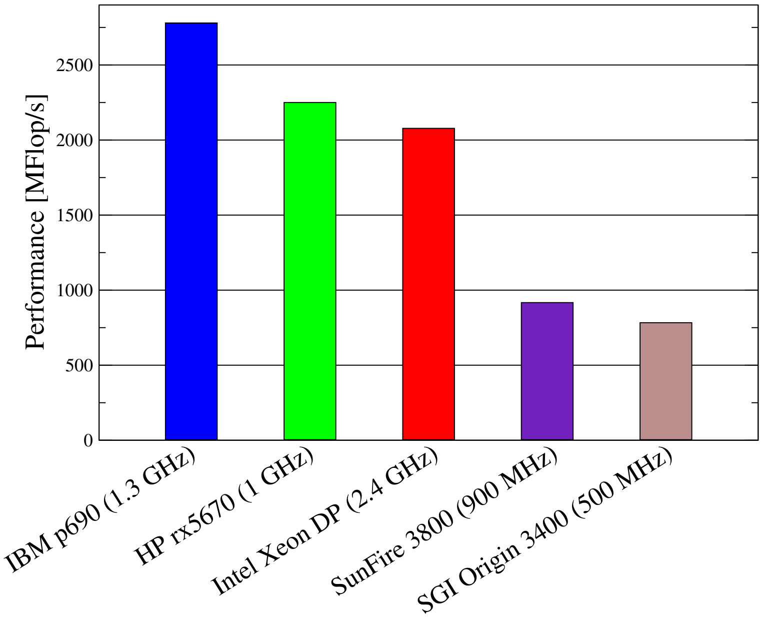

Fig. 4. *Sustained single processor performance for DMRG for a standard benchmark case (see text).*

the ground state is a Mott insulator (MI) with finite charge excitation gap but vanishing spin excitation gap and site-parity $P = -1$ [17]. The corresponding SDW phase has continuous symmetry and hence cannot exhibit true long-range order in 1D.

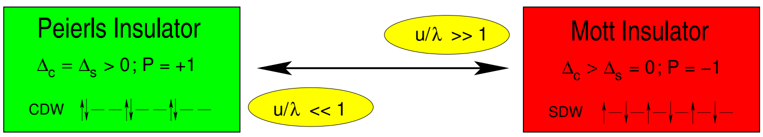

Fig. 5. Schematic phase diagram of the Holstein Hubbard model, where $u = U/4t$ and $\lambda = \varepsilon_p/2t = 2\alpha g^2$ with $\alpha = \omega_0/t$.

An interesting and still controversial question is whether or not only one critical point separates PI and MI phases at T=0 [19, 20]. Furthermore, it is also important to understand how the nature of the PI and MI phases is modified when phonon dynamical effects are taken into account. For a detailed discussion of these problems we refer to recent work [17]. Here we will focus on computational aspects at the determination of ground-state and spectral properties of the HHM.

To spot the transition from the PI to the MI state we first trace the behaviour of the kinetic energy $E_{\rm kin} = -t\langle\sum_{\langle i,j\rangle\sigma} c_{i\sigma}^{\dagger} c_{j\sigma}\rangle$ when increasing the Coulomb on-site repulsion U (see Fig. 6). Since both CDW and SDW correlations reduce the kinetic energy significantly, E_{kin} reaches a maximum when the system crosses from the PI to the MI regime. At the transition point the optical gap closes. In the non-adiabatic strong electron-phonon coupling regime at small u/λ, the electrons are heavily dressed by phonons, forming bipolarons in real space. As a result the system typifies rather a charge-ordered bipolaronic insulator than a traditional band insulator. This is reflected by the number of phonons needed to reach convergence using the ED approach (see inset). On the other hand, at the transition point and within the MI phase, the ground state is basically a zero-phonon state. We emphasize the extremely weak finite-size dependence of the E_{kin} results.

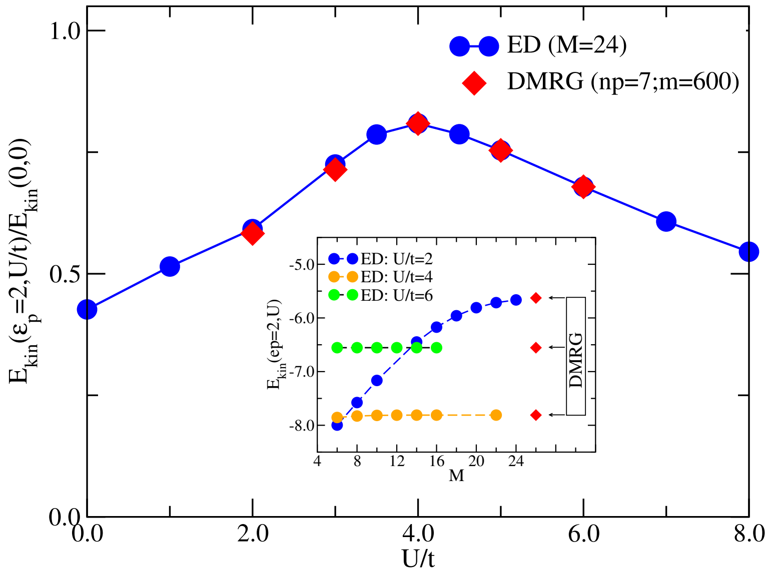

Fig. 6. Kinetic energy as a function of Hubbard interaction U/t with $g^2 = 2$ and $\omega_0/t = 1$. Inset: Comparison of ED and DMRG results for E_{kin} at different U with $g^2 = 2$ and $\omega_0/t = 1$; Convergence of ED results as a function of the cut-off parameter M is demonstrated. The corresponding results from DMRG calculations using $np = 6$ pseudo-sites and $m = 1000$ are represented by stars.

Further insight into the nature of the PI and MI phases can be obtained by calculating staggered charge- and spin-structure factors:

$$S_{\mathrm{c}}(\pi) = \frac{1}{N^2} \sum_{\substack{i,j \\ \sigma\sigma'}} (-1)^{|i-j|} \langle (n_{i\sigma} - \frac{1}{2})(n_{j\sigma'} - \frac{1}{2}) \rangle \quad , \tag{7}$$

$$S_{\mathrm{s}}(\pi) = \frac{1}{N^2} \sum_{i,j} (-1)^{|i-j|} \langle S_i^z S_j^z \rangle \quad , \quad S_i^z = \frac{1}{2}(n_{i\uparrow} - n_{i\downarrow}) \, . \tag{8}$$

Results for $S_{\mathrm{c}}(\pi)$ and $S_{\mathrm{s}}(\pi)$ are given in Fig. 7 for two characteristic Hubbard interactions. Of course, we found pronounced CDW and weak SDW correlations in

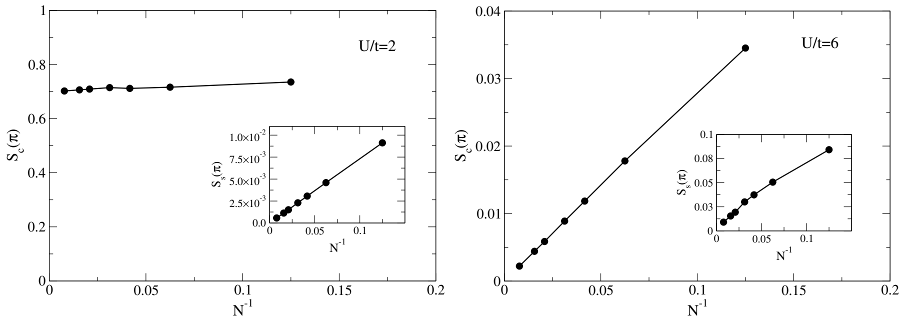

Fig. 7. DMRG study of staggered charge and spin (inset) structure factors as a function of inverse lattice size in the PI (left panel) and MI (right panel) regimes. Generally, $m = 1000$ and 5 boson pseudosites were chosen for the DMRG algorithm, with the exception of the $N = 128$ case where $m = 800$ had to be used for practical reasons with no qualitative loss.

the Peierls distorted state. Increasing U at fixed g and ω_0, Peierls CDW order is strongly suppressed, whereas the spin structure factor becomes enhanced [17]. Most interesting, however, are the different size dependences of these quantities in the PI and MI phases. For the PI, $S_c(\pi)$ shows almost no dependence on the size of the system, which indicates CDW long-range order, whereas $S_c(\pi)$ obviously scales to zero as $N \to \infty$. By contrast, in the MI regime our data provides strong evidence for vanishing spin and charge structures in the thermodynamic limit. Clearly the MI is characterized by short-ranged antiferromagnetic spin correlations but nevertheless the staggered spin-spin correlation function shows a slow (algebraic) decay at large distances. Note that such a finite-size scaling (shown here with lattice sizes up to $N = 128$) is definitely out of the range of ED, i.e. in order to obtain reliable results concerning the behaviour of the infinite system we had to apply the DMRG method. The DMRG memory and CPU-time resources required to compute the kinetic energy and structure factor data presented are moderate, at least in relation to the ED algorithm. Each data point in Fig. 6 takes roughly one CPU day on a modern processor like the Intel Itanium 2, at a few GBytes of memory. The DMRG finite-size analysis of structure factors in Fig. 7 is considerably more demanding, with about 5 to 10 GBytes of memory and several CPU days of computer time.

As an example for the calculation of a dynamical (spectral) quantity, we have determined the regular part of the optical conductivity at $T = 0$,

$$\sigma^{\mathrm{reg}}(\omega) = \frac{\pi}{N} \sum_{m \neq 0} \frac{|\langle \psi_0 | \hat{\jmath} | \psi_m \rangle|^2}{E_m - E_0} \, \delta(\omega - E_m + E_0) \,. \tag{9}$$

Here $|\psi_0\rangle$ and $|\psi_m\rangle$ denote the ground state and excited states, respectively, with corresponding energies E_0 and E_m. The current operator $\hat{\jmath} = -\mathrm{i}et \sum_{i\sigma} (c_{i\sigma}^{\dagger} c_{i+1\,\sigma} - c_{i+1\,\sigma}^{\dagger} c_{i\sigma})$ has finite matrix elements between states of different site-parity only.

Figure 8 displays the optical conductivity of the Holstein-Hubbard model in the MI region, where $\sigma^{\mathrm{reg}}(\omega)$ is dominated by excitations that can be related to those of the pure Hubbard model. In addition phonon side bands appear at multiples of the bare phonon frequency. The optical gap, which represents by its nature a correlation gap, is clearly visible. The optical conductivity can be obtained as well using dynamical DMRG. As our DDMRG code is non-parallel, all calculations were done on one CPU. The DDMRG method yields the optical conductivity (9) convolved with a Lorentzian distribution of width η. The computational effort is smaller for larger broadening η. Of course, the resolution increases as η decreases. As can be seen from the comparison of DDMRG and ED results small values of η are necessary to reproduce the ED peak structure. In fact, DDMRG is better suited for calculating continuous spectral functions in large systems than the discrete spectra of small systems. For the small system investigated here the Lanczos-DMRG approach [21] could be more efficient. We have found that it gets progressively harder (with respect to computing time) to calculate the optical conductivity as the frequency ω increases. This is probably due to the exponentially increasing number of phonon excitations contributing to the optical conductivity at excitation energies larger than ω_0. The DDMRG curves in Fig. 8 were drawn to a point where the required CPU-time is comparable for all values of η (roughly 7 hours on a 2.4 GHz Xeon). Memory requirements for DDMRG are generally quite small. All calculations could be done in less than 200 MBytes of RAM.

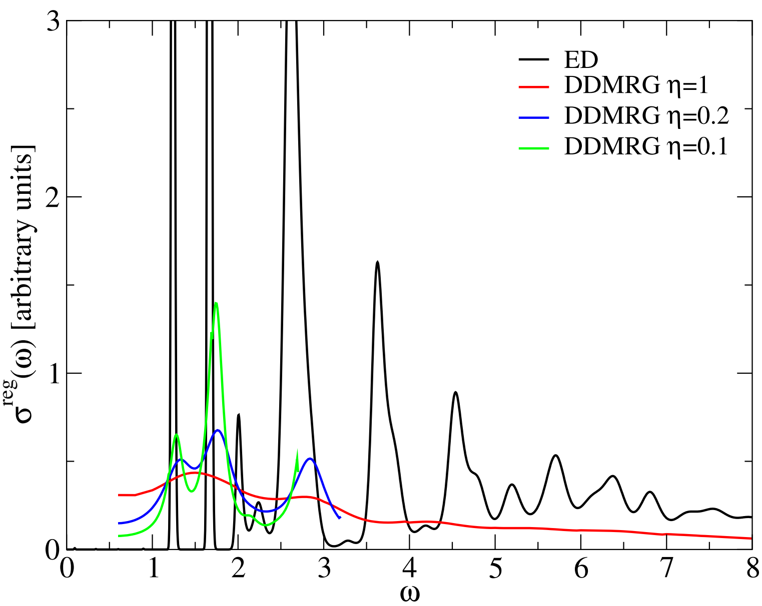

Fig. 8. Optical conductivity for the $N = 8$ site HHM at $U/t = 6$, $g^2 = 2$ and $\omega_0 = 1$. The DDMRG data was calculated at $m = 200$ using three different broadenings $\eta = 1$, $\eta = 0.2$ and $\eta = 0.1$, respectively. Curves for each η were drawn to a point where the CPU-time for the different runs is comparable.

4 Conclusions

In this work we have presented two methods for the calculation of ground state as well as dynamical properties of low-dimensional physical systems. Where Exact Diagonalization (ED) requires vast computing resources already with very small systems, it is nevertheless free of approximations (apart from a well-controlled phononic cutoff). DMRG methods, on the other hand, are variational by design, but in a way that optimizes the accuracy of measured observables. They allow us to investigate much larger systems with moderate computational effort compared to ED. Non-dynamical DMRG can be used very efficiently on modern, large-memory SMP nodes with high peak performance and outperforms even the best massively parallel ED implementations by far. Dynamical DMRG (DDMRG), on the other hand, while still being able to provide important insights into dynamical properties of much larger systems than ED, has difficulty resolving the discrete peak structure of the small system investigated here and suffers up to now from insufficient parallelization potential. Although results from desktop computer systems look promising, a sensible parallelization approach would be in order to take the step towards supercomputing.

As an example for the reliability and predictive power of such large-scale numerical many-body calculations performed on modern supercomputers, we have analyzed ground state and spectral properties of the one-dimensional half-filled Holstein-Hubbard model with respect to the Peierls-insulator to Mott insulator transition. Significant differences in the finite-size scaling behaviour of the staggered charge and spin structure factors allow to distinguish Peierls and Mott insulating states.

Acknowledgments

We thank the RRZE, the RZG (Computing Center Garching), the URZ (Universitätsrechenzentrum Dresden) and the HLRS (High Performance Computing Center Stuttgart) for providing computational resources. The RRZN (Regionales Rechenzentrum Niedersachsen) and the ZIB (Zuse-Institut Berlin) have granted resources on their HLRN supercomputer complex. This work was partially supported by the Bavarian Competence Network for High Performance Computing (G. H.) and the DAAD (H. F.).

References

1. For an overview on several important aspects of strongly correlated electron systems see Science Vol. **288** (2000)
2. A. R. Bishop and B. I. Swanson, Novel Electronic Materials: the MX Family. Los Alamos Science **21**, 133 (1993)
3. J. Hubbard, Electron Correlations in Narrow Energy Bands. Proc. Roy. Soc. London **A 276**, 238–257 (1963); J. Kanamori, Electron Correlation and Ferromagnetism of Transition Metals, Prog. Theor. Phys. **30**, 275–289 (1963)
4. T. Holstein, Studies of Polaron Motion. 1. The Molecular Crystal Model. Ann. Phys. (N.Y.) **8**, 325–342 (1959); Studies of Polaron Motion. 1. The Small Polaron. Ann. Phys. (N.Y.) **8**, 343–389 (1959)

5. G. Wellein and H. Fehske, Self-trapping problem of electrons or excitons in one dimension, Phys. Rev. B **58**, 6208–6218 (1998)
6. H. Fehske, M. Holicki, and A. Weiße, Lattice dynamical effects on the Peierls transition in one-dimensional metals and spin chains. Advances in Solid State Physics, **40**, 235-249 (2000)
7. B. Bäuml, G. Wellein, and H. Fehske, Optical absorption and single-particle excitations in the two-dimensional Holstein-tJ model. Phys. Rev. B **58**, 3663–3676 (1998).
8. G. Wellein, H. Röder, and H. Fehske, Polarons and Bipolarons in Strongly Interacting Electron-Phonon Systems. Phys. Rev. B **33**, 9666–9675 (1996)
9. G. Wellein and H. Fehske, in E. Krause and W. Jäger (Eds.): High Performance Computing in Science and Engineering '99, pp 112-129, 2000. Springer-Verlag Berlin Heidelberg (2000)
   Towards the limits of present-day supercomputers: Exact diagonalization of strongly correlated electron-phonon systems
10. S. R. White, Density Matrix Formulation for Quantum Renormalization Groups. Phys. Rev. Lett. **69**, 2863–2866 (1992)
11. S. R. White, Density-Matrix Algorithms for Quantum Renormalization Groups. Phys. Rev. B **48**, 10345–10356 (1993)
12. R. M. Noack and S. R. White, The Density Matrix Renormalization Group. In: I. Peschel, X. Wang, M. Kaulke and K. Hallberg (eds), Density-Matrix Renormalization: A New Numerical Method in Physics. Lectures of a seminar and workshop, held at the Max-Planck-Institut für Physik Komplexer Systeme, Dresden, Germany, August 24th to September 18th, 1998. Springer, Berlin Heidelberg New York (1999) (Lecture Notes in Physics Vol. 528)
13. E. Jeckelmann and S. R. White, Density-Matrix Renormalization Group Study of the Polaron Problem in the Holstein Model. Phys. Rev. B **57**, 6376–6385 (1998)
14. G. Hager, E. Jeckelmann, H.Fehske, and G. Wellein, Parallelization Strategies for Density Matrix Renormalization Group Algorithms on Shared-Memory Systems. `arXiv:cond-mat/0305463`
15. S. Goedecker and A. Hoisie, Performance Optimization of Numerically Intensive Codes. Siam, Philadelphia (2001)
16. E. Jeckelmann, Dynamical density-matrix renormalization-group method. Phys. Rev. B **66**, 045114 (2002)
17. H. Fehske, G. Wellein, A. Weiße, F. Göhmann, H. Büttner, and A. R. Bishop, Peierls insulator Mott-insulator transition in 1D. Physica B 312-313 , 562-563 (2002); H. Fehske, A. P. Kampf, M. Sekania and G. Wellein, Nature of the Peierls- to Mott-insulator transition in 1D. Eur. Phys. J. B **31**, 11-16 (2003)
18. H. Fehske, G. Wellein, A. P. Kampf, M. Sekania, G. Hager, A. Weiße, H. Büttner, and A. R. Bishop, in S. Wagner et al. (Eds.): High Performance Computing in Science and Engineering Munich 2002, pp. 339-349, 2003. Springer-Verlag Berlin Heidelberg (2003)
   One-dimensional electron-phonon systems: Mott- versus Peierls-insulators
19. M. Fabrizio, A. O. Gogolin, and A. A. Nersesyan, From Band Insulator to Mott Insulator in One Dimension. Phys. Rev. Lett. **83**, 2014-2017 (1999)
20. Ph. Brune, G. I. Japradize, A. P. Kampf, and M. Sekania, Nature of the insulating phases in the half-filled ionic Hubbard model. arXiv:cond-mat/0304697
21. K. Hallberg, Density-matrix algorithm for the calculation of dynamical properties of low-dimensional systems. Phys. Rev. B **52**, R9827-R9830 (1995)